Stochastic Optimal Control Sampling for Diffusion Inverse Problems

Jie Zhang[1], Youmei Qiu[1], Hanling Tian[1], Jingyuan Zhang[1], Xiang Yin[1], and Xiaolin Huang[1,2⋆]

[1] School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai, China
[2] Shanghai Key Laboratory of Flexible Medical Robotics, Tongren Hospital, Institute of Medical Robotics, Shanghai Jiao Tong University, Shanghai, China
{zhangjie20000411,qiuyoumei,hanlingtian,tonyzhang666,yinxiang, xiaolinhuang}@sjtu.edu.cn

Abstract. Benefiting from the strong ability to capture data distributions, diffusion models have become powerful tools for solving image inverse problems. The key is to controllably steer the sampling trajectory toward the measurements while respecting the diffusion prior. In this work, we introduce Stochastic Optimal Control Sampling (SOCS), which models the denoising process as a dynamical system and injects control signals via SOC. Previous SOC-based approach addresses inverse problems by optimizing over the entire trajectory, which is computationally expensive. In contrast, we derive a closed-form control update and apply it at each sampling step, pulling the measurement-consistent clean prediction back onto the denoising flow. In SOCS, we can readily modulate the control strength to align with the diffusion model's native capabilities and thereby enhance perceptual quality. Our method is compatible with a variety of linear stochastic differential equation backbones. Extensive experiments across a broad spectrum of image inverse tasks demonstrate that SOCS achieves accurate measurement-aligned reconstructions with improved visual fidelity and stronger quantitative performance. Code is available at `https://github.com/zjqwq01/SOCS-DIP`.

Keywords: Diffusion model · Inverse problem · Stochastic optimal control

1 Introduction

Image inverse problems [10,11,26] are pervasive in science and engineering, from image super-resolution to more challenging nonlinear degradations. The goal is to reconstruct the ground-truth image from degraded measurements. The central difficulty stems from a degradation operator that is many-to-one and often nonlinear, which renders the solution inherently non-unique. Effective inversion therefore requires strong priors that should be leveraged efficiently. From this perspective, diffusion models [13–15, 33–35] are especially promising for inverse

⋆ Corresponding Author.

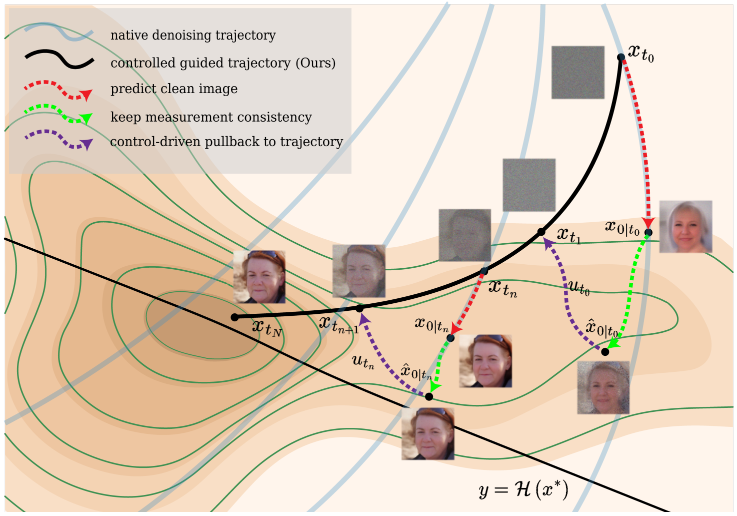

Fig. 1: Illustration of our method (an example on a random inpainting task). We incorporate stochastic optimal control into sampling so that, at every step, the optimal control policy pulls the measurement-consistent clean estimate back onto the denoising trajectory, preserving fidelity while honoring the diffusion prior.

problems. They learn expressive generative priors that capture high-order statistics and geometric structure of natural images, and enable principled posterior sampling by conditioning score-based dynamics on measurements. As a result, diffusion models have become compelling foundations for inverse problems.

To generate images that are both measurement consistent and remarkably faithful, one way is to retrain diffusion models for a specific inverse task [30, 31, 41], which however requires massive data and incurs heavy computational cost. A more efficient way is to inject measurement information into the denoising process of a pretrained diffusion model. The current mainstream of such training-free methods lies within a Bayesian sampling framework, see, e.g., [5,32]. They approximate the conditional score by replacing it with a simple function of the noise-free data likelihood and use such approximation to steer denoising toward the measurements. Enforcing measurement consistency along the trajectory naturally calls for controlling the diffusion process to satisfy the measurements. To design such a control signal in this direction, [1] established the foundational theoretical connection between diffusion processes and stochastic optimal control (SOC) and later [20] applies SOC to diffusion inverse problems by designing a traditional iLQR controller [21,36], which optimizes the entire diffusion path via iterative rollouts; consequently, it requires second-order Taylor updates, which demand Hessians and lead to high memory usage and long runtimes.

In this work, we introduce a new framework to apply SOC in diffusion inverse problems, termed **S**tochastic **O**ptimal **C**ontrol **S**ampling (SOCS). Instead of relying on iLQR-style trajectory controllers, we integrate SOC into the stepwise stochastic routine of the diffusion model, which avoids costly higher-order computations and enables the use of a pretrained unconditional diffusion prior without retraining. As shown in Fig. 1, at each iteration we first obtain a preliminary clean estimate by advancing the reverse-diffusion probability-flow ODE with a few Euler steps, second enforce consistency of this estimate with the measurement, and third map the estimate back to the denoising manifold through a closed-form transport. Our method rests on a natural intuition that an uncontrolled diffusion prior already produces high-quality samples, and excessive control will drive the process away from its native manifold, thereby undermining this ability. In light of this, we adopt controllers with bounded amplitude and provide a series of theoretical foundations to justify this design. Accordingly, SOCS not only penalizes the terminal state but also deliberately and gently modulates the amplitude of intermediate information injection, aligning the denoising process with the model's intrinsic generative capacity. Moreover, our method requires at most first-order Jacobians, yielding wall-clock times comparable to sampling-based baselines while preserving the benefits of SOC guidance.

Empirically, our method delivers substantial gains over existing approaches across a broad spectrum of inverse problems and applies to all linear stochastic differential equation (SDE) formulations [33], including the classical variance-exploding SDE (VE-SDE) and variance-preserving SDE (VP-SDE). Quantitative experiments show that tuning the strength of the injected control markedly improves generative performance. We further extend SOCS to various Latent Diffusion Models (LDMs), including traditional diffusion-based frameworks and state-of-the-art Flow-Matching (FM) variants, demonstrating its compatibility with high-resolution image restoration across different generative paradigms.

2 Background

2.1 Diffusion Models

Diffusion models [13, 14, 33–35] learn a generative prior by inverting a gradual noising process. Let $x_0 \sim p_{\text{data}}$ and choose a noise schedule σ_t that increases from $\sigma_0 = 0$ to $\sigma_T = \sigma_{\max}$, so the forward marginals $p(x_t) \equiv p(x; \sigma_t)$ are Gaussian–smoothed versions of p_{data}. In the SDE formulation, the forward dynamics can be written as $\mathrm{d}x_t = f_t(x_t)\mathrm{d}t + g_t \mathrm{d}w_t$. With f_t the drift and g_t the diffusion scale (e.g., VP or VE choices), the reverse-time SDE relies on the score $\nabla_x \log p(x_t; \sigma_t)$ to correct the trajectory back to the data distribution. Generation transports $x_T \sim \mathcal{N}(0, \sigma_T^2 I)$ back to data by integrating the reverse dynamics, where the score is approximated by a network $s_\theta(x_t, \sigma_t) \approx \nabla_x \log p(x_t; \sigma_t)$. An equivalent training view predicts the injected noise under $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t}\varepsilon$, where $\bar{\alpha}_t$ is the cumulative decay factor and $\varepsilon \sim \mathcal{N}(0, I)$, minimizing $\mathbb{E}_{x_0, \varepsilon, t}\left[w(t)\|\varepsilon - \varepsilon_\theta(x_t, t)\|^2\right]$, which yields $s_\theta(x_t, t) \approx -\varepsilon_\theta(x_t, t)/\sqrt{1 - \bar{\alpha}_t}$. Alternative output formulations that predict clean image x_0 or the noise ϵ are algebraically equivalent and map to the same reverse process via simple reparameterizations that are invertible.

2.2 Bayesian Inverse Problems with Diffusion

Inverse problems seek to recover a clean signal from degraded measurements. Let the degradation process be modeled as $y = \mathcal{H}(x_0) + n$, where $\mathcal{H}$ is the degradation operator, x_0 is the unknown clean data, y is the measurement, and $n \sim \mathcal{N}(0, \beta_y^2 I)$ is additive Gaussian noise. In a Bayesian formulation that treats uncertainty explicitly, the likelihood is $p(y \mid x_0) = \mathcal{N}(\mathcal{H}(x_0), \beta_y^2 I)$, and solving the inverse problem amounts to sampling from the posterior $p(x_0 \mid y)$.

Given access to the measurement y, by Bayes' rule the conditional score can be written as:

$$\nabla_x \log p_t(x_t \mid y) \;=\; \nabla_x \log p_t(x_t) \;+\; \nabla_x \log p_t(y \mid x_t).$$

Here, the term $\nabla_x \log p_t(x_t)$ is readily available from a pretrained diffusion model, whereas the noise-likelihood term $\nabla_x \log p_t(y \mid x_t)$ is generally intractable. A variety of methods [2, 5, 23, 42, 45] have been proposed to estimate or circumvent this term. DPS [5] approximates the likelihood via the posterior mean of the clean signal, replacing $p(y \mid x_t)$ with $p\big(y \mid \hat{x}_0\big)$, where $\hat{x}_0 := \mathbb{E}[x_0 \mid x_t]$. Building on this perspective, [23] argues that the recursive nature of denoising induces a complex, multi-modal posterior that is difficult to approximate with simple forms, and therefore proposes a variational-inference approach to infer the posterior of the data given observations. DAPS [42] mitigates the step-size dependence of conventional denoising by decoupling the continuous process, which enables later steps to effectively correct errors made earlier. In practice, sampling methods developed under the Bayesian framework commonly employ a guidance term of the form $\nabla_{x_t} \|\mathcal{H}(x_t) - y\|_2^2$. Without approximating the noise likelihood term, we jointly model this term together with the model prior as a single control signal and modulate it in a unified way. This yields a guidance behavior similar to those approaches and, from the SOC perspective, establishes the necessity of this term.

2.3 Diffusion with Stochastic Optimal Control

Integrating stochastic optimal control (SOC) into diffusion-model sampling has emerged as a promising guidance paradigm. DIS [1] first established a foundational link between diffusion processes and SOC, motivating practical formulations. Building on this theory, RB-Modulation [28] applies SOC to diffusion-style personalization, framing it as an optimal control problem. Relevant to our setting, UniDB [44] formulates the forward noising of diffusion bridges within an SOC framework on image restoration tasks and introduces a simple controller linked to Doob's h-transform that shows strong detail preservation. However, it requires retraining whenever terminal penalties or other settings change. [20] solves inverse problems by formulating reverse diffusion as an optimal-control problem and derives a practical controller that improves stability and data consistency across operators, although computing and storing Hessian matrices impose substantial costs. To address these challenges, we model the denoising process to obviate retraining and integrate SOC at each step, thereby avoiding higher-order computations and substantially reducing runtime.

3 Method

3.1 Inverse Problems under an SOC Formulation

Inverse problems aim to recover the original signal from noisy measurements y. With a similar goal, SOCS formulates a stochastic optimal control (SOC) problem whose constraint is a reverse-diffusion linear SDE initialized from random noise, with drift $f(x_t, t) = f_t x_t$, where f_t is some scalar-valued function. To align with standard SOC modeling, we let $t = 0$ denote the pure-noise initialization and $t = T$ denote the final clean-image time. The objective steers the reverse denoising trajectory toward a prescribed terminal state by minimizing the squared Euclidean residual $\|\mathcal{H}(x_T) - y\|_2^2$.

Specifically, our SOC problem with $\mathcal{H}$ under linear SDE can be written as:

$$\begin{aligned} &\min_{\mathbf{u}_{t,\gamma}\in\mathcal{U}} \mathbb{E}\left[\int_0^T \tfrac{1}{2}\left\|\mathbf{u}_{t,\gamma}\right\|_2^2 \mathrm{d}t + \tfrac{\gamma}{2}\left\|y - \mathcal{H}(x_T)\right\|_2^2\right] \\ &\text{s.t.} \quad \mathrm{d}x_t = (f_t x_t + g_t \mathbf{u}_{t,\gamma})\,\mathrm{d}t + g_t\,\mathrm{d}w_t. \end{aligned} \tag{1}$$

Interpretation (prior proximity via control energy). By Girsanov's theorem, the control energy upper-bounds the KL deviation of the controlled terminal distribution from the diffusion prior: $D_{\mathrm{KL}}(Q_T\|P_T) \leq \frac{1}{2}\,\mathbb{E}\Big[\int_0^T \|u_{t,\gamma}\|_2^2\,\mathrm{d}t\Big]$. A detailed proof can be found in Appendix.
Interpretation (posterior sampling). The SOC objective in Eq. (1) induces a Gibbs reweighting of the diffusion prior on path space, $\frac{\mathrm{d}Q^*_{\mathrm{path}}}{\mathrm{d}P_{\mathrm{path}}} \propto \exp\big(-\phi(x_T)\big)$, $\phi(x_T) = \frac{\gamma}{2}\|y - \mathcal{H}(x_T)\|_2^2$, hence $q_T^*(x) \propto p_T(x)\exp(-\phi(x))$. Under Gaussian noise, $\phi(x)$ is proportional to the negative log-likelihood, so q_T admits a posterior-sampling interpretation. More details are provided in AppendiX.
Theorem 3.1. Under linear dynamics, additive zero-mean control-independent noise, and a quadratic terminal cost (or a Gauss–Newton linearization of $\mathcal{H}$), the certainty equivalence principle [3,28] implies that the optimal controller remains the same. Thus, we obtain the following SOC problem with deterministic ODE condition:

$$\begin{aligned} &\min_{\mathbf{u}_{t,\gamma}\in\mathcal{U}} \int_0^T \tfrac{1}{2}\left\|\mathbf{u}_{t,\gamma}\right\|_2^2 \mathrm{d}t + \tfrac{\gamma}{2}\left\|y - \mathcal{H}(x_T)\right\|_2^2 \\ &\text{s.t.} \quad \mathrm{d}x_t = (f_t x_t + g_t \mathbf{u}_{t,\gamma})\,\mathrm{d}t. \end{aligned} \tag{2}$$

A key observation is that we can derive a closed-form solution for the SOC problem Eq. (2). Denote $d_{0,\gamma}^{-1} := \big(I + \gamma\, e^{2\bar{f}_T}\,\bar{g}_T^2\, J_H(x_T^u)^\top \circ \mathcal{H}\big)^{-1}$, $\bar{f}_{s:t} := \int_s^t f_z \mathrm{d}z$ and $\bar{g}_{s:t}^2 := \int_s^t e^{-2\bar{f}_z} g_z^2 \mathrm{d}z$, we simplify notation to $\bar{f}_t$ and $\bar{g}_t^2$ for $\bar{f}_{0:t}$ and $\bar{g}_{0:t}^2$. Here $J_H(x_T^u)$ is the Jacobian of $\mathcal{H}(x)$ evaluated at x_T^u, and the symbol $\circ$ denotes operator composition. The optimal controller $\mathbf{u}^*_{t,\gamma}$ could be calculated:

$$\mathbf{u}^*_{t,\gamma} = -g_t \gamma e^{\bar{f}_{t:T}} d_{0,\gamma}^{-1} J_H(x_T^u)^\top \big(e^{\bar{f}_T}\mathcal{H}(x_0) - y\big). \tag{3}$$

Therefore, transition of x_t from x_0 is:

$$x_t = e^{\bar{f}_t} x_0 + \gamma e^{\bar{f}_t} e^{\bar{f}_T} \bar{g}_t^2 d_{0,\gamma}^{-1} J_H(x_T^u)^\top \big(y - e^{\bar{f}_T}\mathcal{H}(x_0)\big). \tag{4}$$

The detailed derivations of Eq. (3) and Eq. (4) are provided in Appendix. The above discussions hold for all linear SDEs, covering both the VE–SDE and VP–SDE. Existing Bayesian sampling methods [5, 42] inherently assume an infinite terminal penalty with γ driven to infinity. In practice, treating γ as a tunable hyperparameter can be more effective. In the next section, using the VP–SDE as a running case, we examine how different approximations to $\mathcal{H}$ and the choice of the terminal penalty affect Eq. (4) and present two main variants of the method.

3.2 Case Study: SOCS-VP-SDE

In this section, we apply our posterior-sampling pipeline to VP–SDE. In the VP–SDE, the forward noising dynamics are $\mathrm{d}x_t = -\frac{1}{2}\beta(t)x_t\mathrm{d}t + \sqrt{\beta(t)}\mathrm{d}w_t$, where the drift term $-\frac{1}{2}\beta(t)x_t$ linearly contracts the state toward the origin, and the diffusion term $\sqrt{\beta(t)}$ injects Gaussian noise under a time-varying schedule. We take $\beta(t) = g(t)^2$.

SOCS-VP-Nonlinear-γ. Eq. (4) makes clear that x_t consists of a baseline term $e^{\bar{f}_t}x_0$ together with an affine data-consistency correction. The latter is formed by back-projecting the measurement residual from measurement space to image space via the Jacobian adjoint $J_{\mathcal{H}}(x_T^u)^\top$, and modulating it with a time-varying gain. Using $\bar{g}_t^2 = e^{-2\bar{f}_t} - 1$, Eq. (4) can be equivalently recast as:

$$
\begin{aligned}
x_t &= e^{\bar{f}_t}x_0 + \gamma e^{\bar{f}_T}\big(e^{-\bar{f}_t} - e^{\bar{f}_t}\big)\hat{x}_T, \\
\hat{x}_T &\triangleq \Big(I + \gamma(1 - e^{2\bar{f}_T})\, J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H}\Big)^{-1} \circ J_{\mathcal{H}}(x_T^u)^\top\big(y - e^{\bar{f}_T}\mathcal{H}(x_0)\big).
\end{aligned} \tag{5}
$$

In practice, the Gauss–Newton inverse of the regularized normal matrix defined in Eq. (5) is typically challenging to handle. We use $\hat{x}_T$ to denote the Gauss–Newton update , and compute it via an implicit solve instead of an explicit matrix inverse. Equivalently, we solve $(I + \gamma(1 - e^{2\bar{f}_T})J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H})(\hat{x}_T) = J_{\mathcal{H}}(x_T^u)^\top(y - e^{\bar{f}_T}\mathcal{H}(x_0))$. Leveraging Langevin dynamics [39] sampling, we refine $\hat{x}_T$ via a stochastic-gradient step on the corresponding data-consistency energy:

$$
\begin{aligned}
\hat{x}_T^{j+1} =& \hat{x}_T^j - \eta\nabla_{\hat{x}_T^j}\big\|\big(I + \gamma(1 - e^{2\bar{f}_T})J_{\mathcal{H}}(x_T^u)^\top\big) \circ \mathcal{H}(\hat{x}_T^j) \\
& - J_{\mathcal{H}}(x_T^u)^\top\big(y - e^{\bar{f}_T}\mathcal{H}(x_0)\big)\big\|^2 + \sqrt{2\eta}\epsilon_j,
\end{aligned} \tag{6}
$$

where $\eta > 0$ is the step size and $\epsilon_j \sim \mathcal{N}(0, I)$. Unlike previous methods [4–6, 42], which motivate the update from a Bayesian posterior estimation and invoke the gradient of the quadratic data-fidelity loss, here an analogous quadratic guidance term arises naturally from the SOC formulation. Since we do not approximate the degradation operator and explicitly retain the terminal penalty coefficient γ, we refer to this sampling scheme as SOCS-VP-Nonlinear-γ.

If we disregard the running cost in the SOC formulation by taking $\gamma \to \infty$, then, since $\lim_{\gamma\to\infty}\gamma d_{0,\gamma}^{-1} = \frac{1}{e^{2\bar{f}_T}\bar{g}_T^2}\big(J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H}\big)^{-1}$. Eq. (4) simplifies:

$$
x_t = e^{\bar{f}_t}\big(1 - \frac{\bar{g}_t^2}{\bar{g}_T^2}\big)x_0 + \frac{e^{\bar{f}_t}\bar{g}_t^2}{e^{\bar{f}_T}\bar{g}_T^2}\big(J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H}\big)^{-1} \circ J_{\mathcal{H}}(x_T^u)^\top(y). \tag{7}
$$

Following the SOCS-VP-Nonlinear-γ spirit, we set $\hat{x}_T = \big(J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H}\big)^{-1} \circ J_{\mathcal{H}}(x_T^u)^\top(y)$. This amounts to solving $J_{\mathcal{H}}(x_T^u)^\top \circ \mathcal{H}(\hat{x}_T) = J_{\mathcal{H}}(x_T^u)^\top(y)$, we enforce $\mathcal{H}(\hat{x}_T) = y$ as a sufficient measurement-consistency condition that satisfies the preceding equation. We then refine $\hat{x}_T$ using Langevin dynamics:

$$\hat{x}_T^{j+1} = \hat{x}_T^j - \eta \nabla_{\hat{x}_T^j} \big\|\mathcal{H}(\hat{x}_T^j) - y\big\|^2 + \sqrt{2\eta}\epsilon_j. \tag{8}$$

According to [42], as $\eta \to 0$ and $j \to \infty$, the iterate $\hat{x}_T^j$ approximately follows $p(\hat{x}_T \mid x_t, y)$. We view this optimization as a special instance of Eq. (6), which we denote as SOCS-VP-Nonlinear-∞.

SOCS-VP-Linear. Furthermore, if we take the regime $\gamma \to \infty$ and adopt a linear approximation of the degradation operator, Eq. (4) reduces to:

$$x_t = e^{\bar{f}_t}\big(1 - \frac{\bar{g}_t^2}{\bar{g}_T^2}\big)x_0 + \frac{e^{\bar{f}_t}\bar{g}_t^2}{e^{\bar{f}_T}\bar{g}_T^2} \underbrace{\big(\mathcal{H}^\top\mathcal{H}\big)^{-1}\mathcal{H}^\top}_{\mathcal{H}^\dagger(\cdot)}(y), \tag{9}$$

where the coefficients in front of x_0 and $\mathcal{H}^\dagger(y)$ are time-dependent constants that can be precomputed offline. Because $\mathcal{H}^\dagger(y)$ assumes exact access to the pseudoinverse, which is rarely available in practice, we adopt a pragmatic two-step scheme. Within the EDM framework, we first produce a coarse prediction $x_T(x_t)$ using a few-step Euler ODE solver, however, this prediction may not fully align with the low-quality input in the degenerate subspace. Following [17, 38], we enforce measurement consistency by projecting onto the affine subspace $\{x \in \mathbb{R}^n : \mathcal{H}(x) = y\}$:

$$\hat{x}_T(x_t) = \big(I - \mathcal{H}^\dagger\mathcal{H}\big)x_T(x_t) + \mathcal{H}^\dagger(y). \tag{10}$$

Details for $\mathcal{H}(\cdot)$ and $\mathcal{H}^\dagger(\cdot)$ are given in the Appendix.

In all cases above, once we obtain the projected estimate $\hat{x}_T(x_t)$ from Eq. (10) in the linear setting, or the MCMC-refined estimate from Eq. (6) and Eq. (8) in the nonlinear or finite-γ setting, we apply SOC to pull the resulting $\hat{x}_T$ back onto the denoising manifold at time t by substituting it into the corresponding closed-form expression for x_t. This continues the reverse-time sampling with measurement consistency and SOC guidance. Under a locally linear and well-conditioned observation model, *SOCS-Linear* regime yielding a stable closed-form affine projection onto $\{x : \mathcal{H}(x) = y\}$. The *SOCS-Nonlinear* regime retains the full $J_{\mathcal{H}}(\cdot)$ with a finite γ, preserving the geometry of the prior manifold under strong degeneracy or model mismatch and enabling multi step likelihood-guided refinement. As emphasized in [42], to enhance robustness, we decouple consecutive states $x_{t+\Delta t}$ and x_t along the entire sampling trajectory so that the sampler can correct global errors that arise early in denoising, rather than restricting the distance. Algorithm 1 provides the detailed procedure.

As shown in **??**, the computational overhead of the Langevin refinements is minor relative to diffusion model evaluations, since the dominant cost typically stems from the neural network rather than the measurement operator $\mathcal{H}$. For completeness, we also present the linear-SDE case instantiated as a VE–SDE in Appendix. Moreover, we extend SOCS to latent diffusion models (LDMs) and flow-matching models (represented by SD3), details are provided in Appendix.

Algorithm 1 SOCS (Stochastic Optimal Control Sampling)

Require: Score model s_θ, measurement $\mathbf{y}$, operator $\mathcal{H}$, time grid $(t_i)_{i=0}^{N}$

Sample $\mathbf{x}_{t_0} \sim \mathcal{N}(\mathbf{0}, I)$
for $i = 0$ **to** $N-1$ **do**
  $\mathbf{x}_{T|t_i} \leftarrow \mathbf{x}_{t_i}^{(T)}(\mathbf{x}_{t_i})$ ▷ few step ODE sampling with s_θ
  if mode = Linear **then**
    $\hat{\mathbf{x}}_{T|t_i} \leftarrow (I - \mathcal{H}^\dagger \mathcal{H})\,\mathbf{x}_{T|t_i} + \mathcal{H}^\dagger \mathbf{y}$ ▷ Eq. (10)
  else if mode = Nonlinear-γ **then**
    for $j = 0$ **to** $N_{\text{inner}} - 1$ **do**
      $\hat{\mathbf{x}}_{T|t_i}^{(j+1)} \leftarrow \textsc{Refine}(\hat{\mathbf{x}}_{T|t_i}^{(j)})$ ▷ Eq. (8) for $\gamma \to \infty$/Eq. (6) for general cases.
    end for
    $\hat{\mathbf{x}}_{T|t_i} \leftarrow \hat{\mathbf{x}}_{T|t_i}^{(N_{\text{inner}})}$
  end if
  Project $\hat{\mathbf{x}}_{T|t_i}$ to SOC manifold, get $\mathbf{x}_{t_{i+1}}$ ▷ Eq. (4)
end for
return $\mathbf{x}_N \leftarrow \hat{\mathbf{x}}_N$

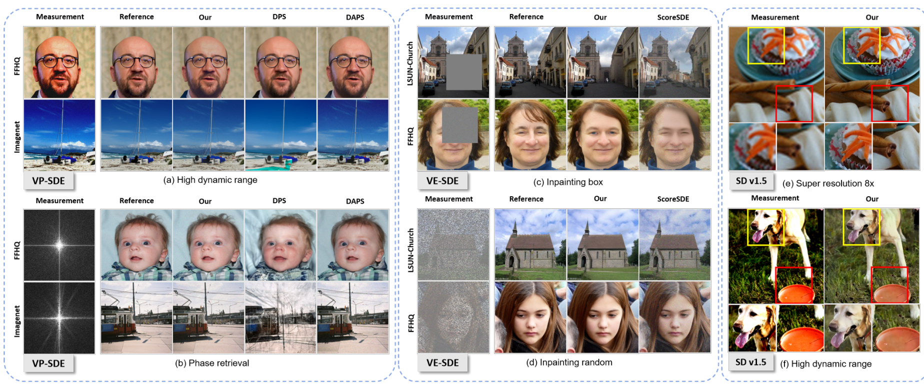

Fig. 2: Representative samples of SOCS. Our method leverages stochastic optimal control (SOC) to provide an efficient, theoretically grounded framework for diffusion inverse problems. In (a)(b), we present VP-SDE results on FFHQ and ImageNet-256, respectively; in (c)(d), we showcase VE-SDE results on LSUN Church and FFHQ; in (e)(f), we show natural images at a resolution of 512.

3.3 Comparison with posterior sampling methods

Existing posterior sampling methods enforce data consistency by steering the reverse denoising process with measurement guidance. This includes approximating $p_t(y \mid x_t)$ with an isotropic Gaussian [5, 32], projection-based corrections [5, 6], and subsequent optimization refinements [4, 25, 40, 42, 45]. Although implemented differently, these approaches ultimately rely on the same core term $\nabla_{x_t}\|\mathcal{H}(x_t) - y\|_2^2$, which injects observational information into prior-driven dynamics. We provide the first stochastic optimal control justification of this term. In SOCS-Nonlinear, the standard gradient guidance emerges as a special case in the limit $\gamma \to \infty$, whereas for finite γ our framework yields an optimal feed-

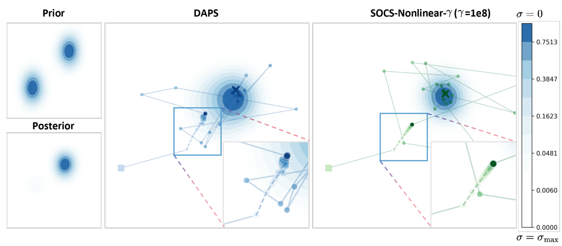

Fig. 3: SOCS vs DAPS on 2D synthetic data. SOCS exhibits a more direct refinement trajectory.

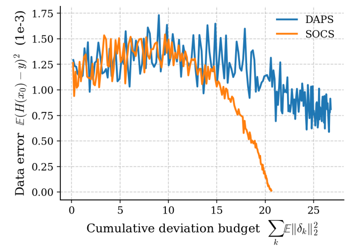

Fig. 4: Data error versus cumulative deviation budget.

back that minimizes control energy under a terminal measurement penalty. Both SOCS and DAPS [42] benefit from trajectory decoupling, Concretely, at each outer step k, they first predict a clean image $x_{T,k}$ from the current state x_k, then run gradient-based Langevin refinement to improve data consistency and obtain $\hat{x}_{T,k}$, and finally project back to the denoising trajectory to produce x_{k+1}. Fig. 3 compares SOCS and DAPS on a two-dimensional nonlinear inverse problem with a Gaussian mixture prior. We visualize the sampling trajectories $\{x_t\}_{t\in[0,T]}$ for both methods, and within the first step we additionally show the inner refinement path $x_{T,0} \to \hat{x}_{T,0}$, where intermediate Langevin iterates are connected by dashed segments. SOCS exhibits a more direct refinement trajectory. In Fig. 4, we further quantify the budget–consistency trade-off using δ_k, which characterizes the optimization trajectory from $x_{T,k} \to \hat{x}_{T,k}$. Benefiting from the optimal control objective, SOCS consistently attains lower data error than DAPS under a smaller cumulative deviation budget $\sum_k \mathbb{E}\|\delta_k\|_2^2$, which cannot be achieved by Langevin refinement driven solely by $\nabla_{x_t}\|\mathcal{H}(x_t) - y\|_2^2$. We provide additional discussions in Appendix.

4 Experiments

4.1 Experimental Setup

We evaluate our three sampling variants using the FFHQ model pretrained by [5] and the ImageNet pretrained model provided by [8]. As discussed in Sec. 3, we adopt the EDM time discretization and noise schedule. For each time t, we estimate $x_T(x_t)$ with a few-step Euler ODE solver.

For all linear tasks, we use 4 ODE-solver NFEs and 250 reverse-denoising steps and for nonlinear tasks, we use 8 ODE-solver NFEs and 500 reverse-denoising steps. Detailed settings for SOCS-Linear, SOCS-Nonlinear, and SOCS-Nonlinear-γ, including model configurations, samplers, denoising scheduler and other hyperparameters are provided in the Appendix. We also report an ablation on the terminal-penalty coefficient γ to investigate its impact on reconstruction quality, along with other ablation studies.

Datasets and metrics. Following convention, we evaluate on FFHQ 256×256 [16] and ImageNet 256×256 [7]. For assessment, we use 100 images from the validation split of each dataset. We include Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), Learned Perceptual Image Patch Similarity (LPIPS) [43] and Fréchet Inception Distance (FID) [12] as evaluation metrics.

Inverse problems. We evaluate our method on a suite of linear and nonlinear tasks. For linear inverse problems, we include: (i) $\times 4$ bicubic super-resolution; (ii) Gaussian deblur with a 61×61 kernel and standard deviation 3.0; (iii) motion blur generated using 61×61 kernels with blur strength 0.5; (iv) inpainting, where box inpainting masks a random 128×128 square and random inpainting masks 70% of pixels; (v) super-resolution + inpainting, which first applies a random $30\% - 70\%$ pixel mask followed by $\times 4$ bicubic downsampling.

We further consider three nonlinear inverse problems: (i) phase retrieval, where we apply the Fourier transform to each image and keep only the magnitude as the measurement. Given the task's inherent difficulty, following prior work [5, 42], we use an oversampling rate of 2.0 which is a standard choice to mitigate the ill-posedness and report the best result over four independent runs; (ii) High dynamic range (HDR) reconstruction, which aims to recover a higher–dynamic-range image from a low–dynamic–range input, with the expansion factor set to 2; (iii) For nonlinear deblurring, we follow the default configuration described in [37] under the same protocol and parameter settings.

For both linear and nonlinear tasks, we report the mean performance (PSNR, SSIM and LPIPS) across 100 validation images. The FID scores are evaluated with the same 100 validation images. To conform with the noisy inverse-problem setting, we corrupt all measurements, both linear and nonlinear, with additive white Gaussian noise of standard deviation $\beta_y = 0.05$. For experiments based on SD-v1.5, we instead use $\beta_y = 0.01$, while for SD3 we set $\beta_y = 0.03$. Complete task specifications and hyperparameter are provided in the Appendix.

Baselines. For VP-SDE, we compare against DPS [5], DDRM [17], DDNM [38], DCDP [22], FPS-SMC [9], DiffPIR [45], DAPS [42] and DOC [20] which also use SOC. Note that DDRM, DDNM, and FPS-SMC are limited to linear inverse problems and do not handle nonlinear cases. We especially include RED-diff [23] for nonlinear experiments. For VE–SDE, we benchmark against MCG [6] and ScoreSDE [35], focusing on the inpainting tasks. In addition, for SD v1.5 based latent diffusion experiments, we report results for LatentDAPS [42] and PSLD [29]. For flow-matching (FM) models, we include PSLD, LatentDAPS, FlowChef [24], and FlowDPS [18] as additional baselines.

4.2 Main Results

We report quantitative results for the linear experiments on the FFHQ and ImageNet datasets in Tab. 1, and for the nonlinear experiments in Tab. 2. Together with the qualitative comparisons in Fig. 2, our method achieves comparable or superior performance across all tasks.

The observed preference between the linear and nonlinear variants follows the conditioning and local linearity of the observation model. When the forward map behaves linearly and remains well-conditioned, the linear variant approximates $J_{\mathcal{H}} \approx \mathcal{H}$ with $\gamma \to \infty$ and performs a stable affine projection onto the

Table 1: Quantitative evaluation on FFHQ and ImageNet across 6 linear tasks. We compare methods on FFHQ (left) and ImageNet (right). The best and second-best results are indicated by bold and underlined marks, respectively.

Task	Method	FFHQ (256×256)				ImageNet (256×256)			
		PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
Super Resolution 4×	SOCS-Linear	28.51	0.797	0.165	**72.99**	**26.42**	**0.733**	0.340	**93.77**
	SOCS-Nonlinear	28.39	0.772	<u>0.150</u>	74.82	26.10	0.672	**0.277**	102.51
	SOCS-Nonlinear-γ	<u>28.66</u>	0.800	0.166	76.86	<u>26.36</u>	<u>0.708</u>	<u>0.283</u>	<u>97.54</u>
	DAPS [42]	28.31	0.768	**0.149**	74.75	25.95	0.667	0.335	112.60
	DPS [5]	24.25	0.677	0.190	84.33	22.56	0.574	0.387	155.01
	DDRM [17]	27.53	0.796	0.180	92.19	24.90	0.673	0.308	183.50
	DDNM [38]	27.93	<u>0.808</u>	0.170	91.15	25.02	0.694	0.306	158.35
	DCDP [22]	26.06	0.629	0.322	118.09	23.83	0.548	0.511	162.58
	FPS-SMC [9]	**29.02**	**0.828**	0.234	106.56	24.64	0.671	0.340	174.98
	DiffPIR [45]	26.48	0.733	0.272	<u>74.62</u>	24.04	0.616	0.308	97.72
	RED-diff [23]	25.46	0.591	0.391	130.71	23.12	0.504	0.546	163.79
	DOC [20]	27.15	0.790	0.173	90.96	-	-	-	-
Inpaint (Box)	SOCS-Linear	24.22	0.794	<u>0.109</u>	53.21	20.05	0.741	0.263	160.89
	SOCS-Nonlinear	24.27	0.742	0.135	55.26	<u>20.40</u>	0.702	0.248	117.96
	SOCS-Nonlinear-γ	24.45	0.826	**0.092**	**43.79**	**20.51**	**0.772**	**0.209**	**113.67**
	DAPS	24.41	0.753	0.117	51.51	20.05	0.702	0.246	130.05
	DPS	23.49	0.804	0.134	62.18	18.97	0.717	0.269	124.78
	DDRM	22.18	0.818	0.126	62.87	18.72	0.739	0.235	125.26
	DDNM	24.14	**0.837**	0.117	48.96	18.46	<u>0.752</u>	0.238	<u>117.25</u>
	DCDP	<u>24.49</u>	0.776	0.206	47.17	17.44	0.656	0.261	167.07
	FPS-SMC	**24.70**	<u>0.828</u>	0.126	<u>45.29</u>	19.83	0.743	<u>0.209</u>	125.73
Inpaint (Random)	SOCS-Linear	29.31	0.834	0.105	77.09	28.44	0.790	0.175	91.92
	SOCS-Nonlinear	<u>30.13</u>	0.820	0.093	55.42	28.05	0.754	**0.102**	55.44
	SOCS-Nonlinear-γ	**30.63**	**0.872**	**0.065**	**52.73**	<u>28.75</u>	<u>0.808</u>	<u>0.108</u>	**53.02**
	DAPS	29.66	0.800	<u>0.090</u>	<u>54.57</u>	**28.76**	0.789	0.145	<u>54.41</u>
	DPS	28.30	0.812	0.132	71.74	26.59	0.736	0.254	98.27
	DDRM	25.80	0.774	0.166	103.58	22.63	0.742	0.271	219.72
	DDNM	29.17	<u>0.850</u>	0.092	64.65	26.55	**0.868**	0.131	99.36
	DCDP	30.12	0.835	0.105	62.09	26.41	0.731	0.167	68.22
	FPS-SMC	27.68	0.824	0.213	104.22	26.50	0.707	0.233	68.96
	DOC	27.70	0.795	0.132	100.51	-	-	-	-
Inpaint&SR	SOCS-Linear	27.21	<u>0.762</u>	0.195	86.36	<u>25.53</u>	<u>0.697</u>	0.384	142.82
	SOCS-Nonlinear	26.75	0.714	0.191	94.79	24.65	0.569	0.346	134.14
	SOCS-Nonlinear-γ	**27.29**	**0.780**	<u>0.180</u>	**82.30**	**25.73**	**0.708**	<u>0.344</u>	<u>123.85</u>
	DAPS	<u>27.22</u>	0.757	**0.166**	<u>82.51</u>	25.48	0.686	**0.309**	**118.82**
	DPS	23.00	0.640	0.213	87.87	21.45	0.533	0.425	166.25
Gaussian Deblurring	SOCS-Linear	27.21	0.785	0.190	98.13	25.38	<u>0.720</u>	**0.275**	111.49
	SOCS-Nonlinear	<u>28.52</u>	0.788	0.179	76.12	26.32	0.708	0.313	**106.00**
	SOCS-Nonlinear-γ	**28.79**	**0.821**	**0.162**	**70.36**	<u>26.67</u>	**0.740**	<u>0.282</u>	110.69
	DAPS	28.46	<u>0.809</u>	<u>0.169</u>	73.54	26.37	0.697	0.297	109.61
	DPS	25.30	0.706	0.215	74.87	22.84	0.587	0.312	115.02
	DDRM	27.68	0.756	0.224	95.25	**27.81**	0.639	0.333	124.55
	DDNM	28.25	0.801	0.220	72.23	23.82	0.657	0.323	211.84
	DCDP	27.73	0.784	0.218	<u>71.74</u>	25.12	0.690	0.429	142.46
	FPS-SMC	27.05	0.770	0.284	122.98	23.09	0.588	0.417	220.59
	DiffPIR	25.42	0.714	0.345	119.64	23.25	0.646	0.340	195.75
	RED-diff	26.34	0.592	0.240	72.09	24.24	0.589	0.292	<u>108.00</u>
Motion Deblurring	SOCS-Linear	26.53	0.724	0.196	93.62	25.38	0.720	0.275	111.48
	SOCS-Nonlinear	<u>29.74</u>	0.790	0.128	70.28	<u>28.36</u>	0.749	<u>0.203</u>	<u>72.83</u>
	SOCS-Nonlinear-γ	**30.39**	<u>0.828</u>	**0.104**	72.32	**28.91**	**0.773**	0.204	90.51
	DAPS	29.49	**0.829**	<u>0.110</u>	**55.66**	27.76	<u>0.765</u>	**0.182**	**64.79**
	DPS	26.48	0.750	0.137	70.14	25.12	0.682	0.272	100.31
	DCDP	25.85	0.736	0.174	87.73	23.71	0.648	0.414	177.72
	FPS-SMC	28.66	0.820	0.139	<u>67.92</u>	26.66	0.680	0.339	87.92
	DiffPIR	25.77	0.725	0.356	131.46	24.70	0.678	0.369	102.73
	DOC	25.87	0.759	0.199	92.76	-	-	-	-

Table 2: Quantitative evaluation on FFHQ and ImageNet across 3 nonlinear tasks. We compare methods on FFHQ (left) and ImageNet (right). The best and second-best results are indicated by bold and underlined marks, respectively.

Task	Method	FFHQ (256×256)				ImageNet (256×256)			
		PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)	PSNR (↑)	SSIM (↑)	LPIPS (↓)	FID (↓)
Phase Retrieval	SOCS-Linear	-	-	-	-	-	-	-	-
	SOCS-Nonlinear	$28.73_{\pm 3.01}$	$0.759_{\pm 0.073}$	$0.119_{\pm 0.070}$	52.43	$24.74_{\pm 6.46}$	$\underline{0.710}_{\pm 0.212}$	$0.263_{\pm 0.200}$	96.70
	SOCS-Nonlinear-γ	$\underline{29.59}_{\pm 2.88}$	$\mathbf{0.808}_{\pm 0.050}$	$\underline{0.072}_{\pm 0.029}$	**44.86**	$\underline{24.90}_{\pm 6.58}$	$\mathbf{0.716}_{\pm 0.221}$	$\underline{0.256}_{\pm 0.207}$	**89.42**
	DAPS	$\mathbf{29.88}_{\pm 2.38}$	$\underline{0.807}_{\pm 0.041}$	$\mathbf{0.069}_{\pm 0.030}$	$\underline{45.33}$	$\mathbf{25.05}_{\pm 7.11}$	$0.687_{\pm 0.230}$	$\mathbf{0.253}_{\pm 0.214}$	$\underline{95.92}$
	DPS	$16.29_{\pm 5.45}$	$0.442_{\pm 0.192}$	$0.388_{\pm 0.180}$	137.81	$15.89_{\pm 3.86}$	$0.445_{\pm 0.159}$	$0.413_{\pm 0.127}$	201.93
	DCDP	$28.06_{\pm 8.64}$	$0.754_{\pm 0.351}$	$0.257_{\pm 0.272}$	69.56	$23.05_{\pm 6.54}$	$0.624_{\pm 0.257}$	$0.290_{\pm 0.232}$	119.47
	RED-diff	$16.83_{\pm 7.84}$	$0.521_{\pm 0.202}$	$0.579_{\pm 0.200}$	209.75	$14.36_{\pm 5.38}$	$0.398_{\pm 0.129}$	$0.579_{\pm 0.129}$	220.99
Nonlinear Deblur	SOCS-Linear	$21.73_{\pm 2.36}$	$0.647_{\pm 0.008}$	$0.286_{\pm 0.081}$	113.90	$20.70_{\pm 3.06}$	$0.564_{\pm 0.150}$	$0.479_{\pm 0.150}$	193.74
	SOCS-Nonlinear	$28.29_{\pm 1.58}$	$0.769_{\pm 0.035}$	$0.169_{\pm 0.051}$	73.36	$26.11_{\pm 2.85}$	$0.704_{\pm 0.080}$	$0.262_{\pm 0.102}$	111.88
	SOCS-Nonlinear-γ	$\underline{28.42}_{\pm 1.63}$	$\mathbf{0.786}_{\pm 0.037}$	$0.157_{\pm 0.050}$	68.20	$\underline{26.81}_{\pm 2.51}$	$\mathbf{0.740}_{\pm 0.069}$	$\mathbf{0.201}_{\pm 0.106}$	80.51
	DAPS	$28.10_{\pm 1.62}$	$0.762_{\pm 0.030}$	$\underline{0.153}_{\pm 0.033}$	69.45	$\mathbf{27.06}_{\pm 2.81}$	$\underline{0.733}_{\pm 0.066}$	$\underline{0.205}_{\pm 0.088}$	83.81
	DPS	$23.05_{\pm 2.24}$	$0.637_{\pm 0.087}$	$0.226_{\pm 0.076}$	90.69	$21.98_{\pm 2.62}$	$0.551_{\pm 0.144}$	$0.443_{\pm 0.149}$	207.94
	DCDP	$27.44_{\pm 2.45}$	$0.770_{\pm 0.068}$	$0.186_{\pm 0.054}$	**58.41**	$26.07_{\pm 0.15}$	$0.620_{\pm 0.113}$	$0.216_{\pm 0.113}$	**53.71**
	RED-diff	$\mathbf{29.16}_{\pm 0.08}$	$\underline{0.783}_{\pm 0.051}$	$\mathbf{0.146}_{\pm 0.055}$	$\underline{63.22}$	$25.08_{\pm 0.68}$	$0.584_{\pm 0.130}$	$0.244_{\pm 0.046}$	$\underline{55.75}$
High Dynamic Range	SOCS-Linear	$25.96_{\pm 1.84}$	$0.767_{\pm 0.058}$	$\underline{0.112}_{\pm 0.033}$	63.87	$22.51_{\pm 3.67}$	$0.744_{\pm 0.109}$	$0.179_{\pm 0.082}$	72.35
	SOCS-Nonlinear	$\underline{26.30}_{\pm 3.24}$	$\underline{0.800}_{\pm 0.079}$	$0.118_{\pm 0.061}$	$\underline{61.54}$	$\underline{24.24}_{\pm 4.29}$	$\underline{0.757}_{\pm 0.117}$	$\underline{0.152}_{\pm 0.094}$	**60.72**
	SOCS-Nonlinear-γ	$24.31_{\pm 3.11}$	$\mathbf{0.825}_{\pm 0.088}$	$\mathbf{0.108}_{\pm 0.073}$	69.05	$21.88_{\pm 4.51}$	$\mathbf{0.768}_{\pm 0.121}$	$\mathbf{0.148}_{\pm 0.088}$	87.64
	DAPS	$\mathbf{26.92}_{\pm 3.35}$	$0.792_{\pm 0.084}$	$0.145_{\pm 0.066}$	**55.00**	$\mathbf{25.21}_{\pm 4.05}$	$0.714_{\pm 0.116}$	$0.167_{\pm 0.096}$	$\underline{62.37}$
	DPS	$24.53_{\pm 5.39}$	$0.660_{\pm 0.180}$	$0.270_{\pm 0.021}$	79.48	$19.83_{\pm 6.46}$	$0.603_{\pm 0.280}$	$0.538_{\pm 0.248}$	179.90
	RED-diff	$21.26_{\pm 2.55}$	$0.617_{\pm 0.080}$	$0.269_{\pm 0.081}$	97.15	$22.85_{\pm 3.31}$	$0.602_{\pm 0.128}$	$0.279_{\pm 0.138}$	87.19

(a) Super resolution by a factor 16 **(b)** 192 × 192 box inpainting

Fig. 5: Sample diversity. We present several diverse samples generated by the SOCS under two sparse measurements. SOCS produces a variety of samples with distinct features, including differences in expression, wearings, and hairstyles.

measurement-consistent set, which cooperates with the diffusion prior and avoids unnecessary iterative refinement. In the presence of strong degeneracy, large null spaces or spatially varying operators, or noticeable model mismatch, the nonlinear variant retains $J_{\mathcal{H}}(\cdot)$ with a finite γ and conducts multi step Langevin guidance that respects the geometry of prior manifold and better resolves ambiguity without introducing projection artifacts.

Nonlinear forward models further accentuate this preference. Nonlinear deblurring involves state-dependent point spread functions that invalidate a global linear approximation, and high dynamic range formation includes saturation and tone-curve effects that break linearity. SOCS-Nonlinear addresses these cases via Jacobian-aware refinement and controlled guidance, enforcing data consistency while maintaining trajectory stability. Phase retrieval is magnitude based and discards phase, so no bounded linear pseudoinverse $\mathcal{H}^\dagger$ exists to map amplitude-only measurements back to the image domain. This makes SOCS-Linear inapplicable, whereas SOCS-Nonlinear remains effective through likelihood-driven

Table 3: Quantitative evaluation on FFHQ at 768×768 on two super-resolution tasks with SD3-medium. The best and second-best results within each type of task are indicated by bold and underlined marks, respectively.

Method	FFHQ(768×768)							
	Super-resolution $\times 12$ (Avgpool)				Super-resolution $\times 12$ (Bicubic)			
	PSNR↑	SSIM↑	LPIPS↓	FID↓	PSNR↑	SSIM↑	LPIPS↓	FID↓
PSLD [29]	25.58	0.547	0.481	148.16	25.60	0.572	0.517	154.78
LatentDAPS [42]	28.94	0.823	<u>0.296</u>	74.43	29.65	0.819	<u>0.276</u>	65.41
FlowChef [24]	29.21	**0.852**	0.341	153.37	29.56	**0.846**	0.346	152.00
FlowDPS [18]	<u>29.41</u>	0.777	0.482	<u>65.00</u>	<u>29.73</u>	0.775	0.482	<u>61.24</u>
SOCS-SD3(Ours)	**30.14**	<u>0.827</u>	**0.259**	**53.61**	**30.45**	<u>0.842</u>	**0.248**	**45.72**

updates. Empirically, as conditioning worsens, SOCS-Nonlinear consistently surpasses SOCS-Linear despite higher cost. We therefore recommend SOCS-Linear for super-resolution and SOCS-Nonlinear for the remaining tasks.

Moreover, our method produces a wider range of plausible samples when the measurements contain less information. For demonstration, we use a larger downsampling factor of 16 for the super resolution task and a larger 192×192 mask for the box inpainting task. The examples in Fig. 5 show that our method generates diverse samples while remaining faithful to the measurements.

Further results on LDMs. To demonstrate the broad versatile applicability of SOCS across various Latent Diffusion Model (LDM) frameworks, we evaluate our method on both the traditional epsilon-prediction based SD v1.5 [27] and the advanced Flow-Matching (FM) based SD3. By operating within the lightweight VAE [19] latent space, these models facilitate high-fidelity image synthesis under complex prompt guidance. For SD v1.5, we implement the method described in Sec. 3 to validate SOCS on multiple inverse problems. Representative samples on FFHQ 256×256 for three linear and one non-linear task are shown in **??**, with 512×512 results presented in Fig. 2. Furthermore, we extend SOCS to FM-based models like SD3-medium, which leverage stable ODE-based sampling trajectories and offer superior scalability for large-scale generative modeling. Under an idealized straight-line interpretation, a flow-matching trajectory can be approximated as a linear ODE with near-constant velocity, enabling a seamless integration of SOCS. Quantitative results in Tab. 3 for two $\times 12$ super-resolution settings (768×768) demonstrate that our method achieves state-of-the-art performance across most metrics. Visual results for SOCS-SD3 are provided in **??**. Detailed prompts, theoretical derivations, and implementation specifics for both LDM paradigms are provided in Appendix and **??**, respectively.

4.3 Additional Discussion

Terminal penalty coefficient of SOCS. Fig. 6 reports quantitative performance as a function of γ over five tasks, and **??** provides qualitative comparisons. We observe a clear trade-off: as $\gamma \to \infty$, the update becomes dominated by the reconstruction term, which can pull samples away from the generative prior and

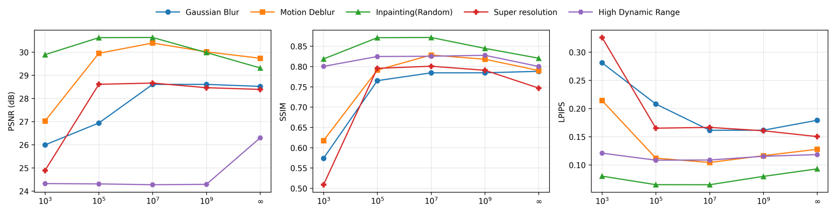

Fig. 6: Quantitative evaluation of image quality metrics as a function of the terminal penalty coefficient γ.

yield oversmoothed outputs; conversely, overly small γ (e.g., $\sim 10^3$) under-injects guidance, leading to noisy super-resolution and weaker measurement consistency in phase retrieval. From the posterior-sampling perspective in App. **??**, γ acts as an effective noise precision balancing the diffusion prior and the measurement likelihood, with $\gamma = 1/\sigma^2$ recovering the exact Gaussian posterior in canonical measurement space. In practice, preconditioning and the r_t-scaling of the data-fidelity term in the Langevin refinement rescale this trade-off, which motivates using numerically larger γ. Importantly, across operators and tasks, performance exhibits a broad plateau for finite γ, where metrics and visual quality vary little. This robustness stems from how γ enters SOCS. In Eq. (6), it appears through a regularized inverse $(I + \gamma A)^{-1}$, so the effective guidance weight $\gamma(I + \gamma A)^{-1}$ saturates once γ is sufficiently large. Consequently, SOCS does not require fine-grained tuning; we recommend $\gamma \in [10^5, 10^9]$ for a favorable balance between measurement fidelity and the generative prior.

Computational efficiency of SOCS. We run SOCS with various configurations to test performance under different computing budgets. Specifically, we evaluate with the number of function evaluations (NFE) of the diffusion models ranging from 400 to 4K, with configurations specified in Appendix.

More ablation studies. We discuss more additional ablation studies in Appendix, covering the number of function evaluations, the number of ODE steps, and the impact of the noise schedule on model performance.

5 Conclusion

In summary, we present SOCS for solving image inverse problems. Our method models the denoising dynamics through the lens of SOC and integrates them into sampling process, enforcing measurement consistency while gently tuning information injection to match the model's native generative capacity. Empirically, SOCS delivers competitive reconstruction quality across a range of challenging inverse problems compared with existing methods.

Acknowledgements

The authors thank the anonymous reviewers for their valuable comments and constructive feedback, which have significantly enhanced the quality of this work. This research was supported by the AI for Science Program of the Shanghai Municipal Commission of Economy and Informatization (Grant No. 2025-GZL-RGZN-BTBX-02026) and the National Natural Science Foundation of China (Grant No. 62376155).

References

1. Berner, J., Richter, L., Ullrich, K.: An optimal control perspective on diffusion-based generative modeling. arXiv preprint arXiv:2211.01364 (2022)
2. Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A., Akyildiz, O.D.: Tweedie moment projected diffusions for inverse problems. arXiv preprint arXiv:2310.06721 (2023)
3. Chen, T., Gu, J., Dinh, L., Theodorou, E.A., Susskind, J., Zhai, S.: Generative modeling with phase stochastic bridges. arXiv preprint arXiv:2310.07805 (2023)
4. Chen, T., Wang, Z., Zhou, M.: Enhancing and accelerating diffusion-based inverse problem solving through measurements optimization. arXiv preprint arXiv:2412.03941 (2024)
5. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687 (2022)
6. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. Advances in Neural Information Processing Systems **35**, 25683–25696 (2022)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
8. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
9. Dou, Z., Song, Y.: Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In: The Twelfth International Conference on Learning Representations (2024)
10. Freeman, W.T., Jones, T.R., Pasztor, E.C.: Example-based super-resolution. IEEE Computer graphics and Applications **22**(2), 56–65 (2002)
11. Geman, S., Geman, D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Transactions on pattern analysis and machine intelligence (6), 721–741 (1984)
12. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in neural information processing systems **33**, 6840–6851 (2020)
14. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. Advances in neural information processing systems **35**, 26565–26577 (2022)

15. Karras, T., Aittala, M., Lehtinen, J., Hellsten, J., Aila, T., Laine, S.: Analyzing and improving the training dynamics of diffusion models. In: In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24174–24184 (2024)
16. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019)
17. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. Advances in neural information processing systems **35**, 23593–23606 (2022)
18. Kim, J., Kim, B.S., Ye, J.C.: Flowdps: Flow-driven posterior sampling for inverse problems. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12328–12337 (2025)
19. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
20. Li, H., Pereira, M.: Solving inverse problems via diffusion optimal control. Advances in Neural Information Processing Systems **37**, 73549–73571 (2024)
21. Li, W., Todorov, E.: Iterative linear quadratic regulator design for nonlinear biological movement systems. In: First International Conference on Informatics in Control, Automation and Robotics. vol. 2, pp. 222–229. SciTePress (2004)
22. Li, X., Kwon, S.M., Liang, S., Alkhouri, I.R., Ravishankar, S., Qu, Q.: Decoupled data consistency with diffusion purification for image restoration. arXiv preprint arXiv:2403.06054 (2024)
23. Mardani, M., Song, J., Kautz, J., Vahdat, A.: A variational perspective on solving inverse problems with diffusion models. In: Proceedings of the 12th International Conference on Learning Representations (ICLR) (2024), `https://openreview.net/forum?id=umG1nU1wZg`, iCLR 2024
24. Patel, M., Wen, S., Metaxas, D.N., Yang, Y.: Steering rectified flow models in the vector field for controlled image generation. arXiv preprint arXiv:2412.00100 (2024)
25. Peng, X., Zheng, Z., Dai, W., Xiao, N., Li, C., Zou, J., Xiong, H.: Improving diffusion models for inverse problems using optimal posterior covariance. In: Proceedings of the 41st International Conference on Machine Learning (ICML). Proceedings of Machine Learning Research, vol. 235, p. —. PMLR (2024), `https://proceedings.mlr.press/v235/peng24a.html`
26. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on pattern analysis and machine intelligence **12**(7), 629–639 (2002)
27. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
28. Rout, L., Chen, Y., Ruiz, N., Kumar, A., Caramanis, C., Shakkottai, S., Chu, W.S.: Rb-modulation: Training-free personalization of diffusion models using stochastic optimal control. arXiv preprint arXiv:2405.17401 (2024)
29. Rout, L., Raoof, N., Daras, G., Caramanis, C., Dimakis, A., Shakkottai, S.: Solving linear inverse problems provably via posterior sampling with latent diffusion models. Advances in Neural Information Processing Systems **36**, 49960–49990 (2023)
30. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 conference proceedings. pp. 1–10 (2022)

31. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. IEEE transactions on pattern analysis and machine intelligence **45**(4), 4713–4726 (2022)
32. Song, J., Vahdat, A., Mardani, M., Kautz, J.: Pseudoinverse-guided diffusion models for inverse problems. In: International Conference on Learning Representations (2023)
33. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. Advances in neural information processing systems **32** (2019)
34. Song, Y., Ermon, S.: Improved techniques for training score-based generative models. Advances in neural information processing systems **33**, 12438–12448 (2020)
35. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)
36. Tassa, Y., Erez, T., Todorov, E.: Synthesis and stabilization of complex behaviors through online trajectory optimization. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. pp. 4906–4913. IEEE (2012)
37. Tran, P., Tran, A.T., Phung, Q., Hoai, M.: Explore image deblurring via encoded blur kernel space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11956–11965 (2021)
38. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. In: ICLR (2023)
39. Welling, M., Teh, Y.W.: Bayesian learning via stochastic gradient langevin dynamics. In: Proceedings of the 28th international conference on machine learning (ICML-11). pp. 681–688 (2011)
40. Wu, Z., Sun, Y., Chen, Y., Zhang, B., Yue, Y., Bouman, K.: Principled probabilistic imaging using diffusion models as plug-and-play priors. Advances in Neural Information Processing Systems **37**, 118389–118427 (2024)
41. Xia, B., Zhang, Y., Wang, S., Wang, Y., Wu, X., Tian, Y., Yang, W., Van Gool, L.: Diffir: Efficient diffusion model for image restoration. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 13095–13105 (2023)
42. Zhang, B., Chu, W., Berner, J., Meng, C., Anandkumar, A., Song, Y.: Improving diffusion inverse problem solving with decoupled noise annealing. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 20895–20905 (2025)
43. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
44. Zhu, K., Pan, M., Ma, Y., Fu, Y., Yu, J., Wang, J., Shi, Y.: Unidb: A unified diffusion bridge framework via stochastic optimal control. arXiv preprint arXiv:2502.05749 (2025)
45. Zhu, Y., Zhang, K., Liang, J., Cao, J., Wen, B., Timofte, R., Van Gool, L.: Denoising diffusion models for plug-and-play image restoration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1219–1229 (2023)