Diffusion Model Based Signal Recovery Under 1-Bit Quantization

Youming Chen, Zhaoqiang Liu*

University of Electronic Science and Technology of China

Abstract

Diffusion models (DMs) have demonstrated to be powerful priors for signal recovery, but their application to 1-bit quantization tasks, such as 1-bit compressed sensing and logistic regression, remains a challenge. This difficulty stems from the inherent non-linear link function in these tasks, which is either non-differentiable or lacks an explicit characterization. To tackle this issue, we introduce Diff-OneBit, which is a fast and effective DM-based approach for signal recovery under 1-bit quantization. Diff-OneBit addresses the challenge posed by non-differentiable or implicit links functions via leveraging a differentiable surrogate likelihood function to model 1-bit quantization, thereby enabling gradient based iterations. This function is integrated into a flexible plug-and-play framework that decouples the data-fidelity term from the diffusion prior, allowing any pretrained DM to act as a denoiser within the iterative reconstruction process. Extensive experiments on the FFHQ, CelebA and ImageNet datasets demonstrate that Diff-OneBit gives high-fidelity reconstructed images, outperforming state-of-the-art methods in both reconstruction quality and computational efficiency across 1-bit compressed sensing and logistic regression tasks.

Introduction

In linear measurement models (Candès, Romberg, and Tao 2006; Donoho 2006; Foucart and Rauhut 2013), the goal is to recover a high-dimensional signal $\mathbf{x}^* \in \mathbb{R}^N$ from noisy linear measurements:

$$\mathbf{y} = \mathbf{A}\mathbf{x}^* + \boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_M]^\top \in \mathbb{R}^{M\times N}$ is the forward matrix, $\boldsymbol{\epsilon} = [\epsilon_1, \dots, \epsilon_M]^\top \in \mathbb{R}^M$ is the noise vector, and $\mathbf{y} = [y_1, \dots, y_M]^\top \in \mathbb{R}^M$ represents the acquired measurements. To address this typically ill-posed measurement model, especially in the high-dimensional regime where $M \ll N$, a series of works has developed effective recovery methods by exploiting hand-crafted structural priors of $\mathbf{x}^*$, such as sparsity or low-rankness (Chen et al. 2013; Li 2013; Nguyen and Tran 2012, 2013; Xu, Caramanis, and Mannor 2012).

However, the assumption of infinite measurement precision in linear measurement models is often unrealistic in hardware implementation. In practice, measurements are quantized to a finite number of bits. In particular, 1-bit quantization is valued for its low hardware cost and robustness to non-linear distortions (Boufounos and Baraniuk 2008; Boufounos 2010). Both 1-bit compressed sensing (CS) and logistic regression are examples of the 1-bit quantized measurement model. For 1-bit CS, measurements are quantized to their signs:

$$\mathbf{y} = \text{sign}(\mathbf{A}\mathbf{x}^* + \boldsymbol{\epsilon}). \tag{2}$$

For logistic regression, the likelihood of each binary outcome $y_i \in \{-1, +1\}$ can be expressed as follows:

$$\mathbb{P}(y_i = 1) = \frac{1}{1 + e^{-\mathbf{a}_i^\top \mathbf{x}^*}} = \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{2}\right). \tag{3}$$

1-bit quantized measurement models have also been widely studied under hand-crafted structural priors such as sparsity and low-rankness (Plan and Vershynin 2012; Yan, Yang, and Osher 2012; Plan and Vershynin 2013; Jacques et al. 2013; Davenport et al. 2014; Knudson, Saab, and Ward 2016; Plan, Vershynin, and Yudovina 2017).

Priors have shifted from hand-crafted designs to leveraging the expressive capabilities of conventional generative models, such as variational autoencoders (VAEs) and generative adversarial networks (GANs) (Bora et al. 2017; Liu and Scarlett 2020b). Recently, diffusion models (DMs) (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020) stand out as powerful generative priors for capturing the statistical properties of natural images. Reconstruction methods using DMs generally fall into two categories: Supervised methods, which train neural networks end-to-end on paired clean and degraded images for specific measurement models, and unsupervised methods, which utilize pretrained unconditional diffusion models and knowledge of the degradation process. Recently, unsupervised methods have been favored for their versatility. They typically integrate a data-fidelity term with a DM-based prior within a gradient based optimization or sampling framework, achieving reliable reconstruction quality across a wide range of linear and non-linear measurement models (Kadkhodaie and Simoncelli 2021; Jalal et al. 2021a; Choi et al. 2021; Chung et al. 2022, 2023; Zhu et al. 2023; Kawar et al. 2022; Wang, Yu, and Zhang 2023; Song et al. 2023, 2024; Mardani et al. 2024; Zhang et al. 2025; Zheng, Li, and Liu 2025).

*Corresponding author.

Specifically, under the assumption that the nonlinear forward operator $\mathcal{A} : \mathbb{R}^N \to \mathbb{R}^M$ is differentiable, approaches such as DPS (Chung et al. 2023), DiffPIR (Zhu et al. 2023), and DAPS (Zhang et al. 2025) can handle nonlinear measurement models of the form $\mathbf{y} = \mathcal{A}(\mathbf{x}^*) + \boldsymbol{\epsilon}$ through gradient based iterative procedures. However, these methods are not directly applicable to 1-bit CS or logistic regression, as the corresponding forward operators are either non-differentiable or lack an explicit formulation.

Contributions To address this limitation, we draw inspiration from plug-and-play (PnP) frameworks (Venkatakrishnan, Bouman, and Wohlberg 2013; Chan, Wang, and Elgendy 2016; Sreehari et al. 2016; Kamilov, Mansour, and Wohlberg 2017) and propose Diff-OneBit, which is a PnP method designed to tackle the 1-bit CS problem using pretrained DMs. To overcome the challenge posed by the non-differentiable link function in 1-bit CS, our approach decouples the data and prior terms via half-quadratic splitting (HQS) (Geman and Yang 1995) and introduces a differentiable surrogate function for the data-fidelity term, solving them iteratively within the diffusion sampling framework. We further adapt Diff-OneBit to address the logistic regression problem. Evaluations of Diff-OneBit on 1-bit CS and logistic regression tasks using the FFHQ and CelebA datasets demonstrate that it achieves superior image reconstruction quality with greater efficiency compared to competing methods.

Related work

In this subsection, we present an overview of relevant existing studies. These studies can broadly be categorized into (i) Signal recovery with conventional generative models, and (ii) Signal recovery with diffusion models.

Signal recovery with conventional generative models: Building on the pioneering work of (Bora et al. 2017), prior research has explored signal recovery problems using conventional generative models. Subsequent studies have introduced novel architectures, including untrained deep image priors (Van Veen et al. 2018), underparameterized deep decoders (Heckel and Hand 2019), and invertible networks designed to eliminate representation error (Asim et al. 2020). Additionally, several methods with provable guarantees have been developed, such as those based on Langevin dynamics (Nguyen, Jagatap, and Hegde 2021) and posterior sampling (Jalal et al. 2021b). The application scope of these priors has expanded to address problems like phase retrieval (Hand, Leong, and Voroninski 2018; Liu, Ghosh, and Scarlett 2021; Liu, Wang, and Liu 2022), generative principal component analysis (Liu et al. 2022; Chen, Ng, and Liu 2025) and generalized eigenvalue problems (Liu, Li, and Chen 2024). This progress has been supported by extensive studies confirming the adversarial robustness of these methods (Genzel, Macdonald, and März 2022) and comprehensive surveys of the field (Ongie et al. 2020; Scarlett et al. 2023).

Notably, there have been specific investigations into 1-bit CS using conventional generative models. For instance, Qiu, Wei, and Yang (2020) employ ReLU generative networks with dithering techniques to enable signal recovery. Liu et al. (2020) establish recovery guarantees for 1-bit CS using conventional generative models. Furthermore, single-index models encompassing 1-bit quantized measurement models have been studied under conventional generative priors (Wei, Yang, and Wang 2019; Liu and Scarlett 2020a; Liu and Han 2022; Liu and Liu 2022; Chen et al. 2023). However, these works focus on recovery guarantees related to optimal solutions for empirical risk minimization problems or approaches that approximately optimize over the range of conventional single-step generative models, without considering plug-and-play frameworks.

Signal recovery with diffusion models: A family of unsupervised methods using pretrained DMs excels at solving signal recovery problems by adapting the reverse diffusion process. For example, Score-ALD (Jalal et al. 2021a) utilizes Langevin dynamics for posterior sampling, while MCG (Chung et al. 2022) adds manifold constraints in reverse sampling. ILVR (Choi et al. 2021) refines each step based on reference images, and DDRM (Kawar et al. 2022) operates in the spectral domain of the forward operator. DiffPIR (Zhu et al. 2023) integrates plug-and-play frameworks. DPS (Chung et al. 2023) extends posterior sampling to noisy and non-linear problems. DDNM (Wang, Yu, and Zhang 2023) introduces null-space guidance and ΠGDM (Song et al. 2023) introduces pseudoinverse guidance in sampling. ReSample (Song et al. 2024) utilizes hard data consistency for latent diffusion models. RED-diff (Mardani et al. 2024) offers a variational perspective connecting diffusion to regularization by denoising. DAPS (Zhang et al. 2025) employs a decoupled noise annealing process. However, the aforementioned methods are primarily designed for linear measurement models, while methods capable of handling non-linear measurement models, such as DPS, DiffPIR, DAPS, RED-diff, and ΠGDM, are mainly designed for explicit and differentiable forward operators and cannot effectively handle quantized measurements. In the following, we discuss recent DM-based methods aiming for quantized signal recovery.

In QCS-SGM (Meng and Kabashima 2023), a DM serves as a prior for recovering signals from noisy quantized measurements through posterior sampling, integrating a noise-perturbed pseudo-likelihood score with the score of the DM. However, QCS-SGM relies on annealed Langevin dynamics, which demands a large number of iterations and incurs substantial computational overhead. SIM-DMIS (Tang et al. 2025) investigates single index models that encompass 1-bit quantized measurement models with a diffusion prior. The approach first obtains an appropriate initial estimate, then reconstructs the underlying signal using a single round of unconditional sampling combined with partial inversion of the DM, which is an operation that approximates projecting the initial estimate onto the range the DM.

Background

Diffusion models

DMs are generative models that map data to noise through a pre-defined forward diffusion process (Ho, Jain, and Abbeel 2020; Karras et al. 2022; Song and Ermon 2019, 2020; Song

et al. 2021). Specifically, the forward process perturbs a data sample $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$ over a continuous time variable $t \in [0, T]$ according to a Gaussian transition kernel:

$$q_{0t}(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I}), \tag{4}$$

where α_t and σ_t are time-dependent functions controlling the signal and noise schedules. Let p_t denotes the marginal distribution of $\mathbf{x}_t$. As $t \to T$, the perturbed distribution $p_t(\mathbf{x}_t)$ converges to an isotropic Gaussian.

To generate a sample, DMs learn to reverse this process, starting from pure Gaussian noise $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and iteratively denoising it to produce a sample from the data distribution. Specifically, a neural network $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)$ is trained to predict the noise component $\boldsymbol{\epsilon}$ from the noisy state $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}$ by minimizing the following objective (Song and Ermon 2019; Song, Meng, and Ermon 2021):

$$\mathbb{E}_{\mathbf{x}_0 \sim p_0, \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I}), t \sim \mathcal{U}(0,T)}[\|\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) - \boldsymbol{\epsilon}\|_2^2]. \tag{5}$$

The pretrained noise predictor $\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)$ provides an estimate of the score function via the relation $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \approx -\boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)/\sigma_t$. Thus, sampling from DMs can be performed by numerically solving a stochastic differential equation (SDE) (Song et al. 2021):

$$\mathrm{d}\mathbf{x}_t = \left[f(t)\mathbf{x}_t + \frac{g^2(t)}{\sigma_t} \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t) \right] \mathrm{d}t + g(t)\mathrm{d}\mathbf{w}_t, \tag{6}$$

where $\mathbf{w}_t$ represents a standard Wiener process, and $f(t)$ and $g(t)$ are the drift and diffusion coefficients, respectively. In Variance Preserving (VP) SDE, we have $f(t) = \mathrm{d}\log\alpha_t/\mathrm{d}t$, $g^2(t) = \mathrm{d}\sigma_t^2/\mathrm{d}t - 2f(t)\sigma_t^2$. Alternatively, sampling can be performed deterministically by numerically solving the corresponding probability flow ordinary differential equation (ODE):

$$\frac{\mathrm{d}\mathbf{x}_t}{\mathrm{d}t} = f(t)\mathbf{x}_t + \frac{g^2(t)}{2\sigma_t} \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\mathbf{x}_t, t). \tag{7}$$

Conditional diffusion models for signal recovery

For signal recovery with measurements $\mathbf{y}$, the objective is to sample the clean signal $\mathbf{x}_0$ from the posterior distribution $p(\mathbf{x}_0|\mathbf{y})$. DMs achieve this by conditioning the reverse sampling process, modifying Eq. (7) with the posterior score $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathbf{y})$:

$$\frac{\mathrm{d}\mathbf{x}_t}{\mathrm{d}t} = f(t)\mathbf{x}_t - \frac{g^2(t)}{2} \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathbf{y}). \tag{8}$$

For a linear measurement model with additive Gaussian noise $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$ as in Eq. (1), the conditional score is directly related to the conditional expectation as follows (Daras et al. 2024):

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathbf{y}) = \frac{\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}] - \mathbf{x}_t}{\sigma^2}. \tag{9}$$

DiffPIR approximates this conditional expectation using HQS, which decouples the data-fidelity and prior terms as follows:

$$\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}] \approx \arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \frac{\mu_t}{2}\|\mathbf{x} - \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]\|_2^2. \tag{10}$$

Here, $\mu_t := \lambda(\alpha_t^2/\sigma_t^2)$ with $\lambda > 0$ being a hyperparameter is a penalty parameter and $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ can be directly obtained using Tweedie's formula (Daras et al. 2024). Alternatively, approaches such as DPS approximate $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t)$ using Bayes' rule with $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|\mathbf{y}) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$, where the term $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ is simply the unconditional score function. However, these frameworks cannot be directly extended to handle 1-bit quantized measurement models for which both $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t, \mathbf{y}]$ and $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{y}|\mathbf{x}_t)$ are difficult to approximate.

Method

Our goal is to address 1-bit quantized measurement models using a DM as a powerful generative prior. We formulate the reconstruction as a maximum a posterior (MAP) estimation problem. Given the 1-bit measurements $\mathbf{y}$, we seek to find the optimal vector $\hat{\mathbf{x}}$ that maximizes the posterior likelihood:

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} \mathbb{P}(\mathbf{x}|\mathbf{y}) = \arg\max_{\mathbf{x}} \mathbb{P}(\mathbf{y}|\mathbf{x})\mathbb{P}(\mathbf{x}) \tag{11}$$

$$= \arg\min_{\mathbf{x}} \underbrace{-\log \mathbb{P}(\mathbf{y}|\mathbf{x})}_{\text{Data Fidelity Term}} + \underbrace{(-\log \mathbb{P}(\mathbf{x}))}_{\text{Prior Regularization Term}}. \tag{12}$$

In the following, We discuss the approaches designed for 1-bit CS and logistic regression, which are two popular instances of 1-bit quantization problems.

Approach for 1-bit CS

For the 1-bit CS model in Eq. (2), we begin by modeling the underlying measurement process before quantization. For any $\mathbf{x} \in \mathbb{R}^n$, letting $\tilde{b}_i = \mathbf{a}_i^\top \mathbf{x} + \epsilon_i$ be the corrupted measurement. The 1-bit measurement $y_i \in \{-1, +1\}$ is the result of applying the sign function to this noisy value:

$$y_i = \text{sign}(\tilde{b}_i). \tag{13}$$

The probability of observing $y_i = +1$ is the probability that the noisy measurement is non-negative:

$$\mathbb{P}(y_i = +1|\mathbf{x}) = \mathbb{P}(\tilde{b}_i \geq 0) = \mathbb{P}(\mathbf{a}_i^\top \mathbf{x} + \epsilon_i \geq 0) \tag{14}$$

$$= \mathbb{P}(\epsilon_i \geq -\mathbf{a}_i^\top \mathbf{x}). \tag{15}$$

Similarly, the probability of observing $y_i = -1$ is:

$$\mathbb{P}(y_i = -1|\mathbf{x}) = \mathbb{P}(\tilde{b}_i < 0) = \mathbb{P}(\mathbf{a}_i^\top \mathbf{x} + \epsilon_i < 0) \tag{16}$$

$$= \mathbb{P}(\epsilon_i < -\mathbf{a}_i^\top \mathbf{x}). \tag{17}$$

Eqs. (15) and (17) can be expressed using the cumulative distribution function (CDF). Assuming that the additive noise is Gaussian with $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$, then ϵ_i/σ follows the standard normal distribution, with its CDF defined as $\Phi(z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-u^2/2} \mathrm{d}u$, and we obtain

$$\mathbb{P}(y_i = 1|\mathbf{x}) = \mathbb{P}\left(\frac{\epsilon_i}{\sigma} \geq -\frac{\mathbf{a}_i^\top \mathbf{x}}{\sigma}\right) = \Phi\left(\mathbf{a}_i^\top \mathbf{x}/\sigma\right). \tag{18}$$

Applying the symmetry of the Gaussian distribution, we obtain $\mathbb{P}(\epsilon_i \geq v) = 1 - \mathbb{P}(\epsilon_i < v) = P(\epsilon_i \leq -v)$, and thus

$$\mathbb{P}(y_i = -1|\mathbf{x}) = \mathbb{P}\left(\frac{\epsilon_i}{\sigma} < -\frac{\mathbf{a}_i^\top \mathbf{x}}{\sigma}\right) = \Phi\left(-\mathbf{a}_i^\top \mathbf{x}/\sigma\right). \tag{19}$$

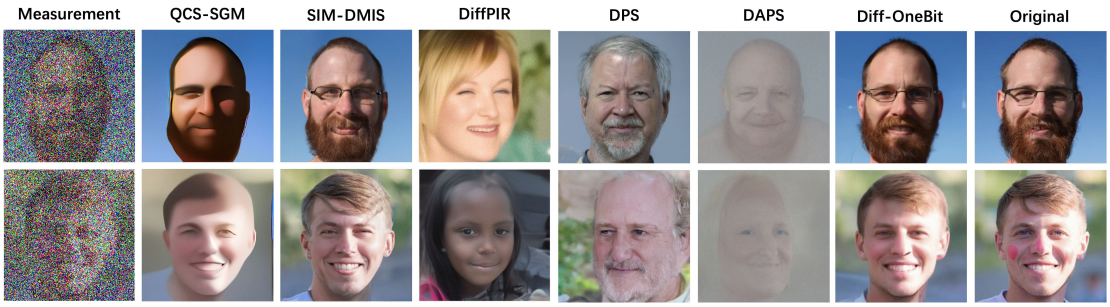

Figure 1. Qualitative results of 1-bit CS on FFHQ images. We compare Diff-OneBit with QCS-SGM, SIM-DMIS and DiffPIR. Input images have an addtive Gaussian noise of $\sigma = 0.5$.

Combining these two cases into a single one, yielding:

$$\mathbb{P}(y_i|\mathbf{x}) = \Phi\left(\mathbf{a}_i^\top \mathbf{x}/\sigma\right)^{\frac{1+y_i}{2}} \cdot \Phi\left(-\mathbf{a}_i^\top \mathbf{x}/\sigma\right)^{\frac{1-y_i}{2}}. \quad (20)$$

Then, we have the following for $\mathbf{y}$:

$$\mathbb{P}(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{M} \Phi\left(\mathbf{a}_i^\top \mathbf{x}/\sigma\right)^{\frac{1+y_i}{2}} \cdot \Phi\left(-\mathbf{a}_i^\top \mathbf{x}/\sigma\right)^{\frac{1-y_i}{2}}. \quad (21)$$

This expression serves as our differentiable surrogate likelihood. For the MAP objective, we require its negative logarithm. Taking the negative logarithm of Eq. (21) yields our final data-fidelity function as follows:

$$\mathcal{L}(\mathbf{x};\mathbf{y}) = -\sum_{i=1}^{M} \left[\frac{1+y_i}{2} \log \Phi\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{\sigma}\right) + \frac{1-y_i}{2} \log \Phi\left(-\frac{\mathbf{a}_i^\top \mathbf{x}}{\sigma}\right)\right]. \quad (22)$$

With a tractable data-fidelity term, we now address the MAP objective in Eq. (12). HQS introduces an auxiliary variable $\mathbf{z}$ to decouple the data-fidelity and prior terms, reformulating the objective as

$$\arg\min_{\mathbf{x},\mathbf{z}} \mathcal{L}(\mathbf{x};\mathbf{y}) - \log p(\mathbf{z}) \quad \text{subject to} \quad \mathbf{x} = \mathbf{z}. \quad (23)$$

This constrained optimization problem is typically addressed by iteratively updating $\mathbf{x}$ and $\mathbf{z}$ through two subproblems, synchronized with the timestep t to ensure consistency with its reverse process schedule.

The $\mathbf{x}$-update rule enforces consistency with the observed 1-bit measurements using a differentiable surrogate likelihood. The subproblem is of the following form:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \left(\mathcal{L}(\mathbf{x};\mathbf{y}) + \frac{\mu}{2}\|\mathbf{x} - \hat{\mathbf{z}}\|_2^2\right), \quad (24)$$

where $\mu > 0$ is the penalty parameter and $\hat{\mathbf{z}}$ is an estimated vector. The differentiability of $\mathcal{L}(\cdot;\mathbf{y})$ enables efficient optimization via gradient methods like Adam (Kingma and Ba 2014). The $\mathbf{z}$-update solves:

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \left(-\log p(\mathbf{z}) + \frac{\mu}{2}\|\mathbf{z} - \hat{\mathbf{x}}\|_2^2\right). \quad (25)$$

The optimization problem in Eq. (25) can be thought of as a denoising problem, and its optimal solution can be approximated using a DM through the single-step Tweedie's formula, see, e.g., (Zhu et al. 2023; Venkatakrishnan, Bouman, and Wohlberg 2013), for which we provide a detailed analysis in the appendix.

Based on the above discussions, our Diff-OneBit approach is illustrated in Figure 3 and detailed in Algorithm 1.

Adaption for logistic regression

We can adapt our Diff-OneBit approach designed for 1-bit CS to logistic regression. For the logistic regression model with 1-bit observation $y_i \in \{-1, +1\}$, the exact likelihood of an observation can be directly expressed using the hyperbolic tangent function. The probabilities for the two possible outcomes are:

$$\mathbb{P}(y_i = +1|\mathbf{x}) = \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{2}\right), \quad (26)$$

$$\mathbb{P}(y_i = -1|\mathbf{x}) = \frac{1}{2} - \frac{1}{2}\tanh\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{2}\right). \quad (27)$$

Following the structure of the 1-bit model, we combine both cases into a single likelihood expression for $\mathbb{P}(y_i|\mathbf{x})$. Analogous to the 1-bit CS formulation in Eq. (21), the exponents select the appropriate probability for the binary outcome:

$$\mathbb{P}(\mathbf{y}|\mathbf{x}) = \prod_{i=1}^{M} \mathbb{P}(y_i|\mathbf{x}). \quad (28)$$

Taking the negative logarithm yields the data-fidelity function, which corresponds to the standard binary cross-entropy loss:

$$\mathcal{L}(\mathbf{x};\mathbf{y}) = -\log \mathbb{P}(\mathbf{y}|\mathbf{x}) \quad (29)$$

$$= -\sum_{i=1}^{M} \left[\frac{1+y_i}{2} \log\left(\frac{1}{2} + \frac{1}{2}\tanh\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{2}\right)\right) + \frac{1-y_i}{2} \log\left(\frac{1}{2} - \frac{1}{2}\tanh\left(\frac{\mathbf{a}_i^\top \mathbf{x}}{2}\right)\right)\right]. \quad (30)$$

The function in Eq. (30) is convex and smooth, making it well-suited for gradient based optimization within our MAP framework. Unlike the surrogate likelihood in Eq. (22), this formulation represents the exact model likelihood and requires no hyperparameters such as the noise level σ.

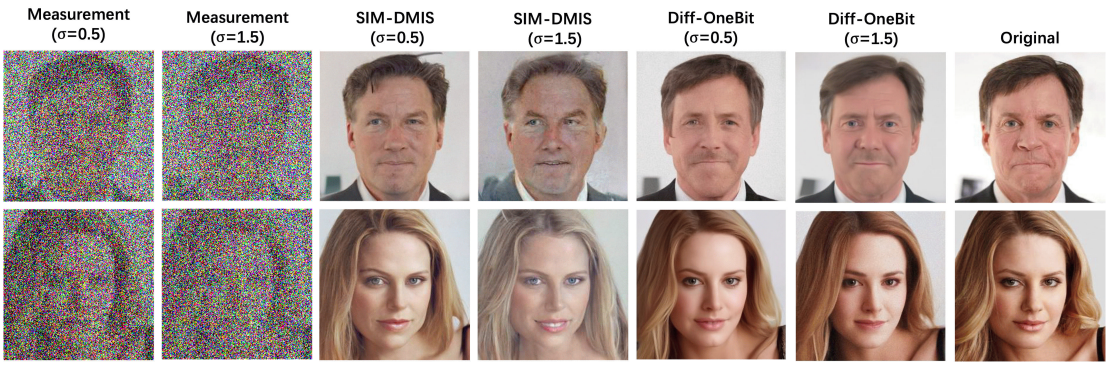

Figure 2. Qualitative results of 1-bit CS on CelebA images. We compare Diff-OneBit with SIM-DMIS. Input images have an additive Gaussian noise of $\sigma = 0.5$ and $\sigma = 1.5$.

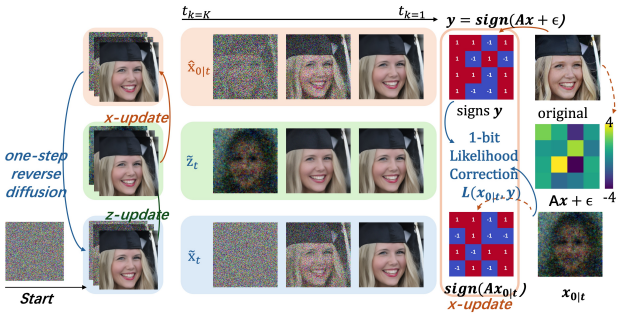

Figure 3. An illustration of the Diff-OneBit sampling step. At each reverse diffusion time t, we start with a noisy sample $\tilde{\mathbf{x}}_t$. First, the diffusion model acts as a denoiser to predict a preliminary clean image $\tilde{\mathbf{z}}_t$ (the prior update). This estimate is then corrected to a guided version $\hat{\mathbf{x}}_{0|t}$ by solving a data-consistency sub-problem using our differentiable surrogate likelihood. Finally, this guided estimate is used to compute the next state $\tilde{\mathbf{x}}_{t-1}$, advancing one step in the guided reverse diffusion.

Algorithm 1: Diff-OneBit

Require: Measurements $\mathbf{y}$, forward matrix $\mathbf{A}$, pretrained noise prediction network $\boldsymbol{\epsilon}_{\boldsymbol{\theta}}$, maximum number of iterations K, data-fidelity function $\mathcal{L}(\cdot; \mathbf{y})$, schedules $\{\alpha_{t_k}, \sigma_{t_k}\}_{k=0}^{K}$, penalty coefficient λ, pre-calculated $\mu_k = \lambda(\alpha_{t_k}^2/\sigma_{t_k}^2)$

Ensure: Reconstructed signal $\hat{\mathbf{x}}_0$

Sample $\tilde{\mathbf{x}}_{t_K} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $k = K, \ldots, 1$ **do**
  $\tilde{\mathbf{z}}_{t_k} \leftarrow (\tilde{\mathbf{x}}_{t_k} - \sigma_{t_k}\boldsymbol{\epsilon}_\theta(\tilde{\mathbf{x}}_{t_k}, t_k))/\alpha_{t_k}$
  $\hat{\mathbf{x}}_{0|t_k} \leftarrow \arg\min_{\mathbf{x}} \left(\mathcal{L}(\mathbf{x}; \mathbf{y}) + \frac{\mu_k}{2}\|\mathbf{x} - \tilde{\mathbf{z}}_{t_k}\|_2^2\right)$
  $\tilde{\boldsymbol{\epsilon}}_{t_k} \leftarrow (\tilde{\mathbf{x}}_{t_k} - \alpha_{t_k}\hat{\mathbf{x}}_{0|t_k})/\sigma_{t_k}$
  $\tilde{\mathbf{x}}_{t_{k-1}} \leftarrow \alpha_{t_{k-1}}\hat{\mathbf{x}}_{0|t_k} + \sigma_{t_{k-1}}\tilde{\boldsymbol{\epsilon}}_{t_k}$
end for
return $\tilde{\mathbf{x}}_{t_0}$

Experiments

In this section, we demonstrate the performance of our Diff-OneBit approach on 1-bit CS and logistic regression tasks.

Experimental setup. Our experiments employ pretrained DMs that correspond to image datasets at a resolution of 256×256. Specifically, for FFHQ dataset (Karras, Laine, and Aila 2019), we utilized the score model (Song and Ermon 2020) for QCS-SGM and the denoising model provided by (Chung et al. 2023) for Diff-OneBit, DPS, DAPS, DiffPIR and SIM-DMIS. For CelebA (Karras et al. 2018) datasets, we used the models pretrained by (Ho, Jain, and Abbeel 2020).[1] For ImageNet (Deng et al. 2009), we utilized the models pretrained by (Dhariwal and Nichol 2021). All the experiments are conducted on NVIDIA Tesla V100.

Measurement tasks: In all experiments, we evaluate our method on the validation set of 100 images on 1-bit CS and logistic regression. We have a forward matrix with i.i.d. entries drawn from a normal distribution $\mathcal{N}(0, 1/M)$. Measurements are evaluated under Gaussian noises with a measurement ratio of $M/N = 1/16$.

Baselines and metrics. We compare our proposed Diff-OneBit approach with several state-of-the-art methods. The primary baselines are specialized solvers for quantized measurements, including QCS-SGM and SIM-DMIS. Additionally, we include DiffPIR, DPS, and DAPS, that can handle non-linear measurement models, in our comparison. However, their non-linear frameworks are incompatible with 1-bit quantization, as the link function is non-differentiable or cannot be explicitly expressed. Thus, for our comparison, we use their linear versions, ignoring the non-linear link function. Reconstruction quality is evaluated using a comprehensive set of metrics. For fidelity and faithfulness to the ground truth, we use the peak signal-to-noise ratio (PSNR) and the structural similarity index measure (SSIM). For perceptual quality, we use the Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018) score and the Fréchet Inception Distance (FID) (Heusel et al. 2017).

[1]Since QCS-SGM does not provide a configuration file for reproducing results on the CelebA dataset, and we have not yet trained a score model on CelebA, we only tested the performance of QCS-SGM on the FFHQ dataset.

FFHQ		1-bit CS ($\sigma = 0.5$)				1-bit CS ($\sigma = 1.5$)			
Method	NFE	PSNR↑	SSIM↑	LPIPS↓	FID↓	PSNR↑	SSIM↑	LPIPS↓	FID↓
Diff-OneBit	20	**22.05** $\pm$ 1.88	**0.64** $\pm$ .07	**0.34** $\pm$ .05	**85.40**	**20.25** $\pm$ 1.01	**0.53** $\pm$ .06	**0.44** $\pm$ .05	**96.97**
SIM-DMIS	150	18.97 $\pm$ 2.36	0.54 $\pm$ .09	0.42 $\pm$ .05	110.05	14.32 $\pm$ 1.85	0.41 $\pm$ .08	0.56 $\pm$ .06	159.38
DiffPIR	100	9.88 $\pm$ 1.41	0.13 $\pm$ .04	0.78 $\pm$ .05	219.24	9.79 $\pm$ 1.39	0.12 $\pm$ .04	0.78 $\pm$ .05	226.36
DPS	1000	12.37 $\pm$ 1.70	0.31 $\pm$ .08	0.64 $\pm$ .07	132.84	11.68 $\pm$ 1.47	0.29 $\pm$ .08	0.65 $\pm$ .06	133.02
DAPS	1000	11.30 $\pm$ 1.60	0.37 $\pm$ .08	0.69 $\pm$ .05	203.12	11.39 $\pm$ 1.54	0.34 $\pm$ .08	0.73 $\pm$ .05	224.04
QCS-SGM	11555	18.88 $\pm$ 2.30	0.52 $\pm$ .10	0.55 $\pm$ .06	211.85	16.84 $\pm$ 0.98	0.47 $\pm$ .08	0.64 $\pm$ .06	297.27
CelebA		1-bit CS ($\sigma = 0.5$)				1-bit CS ($\sigma = 1.5$)			
Method	NFE	PSNR↑	SSIM↑	LPIPS↓	FID↓	PSNR↑	SSIM↑	LPIPS↓	FID↓
Diff-OneBit	20	**22.48** $\pm$ 1.80	**0.66** $\pm$ .09	**0.31** $\pm$ .05	**66.21**	**20.75** $\pm$ 1.03	**0.55** $\pm$ .07	**0.40** $\pm$ .04	**71.81**
SIM-DMIS	150	18.88 $\pm$ 2.50	0.56 $\pm$ .09	0.38 $\pm$ .05	75.93	14.78 $\pm$ 2.36	0.43 $\pm$ .09	0.51 $\pm$ .06	102.75
DiffPIR	100	9.20 $\pm$ 1.51	0.13 $\pm$ .05	0.81 $\pm$ .10	184.89	9.12 $\pm$ 1.49	0.12 $\pm$ .05	0.82 $\pm$ .10	185.24
DPS	1000	12.94 $\pm$ 1.82	0.36 $\pm$ .09	0.56 $\pm$ .07	97.01	11.78 $\pm$ 1.71	0.31 $\pm$ .08	0.60 $\pm$ .06	105.19
DAPS	1000	11.33 $\pm$ 1.62	0.38 $\pm$ .08	0.67 $\pm$ .05	186.51	11.32 $\pm$ 1.53	0.35 $\pm$ .08	0.72 $\pm$ .05	211.46
ImageNet		1-bit CS ($\sigma = 0.5$)				1-bit CS ($\sigma = 1.5$)			
Method	NFE	PSNR↑	SSIM↑	LPIPS↓	FID↓	PSNR↑	SSIM↑	LPIPS↓	FID↓
Diff-OneBit	20	**20.56** $\pm$ 2.34	**0.52** $\pm$ .16	**0.47** $\pm$ .09	**92.71**	**18.23** $\pm$ 1.72	**0.43** $\pm$ .15	**0.55** $\pm$ .07	**121.81**
SIM-DMIS	150	17.37 $\pm$ 2.68	0.41 $\pm$ .15	0.50 $\pm$ .08	95.72	14.87 $\pm$ 2.79	0.30 $\pm$ .15	0.61 $\pm$ .09	138.37

Table 1. Quantitative results of 1-bit CS on FFHQ, CelebA and ImageNet images. We compare our method with DiffPIR, DPS, DAPS QCS-SGM and SIM-DMIS. Input images have an additive Gaussian noise of $\sigma = 0.5$ and $\sigma = 1.5$.

NFE	PSNR↑	SSIM↑	LPIPS↓
20	22.05 $\pm$ 1.88	0.64 $\pm$ .07	0.34 $\pm$ .05
50	22.39 $\pm$ 1.34	0.65 $\pm$ .06	0.33 $\pm$ .05
100	22.65 $\pm$ 1.08	0.65 $\pm$ .06	0.32 $\pm$ .05

Table 2. Quantitative results of 1-bit CS on FFHQ images under different NFEs. We evaluate our method under 20, 50 and 100 NFEs with additive Gaussian noise of $\sigma = 0.5$.

Quantitative results

For 1-bit CS, the main quantitative results for the FFHQ, CelebA and ImageNet datasets are presented in Table 1. Our proposed method, Diff-OneBit, consistently outperforms all baseline methods across all evaluation metrics, demonstrating its superior reconstruction accuracy. For logistic regression, the main quantitative results on the FFHQ dataset are presented in Table 3. We compare our method with DPS, DAPS, and DiffPIR, which are capable of handling non-linear measurements. We also compare with SIM-DMIS, as logistic regression is a special case of SIM.

Our experiments demonstrate that existing solvers for linear or non-linear measurement models perform poorly in 1-bit quantization tasks. Tables 1 and 3 show that DiffPIR, DPS, and DAPS significantly underperforms compared to specialized methods and our proposed Diff-OneBit approach. These results demonstrate that optimization and gradient based methods, designed for linear or non-linear problems, are ineffective for 1-bit measurement models.

FFHQ	logistic regression		
Method	PSNR↑	SSIM↑	LPIPS↓
Diff-OneBit	**19.99** $\pm$ 1.94	**0.56** $\pm$ .08	**0.42** $\pm$ .05
SIM-DMIS	16.03 $\pm$ 2.34	0.46 $\pm$ .09	0.49 $\pm$ .06
DiffPIR	9.84 $\pm$ 1.40	0.13 $\pm$ .04	0.79 $\pm$ .05
DPS	15.03 $\pm$ 2.37	0.43 $\pm$ .11	0.53 $\pm$ .08
DAPS	13.43 $\pm$ 1.64	0.40 $\pm$ .08	0.63 $\pm$ .07

Table 3. Quantitative results of logistic regression on FFHQ images. We compare Diff-OneBit with DAPS, DPS, DiffPIR, and SIM-DMIS without additional Gaussian noise, setting NFE to 20 for Diff-OneBit, 100 for DiffPIR, 150 for SIM-DMIS, and 1000 for DPS and DAPS.

Qualitative results

For 1-bit CS, Figures 1 and 2 provide visual comparison of the reconstruction results on FFHQ and CelebA images for the 1-bit CS.[2] Images reconstructed by Diff-OneBit exhibit significantly fewer artifacts and preserve finer textural details and structural integrity. In contrast, specialized solvers like QCS-SGM produce reconstructions with noticeable smoothing, while SIM-DMIS results in loss of high-frequency details. DiffPIR, DPS, and DAPS fail to capture the semantic content of images from 1-bit measurements using gradient information. For logistic regression, Figure 6 provides visual comparison of the reconstruction results on FFHQ images for logistic regression. Images reconstructed by our Diff-OneBit approach exhibit significantly fewer ar-

[2]We use $\mathbf{A}^\top \mathbf{y}$ to denote the measurement, ensuring dimensional consistency.

Figure 4. Qualitative results of 1-bit CS on FFHQ images. We compare Diff-OneBit under varying Gaussian noise levels in 1-bit CS.

Figure 5. Qualitative results of logistic regression on FFHQ images. We compare Diff-OneBit under varying Gaussian noise levels in logistic regression.

tifacts and preserve finer textural details and structural integrity.

Moreover, Figures 4 and 5 provide a visual comparison of reconstructed FFHQ images for 1-bit CS and logistic regression based on the Diff-OneBit method under varying Gaussian noise levels, specifically with $\sigma \in \{0.5, 1.0, 1.5, 2.0\}$ for 1-bit CS and $\sigma \in \{0.0, 0.5, 1.0, 1.5\}$ for logistic regression. These figures demonstrate that our approach consistently achieves high-fidelity image recovery across different noise intensities, showing robust performance in preserving details, for which detailed quantitative results are provided in the appendix.

Ablation studies and further analysis

To validate the design choices and flexibility of our method, we conduct further analysis.

Experiment under varying NFEs. We evaluate our method by varying NFEs from 20 to 100 for 1-bit CS on 100 FFHQ images with Gaussian noise $\sigma = 0.5$. Table 2 demonstrates that our method is efficient even at low NFEs. Moreover, the reconstruction quality further improves as NFE increases from 20 to 100, showing the scalability of our method.

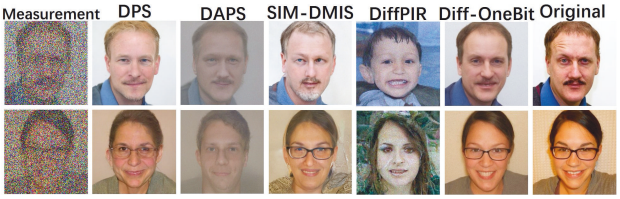

Figure 6. Qualitative results of logistic regression on FFHQ images. We compare Diff-OneBit with DPS, DAPS, DiffPIR and SIM-DMIS.

Figure 7. Effect of hyperparameters λ. Reconstructed images from the same 1-bit CS measurement with input images corrupted by additive Gaussian noise of $\sigma = 0.5$.

Comparison with DiffPIR. Our method, similar to DiffPIR, employs the HQS approach but differs in two major aspects. First, DiffPIR lacks specific modeling for 1-bit measurements, resulting in inferior reconstruction, as shown in our experiments. Then, unlike DiffPIR, which adds Gaussian noise during sampling, our algorithm performs better without random noise. The impact of varying random noise levels is presented in the appendix. We adopt similar hyperparameters, with λ controlling the data-fidelity term via $\mu_t = \lambda(\alpha_t^2/\sigma_t^2)$. Figure 7 demonstrates that our method, Diff-OneBit, exhibits robust performance across a range of λ values on the FFHQ dataset. However, overly strong ($\lambda < 0.0001$) or weak ($\lambda \geq 0.1$) guidance causes noisy or overly smooth reconstructions, respectively.

Computational efficiency. Our algorithm achieves high-quality reconstructions using only 20 NFEs. However, the gradient descent optimization introduces additional computational overhead. To ensure a fair performance comparison, we provide a detailed analysis of the time overhead in the appendix.

Conclusion

In this paper, we address the challenge of signal recovery under 1-bit quantization in the context of 1-bit compressed sensing and logistic regression using diffusion models. We propose Diff-OneBit, an approach that integrates a differentiable surrogate likelihood function with a plug-and-play optimization scheme, leveraging a pretrained diffusion model. Experimental results demonstrate that Diff-OneBit achieves state-of-the-art reconstruction quality on the FFHQ and CelebA datasets, offering high computational efficiency and versatility across 1-bit quantized measurement models.

References

Asim, M.; Daniels, M.; Leong, O.; Ahmed, A.; and Hand, P. 2020. Invertible generative models for inverse problems: Mitigating representation error and dataset bias. In *ICML*.

Bora, A.; Jalal, A.; Price, E.; and Dimakis, A. G. 2017. Compressed sensing using generative models. In *ICML*.

Boufounos, P. T. 2010. Reconstruction of sparse signals from distorted randomized measurements. In *ICASSP*.

Boufounos, P. T.; and Baraniuk, R. G. 2008. 1-bit compressive sensing. In *CISS*.

Candès, E. J.; Romberg, J.; and Tao, T. 2006. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2): 489–509.

Chan, S. H.; Wang, X.; and Elgendy, O. A. 2016. Plug-and-play ADMM for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1): 84–98.

Chen, J.; Ng, M. K.; and Liu, Z. 2025. Solving Quadratic Systems With Full-Rank Matrices Using Sparse or Generative Priors. *IEEE Transactions on Signal Processing*.

Chen, J.; Scarlett, J.; Ng, M.; and Liu, Z. 2023. A unified framework for uniform signal recovery in nonlinear generative compressed sensing. *NeurIPS*.

Chen, Y.; Jalali, A.; Sanghavi, S.; and Caramanis, C. 2013. Low-rank matrix recovery from errors and erasures. *IEEE Transactions on Information Theory*, 59(7): 4324–4337.

Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models. In *ICCV*.

Chung, H.; Kim, J.; McCann, M. T.; Klasky, M. L.; and Ye, J. C. 2023. Diffusion Posterior Sampling for General Noisy Inverse Problems. In *ICLR*.

Chung, H.; Sim, B.; Ryu, D.; and Ye, J. C. 2022. Improving diffusion models for inverse problems using manifold constraints. In *NeurIPS*.

Daras, G.; Chung, H.; Lai, C.-H.; Mitsufuji, Y.; Ye, J. C.; Milanfar, P.; Dimakis, A. G.; and Delbracio, M. 2024. A survey on diffusion models for inverse problems. *arXiv preprint arXiv:2410.00083*.

Davenport, M. A.; Plan, Y.; Van Den Berg, E.; and Wootters, M. 2014. 1-bit matrix completion. *Information and Inference: A Journal of the IMA*, 3(3): 189–223.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *CVPR*.

Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat GANs on image synthesis. In *NeurIPS*.

Donoho, D. L. 2006. Compressed sensing. *IEEE Transactions on information theory*, 52(4): 1289–1306.

Foucart, S.; and Rauhut, H. 2013. An invitation to compressive sensing. In *A mathematical introduction to compressive sensing*, 1–39.

Geman, D.; and Yang, C. 1995. Nonlinear image recovery with half-quadratic regularization. *IEEE Transactions on Image Processing*, 4(7): 932–946.

Genzel, M.; Macdonald, J.; and März, M. 2022. Solving inverse problems with deep neural networks–robustness included? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1): 1119–1134.

Hand, P.; Leong, O.; and Voroninski, V. 2018. Phase retrieval under a generative prior. In *NeurIPS*.

Heckel, R.; and Hand, P. 2019. Deep Decoder: Concise Image Representations from Untrained Non-convolutional Networks. In *ICLR*.

Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. *NeurIPS*.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *NeurIPS*.

Jacques, L.; Laska, J. N.; Boufounos, P. T.; and Baraniuk, R. G. 2013. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Transactions on Information Theory*, 59(4): 2082–2102.

Jalal, A.; Arvinte, M.; Daras, G.; Price, E.; Dimakis, A. G.; and Tamir, J. 2021a. Robust compressed sensing MRI with deep generative priors. In *NeurIPS*.

Jalal, A.; Karmalkar, S.; Dimakis, A.; and Price, E. 2021b. Instance-Optimal Compressed Sensing via Posterior Sampling. In *ICML*.

Kadkhodaie, Z.; and Simoncelli, E. 2021. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. In *NeurIPS*.

Kamilov, U. S.; Mansour, H.; and Wohlberg, B. 2017. A plug-and-play priors approach for solving nonlinear imaging inverse problems. *IEEE Signal Processing Letters*, 24(12): 1872–1876.

Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*.

Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. In *NeurIPS*.

Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *CVPR*.

Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising diffusion restoration models. *NeurIPS*.

Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Knudson, K.; Saab, R.; and Ward, R. 2016. One-bit compressive sensing with norm estimation. *IEEE Transactions on Information Theory*, 62(5): 2748–2758.

Li, X. 2013. Compressed sensing and matrix completion with constant proportion of corruptions. *Constructive Approximation*, 37(1): 73–99.

Liu, J.; and Liu, Z. 2022. Non-iterative recovery from nonlinear observations using generative models. In *CVPR*.

Liu, Z.; Ghosh, S.; and Scarlett, J. 2021. Towards Sample-Optimal Compressive Phase Retrieval with Sparse and Generative Priors. In *NeurIPS*.

Liu, Z.; Gomes, S.; Tiwari, A.; and Scarlett, J. 2020. Sample complexity bounds for 1-bit compressive sensing and binary stable embeddings with generative priors. In *ICML*.

Liu, Z.; and Han, J. 2022. Projected gradient descent algorithms for solving nonlinear inverse problems with generative priors. In *IJCAI*.

Liu, Z.; Li, W.; and Chen, J. 2024. Generalized Eigenvalue Problems with Generative Priors. In *NeurIPS*.

Liu, Z.; Liu, J.; Ghosh, S.; Han, J.; and Scarlett, J. 2022. Generative Principal Component Analysis. In *ICLR*.

Liu, Z.; and Scarlett, J. 2020a. The generalized Lasso with nonlinear observations and generative priors. *NeurIPS*.

Liu, Z.; and Scarlett, J. 2020b. Information-theoretic lower bounds for compressive sensing with generative models. *IEEE Journal on Selected Areas in Information Theory*, 1(1): 292–303.

Liu, Z.; Wang, X.; and Liu, J. 2022. Misspecified phase retrieval with generative priors. In *NeurIPS*.

Mardani, M.; Song, J.; Kautz, J.; and Vahdat, A. 2024. A Variational Perspective on Solving Inverse Problems with Diffusion Models. In *ICLR*.

Meng, X.; and Kabashima, Y. 2023. Quantized Compressed Sensing with Score-Based Generative Models. In *ICLR*.

Nguyen, N. H.; and Tran, T. D. 2012. Robust Lasso with missing and grossly corrupted observations. *IEEE Transactions on Information Theory*, 59(4): 2036–2058.

Nguyen, N. H.; and Tran, T. D. 2013. Exact Recoverability From Dense Corrupted Observations via ℓ_1-Minimization. *IEEE Transactions on Information Theory*, 59(4): 2017–2035.

Nguyen, T. V.; Jagatap, G.; and Hegde, C. 2021. Provable Compressed Sensing with Generative Priors via Langevin Dynamics. *https://arxiv.org/2102.12643*.

Ongie, G.; Jalal, A.; Metzler, C. A.; et al. 2020. Deep learning techniques for inverse problems in imaging. *IEEE J. Sel. Areas Inf. Theory*, 1(1): 39–56.

Plan, Y.; and Vershynin, R. 2012. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1): 482–494.

Plan, Y.; and Vershynin, R. 2013. One-bit compressed sensing by linear programming. *Communications on Pure and Applied Mathematics*, 66(8): 1275–1297.

Plan, Y.; Vershynin, R.; and Yudovina, E. 2017. High-dimensional estimation with geometric constraints. *Information and Inference: A Journal of the IMA*, 6(1): 1–40.

Qiu, S.; Wei, X.; and Yang, Z. 2020. Robust one-bit recovery via ReLU generative networks: Near-optimal statistical rate and global landscape analysis. In *ICML*.

Scarlett, J.; Heckel, R.; Rodrigues, M. R.; Hand, P.; and Eldar, Y. C. 2023. Theoretical perspectives on deep learning methods in inverse problems. *IEEE J. Sel. Areas Inf. Theory*, 3(3): 433–453.

Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*.

Song, B.; Kwon, S. M.; Zhang, Z.; Hu, X.; Qu, Q.; and Shen, L. 2024. Solving Inverse Problems with Latent Diffusion Models via Hard Data Consistency. In *ICLR*.

Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *ICLR*.

Song, J.; Vahdat, A.; Mardani, M.; and Kautz, J. 2023. Pseudoinverse-guided diffusion models for inverse problems. In *ICLR*.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*.

Song, Y.; and Ermon, S. 2020. Improved techniques for training score-based generative models. In *NeurIPS*.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *ICLR*.

Sreehari, S.; Venkatakrishnan, S. V.; Wohlberg, B.; Buzzard, G. T.; Drummy, L. F.; Simmons, J. P.; and Bouman, C. A. 2016. Plug-and-play priors for bright field electron tomography and sparse interpolation. *IEEE Transactions on Computational Imaging*, 2(4): 408–423.

Tang, A.; Chen, Y.; Xue, S.; and Liu, Z. 2025. Learning Single Index Models with Diffusion Priors. In *ICML*.

Van Veen, D.; Jalal, A.; Soltanolkotabi, M.; et al. 2018. Compressed sensing with deep image prior and learned regularization. *https://arxiv.org/1806.06438*.

Venkatakrishnan, S. V.; Bouman, C. A.; and Wohlberg, B. 2013. Plug-and-play priors for model based reconstruction. In *GlobalSIP*.

Wang, Y.; Yu, J.; and Zhang, J. 2023. Zero-Shot Image Restoration Using Denoising Diffusion Null-Space Model. In *ICLR*.

Wei, X.; Yang, Z.; and Wang, Z. 2019. On the statistical rate of nonlinear recovery in generative models with heavy-tailed data. In *ICML*.

Xu, H.; Caramanis, C.; and Mannor, S. 2012. Outlier-robust PCA: The high-dimensional case. *IEEE Transactions on Information Theory*, 59(1): 546–572.

Yan, M.; Yang, Y.; and Osher, S. 2012. Robust 1-bit compressive sensing using adaptive outlier pursuit. *IEEE Transactions on Signal Processing*, 60(7): 3868–3875.

Zhang, B.; Chu, W.; Berner, J.; Meng, C.; Anandkumar, A.; and Song, Y. 2025. Improving diffusion inverse problem solving with decoupled noise annealing. In *CVPR*.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*.

Zheng, Y.; Li, W.; and Liu, Z. 2025. Integrating Intermediate Layer Optimization and Projected Gradient Descent for Solving Inverse Problems with Diffusion Models. In *ICML*.

Zhu, Y.; Zhang, K.; Liang, J.; Cao, J.; Wen, B.; Timofte, R.; and Van Gool, L. 2023. Denoising diffusion models for plug-and-play image restoration. In *CVPRW*.