ANTITHETIC NOISE IN DIFFUSION MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

We initiate a systematic study of antithetic initial noise in diffusion models, discovering that pairing each noise sample with its negation consistently produces strong negative correlation. This universal phenomenon holds across datasets, model architectures, conditional and unconditional sampling, and even other generative models such as VAEs and Normalizing Flows. To explain it, we combine experiments and theory and propose a symmetry conjecture that the learned score function is approximately affine antisymmetric (odd symmetry up to a constant shift), supported by empirical evidence. This negative correlation leads to substantially more reliable uncertainty quantification with up to 90\% narrower confidence intervals. We demonstrate these gains on tasks including estimating pixel-wise statistics and evaluating diffusion inverse solvers. We also provide extensions with randomized quasi-Monte Carlo noise designs for uncertainty quantification, and explore additional applications of the antithetic noise design to improve image editing and generation diversity. Our framework is training-free, model-agnostic, and adds no runtime overhead. Code is available at https://anonymous.4open. science/r/Antithetic-Noise-in-Diffusion-Models-8B54.

1 Introduction

Diffusion models have set the state of the art in photorealistic image synthesis, high-fidelity audio, and video generation (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b; Kong et al., 2021); they also power applications such as text-to-image generation (Rombach et al., 2022), image editing and restoration (Meng et al., 2022), and inverse problem solving (Song et al., 2022).

For many pretrained diffusion models, sampling relies on three elements: the network weights, the denoising schedule, and the initial Gaussian noise. Once these are fixed, sampling is often deterministic: the sampler transforms the initial noise into an image via successive denoising passes.

Much of the literature improves the first two ingredients and clusters into two strands: (i) architectural and training developments, which improve sample quality and scalability through backbone or objective redesign (e.g., EDM (Karras et al., 2022), Latent Diffusion Models (Rombach et al., 2022), DiT (Peebles & Xie, 2023)), (ii) accelerated sampling, which reduces the number of denoising steps while retaining high-quality generation (e.g., DDIM (Song et al., 2021a), Consistency Models (Song et al., 2023), DPM-Solver++ (Lu et al., 2023), and progressive distillation (Salimans & Ho, 2022)).

However, the third ingredient—the *initial Gaussian noise*—has received comparatively little attention. Prior work has optimized initial noise for generation quality, editing, controllability, or inverse problem solving (Guo et al., 2024; Qi et al., 2024; Zhou et al., 2024; Ban et al., 2025; Chen et al., 2024a; Eyring et al., 2024; Song et al., 2025; Wang et al., 2024; Chihaoui et al., 2024). However, most of these efforts are task-specific. A systematic understanding of how noise itself shapes diffusion model outputs is still missing.

Our perspective is orthogonal to prior work. Our central discovery is both *simple and universal*: pairing every Gaussian noise z with its negation -z—known as *antithetic sampling* (Owen, 2013)—consistently produces samples that are **strongly negatively correlated**. This phenomenon holds regardless of architecture, dataset, sampling schedule, and both conditional and unconditional sampling. It further extends to other generative models such as VAEs and Normalizing Flows. We explain this phenomenon through both experiments and theory. This leads to a symmetry conjecture that the score function is approximately affine antisymmetric, providing a new structural insight supported by empirical evidence.

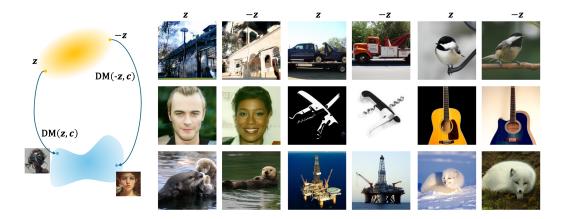


Figure 1: Use antithetic noise -z and z (with condition c) to generate visually "opposite" images.

This universal property has direct impact on uncertainty quantification and also enables additional applications:

- (i) Sharper uncertainty quantification. Antithetic pairs naturally act as control variates, enabling significant variance reduction and thus sharper uncertainty quantification. Our antithetic estimator delivers up to 90% tighter confidence intervals and cuts computation cost by more than 100 times. The efficiency gain immediately leads to huge cost savings in a variety of tasks, including bias detection in generation and diffusion inverse solver evaluation, as we demonstrate in later experiments.
- (ii) Other applications. Because each antithetic noise pair drives reverse-diffusion trajectories toward distant regions of the data manifold (see Figure 1), the paired-sampling scheme increases diversity "for free" while preserving high image quality, as confirmed by SSIM and LPIPS in our experiments. Moreover, algorithms that rely on intermediate sampling or approximation steps also benefit from the improved reliability provided by antithetic noise. As an illustration, we present an image editing example in C.2 and show that our antithetic design serves as a *plug-and-play* tool to improve performance at no additional cost.

Building on the antithetic pairs, we generalize the idea to apply quasi-Monte Carlo (QMC) and randomized QMC (RQMC). The resulting RQMC estimator often delivers further variance reduction. Although QMC has been widely used in computer graphics (Keller, 1995; Waechter & Keller, 2011), quantitative finance (Joy et al., 1996; L'Ecuyer, 2009), and Bayesian inference (Buchholz et al., 2018; Liu & Owen, 2021), this is, to our knowledge, its first application to diffusion models.

In summary, we discover a universal property of initial noise, reveal a new symmetry in score networks, and demonstrate concrete benefits in various practical applications. This positions initial noise manipulations as a simple, training-free tool for advancing generative modeling.

The remainder of the paper is organized as follows. Section 2 defines the problem and outlines our motivation. Section 3 presents our central finding that antithetic noise pairs produce strongly negatively correlated outputs. We offer both theoretical and empirical explanations for this phenomenon, and present the *symmetry conjecture* in Section 3.2. Section 4 develops estimators and their confidence intervals via antithetic sampling, and extends the approach to QMC. Section 5 reports experiments on the aforementioned applications. Section 6 discusses our method and outlines future directions. Appendix B, C, and D include proofs, additional experiments and detailed setups, and supplemental visualizations, respectively.

2 SETUP, MOTIVATION, AND RELATED WORKS

Unconditional diffusion model: A diffusion model aims to generate samples from an unknown data distribution p_0 . It first *noises* data towards a standard Gaussian progressively, then learns to reverse the process so that Gaussian noise can be *denoised* step-by-step back into target samples. The forward process simulates a stochastic differential equation (SDE): $d\mathbf{x}_t = \mu(\mathbf{x}_t, t)dt + \sigma_t d\mathbf{w}_t$, where $\{\mathbf{w}_t\}_{t=0}^T$ denotes the standard Brownian motion, and $\mu(\mathbf{x}, t), \sigma_t$ are chosen by the users. Let

 p_t denote the distribution of \mathbf{x}_t . Song et al. (2021b) states that if we sample $\mathbf{y}_T \sim p_T$ and simulate the probability-flow ordinary differential equation (PF-ODE) backward from time T to 0 as

$$d\mathbf{y}_t = \left(-\mu(\mathbf{y}_t, t) - \frac{1}{2}\sigma_t^2 \nabla \log p(\mathbf{y}_t, t)\right) dt, \tag{1}$$

then for every time t the marginal distribution of \mathbf{y}_t coincides with p_t . Thus, in an idealized world, one can perfectly sample from p_0 by simulating the PF-ODE (1).

In practice, the score function $\nabla \log p_t(\mathbf{x},t)$ is unavailable, and a neural network $\epsilon_{\theta}^{(t)}(\mathbf{x})$ is trained to approximate it, where θ denotes its weights. Therefore, one can generate new samples from p_0 by first sampling a Gaussian noise and simulating the PF-ODE (1) through a numerical integrator from T to 0. For example, DDIM (Song et al., 2021a) has the (discretized) forward process as $\mathbf{x}_k \mid \mathbf{x}_{k-1} \sim \mathbb{N}(\sqrt{\alpha_k}\mathbf{x}_{k-1}, (1-\alpha_k)I)$ for $k=0,\ldots,T-1$, and backward sampling process as

$$\mathbf{y}_{T} \sim \mathbb{N}(0, I), \ \mathbf{y}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{y}_{t} - \sqrt{1 - \alpha_{t}} \, \epsilon_{\theta}^{(t)}(\mathbf{y}_{t})}{\sqrt{\alpha_{t}}} \right) + \sqrt{1 - \alpha_{t-1}} \, \epsilon_{\theta}^{(t)}(\mathbf{y}_{t}). \tag{2}$$

Once θ is fixed, the randomness in DDIM sampling comes solely from the initial Gaussian noise.

We remark that samples from p_0 can also be drawn by simulating the backward SDE with randomness at each step, as in the DDPM sampler (Ho et al., 2020; Song et al., 2021b). Throughout the main text, we focus on deterministic samplers such as DDIM to explain our idea. In Appendix C.4, we present additional experiments showing that our findings also extend to stochastic samplers like DDPM.

Text–conditioned latent diffusion: In Stable Diffusion and its successors SDXL and SDXL Turbo (Podell et al., 2024; Sauer et al., 2024), a pretrained VAE first compresses each image to a latent tensor z. A text encoder embeds the prompt as $c \in \mathbb{R}^m$. During training, the network $\epsilon_{\theta}^{(t)}$ receives (z_t,c) and learns $\nabla_{z_t} \log p_t(z_t \mid c)$. At generation time we draw latent noise $z_T \sim \mathbb{N}(0,I)$ and run the reverse diffusion from t=T to 0, yielding a denoised latent z_0 , which the decoder maps back to pixels. Given a prompt c and an initial Gaussian noise, the sampler produces an image from $p_0(\cdot \mid c)$.

Diffusion posterior sampling: Diffusion model can also be used as a prior in inverse problems. Suppose we observe only a partial or corrupted measurement $\mathbf{y}_{obs} = \mathcal{A}(\mathbf{x}) + \text{noise from an unknown signal } \mathbf{x}$. Diffusion posterior sampling aims to sample from the posterior distribution $p(\mathbf{x} \mid \mathbf{y}_{obs}) \propto p(\mathbf{y}_{obs} \mid \mathbf{x}) p(\mathbf{x})$, where $p(\mathbf{x})$ is given by a pretrained diffusion model. Given initial noise z and observed \mathbf{y}_{obs} , diffusion posterior samplers apply a sequence of denoising steps from T to 0, each step resembling (2) but incorporating \mathbf{y}_{obs} (Chung et al., 2023a; Song et al., 2022; 2024).

2.1 MOTIVATION

Beyond examining how variations in the initial noise affect the outputs, several technical considerations motivate our focus on the antithetic noise pair (z, -z) for diffusion models. Let DM denote certain diffusion model's mapping from an initial noise vector to a generated image sample.

- **Preserving quality:** Since $z \sim \mathbb{N}(0, I)$ implies $-z \sim \mathbb{N}(0, I)$, the initial noises z and -z share the same marginal distribution. Consequently, $\mathrm{DM}(z) \stackrel{d}{=} \mathrm{DM}(-z)$, so all per-sample statistics remain unchanged. In other words, negating the initial noise does not degrade generation quality.
- Maximal separation in the noise space: High-dimensional Gaussians concentrate on the sphere of radius $\sqrt{\text{dimension}}$ (Vershynin, 2018), so any draw z and its negation -z lie at opposite poles of that hypersphere. This antipodal pairing represents the maximum possible perturbation in noise space, making it a natural extreme test of the sampler's behavior.
- Measuring (non-)linearity: Recent work has examined how closely score networks approximate linear operators. Empirical studies across various diffusion settings show that these networks behave as locally linear maps; this behavior forms the basis for new controllable sampling schemes (Chen et al., 2024b; Li et al., 2024b; Song et al., 2025). Correlation is a natural choice to measure linearity: if a score network were exactly linear, feeding it with noises z and -z would yield a correlation of -1 between $\mathrm{DM}(z)$ and $\mathrm{DM}(-z)$. Thus, the difference between the observed correlation and -1 provides a direct measure of the network's departure from linearity.

3 NEGATIVE CORRELATION FROM ANTITHETIC NOISE

3.1 Pixel-wise correlations: antithetic vs. independent noise

We compare the similarity between paired images generated under two sampling schemes: PN (positive vs. negative, z vs. -z) and RR (random vs. random, z_1 vs. z_2) under three settings: (1) unconditional diffusion models, (2) class- or prompt-based conditional diffusion models, and (3) generative models beyond diffusion.

For diffusion models, we evaluate both unconditional and conditional cases using publicly available pre-trained checkpoints, with implementation details described in Section 5.1. Notably, we include both traditional U-Net architecture and transformer-based DiT (Peebles & Xie, 2023). For generative models beyond diffusion, we select two representative baselines: an unconditional VAE on MNIST and a conditional Glow model (Kingma & Dhariwal, 2018) on CIFAR-10. The experimental details for VAE and Glow are provided in Appendix C.1.

To quantify similarity, we use two metrics: the standard Pearson correlation and a centralized Pearson correlation. Let $x_{i,1}$ and $x_{i,2}$ denote the flattened pixel values of the two images in the i-th generated pair. The standard Pearson correlation is computed directly between $x_{i,1}$ and $x_{i,2}$. To correct for dataset-level or class-level bias, we also define a centralized correlation. For a dataset, prompt, or class with K pairs generated, we compute the mean $\mu_c = \sum_{i=1}^K (x_{i,1} + x_{i,2})/2K$. The centralized correlation of the i-th pair is defined as the standard Pearson correlation between the centralized images $x_{i,1} - \mu_c$ and $x_{i,2} - \mu_c$. For each comparison, t-test is conducted to assess statistical significance. In all experiments, the resulting p-values are negligible ($< 10^{-10}$), which confirms significance; hence, they are omitted from the presentation.

Additional metrics such as the Wasserstein distance are presented in Appendix C.4.

Results: Table 1 summarizes the statistics in all classes of models. PN pairs consistently show significantly stronger negative correlations than RR pairs, with this contrast also visually evident in their histograms (Figure 2). In addition, centralization strengthens the negative correlation, since it removes shared patterns. For example, in CelebA-HQ, centralization removes global facial structure, while in DiT and Glow, it removes class-specific patterns.

The same behavior appears in both DDPM models and diffusion posterior samplers, which we report in Appendix C.4. These results demonstrate that the negative correlation resulting from antithetic sampling is a universal phenomenon across diverse architectures and conditioning schemes.

Model	Standard (Correlation	Centralized Correlation		
	PN	RR	PN	RR	
LSUN-Church	-0.62 (0.11)	0.08 (0.17)	-0.77 (0.07)	0.00 (0.17)	
CIFAR-10	-0.76 (0.13)	0.05 (0.24)	-0.86 (0.07)	0.00(0.23)	
CelebA-HQ	-0.34 (0.19)	0.25 (0.20)	-0.78 (0.06)	-0.01 (0.20)	
SD 1.5	-0.47 (0.05)	0.05 (0.04)	-0.54 (0.08)	0.00 (0.01)	
DiT	-0.07 (0.27)	0.26 (0.18)	-0.45 (0.08)	0.00 (0.03)	
VAE	0.21 (0.12)	0.42 (0.17)	-0.41 (0.12)	-0.00 (0.24)	
Glow	-0.52 (0.02)	0.08 (0.05)	-0.57 (0.02)	-0.01 (0.02)	

3.2 EXPLANATORY EXPERIMENTS: TEMPORAL CORRELATIONS & SYMMETRY CONJECTURE

We aim to explain the strong negative correlation between DM(Z) and DM(-Z) for $Z \sim \mathbb{N}(0, I)$. Because DM performs iterative denoising through the score network $\epsilon_{\theta}^{(t)}$ (see (2)), we visualize how these correlations evolve over diffusion time-steps in Figure 3.

Throughout the transition from noises to samples, the PN correlation of $\epsilon_{\theta}^{(t)}$ —indicated by the orange, blue, and green dashed lines in Figure 3—starts at -1, stays strongly negative, and only climbs slightly in the final steps. This nearly -1 correlation is remarkable, as it suggests that the score network ϵ_{θ}

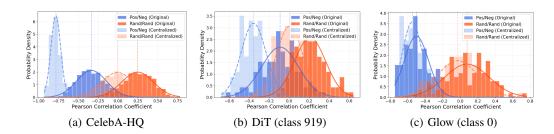


Figure 2: Histograms of standard and centralized Pearson correlation coefficients for CelebA-HQ, and DiT class and Glow in single class. Dashed lines indicate the average.

learns an approximately affine antisymmetric function. To explain, we first state the following results which shows -1 correlation is equivalent to affine antisymmetric, with proof in Appendix B.1.

Lemma 1. Let $Z \sim \mathbb{N}(0, I_d)$, suppose a map $f: \mathbb{R}^d \to \mathbb{R}$ satisfies $\operatorname{Corr}(f(Z), f(-Z)) = -1$, then f must be affine antisymmetric at (0,c) for some c, i.e., $f(\mathbf{x}) + f(-\mathbf{x}) = 2c$ for every \mathbf{x} .

Lemma 1 and Figure 3 leads us to the following conjecture:

Conjecture For each time step t, the score network $\epsilon_{\theta}^{(t)}$ is approximately affine antisymmetric at $(0, c_t)$, i.e.,

$$\epsilon_{\theta}^{(t)}(\mathbf{x}) + \epsilon_{\theta}^{(t)}(-\mathbf{x}) \approx 2\mathbf{c}_t$$

for some fixed vector \mathbf{c}_t that depends on t.

223

224 225

226

227 228 229

230

231

232

233 234

235

236

237 238

239 240

241

242

243

244

245

246

247

248

249

250

251

252

253

254

255

256

257

258

259

260

261

262

263

264

265

266

267

268

269

The conjecture implies that the score network has learned an "almost odd" function up to an additive shift. This is consistent with both theory and observations for large t, where $p_t(\mathbf{x})$ is close to standard Gaussian, thus the score function is almost linear in x. Nonetheless, our conjecture is far more general: it applies to every timestep t and can accommodate genuinely nonlinear odd behaviors (for example, functions like $\sin x$).

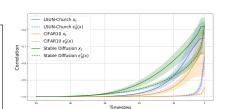


Figure 3: Correlation of x_t (solid) and $\epsilon_{\theta}^{(t)}$ (dash) between antithetic (PN) pairs. Step 50 is the initial noise and Step 0 is the generated image. Shaded bands show ± 1 std. dev.

This conjecture combined with the DDIM update rule (2) immediately explains the negative correlation: The one step iteration of F_t in (2) can be written as a linear combination between x and the output of the score network: $F_t(\mathbf{x}) = a_t \mathbf{x} + b_t \epsilon_{\theta}^{(t)}(\mathbf{x})$. Given the conjecture, $F_t(-\mathbf{x}) \approx -a_t \mathbf{x} - b_t \epsilon_{\theta}^{(t)}(\mathbf{x}) +$ $2b_t \mathbf{c}_t = -F_t(\mathbf{x}) + 2b_t \mathbf{c}_t$, so the onestep DDIM update is affine antisymmetric at $(0, b_t \mathbf{c}_t)$. Thus, beginning with a strongly negatively correlated pair, each DDIM update preserves that strong negative correlation all the way through to the final output.

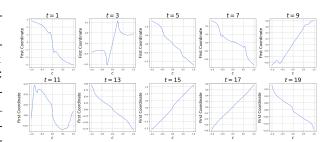


Figure 4: First-coordinate output of the pretrained score network on CIFAR10 as a function of the interpolation scalar c for a 50-step DDIM at $t = 1, 3, \dots, 19$.

To (partly) validate our conjecture, we perform the following experiment. Using a pretrained CIFAR-10 score network with a 50-step DDIM sampler, we picked time steps $t = 1, 3, \dots, 19$ (where smaller t is closer to the image and larger t is closer to pure noise). For each t, we sample $\mathbf{x} \sim \mathbb{N}(0, I_d)$, evaluated the first coordinate of $\epsilon_{\theta}^{(t)}(c\mathbf{x})$ as c varied from -1 to 1 (interpolating from $-\mathbf{x}$ to \mathbf{x}), and plotted the result in Figure 4. We exclude $t \ge 20$, where the curve is very close to a straight line.

Figure 4 confirms that, at every step t, the first coordinate of the score network is indeed overall affine antisymmetric. Although the curves display nonlinear oscillations at small t, the deviations on either side are approximately mirror images, and for larger t the mapping is almost a straight line. The symmetry center c_t usually lies near zero, but not always—for instance, at t=11, $c_t\approx 0.05$ even though the function spans from -0.05 to 0.15. In Appendix C.5, we present further validation experiments using alternative coordinates, datasets, and a quantitative metric called the antisymmetry score. Despite the distinct function shapes across coordinates, the conjectured symmetry still persists.

We further evaluate the antisymmetry conjecture across additional coordinates and datasets. On CIFAR-10 and Church, we assess how closely one-dimensional slices of the network outputs behave like affine-antisymmetric functions. To this end, we introduce the *affine antisymmetry score*, which quantifies the degree of antisymmetry. Across datasets, the resulting antisymmetry scores are consistently high (mean values above 0.99 and even the lower quantiles still near 0.9). These results provide strong empirical support for the conjecture; full experimental details and plots are given in Appendix C.5.

Finally, by applying existing results on the Ornstein–Uhlenbeck process, we provide a theoretical confirmation of our conjecture in Appendix B.2 for the regime of large t.

Both the conjecture and its empirical support could be of independent interest. They imply that the score network learns a function with strong symmetry, even though that symmetry is not explicitly enforced. The conjecture matches the Gaussian linear score approximation in the high-noise regime, which has proven effective (Wang & Vastola, 2024). In the intermediate-to-low noise regime, a single Gaussian approximation performs poorly—this is clear in Figure 4 for t=3,7,11, where the score is strongly nonlinear—yet our conjectured approximate symmetry still holds.

Besides explaining the source of negative correlation, we believe this result provides new insight into the structure of diffusion models, which are often treated as black-box functions. Leveraging this finding for algorithms and applications is an important direction for future work.

For readers interested in a supplementary theoretical perspective, we also provide an expository discussion in Appendix B.4 that connects negative correlation to the FKG inequality.

4 UNCERTAINTY QUANTIFICATION

In uncertainty quantification for a diffusion model DM, the goal is typically to estimate expectations of the form $\mathbb{E}_{z \sim \mathbb{N}(0,I)}\big[S\big(\mathrm{DM}(z)\big)\big]$, where S is a statistic of user's interest. Our goal is to leverage negative correlation to design an estimator with higher accuracy at fixed computation cost of inference.

Standard Monte Carlo: To approximate this expectation, the simplest approach is to draw N independent noises $z_1,\ldots,z_N\sim\mathbb{N}(0,I)$, calculate $S_i=S\big(\mathrm{DM}(z_i)\big)$ for each i and form the standard Monte Carlo (MC) estimator $\hat{\mu}_N^{\mathrm{MC}}\coloneqq\sum_{i=1}^NS_i/N$ by taking their average. A $(1-\alpha)$ confidence interval, denoted as $\mathrm{CI}_N^{\mathrm{MC}}(1-\alpha)$ is then $\hat{\mu}_N^{\mathrm{MC}}\pm z_{1-\alpha/2}\sqrt{(\hat{\sigma}_N^{\mathrm{MC}})^2/N}$, where $\hat{\sigma}_N^2$ is the sample variance of S_1,S_2,\ldots,S_N and $z_{1-\alpha/2}$ is the $(1-\alpha/2)$ -quantile of the standard normal. A formal guarantee of the above construction is given in Proposition 1 in Appendix B.3.

Antithetic Monte Carlo: The observed negative correlation motivates an improved estimator. Let N=2K be even, users can generate K pairs of antithetic noise $(z_1,-z_1),\ldots,(z_K,-z_K)$. Define $S_i^+=S(\mathrm{DM}(z_i)), S_i^-=S(\mathrm{DM}(-z_i)),$ and let $\bar{S}_i=0.5(S_i^++S_i^-)$ be their average. Our Antithetic Monte Carlo (AMC) estimator is $\hat{\mu}_N^{\mathrm{AMC}}:=\sum_{i=1}^K \bar{S}_i/K,$ and confidence interval $\mathrm{CI}_N^{\mathrm{AMC}}(1-\alpha)$ is $\hat{\mu}_N^{\mathrm{AMC}}\pm z_{1-\alpha/2}\sqrt{2(\hat{\sigma}_N^{\mathrm{AMC}})^2/N},$ where $(\hat{\sigma}_N^{\mathrm{AMC}})^2$ is the sample variance of $\bar{S}_1,\ldots,\bar{S}_K$.

The intuition is simple: if the pair (S_i^+, S_i^-) remains negatively correlated, then averaging antithetic pairs can reduce variance by partially canceling out opposing errors—since negative correlation suggests that when one estimate exceeds the true value, the other is likely to fall below it.

The AMC estimator is unbiased, and $\operatorname{CI}_N^{\operatorname{AMC}}(1-\alpha)$ achieves correct coverage as the sample size increases. Let $\rho:=\operatorname{Corr}(S_i^+,S_i^-)$. We can prove that AMC's standard error and its confidence-interval width are each equal to the Monte Carlo counterparts multiplied by the factor $\sqrt{1+\rho}$. Thus, when $\rho<0$, AMC produces **provably lower variance and tighter confidence intervals** than MC, with greater gains as ρ becomes more negative. See Appendix B.3 for a formal statement and proof.

Since both methods have the same computational cost, the variance reduction from negative correlation and the antithetic design yields a direct and cost-free improvement.

Quasi-Monte Carlo: The idea of using negatively correlated samples has been widely adopted in Monte Carlo methods (Craiu & Meng, 2005) and statistical risk estimation (Liu et al., 2024) for improved performance. This principle of variance reduction can be further extended to QMC methods. As an alternative to Monte Carlo, QMC constructs a deterministic set of N samples that have negative correlation and provide a more balanced coverage of the sample space. QMC samples can also be randomized, resulting in RQMC methods. RQMC maintains the marginal distribution of each sample while preserving the low-discrepancy properties of QMC. Importantly, repeated randomizations allow for empirical error estimation and confidence intervals (L'Ecuyer et al., 2023).

Under regularity conditions (Niederreiter, 1992), QMC and RQMC can improve the error *convergence* rate from $O(N^{-1/2})$ to $O(N^{-1+\epsilon})$ for any $\epsilon>0$, albeit ϵ may absorb a dimension-dependent factor $(\log N)^d$. For sufficiently smooth functions, RQMC can achieve rates as fast as $O(N^{-3/2+\epsilon})$ (Owen, 1997b;a). Moreover, the space-filling property of QMC may also promote greater sample diversity.

Although the dimension of the Gaussian noise is higher than the typical regime where QMC is expected to be most effective, RQMC still further shrinks our confidence intervals by several-fold compared with standard MC. This hints that the image generator's effective dimension is much lower than its ambient one, allowing QMC methods to stay useful. Similar gains arise when QMC methods deliberately exploit low-dimensional structure in practical problems (Wang & Fang, 2003; Xiao & Wang, 2019; Liu & Owen, 2023). Developing ways to identify and leverage this structure more systematically is an appealing avenue for future work.

5 EXPERIMENT

Sections 5.1 and 5.2 present two uncertainty quantification applications: estimating pixel-wise statistics and evaluating diffusion inverse problem solvers. The latter considers two popular algorithms across a range of tasks and datasets. For each task, we apply MC, AMC, and, when applicable, QMC estimators as described in Section 4. For each estimator, we report the 95% confidence interval (CI) width and its relative efficiency. The relative efficiency of a new estimator (AMC or QMC) is defined as the squared ratio of the MC CI width to that of the estimator, $(\text{CI}_{MC}/\text{CI}_{new \, estimator})^2$. It reflects how many times more MC samples are required to achieve the same accuracy as our new estimator. Section 5.3 shows that antithetic noise produces more diverse images than independent noise while preserving quality, and Appendix C.2 explores an image editing example.

5.1 PIXEL-WISE STATISTICS

We begin by evaluating four pixel-level statistics. These include the (i) pixel value mean, (ii) perceived brightness, (iii) contrast, and (iv) the image centroid, with definitions in Appendix C.6. These statistics are actively used in diffusion workflows for diagnosing artifacts and assessing reliability, with applications in detecting signal leakage (Lin et al., 2024; Everaert et al., 2024), out-of-distribution detection (Le Bellier & Audebert, 2024), identifying artifact-prone regions (Kou et al., 2024), and improving the reliability of weather prediction (Li et al., 2024a).

Setup: We study both unconditional and conditional diffusion models; details are in Appendix C.3:

For *unconditional diffusion models*, we evaluate pre-trained models on CIFAR-10 (Krizhevsky et al., 2009), CelebA-HQ (Xia et al., 2021), and LSUN-Church (Yu et al., 2015). For each dataset, 1,600 image pairs are generated under both PN and RR noise sampling with 50 DDIM steps.

For *conditional diffusion models*, we evaluate Stable Diffusion 1.5 on 200 prompts from the Pick-a-Pic (Kirstain et al., 2023) and DrawBench (Saharia et al., 2022), and DiT (Peebles & Xie, 2023) on 32 ImageNet classes. For each class or prompt, 100 PN and RR pairs are generated with 20 DDIM steps.

Implementation: To ensure fair comparisons under equal sample budget:

For *unconditional diffusion models*: MC uses 3,200 independent samples, AMC uses 1,600 antithetic pairs, and QMC employs a Sobol' point set of size 64 with 50 independent randomizations. Due to the dimensionality limits, QMC is restricted to CIFAR-10.

For *conditional diffusion models*: for each class or prompt, MC uses 200 independent samples, AMC uses 100 antithetic pairs, and QMC employs 8 Sobol' points with 25 randomizations.

CIs for MC and AMC are constructed as described in Section 4. For QMC, we use a student-t interval to ensure reliable coverage (L'Ecuyer et al., 2023), with details in Appendix C.6.

Results: Across all statistics, both AMC and QMC have much shorter CIs than MC, with relative efficiencies ranging from 3.1 to 136. This strongly suggests our estimators can dramatically reduce cost for uncertainty estimates. AMC and QMC estimators yield comparable results. The performance of QMC depends on how the total budget is allocated between the size of the QMC point set and the number of random replicates, a trade-off we explore in further details in Appendix C.6.

Table 2: CI lengths and efficiency $(CI_{MC}/CI)^2$ (in parenthese	s), using MC as baseling	e.
--	--------------------------	----

Dataset		Brightness	Pixel Mean	Contrast	Centroid
CIFAR10	MC	2.00	2.04	1.08	0.11
	QMC	0.35 (32.05)	0.39 (26.94)	0.22 (23.43)	0.04 (7.35)
	AMC	0.35 (32.66)	0.39 (27.12)	0.23 (22.05)	0.04 (6.96)
CelebA	MC	1.77	1.76	0.60	0.60
	AMC	0.26 (47.41)	0.31 (33.15)	0.19 (10.18)	0.20 (7.82)
Church	MC	1.66	1.64	1.02	0.82
	AMC	0.14 (134.13)	0.16 (103.80)	0.20 (27.16)	0.26 (9.84)
Stable Diff.	MC	3.86	3.95	2.18	4.03
	QMC	1.36 (8.07)	1.43 (7.60)	1.03 (4.46)	1.88 (4.54)
	AMC	0.90 (18.53)	1.00 (15.53)	0.89 (6.25)	1.62 (6.15)
DiT	MC	7.52	7.74	3.32	3.15
	QMC	3.44 (4.78)	3.49 (4.91)	1.82 (3.31)	1.69 (3.46)
	AMC	3.13 (5.79)	3.12 (6.16)	1.72 (3.74)	1.74 (3.30)

5.2 EVALUATING DIFFERENT DIFFUSION INVERSE PROBLEM SOLVERS

Bayesian inverse problems are important in medical imaging, remote sensing, and astronomy. Recently, many methods have been proposed that leverage diffusion priors to solve inverse problems. For instance, Zheng et al. (2025) surveys 14 such methods, yet this is still far from exhaustive. As a result, comprehensive benchmarking is costly. Meanwhile, most existing works on new diffusion inverse solvers (DIS) largely ignore uncertainty quantification, though it is critical for fair evaluation. This raises a key challenge: how to efficiently evaluate posterior sampling algorithms with reliable estimates of reconstruction quality and uncertainty. Our approach aims to efficiently address this gap.

We show that our antithetic estimator can save substantial computational cost for quantifying uncertainty by requiring far fewer samples than standard estimators. We focus on two popular DISs, Diffusion Posterior Sampling (DPS) (Chung et al., 2023a) and Decomposed Diffusion Sampler (DDS) (Chung et al., 2023b), and evaluate them on a range of tasks described below.

5.2.1 Human Face Reconstruction

We evaluate DPS across three inverse problems: inpainting, Gaussian deblurring, and super-resolution. Reconstruction error is measured using both PSNR and L_1 distance relative to the ground truth image. For all tasks, we use 200 human face images from the CelebA-HQ dataset. For each corrupted image, we generate 50 DPS reconstructions using both PN and RR noise pairs to compare their estimators. Operator-specific configurations (e.g., kernel sizes, mask ratios) are provided in Appendix C.6.

Results: As shown in Table 3, across all three tasks, AMC achieves substantially shorter confidence intervals than standard MC. This implies AMC yields efficiency gains, ranging from 54% - 84% in inpainting, 41% - 54% in super-resolution, and 34% - 56% in deblurring.

The reconstruction metrics $(L_1, PSNR)$ differ by less than 0.2% between the two estimators, showing that they produce consistent results. Meanwhile, our method offers a clear advantage in efficiency.

Table 3: Comparison of AMC vs. MC across tasks in DPS with efficiency (CI_{MC}/CI)² in parentheses

	L1		PSNR		
Task	AMC CI	MC CI	AMC CI	MC CI	
Inpainting Super-resolution Gaussian Deblur	0.83 (1.54) 1.01 (1.41) 0.89 (1.34)	1.03 1.20 1.03	0.23 (1.84) 0.24 (1.54) 0.23 (1.56)	0.31 0.30 0.29	

5.2.2 MEDICAL IMAGE RECONSTRUCTION

We use the DDS (Chung et al., 2023b) as the reconstruction backbone and apply our antithetic noise initialization to estimate reconstruction confidence intervals. Given a single measurement, we generate 100 reconstruction pairs using 50-step DDS sampling. Reconstruction quality is evaluated against ground truth using L_1 distance and PSNR. Dataset details can be found in Appendix C.6.4.

Results: Both MC and AMC estimators give consistent estimates, with estimated L_1 error of 0.0194 and 0.0195, and PSNR values of 31.48 and 31.45, respectively. Similarly, we observe that AMC produces much shorter confidence intervals. The CI lengths decrease with relative efficiency of 1.44 for L_1 and 1.37 for PSNR. This shows that antithetic initialization consistently reduces estimator variance with no extra cost without degrading reconstruction fidelity in large-scale inverse problems.

Conclusion: These two experiments show that a single antithetic estimator reduces evaluation costs for diffusion inverse solvers by 34%–84%, a saving that becomes especially valuable given the large number of solvers in the literature. This makes large-scale benchmarking far more practical.

5.3 DIVERSITY IMPROVEMENT

Many diffusion tools generate a few images from one prompt to let users select a preferred one. However, these images may look similar with randomly sampled initial noise (Marwood et al., 2023; Sadat et al., 2024). We show that antithetic noise produces more diverse images than independent noise, without reducing image quality or increasing computational cost. Diversity is measured using pairwise SSIM (lower = more diverse) and LPIPS (Zhang et al., 2018) (higher = more diverse).

We evaluate the diversity metrics for both unconditional and conditional diffusion on image pairs generated following the same setup in Section 5.1. As shown in Table 4, antithetic noise pairs consistently lead to higher diversity, as indicated by lower SSIM and higher LPIPS scores. Importantly, image quality remains stable for both noise types. We expect our method can be easily integrated into existing diversity optimization techniques to further improve their performance.

Table 4: Average percentage improvement of PN pairs over RR pairs on SSIM and LPIPS.

Metric	Un	conditional Diffu	Condition	nal Diffusion	
	CIFAR-10	LSUN-Church	CelebA-HQ	SD1.5	DiT
SSIM (%) LPIPS (%)	88.78 6.69	45.69 3.54	36.78 15.14	28.32 5.78	23.99 10.62

6 DISCUSSION AND FUTURE DIRECTIONS

We find that antithetic initial noise yields negatively correlated samples. This is a robust property that generalizes across every model that we have tested. The negative correlation is especially useful for uncertainty quantification, where it brings huge variance reduction and cost savings. This makes our approach highly effective for evaluation and benchmarking tasks. Meanwhile, this finding can serve as a simple plug-in to improve related tasks, such as diversity improvement and image editing. A limitation of our work is that while the symmetry conjecture is supported by both empirical and theoretical evidence, it remains open in full generality. Looking forward, future directions include systematically studying the symmetry conjecture; integrating antithetic noise with existing noise-optimization methods; and further leveraging QMC to improve sampling and estimation.

ETHICS STATEMENT

This work does not involve human subjects or personally identifiable data collected by the authors. All datasets used are publicly available and distributed under licenses permitting academic research.

Our medical image reconstruction experiments rely on benchmark datasets that are fully de-identified before release (e.g., publicly available MRI data). No new data collection was performed, and no protected health information is included. The experiments are purely methodological, intended to evaluate reconstruction fidelity and uncertainty quantification techniques, and make no clinical claims or diagnostic recommendations.

Our contributions are methodological, focusing on sampling and uncertainty quantification in generative models. While such methods could be applied broadly, including in sensitive domains, we emphasize that our work aims to improve reliability, transparency, and efficiency in evaluation.

REPRODUCIBILITY STATEMENT

Full experimental Section and Appendix details can be found in The code available https://anonymous.4open.science/r/ is at Antithetic-Noise-in-Diffusion-Models-8B54.

REFERENCES

- Rishabh Agarwal, Levi Melnick, Nicholas Frosst, Xuezhou Zhang, Ben Lengerich, Rich Caruana, and Geoffrey E Hinton. Neural additive models: Interpretable machine learning with neural nets. *Advances in Neural Information Processing Systems*, 34:4699–4711, 2021.
- Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 348. Springer Science & Business Media, 2013.
- Yuanhao Ban, Ruochen Wang, Tianyi Zhou, Boqing Gong, Cho-Jui Hsieh, and Minhao Cheng. The crystal ball hypothesis in diffusion models: Anticipating object positions from initial noise. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=GpdO9r73xT.
- Alexander Buchholz, Florian Wenzel, and Stephan Mandt. Quasi-Monte Carlo variational inference. In *International Conference on Machine Learning*, pp. 668–677. PMLR, 2018.
- Changgu Chen, Libing Yang, Xiaoyan Yang, Lianggangxu Chen, Gaoqi He, Changbo Wang, and Yang Li. FIND: Fine-tuning initial noise distribution with policy optimization for diffusion models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 6735–6744, 2024a.
- Siyi Chen, Huijie Zhang, Minzhe Guo, Yifu Lu, Peng Wang, and Qing Qu. Exploring low-dimensional subspace in diffusion models for controllable image editing. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024b. URL https://openreview.net/forum?id=50aOEfb2km.
- Hamadi Chihaoui, Abdelhak Lemkhenter, and Paolo Favaro. Blind image restoration via fast diffusion inversion. *Advances in Neural Information Processing Systems*, 37:34513–34532, 2024.
- Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2023a. URL https://openreview.net/forum?id=OnD9zGAGT0k.
- Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems. *arXiv preprint arXiv:2303.05754*, 2023b.
- Radu V Craiu and Xiao-Li Meng. Multiprocess parallel antithetic coupling for backward and forward Markov chain Monte Carlo. *Annals of Statistics*, pp. 661–697, 2005.

Martin Nicolas Everaert, Athanasios Fitsios, Marco Bocchio, Sami Arpa, Sabine Süsstrunk, and Radhakrishna Achanta. Exploiting the signal-leak bias in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 4025–4034, 2024.

- Luca Eyring, Shyamgopal Karthik, Karsten Roth, Alexey Dosovitskiy, and Zeynep Akata. Reno: Enhancing one-step text-to-image models through reward-based noise optimization. *Advances in Neural Information Processing Systems*, 37:125487–125519, 2024.
- Cees M Fortuin, Pieter W Kasteleyn, and Jean Ginibre. Correlation inequalities on some partially ordered sets. *Communications in Mathematical Physics*, 22:89–103, 1971.
- Xiefan Guo, Jinlin Liu, Miaomiao Cui, Jiankai Li, Hongyu Yang, and Di Huang. Initno: Boosting text-to-image diffusion models via initial noise optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9380–9389, 2024.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Corwin Joy, Phelim P Boyle, and Ken Seng Tan. Quasi-Monte Carlo methods in numerical finance. *Management science*, 42(6):926–938, 1996.
- Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022.
- Alexander Keller. A quasi-Monte Carlo algorithm for the global illumination problem in the radiosity setting. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing: Proceedings of a conference at the University of Nevada, Las Vegas, Nevada, USA, June 23–25, 1994*, pp. 239–251. Springer, 1995.
- Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=G5RwHpBUv0.
- Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1):e190007, 2020.
- Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=a-xFK8Ymz5J.
- Siqi Kou, Lei Gan, Dequan Wang, Chongxuan Li, and Zhijie Deng. Bayesdiff: Estimating pixel-wise uncertainty in diffusion via bayesian inference. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=YcM6ofShwY.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. Flowedit: Inversion-free text-based editing using pre-trained flow models. *arXiv preprint arXiv:2412.08629*, 2024.
- Georges Le Bellier and Nicolas Audebert. Detecting out-of-distribution earth observation images with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 481–491, 2024.
- Lizao Li, Robert Carver, Ignacio Lopez-Gomez, Fei Sha, and John Anderson. Generative emulation of weather forecast ensembles with diffusion models. *Science Advances*, 10(13):eadk4489, 2024a.

Xiang Li, Yixiang Dai, and Qing Qu. Understanding generalizability of diffusion models requires rethinking the hidden gaussian structure. *Advances in neural information processing systems*, 37: 57499–57538, 2024b.

- Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 5404–5411, 2024.
 - Sifan Liu and Art B Owen. Quasi-Monte Carlo quasi-Newton in variational Bayes. *Journal of Machine Learning Research*, 22(243):1–23, 2021.
- Sifan Liu and Art B Owen. Preintegration via active subspace. *SIAM Journal on Numerical Analysis*, 61(2):495–514, 2023.
- Sifan Liu, Snigdha Panigrahi, and Jake A. Soloff. Cross-validation with antithetic Gaussian randomization. *arXiv preprint arXiv:2412.14423*, 2024.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver++: Fast solver for guided sampling of diffusion probabilistic models, 2023. URL https://openreview.net/forum?id=4vGwQqviud5.
- Pierre L'Ecuyer. Quasi-Monte Carlo methods with applications in finance. *Finance and Stochastics*, 13:307–349, 2009.
- Pierre L'Ecuyer, Marvin K. Nakayama, Art B. Owen, and Bruno Tuffin. Confidence intervals for randomized quasi-Monte Carlo estimators. *In 2023 Winter Simulation Conference*, pp. 445–446, 2023.
- David Marwood, Shumeet Baluja, and Yair Alon. Diversity and diffusion: Observations on synthetic image distributions with stable diffusion. *arXiv preprint arXiv:2311.00056*, 2023.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=aBsCjcPu_tE.
- H. Niederreiter. *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Philadelphia, PA, 1992.
- A. B. Owen. Scrambled net variance for integrals of smooth functions. *Annals of Statistics*, 25(4): 1541–1562, 1997a.
- A. B. Owen. Monte Carlo variance of scrambled net quadrature. *SIAM Journal of Numerical Analysis*, 34(5):1884–1910, 1997b.
- Art B. Owen. *Monte Carlo Theory, Methods and Examples*. https://artowen.su.domains/mc/, 2013.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=di52zR8xgf.
- Zipeng Qi, Lichen Bai, Haoyi Xiong, and Zeke Xie. Not all noises are created equally: Diffusion noise selection and optimization. *arXiv preprint arXiv:2407.14041*, 2024.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 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.

Seyedmorteza Sadat, Jakob Buhmann, Derek Bradley, Otmar Hilliges, and Romann M. Weber. CADS: Unleashing the diversity of diffusion models through condition-annealed sampling. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=zMoNrajk2X.

- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=TIdIXIpzhoI.
- Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2024.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. pmlr, 2015.
- Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse problems with latent diffusion models via hard data consistency. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=j8hdRqOUhN.
- Bowen Song, Zecheng Zhang, Zhaoxu Luo, Jason Hu, Wei Yuan, Jing Jia, Zhengxu Tang, Guanyang Wang, and Liyue Shen. Ccs: Controllable and constrained sampling with diffusion models via initial noise perturbation. *arXiv preprint arXiv:2502.04670*, 2025.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a. URL https://openreview.net/forum?id=St1giarCHLP.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL https://openreview.net/forum?id=PxTIG12RRHS.
- Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=vaRCHVj0uGI.
- Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, Dhruv Nair, Sayak Paul, William Berman, Yiyi Xu, Steven Liu, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. https://github.com/huggingface/diffusers, 2022.
- Carsten Waechter and Alexander Keller. Quasi-Monte Carlo light transport simulation by efficient ray tracing, May 31 2011. US Patent 7,952,583.
- Binxu Wang and John Vastola. The unreasonable effectiveness of gaussian score approximation for diffusion models and its applications. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=I0uknSHM2j.
- Hengkang Wang, Xu Zhang, Taihui Li, Yuxiang Wan, Tiancong Chen, and Ju Sun. Dmplug: A plug-in method for solving inverse problems with diffusion models. *Advances in Neural Information Processing Systems*, 37:117881–117916, 2024.

- Xiaoqun Wang and Kai-Tai Fang. The effective dimension and quasi-Monte Carlo integration. *Journal of Complexity*, 19:101–124, 2003.
- Weihao Xia, Yujiu Yang, Jing-Hao Xue, and Baoyuan Wu. Tedigan: Text-guided diverse face image generation and manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2021.
- Ye Xiao and Xiaoqun Wang. Enhancing quasi-Monte Carlo simulation by minimizing effective dimension for derivative pricing. *Computational Economics*, 54(1):343–366, 2019.
- Seungil You, David Ding, Kevin Canini, Jan Pfeifer, and Maya Gupta. Deep lattice networks and partial monotonic functions. *Advances in neural information processing systems*, 30, 2017.
- Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv* preprint arXiv:1506.03365, 2015.
- Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv* preprint arXiv:1811.08839, 2018.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 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.
- Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy Feng, Caifeng Zou, Yu Sun, Nikola Borislavov Kovachki, Zachary E Ross, Katherine Bouman, and Yisong Yue. Inversebench: Benchmarking plug-and-play diffusion priors for inverse problems in physical sciences. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=U3PBITXNG6.
- Zikai Zhou, Shitong Shao, Lichen Bai, Zhiqiang Xu, Bo Han, and Zeke Xie. Golden noise for diffusion models: A learning framework. *arXiv preprint arXiv:2411.09502*, 2024.

A USE OF LLMS

LLMs are used only for writing, editing, or formatting purposes and do not affect the core methodology, contributions, and originality of the paper.

B Proofs

B.1 Proofs in Section 3.2

Proof of Lemma 1. It suffices to show Var(f(Z) + f(-Z)) = 0. By definition:

$$\operatorname{Corr}(f(-Z),f(Z)) = \frac{\operatorname{Cov}(f(-Z),f(Z))}{\sqrt{\operatorname{Var}(f(Z))\operatorname{Var}(f(-Z))}} = -1.$$

Therefore

$$\begin{split} \operatorname{Var}(f(Z) + f(-Z)) &= \operatorname{Var}(f(Z)) + \operatorname{Var}(f(-Z)) + 2\operatorname{Cov}(f(Z), f(-Z)) \\ &= \operatorname{Var}(f(Z)) + \operatorname{Var}(f(-Z)) - 2\sqrt{\operatorname{Var}(f(Z))\operatorname{Var}(f(-Z))} \\ &= \left(\sqrt{\operatorname{Var}(f(Z))} - \sqrt{\operatorname{Var}(f(-Z))}\right)^2. \end{split}$$

On the other hand, since $-Z \sim \mathbb{N}(0,I)$ given $Z \sim \mathbb{N}(0,1)$, we have $\mathrm{Var}(f(Z)) = \mathrm{Var}(f(-Z))$ and then $\mathrm{Var}(f(Z) + f(-Z)) = 0$, as claimed.

Background: The Ornstein–Uhlenbeck (OU) process. The OU process $(X_t)_{t\geq 0}$ in \mathbb{R}^d solves the SDE

$$dX_t = -X_t dt + \sqrt{2} dB_t, \qquad X_0 \sim \mu_0,$$

where (B_t) is standard Brownian motion in \mathbb{R}^d . This is a Gaussian Markov process with stationary distribution $\gamma = \mathcal{N}(0, I_d)$ and generator

$$Lf(x) = \Delta f(x) - x \cdot \nabla f(x).$$

For any initial law μ_0 , let μ_t denote the law of X_t . Then $\mu_t = \mu_0 P_t$, where $(P_t)_{t \ge 0}$ is the *OU* semigroup defined by

$$(P_t f)(x) = \mathbb{E}[f(X_t) \mid X_0 = x].$$

Write the *score* of μ_t as $s_t(x) := \nabla \log p_t(x)$, where p_t is Lebesgue density of μ_t , we have

Theorem 1 (Score converges to the Gaussian score). Under the setup above,

$$\lim_{t \to \infty} \mathbb{E}_{\mu_t} [\|s_t(X_t) + X_t\|^2] = 0.$$

Moreover, the convergence is quantified by

$$\mathbb{E}_{\mu_t} [\|s_t(X_t) + X_t\|^2] = \mathcal{I}(\mu_t \mid \gamma) \le e^{-2t} \mathcal{I}(\mu_0 \mid \gamma).$$

Proof. Recall the definition of relative Fisher information

$$\mathcal{I}(\mu_t \mid \gamma) := \mathbb{E}_{\mu_t} [\|\nabla \log(p_t/\gamma)\|^2].$$

Since $\nabla \gamma(x) = -x$, we have:

$$\mathcal{I}(\mu_t \mid \gamma) = \mathbb{E}_{\mu_t} [\|s_t(X_t) + X_t\|^2].$$

Moreover, it is well known (e.g. Chapter 5 of Bakry et al. (2013)) that $\mathcal{I}(\mu_t \mid \gamma) \leq \exp(-2t)\mathcal{I}(\mu_0 \mid \gamma)$, this concludes our proof.

In variance-preserving diffusion models, the forward noising process is exactly the Ornstein-Uhlenbeck semigroup. Theorem 1 implies that the true score converges to the Gaussian score -x. Consequently, one-dimensional slices of the score become asymptotically affine antisymmetric, confirming our symmetry conjecture in the high-noise limit.

The next corollary shows, the 1-step DDIM update has nearly -1 correlation at large t.

Corollary 1 (Correlation of 1-step DDIM). With all the setup the same as above, assuming we have a score network $\epsilon^{(t)}$ approximating the score function. Let $\eta_t := \max\{\mathbb{E}_{p_t} \| \epsilon^{(t)}(X) - s_t(X) \|^2, \mathbb{E}_{p_t} \| \epsilon^{(t)}(-X) - s_t(-X) \|^2\}$ be the expected (symmetrized) error. Consider one-step DDIM update $F_t(x) = a_t x + b_t \epsilon_a^{(t)}(x) := (F_{t,1}(x), \dots, F_{t,d}(x)) \in \mathbb{R}^d$ as defined in (2). Let

$$v_{t,i} := \min\{\operatorname{Var}_{p_t} \big(F_{t,i}(X)\big), \operatorname{Var}_{p_t} \big(F_{t,i}(-X)\big)\}.$$

Then

$$\left|\operatorname{Corr}(F_{t,i}(X), F_{t,i}(-X)) + 1\right| \leq \frac{2|b_t|}{\sqrt{v_{t,i}}} \left(\sqrt{\eta_t} + e^{-t}\sqrt{\mathcal{I}(\mu_0 \mid \gamma)}\right).$$

Proof. Theorem 1 shows $\mathcal{I}(\mu_t \mid \gamma) \leq \exp(-2t)\mathcal{I}(\mu_0 \mid \gamma)$. Write $\epsilon_{\theta}^{(t)} = s_t + r_t$ and $s_t = -x + \Delta_t$, so that $\mathbb{E}_{p_t} \|r_t(X)\|^2 \leq \eta_t$ and $\mathbb{E}_{p_t} \|\Delta_t(X)\|^2 = \mathcal{I}(\mu_t \mid \gamma)$. Then

$$F_t(x) = (a_t - b_t)x + b_t(\Delta_t(x) + r_t(x)), \qquad F_t(-x) = -F_t(x) + E_t(x),$$

with $E_t(x) := b_t[\Delta_t(-x) + \Delta_t(x) + r_t(-x) + r_t(x)]$. By triangle inequality of $L^2(p_t)$ norm,

$$||E_t||_{L^2(p_t)} \le |b_t| (e^{-t} \sqrt{\mathcal{I}(\mu_0 \mid \gamma)} + e^{-t} \sqrt{\mathcal{I}(\mu_0 \mid \gamma)} + 2\sqrt{\eta_t}).$$

Using

$$\operatorname{Corr}(F_{t,i}(X), F_{t,i}(-X)) = -1 + \frac{\operatorname{Cov}(F_{t,i}(X), E_{t,i}(X))}{\sqrt{\operatorname{Var}(F_{t,i}(X)) \operatorname{Var}(F_{t,i}(-X))}}$$

and Cauchy-Schwarz with the variance lower bounds gives

$$\left| \operatorname{Corr} \left(F_t(X), F_t(-X) \right) + 1 \right| \le \frac{\| E_t \|_{L^2(p_t)}}{\sqrt{v_{t,i}}} \le \frac{2 |b_t|}{\sqrt{v_{t,i}}} (\sqrt{\eta_t} + e^{-t} \sqrt{\mathcal{I}(\mu_0 \mid \gamma)},$$

as claimed. \Box

In practice, if $v_{t,i} > 0$, the correlation is close to -1 for large t; its deviation from -1 is on the order of the neural network approximation error plus an exponentially decaying term Ce^{-t} .

B.3 Proofs in Section 4

B.3.1 MONTE CARLO ESTIMATOR:

Proposition 1. Denote $\mathbb{E}_{z \sim \mathbb{N}(0,I)}[S(DM(z))]$ by μ and $\operatorname{Var}_{z \sim \mathbb{N}(0,I)}[S(DM(z))]$ by σ^2 . Assuming $\sigma^2 < \infty$, then we have:

- $\hat{\mu}_N^{MC} \to \mu$ and $(\hat{\sigma}_N^{MC})^2 \to \sigma^2$ almost surely, and $\mathbb{P}\left(\mu \in CI_N^{MC}(1-\alpha)\right) \to 1-\alpha$, both as $N \to \infty$.
- $\mathbb{E}[(\hat{\mu}_N^{MC} \mu)^2] = \operatorname{Var}[\hat{\mu}_N^{MC}] = \sigma^2/N.$

In words, the above proposition shows: (i) *Correctness*: The standard Monte Carlo estimator $\hat{\mu}_N^{\text{MC}}$ converges to the true value, and the coverage probability of $\text{CI}_N^{\text{MC}}(1-\alpha)$ converges to the nominal level $(1-\alpha)$. (ii) *Reliability*: The expected squared error of $\hat{\mu}_N^{\text{MC}}$ equals the variance of a single sample divided by sample size. The confidence interval has width approximately $2\sigma z_{1-\alpha/2}/\sqrt{N}$.

Proof of Proposition 1. Let $S_i := S(\mathrm{DM}(z_i))$ for i = 1, ..., N. Because the noises z_i are drawn independently from $\mathbb{N}(0, I)$, the random variables $S_1, S_2, ...$ are independent and identically distributed with mean μ and variance $\sigma^2 < \infty$.

Consistency. By the strong law of large numbers,

$$\hat{\mu}_N^{\text{MC}} = \frac{1}{N} \sum_{i=1}^N S_i \xrightarrow{\text{a.s.}} \mu.$$

The sample variance estimator

$$(\hat{\sigma}_N^{\text{MC}})^2 = \frac{1}{N-1} \sum_{i=1}^N (S_i - \hat{\mu}_N^{\text{MC}})^2 = \frac{1}{N-1} \sum_{i=1}^N S_i^2 - \frac{N}{N-1} (\hat{\sigma}_N^{\text{MC}})^2.$$

The first term converges to $\mathbb{E}[S_1^2]$ almost surely by the law of large numbers, the second term converges to μ^2 almost surely as shown above. Therefore $(\hat{\sigma}_N^{\text{MC}})^2 \xrightarrow{\text{a.s.}} \sigma^2$.

Asymptotic normality and coverage. The classical *central limit theorem* states that

$$\sqrt{N} \frac{\hat{\mu}_N^{\text{MC}} - \mu}{\sigma} \xrightarrow{d} \mathbb{N}(0, 1).$$

Replacing the unknown σ by the consistent estimator $\hat{\sigma}_N^{\text{MC}}$ and applying *Slutsky's theorem* yields

$$\sqrt{N} \, \frac{\hat{\mu}_N^{\text{MC}} - \mu}{\hat{\sigma}_N^{\text{MC}}} \, \xrightarrow{d} \, \mathbb{N}(0, 1).$$

Hence

$$\mathbb{P}\big(\mu \in \mathrm{CI}_N^{\mathrm{MC}}(1-\alpha)\big) = \mathbb{P}\!\Big(\big|\sqrt{N}(\hat{\mu}_N^{\mathrm{MC}}-\mu)/\hat{\sigma}_N^{\mathrm{MC}}\big| \leq z_{1-\alpha/2}\Big) \ \longrightarrow \ 1-\alpha.$$

Mean-squared error. Because $\hat{\mu}_N^{\text{MC}}$ is the average of N i.i.d. variables,

$$\operatorname{Var}[\hat{\mu}_N^{\mathrm{MC}}] = \frac{\sigma^2}{N}.$$

Moreover, since $\hat{\mu}_N^{\text{MC}}$ is unbiased $(\mathbb{E}[\hat{\mu}_N^{\text{MC}}] = \mu)$, its mean-squared error

$$\mathbb{E}\big[(\hat{\mu}_N^{\text{MC}} - \mu)^2\big] = \mathbb{E}\big[(\hat{\mu}_N^{\text{MC}} - \mathbb{E}[\hat{\mu}_N^{\text{MC}}])^2\big] = \frac{\sigma^2}{N}.$$

This completes the proof.

B.3.2 ANTITHETIC MONTE CARLO ESTIMATOR:

Proposition 2. Denote $\mathbb{E}[S(DM(z))]$ by μ , Var[S(DM(z))] by σ^2 , and $Cov(S_1^+, S_1^-)$ by ρ . Assuming $\sigma^2 < \infty$, then we have:

- $\hat{\mu}_N^{AMC} \to \mu$ and $(\hat{\sigma}_N^{AMC})^2 \to (1+\rho)\sigma^2/2$ almost surely as $N \to \infty$,
- $\mathbb{E}[(\mu_N^{AMC} \mu)^2] = \text{Var}[\mu_N^{AMC}] = \sigma^2(1+\rho)/N.$
- $\mathbb{P}\left(\mu \in CI_N^{AMC}(1-\alpha)\right) \to 1-\alpha \text{ as } N \to \infty.$

Proof. Let N=2K and generate independent antithetic noise pairs $(z_i,-z_i)$ for $i=1,\ldots,K$. Recall

$$S_i^+ = S(DM(z_i)), \quad S_i^- = S(DM(-z_i)), \quad \bar{S}_i = \frac{1}{2}(S_i^+ + S_i^-), \quad i = 1, \dots, K.$$

Because the pairs are independent and identically distributed (i.i.d.), the random variables $\bar{S}_1, \dots, \bar{S}_K$ are i.i.d. with

$$\mathbb{E}[\bar{S}_1] = \mu$$

and

$$\operatorname{Var}[\bar{S}_1] = \operatorname{Var}\left[\frac{1}{2}(S_1^+ + S_1^-)\right]$$
 (3)

$$= \left(\frac{1}{2}\right)^2 \operatorname{Var}\left[S_i^+ + S_i^-\right] \tag{4}$$

$$= \frac{1}{4} \left(\text{Var}[S_i^+] + \text{Var}[S_i^-] + 2 \text{ Cov}(S_i^+, S_i^-) \right)$$
 (5)

$$=\frac{1}{4}\left(\sigma^2+\sigma^2+2\rho\sigma^2\right) \tag{6}$$

$$=\frac{1+\rho}{2}\,\sigma^2.\tag{7}$$

where we have written $\mathrm{Cov}(S_i^+, S_i^-) = \rho \, \sigma^2$ as in the statement. Set $\hat{\mu}_N^{\mathrm{AMC}} = K^{-1} \sum_{i=1}^K \bar{S}_i$ and let $(\hat{\sigma}_N^{\mathrm{AMC}})^2$ be the sample variance of $\bar{S}_1, \ldots, \bar{S}_K$. All three claims in Proposition 2 follow from Proposition 1.

B.4 AN ALTERNATIVE EXPLANATION VIA THE FKG INEQUALITY

We also provide an expository discussion highlighting the FKG connection behind negative correlation. Consider the univariate case where a scalar Gaussian z is fed through a sequence of one-dimensional linear maps and ReLU activations. Let F be the resulting composite function. One can show that, regardless of the number of layers or the linear coefficients used, $\mathrm{Corr} \big(F(z), F(-z) \big) \leq 0$. The proof relies on the univariate FKG inequality (Fortuin et al., 1971), a well-known result in statistical physics and probability theory.

We generalize this result to higher dimensions via *partial monotonicity*, under which negative correlation still holds.

We first formally state and prove the claim for the univariate case in Section B.4.

Proposition 3 (Univariate case). Let $z \sim \mathbb{N}(0,1)$ and let $F : \mathbb{R} \to \mathbb{R}$ be the output of any one-dimensional feed-forward network obtained by alternating scalar linear maps and ReLU activations:

$$h_0(z) = z,$$
 $h_{\ell}(z) = \text{ReLU}(w_{\ell} h_{\ell-1}(z) + b_{\ell}), \ \ell = 1, \dots, L,$ $F(z) = h_L(z).$

For all choices of depths L and coefficients $\{w_{\ell}, b_{\ell}\}_{\ell=1}^{L}$,

$$Corr(F(z), F(-z)) \leq 0.$$

Proof. An important observation is that monotonicity is preserved under composition: combining one monotonic function with another produces a function that remains monotonic.

Each scalar linear map $x \mapsto w_\ell x + b_\ell$ is monotone: it is non-decreasing if $w_\ell \ge 0$ and non-increasing if $w_\ell < 0$. The ReLU map $x \mapsto \text{ReLU}(x) = \max\{0, x\}$ is non-decreasing. Hence the final function F is monotone. Without loss of generality, we assume F is non-decreasing.

FKG inequality guarantees $\operatorname{Cov}_{Z \sim \mathbb{N}(0,1)}(f(Z),g(Z)) \geq 0$ provided that f,g are non-decreasing. Therefore, $\operatorname{Cov}_{Z \sim \mathbb{N}(0,1)}(F(Z),F(-Z)) = -\operatorname{Cov}_{Z \sim \mathbb{N}(0,1)}(F(Z),-F(-Z)) \leq 0$. Since

$$\operatorname{Corr}(F(z), F(-z)) = \frac{\operatorname{Cov}_{Z \sim \mathbb{N}(0,1)}(F(Z), F(-Z))}{\sqrt{\operatorname{Var}(F(z))\operatorname{Var}(F(-z))}},$$

we have $Corr(F(z), F(-z)) \le 0$.

 To generalize the result to higher dimension, we define a function $f: \mathbb{R}^m \to \mathbb{R}$ to be partially monotone if for each coordinate j, holding all other inputs fixed, the map $t \mapsto f(x_1,\ldots,x_{j-1},t,x_{j+1},\ldots,x_m)$ is either non-decreasing or non-increasing. Mixed monotonicity is allowed. For example, f(x,y) = x - y is non-decreasing in x and non-increasing in y, yet still qualifies as partially monotone. We have the following result:

Proposition 4. For a diffusion model $DM : \mathbb{R}^d \to \mathbb{R}^m$ and a summary statistics $S : \mathbb{R}^m \to \mathbb{R}$, if the joint map is partially monotone, then $Corr(S \circ DM(Z)), S \circ DM(-Z)) \leq 0$.

Now we prove the general case:

Proof of Proposition 4. Let $G := S \circ DM$. For each coordinate $j \in [d]$ fix a sign

$$s_j = \begin{cases} +1, & \text{if } G \text{ is non-decreasing in } x_j, \\ -1, & \text{if } G \text{ is non-increasing in } x_j. \end{cases}$$

Write $s=(s_1,\ldots,s_d)\in\{\pm 1\}^d$ and define, for any $z\in\mathbb{R}^d$,

$$\widetilde{G}(z) = \widetilde{G}(z_1, \dots, z_d) := G(s_1 z_1, \dots, s_d z_d).$$

Similarly, define

$$\widetilde{H}(z) := \widetilde{G}(-z) = G(-s_1 z_1, \dots, -s_d z_d).$$

By construction, \widetilde{G} is coordinate-wise non-decreasing and \widetilde{H} is coordinate-wise non-increasing.

Because each coordinate of $Z \sim \mathbb{N}(0, I_d)$ is symmetric, the random vectors $(s_1 Z_1, \dots, s_d Z_d)$ and (Z_1, \dots, Z_d) have the same distribution. Hence

$$\operatorname{Cov}(G(Z), G(-Z)) = \operatorname{Cov}(\widetilde{G}(Z), \widetilde{H}(Z)).$$

The multivariate FKG inequality (Fortuin et al., 1971) for product Gaussian measures states that, when U and V are coordinate-wise non–decreasing, $\mathrm{Cov}\big(U(Z),V(Z)\big)\geq 0$. Apply it to the pair $\big(\widetilde{G},-\widetilde{H}\big)$: both components are non–decreasing, so

$$\operatorname{Cov} \bigl(\widetilde{G}(Z), -\widetilde{H}(Z) \bigr) \ \geq \ 0 \ \implies \ \operatorname{Cov} \bigl(\widetilde{G}(Z), \widetilde{H}(Z) \bigr) \ \leq \ 0 \ \Leftrightarrow \ \operatorname{Cov} \bigl(G(Z), G(-Z) \bigr) \ \leq \ 0.$$

Therefore
$$Corr(G(Z), G(-Z)) \leq 0$$
, as claimed.

Proposition 4 relies on checking partial-monotonicity. If S is partially monotone (including any linear statistic), then the conditions of Proposition 4 are satisfied by, e.g., Neural Additive Models (Agarwal et al., 2021) and Deep Lattice Networks (You et al., 2017). Unfortunately, partial-monotonicity is in general hard to verify.

While popular diffusion architectures like DiT and U-Net lack partial monotonicity, we include Proposition 4 as an expository attempt to highlight the FKG connection behind negative correlation.

C ADDITIONAL EXPERIMENTS

C.1 GENERATIVE MODELS BEYOND DIFFUSION

A mentioned in Section 3, we evaluate correlation on a VAE and flow-based models. Unlike the diffusion models, which we use public pre-trained checkpoints, both of these models required explicit training before evaluation. Here we describe the training setup and generation procedure.

VAE: We train the unconditional VAE on MNIST following the publicly available implementation provided in . The VAE consists of a simple convolutional encoder–decoder architecture with a Gaussian latent prior. After training, we generate 1,600 paired samples under PN and RR schemes, respectively.

Glow: We train the class-conditional Glow model (Kingma & Dhariwal, 2018) on CIFAR-10 using the normflows library. The architecture follows a multiscale normalizing flow design and Glow blocks. After training, we generated 100 PN and 100 RR pairs per class.

C.2 IMAGE EDITING

We apply our antithetic initial noise strategy to the image editing algorithm FlowEdit (Kulikov et al., 2024), which edits images toward a target text prompt using pre-trained flow models.

In Algorithm 1 of FlowEdit, at each timestep, the algorithm samples $n_{\rm avg}$ random noises $Z \sim \mathcal{N}(0,1)$ to create noisy versions of the source image, computes velocity differences between source and target conditions, and averages these directions to drive the editing process.

In the $n_{\text{avg}} = 2$ setting, we replace the two independent random samples with antithetic noises: for each $Z \sim \mathcal{N}(0, I)$ we also use -Z and average the two velocity updates.

We compare on 76 prompts provided in FlowEdit's official GitHub repository. For each prompt, we generate 10 images using both PN and RR. All other parameters follow the repository defaults. We evaluate performance using CLIP(semantic text–image alignment; higher is better) and LPIPS (perceptual distance to the source; lower is better), which jointly measure text adherence and structure preservation.

As a result, PN improves the mean CLIP score, winning in 56.59% of all pairwise comparisons. It also reduces LPIPS, winning in 81.58% of all pairwise comparisons.

C.3 IMPLEMENTATION DETAIL

We use pretrained models from Hugging Face's Diffusers library (von Platen et al., 2022): google/ddpm-church-256, google/ddpm-cifar10-32, and google/ddpm-celebahq-256 for unconditional diffusion; Stable Diffusion v1.5 for text-to-image; and the original repository from (Chung et al., 2023a) for guided generation. Experiments were run on eight NVIDIA L40 GPUs. The most intensive setup—Stable Diffusion—takes about five minutes to generate 100 images for a single prompt.

C.4 ADDITIONAL EXPERIMENTS ON PIXEL-WISE SIMILARITY

C.4.1 DDPM

Antithetic sampling setup: Unlike DDIM, where the sampling trajectory is deterministic once the initial noise is fixed, DDPM adds a random Gaussian noise at every timestep. The update function in DDPM is

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} (x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_{\theta}(x_t, t)) + \sigma_t z_t$$

and $z_t \sim \mathbb{N}(0,1)$ if t > 1, else z = 0. Therefore, in DDPM, antithetic sampling requires not only negating the initial noise but also negating *every noise* z_t added at each iteration.

Table 5 and table 6 report the standard Pearson correlations and centralized correlation between pixel values of image pairs produced using the same pre-trained models under PN and RR noise schemes using DDPM (default 1000 steps) across CIFAR10, CelebA-HQ, and LSUN-Church. For each dataset, we follow the same setup explained in Section 3. We calculate the pixelwise correction with 1600 antithetic noise pairs and 1600 independent noises.

Consistent with the behavior observed in DDIM, PN pairs in DDPM samples also exhibit negative correlation. Using the standard Pearson correlation, the mean correlation for PN pairs is strongly negative in all datasets—CIFAR10 (-0.73), Church (-0.45), and in the more identity-consistent CelebA (-0.18).

The centralized correlation analysis further sharpens this contrast: mean PN correlations are substantially lower (CIFAR10 -0.80, CelebA -0.67, Church -0.65). These results confirm again that PN noise pairs consistently introduce strong negative dependence, while RR pairs remain close to uncorrelated or weakly positive.

Table 5: DDPM standard Pearson correlation coefficients for PN and RR pairs

Statistic	CIFAR10		CelebA		Church	
	PN	RR	PN	RR	PN	RR
Mean	-0.73	0.04	-0.18	0.29	-0.45	0.11
Min	-0.95	-0.61	-0.80	-0.50	-0.90	-0.60
25th percentile	-0.81	-0.12	-0.31	0.16	-0.55	-0.02
75th percentile	-0.66	0.20	-0.05	0.44	-0.36	0.23
Max	-0.21	0.80	0.44	0.78	0.07	0.80

Table 6: DDPM centralized Pearson correlation coefficients for PN and RR pairs

Statistic	CIFAR10		CelebA		Church	
	PN	RR	PN	RR	PN	RR
Mean	-0.80	0.01	-0.67	-0.00	-0.65	-0.00
Min	-0.96	-0.61	-0.89	-0.60	-0.93	-0.56
25th percentile	-0.87	-0.15	-0.73	-0.15	-0.72	-0.12
75th percentile	-0.75	0.16	-0.62	0.15	-0.60	0.11
Max	-0.23	0.76	0.14	0.64	-0.32	0.72

C.4.2 DIFFUSION POSTERIOR SAMPLING (DPS)

The same pattern persists in the setting of using the diffusion model as a prior for posterior sampling, too, which has been utilized to solve various inverse problems, such as inpainting, super-resolution, and Gaussian deblurring.

Since there is a ground truth image available in the image restoration task, the standard pixel-wise correlation is calculated using the difference between reconstructed images and the corresponding ground truth, and the centralized correlation is calculated using the same definition described in 3. Although the overall standard correlation values are shifted up, due to the deterministic conditioning,

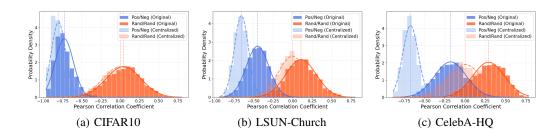


Figure 5: Pearson Correlation histograms for PN and RR pairs across three datasets using DDPM. Dashed lines indicate the mean Pearson correlation for each group.

the posterior nature of sampling—PN pairs still shows significantly lower correlations than RR pairs across all tasks.

For the standard Pearson correlation, the mean PN correlations range from 0.20 to 0.27, while RR correlations consistently lie above 0.50. For the centralized correlation, PN correlations are strongly negative across all tasks (means around -0.72). In contrast, RR pairs remain centered near zero (mean correlations around -0.01 to -0.02).

Table 7: DPS Pearson correlation coefficients for PN and RR pairs

Statistic	Inpai	nting	Gaussian Deblur		Super-resolution		
	PN	RR	PN	RR	PN	RR	
Mean	0.28	0.57	0.27	0.57	0.20	0.53	
Min	-0.14	0.05	-0.22	0.13	-0.34	0.01	
25th percentile	0.19	0.52	0.17	0.51	0.10	0.47	
75th percentile	0.36	0.63	0.36	0.63	0.30	0.60	
Max	0.66	0.83	0.62	0.83	0.64	0.81	

Table 8: DPS centralized correlation coefficients for PN and RR pairs

Statistic	Inpainting		Gaussian Blur		Super-resolution	
	PN	RR	PN	RR	PN	RR
Mean	-0.72	-0.02	-0.71	-0.01	-0.72	-0.01
Min	-0.86	-0.43	-0.84	-0.36	-0.87	-0.35
25th percentile	-0.76	-0.07	-0.75	-0.07	-0.76	-0.08
75th percentile	-0.69	0.04	-0.67	0.05	-0.69	0.05
Max	-0.14	0.47	0.01	0.48	-0.14	0.43

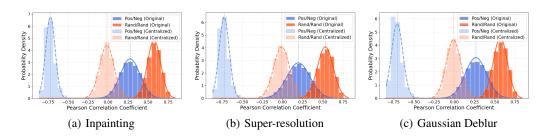


Figure 6: Pearson Correlation histograms for PN and RR pairs across three tasks in DPS. Dashed lines indicate the mean Pearson correlation for each group.

C.4.3 Wasserstein distance

To complement the correlation-based analyses in the main text, we also evaluate similarity using the Wasserstein distance, a measure of distributional discrepancy. It quantifies the minimal "effort" required to transform one probability distribution into another, which means lower Wasserstein values indicate closer alignment between the two distributions, while higher values indicate larger differences.

To calculate Wasserstein distances, we treat each generated image pair as a sample from a distribution under the sampling scheme, PN or RR. As shown in Table 9, PN consistently exhibits larger Wasserstein distances than RR across nearly all models, and the differences are statistically significant at all $p < 10^{-10}$. This implies that antithetic initial noises lead to more divergent distributions than random sampling and confirms our results from the correlation analysis.

Table 9: Wasserstein Distance, shown are means (SD) with corresponding t-statistics and p-values.

Model	Wasserstei	n Distance	t-stats (p)
	PN	RR	
LSUN-Church	0.16 (0.10)	0.12 (0.07)	t = 12.09, p = 0
CIFAR-10	0.19 (0.14)	0.15 (0.09)	t = 8.30, p = 0
CelebA-HQ	0.17 (0.11)	0.12 (0.07)	t = 14.50, p = 0
SD 1.5	0.10 (0.06)	0.09 (0.04)	t = 33.11, p = 0
DiT	0.19 (0.15)	0.14 (0.12)	t = 14.17, p = 0
VAE	0.07 (0.05)	0.05 (0.03)	t = 13.24, p = 0
Glow	0.15 (0.09)	0.12 (0.08)	t = 7.92, p = 0

C.5 ADDITIONAL EXPERIMENTS ON THE SYMMETRY CONJECTURE

We run additional experiments to validate the conjecture in Section 3.2. Using a pretrained score network and a 50-step DDIM sampler, we evaluated both CIFAR-10 and Church. For each dataset, we selected five random coordinates in the $C \times H \times W$ tensor (channel, height, width). At every chosen coordinate we examined the network output at time steps $t=1,\ldots,20$ (small t is close to the final image, large t is close to pure noise). For each t we drew a standard Gaussian sample $\mathbf{x} \sim \mathbb{N}(0,I_d)$ and computed the value of $\epsilon_{\theta}^{(t)}(c\mathbf{x})$ at the selected coordinate. The resulting plots appear in Figures 7–16.

To measure how much a one-dimensional function resembles an antisymmetric shape, we introduce the *affine antisymmetry score*

$$AS(f) := 1 - \frac{\int_{-1}^{1} (0.5f(-x) + 0.5f(x) - \bar{f})^{2}}{\int_{-1}^{1} (f(x) - \bar{f})^{2}}$$

where $\bar{f} := \int_{-1}^{1} f dx/2$ is the average value of f on [-1, 1].

The integral's numerator is the squared average distance between the antithetic mean 0.5 f(-x) + 0.5 f(x) and the overall mean \bar{f} , while the denominator is the full variance of f over the interval.

The antisymmetry score is well-defined for every non-constant function f; it takes values in the range [0,1] and represents the fraction of the original variance that is eliminated by antithetic averaging. The score attains $\mathrm{AS}(f)=1$ exactly when the antithetic sum f(x)+f(-x) is constant, that is, when f is perfectly affine-antisymmetric. Conversely, $\mathrm{AS}(f)=0$ if and only if f(-x)=f(x)+c for some constant c, meaning f is affine-symmetric.

For each dataset we have 100 (5 coordinates \times 20 time steps) scalar functions; the summary statistics of their AS scores are listed in Table 10. Both datasets have very high AS scores: the means exceed 0.99 and the 10% quantiles are above 0.97, indicating that antithetic pairing eliminates nearly all variance in most cases. Even the lowest scores (about 0.77) still remove more than three-quarters of the variance.

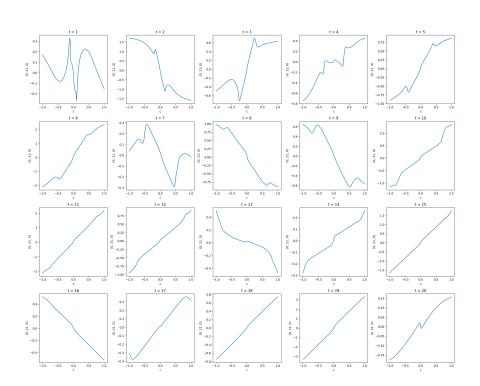


Figure 7: CIFAR10: Coordinate (0, 21, 0)

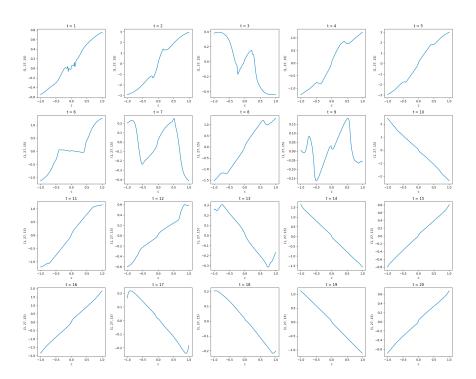


Figure 8: CIFAR10: Coordinate (1, 27, 15)

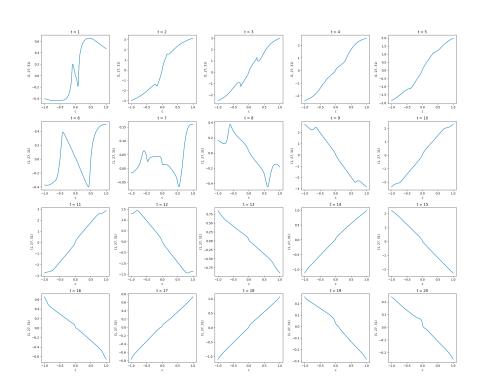


Figure 9: CIFAR10: Coordinate (1, 27, 31)

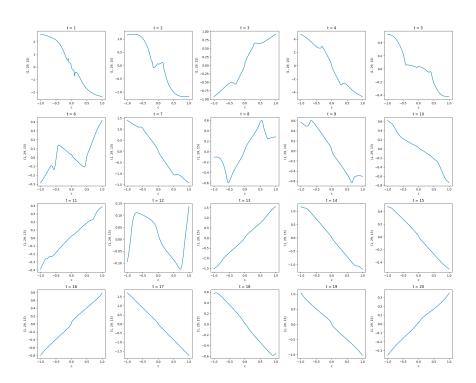


Figure 10: CIFAR10: Coordinate (1, 29, 15)

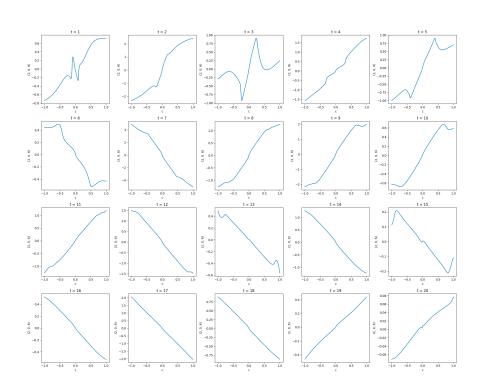


Figure 11: CIFAR10: Coordinate (2,0,6)

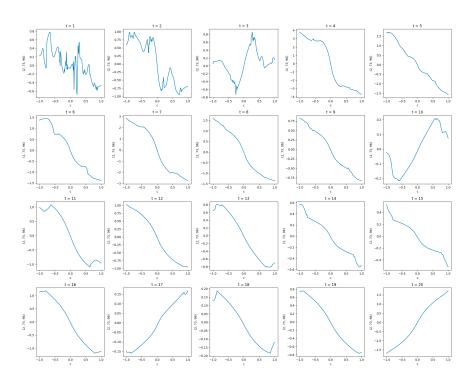


Figure 12: Church: Coordinate (2, 73, 76)

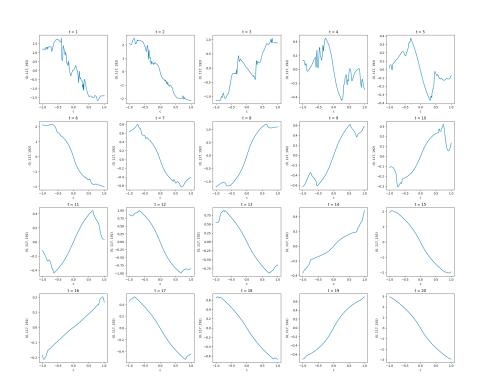


Figure 13: CIFAR10: Coordinate (0, 117, 192)

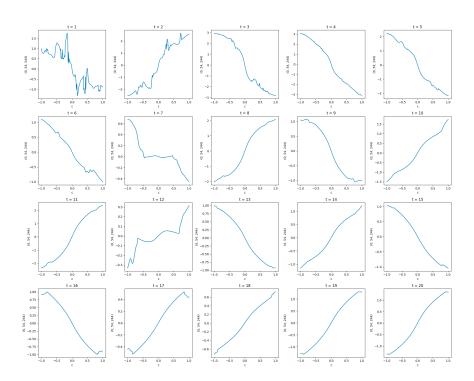


Figure 14: Church: Coordinate (0, 54, 244)

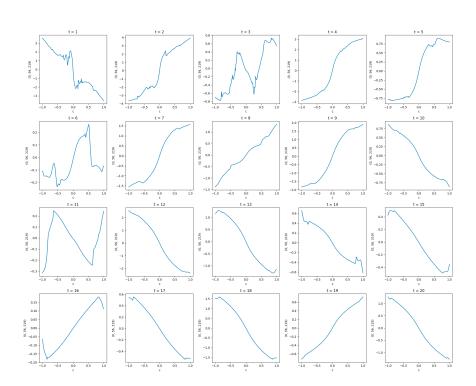


Figure 15: CIFAR10: Coordinate (0, 56, 219)

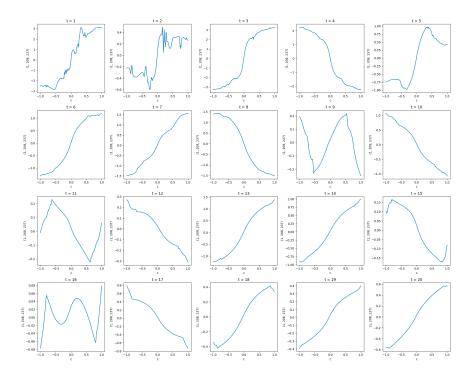


Figure 16: Church: Coordinate (1, 208, 237)

Table 1	10: <i>A</i>	Affine	antisy	ymmetry	score
---------	--------------	--------	--------	---------	-------

Dataset	CIFAR10	Church		
Mean	0.9932	0.9909		
Min	0.7733	0.7710		
1% quantile	0.9104	0.8673		
2% quantile	0.9444	0.8768		
5% quantile	0.9690	0.9624		
10% quantile	0.9865	0.9750		
Median	0.9992	0.9995		

C.6 ADDITIONAL EXPERIMENTS FOR UNCERTAINTY QUANTIFICATION

C.6.1 UNCERTAINTY QUANTIFICATION

The image metrics used in uncertainty quantification are used to capture different aspects of pixel intensity, color distribution, and perceived brightness.

Mean pixel value is defined as the average of all pixel intensities across the image (including all channels). Formally, for an image $I \in \mathbb{R}^{C \times H \times W}$, the mean is computed as

$$\mu = \frac{1}{HWC} \sum_{i=1}^{C} \sum_{j=1}^{H} \sum_{k=1}^{W} I_{i,j,k}.$$

Brightness is calculated using the standard CIE formula: $0.299 \cdot R + 0.587 \cdot G + 0.114 \cdot B$ to produce a grayscale value at *each pixel*, where R, G, B is the red, green, and blue color value of the given pixel. It is widely used in video and image compression, and approximates human visual sensitivity to color. The brightness of an image is then the average of the grayscale value across all pixels.

Contrast is computed as the difference in average pixel intensity between the top and bottom halves of an image. Let $x \in [0,1]^{C \times H \times W}$ be the normalized image, we define contrast as $100 \cdot (\mu_{\text{top}} - \mu_{\text{bottom}})$, where μ_{top} and μ_{bottom} are the average intensities over the top and bottom halves, respectively.

Centroid measures the coordinate of the brightness-weighted center of mass of the image. For scalar comparison, we focus on the vertical component of the centroid to assess spatial uncertainty. After converting to grayscale $M \in \mathbb{R}^{H \times W}$ by averaging across channels, we treat the image as a 2D intensity distribution and compute the vertical centroid as

$$\frac{1}{\sum_{i=1}^{H} \sum_{j=1}^{W} M_{i,j}} \sum_{i=1}^{H} \sum_{j=1}^{W} i \cdot M_{i,j},$$

where i denotes the row index.

C.6.2 QMC EXPERIMENTS

RQMC confidence interval construction: For the RQMC experiments, we consider R independent randomization of a Sobol' point set of size n, with a fixed budget of N=Rn function evaluations. Denote the RQMC point set in the r-th replicate as $\{\mathbf{u}_{r,k}\}_{1\leq k\leq n}\subset [0,1]^d$, where $\mathbf{u}_{r,k}\sim \mathrm{Unif}([0,1]^d)$. Applying the Gaussian inverse cumulative distribution function Φ^{-1} to each coordinate of $\mathbf{u}_{r,k}$ transforms the uniform samples to standard normal samples. Consequently, the estimate in each replicate is given by

$$\hat{\mu}_r^{\text{QMC}} = \frac{1}{n} \sum_{k=1}^n S(\text{DM}(\Phi^{-1}(\mathbf{u}_{r,k}))), \quad r = 1, \dots, R.$$

The overall point estimate is their average

$$\hat{\mu}_N^{\text{QMC}} = \frac{1}{R} \sum_{r=1}^R \hat{\mu}_r^{\text{QMC}}.$$

Let $(\hat{\sigma}_R^{\rm QMC})^2$ be the sample variance of $\hat{\mu}_1^{\rm QMC},\ldots,\hat{\mu}_R^{\rm QMC}$. The Student t confidence interval is given by

$$\mathrm{CI}_N^{\mathrm{QMC}}(1-\alpha) \; = \; \hat{\mu}_N^{\mathrm{QMC}} \; \pm \; t_{R-1,\, 1-\alpha/2} \, \frac{\hat{\sigma}_R^{\mathrm{QMC}}}{\sqrt{R}}, \label{eq:circumstant}$$

where $t_{R-1,\,1-\alpha/2}$ is the $(1-\alpha/2)$ -quantile of the t-distribution with R-1 degrees of freedom. If the estimates $\hat{\mu}_r^{\rm QMC}$ are normally distributed, this confidence interval has exact coverage probability $1-\alpha$. In general, the validity of Student t confidence interval is justified by CLT. An extensive numerical study by L'Ecuyer et al. (2023) demonstrates that Student t intervals achieve the desired coverage empirically.

Exploring different configurations of R and n This experiment aims to understand how different configurations of R, the number of replicates, and n, the size of the QMC point set, affect the CI length of the RQMC method. The total budget of function evaluations is fixed at Rn=3200, to be consistent with the AMC and MC experiments. We consider the four image metrics used in Section 4 and one additional image metric, MUSIQ.

For RQMC, each configuration was repeated five times, and the results were averaged to ensure stability. AMC and MC each consist of a single run over 3200 images. All experiments are conducted using the CIFAR10 dataset.

The results are shown in Table 11 and underlined values indicate the best CI length among the three QMC configurations, while bold values indicate the best CI length across all methods, including MC and AMC. For the brightness and pixel mean metrics, the configuration with point set size n=64 and number of replicates R=50 reduces CI length the most. In contrast and centroid, the configuration with larger n has a better CI length. For MUSIQ, a more complex metric, changes in CI lengths across configurations are subtle, and the configuration with the largest R has the shortest CI length. While no single configuration consistently advantages, all RQMC and antithetic sampling methods outperform plain MC.

Table 11: Average 95% CI length and relative efficiencies (vs MC baseline). The first three rows are for RQMC methods, with the configuration of $R \times n$ indicated in the first column.

$R \times n$	Brightness		Mean		Contrast		Centroid		MUSIQ	
	CI	Eff.	CI	Eff.	CI	Eff.	CI	Eff.	CI	Eff.
25 × 128	0.37	29.81	0.40	25.71	0.22	24.57	0.04	8.17	0.13	1.11
50×64	0.35	32.05	0.39	26.94	0.22	23.43	0.04	7.35	0.13	1.13
200×16	0.47	18.38	0.49	17.30	0.29	14.20	0.05	5.84	0.12	1.21
AMC	0.35	32.66	0.39	27.12	0.23	22.05	0.04	6.96	0.13	1.06
MC	2.00	-	2.04	-	1.08	-	0.11	-	0.13	-

C.6.3 DPS EXPERIMENT IMPLEMENTATION

We evaluate the confidence interval length reduction benefits of antithetic initial noise across three common image inverse problems: super-resolution, Gaussian deblurring, and inpainting. For each task, the forward measurement operator is applied to the true image, and noisy observations are generated using Gaussian noise with $\sigma=0.05$. We use the official implementation of DPS (Chung et al., 2023a) with the following parameters:

- **Super-resolution** uses the *super_resolution* operator with an input shape of (1, 3, 256, 256) and an upsampling scale factor of 2. This models the standard bicubic downsampling scenario followed by Gaussian noise corruption.
- Gaussian deblurring employs the *gaussian_blur* operator, again with a kernel size of 61 but with intensity set to 1, which is intended to test the variance reduction in a simpler inverse scenario.

• **Inpainting** is set up using the *inpainting* operator with a random binary mask applied to the input image. The missing pixel ratio is drawn from a uniform range between 30% and 70%, and the image size is fixed at 256 × 256. A higher guidance scale (0.5) is used to compensate for the sparsity of observed pixels.

As shown in Table 3, AMC consistently achieves shorter CI lengths than MC across all tasks and metrics without sacrificing reconstruction quality, implied by the L1 and PSNR metrics to measure the difference between reconstructed images and ground truth images.

C.6.4 DDS DATASET

Data used in the DDS experiment on uncertainty quantification (Section 4) were obtained from the NYU fastMRI Initiative database (Knoll et al., 2020; Zbontar et al., 2018). A listing of NYU fastMRI investigators, subject to updates, can be found at: fastmri.med.nyu.edu. The primary goal of fastMRI is to test whether machine learning can aid in the reconstruction of medical images.

D MORE VISUALIZATIONS

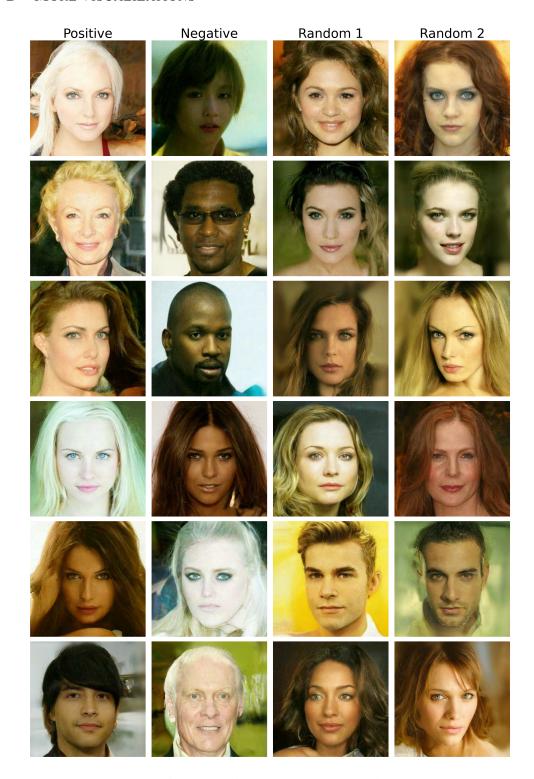


Figure 17: CelebA-HQ Image Generated

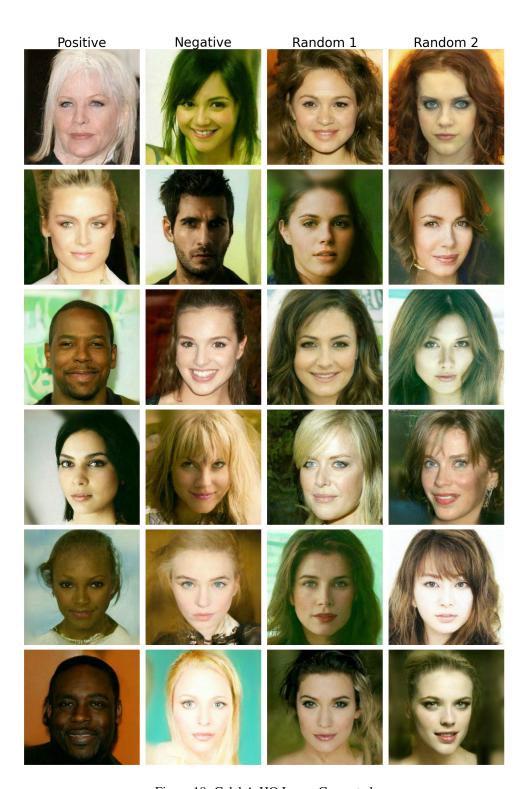


Figure 18: CelebA-HQ Image Generated

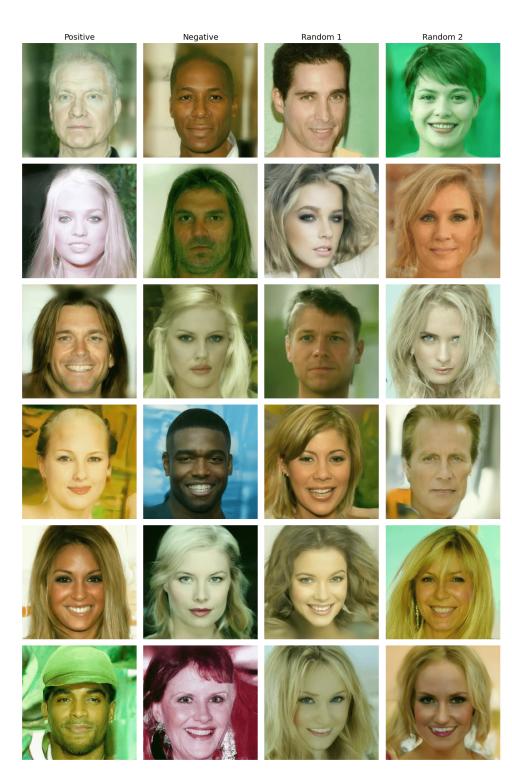


Figure 19: CelebA-HQ Image Generated



Figure 20: DiT Class 974: geyser



Figure 21: DiT Class 387: lesser panda, red panda

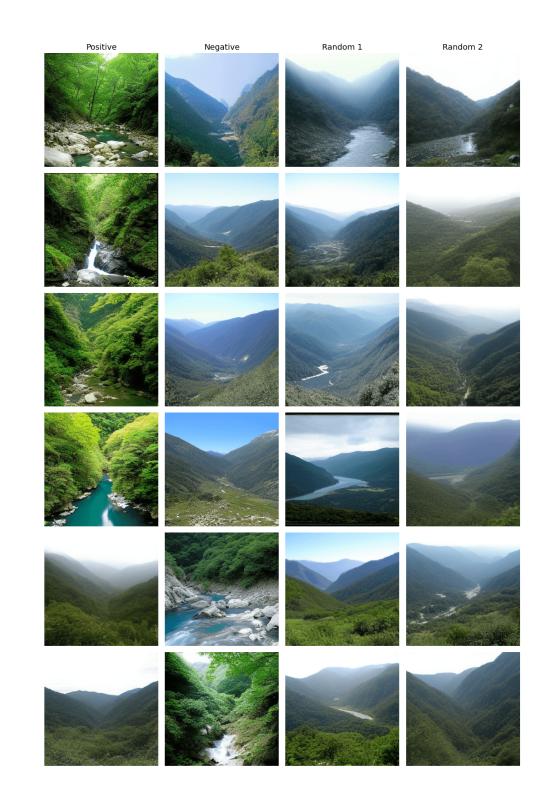


Figure 22: DiT Clas 979: valley, vale



Figure 23: DiT Class 388: giant panda, panda

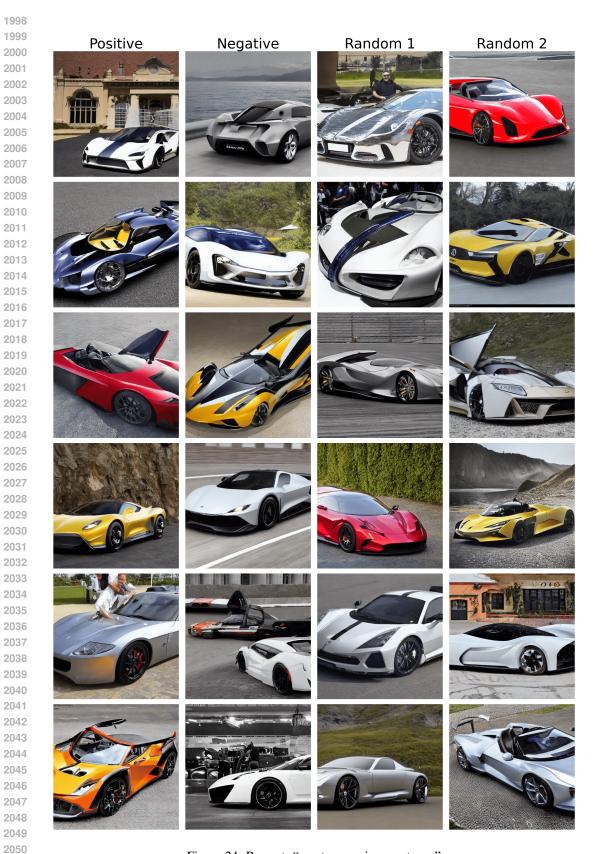


Figure 24: Prompt: "most expensive sports car"