

# DPAC: Distribution-Preserving Adversarial Control for Diffusion Sampling

Han-Jin Lee   Han-Ju Lee   Jin-Seong Kim   Seok-Hwan Choi

Yonsei University Department of Computer Science

{han-.jin, hanleju, js.kim, sh.choi}@yonsei.ac.kr

## Abstract

*Adversarially guided diffusion sampling often achieves the target class, but sample quality degrades as deviations between the adversarially controlled and nominal trajectories accumulate. We formalize this degradation as a path-space Kullback-Leibler divergence (path-KL) between controlled and nominal (uncontrolled) diffusion processes, thereby showing via Girsanov’s theorem that it exactly equals the control energy. Building on this stochastic optimal control (SOC) view, we theoretically establish that minimizing this path-KL simultaneously tightens upper bounds on both the 2-Wasserstein distance and Fréchet Inception Distance (FID), revealing a principled connection between adversarial control energy and perceptual fidelity. From a variational perspective, we derive a first-order optimality condition for the control: among all directions that yield the same classification gain, the component tangent to iso-(log-)density surfaces (i.e., orthogonal to the score) minimizes path-KL, whereas the normal component directly increases distributional drift. This leads to **DPAC** (Distribution-Preserving Adversarial Control), a diffusion guidance rule that projects adversarial gradients onto the tangent space defined by the generative score geometry. We further show that in discrete solvers, the tangent projection cancels the  $O(\Delta t)$  leading error term in the Wasserstein distance, achieving an  $O(\Delta t^2)$  quality gap; moreover, it remains second-order robust to score or metric approximation. Empirical studies on ImageNet-100 validate the theoretical predictions, confirming that DPAC achieves lower FID and estimated path-KL at matched attack success rates.*

## 1. Introduction

Denoising diffusion models [11, 20] represent the state-of-the-art in generative modeling [6, 16], largely due to their amenability to guidance [10]. This control mechanism can steer generation towards prompts or class labels, and can also be used for evaluating model robustness by generating Universal Adversarial Examples (UAEs) [4, 5, 22].

However, existing gradient-based guidance methods like AdvDiff [5] suffer from a fundamental, unresolved flaw. As guidance strength increases to maximize the Attack Success Rate (ASR), the sample quality catastrophically collapses. The resulting high-FID, artifact-laden images are invalid as “adversarial examples,” which must be both effective (high ASR) and realistic (low FID).

In this paper, we diagnose and solve this instability. We first identify that the quality collapse is caused by a “normal” (score-parallel) component of the guidance gradient. This component, while effective for the adversarial task, aggressively pushes the sampling trajectory off the data manifold. Based on this diagnosis, we propose **DPAC (Distribution-Preserving Adversarial Control)**, a new guidance framework built on the principle of *tangential control*. DPAC uses a geometric projection  $\Pi_{\perp}$  to surgically remove this harmful normal component, isolating the “tangential” gradient that steers generation within the data manifold.

Our experiments on ImageNet (256x256) provide decisive validation (Fig. 2). The results show DPAC resolves the core ASR-FID trade-off: at high guidance strengths where AdvDiff’s quality catastrophically collapses (FID 69.37), DPAC remains stable and artifact-free (FID 44.89). Moreover, our method is significantly more efficient. DPAC achieves a superior peak fidelity (FID 33.90) using only one-third the energy (guidance strength) of the baseline’s inferior optimum (FID 34.66). This  $\approx 66\%$  efficiency gain empirically verifies our theory that removing the harmful component leaves a more effective control signal.

## 2. Background

### 2.1. Diffusion Models as SDEs

Denoising diffusion models [11, 18, 20] define a forward noising process that gradually perturbs data  $x_0 \sim p_{\text{data}}$  into noise via a Stochastic Differential Equation (SDE):

$$dX_t = f(X_t, t) dt + g(t) dW_t, \quad t \in [0, 1]. \quad (1)$$

Here,  $f(\cdot, t)$  is a linear drift,  $g(t)$  is the diffusion coefficient, and  $W_t$  is a standard Wiener process. As  $t \rightarrow 1$ ,  $X_t$  converges to a simple prior, typically  $\mathcal{N}(0, I)$ .

A key insight [1, 20] is that this process is reversible. We can generate new data by solving the corresponding reverse-time SDE from  $t = 1$  to  $t = 0$ :

$$dX_t = [f(X_t, t) - g(t)^2 s_\theta(X_t, t)] dt + g(t) d\bar{W}_t, \quad (2)$$

where  $d\bar{W}_t$  is a reverse-time Wiener process. The only unknown term is the **score function**  $s_\theta(X_t, t) = \nabla_{X_t} \log p_t(X_t)$ , which points in the direction of highest data density. This function is approximated by a time-conditioned neural network  $s_\theta$  (e.g., a U-Net) trained via denoising score-matching [11].

## 2.2. Adversarial Control and its Instability

The reverse SDE provides a natural mechanism for control. An adversarial attack (or UAE generation) steers the sampling process to fool a target classifier  $v_\phi$ . This is achieved by defining a controlled SDE (as detailed in Sec. 3, Eq. (4)), which injects an additional control drift  $u_t$ :

$$dX_t = [\text{Original Drift}] + g(t)u_t(X_t, t) dt + g(t) d\bar{W}_t. \quad (3)$$

This SDE formulation provides a theoretical lens to interpret existing heuristic methods. For example, the guidance strategy used in AdvDiff [5] is equivalent to setting the control  $u_t$  as the raw classifier gradient,  $u_t \propto \nabla_{X_t} \ell_{\text{tar}}(v_\phi(X_t); y_{\text{tar}})$ , where  $y_{\text{tar}}$  is the target class and  $\ell_{\text{tar}}$  is the attack objective function (e.g., cross-entropy loss). This successfully increases the ASR.

However, this raw gradient control is the source of the instability. As we formalize in Sec. 3, any gradient vector  $u_t$  can be decomposed into two components relative to the data manifold at  $X_t$ : (1) a **tangential component** ( $u_\perp$ ), which steers the sample along the iso-density surface (preserving the density  $p_t$ ), and (2) a **normal component** ( $u_\parallel$ ), which is parallel to the score  $s_\theta$  and steers the sample *off* the manifold (distorting the density  $p_t$ ). Standard guidance (e.g., AdvDiff) uses this raw vector  $u_t$  indiscriminately, thus injecting the harmful normal component  $u_\parallel$ . This distortion explicitly pushes the sampling trajectory away from the true data distribution, which manifests as the catastrophic FID collapse observed in our experiments (Fig. 2a). Our work (Sec. 4) is motivated by the need to surgically remove this harmful component.

## 3. Theoretical Foundations

### 3.1. Reverse SDE Formulation

We begin by formalizing the concepts from Sec. 2. Starting from the score-based reverse SDE (Eq. (2)) and the general control framework (Eq. (3)), we define our complete adversarially guided sampling process as:

$$dX_t = [f(X_t, t) - g(t)^2 s_\theta(X_t, t | y_{\text{gt}})] dt + g(t)u_t(X_t, t) dt + g(t) d\bar{W}_t, \quad (4)$$

where  $s_\theta(\cdot, t | y_{\text{gt}})$  is the clean-conditioned score function for a ground-truth label  $y_{\text{gt}}$  [10], often approximated in practice with classifier-free guidance (CFG), and  $u_t$  is the adversarial control. The path distributions of the uncontrolled ( $u_t \equiv 0$ ) and controlled processes are denoted by  $\mathcal{P}^0$  and  $\mathcal{P}^u$ , respectively. The corresponding marginal densities at each time  $t$  are written  $p_t^0$  and  $p_t^u$ . The control term  $g(t)u_t dt$  perturbs the reverse diffusion only through its drift, and this specific scaling by  $g(t)$  ensures that the associated path-space KL divergence (path-KL) takes a simple quadratic form, a standard result derived from Girsanov’s Theorem [12]. Intuitively,  $u_t$  injects a small, continuous steering signal that gradually shifts the sampling trajectory toward a target label  $y_{\text{tar}}$  while maintaining consistency with the clean manifold defined by  $y_{\text{gt}}$ .

**Standing assumptions and notation.** Throughout, we assume  $g(t) > 0$  and  $f(\cdot, t)$ ,  $s_\theta(\cdot, t | y_{\text{gt}})$  are globally Lipschitz in  $x$  uniformly in  $t$  and of linear growth; these are standard conditions to guarantee the existence and uniqueness of a strong solution for the SDE [15]. We also assume that  $u_t$  is progressively measurable (i.e., non-anticipating) with  $\mathbb{E} \int_0^1 \|u_t\|_2^2 dt < \infty$ . This ensures that the Novikov condition holds, which validates the application of Girsanov’s Theorem and ensures the path-KL divergence is finite. All processes are defined on the same filtered probability space, and expectations w.r.t. a path law  $\mathcal{P}$  are written  $\mathbb{E}_{\mathcal{P}}[\cdot]$ . For brevity, we write  $s_t := s_\theta(x, t | y_{\text{gt}})$  and  $g_t := g(t)$ .

### 3.2. Path-Space Divergence and Perceptual Bounds

To quantify how strongly the control  $u_t$  perturbs the reverse-time diffusion process, we analyze the KL divergence between the path distributions  $\mathcal{P}^0$  and  $\mathcal{P}^u$ . By Girsanov’s theorem [12], under Novikov’s condition, the relative entropy between them equals the cumulative control energy:

$$KL(\mathcal{P}^u \| \mathcal{P}^0) = \frac{1}{2} \mathbb{E}_{\mathcal{P}^u} \int_0^1 \|u_t(X_t, t)\|_2^2 dt. \quad (5)$$

Hence, the path-KL directly measures the total energy injected by the control. Because it aggregates the squared deviation between the controlled and nominal trajectories, it can be interpreted as a quantitative measure of sample distortion. As shown in Fig. 1 (left), the controlled trajectory (red) deviates from the nominal diffusion (blue) in proportion to its control energy. (The equivalence in Eq. (5) relies on the drift scaling by  $g(t)$  in Eq. (4) and the nondegeneracy condition  $g(t) > 0$ .)

The path-space divergence also upper-bounds the discrepancy of terminal distributions. The marginalization inequality

$$KL(p_t^u \| p_t^0) \leq KL(\mathcal{P}^u \| \mathcal{P}^0), \quad \forall t \in [0, 1], \quad (6)$$






Figure 1. **Toy 1D diffusion illustrating the link between path-KL and distribution-preserving control.** Shown are three panels from left to right. **Left:** samples from uncontrolled ( $P^0$ ) and controlled ( $P^u$ ) 1D diffusions with  $x_0 \sim \mathcal{N}(0, 0.5^2)$ ,  $dx_t = \log(t) dt + dW_t$  (uncontrolled), and  $dx_t = [\log(t) - 2 \log(t)] dt + dW_t$  (controlled,  $u_t = -2 \log t$ ). Thin lines show sample trajectories and thick lines their means, indicating that control shifts the mean path while keeping stochastic spread. **Center:** cumulative empirical  $KL(P^u || P^0)$  (blue) and control energy  $\frac{1}{2} \sum_k u(t_k)^2 \Delta t$  (orange) coincide, verifying  $KL(P^u || P^0) = \frac{1}{2} \mathbb{E}_{P^u} \int_0^1 \|u_t\|^2 dt$  (Girsanov identity; final values 3.90, 3.92, 3.87). **Right:** iso-density contours of a 2D Gaussian mixture showing normal control (red,  $u_t^{\text{nor}} \propto -\nabla \log p_t$ ) that changes density versus tangential control (green,  $u_t^{\text{tan}}$ ) that moves along level sets. Together, these panels visualize that reducing path-KL (control energy) yields density-preserving, tangential dynamics.

implies that minimizing the path-KL implicitly limits the KL divergence at every time step. If the clean terminal density  $p_0^0$  satisfies the Talagrand  $T_2(C)$  inequality [14], then the Wasserstein distance obeys

$$W_2^2(p_0^u, p_0^0) \leq 2C KL(p_0^u || p_0^0) \leq 2C KL(P^u || P^0). \quad (7)$$

Moreover, let  $\phi$  be the  $L$ -Lipschitz feature embedding used for FID calculation [9] (hence  $W_2(\phi_{\#} p, \phi_{\#} q) \leq L W_2(p, q)$ ). Letting  $\varepsilon_{\text{gauss}}$  denote the Gaussian approximation error, the triangle inequality combined with the  $L$ -Lipschitz property of  $\phi$  yields an upper bound on the square root of the FID:

$$\begin{aligned} \sqrt{FID(\phi_{\#} p_0^u, \phi_{\#} p_0^0)} &\leq L W_2(p_0^u, p_0^0) + \varepsilon_{\text{gauss}} \\ &\leq \underbrace{L \sqrt{2C}}_{=: \tilde{K}} \sqrt{KL(P^u || P^0)} + \varepsilon_{\text{gauss}}. \end{aligned} \quad (8)$$

Thus, Eq. (6), Eq. (8) establish a direct link between the control energy ( $KL(P^u || P^0)$ ) and the perceptual fidelity (FID). This theoretical result is central to our goal: to generate high-fidelity UAEs, the adversarial control  $u_t$  must be designed to minimize this path-KL divergence, which motivates the decomposition method introduced in Sec. 4. Empirically, we also observe that as this path-space energy increases, the perceptual quality (FID) degrades proportionally, demonstrating that the path-KL serves as a unified indicator of both geometric and perceptual deviation. As shown in Fig. 1 (middle), the cumulative path-KL closely matches the control energy, verifying the Girsanov identity in Eq. (5).

### 3.3. Tangential Control and First-Order Optimality

We now investigate how to design an adversarial control  $u_t$  that minimally changes the data distribution while maximizing the target classification gain. The Fokker-Planck equation associated with Eq. (4) is [15, 20]

$$\partial_t p_t = -\nabla \cdot (\mu p_t) - \nabla \cdot (g_t u_t p_t) + \frac{1}{2} \nabla \cdot (g_t^2 \nabla p_t), \quad (9)$$

where  $\mu = f - g^2 s$ . Linearizing with respect to a small control variation yields the Gâteaux derivative

$$\partial_t (\delta p_t) = -g_t \nabla \cdot (p_t u_t). \quad (10)$$

Thus, the condition  $\nabla \cdot (p_t u_t) = 0$  ensures that the density remains constant to first order. This means  $u_t$  is tangential to the density manifold and does not alter its mass distribution. As illustrated in Fig. 1 (right), the tangential component flows along the iso-density contours (green), whereas the normal component (red, aligned with the score  $\nabla \log p_t$ ) pushes mass toward or away from density peaks.

Formally, based on a weighted Hodge-Helmholtz decomposition [2], any vector field  $u_t$  can be uniquely decomposed relative to the metric  $G_t$  as:

$$u_t = \nabla \psi_t + v_t, \quad \nabla \cdot (p_t v_t) = 0, \quad (11)$$

where  $\nabla \psi_t$  and  $v_t$  correspond to the normal (density-changing) and tangential (density-preserving), respectively. These components are orthogonal under the  $L^2(p_t, G_t)$  inner product  $\langle v, w \rangle_{L^2(p_t, G_t)} = \mathbb{E}_{p_t}[v^\top G_t w]$ . We note that for full consistency with the metric  $G_t$ , the divergence-free subspace would strictly be defined by  $\{v : \nabla \cdot (p_t G_t v) = 0\}$ . For clarity and simplicity, we adopt the  $\nabla \cdot (p_t v_t) = 0$  constraint, as both definitions coincide when  $G_t \approx I$ .

Let  $\ell_{\text{tar}}(x_0; y_{\text{tar}})$  denote the target classification loss at the terminal time  $t = 0$ . To maximize this objective, we seek the direction of steepest ascent for its expectation,  $\mathbb{E}[\ell_{\text{tar}}(X_0)]$ . The first-order variation of this expectation under a small perturbation  $u_t \mapsto u_t + \varepsilon v_t$  can be found using the adjoint method [3]. This involves defining a potential function  $\psi(x, t)$  that solves the backward Kolmogorov equation:

$$-\partial_t \psi = \mu \cdot \nabla \psi + \frac{1}{2} g_t^2 \Delta \psi, \quad \psi(\cdot, 0) = \ell_{\text{tar}}(\cdot), \quad (12)$$

where  $\psi(x, t)$  represents the expected future loss from state  $x$  at time  $t$ . This yields the sensitivity field  $h_t := g_t \nabla \psi(x, t)$  which allows us to express the change in loss as

$$\delta \mathbb{E}[\ell_{\text{tar}}(X_0)] = \varepsilon \int_0^1 \mathbb{E}[\langle h_t, v_t \rangle] dt. \quad (13)$$

When  $g_t$  is small or diffusion is ignored, the heuristic ODE adjoint  $-\dot{\lambda}_t = (\partial_x \mu)^\top \lambda_t$  offers an equivalent shorthand; we adopt the PDE-based form for rigor. Thus, this analysis provides the optimal adversarial direction: Eq. (13) shows that the target loss is maximally increased when the control  $v_t$  is aligned with the sensitivity field  $h_t$ .

We now minimize the total energy  $\frac{1}{2} \int_0^1 \|u_t\|_{L^2(p_t, G_t)}^2 dt$  subject to a fixed gain  $\int_0^1 \langle h_t, u_t \rangle_{L^2(p_t, G_t)} dt = \Gamma$  and the constraint  $\nabla \cdot (p_t u_t) = 0$ . The unique minimizer is the orthogonal projection operator, denoted  $\Pi_{\mathcal{T}_t}^{G_t}$ , onto the tangential subspace  $\mathcal{T}_t = \{v : \nabla \cdot (p_t v) = 0\}$ :

$$u_t^{\text{tan}} = \Pi_{\mathcal{T}_t}^{G_t}(G_t^{-1} h_t). \quad (14)$$

Because the normal component contributes only additional energy, the tangential projection is the first-order optimal solution. In practice, this exact projection (Eq. (14)) is computationally intractable. We approximate it by implementing a simpler projection that only removes the component along the score direction  $s_t$ . For clarity, we first define the parallel ( $\Pi_{\parallel}$ ) and perpendicular ( $\Pi_{\perp}$ ) projection operators relative to a vector  $s$  and metric  $G$  as:

$$\Pi_{\parallel}^{(s, G)} u := \frac{\langle u, s \rangle_G}{\langle s, s \rangle_G} s, \quad \Pi_{\perp}^{(s, G)} u := u - \Pi_{\parallel}^{(s, G)} u. \quad (15)$$

Our computationally efficient surrogate is then given by  $u_t^{\text{proj}} = \Pi_{\perp}^{(s_t, G_t)} w_t$ , which is written explicitly as:

$$u_t^{\text{proj}} = w_t - \frac{\langle w_t, s_t \rangle_{G_t}}{\langle s_t, s_t \rangle_{G_t}} s_t, \quad (16)$$

where  $w_t = G_t^{-1} h_t$ , and  $\langle a, b \rangle_{G_t} := a^\top G_t b$ . This enforces near-orthogonality between the adversarial control and the generative score field. Exact density preservation is achieved by the divergence-free projection (Eq. (14)); Eq. (16) is its computationally efficient first-order surrogate.

### 3.4. Discrete Bound and Robustness

We now examine the implications of this structure for discrete samplers such as DDIM [18] or DDPM [11]. We adopt a shorthand for discrete time steps, where the subscript  $k$  implies evaluation at  $(x_k, t_k)$ . Let  $\mu_k := \mu(x_k, t_k)$ ,  $g_k := g(t_k)$ ,  $u_k := u_{t_k}$ , with step size  $\Delta t_k$ ; we adopt shared noise coupling:

$$x_{k-1} = x_k + \mu_k \Delta t_k + g_k u_k \Delta t_k + g_k \sqrt{\Delta t_k} \epsilon_k, \quad (17)$$

where  $\epsilon_k \sim \mathcal{N}(0, I)$  is the standard Gaussian noise. We write  $\Delta t_{\text{max}} := \max_k \Delta t_k$ , and assume shared-noise coupling (same  $\epsilon_k$ ) for the controlled and nominal chains.

Under Lipschitz continuity of  $f$ ,  $s$  and uniform boundedness conditioning of  $G_t$ , we state the discrete  $W_2$  bound:

$$W_2^2(\hat{p}_0^u, \hat{p}_0^0) \leq C_1 \sum_{k=1}^K \|\Pi_{\parallel}^{(s_k, G_k)} u_k\|_{G_k}^2 \Delta t_k + C_2 \Delta t_{\text{max}}^2, \quad (18)$$

for constants  $C_1, C_2$  depending only on the Lipschitz and conditioning bounds. The dominant, step-size invariant term is entirely driven by the normal component  $\Pi_{\parallel} u_{t_k}$ , whereas tangential control cancels it and leaves only the  $O(\Delta t_{\text{max}}^2)$  discretization remainder. Thus, distribution-preserving control improves fidelity by removing the non-vanishing leading term in Eq. (18). Using synchronous coupling of the controlled and nominal chains with shared noise, one can show that the one-step deviation obeys  $\delta_{k-1} = \delta_k + (\mu_k^u - \mu_k^0) \Delta t_k + g_k u_k \Delta t_k$ . Under Lipschitz continuity of  $(f, s)$ , this difference grows at most linearly in  $|\delta_k|$ , and the first-order term depends solely on the score-parallel component  $\Pi_{\parallel}^{(s_k, G_k)} u_k$ . Consequently, the tangential control cancels this  $O(\Delta t)$  term, leaving only the discretization remainder  $O(\Delta t_{\text{max}}^2)$  in Eq. (18) (see supplementary materials Sec. 10.4 for a detailed proof).

Finally, we state the robustness bound for our method. When the score, metric, or gradient fields are perturbed by bounded estimation errors  $\varepsilon_t$  (score),  $\Delta_t$  (metric), and  $e_t$  (sensitivity), the resulting increase in path energy remains second-order:

$$KL(\mathcal{P}^{\tilde{u}^{\text{tan}}} \| \mathcal{P}^0) - KL(\mathcal{P}^{u^{\text{tan}}} \| \mathcal{P}^0) \leq K \int_0^1 (\|\varepsilon_t\|^2 + \|\Delta_t\|^2 + \|e_t\|^2) dt, \quad (19)$$

where  $K$  depends only on the Lipschitz constants of  $(f, s)$  and the spectral bounds of  $G_t$ . This robustness property ensures that the proposed control remains stable under approximate metrics or imperfect gradient estimates, providing a strong guarantee for practical implementations. This follows from the strong convexity of the energy functional  $\mathcal{E}(u) = \frac{1}{2} \int \|u_t\|_{L^2(p_t, G_t)}^2 dt$ , which makes the tangential

projection  $u_t^{\text{tan}} = \Pi_{\mathcal{T}_t}^{G_t}(G_t^{-1}h_t)$  non-expansive. Perturbing the score, metric, or sensitivity fields by  $(\varepsilon_t, \Delta_t, e_t)$  inflates the energy by at most a constant multiple of their squared norms (via Pythagorean/Bregman inequality), yielding the second-order bound in Eq. (19) (see supplementary materials Sec. 10.5 for a detailed proof). Together, these results confirm that minimizing the path-space KL yields a theoretically grounded principle for constructing adversarial guidance that both achieves the target classification and preserves the visual fidelity of diffusion-generated samples. This principle directly leads to the practical algorithm described in Sec. 4.

## 4. Method: DPAC

### 4.1. The Practical DPAC Projection Rule

Our method, DPAC (Distribution-Preserving Adversarial Control), implements the theoretical principle from Sec. 3 using a practical, computationally efficient surrogate. Our theory showed that the optimal, minimum-energy control  $u_t^{\text{tan}}$  is the projection of the theoretical sensitivity field  $h_t$  onto the tangential subspace  $\mathcal{T}_t$  (Eq. (14)). However, both this ideal projection  $\Pi_{\mathcal{T}_t}^{G_t}$  and the theoretical sensitivity  $h_t$  are computationally intractable.

We therefore make two practical approximations. First, we approximate the theoretical sensitivity field with a one-step lookahead gradient  $w_k$ , computed via our discrete sensitivity oracle  $\mathcal{G}_{k-1}$ :

$$w_k = \mathcal{G}_{k-1}(x_{k-1}^{\text{clean}}) := \nabla_{x_{k-1}^{\text{clean}}} \ell_{\text{tar}}(x_{k-1}^{\text{clean}}) \quad (20)$$

where  $\ell_{\text{tar}}$  is the target loss. This gradient  $w_k$  (which represents sensitivity at  $t_{k-1}$ ) is computed at each step  $k$  where the guidance scale  $\eta_k$  is active. Second, we implement the efficient pointwise surrogate derived in Eq. (16). This approximates the ideal projection by removing the score-parallel component from our practical gradient  $w_k$ :

$$u_k^* = w_k - \frac{\langle w_k, s_k \rangle_{G_k}}{\langle s_k, s_k \rangle_{G_k} + \epsilon} s_k, \quad (21)$$

with a tiny  $\epsilon \sim 10^{-8}$  for numerical stability. This provides the unscaled tangential direction  $u_k^*$ . We then apply a nonnegative scalar schedule  $\eta_k$  to control the guidance strength:  $u_k^{(\eta)} := \eta_k u_k^*$ . Consistent with our discrete bound (Eq. (18)), we use a *late-window* schedule (e.g., last 20% of steps)[5] to reduce early trajectory drift.

### 4.2. Metric Choices and Stability

Our practical implementation (Algorithm 1) requires computing an inner product between an  $x$ -space gradient ( $w \propto \nabla_x \ell$ ) and an  $\epsilon$ -space score ( $s_k$ ). This is an approximation of the true, theoretically-pure projection. We therefore test two practical, computationally cheap choices for the metric  $G_t$  used in this inner product  $\langle \cdot, \cdot \rangle_{G_t}$ .

**(A) Identity (Euclidean):**  $G_t = I$ . This is the simplest choice. It computes a standard Euclidean inner product, effectively treating the  $x$ -space gradient and  $\epsilon$ -space score as commensurate.

**(B) Noise-Scaled:**  $G_t = (1 - \alpha_t)^{-1}I$ . Here,  $\alpha_t$  denotes the cumulative signal power at time  $t$  (i.e.,  $\bar{\alpha}_t$  in standard DDPM [11] notation). This metric is motivated by the scaling of the score function in the original SDE theory. It aims to create a theoretically-consistent inner product by weighting the dot product by the inverse noise variance  $(1 - \alpha_t)^{-1}$ . We empirically evaluate its effect in Sec. 5.3.

**Stability via Normalization.** The raw gradient  $w = \mathcal{G}_{k-1}(x_{k-1}^{\text{clean}})$  (Algorithm 1, line 8) can have an arbitrarily large magnitude. Directly applying this vector (even after projection) leads to numerical instability and sample collapse (see Sec. 5.3). A simple clipping approach is insufficient, as the perturbation magnitude would remain coupled to the raw gradient scale. To ensure stable guidance, our algorithm follows a critical Project-then-Normalize sequence which disentangles the guidance direction from its magnitude. We first project the gradient  $w$  to compute the (unscaled) tangential direction  $u^*$  (line 9), and then normalize it to a unit vector  $u_{\text{hat}}$  (line 10). This ensures that the final step size of the perturbation (line 11) is *exclusively* controlled by the scalar schedule  $\eta_k$ . This disentanglement is the key to preventing the numerical collapse observed in naive guidance implementations.

### 4.3. Stable Injection via Denoise-then-Perturb

While our core theory is built upon modifying the SDE drift (Eq. (17)), we found this approach to be numerically unstable in discrete samplers. As  $g_t \Delta t \rightarrow 0$ , the drift modification leads to vanishing control signals.

Instead, we implement our tangential guidance via a robust **Denoise-then-Perturb** (PGD-style) mechanism, which is detailed in Algorithm 1. For each step  $k$ , we first compute the standard clean denoised output  $x_{k-1}^{\text{clean}}$  via the base sampler  $\Phi_{\text{base}}$ , and then apply the guided perturbation directly in  $x$ -space:

$$x_{k-1}^{\text{clean}} = \Phi_{\text{base}}(x_k, s_k, t_k) \quad (22)$$

$$x_{k-1} = x_{k-1}^{\text{clean}} + \eta_k \cdot u_{\text{hat}} \quad (23)$$

where  $u_{\text{hat}}$  is the normalized tangential direction (Algorithm 1, line 10) derived from the gradient  $w_k$  (Eq. (20)). This PGD-style injection (Eq. (23)) is a stable discrete surrogate that allows the scale  $\eta_k$  to directly control the perturbation magnitude, thus achieving an effective  $O(1)$  displacement per step, regardless of the small size of  $g_t \Delta t$ .

### 4.4. Complexity and Implementation Notes

Let  $C_{\text{score}}$  be the cost of one score evaluation and  $C_{\text{grad}}$  be the cost of one gradient (sensitivity) query.

---

**Algorithm 1** DPAC-Guided Sampling (PGD-Style)

---

**Require:**  $x_K, s_\theta, \Phi_{\text{base}}, \mathcal{G}_{k-1}$  (Gradient Oracle),  $\eta(\cdot)$ **Ensure:**  $x_0$ 

```
1:  $x \leftarrow x_K$ 
2: for  $i = 0$  to  $K - 1$  do    ▷ Loop over progress index  $i$ 
3:    $t_k \leftarrow \text{steps}[K - 1 - i]$ 
4:    $s_k \leftarrow s_\theta(x, t_k, y_{\text{gt}})$     ▷ 1. Get clean score
5:    $x_{\text{clean}} \leftarrow \Phi_{\text{base}}(x, s_k, t_k)$  ▷ 2. Denoise (DDIM step)
6:    $\eta_k \leftarrow \eta(i)$     ▷ 3. Get guidance scale
7:    $w \leftarrow \mathcal{G}_{k-1}(x_{\text{clean}})$     ▷ 4. Compute gradient
8:    $G_k \leftarrow \text{GetMetric}(t_k, s_k)$ 
9:    $u^* \leftarrow w - \frac{\langle w, s_k \rangle_{G_k}}{\langle s_k, s_k \rangle_{G_k} + \epsilon} s_k$     ▷ 5. Project gradient
10:   $u_{\text{hat}} \leftarrow u^* / (\|u^*\| + \epsilon)$     ▷ 6. Normalize direction
11:   $x \leftarrow x_{\text{clean}} + \eta_k \cdot u_{\text{hat}}$     ▷ 7. Perturb (Inject control)
12: return  $x_0$ 
```

---

**Runtime.** Per step, DPAC adds: (i) one sensitivity query  $w_k = \mathcal{G}_{k-1}(x_{k-1}^{\text{clean}})$  (see Eq. (20)); (ii) one inner product (for the projection, Eq. (21)); (iii) vector normalization and PGD perturbation. Since the sensitivity query  $w_k$  (which requires a forward and backward pass, potentially including a VAE decode) is the dominant cost, the per-step overhead during the active guidance window is roughly  $C_{\text{grad}} + C_{\text{score}}$ .

**Control Energy Estimators.** Since our final implementation (Algorithm 1) uses a PGD-style injection, the Girsanov identity (cf. Sec. 3) no longer directly applies. We therefore report the Cumulative Perturbation Energy (CPE) as our practical proxy for path energy. This metric measures the energy of the intervention vector *prior* to the normalization step (Algorithm 1, line 10), defined as:

$$\hat{\mathcal{E}}_{\text{CPE}} = \frac{1}{2} \sum_k \|\eta_k u_k\|_2^2, \quad (24)$$

where  $u_k$  is the unnormalized control vector:  $u_k = u_k^*$  (for DPAC, Eq. (21)) or  $u_k = w_k$  (for AdvDiff, Eq. (20)). We share the noise  $\{\epsilon_k\}$  across methods (common random numbers) to reduce variance in  $\hat{\mathcal{E}}_{\text{CPE}}$  and align with our bounds.

**Robustness.** The use of normalization  $u_{\text{hat}}$  and the PGD-style update inherently prevents the catastrophic numerical explosion caused by raw, unnormalized gradients. Under score/metric/sensitivity errors  $(\epsilon_t, \Delta_t, e_t)$ , the stability is maintained.

**Defaults.** Unless stated otherwise: We use PGD injection, linear  $\eta_k$  ramp in the last 20% of steps, and  $G_t = I$  for the main DPAC result.

## 5. Experiments

Our experiments are designed to empirically validate our core theoretical claims: (1) that standard gradient guidance,

as exemplified by methods like AdvDiff [5], injects a score-parallel component that destabilizes the sampling trajectory and leads to distributional collapse (high FID); and (2) that our proposed DPAC (tangential control) resolves this issue by projecting away this harmful component. We show that this theoretical correction yields simultaneous, consistent improvements in stability (FID), efficiency (energy), and effectiveness (best attainable FID).

### 5.1. Experimental Setup

**Baselines.** We conduct a direct, controlled comparison between two guidance strategies, both implemented using an identical Denoise-then-Perturb (PGD-style) injection mechanism. This isolates the impact of the *control direction* itself.

- **AdvDiff (Baseline):** We re-implement the core guidance mechanism of AdvDiff [5]. This method uses the raw classifier gradient  $w_k = \nabla_{x_{k-1}^{\text{clean}}} \ell_{\text{tar}}(x_{k-1}^{\text{clean}})$  as the control direction (i.e., gradient ascent to maximize loss). This injects a score-parallel component (normal to iso-density surfaces), which our theory predicts will distort the distribution.
- **DPAC (Ours):** The tangential control method from Sec. 4. It uses the  $G_k$ -inner-product projection of the same ascent direction  $w_k$  onto the subspace orthogonal to the score  $s_k$ , i.e.,  $u_{\text{hat}} \propto \Pi_{\perp}^{(s_k, G_k)}(w_k)$  as in Eq. (16). We test two variants for  $G_k$ : the identity ( $G_k = I$ ) and the noise-scaled metric ( $G_k = (1 - \alpha_k)^{-1} I$ ).

Note that diffusion-based UAE generation remains a nascent subfield. Among existing approaches, AdvDiff [5] is currently the only method that directly performs adversarial guidance within the diffusion trajectory for universal attacks. Other related works such as DiffPGD [22] or AdvDiffuser [4] adopt diffusion priors merely as regularizers or outer-loop optimizers rather than as in-trajectory controls. Therefore, AdvDiff serves as the only directly comparable baseline for our study.

**Implementation Details.** We use the pre-trained Latent Diffusion Model (LDM) [16] checkpoint (cin256-v2) on ImageNet-1K. For efficient evaluation, all experiments are conducted on the ImageNet-100 (IN-100) subset (classes 0-99). We use a pre-trained ResNet50 [8] (V1 legacy, ‘pretrained=True’) as the classifier  $v_\phi$  for both guidance and ASR evaluation. The CFG scale for  $y_{\text{gt}}$  is fixed to 3.0. Unless otherwise stated, the target is  $y_{\text{tar}} = (y_{\text{gt}} + 1) \pmod{1000}$ ; ASR is computed on the full 1K head (R50-V1), while FID/IS use IN-100 validation statistics. All sampling uses DDIM with 200 steps (main results) or 50 steps (ablation) on 1300 samples (13 per class). While 1300 samples is limited for absolute FID/IS estimation, our primary goal is the *relative* comparison between guidance strategies. To this end, we use shared-noise coupling across methods (as discussed in Sec. 4.4). This isolates the impact






Figure 2. **Quantitative validation of DPAC on ImageNet-100 (200 steps).** (a) **Stability:** AdvDiff (red) suffers catastrophic FID collapse ( $39.9 \rightarrow 69.37$ ) at high guidance scales ( $\eta_k$ ). DPAC (blue) remains robustly stable. (b) **Effectiveness & Efficiency:** A direct comparison of the **best FID** each method achieved. DPAC (blue) achieves a **superior peak fidelity (FID 33.90)** while using only **one-third** of the energy ( $\mathcal{E}=54.0$ ) that AdvDiff (red) required for its worse optimum (FID 34.66 at  $\mathcal{E}=160.0$ ). (c) **Theoretical Validation:** At all scales, DPAC consistently uses  $\approx 66\%$  less energy, empirically validating our theory.

of the control  $u_k$  from stochastic sampling variance, aligning with our theoretical bounds (Eq. (18)) and ensuring a robust relative comparison. The control signal is injected during the last 20% of steps (late-window) with a linear ramp-up of  $\eta_k$ .

**Metrics.** We evaluate three key aspects:

- **Effectiveness (ASR, R50-V1):** Attack Success Rate—the percentage of samples classified as  $y_{\text{tar}}$  by the ResNet-50 (V1) classifier.
- **Fidelity & Stability (FID/IS):** We use Fréchet Inception Distance (FID) [9] (lower is better) and Inception Score (IS) [17] (higher is better).
- **Efficiency (CPE):** We report CPE, as defined in Eq. (24) (Sec. 4.4). This metric measures the energy of the unnormalized control vector  $u_k$  and aligns with our theoretical analysis (cf. Sec. 5.2).

## 5.2. Core Validation

Our main results at 200 DDIM steps (S200) (Tab. 1 and Fig. 2) provide strong evidence for these claims.

**Stability (Solving Collapse).** Fig. 2(a) provides the clearest view of the stability crisis. As the guidance scale  $\eta_k$  increases, AdvDiff (red) suffers a catastrophic failure. At  $\eta_k = 10.0$ , its fidelity collapses, with the FID exploding from 39.9 to 69.37. This empirically confirms our hypothesis: the unconstrained score-parallel component destabilizes the sampling trajectory. In stark contrast, DPAC (blue) remains robust, maintaining a stable FID of 44.89.

**Efficiency (Validating Theory).** This stability gain is accompanied by a dramatic and consistent improvement in energy efficiency. As shown in Fig. 2(c), DPAC consistently consumes  $\approx 66\%$  less path energy than AdvDiff across the entire S200 run (e.g., 1350.4 vs 4000.0 at  $\eta_k = 10.0$ ). This near-constant reduction matches the geometric prediction from our theory (Sec. 3): the ratio  $\mathcal{E}_{\text{DPAC}}/\mathcal{E}_{\text{AdvDiff}} \approx$

$0.34$  corresponds to  $\sin^2 \theta$  (the score-orthogonal share retained by DPAC), while the baseline’s score-parallel share is  $\cos^2 \theta \approx 0.66$ .

**Effectiveness (Peak Fidelity & Cost).** Fig. 2(b) directly compares the best possible fidelity each method could achieve. AdvDiff (red) achieves its optimal fidelity (FID 34.66) at a high energy cost of  $\mathcal{E} = 160.0$ . DPAC (blue) achieves a superior optimal fidelity (FID 33.90) using only one-third of the energy ( $\mathcal{E} = 54.02$ ). This demonstrates that DPAC is not a compromise: it achieves a better peak quality for a fraction of the cost, cleanly separating fidelity from guidance strength.

## 5.3. Ablation Studies

We conduct ablation studies using 50 DDIM steps (S50) to verify our design choices.

**Robustness to Sampling Steps.** The S50 results confirm our findings from S200. DPAC consistently saves  $\approx 64.8\%$  energy (e.g., 87.96 vs. 250.0 at  $\eta_k = 5.0$ ) and avoids the baseline’s FID collapse (35.05 vs. 39.92 at  $\eta_k = 10.0$ ). This confirms our method’s advantages are robust to the number of DDIM steps. We ablate the choice of the inner product matrix  $G_k$  used for the orthogonal projection, comparing our default  $G_k = I$  (identity) against the theoretically-motivated  $G_k = (1 - \alpha_k)^{-1}I$  (noise\_scaled). As shown in Tab. 2, the choice of metric has a negligible impact on the final results. Both metrics achieve a nearly identical ASR-FID trade-off. This suggests that while  $G_k = (1 - \alpha_k)^{-1}I$  is geometrically precise, the practical choice of  $G_k = I$  is sufficient, simpler, and achieves the same state-of-the-art stability.

**Further Analysis.** We provide comprehensive additional experiments and in-depth analyses to further validate our claims in the Supplementary Material (Sec. 13).

Table 1. **Quantitative comparison on ImageNet-100 (200 steps) across all guidance scales.** DPAC achieves a superior optimal FID (33.90) at one-third the energy cost of AdvDiff’s optimum (FID 34.66). Crucially, AdvDiff suffers catastrophic FID collapse ( $39.9 \rightarrow 69.37$ ) at high scales, while DPAC remains stable (44.89). At every scale, DPAC is more efficient, consuming  $\approx 66\%$  less Path Energy.

| Metric                                   | Scale ( $\eta_k$ ) | AdvDiff (Baseline)          |                      |                   |                         | DPAC (Ours, $G_k = I$ )     |                      |                   |                         |
|------------------------------------------|--------------------|-----------------------------|----------------------|-------------------|-------------------------|-----------------------------|----------------------|-------------------|-------------------------|
|                                          |                    | ASR (R50–V1) ( $\uparrow$ ) | FID ( $\downarrow$ ) | IS ( $\uparrow$ ) | Energy ( $\downarrow$ ) | ASR (R50–V1) ( $\uparrow$ ) | FID ( $\downarrow$ ) | IS ( $\uparrow$ ) | Energy ( $\downarrow$ ) |
| Optimal FID $\rightarrow$<br>& Stability | 0.5                | 97.1%                       | 35.35                | 33.30             | 10.0                    | 50.1%                       | 35.74                | <b>34.29</b>      | <b>3.38</b>             |
|                                          | 1.0                | <b>100.0%</b>               | 34.98                | 32.19             | 40.0                    | 97.9%                       | 34.58                | <b>33.04</b>      | <b>13.50</b>            |
|                                          | 2.0                | <b>100.0%</b>               | <i>34.66</i>         | 30.72             | <i>160.0</i>            | <b>100.0%</b>               | <b>33.90</b>         | <b>31.48</b>      | <b>54.02</b>            |
|                                          | 5.0                | 100.0%                      | 39.90                | 24.77             | 1000.0                  | 99.9%                       | 35.09                | <b>28.46</b>      | <b>337.61</b>           |
|                                          | 10.0               | 100.0%                      | <b>69.37</b>         | 20.62             | 4000.0                  | 99.6%                       | <b>44.89</b>         | <b>22.86</b>      | <b>1350.43</b>          |

Table 2. **Ablation on Inner Product Metric  $G_k$  (S50).** Both metrics yield a nearly identical ASR–FID trade-off, validating the simpler choice  $G_k = I$ .

| Method                                | Scale ( $\eta_k$ ) | ASR (R50–V1) ( $\uparrow$ ) | FID ( $\downarrow$ ) |
|---------------------------------------|--------------------|-----------------------------|----------------------|
| DPAC ( $G_k = I$ )                    | 5.0                | 93.3%                       | 34.46                |
| DPAC ( $G_k = (1 - \alpha_k)^{-1}I$ ) | 5.0                | 92.6%                       | 34.54                |
| DPAC ( $G_k = I$ )                    | 10.0               | 99.3%                       | 35.05                |
| DPAC ( $G_k = (1 - \alpha_k)^{-1}I$ ) | 10.0               | 99.3%                       | 34.74                |

## 6. Related Works

**Adversarial Examples.** Classical  $\ell_p$ -bounded attacks such as FGSM [7] and PGD [13] optimize a perturbation  $\delta$  for a given input  $x$ . Beyond norm-bounded threats, unrestricted or generative adversarial examples synthesize adversarial images from scratch using generative models (e.g., AC-GAN) [19] or produce perturbations via GANs [21].

**Guidance for Diffusion Models.** Diffusion guidance methods like Classifier Guidance (CG) [6] and Classifier-Free Guidance (CFG) [10] steer sampling using classifier gradients or unconditional scores, respectively, to trade off fidelity and diversity. While both mechanisms are effective, increasing the guidance scale is known to improve alignment at the expense of sample quality (e.g., FID), revealing an intrinsic guidance–fidelity tension [6, 10].

### Diffusion + PGD: Gradient Injection During Sampling.

A body of work integrates PGD-style optimization with diffusion to enhance realism and controllability. For example, Diff-PGD [22] leverages a diffusion prior to keep attacks on-manifold, while AdvDiffuser [4] conducts step-wise optimization to perturb less-salient regions.

**Generative Adversarial Examples with Diffusion.** AdvDiff [5] stands as the primary baseline, generating UAEs by injecting raw classifier gradients into the continuous diffusion sampling process. While this method suffers from a severe fidelity collapse that prior work treats as an empirical trade-off, our work (DPAC) is the first to diagnose and solve this instability via tangential projection. Other methods like Diff-PGD [22] and AdvDiffuser [4] are not directly

comparable, as they apply gradient optimization \*outside\* the reverse SDE dynamics, unlike the continuous guidance setting of AdvDiff and DPAC.

## 7. Discussion

Our analysis assumes non-degenerate diffusion  $g(t) > 0$ , globally Lipschitz drift and score with linear growth, and a  $T_2(C)$  tail for  $p_0^0$ , which together guarantee a strong reverse SDE solution, the Girsanov identity equating path–KL with control energy, and the information–transport bounds summarized in Sec. 3. The exact  $p_t$ -divergence-free projector is intractable in high dimensions; our practical rule removes only the score-parallel component and is therefore a first-order surrogate for the true  $p_t$ -tangential field, so that with approximate scores or finite-step discretization a small density-changing residual can remain. This motivates the late-window schedule and the project-then-normalize update in Sec. 4, which curb early manifold drift and decouple direction from magnitude. Because the denoise-then-perturb injection breaks the exact path–KL identity, we report CPE as a proxy and use shared-noise coupling to reduce variance, which aligns with the FID/energy trends observed in Sec. 5. Finally, because DPAC enables high-fidelity targeted UAEs, deployments should include provenance signals and audit logging to mitigate misuse.

## 8. Conclusion

We cast adversarial guidance for diffusion as path–KL minimization and proved that reducing control energy tightens terminal  $W_2$ /FID bounds, that the  $p_t$ -tangential component is the first-order optimal direction under a fixed gain, that in discrete samplers the tangential rule cancels the  $O(\Delta t)$  leading term in the discrete  $W_2$  bound, and that the solution is second-order robust to score, metric, and gradient errors. We instantiated these ideas with a score-orthogonal projector and a stable denoise-then-perturb update, achieving on ImageNet-100 a consistent reduction of FID and estimated path energy at matched ASR, eliminating the collapse seen with raw gradient guidance. The results suggest that target attainment and fidelity preservation can be unified by an energy-minimization principle.

## References

- [1] Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. [2](#)
- [2] Harsh Bhatia, Gregory Norgard, Valerio Pascucci, and Peer-Timo Bremer. The helmholtz-hodge decomposition—a survey. *IEEE Transactions on Visualization and Computer Graphics*, 19(8):1386–1404, 2013. [3](#)
- [3] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, page 6572–6583, Red Hook, NY, USA, 2018. Curran Associates Inc. [4](#)
- [4] Xinquan Chen, Xitong Gao, Juanjuan Zhao, Kejiang Ye, and Cheng-Zhong Xu. Advdiffuser: Natural adversarial example synthesis with diffusion models. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4539–4549, 2023. [1](#), [6](#), [8](#)
- [5] Xuelong Dai, Kaisheng Liang, and Bin Xiao. Advdiff: Generating unrestricted adversarial examples using diffusion models. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part XLVI*, page 93–109, Berlin, Heidelberg, 2024. Springer-Verlag. [1](#), [2](#), [5](#), [6](#), [8](#)
- [6] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2021. Curran Associates Inc. [1](#), [8](#)
- [7] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [8](#)
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [6](#)
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 6629–6640, Red Hook, NY, USA, 2017. Curran Associates Inc. [3](#), [7](#)
- [10] Jonathan Ho. Classifier-free diffusion guidance. *ArXiv*, abs/2207.12598, 2022. [1](#), [2](#), [8](#)
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [1](#), [2](#), [4](#), [5](#)
- [12] Gopinath Kallianpur and Rajeeva L. Karandikar. *Girsanov’s Theorem*, pages 95–101. Birkhäuser Boston, Boston, MA, 2000. [2](#)
- [13] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [8](#)
- [14] Michael Molloy and Bruce Reed. *A Closer Look at Talagrand’s Inequality*, pages 231–236. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002. [3](#)
- [15] Bernt Oksendal. *Stochastic differential equations (3rd ed.): an introduction with applications*. Springer-Verlag, Berlin, Heidelberg, 1992. [2](#), [3](#)
- [16] Robin Rombach, A. Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10674–10685, 2021. [1](#), [6](#)
- [17] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 2234–2242, Red Hook, NY, USA, 2016. Curran Associates Inc. [7](#)
- [18] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#), [4](#)
- [19] Yang Song, Rui Shu, Nate Kushman, and Stefano Ermon. Constructing unrestricted adversarial examples with generative models. In *Neural Information Processing Systems*, 2018. [8](#)
- [20] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [1](#), [2](#), [3](#)
- [21] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, page 3905–3911. AAAI Press, 2018. [8](#)
- [22] Haotian Xue, Alexandre Araujo, Bin Hu, and Yongxin Chen. Diffusion-based adversarial sample generation for improved stealthiness and controllability. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2023. Curran Associates Inc. [1](#), [6](#), [8](#)

# DPAC: Distribution-Preserving Adversarial Control for Diffusion Sampling

## Supplementary Material

### 9. Notation and Assumptions

This section summarizes the notation used throughout the main paper and supplementary materials, and explicitly states the standing assumptions required for our theoretical analysis.

#### 9.1. Notation

- $t \in [0, 1]$ : Continuous time index, where  $t = 0$  corresponds to data and  $t = 1$  corresponds to noise.
- $X_t$ : State of the diffusion process at time  $t$ .  $x_0 \sim p_{\text{data}}$  is the clean data.
- $p_t(X_t)$ : Marginal density of  $X_t$ .
- $f(X_t, t)$ : Drift coefficient of the forward SDE.
- $g(t)$ : Diffusion coefficient of the forward SDE.
- $W_t, \bar{W}_t$ : Standard Wiener processes (forward and reverse).
- $s_\theta(X_t, t)$ : The score function  $\nabla_{X_t} \log p_t(X_t)$ , approximated by a neural network  $s_\theta$ .
- $y_{\text{gt}}, y_{\text{tar}}$ : Ground-truth label and target (adversarial) label, respectively.
- $\mu(X_t, t) = f(X_t, t) - g(t)^2 s_\theta(X_t, t|y_{\text{gt}})$ : The drift of the nominal (uncontrolled) reverse SDE.
- $u_t = u_t(X_t, t)$ : The adversarial control drift term injected into the reverse SDE.
- $\mathcal{P}^0, \mathcal{P}^u$ : Path-space measures (distributions over trajectories) for the uncontrolled ( $u_t \equiv 0$ ) and controlled processes, respectively.
- $p_t^0, p_t^u$ : Marginal densities at time  $t$  for the uncontrolled and controlled processes.
- $KL(\mathcal{P}^u || \mathcal{P}^0)$ : The Kullback-Leibler (KL) divergence between the two path-space measures. This is referred to as “path-KL”.
- $W_2(\cdot, \cdot)$ : The 2-Wasserstein distance.
- $\phi$ : An L-Lipschitz feature embedding used for FID calculation.
- $\psi(x, t)$ : The potential function (expected future loss) from state  $x$  at time  $t$ , used in the adjoint method.
- $h_t = g_t \nabla \psi(x, t)$ : The sensitivity field, representing the optimal adversarial direction.
- $G_t$ : A metric tensor, used to define the inner product  $\langle \cdot, \cdot \rangle_{G_t}$  and geometry.
- $\Pi_{\parallel}^{(s, G)}, \Pi_{\perp}^{(s, G)}$ : Projection operators onto the subspaces parallel and perpendicular to a vector  $s$  with respect to the metric  $G$ .
- $\mathcal{T}_t = \{v : \nabla \cdot (p_t v) = 0\}$ : The tangential (density-preserving) subspace.
- $k$ : Discrete time step index.
- $w_k$ : The practical gradient computed via the sensitivity

oracle  $\mathcal{G}_{k-1}$ .

- $u_k^*, u_{\text{hat}}$ : The unscaled and normalized tangential direction vectors in the discrete implementation.
- $\eta_k$ : The scalar guidance scale schedule.
- $\mathcal{E}_{CPE}$ : Cumulative Perturbation Energy, the practical proxy for path-KL.

#### 9.2. Standing Assumptions

Our theoretical results in Sec. 3 rely on the following standard assumptions from stochastic optimal control and SDE literature:

1. **(Non-degeneracy)** The diffusion coefficient is strictly positive:  $g(t) > 0$  for  $t \in [0, 1]$ . This ensures the SDE is non-degenerate and path-space measures are equivalent (Girsanov’s Theorem applies).
2. **(SDE Regularity)** The drift  $f(\cdot, t)$  and the score function  $s_\theta(\cdot, t|y_{\text{gt}})$  are globally Lipschitz continuous in  $x$  uniformly in  $t$ , and satisfy linear growth conditions. This guarantees the existence and uniqueness of a strong solution for the SDEs.
3. **(Finite Control Energy)** The adversarial control  $u_t$  is progressively measurable and satisfies the finite energy condition:  $\mathbb{E} \int_0^1 \|u_t\|_2^2 dt < \infty$ . This ensures Novikov’s condition holds, which validates the application of Girsanov’s Theorem and guarantees the path-KL (Eq. (5)) is finite.
4. **(Information-Transport)** The clean data distribution at  $t = 0$ ,  $p_0^0$ , satisfies the Talagrand  $T_2(C)$  inequality for some constant  $C$ . This allows us to bound the  $W_2$  distance by the KL divergence (Eq. (7)).
5. **(FID Embedding)** The FID feature embedding  $\phi$  is L-Lipschitz continuous. This relates the  $W_2$  distance in the data space to the  $W_2$  distance in the feature space (Eq. (8)).
6. **(Discrete Bound Conditions)** For the discrete  $W_2$  bound (Eq. (18)), we assume Lipschitz continuity of  $f$  and  $s$ , and uniform boundedness and conditioning of the metric  $G_t$ .
7. **(Robustness Bound Conditions)** For the robustness bound (Eq. (19)), we assume the estimation errors for the score ( $\epsilon_t$ ), metric ( $\Delta_t$ ), and sensitivity ( $e_t$ ) are bounded.

### 10. Detailed Proofs

This section provides detailed derivations and proofs for the theoretical claims made in Sec. 3 of the main paper. We use the notation and assumptions defined in Sec. 9.



### 10.1. Girsanov Identity

We provide a line-by-line derivation of the identity stated in Eq. (5).

**Uncontrolled and controlled reverse SDEs.** Let the nominal (uncontrolled) reverse SDE be

$$dX_t = \mu(X_t, t) dt + g(t) d\bar{W}_t, \quad (25)$$

where  $\mu_t = f_t - g_t^2 s_\theta$ . With control injected as a drift scaled by  $g(t)$ ,

$$dX_t = [\mu(X_t, t) + g(t)u_t] dt + g(t) d\bar{W}_t. \quad (26)$$

**Assumptions and Radon–Nikodym derivative.** Under Sec. 9.2 (non-degeneracy  $g(t) > 0$  and Novikov/finite energy), the path laws  $\mathcal{P}^0$  and  $\mathcal{P}^u$  are equivalent, and Girsanov’s theorem gives

$$\frac{d\mathcal{P}^u}{d\mathcal{P}^0} = \exp\left(\int_0^1 \langle u_t, d\bar{W}_t \rangle - \frac{1}{2} \int_0^1 \|u_t\|_2^2 dt\right). \quad (27)$$

**Path–KL as control energy.** Using Eq. (27) and defining the path KL as  $KL(\mathcal{P}^u \|\mathcal{P}^0) = \mathbb{E}_{\mathcal{P}^u}[\log \frac{d\mathcal{P}^u}{d\mathcal{P}^0}]$ ,

$$KL(\mathcal{P}^u \|\mathcal{P}^0) = \mathbb{E}_{\mathcal{P}^u} \left[ \int_0^1 \langle u_t, d\bar{W}_t \rangle - \frac{1}{2} \int_0^1 \|u_t\|_2^2 dt \right] \quad (28)$$

$$= \mathbb{E}_{\mathcal{P}^u} \left[ \int_0^1 \langle u_t, d\bar{W}_t^u + u_t dt \rangle - \frac{1}{2} \int_0^1 \|u_t\|_2^2 dt \right] \quad (29)$$

$$= \underbrace{\mathbb{E}_{\mathcal{P}^u} \left[ \int_0^1 \langle u_t, d\bar{W}_t^u \rangle \right]}_{=0} + \frac{1}{2} \mathbb{E}_{\mathcal{P}^u} \left[ \int_0^1 \|u_t\|_2^2 dt \right] \quad (30)$$

$$= \frac{1}{2} \mathbb{E}_{\mathcal{P}^u} \int_0^1 \|u_t\|_2^2 dt. \quad (31)$$

Here  $d\bar{W}_t^u := d\bar{W}_t - u_t dt$  is a Wiener process under  $\mathcal{P}^u$  by Girsanov, and the stochastic integral has zero expectation under  $\mathcal{P}^u$ . *Remark.* The specific scaling by  $g(t)$  in Eq. (26) is chosen so that Eq. (27) leads to the simple quadratic form in Eq. (31).error term.

### 10.2. KL–FID Connection

**Path–KL  $\rightarrow$  marginal KL (Eq. (6)).** Marginalization is a measurable map from a path to its state  $X_t$ , so the data-processing inequality yields, for all  $t \in [0, 1]$ ,

$$KL(p_t^u \| p_t^0) \leq KL(\mathcal{P}^u \|\mathcal{P}^0). \quad (32)$$

In particular, this holds at the terminal time  $t=0$ .

**Marginal KL  $\rightarrow W_2$  (Eq. (7)).** Assuming the clean terminal distribution satisfies Talagrand  $T_2(C)$ ,

$$W_2^2(p_0^u, p_0^0) \leq 2C KL(p_0^u \| p_0^0) \leq 2C KL(\mathcal{P}^u \|\mathcal{P}^0), \quad (33)$$

where the last inequality follows by combining Eq. (32) with  $t=0$ .

**$W_2 \rightarrow$  FID (Eq. (8)).** Let  $\phi$  be the  $L$ -Lipschitz feature embedding used in FID. Then

$$W_2(\phi_{\#} p_0^u, \phi_{\#} p_0^0) \leq L W_2(p_0^u, p_0^0). \quad (34)$$

Since FID computes the  $W_2$  between Gaussian fits in feature space, let  $\varepsilon_{\text{gauss}}$  denote the approximation gap; by the triangle inequality,

$$\begin{aligned} \sqrt{\text{FID}(\phi_{\#} p_0^u, \phi_{\#} p_0^0)} &\leq W_2(\phi_{\#} p_0^u, \phi_{\#} p_0^0) + \varepsilon_{\text{gauss}} \\ &\leq L W_2(p_0^u, p_0^0) + \varepsilon_{\text{gauss}} \\ &\leq L\sqrt{2C} \sqrt{KL(\mathcal{P}^u \|\mathcal{P}^0)} + \varepsilon_{\text{gauss}}, \end{aligned} \quad (35)$$

where the last step uses Eq. (33). Letting  $\tilde{K} := L\sqrt{2C}$ ,

$$\sqrt{\text{FID}(\phi_{\#} p_0^u, \phi_{\#} p_0^0)} \leq \tilde{K} \sqrt{KL(\mathcal{P}^u \|\mathcal{P}^0)} + \varepsilon_{\text{gauss}}. \quad (36)$$

Equations (32), (33) and (35) recover the main-text chain Eqs. (6) to (8): minimizing the path–KL (control energy) tightens an explicit upper bound on  $W_2$  and hence on FID.

### 10.3. Weighted Hodge and First-Order Optimality

**Claim and context (summary first).** Under the linearized Fokker–Planck dynamics (Eqs. (9) and (10)), the first-order density-preserving (tangential) constraint is  $\nabla \cdot (p_t u_t) = 0$ . In the Hilbert space  $L^2(p_t, G_t)$  with inner product  $\langle v, w \rangle_{L^2(p_t, G_t)} = \mathbb{E}_{p_t}[v^\top G_t w]$ , the  $p_t$ -weighted Hodge decomposition (Eq. (11)) writes any control as  $u_t = \nabla \psi_t + v_t$  with  $v_t \in \mathcal{T}_t := \{v : \nabla \cdot (p_t v) = 0\}$  and  $\nabla \psi_t \perp v_t$ . For the constrained program

$$\begin{aligned} \min_u \quad & \mathcal{E}(u) := \frac{1}{2} \int_0^1 \|u_t\|_{L^2(p_t, G_t)}^2 dt \\ \text{s.t.} \quad & \begin{cases} \int_0^1 \langle h_t, u_t \rangle_{L^2(p_t, G_t)} dt = \Gamma, \\ \nabla \cdot (p_t u_t) = 0, \end{cases} \end{aligned} \quad (37)$$

the unique minimizer is the orthogonal projection

$$u_t^* = \Pi_{\mathcal{T}_t}^{G_t}(G_t^{-1} h_t), \quad (38)$$

i.e., the optimal control has zero normal component and lives entirely in the tangential subspace, matching the main-text statement (Eq. (14)). Here  $h_t$  is the adjoint (sensitivity) field defined via the backward Kolmogorov equation (Eq. (12)).

**Weighted Hodge decomposition and energy split.** By orthogonality in  $L^2(p_t, G_t)$ ,

$$\begin{aligned} u_t &= \nabla \psi_t + v_t, \quad v_t \in \mathcal{T}_t, \\ \mathcal{E}(u) &= \frac{1}{2} \int_0^1 \|\nabla \psi_t\|_{L^2(p_t, G_t)}^2 dt + \frac{1}{2} \int_0^1 \|v_t\|_{L^2(p_t, G_t)}^2 dt. \end{aligned} \quad (39)$$

Because feasible controls satisfy  $u_t \in \mathcal{T}_t$ , the normal part  $\nabla \psi_t$  is infeasible and, more generally, any residual normal component can only *increase* energy by Eq. (40) without improving feasibility.

**Reduction to a projected least-norm problem.** The constraints in Eq. (37) reduce the search to the closed subspace  $\mathcal{T}_t$  with a linear gain functional  $u \mapsto \int_0^1 \langle h_t, u_t \rangle_{L^2(p_t, G_t)} dt$ . By Riesz representation in  $L^2(p_t, G_t)$ , the least-norm solution for a fixed gain  $\Gamma$  is aligned with the orthogonal projection of the representer onto  $\mathcal{T}_t$ :

$$\begin{aligned} u_t^* &= \lambda \Pi_{\mathcal{T}_t}^{G_t}(G_t^{-1} h_t), \\ \lambda > 0 \text{ chosen so that } \int_0^1 \langle h_t, u_t^* \rangle_{L^2(p_t, G_t)} dt &= \Gamma. \end{aligned} \quad (41)$$

Ignoring the scalar  $\lambda$  (directional statement), this is precisely Eq. (38) and the main-text formula (Eq. (14)).

**PDE (Lagrangian) viewpoint.** Equivalently, form the Lagrangian with a scalar multiplier  $\lambda$  for the gain and a pressure-like potential  $\pi_t$  for the divergence constraint:

$$\begin{aligned} \mathcal{L}(u, \lambda, \pi) &= \frac{1}{2} \int_0^1 \|u_t\|_{L^2(p_t, G_t)}^2 dt \\ &\quad - \lambda \int_0^1 \langle h_t, u_t \rangle_{L^2(p_t, G_t)} dt \\ &\quad + \int_0^1 \int \pi_t \nabla \cdot (p_t u_t) dx dt. \end{aligned} \quad (42)$$

Stationarity in  $u_t$  gives  $G_t u_t - \lambda h_t + \nabla \pi_t = 0$ , i.e.,

$$u_t = G_t^{-1}(\lambda h_t - \nabla \pi_t). \quad (43)$$

Imposing  $\nabla \cdot (p_t u_t) = 0$  selects the  $p_t$ -divergence-free component of  $G_t^{-1} h_t$ , hence Eq. (38). This derivation also makes explicit that any gradient (normal) component arises from  $\nabla \pi_t$  and is removed by the constraint.

**Metric note and practical surrogate.** For full metric consistency the true tangential space is  $\{v : \nabla \cdot (p_t G_t v) = 0\}$ ; we adopt  $\nabla \cdot (p_t v) = 0$  for clarity, which coincides when  $G_t \approx I$  (as in the main text). The exact projector  $\Pi_{\mathcal{T}_t}^{G_t}$  is intractable in high dimensions; this motivates the score-orthogonal surrogate used in practice (Eq. (16)) and the late-window schedule in the method section, which approximate the theoretical tangential control at low cost.

## 10.4. Discrete Wasserstein Bound

**Statement (main-text Eq. (18)).** Let  $\{t_k\}_{k=0}^K$  be a partition of  $[0, 1]$  with  $\Delta t_k = t_k - t_{k-1}$  and  $\Delta t_{\max} = \max_k \Delta t_k$ . Consider the synchronous-noise discretization in Eq. (17):

$$x_{k-1}^0 = x_k^0 + \mu(x_k^0, t_k) \Delta t_k + g_k \sqrt{\Delta t_k} \varepsilon_k, \quad (44)$$

$$x_{k-1}^u = x_k^u + \mu(x_k^u, t_k) \Delta t_k + g_k u_k \Delta t_k + g_k \sqrt{\Delta t_k} \varepsilon_k, \quad (45)$$

where  $\varepsilon_k \sim \mathcal{N}(0, I)$  are shared across the two chains,  $g_k := g(t_k)$ , and  $\mu = f - g^2 s$ . Assume (i)  $\mu(\cdot, t)$  is  $L$ -Lipschitz in  $x$  uniformly in  $t$  and of linear growth, (ii)  $\|G_k\|$  and  $\|G_k^{-1}\|$  are uniformly bounded for all  $k$  ( $mI \preceq G_k \preceq MI$ ), and (iii)  $g_{\min} \leq g_k \leq g_{\max}$ . Then there exist constants  $C_1, C_2 > 0$  (depending only on  $L, m, M, g_{\max}$ ) such that

$$W_2^2(\hat{p}_0^u, \hat{p}_0^0) \leq C_1 \sum_{k=1}^K \|\Pi_{\parallel}^{(s_k, G_k)} u_k\|_{G_k}^2 \Delta t_k + C_2 \Delta t_{\max}^2, \quad (46)$$

where  $s_k := s_\theta(x_k^0, t_k)$  (or any 1st-order consistent choice) and  $\Pi_{\parallel}^{(s_k, G_k)}$  is the  $G_k$ -orthogonal projector onto  $\text{span}\{s_k\}$ .

**Coupling and one-step deviation.** Let  $\delta_k := x_k^u - x_k^0$ . Subtracting Eq. (44) from Eq. (45) cancels the noise:

$$\delta_{k-1} = \delta_k + [\mu(x_k^u, t_k) - \mu(x_k^0, t_k)] \Delta t_k + g_k u_k \Delta t_k. \quad (47)$$

By Lipschitzness of  $\mu$ ,  $\|\mu(x_k^u, t_k) - \mu(x_k^0, t_k)\| \leq L \|\delta_k\|$ .

**Linearization: isolating the first-order source.** Using the mean-value form,

$$\begin{aligned} \mu(x_k^u, t_k) - \mu(x_k^0, t_k) &= J_k \delta_k + r_k, \\ \|J_k\| &\leq L, \quad \|r_k\| \leq c_1 \|\delta_k\|^2, \end{aligned} \quad (48)$$

for a constant  $c_1$ . Inserting into Eq. (47),

$$\delta_{k-1} = (I + J_k \Delta t_k) \delta_k + g_k u_k \Delta t_k + r_k \Delta t_k. \quad (49)$$

**Normal/tangential split and one-step  $G_k$ -bound.** We now decompose  $u_k = u_k^{\parallel} + u_k^{\perp}$  and analyze how each component contributes to the squared  $G_k$ -norm of Eq. (49).

$$\begin{aligned} \|\delta_{k-1}\|_{G_k}^2 &= \|(I + J_k \Delta t_k) \delta_k + g_k u_k^{\parallel} \Delta t_k \\ &\quad + g_k u_k^{\perp} \Delta t_k + r_k \Delta t_k\|_{G_k}^2 \\ &= \|(I + J_k \Delta t_k) \delta_k + g_k u_k^{\parallel} \Delta t_k\|_{G_k}^2 \\ &\quad + 2g_k \Delta t_k \langle (I + J_k \Delta t_k) \delta_k + g_k u_k^{\parallel} \Delta t_k, u_k^{\perp} \rangle_{G_k} \\ &\quad + g_k^2 \|u_k^{\perp}\|_{G_k}^2 \Delta t_k^2 \\ &\quad + O(\|\delta_k\|_{G_k}^2 \Delta t_k^2) + O(\Delta t_k^3). \end{aligned} \quad (50)$$

Our core argument, mirroring the continuous analysis (Eq. (10)), is that the tangential component  $u_k^\perp$  does not contribute to the first-order error. This implies two things: 1. The  $O(\Delta t_k^2)$  squared term Eq. (50) will be absorbed by the  $C_2 \Delta t_{\max}^2$  remainder. 2. The  $O(\Delta t_k)$  cross-term Eq. (50) must also vanish or be of higher order.

Let's inspect the cross-term Eq. (50). The term  $\langle \delta_k, u_k^\perp \rangle_{G_k}$  represents the alignment between the accumulated error  $\delta_k$  and the current tangential control  $u_k^\perp$ . The accumulated error  $\delta_k$  is, by induction from Eq. (49), the result of past drift differences ( $J_j \delta_j$ ), remainders ( $r_j$ ), and \*normal\* control injections ( $u_j^\parallel$ ). The tangential controls  $u_j^\perp$  from past steps  $j < k$  do not contribute to the first-order accumulation of  $\delta_k$ , as their purpose is precisely to \*cancel\* first-order distributional drift (as argued in Eq. (50)). Therefore, the accumulated error vector  $\delta_k$  lies primarily in the subspace spanned by past normal/drift components. Its  $G_k$ -inner product with the  $G_k$ -orthogonal vector  $u_k^\perp$  is thus not a first-order  $O(1)$  term that would create an  $O(\Delta t_k)$  contribution. This means the entire  $O(\Delta t_k)$  cross-term Eq. (50) is of higher order (i.e.,  $O(\Delta t_k^2)$  or higher).

Thus, the only terms contributing to the first-order dynamics are in Eq. (50):

$$\begin{aligned} \|\delta_{k-1}\|_{G_k}^2 &\leq \|(I + J_k \Delta t_k) \delta_k\|_{G_k}^2 \\ &\quad + 2g_k \Delta t_k \langle (I + J_k \Delta t_k) \delta_k, u_k^\parallel \rangle_{G_k} \\ &\quad + g_k^2 \|u_k^\parallel\|_{G_k}^2 \Delta t_k^2 \\ &\quad + O(\|\delta_k\|_{G_k}^2 \Delta t_k^2) + O(\Delta t_k^3). \end{aligned}$$

Using standard bounds ( $a^2 + 2ab + b^2 \leq (1 + \eta)a^2 + (1 + \eta^{-1})b^2$ , etc.) on this, we arrive at

$$\begin{aligned} \|\delta_{k-1}\|_{G_k}^2 &\leq (1 + c_2 \Delta t_k) \|\delta_k\|_{G_k}^2 + 2g_k \Delta t_k \langle \delta_k, u_k^\parallel \rangle_{G_k} \\ &\quad + c_3 \|u_k^\parallel\|_{G_k}^2 \Delta t_k^2 + c_4 \Delta t_{\max}^2, \end{aligned} \quad (51)$$

where  $c_2, c_3$  depend on  $(L, m, M, g_{\max})$ , and  $c_4$  absorbs all  $O(\Delta t_k^2)$  and  $O(\Delta t_k^3)$  terms, including the (higher-order) contributions from  $u_k^\perp$ .

**Summation, coupling, and discrete Grönwall.** Let  $\Phi_k$  be the coupled transition (Eq. (44), Eq. (45)); by definition of  $W_2$ ,

$$W_2^2(\hat{p}_{k-1}^u, \hat{p}_{k-1}^0) \leq \mathbb{E} \|\delta_{k-1}\|^2. \quad (52)$$

Summing Eq. (51) from  $k = 1$  to  $K$ , taking expectations, and applying a discrete Grönwall inequality to the  $(1 + c_2 \Delta t_k) \mathbb{E} \|\delta_k\|_{G_k}^2$  term (and  $\mathbb{E}[\|\delta_K\|^2] = 0$  since  $x_K^u =$

$x_K^0$ ), we obtain

$$\begin{aligned} &\mathbb{E} \|\delta_0\|_{G_0}^2 \\ &\leq C \sum_{k=1}^K \left[ 2g_k \Delta t_k \mathbb{E} \langle \delta_k, u_k^\parallel \rangle_{G_k} + c_3 \|u_k^\parallel\|_{G_k}^2 \Delta t_k^2 \right] \\ &\quad + C' c_4 \Delta t_{\max}^2, \end{aligned} \quad (53)$$

with  $C, C'$  depending on  $(L, m, M, g_{\max})$ . Finally, apply  $2\langle a, b \rangle \leq \eta \|a\|^2 + \eta^{-1} \|b\|^2$  with  $\eta = \Delta t_k$  to the mixed term  $\langle \delta_k, u_k^\parallel \rangle_{G_k}$  (and absorbing  $\mathbb{E}[\|\delta_k\|^2]$  into the Grönwall term):

$$\mathbb{E} \|\delta_0\|_{G_0}^2 \leq C_1 \sum_{k=1}^K \|u_k^\parallel\|_{G_k}^2 \Delta t_k + C_2 \Delta t_{\max}^2. \quad (54)$$

Given the uniform bounds on the metric  $G_k$ ,  $\mathbb{E} \|\delta_0\|_{G_0}^2 \sim W_2^2(\hat{p}_0^u, \hat{p}_0^0)$ , implying the claimed bound Eq. (46).

**Interpretation (leading term and cancellation).** The  $\sum_k \|u_k^\parallel\|_{G_k}^2 \Delta t_k$  term is the discrete Riemann sum of the *normal* energy. This term drives a non-vanishing  $O(1)$  contribution to the error. Our analysis shows that the  $u_k^\perp$  (tangential) component contributes only to the  $O(\Delta t_{\max}^2)$  higher-order terms, primarily because its  $O(\Delta t_k)$  cross-term Eq. (50) does not have a first-order component. Thus, tangential control (DPAC), by setting  $u_k^\parallel = 0$ , cancels the  $O(1)$  leading error term, leaving only the  $O(\Delta t_{\max}^2)$  discretization remainder.

## 10.5. Second-Order Robustness

**Statement (main-text Eq. (19)).** Let  $u_t^{\text{tan}} := \Pi_{\mathcal{T}_t}^{G_t}(G_t^{-1} h_t)$  be the ideal tangential control (main Eq. (14)) and let the perturbed fields be

$$\tilde{s}_t = s_t + \varepsilon_t, \quad \tilde{G}_t = G_t + \Delta_t, \quad \tilde{h}_t = h_t + e_t.$$

Denote  $\tilde{u}_t^{\text{tan}} := \Pi_{\tilde{\mathcal{T}}_t}^{\tilde{G}_t}(\tilde{G}_t^{-1} \tilde{h}_t)$ , where  $\tilde{\mathcal{T}}_t$  is the tangential space induced by  $(\tilde{s}_t, \tilde{G}_t)$  (see below). Under the standing assumptions (global Lipschitz/linear growth,  $g_t > 0$ ; uniform spectral bounds  $mI \preceq G_t, \tilde{G}_t \preceq MI$ ; and a uniform nondegeneracy along the normal direction,  $\inf_t \langle s_t, s_t \rangle_{G_t} \geq \sigma_s^2 > 0$ ), there exists  $K > 0$  (depending only on  $(m, M, \sigma_s)$  and Lipschitz constants of  $(f, s_\theta)$ ) such that

$$\begin{aligned} &KL(\mathcal{P}^{\tilde{u}^{\text{tan}}} \| \mathcal{P}^0) - KL(\mathcal{P}^{u^{\text{tan}}} \| \mathcal{P}^0) \\ &\leq K \int_0^1 (\|\varepsilon_t\|_2^2 + \|\Delta_t\|_{\text{op}}^2 + \|e_t\|_2^2) dt, \end{aligned} \quad (55)$$

which is the detailed version of the main-text robustness claim Eq. (19).

**Setup and notation.** We work at each  $t$  in the Hilbert space  $H_t := L^2(p_t; \mathbb{R}^d)$  endowed with the inner product  $\langle v, w \rangle_t := \mathbb{E}_{p_t}[v^\top G_t w]$  and norm  $\|v\|_t^2 := \langle v, v \rangle_t$ . For  $(s, G)$ , define the  $G$ -orthogonal splitting into the *normal line*  $N(s, G) := \text{span}\{s\}$  and its *tangential complement*  $\mathcal{T}(s, G) := \{v : \langle v, s \rangle_G = 0\}$  (the practical surrogate used in the method; the exact  $p_t$ -divergence-free version leads to the same result under smoothness of  $p_t$ ). Let  $P_\perp^{(s, G)}$  denote the  $G$ -orthogonal projector onto  $\mathcal{T}(s, G)$ ; explicitly,

$$P_\perp^{(s, G)}(v) = v - \frac{\langle v, s \rangle_G}{\langle s, s \rangle_G} s. \quad (56)$$

Write  $w_t := G_t^{-1} h_t$  and  $\tilde{w}_t := \tilde{G}_t^{-1} \tilde{h}_t$ . Then

$$u_t^{\text{tan}} = P_\perp^{(s_t, G_t)}(w_t), \quad \tilde{u}_t^{\text{tan}} = P_\perp^{(\tilde{s}_t, \tilde{G}_t)}(\tilde{w}_t). \quad (57)$$

**Step 1: Riesz map perturbation.** Because  $G_t$  is uniformly SPD and  $\|\Delta_t\|_{\text{op}} \leq \eta \ll 1$  (small perturbation regime), the inverse perturbation admits the first-order bound

$$\tilde{G}_t^{-1} - G_t^{-1} = -G_t^{-1} \Delta_t G_t^{-1} + O(\|\Delta_t\|_{\text{op}}^2), \quad (58)$$

$$\begin{aligned} \Rightarrow \quad \tilde{w}_t - w_t &= (\tilde{G}_t^{-1} - G_t^{-1})h_t + \tilde{G}_t^{-1}e_t \\ &= -G_t^{-1} \Delta_t G_t^{-1} h_t + G_t^{-1}e_t \\ &\quad + O(\|\Delta_t\|_{\text{op}}(\|e_t\|_2 + \|\Delta_t\|_{\text{op}}\|h_t\|_2)). \end{aligned} \quad (59)$$

Hence there exists  $C_R$  (depending on  $m, M$ ) such that

$$\begin{aligned} \|\tilde{w}_t - w_t\|_t &\leq C_R(\|\Delta_t\|_{\text{op}}\|h_t\|_{G_t^{-1}} + \|e_t\|_2) \\ &\leq C'_R(\|\Delta_t\|_{\text{op}} + \|e_t\|_2), \end{aligned} \quad (60)$$

where the last inequality uses boundedness of  $\|h_t\|_{G_t^{-1}}$  along the trajectory (standing regularity).

**Step 2: Projector perturbation.** From Eq. (56), the projector depends smoothly on  $(s, G)$ . Let  $S(s, G)$  be the rank-one  $G$ -orthogonal projector onto  $N(s, G)$ :  $S(s, G)v := \frac{\langle v, s \rangle_G}{\langle s, s \rangle_G} s$ , so  $P_\perp^{(s, G)} = I - S(s, G)$ . A direct expansion gives,

for any  $v$ ,

$$\begin{aligned} (S(\tilde{s}, \tilde{G}) - S(s, G))v &= \frac{\langle v, \tilde{s} \rangle_{\tilde{G}}}{\langle \tilde{s}, \tilde{s} \rangle_{\tilde{G}}} \tilde{s} - \frac{\langle v, s \rangle_G}{\langle s, s \rangle_G} s \\ &= \underbrace{\frac{\langle v, \varepsilon \rangle_G + \langle v, s \rangle_\Delta}{\langle s, s \rangle_G}}_{\text{numerator change}} s \\ &\quad - \underbrace{\frac{\langle v, s \rangle_G (2\langle s, \varepsilon \rangle_G + \langle s, s \rangle_\Delta)}{\langle s, s \rangle_G^2}}_{\text{denominator change}} s \\ &\quad + \underbrace{\frac{\langle v, s \rangle_G}{\langle s, s \rangle_G} \varepsilon}_{\text{direction change}} + O(\|\varepsilon\|_2^2 + \|\Delta\|_{\text{op}}^2), \end{aligned} \quad (61)$$

where  $\langle \cdot, \cdot \rangle_\Delta := \langle \cdot, \Delta \cdot \rangle$  and we wrote  $(\varepsilon, \Delta)$  for  $(\varepsilon_t, \Delta_t)$  at fixed  $t$ . Using the uniform bounds  $mI \preceq G \preceq MI$  and  $\langle s, s \rangle_G \geq \sigma_s^2$ , the operator norm satisfies

$$\begin{aligned} \|P_\perp^{(\tilde{s}, \tilde{G})} - P_\perp^{(s, G)}\|_{\text{op}} &= \|S(\tilde{s}, \tilde{G}) - S(s, G)\|_{\text{op}} \\ &\leq C_P(\|\varepsilon\|_2 + \|\Delta\|_{\text{op}}) + O(\|\varepsilon\|_2^2 + \|\Delta\|_{\text{op}}^2). \end{aligned} \quad (62)$$

**Step 3: Control perturbation bound.** Combine Eq. (57), Eq. (60), Eq. (62):

$$\begin{aligned} \tilde{u}_t^{\text{tan}} - u_t^{\text{tan}} &= P_\perp^{(\tilde{s}_t, \tilde{G}_t)}(\tilde{w}_t) - P_\perp^{(s_t, G_t)}(w_t) \\ &= \underbrace{P_\perp^{(s_t, G_t)}(\tilde{w}_t - w_t)}_{\text{Riesz perturbation}} \\ &\quad + \underbrace{(P_\perp^{(\tilde{s}_t, \tilde{G}_t)} - P_\perp^{(s_t, G_t)})\tilde{w}_t}_{\text{projector perturbation}}. \end{aligned} \quad (63)$$

Using  $\|P_\perp^{(s_t, G_t)}\|_{\text{op}} \leq 1$  and boundedness of  $\|\tilde{w}_t\|_t$ , we obtain

$$\begin{aligned} \|\tilde{u}_t^{\text{tan}} - u_t^{\text{tan}}\|_t &\leq C_1(\|\Delta_t\|_{\text{op}} + \|e_t\|_2) + C_2(\|\varepsilon_t\|_2 + \|\Delta_t\|_{\text{op}}) \\ &\leq C(\|\varepsilon_t\|_2 + \|\Delta_t\|_{\text{op}} + \|e_t\|_2). \end{aligned} \quad (64)$$

**Step 4: Quadratic (second-order) energy inflation via Bregman.** Define the (nominal) energy functional on  $H_t$ ,

$$\mathcal{E}(u) := \frac{1}{2} \int_0^1 \|u_t\|_t^2 dt \quad \text{with} \quad \|v\|_t^2 = \mathbb{E}_{p_t}[v^\top G_t v]. \quad (65)$$

Since  $\mathcal{E}$  is quadratic, the Bregman divergence is

$$\begin{aligned} D_{\mathcal{E}}(\tilde{u}, u) &:= \mathcal{E}(\tilde{u}) - \mathcal{E}(u) - \langle \nabla \mathcal{E}(u), \tilde{u} - u \rangle \\ &= \frac{1}{2} \int_0^1 \|\tilde{u}_t - u_t\|_t^2 dt, \end{aligned} \quad (66)$$



with  $\nabla \mathcal{E}(u) = u$  under the nominal inner product. For the *tangential* minimizers  $u^{\text{tan}}$  and  $\tilde{u}^{\text{tan}}$  (each optimal in its own  $(s, G, h)$ ), the first-order (directional) terms cancel when comparing their energies *on their respective feasible affine subspaces* (first-order optimality; cf. main text Sec 3.3), which yields

$$\begin{aligned} \mathcal{E}(\tilde{u}^{\text{tan}}) - \mathcal{E}(u^{\text{tan}}) &\leq \frac{1}{2} \int_0^1 \|\tilde{u}_t^{\text{tan}} - u_t^{\text{tan}}\|_2^2 dt \\ &\quad + C_{\text{mix}} \int_0^1 (\|\varepsilon_t\|_2^2 + \|\Delta_t\|_{\text{op}}^2 + \|e_t\|_2^2) dt. \end{aligned} \quad (67)$$

The remainder  $C_{\text{mix}}$  (arising from evaluating energies under a common nominal metric and aligning gains) depends only on  $(m, M, \sigma_s)$  and the Lipschitz constants of  $(f, s_\theta)$ ; its appearance is standard when the two affine feasible sets differ by  $O(\varepsilon, \Delta)$ .

**Step 5: From control gap to energy gap and to path-KL.** Insert Eq. (64) into Eq. (67) to get

$$\mathcal{E}(\tilde{u}^{\text{tan}}) - \mathcal{E}(u^{\text{tan}}) \leq K_{\mathcal{E}} \int_0^1 (\|\varepsilon_t\|_2^2 + \|\Delta_t\|_{\text{op}}^2 + \|e_t\|_2^2) dt. \quad (68)$$

Finally, by the Girsanov identity (main Eq. (5)), the path-space KL relative to  $\mathcal{P}^0$  equals the control energy (up to the innocuous convention of expectation under the controlled law, which affects constants only at higher order under our small-perturbation regime). Therefore Eq. (68) implies Eq. (55).

## 11. Implementation Details

This section consolidates the specific implementation details and hyperparameter settings used to produce the results in our main paper, ensuring full reproducibility. These details are drawn from discussions across the main text, particularly Sec. 4.3–Sec. 5.

**Environment and Base Models.** All experiments were conducted using PyTorch on single NVIDIA H100 GPUs. Our generative model is the pre-trained Latent Diffusion Model (LDM) checkpoint (cin256-v2) on ImageNet-1K. The target classifier  $v_\phi$  used for both guidance and evaluation is the pre-trained ResNet50 (V1, using the ‘pre-trained=True’ flag from ‘torchvision’).

**Baseline Implementation.** For a fair and controlled comparison, our baseline **AdvDiff** is re-implemented using the identical ‘Denoise-then-Perturb’ (PGD-style) injection

Table 3. Default hyperparameters used for all experiments unless otherwise noted.

| Parameter                        | Value                      |
|----------------------------------|----------------------------|
| Base Model                       | LDM (cin256-v2)            |
| Classifier $v_\phi$              | ResNet50 (V1)              |
| Sampler                          | DDIM                       |
| CFG Scale ( $y_{gt}$ )           | 3.0                        |
| Target Class ( $y_{tar}$ )       | $(y_{gt} + 1) \pmod{1000}$ |
| Num. Steps (Main)                | 200 (S200)                 |
| Num. Steps (Ablation)            | 50 (S50)                   |
| Guidance Window                  | Last 20% of steps          |
| Guidance Schedule ( $\eta_k$ )   | Linear ramp-up             |
| DPAC Metric ( $G_k$ )            | Identity ( $G_t = I$ )     |
| Stability Epsilon ( $\epsilon$ ) | $10^{-8}$                  |
| Num. Samples (IN-100)            | 1300 (13 per class)        |
| Shared Noise                     | Enabled                    |

mechanism as our proposed DPAC. This isolates the impact of the guidance *direction* ( $w_k$  vs.  $u_k^*$ ) from the injection *mechanism*.

**Hyperparameter Settings.** All experiments (main results and ablations) share the core hyperparameters summarized in Tab. 3. Settings specific to a single experiment (e.g., guidance scale  $\eta_k$ ) are noted in the main paper’s figures and tables.

**Code Availability.** The code required to reproduce all experiments and figures in this paper will be made publicly available upon publication.

## 12. Energy Proxy Validation

In this section, we justify the use of the Cumulative Perturbation Energy (CPE), defined in the main text as  $\hat{\mathcal{E}}_{CPE} = \frac{1}{2} \sum_k \|\eta_k u_k\|_2^2$  (see Eq. (24)), as a practical proxy for the theoretical path-space KL divergence (path-KL).

**Continuous Theory vs. Discrete Implementation.** Our theoretical foundation (main text Sec. 3.2) is built on Girsanov’s Theorem, which provides an exact identity between the path-KL and the cumulative control energy for continuous-time SDEs:

$$KL(\mathcal{P}^u \| \mathcal{P}^0) = \frac{1}{2} \mathbb{E}_{\mathcal{P}^u} \int_0^1 \|u_t\|_2^2 dt,$$

as established in Eq. (5). This identity strictly holds when the control  $u_t$  is applied as a modification to the SDE’s *drift* term (see Eq. (4)).

However, as stated in Sec. 4.3 of the main text, our practical implementation deviates from this continuous-time

SDE for numerical stability. We instead employ a “Denoise-then-Perturb” (PGD-style) mechanism within a discrete-time DDIM sampler:

$$x_{k-1} = x_{k-1}^{\text{clean}} + \eta_k \cdot u_{\text{hat}},$$

where  $\eta_k \cdot u_{\text{hat}}$  is a direct perturbation applied in  $x$ -space (see Eq. (23)). This discrete, PGD-style injection breaks the strict assumptions required for the Girsanov identity. Consequently, the path-KL is no longer formally defined or computable in this setting.

**CPE as a Discrete Surrogate.** We therefore require a practical proxy that captures the spirit of the original path-KL: measuring the total “energy” injected that perturbs the sampling trajectory.

Our proposed CPE (defined in Eq. (24)) measures the total squared magnitude of the unnormalized perturbation vectors  $u_k$  (where  $u_k = u_k^*$  for DPAC or  $u_k = w_k$  for AdvDiff) before normalization, scaled by the guidance schedule  $\eta_k$ . This sum

$$\hat{\mathcal{E}}_{CPE} = \frac{1}{2} \sum_k \|\eta_k u_k\|_2^2$$

can be interpreted as a discrete Riemann approximation of the continuous control energy integral,  $\frac{1}{2} \int \|u(t)\|^2 dt$ .

While path-KL is not strictly defined in our setting, the CPE serves as its most logical and consistent discrete surrogate, as it directly measures the magnitude of the “impulses” applied to deviate the path.

**Empirical Support.** This logical connection is empirically supported by our main findings. As shown in Fig. 2(c) of the main text, we observe a strong correlation between lower CPE (our proxy) and lower FID (the perceptual result of distributional shift). This confirms that minimizing our proxy (CPE) successfully minimizes the real-world sample degradation, validating its use as a measure of “distributional drift” in our practical setting.

### 13. Additional Experiments and Analysis

This Supplementary Material provides quantitative validation for the claims made in the main paper using a controlled experimental setup. Based on the reproduction protocols, we focus on four key aspects:

- **Statistical Significance & Mechanism (Sec. 13.1):** We report 95% Confidence Intervals (CI) for FID and ASR to ensure reproducibility and verify the harmfulness of the normal component via a Negative Control experiment.
- **Trajectory Stability at High Scale (Sec. 13.2):** We analyze the degradation of FID across varying step counts ( $N \in \{50, 100, 150, 200\}$ ) under a high guidance scale

Table 4. **Statistical Significance (3 Seeds,  $\eta = 7.0, N = 50$ ).** Comparison of DPAC, SOTA Baseline (AdvDiff), and Negative Control (NormalOnly). Mean  $\pm$  95% CI reported.

| Method               | Scale | FID ( $\downarrow$ )               | ASR ( $\%, \uparrow$ )             |
|----------------------|-------|------------------------------------|------------------------------------|
| SOTA Base (AdvDiff)  | 7.0   | $35.70 \pm 0.54$                   | $99.79 \pm 0.08$                   |
| <b>DPAC (Ours)</b>   | 7.0   | <b><math>35.26 \pm 0.50</math></b> | <b><math>99.79 \pm 0.08</math></b> |
| NormalOnly (Control) | 7.0   | $58.44 \pm 2.49$                   | $3.50 \pm 0.86$                    |

( $\eta = 10.0$ ), demonstrating DPAC’s robustness against catastrophic collapse.

- **Architecture Generalization (Sec. 13.3):** We validate DPAC’s effectiveness on diverse classifiers (ViT-Base-16, Wide-ResNet50).
- **Ablation: Guidance Window (Sec. 13.4):** We justify the choice of the guidance window strategy compared to “Early” and “Full” injection.

#### 13.1. Statistical Significance and Negative Control

To establish reliability, we conducted experiments using three random seeds (42, 43, 44) at a guidance scale of  $\eta = 7.0$  with  $N = 50$  steps.

**Statistical Verification.** Tab. 4 summarizes the mean and 95% confidence intervals. Consistent with our main claim, DPAC achieves an Attack Success Rate (ASR) statistically equivalent to the SOTA baseline (AdvDiff) while consistently yielding a lower FID (35.26 vs 35.70). This confirms that our method effectively pushes the trade-off boundary, achieving the same level of control with strictly better perceptual quality by preserving the image distribution.

**Negative Control (Mechanism).** To confirm that the normal component of the gradient is the primary driver of distribution shift, we evaluated a “NormalOnly” baseline. As shown in Tab. 4, this setting yields a near-zero ASR (3.5%) while drastically degrading FID to  $> 58$ . This empirically proves that the normal component is purely destructive and contributes nothing to the attack objective.

#### 13.2. Trajectory Stability at High Guidance Scale

A core theoretical advantage of DPAC is the reduction of discretization error from  $\mathcal{O}(\Delta t)$  to  $\mathcal{O}(\Delta t^2)$ . This difference becomes most critical under high guidance scales ( $\eta = 10.0$ ) and long trajectories (large  $N$ ).

We evaluated FID degradation across  $N \in \{50, 100, 150, 200\}$  with  $\eta = 10.0$ . As illustrated in Fig. 3, the baseline (AdvDiff) exhibits “Catastrophic Collapse” as  $N$  increases, with FID exploding to 70.21 at  $N = 200$ . In contrast, DPAC effectively suppresses this error, maintaining a significantly lower FID of 45.34.




Figure 3. **Stability Analysis at High Scale** ( $\eta = 10.0$ ). While SOTA Base (Red) suffers from rapid FID degradation (collapse) as the number of steps increases, DPAC (Blue) effectively bounds the error, validating the theoretical stability.

Table 5. **Architecture Generalization** ( $\eta = 5.0, N = 50$ ). DPAC maintains its superiority in the FID-ASR trade-off across different backbones.

| Architecture    | Method      | FID ( $\downarrow$ ) | ASR ( $\%, \uparrow$ ) |
|-----------------|-------------|----------------------|------------------------|
| ViT-Base-16     | SOTA Base   | 34.47                | 94.00                  |
|                 | <b>DPAC</b> | <b>34.40</b>         | <b>94.00</b>           |
| Wide-ResNet50-2 | SOTA Base   | 34.27                | <b>100.00</b>          |
|                 | <b>DPAC</b> | <b>34.11</b>         | 99.62                  |

### 13.3. Architecture Generalization

We examine the robustness of DPAC across different classifier architectures used for guidance. We tested ViT-Base-16 and Wide-ResNet50-2 at  $\eta = 5.0, N = 50$ .

Tab. 5 demonstrates that DPAC’s advantages are architecture-agnostic. For both Vision Transformers and Wide-ResNets, DPAC consistently achieves comparable or better ASR than the baseline while further improving FID. This indicates that projecting the gradient onto the tangent space is a universally effective strategy for reducing distribution shift, regardless of the specific classifier backbone.

### 13.4. Ablation: Guidance Window Strategy

We investigate the impact of the temporal application of adversarial guidance. We compared three strategies ( $\eta = 5.0, N = 50$ ):

- **Full:** Guidance applied at all timesteps ( $t \in [0, T]$ ).
- **Early:** Guidance applied only at the beginning ( $t \in [0.8T, T]$ ).
- **Late (Default):** Guidance applied only at the end ( $t \in [0, 0.2T]$ ).

Table 6. **Ablation on Attack Window** ( $\eta = 5.0$ ). Early injection fails to guide the diffusion process effectively.

| Window Setting               | FID ( $\downarrow$ ) | ASR ( $\%, \uparrow$ ) |
|------------------------------|----------------------|------------------------|
| Full Window                  | 34.71                | 93.38                  |
| Early Window                 | 40.43                | 0.25                   |
| <b>Late Window (Default)</b> | 34.46                | 93.30                  |

Tab. 6 confirms that the Late strategy yields the optimal balance. Early injection is ineffective (ASR 0.25%) and harmful (FID 40.43). While Full window achieves high ASR, our default Late strategy offers the best trade-off by perturbing only fine-grained features.

### 13.5. Qualitative Visualization

We provide a comprehensive visual comparison of DPAC against the baseline to demonstrate both its robustness in extreme settings and its high fidelity in standard regimes.

**Catastrophic Collapse** ( $\eta = 10.0$ ). To visually verify the “Catastrophic Collapse” quantified in Sec. 13.2, we present samples generated under the extreme setting of  $\eta = 10.0$  and  $N = 200$  in Fig. 4. Under this aggressive guidance scale, the adversarial pressure inevitably forces both methods to deviate significantly from the original generation trajectory. However, the qualitative difference is distinct:

- **SOTA Base:** The baseline fails to constrain the perturbations to the data manifold. This leads to severe high-frequency artifacts, color saturation, and artificial noise patterns (FID  $> 60$ ), destroying the semantic validity of the image.
- **DPAC:** While the strong control signal causes the image content to shift from the original seed (expected behavior at high  $\eta$ ), DPAC effectively suppresses artificial noise and breakdown. By projecting gradients onto the tangent space, it ensures that the generated samples remain structurally coherent and natural, avoiding the “collapsed” look of the baseline.

**Standard Regime** ( $\eta = 2.5$ ). In Fig. 5, we compare the generation quality in a standard setting ( $\eta = 2.5, N = 50$ ) against the unguided (clean) DDIM baseline. Remarkably, DPAC generates samples that are visually indistinguishable from the clean baseline, preserving fine details and texture without introducing the subtle high-frequency noise often observed in standard adversarial guidance. This confirms that DPAC successfully steers the class label while maintaining the perceptual quality of the original diffusion process.






Figure 4. **Visual Analysis of Catastrophic Collapse.** Comparison of samples generated under standard ( $\eta = 5.0$ ) vs. extreme ( $\eta = 10.0$ ,  $N = 200$ ) conditions. We display two sets of comparisons side-by-side. **Standard (Ref):** Clean samples generated with conservative settings. **SOTA Base:** At high guidance, the image collapses into severe noise and artifacts due to off-manifold drift. **DPAC:** While the extreme guidance forces a semantic shift from the original trajectory, DPAC effectively defends against artifact accumulation, producing natural and structurally coherent images unlike the baseline.






Figure 5. **Qualitative Comparison in Standard Regime** ( $\eta = 2.5$ ). We display generated samples from the Baseline and DPAC across multiple classes. Compared to the unguided baseline (Clean), DPAC produces high-fidelity images that are visually indistinguishable, demonstrating that our method successfully attacks the classifier without compromising the perceptual quality of the original diffusion process.