

Constrained Particle Seeking: Solving Diffusion Inverse Problems with Just Forward Passes

Anonymous submission

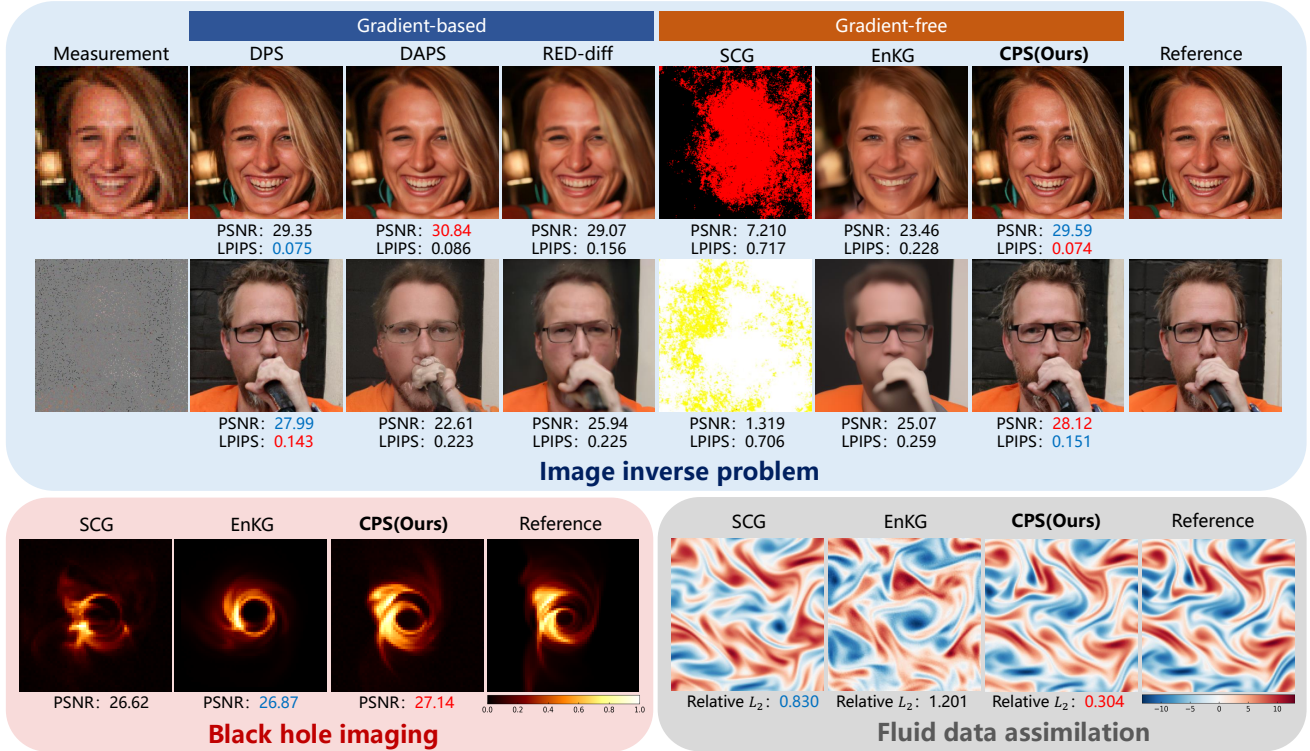


Figure 1: **Qualitative comparison between CPS and baseline methods.** The proposed CPS relies solely on the forward pass of the observation model to solve both image and scientific inverse problems. It performs competitively with advanced gradient-based methods and significantly outperforms gradient-free methods. The best results are highlighted in red, and the second-best results are shown in blue.

Abstract

Diffusion models have gained prominence as powerful generative tools for solving inverse problems due to their ability to model complex data distributions. However, existing methods typically rely on complete knowledge of the forward observation process to compute gradients for guided sampling, limiting their applicability in scenarios where such information is unavailable. In this work, we introduce *Constrained Particle Seeking (CPS)*, a novel gradient-free approach that leverages all candidate particle information to actively search for the optimal particle while incorporating constraints aligned with high-density regions of the unconditional prior. Unlike previous methods that passively select promising candidates, CPS reformulates the inverse problem as a constrained optimization task, enabling more flexible and efficient particle seeking. We demonstrate that CPS can effectively solve both image and scientific inverse problems, achieving results com-

parable to gradient-based methods while significantly outperforming gradient-free alternatives.

Introduction

The inverse problem is a fundamental challenge across various fields, including computational imaging (Tonolini et al. 2020; McCann, Jin, and Unser 2017), medicine (Song et al. 2021), geophysics (Schwarzbach, Börner, and Spitzer 2011), and fluid dynamics (Li et al. 2020). Its goal is to recover the original signal, $\mathbf{x} \in \mathbb{R}^d$, from indirect and noisy observations, $\mathbf{y} \in \mathbb{R}^m$, typically modeled as:

$$\mathbf{y} = \mathcal{H}(\mathbf{x}) + \boldsymbol{\eta} \quad (1)$$

where $\mathcal{H}(\cdot) : \mathbb{R}^d \rightarrow \mathbb{R}^m$ represents the forward measurement operator and $\boldsymbol{\eta}$ denotes additive noise. This problem is

inherently ill-posed ($m < d$), meaning multiple potential solutions may satisfy the given observations. To resolve this, appropriate priors and regularization techniques are essential to mitigate ill-posedness and ensure a meaningful solution.

Recent advances in generative modeling, particularly diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2020; Lipman et al. 2022; Rombach et al. 2022), have proven effective as plug-and-play priors for solving inverse problems, offering significant advantages over traditional manual priors (Tibshirani 1996; Iordache, Bioucas-Dias, and Plaza 2012). Existing approaches rely on gradients of the observation process or pseudo-inverse information to guide the unconditional sampling process, ensuring that the generated results align with observations (Zhang et al. 2025; Wang, Yu, and Zhang 2022; Chung et al. 2022; Kawar et al. 2022a; Mardani et al. 2023). However, this dependency on privileged information limits the applicability of these methods, especially in real-world scenarios. For example, in some scenarios, forward models often involve numerical simulations that are highly nonlinear or computationally expensive, making gradient computation challenging (Oliver, Reynolds, and Liu 2008; Gillhofer et al. 2019; Duong et al. 2020; Kawar et al. 2022b).

Several efforts have explored the use of pre-trained diffusion models to solve inverse problems with black-box access to the forward model. One such approach is SCG (Huang et al. 2024), which employs a strategy akin to rejection sampling. At each step of the reverse denoising process, SCG samples multiple candidate particles and retains only the one that best aligns with the observation, discarding the others. While this method is simple and intuitive, it suffers from inefficiency due to the discarding of many candidates at each step. Another notable method, EnKG (Zheng et al. 2025a), is inspired by ensemble Kalman filtering (Katzfuss, Stroud, and Wikle 2016). EnKG initializes and maintains a set of particles to represent the state, progressively refining it based on observations. However, EnKG requires maintaining thousands of particles in high-dimensional inverse problems, resulting in significant computational overhead. Additionally, DPG (Tang et al. 2024) avoids the calculation of derivatives by using policy gradients, but it necessitates adding noise perturbations to clean samples, which can push it outside the data manifold. More importantly, these methods often underperform relative to gradient-based approaches when gradients are available.

To overcome these limitations, we propose **Constrained Particle Seeking (CPS)**, a principled approach that synthesizes information from all candidate particles to identify an optimal solution. We argue that each candidate particle, rather than being discarded, offers valuable local insights that can collectively guide the denoising process. In CPS, we reframe the diffusion inverse problem as a constrained optimization task. The objective is to find a particle that optimally aligns with the observation, where the alignment is measured by a local surrogate model of the forward measurement. This surrogate is robustly estimated using the full set of candidate particles. Concurrently, we introduce a constraint that ensures the resulting particle remains within the

high-probability density region of the reverse diffusion process, thereby preserving consistency with the diffusion prior. By leveraging all available candidates, CPS enhances sampling efficiency and achieves a superior balance between observation fidelity and prior knowledge. Furthermore, we incorporate a *Restart* strategy (Xu et al. 2023; Lugmayr et al. 2023) to iteratively correct for the cumulative errors inherent in the sequential sampling process, which significantly improves the robustness and accuracy of the final solution.

We evaluated CPS on various tasks, including benchmark image inverse problems and two scientific inverse problems: black hole imaging and fluid data assimilation, both characterized by highly nonlinear measurement processes. Our results show that CPS not only matches the performance of popular gradient-based methods in image inverse problems but also significantly enhances efficiency in forward processes where gradients are unavailable.

Preliminaries and Task Formulation

Diffusion Models

The diffusion model (Ho, Jain, and Abbeel 2020; Song et al. 2020) defines a stochastic process $\{\mathbf{x}_t\}_{t \in [0, T]}$ that progressively transforms data $\mathbf{x}_0 \sim p_{\text{data}}$ into noise. This process is governed by the following differential equation:

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t) dt + g(t) d\mathbf{w} \quad (2)$$

where $\mathbf{f}(\cdot, \cdot)$ is the drift coefficient, $g(\cdot)$ is the diffusion coefficient, and \mathbf{w} represents the standard Wiener process. Over time, the data will gradually be perturbed into a Gaussian density $\mathbf{x}_T \sim \mathcal{N}(0, \sigma_T^2 I)$. To generate data from noise, diffusion models reverse this process by denoising, as described by:

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)] dt + g(t) d\bar{\mathbf{w}} \quad (3)$$

where $\bar{\mathbf{w}}$ denotes the reverse-time Wiener process, and $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$ is the score, which represents the gradient of the log-likelihood of the marginal distribution at time t . The score can be approximated using a neural network by minimizing the denoising score matching loss (Vincent 2011).

Diffusion Posterior Sampling (DPS)

Several methods have been proposed to leverage diffusion models as plug-and-play priors for solving inverse problems. A notable method is diffusion posterior sampling (DPS) (Chung et al. 2022). Specifically, DPS replaces the unconditional score in Eq. 3 with a conditional score:

$$\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t | \mathbf{y}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{y} | \mathbf{x}_t) \quad (4)$$

This modification aims to enable sampling from the posterior distribution $p(\mathbf{x}_0 | \mathbf{y})$ given the observation. However, the conditional score involves the intractable integral:

$$p(\mathbf{y} | \mathbf{x}_t) = \int p(\mathbf{y} | \mathbf{x}_0) p(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0 = \mathbb{E}_{p(\mathbf{x}_0 | \mathbf{x}_t)} [p(\mathbf{y} | \mathbf{x}_0)] \quad (5)$$

This integral is computationally expensive, often requiring multiple clean samples of the full trajectory. To address this, DPS approximates the guidance term as follows:

$$p(\mathbf{y} | \mathbf{x}_t) \approx p(\mathbf{y} | \hat{\mathbf{x}}_{0|t}) = -\frac{1}{\sigma_y^2} \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{H}(\hat{\mathbf{x}}_{0|t})\|^2 \quad (6)$$

where $\hat{x}_{0|t} = \mathbb{E}[x_0|x_t]$ can be computed via Tweedie’s formula with one-step denoising (Efron 2011), and σ_y^2 is the variance of noise in the observation. By employing this approximation, DPS solves the inverse problem without the need for training or finetuning. However, DPS requires a clear gradient of the observation operator, which limits its applicability in scenarios where such gradients are difficult to compute or unavailable.

Inverse Problem as Stochastic Control

We adopt the task formulation proposed by SCG (Huang et al. 2024) and frame the solution to the inverse problem within a stochastic control framework. To guide the denoised trajectory toward a target state that satisfies the observation, we introduce a control signal u_t at each time point t in Eq. 3. This transforms the inverse problem into a stochastic optimal control problem (Pavon 1989), where the terminal cost and transient cost are expressed as follows:

$$C = \mathbb{E} \left[\Phi(x_0) + \int_0^T \frac{\alpha}{2} \|u_t\|^2 |x_T \right] \quad (7)$$

where $\Phi(x_0)$ represents the terminal cost, and α is a coefficient associated with the transient cost. In the context of the inverse problem, we define it as $\Phi(x_0) = \|y - \mathcal{H}(x_0)\|^2$. By applying path integral control theory (Theodorou, Buchli, and Schaal 2010; Williams, Aldrich, and Theodorou 2017; Zhang and Chen 2023), we obtain the following optimal policy:

$$u_t^* dt = \frac{\mathbb{E}_{p(x_0|x_t)} \left[\exp\left(\frac{-\Phi(x_0)}{\alpha}\right) d\bar{w} \right]}{\mathbb{E}_{p(x_0|x_t)} \left[\exp\left(\frac{-\Phi(x_0)}{\alpha}\right) \right]} \quad (8)$$

In this equation, the sample form $p(x_0|x_t)$ is obtained via the uncontrolled process, i.e., Eq. 3. As can be seen, Eq. 8 still faces the same challenge as in Eq. 5, where computing the expected term requires simulating the entire trajectory to obtain x_0 . SCG (Huang et al. 2024) uses an intuitive approximation to estimate the optimal policy. They first consider the limit where $\alpha \rightarrow 0$, in which case Eq. 8 degenerates into selecting the $d\bar{w}$ that minimizes $\Phi(x_0)$. Next, they apply an approach similar to DPS, using one-step denoising results $\Phi(\hat{x}_{0|t})$ to estimate $\Phi(x_0)$. In practice, SCG samples multiple candidates at the diffusion reverse kernel $x_t^1, \dots, x_t^n \sim p(x_t|x_{t+1}) = \mathcal{N}(\mu_t, \sigma_t^2 I)$ of the DDIM solver (Song, Meng, and Ermon 2020) and heuristically retains the particle whose clean estimation is closest to y as the next state. The other particles are discarded.

Methods

In this section, we first present an empirical analysis to highlight the inefficiencies of the SCG method. Building on these insights, we introduce the **Constrained Particle Seeking (CPS)** framework, which utilizes all available candidate location information to identify the optimal particle. Finally, we explain how CPS can be combined with a *Restart* strategy to reduce cumulative errors and further enhance the robustness of the solution.

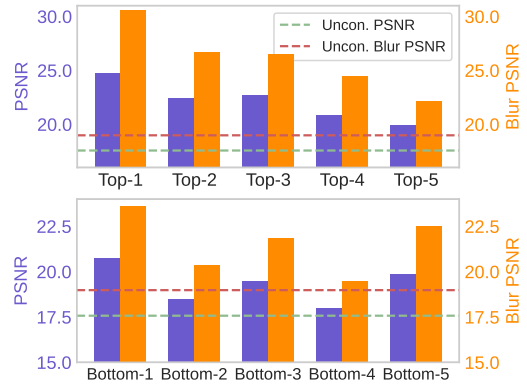


Figure 2: Empirical results of black hole imaging for different particle selections in SCG. The top shows the results obtained by selecting the first few particles in the ranking, while the bottom demonstrates the effect of selecting the last few particles and reversing the sign of their noise.

Leveraging Information from All Particles

While conceptually simple, methods like SCG (Huang et al. 2024) suffer from suboptimal particle efficiency by discarding all but one candidate during the reverse denoising process. We hypothesize that a significant amount of valuable information is lost in this process, as most particles contain insights that can aid in guiding the denoising trajectory. To test this hypothesis, we conduct an empirical study using the black hole imaging inverse problem (Zheng et al. 2025b).

Our results, depicted in Figure 2 (top), provide strong evidence for this hypothesis. Not only does the top-1 particle effectively guide the denoising process towards the observed signal, but the top-2 and top-3 particles also contribute significantly, each outperforming unconditioned sampling. Furthermore, we find a surprising result: even the five worst-performing particles (bottom-1 to bottom-5), when their noise signs are reversed, can still steer the denoising process in the correct direction, as illustrated in Figure 2 (bottom). These findings strongly suggest that information useful for guiding the diffusion process is not confined to a single particle but is distributed across the entire ensemble. Inspired by these empirical observations, we propose a novel method to systematically leverage information from all candidates.

Constrained Particle Seeking (CPS)

The SCG approach can be understood as first sampling from the unconditional transition kernel $p(x_t|x_{t+1})$ and then passively selecting the most promising candidates. In contrast, we take a more active approach by seeking promising particles while incorporating additional constraints that align with high-density regions of the unconditional prior. We reformulate this idea as the following constrained optimization problem:

$$x_t^* = \arg \min_{x_t} \Phi(\hat{x}_{0|t}) \quad \text{s.t. } x_t \sim p(x_t|x_{t+1}) \quad (9)$$

This optimization framework offers greater flexibility, allowing for more effective particle seeking, as it is no longer confined to selecting particles from the sampled set alone.

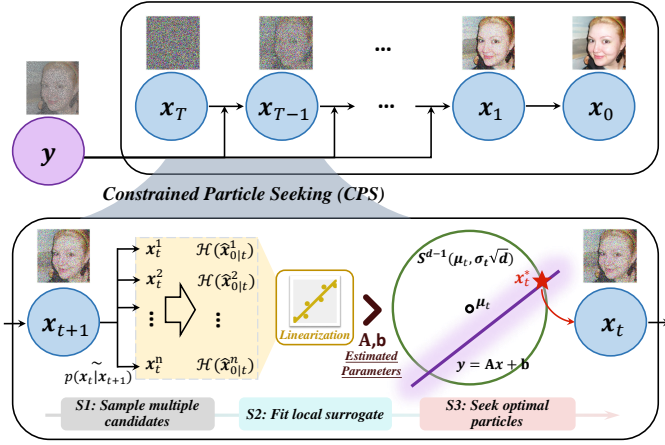


Figure 3: **Overview of Constrained Particle Seeking (CPS) for inverse problem solving.** At each time step, the process involves three steps. *Step 1:* Sample multiple candidate particles; *Step 2:* Fit the forward process surrogate using these samples; *Step 3:* Use the surrogate to determine the optimal particle on the hypersphere, as the next state.

Solving this optimization problem presents two main challenges. First, we must address how to minimize the objective function, particularly when dealing with black-box observation processes. Second, we need to implement the constraint of the unconditional kernel in practice. In the following, we will outline our solutions to these challenges.

Surrogate of observation process. Optimizing the terminal cost typically requires estimating its gradient with respect to \mathbf{x}_t , which can be derived as follows:

$$\nabla_{\mathbf{x}_t} \Phi(\hat{\mathbf{x}}_{0|t}) = \left(\frac{\partial \mathcal{H}(\hat{\mathbf{x}}_{0|t})}{\partial \mathbf{x}_t} \right)^\top (\mathcal{H}(\hat{\mathbf{x}}_{0|t}) - \mathbf{y}) \quad (10)$$

However, computing the term $\frac{\partial \mathcal{H}(\hat{\mathbf{x}}_{0|t})}{\partial \mathbf{x}_t}$ is challenging, as we do not have direct access to the gradient $\mathcal{H}(\cdot)$. To address this, we apply statistical linearization (Roberts and Spanos 2003) to estimate a local surrogate around the region defined by $p(\mathbf{x}_t|\mathbf{x}_{t+1})$. Specifically, we first sample n particles $\mathbf{x}_t^1, \dots, \mathbf{x}_t^n \sim p(\mathbf{x}_t|\mathbf{x}_{t+1})$, as done in SCG, and estimate the corresponding values $\mathcal{H}(\hat{\mathbf{x}}_{0|t}^1), \dots, \mathcal{H}(\hat{\mathbf{x}}_{0|t}^n)$. We then fit these values to the best-fitting linear model with parameters (\mathbf{A}, \mathbf{b}) . This leads to the following optimization problem:

$$\min_{\mathbf{A}, \mathbf{b}} \mathbb{E}_{\mathbf{x}_t \sim p(\mathbf{x}_t|\mathbf{x}_{t+1})} \|\mathcal{H}(\hat{\mathbf{x}}_{0|t}) - (\mathbf{A}\mathbf{x}_t + \mathbf{b})\|^2 \quad (11)$$

The closed-form solution to this problem is given by:

$$\begin{aligned} \mathbf{A} &= \text{Cov}(\mathcal{H}(\hat{\mathbf{x}}_{0|t}), \mathbf{x}_t) \text{Cov}(\mathbf{x}_t, \mathbf{x}_t)^{-1} \\ \mathbf{b} &= \mathbb{E}[\mathcal{H}(\hat{\mathbf{x}}_{0|t})] - \mathbf{A}\mathbb{E}[\mathbf{x}_t] \end{aligned} \quad (12)$$

Given that the transition kernel of the diffusion model is an isotropic Gaussian density, we represent its mean and variance as μ_t and σ_t^2 , respectively. This enables us to further simplify the expressions. Next, we perform empirical estimation using the sampled particles $\mathbf{x}_t^1, \dots, \mathbf{x}_t^n$, yielding the

following expression for \mathbf{A} and \mathbf{b} ,

$$\begin{aligned} \mathbf{A} &= \frac{1}{n\sigma_t^2} \sum_{i=1}^n (\mathcal{H}(\hat{\mathbf{x}}_{0|t}^i) - \bar{\mathcal{H}})(\mathbf{x}_t^i - \mu_t)^\top \\ \mathbf{b} &= \bar{\mathcal{H}} - \mathbf{A}\mu_t \end{aligned} \quad (13)$$

where $\bar{\mathcal{H}}$ denotes $\frac{1}{n} \sum_{i=1}^n \mathcal{H}(\hat{\mathbf{x}}_{0|t}^i)$ for simplicity. At this point, we can use the surrogate model to replace the non-differentiable part of the terminal cost. Furthermore, we observe that \mathbf{A} in Eq. 13 is effectively a weighted linear combination of the sampled particles, integrating all candidate information. This enables us to select particles that outperform any individual sample, thus facilitating a more effective identification of the next state.

Constraint of transition kernel. In the previous section, we introduced the use of a linear surrogate to estimate terminal costs, enabling us to optimize the objective. Eq. 9 also imposes the condition that the optimization variables be sampled from an unconditional distribution, ensuring that the surrogate remains within a credible range and that the model does not deviate from the prior of the base model during inference.

The transition kernel $p(\mathbf{x}_t|\mathbf{x}_{t+1})$ is a Gaussian distribution with a tractable log-likelihood function and it tends to push \mathbf{x}_t toward the mean. As shown by Menon et al. (2020), the norm of samples from a d -dimensional Gaussian distribution follows a χ^d distribution. For the high-dimensional spaces common in diffusion models, the mass of the χ^d distribution is highly concentrated around the value $\sigma_t\sqrt{d}$, forming a hypersphere $S^{d-1}(\mu_t, \sigma_t\sqrt{d})$ (Samuel et al. 2023). Diffusion models generally prefer input data that follows a norm-based distribution aligned with this characteristic, rather than being centered around the mean. Therefore, we aim to keep \mathbf{x}_t within the high-mass region of $p(\mathbf{x}_t|\mathbf{x}_{t+1})$. With this in mind, the original optimization problem (Eq. 9) can be interpreted as follows, using the surrogate model:

$$\min_{\mathbf{x}_t} \|\mathbf{y} - (\mathbf{A}\mathbf{x}_t + \mathbf{b})\|^2 \quad \text{s.t. } \mathbf{x}_t \in S^{d-1}(\mu_t, \sigma_t\sqrt{d}) \quad (14)$$

The aforementioned constrained optimization problem can be solved using the Lagrange multiplier method (Bertsekas 2014), leading to the following system of equations for the optimal solution:

$$\begin{cases} \mathbf{A}^\top(\mathbf{A}\mathbf{x}_t + \mathbf{b} - \mathbf{y}) + \lambda(\mathbf{x}_t - \mu_t) = 0 \\ \|\mathbf{x}_t - \mu_t\|^2 - \sigma_t^2 d = 0 \end{cases} \quad (15)$$

where λ is the Lagrange multiplier. Solving this system directly involves matrix inversion and root finding, which becomes challenging in high-dimensional spaces. We note that the σ_t of the reverse kernel is typically small, leading to constraints that dominate when seeking the optimal solution. As a result, the multiplier λ is expected to be sufficiently large. Under these conditions, we can analytically derive the asymptotic form of the optimal solution:

$$\mathbf{x}_t^* \approx \mu_t + \sigma_t\sqrt{d} \frac{\mathbf{A}^\top(\mathbf{y} - \bar{\mathcal{H}})}{\|\mathbf{A}^\top(\mathbf{y} - \bar{\mathcal{H}})\|} \quad (16)$$




Figure 4: Evolution of Jensen gap of CPS and EnKG for $4\times$ super-resolution on the FFHQ dataset.

A detailed derivation can be found in Appendix A. Ultimately, the solution \mathbf{x}_t^* can be viewed as the optimal particle synthesized by combining all candidate information. In comparison to SCG, which merely selects the particle with the minimum terminal cost, our method is more principled and efficient. Figure 3 provides an overview of our approach.

Note that EnKG (Zheng et al. 2025a) employs a similar linearization strategy. However, EnKG performs linearization over the entire global space, which can lead to inaccurate approximations and difficulties when inverting the covariance matrix (Eq.12). In contrast, our approach utilizes local linearization around $p(\mathbf{x}_t|\mathbf{x}_{t+1})$. In Figure 4, we compute the mean absolute error (MAE) of the Jensen gap (Gao, Sitharam, and Roitberg 2017) for the function $\mathcal{H}(\mathbb{E}[\mathbf{x}_0|\cdot])$ using sampled particles to assess the degree of deviation from linear behavior. Our results demonstrate that EnKG exhibits significant errors in the early stages of generation, whereas CPS consistently maintains small errors. Additionally, CPS seeks one optimal particle at each step by synthesizing candidates, without explicitly maintaining the entire particle set as in EnKG. Our ablation in the experiment section reveals that CPS is significantly more efficient than those generated by the latter.

Improving Robustness with Restart

In practice, we observe that different initial noise conditions can lead to varied outcomes in image restoration tasks. When the initial noise is suboptimal, large adjustments are needed to correct for the global mismatch. On top of that, the one-step estimation from Tweedie’s formula can introduce approximation errors that hurt performance.

To address these challenges, we incorporate the recently proposed Restart strategy (Xu et al. 2023; Lugmayr et al. 2023). This method provides a powerful way to mitigate the accumulation of errors throughout the sampling process. The procedure is a simple, two-step iteration: at time step t , the sample is re-noised according to Eq. 2, then returned to the previous timestep $t + 1$. This re-noised sample is subsequently reused with CPS for guided denoising. Such a step-wise approach allows for the gradual correction of cumulative errors, thereby enhancing the robustness of CPS and ensuring more reliable solutions. The full sampling process, including the Restart strategy, is detailed in Algorithm 1.

Algorithm 1: Sampling procedure for CPS

Require: Observation \mathbf{y} , initial noise $\mathbf{x}_T \sim \mathcal{N}(0, I)$, pre-trained diffusion model $p(\mathbf{x}_t|\mathbf{x}_{t+1}) = \mathcal{N}(\mu_t, \sigma_t I)$, data dimension d , and number of restart N_r

- 1: **for** $t = T - 1, \dots, 0$ **do**
- 2: **for** $r = N_r, \dots, 1$ **do**
 - ▷ Sample from reverse kernel
- 3: $\mathbf{x}_t^1, \dots, \mathbf{x}_t^n \sim p(\mathbf{x}_t|\mathbf{x}_{t+1})$
 - ▷ Fit the forward process surrogate
- 4: $\mathbf{A} \leftarrow \frac{1}{n\sigma_t^2} \sum_{i=1}^n (\mathcal{H}(\hat{\mathbf{x}}_{0|t}^i) - \bar{\mathcal{H}})(\mathbf{x}_t^i - \mu_t)^\top$
 - ▷ Solve optimal particle as next state
- 5: $\mathbf{x}_t \leftarrow \mu_t + \sigma_t \sqrt{d} \frac{\mathbf{A}^\top (\mathbf{y} - \bar{\mathcal{H}})}{\|\mathbf{A}^\top (\mathbf{y} - \bar{\mathcal{H}})\|}$
- 6: **if** $r > 1$ **then**
 - ▷ Re-noise for restart
- 7: $\mathbf{x}_{t+1} \leftarrow p(\mathbf{x}_{t+1}|\mathbf{x}_t)$

Ensure: \mathbf{x}_0

Experiments

To validate the effectiveness of our method, we conduct experiments on image inverse problems and two scientific inverse problems. Our baseline includes several recently proposed gradient-free methods for solving inverse problems, such as SCG (Huang et al. 2024), DPG (Tang et al. 2024), and EnKG (Zheng et al. 2025a). All gradient-free methods are evaluated using a comparable number of forward passes to ensure a fair comparison. Additionally, we incorporate advanced gradient-based methods for comparison in the linear image problem, including DDRM (Kawar et al. 2022a), DPS (Chung et al. 2022), IIGDM (Song et al. 2023), RED-diff (Mardani et al. 2023), and DAPS (Zhang et al. 2025). The formulations of the inverse problems tested are detailed in Appendix B.1. Further implementation details can be found in Appendix B.2.

Image Inverse Problem

Experimental Setup: We begin by evaluating our method on the image inverse problem. We sample 100 images from the FFHQ dataset (Karras, Laine, and Aila 2019) for validation and use the checkpoint provided by Chung et al. (2022). All images were resized to 256×256 pixels, and their quality was assessed using three standard image restoration metrics: PSNR, SSIM, and LPIPS (Zhang et al. 2018).

We focus on three linear inverse problems: 95% missing random mask completion, $4\times$ super-resolution, and deblurring with a Gaussian blur kernel of standard deviation 3 (61×61). To assess the practicality of gradient-free methods, we also include JPEG restoration tasks (Song et al. 2023), where the forward model is non-differentiable. For the JPEG restoration, we consider the quality factor (QF) which is 5. All measurements included additive Gaussian noise with $\sigma_y = 0.05$.

Result: Table 1 and Figure 5 present the quantitative results and visualizations, respectively. Our findings show that SCG fails to handle image inverse problems effectively. In contrast, CPS significantly outperforms gradient-free methods and competes well with gradient-based methods. To the best

Method	Inpaint (95%)			Super-resolution ($\times 4$)			Deblur (gauss)			JPEG (QF=5)			
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	
Gradient-based	DDRM	21.65	0.693	0.270	28.02	0.828	0.213	24.69	0.753	0.187	-	-	-
	DPS	22.49	0.701	0.236	27.65	0.822	0.120	25.59	0.790	0.136	-	-	-
	Π GDM	23.13	0.774	0.231	27.25	0.829	0.135	26.76	0.812	0.141	18.25	0.593	0.304
	RED-diff	24.43	0.807	0.222	26.85	0.835	0.191	28.01	0.853	0.164	-	-	-
	DAPS	22.71	0.754	0.226	28.53	0.840	0.132	28.70	0.843	0.133	-	-	-
Gradient-free	SCG	6.15	0.321	0.807	6.16	0.319	0.805	7.01	0.322	0.792	6.17	0.362	0.825
	DPG	19.83	0.651	0.299	24.81	0.758	0.187	24.11	0.740	0.196	22.52	0.734	0.246
	EnKG	22.64	0.767	0.286	21.72	0.737	0.293	20.84	0.731	0.334	20.89	0.704	0.257
	CPS (Ours)	24.90	0.794	0.161	27.70	0.834	0.115	27.21	0.822	0.128	23.23	0.784	0.223

Table 1: Quantitative results of different methods on FFHQ, including three linear and a non-differentiable inverse problems. Best results are highlighted as **first**, **second** and **third**. We mark “-” for those not applicable.




Figure 5: Qualitative results of image inverse problems on FFHQ: from top to bottom, inpainting, super-resolution, deblurring, and JPEG restoration.

of our knowledge, Π GDM is the only method capable of performing JPEG restoration through pseudo-inverse construction (Song et al. 2023). However, we observe that Π GDM struggles in noisy scenes, with its performance being notably poor. This emphasizes the importance of further developing gradient-free inversion methods.

Method	100%		10%		3%	
	PSNR \uparrow	BPSNR \uparrow	PSNR \uparrow	BPSNR \uparrow	PSNR \uparrow	BPSNR \uparrow
SCG	25.45	30.56	24.97	31.45	24.65	31.28
DPG	14.51	16.75	13.57	16.35	12.16	13.71
EnKG	23.03	25.67	22.77	25.18	21.41	25.14
CPS (Ours)	25.74	31.84	25.34	31.54	25.03	31.81

Table 2: Comparisons of gradient-free methods on black hole imaging problem, including three observation time ratios: 100%, 10% and 3%.

Black Hole Imaging

Experimental Setup: Black hole imaging presents a highly nonlinear and ill-posed forward model due to sparse observations. Measurements are obtained via Very Long Baseline Interferometry (VLBI), where each telescope pair provides visibility data at specific times. To reduce errors from at-




Figure 6: Qualitative results of black hole imaging.

mospheric turbulence and thermal noise, multiple visibility data points are typically combined into *closure phases* and *log closure amplitudes* (Blackburn et al. 2020). In addition to these constraints, the *total flux* of the black hole is often incorporated into the likelihood function. For evaluation, we assume black-box access to the forward model. We employ a pretraining diffusion model from Zheng et al. (2025b) and evaluate it on 100 resized 64×64 images from the test set. We consider the settings of three observation time ratios: 100%, 10% and 3%. The reconstruction quality is measured by calculating the PSNR between the reconstructed image and the ground truth, along with the Blur PSNR (BPSNR), which is smoothed using a filter to assess low-frequency reconstruction performance.

Result: Figure 6 compares black hole images reconstructed using the CPS method and a gradient-free method. The CPS results are visually more consistent with the ground truth. Table 2 presents the quantitative comparison, with CPS achieving the best performance, followed closely by SCG. In contrast, DPG performs poorly when faced with such highly nonlinear problems.

Fluid Data Assimilation

Experimental Setup: Fluid dynamics are typically modeled by the Navier-Stokes (NS) equations, a set of nonlinear partial differential equations (PDEs) (Chandler and Kerswell 2013). In this experiment, we perform fluid assimilation on sparse and noisy vorticity measurements at




Figure 7: Comparison of particle efficiency of gradient-free method across three tasks.

Method	$\times 2$	$\times 4$	$\times 8$
	Relative $L_2 \downarrow$	Relative $L_2 \downarrow$	Relative $L_2 \downarrow$
SCG	1.087	1.095	1.123
DPG	0.491	0.592	0.836
EnK	0.320	0.528	0.821
CPS (Ours)	0.295	0.469	0.684

Table 3: Comparisons of gradient-free methods on fluid data assimilation, including three types of sparse observations: $2\times$, $4\times$, and $8\times$ downscaling.

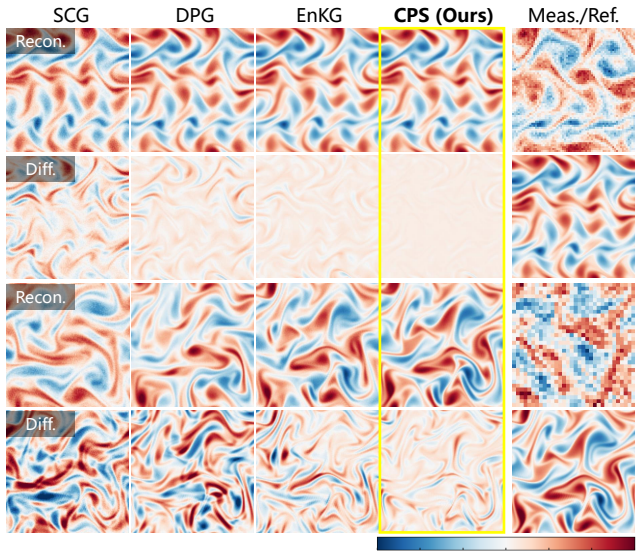


Figure 8: Qualitative results of recovering initial vorticity fields. Odd rows show the reconstruction, even rows show the difference from the reference.

$T = 1$, with the goal of recovering the 2D initial vorticity field. Given that PDE numerical solvers may require thousands of discrete time steps, obtaining accurate numerical derivatives through automatic differentiation presents significant challenges. This makes the problem an ideal candidate for validating gradient-free methods. We use the pre-trained checkpoint provided by Zheng et al. (2025b) and 10 vorticity fields with size 128×128 for testing. The result is evaluated using the relative L_2 error. We consider three settings to simulate sparse measurements, including $2\times$, $4\times$, and $8\times$ downscaling. All measurements included noise with $\sigma_y = 2.0$.

Result: Figure 8 presents both a simple example and a challenging visualization result. In the simple example, all methods are able to restore the overall structure effectively. However, in more difficult scenarios, CPS significantly outperforms the baseline method by recovering the initial vorticity

field pattern from sparse and noisy measurements. SCG and DPG often fail under these conditions. EnK performs second best, but with noticeable errors in some local areas. The quantitative results in Table 3 further support our findings.

Ablation Study

Particle Efficiency Comparison: Gradient-free methods typically require sampling multiple particles. In this study, we evaluate the particle efficiency of these methods using particle numbers of 8, 16, 32, 64, and 128. As shown in Figure 7, CPS demonstrates significantly higher particle efficiency compared to the baseline method, maintaining stability and competitiveness across various tasks. Even with just 8 particles, CPS performs well. In contrast, EnK struggles with small particle sets and requires a larger number of particles to function effectively. DPG faces difficulties with highly ill-posed black hole imaging, while SCG fails to address image-related problems.

Effectiveness of Restart: We investigate the impact of the Restart strategy in CPS, using different random seeds to initialize the noise. As shown in Figure 9, the Restart strategy not only improves overall performance but also enhances the stability and robustness of our algorithm. It effectively addresses the issue of sampling from poor initial noise (Column 3), preventing failures during the inversion process.




Figure 9: Ablation study of the Restart strategy with different random seed initializations.

Conclusion

In this work, we introduced *Constrained Particle Seeking (CPS)*, a novel gradient-free method for solving inverse problems with diffusion models. CPS uses information from all candidate particles and applies constraints based on high-density regions of the prior, enables efficient particle seeking. Our experiments show that CPS outperforms other gradient-free methods and achieves results similar to gradient-based approaches. These results demonstrate CPS’s potential as a reliable solution when gradients of the forward process are not available.

References

- Bertsekas, D. P. 2014. *Constrained optimization and Lagrange multiplier methods*. Academic press.
- Blackburn, L.; Pesce, D. W.; Johnson, M. D.; Wielgus, M.; Chael, A. A.; Christian, P.; and Doeleman, S. S. 2020. Closure statistics in interferometric data. *The Astrophysical Journal*, 894(1): 31.
- Chandler, G. J.; and Kerswell, R. R. 2013. Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *Journal of Fluid Mechanics*, 722: 554–595.
- Chung, H.; Kim, J.; McCann, M. T.; Klasky, M. L.; and Ye, J. C. 2022. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*.
- Duong, C. N.; Truong, T.-D.; Luu, K.; Quach, K. G.; Bui, H.; and Roy, K. 2020. Vec2face: Unveil human faces from their blackbox features in face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6132–6141.
- Efron, B. 2011. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496): 1602–1614.
- Gao, X.; Sitharam, M.; and Roitberg, A. E. 2017. Bounds on the jensen gap, and implications for mean-concentrated distributions. *arXiv preprint arXiv:1712.05267*.
- Gillhofer, M.; Ramsauer, H.; Brandstetter, J.; Schäfl, B.; and Hochreiter, S. 2019. A GAN based solver of black-box inverse problems. In *NeurIPS 2019 Workshop on Solving Inverse Problems with Deep Networks*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, Y.; Ghatare, A.; Liu, Y.; Hu, Z.; Zhang, Q.; Sastry, C. S.; Gururani, S.; Oore, S.; and Yue, Y. 2024. Symbolic Music Generation with Non-Differentiable Rule Guided Diffusion. In *International Conference on Machine Learning*, 19772–19797. PMLR.
- Iordache, M.-D.; Bioucas-Dias, J. M.; and Plaza, A. 2012. Total variation spatial regularization for sparse hyperspectral unmixing. *IEEE Transactions on Geoscience and Remote Sensing*, 50(11): 4484–4502.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4401–4410.
- Katzfuss, M.; Stroud, J. R.; and Wikle, C. K. 2016. Understanding the ensemble Kalman filter. *The American Statistician*, 70(4): 350–357.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022a. Denoising diffusion restoration models. *Advances in neural information processing systems*, 35: 23593–23606.
- Kawar, B.; Song, J.; Ermon, S.; and Elad, M. 2022b. Jpeg artifact correction using denoising diffusion restoration models. *arXiv preprint arXiv:2209.11888*.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; and Anandkumar, A. 2020. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. R. 2023. Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11461–11471.
- Mardani, M.; Song, J.; Kautz, J.; and Vahdat, A. 2023. A variational perspective on solving inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391*.
- McCann, M. T.; Jin, K. H.; and Unser, M. 2017. Convolutional neural networks for inverse problems in imaging: A review. *IEEE Signal Processing Magazine*, 34(6): 85–95.
- Menon, S.; Damian, A.; Hu, S.; Ravi, N.; and Rudin, C. 2020. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2437–2445.
- Oliver, D. S.; Reynolds, A. C.; and Liu, N. 2008. *Inverse theory for petroleum reservoir characterization and history matching*.
- Pavon, M. 1989. Stochastic control and nonequilibrium thermodynamical systems. *Applied Mathematics and Optimization*, 19(1): 187–202.
- Roberts, J. B.; and Spanos, P. D. 2003. *Random vibration and statistical linearization*. Courier Corporation.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Samuel, D.; Ben-Ari, R.; Darshan, N.; Maron, H.; and Chechik, G. 2023. Norm-guided latent space exploration for text-to-image generation. *Advances in Neural Information Processing Systems*, 36: 57863–57875.
- Schwarzbach, C.; Börner, R.-U.; and Spitzer, K. 2011. Three-dimensional adaptive higher order finite element simulation for geo-electromagnetics—a marine CSEM example. *Geophysical Journal International*, 187(1): 63–74.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, J.; Vahdat, A.; Mardani, M.; and Kautz, J. 2023. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*.
- Song, Y.; Shen, L.; Xing, L.; and Ermon, S. 2021. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

Tang, H.; Xie, T.; Feng, A.; Wang, H.; Zhang, C.; and Bai, Y. 2024. Solving general noisy inverse problem via posterior sampling: A policy gradient viewpoint. In *International Conference on Artificial Intelligence and Statistics*, 2116–2124. PMLR.

Theodorou, E.; Buchli, J.; and Schaal, S. 2010. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11: 3137–3181.

Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1): 267–288.

Tonolini, F.; Radford, J.; Turpin, A.; Faccio, D.; and Murray-Smith, R. 2020. Variational inference for computational imaging inverse problems. *Journal of Machine Learning Research*, 21(179): 1–46.

Vincent, P. 2011. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7): 1661–1674.

Wang, Y.; Yu, J.; and Zhang, J. 2022. Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490*.

Williams, G.; Aldrich, A.; and Theodorou, E. A. 2017. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40(2): 344–357.

Xu, Y.; Deng, M.; Cheng, X.; Tian, Y.; Liu, Z.; and Jaakkola, T. 2023. Restart sampling for improving generative processes. *Advances in Neural Information Processing Systems*, 36: 76806–76838.

Zhang, B.; Chu, W.; Berner, J.; Meng, C.; Anandkumar, A.; and Song, Y. 2025. Improving diffusion inverse problem solving with decoupled noise annealing. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 20895–20905.

Zhang, Q.; and Chen, Y. 2023. Path Integral Sampler: A Stochastic Control Approach For Sampling. In *International Conference on Learning Representations*.

Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.

Zheng, H.; Chu, W.; Wang, A.; Kovachki, N. B.; Baptista, R.; and Yue, Y. 2025a. Ensemble Kalman Diffusion Guidance: A Derivative-free Method for Inverse Problems. *Transactions on Machine Learning Research*.

Zheng, H.; Chu, W.; Zhang, B.; Wu, Z.; Wang, A.; Feng, B.; Zou, C.; Sun, Y.; Kovachki, N. B.; Ross, Z. E.; et al. 2025b. InverseBench: Benchmarking Plug-and-Play Diffusion Priors for Inverse Problems in Physical Sciences. In *The Thirteenth International Conference on Learning Representations*.

Reproducibility Checklist

1. General Paper Structure

- 1.1. Includes a conceptual outline and/or pseudocode description of AI methods introduced (yes/partial/no/NA) **yes**
- 1.2. Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results (yes/no) **yes**
- 1.3. Provides well-marked pedagogical references for less-familiar readers to gain background necessary to replicate the paper (yes/no) **yes**

2. Theoretical Contributions

- 2.1. Does this paper make theoretical contributions? (yes/no) **yes**

If yes, please address the following points:

- 2.2. All assumptions and restrictions are stated clearly and formally (yes/partial/no) **yes**
- 2.3. All novel claims are stated formally (e.g., in theorem statements) (yes/partial/no) **yes**
- 2.4. Proofs of all novel claims are included (yes/partial/no) **yes**
- 2.5. Proof sketches or intuitions are given for complex and/or novel results (yes/partial/no) **yes**
- 2.6. Appropriate citations to theoretical tools used are given (yes/partial/no) **yes**
- 2.7. All theoretical claims are demonstrated empirically to hold (yes/partial/no/NA) **yes**
- 2.8. All experimental code used to eliminate or disprove claims is included (yes/no/NA) **NA**

3. Dataset Usage

- 3.1. Does this paper rely on one or more datasets? (yes/no) **yes**

If yes, please address the following points:

- 3.2. A motivation is given for why the experiments are conducted on the selected datasets (yes/partial/no/NA) **yes**
- 3.3. All novel datasets introduced in this paper are included in a data appendix (yes/partial/no/NA) **NA**
- 3.4. All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no/NA) **NA**
- 3.5. All datasets drawn from the existing literature (potentially including authors' own previously pub-

lished work) are accompanied by appropriate citations (yes/no/NA) **yes**

- 3.6. All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available (yes/partial/no/NA) **yes**
- 3.7. All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying (yes/partial/no/NA) **NA**

4. Computational Experiments

- 4.1. Does this paper include computational experiments? (yes/no) **yes**

If yes, please address the following points:

- 4.2. This paper states the number and range of values tried per (hyper-) parameter during development of the paper, along with the criterion used for selecting the final parameter setting (yes/partial/no/NA) **yes**
- 4.3. Any code required for pre-processing data is included in the appendix (yes/partial/no) **yes**
- 4.4. All source code required for conducting and analyzing the experiments is included in a code appendix (yes/partial/no) **yes**
- 4.5. All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no) **yes**
- 4.6. All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes/partial/no) **yes**
- 4.7. If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results (yes/partial/no/NA) **NA**
- 4.8. This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks (yes/partial/no) **yes**
- 4.9. This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics (yes/partial/no) **yes**
- 4.10. This paper states the number of algorithm runs used to compute each reported result (yes/no) **yes**
- 4.11. Analysis of experiments goes beyond single-

dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information (yes/no) **yes**

- 4.12. The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank) (yes/partial/no) **yes**
- 4.13. This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments (yes/partial/no/NA) **yes**