CACFG: CURVATURE-AWARE CLASSIFIER-FREE GUIDANCE AND OPTIMAL CONTROL

Anonymous authors
Paper under double-blind review

ABSTRACT

Diffusion models generate samples by learning to reverse a fixed corruption process, and classifier-free guidance (CFG) is the standard mechanism for conditioning this process on a desired class or prompt. CFG can be applied at varying guidance strengths, and while higher strengths improve image quality and conditional alignment, too high a guidance strength can degrade image quality and diversity. Furthermore, CFG violates principled diffusion sampling dynamics, and existing explanations for why it works despite the violation disagree on the underlying theory or do not extend to deterministic samplers used in practice. We address both these issues. We first frame CFG sampling as a continuous-time optimal control problem, treating the sampling trajectory as a sequence of controls chosen to maximise the probability of the desired condition. Solving the resulting Hamilton–Jacobi–Bellman equation shows that CFG is recovered under specific path costs when using an unconstrained control set. We argue this lack of constraint is responsible for CFG's failure at high guidance strengths, since it permits the sampling path to move arbitrarily far from the current image estimate. To fix this, we propose curvature-aware CFG (CACFG), which constrains the control set to a hypersphere informed by the Gaussian regularisation used when training variational autoencoders. We show that the control inputs produced by CFG sampling routinely violate this bound, and that across diffusion models, datasets, and guidance schedules, CACFG achieves superior generative quality at mid-to-high guidance strengths with a less severe quality-diversity tradeoff than regular CFG.

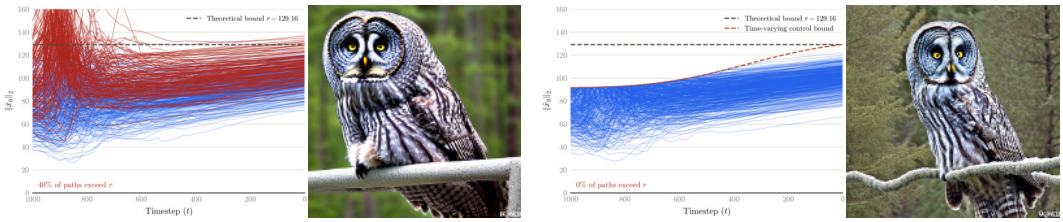
1 INTRODUCTION

Diffusion models generate samples by learning to iteratively reverse a fixed process that gradually corrupts data with noise (Ho et al., 2020; Song et al., 2021b). *Classifier-free guidance (CFG)* (Ho & Salimans, 2022) provides a mechanism to control the content of generated images by using a variable guidance strength parameter to target a tilted conditional distribution. However, this tilted distribution is not the marginal of any true diffusion process, and excessive CFG guidance can reduce image diversity and degrade image quality by introducing visual artefacts (Sadat et al., 2025; Karras et al., 2024). Despite this, CFG is frequently used in practice, and practitioners select the largest guidance strengths that do not harm image quality (Dhariwal & Nichol, 2021; Ho & Salimans, 2022).

Previous work has proposed methods of fixing CFG artefacts: *CFG++* (Chung et al., 2025) and *Interval-CFG (IntCFG)* (Kynkäänniemi et al., 2024) each attribute CFG's artefacts to the sampling trajectory departing the image manifold, correcting for this by varying guidance strength across the trajectory. A separate line of work asks why CFG works *despite* violating diffusion sampling dynamics. Hypotheses attribute its effectiveness to a shared bias in the conditional and unconditional denoisers (Karras et al., 2024), an implicit correction for the poorly-approximated tilted distribution (Moufad et al., 2025), or an equivalence with a predictor-corrector scheme in the continuous-time SDE limit (Bradley & Nakkiran, 2025). However, these explanations disagree on the mechanism responsible, and they do not rigorously extend to commonly used deterministic samplers.

We address this gap by reformulating CFG sampling as a continuous-time optimal control problem: rather than fixing the sampling trajectory as a function of stochastic transition kernels, we treat it as

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.



(a) CFG (b) CACFG

Figure 1: Tweedie's magnitude ($\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2$) of diffusion sampling trajectories across time with $\omega = 7.5$ and using (a) CFG (Ho & Salimans, 2022) and (b) CACFG (ours), alongside a representative sample generated by each method. Note that 40% of CFG sampling paths violate the theoretical 95% upper magnitude bound. CACFG restricts permissible controls to lie within a time-varying bounded control set.

a sequence of controls chosen to maximise the terminal reward $p(\xi \mid \boldsymbol{x}_0)$, i.e., how probable it is that the sample is of the desired condition/class ξ. Solving the resulting *Hamilton–Jacobi–Bellman (HJB) equation* shows that the optimal control recovers CFG when the control set is unconstrained and the traversal penalty is expressed in the same scale as the diffusion training loss. To the best of our knowledge, we are the first to frame CFG as a deterministic optimal control problem in this way. Additionally, existing guidance scheduling methods such as *CFG++* and *IntCFG* also emerge from this same optimal control formulation under alternative choices of the traversal penalty (Appendix B). By examining the path costs of this optimisation problem, we find that the CFG diffusion sampling trajectory is encouraged to: align the conditional and CFG scores at early-to-mid sampling times; converge towards a local maximum at late sampling times; and avoid traversing near a maximum too early. This provides significant insight into how CFG works, and why it exhibits certain pathologies.

Our optimal control formulation shows that traditional CFG imposes *no* constraint on the control set. This permits the sampling path to move arbitrarily far from the current image estimate given sufficiently low traversal penalty. We argue this is directly responsible for CFG's failure modes in high-guidance regimes. To address this, we propose a constrained control set and corresponding curvature-aware CFG scheme (CACFG) that relies on a property latent diffusion models have by construction: the training objective of their variational autoencoder (VAE) is explicitly regularised towards an isotropic Gaussian. We hypothesise that this forces most encodings to lie within a hypersphere of an estimable radius, and we show empirically that the Tweedie estimates produced by traditional CFG sampling frequently violate this theoretical upper bound, while those of CACFG do not (Figure 1). We find that across diffusion models, datasets, and guidance schedules, CACFG has superior generative quality at mid-to-high guidance strengths as measured by various quality metrics.

Our core contributions are as follows: (i) We show that diffusion guidance can be interpreted as a continuous-time control problem; (ii) We show that classifier-free guidance (CFG) is one solution to this control problem, and that existing guidance scheduling methods, such as CFG++ (Chung et al., 2025) and IntCFG (Kynkäänniemi et al., 2024), arise from the same framework under alternative choices of the traversal penalty; (iii) We propose a method of using the Gaussian regularisation term of the VAE training loss to infer the structure of the embedded image manifold; (iv) We propose a constrained control set for the continuous-time optimisation problem that results in a diffusion guidance (CACFG) that provides superior conditional generation with a less severe quality-diversity tradeoff than unconstrained guidance methods, at negligible additional computational cost.

2 Preliminaries

Denoising diffusion models generate data by learning to reverse a fixed diffusion (noising) process. We write $\boldsymbol{x}_t \in \mathbb{R}^m$ for $t \in [0, T]$ to denote the state acted on by this process. The diffusion itself is driven by the stochastic differential equation (SDE) (Ho et al., 2020; Song et al., 2021b),

$$d\boldsymbol{x}_t = f(t)\boldsymbol{x}_t\, dt + g(t)\, d\boldsymbol{B}_t, \quad \forall\, t \in [0, T], \tag{1}$$

where $f(t), g(t) : \mathbb{R} \to \mathbb{R}$ are continuous and $\cdot d\boldsymbol{B}_t$ denotes an Itô integral with respect to the m-dimensional Brownian motion process, $\boldsymbol{B}_t$. The form of f and g is chosen such that (1) gradually erases structure in $\boldsymbol{x}_0$ as $t \to T$, leaving a state that is approximately isotropic Gaussian at termination,

$p(\boldsymbol{x}_T) \approx \mathcal{N}(\mathbf{0}, \sigma_T^2 I)$. Modern image diffusion models leverage both unconditional and conditional modes, where the initial measure of the state X_0 is conditioned on a measurable event ξ of an auxiliary variable Y that takes values in Ω_Y, i.e. $P(\boldsymbol{X}_0|\xi)$. This allows for $\xi = \Omega_Y$ to denote unconditional generation, and $\xi \neq \Omega_Y$ to describe conditional generation with target ξ. We formally define the above diffusion as the *canonical image diffusion* in Definition E.1, and in Theorem E.1 we show that it admits the transition kernel $p(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathcal{N}\left(\alpha_t \boldsymbol{x}_0,\ \beta_t^2 I^{(m\times m)}\right)$ with $\alpha_t, \beta_t : \mathbb{R} \to \mathbb{R}$ and where $I^{(m\times m)}$ denotes the m-dimensional identity.

Generating a sample amounts to reversing (1), which can be done stochastically via Anderson's reverse-time SDE (Anderson, 1982), or (more commonly) deterministically via the probability-flow ODE (Song et al., 2021b),

$$d\boldsymbol{x}_t/dt = f(t)\boldsymbol{x}_t - g(t)^2 \nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t \mid \xi)/2, \tag{2}$$

both of which require an estimate of the score, $\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t \mid \xi)$. This is obtained by training a denoiser $\boldsymbol{\epsilon}_\theta : \mathbb{R}^m \times \mathbb{R} \times \mathcal{Y} \to \mathbb{R}^m$ to recover $\boldsymbol{x}_0$ from a noised observation using the objective,

$$\min_\theta\ \mathbb{E}_{\boldsymbol{x}_0 \sim P(\boldsymbol{x}_0|\xi),\ \boldsymbol{x}_t \sim p(\boldsymbol{x}_t|\boldsymbol{x}_0),\ t\sim p(t)} \left[\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0\|_2^2\right], \quad \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi; \theta) \triangleq \frac{\boldsymbol{x}_t - \beta_t \boldsymbol{\epsilon}_\theta(t, \boldsymbol{x}_t, \xi)}{\alpha_t}. \tag{3}$$

For a minimiser $\theta^\star$ of (3), the score is proportional to the denoiser (see Theorem E.2),

$$\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t \mid \xi) = -\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi)/\beta_t, \tag{4}$$

so either reverse-process formulation (SDE or ODE) can be run with $\boldsymbol{\epsilon}_{\theta^\star}$ in place of the score. Additionally, by *Tweedie's formula* (Efron, 2011) the optimal denoiser's output is equal to the posterior mean (see Theorem E.3),

$$\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi) = (\boldsymbol{x}_t - \beta_t \boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi))/\alpha_t = \mathbb{E}[\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi], \tag{5}$$

and note that we are taking $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi) \triangleq \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi; \theta^\star)$ for brevity.

Classifier-free guidance (CFG). In many cases the conditional score alone is not sufficient to reconstruct high-quality images of the desired class. Although $\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi)$ recovers $\nabla_{\boldsymbol{x}_t} \log p(\boldsymbol{x}_t \mid \xi)$ exactly at optimality (Theorem E.2), faithfully sampling $p(\boldsymbol{x}_0 \mid \xi)$ is not the same as adhering strongly to ξ: a specific conditioning event is typically supported by a small number of training samples compared to the entire training dataset. Thus, the model constructs a noisy estimate of $P(\boldsymbol{X}_0 \mid \xi)$, which places significant probability mass over samples weakly consistent with ξ (Dhariwal & Nichol, 2021; Ho & Salimans, 2022).

To strengthen class adherence, guidance methods target a tilted distribution in place of $p(\boldsymbol{x}_t \mid \xi)$,

$$\bar{p}(\boldsymbol{x}_t \mid \xi) \propto p(\xi \mid \boldsymbol{x}_t)^\omega\, p(\boldsymbol{x}_t), \qquad \omega \in \mathbb{R}_{>0}. \tag{6}$$

Increasing ω concentrates probability mass around these maximal-density regions, whilst leaving its maximisers unchanged. Differentiating the log of (6) and substituting (4) at $\xi = \Omega_Y$ gives

$$\nabla_{\boldsymbol{x}_t} \log \bar{p}(\boldsymbol{x}_t \mid \xi) = \omega \nabla_{\boldsymbol{x}_t} \log p(\xi \mid \boldsymbol{x}_t) - \boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \Omega_Y)/\beta_t \tag{7}$$

$$= -\Big[\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \Omega_Y) + \omega\big(\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi) - \boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \Omega_Y)\big)\Big]/\beta_t, \tag{8}$$

where the second line eliminates $\nabla_{\boldsymbol{x}_t} \log p(\xi \mid \boldsymbol{x}_t)$ via Bayes' rule and the score identity in (4). *Classifier guidance* (Dhariwal & Nichol, 2021) evaluates (7) by training a separate model $p_\theta(\xi \mid \boldsymbol{x}_t)$; *classifier-free guidance* (Ho & Salimans, 2022) instead uses (8), interpolating the class-conditional and unconditional noise models defined above. Note that this guidance does not represent a true diffusion process: $\bar{p}(\boldsymbol{x}_t \mid \xi)$ is not the marginal induced by diffusing $\boldsymbol{x}_0 \sim \bar{P}(\boldsymbol{X}_0 \mid \xi)$, so following this score abandons the reverse-time SDE/flow ODE dynamics of (2) (Karras et al., 2024); however, it empirically improves perceptual sample quality (Dhariwal & Nichol, 2021; Ho & Salimans, 2022).

Latent diffusion models. Let $\boldsymbol{Z} \in \mathbb{R}^n$ be an $\mathbb{R}^n$-valued random variable representing a natural image in n-dimensional pixel-space, sharing the auxiliary (label) variable Y defined above (and in Definition E.1). As above, for $\xi \in \mathcal{Y}$ we write $P(\boldsymbol{Z} \mid \xi)$ for the conditional law, and $\boldsymbol{z} \sim P(\boldsymbol{Z} \mid \xi)$ for a realised image. Taking $\boldsymbol{X}_0 \triangleq \boldsymbol{Z}$ (so $m = n$) equates the diffusion state with the image directly, giving $P(\boldsymbol{X}_0 \mid \xi) = P(\boldsymbol{Z} \mid \xi)$ and recovering ordinary *pixel-space diffusion*.

Latent diffusion models (Rombach et al., 2022) instead diffuse over a compressed representation of $\boldsymbol{z}$, obtained from a variational autoencoder (VAE) (Kingma & Welling, 2014). The encoder maps

a realised image to a distribution of latents, $q_\phi(\boldsymbol{x}_0 \mid \boldsymbol{z}) \triangleq \mathcal{N}\left(\boldsymbol{x}_0;\, \boldsymbol{\mu}_\phi(\boldsymbol{z}),\, \mathrm{diag}\big(\boldsymbol{\sigma}_\phi(\boldsymbol{z})^2\big)\right)$, and the decoder gives the reconstruction likelihood, $p_\psi(\boldsymbol{z} \mid \boldsymbol{x}_0) \triangleq \mathcal{N}\left(\boldsymbol{z};\, \boldsymbol{\nu}_\psi(\boldsymbol{x}_0),\, \mathrm{diag}\big(\boldsymbol{\gamma}_\psi(\boldsymbol{x}_0)^2\big)\right)$. The density of $\boldsymbol{X}_0 \mid \xi$ can then be obtained via $p(\boldsymbol{x}_0 \mid \xi) \triangleq \int_{\mathbb{R}^n} q_\phi(\boldsymbol{x}_0 \mid \boldsymbol{z})\, dP(\boldsymbol{z} \mid \xi)$.

The encoder and decoder are trained jointly, over parameters ϕ, ψ respectively, to minimise the negative evidence lower bound

$$\mathcal{L}_{\text{VAE}}(\phi, \psi) \triangleq \mathbb{E}_{\boldsymbol{z} \sim P(\boldsymbol{Z})}\Big[\underbrace{-\mathbb{E}_{q_\phi(\boldsymbol{x}_0 \mid \boldsymbol{z})}[\log p_\psi(\boldsymbol{z} \mid \boldsymbol{x}_0)]}_{\text{reconstruction}} + \underbrace{D_{\text{KL}}\Big(q_\phi(\boldsymbol{x}_0 \mid \boldsymbol{z}) \,\Big\|\, \mathcal{N}(\boldsymbol{0}, I^{(m \times m)})\Big)}_{\text{regularisation}} \Big], \tag{9}$$

following Kingma & Welling (2014), where $P(\boldsymbol{Z}) \triangleq P(\boldsymbol{Z} \mid \Omega_Y)$ is the unconditional image law. Note that the KL term regularises the encoder distribution towards the same isotropic standard normal that the forward process of (1) itself converges to as $t \to T$–a fact we leverage later.

3 Related Work

Chung et al. (2025) attributes CFG artefacts such as colour oversaturation and broken invertibility to an off-manifold effect: since $\omega > 1$ extrapolates beyond the conditional and unconditional denoiser outputs, the resulting noise estimate can fall off the data manifold. By reformulating guidance as a manifold-constrained score-matching problem, Chung et al. (2025) appeal to the linear manifold hypothesis (Chung et al., 2022) to compute a *manifold constrained gradient* to force the sample to remain on the theorised linear image manifold. Alternatively, *Interval-CFG (IntCFG)* (Kynkäänniemi et al., 2024) hypothesises that CFG artefacts can be attributed to *when* guidance is applied: guidance at high noise levels collapses diversity by prematurely committing the trajectory to a small set of modes, while guidance at low noise levels has minimal effect. Thus, IntCFG disables guidance outside an interval of intermediate noise levels. Though motivated differently, both methods let the guidance scale of (8) vary across the sampling trajectory, a fact we make concrete in Appendix B. Alternatively, techniques such as *Adaptive Projected Guidance (APG)* (Sadat et al., 2025) reshape the CFG update itself. However, none of this work provides a rigorous explanation as to why CFG is able to generate conditional images, *despite* violating diffusion sampling dynamics.

Karras et al. (2024) attributes CFG's effectiveness to a shared model bias: a capacity-limited denoiser over-covers sparsely populated regions of the data manifold, and because the unconditional model exhibits this bias more severely than the conditional one, their ratio acts as an attractive force that pulls samples towards the image manifold. Moufad et al. (2025) show that the tilted distribution of (6) can be recovered by augmenting CFG with a corrective Rényi-divergence term, and they argue that CFG is just poorly approximating the true tilted distribution. Bradley & Nakkiran (2025) instead prove that in the continuous-time SDE limit, the stochastic (DDPM) variant (Ho et al., 2020) of CFG is identical to a predictor-corrector scheme that pairs DDIM (Song et al., 2021a) denoising with a Langevin-dynamics corrector on the tilted distribution. However, this does not account for deterministic (flow ODE) variants of CFG, and the explanations of Karras et al. (2024) and Moufad et al. (2025) disagree on the key mechanisms.

4 Reframing CFG as an Optimal Control

We first consider the tilted density of (6), and note that raising $p(\xi \mid \boldsymbol{x}_t)$ to any power $\omega > 0$ leaves its maximiser (w.r.t. $\boldsymbol{x}_t$) unchanged. Thus, the CFG update of (8) can be viewed as naively moving towards a local maximum of $p(\xi \mid \boldsymbol{x}_t)$ at each timestep; however, this does not guarantee maximisation with respect to the terminal state, $\boldsymbol{x}_0$. Treating CFG sampling as purely an optimisation problem leads to the natural question: *does CFG sampling construct a path that maximises* $p(\xi \mid \boldsymbol{x}_0)$ *subject to some reasonable path constraints?* More concretely, we wish to understand whether CFG constructs trajectories that terminate at a state, $\boldsymbol{x}_0$, that *maximise* $p(\xi \mid \boldsymbol{x}_0)$ subject to interpretable path costs that punish/reward movement from the initial position, $\boldsymbol{x}_T$.

We treat the sequence of denoising targets used along the trajectory as free decisions rather than following a fixed set of transition kernels, and we introduce a time-dependent control variable $\boldsymbol{u}_t \in U(t) \subseteq \mathbb{R}^m$ taking values in a nonempty, closed, and *convex* time-varying control set $U(t)$. Additionally, note that maximising $p(\xi|\boldsymbol{x}_0)$ w.r.t. $\boldsymbol{x}_0$ is equivalent to maximising $p(\boldsymbol{x}_0 \mid \xi)/p(\boldsymbol{x}_0)$

via Bayes' rule. Under this framework, diffusion guidance is the solution to

$$\max_{\boldsymbol{u}_{[t_l,T]}} \left\{ \int_{t_l}^{T} L(s, \boldsymbol{x}_s, \boldsymbol{u}_s)\, ds + \phi(t_l, \boldsymbol{x}_{t_l}) \right\}, \qquad \dot{\boldsymbol{x}}_s = \boldsymbol{b}(s, \boldsymbol{x}_s, \boldsymbol{u}_s), \tag{10}$$

with running reward $L : \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$, dynamics constraint $\boldsymbol{b} : \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}^m$, and terminal reward $\phi(t_l, \boldsymbol{x}_{t_l})$, where $\phi(s, \boldsymbol{x}_s) \triangleq \log\left(\frac{p(\boldsymbol{x}_s|\xi)}{p(\boldsymbol{x}_s)}\right)\ \forall\, s \in [t_l, T]$. (10) defines a reverse-time, continuous-time maximisation problem, and we write

$$\hat{V}(t, \boldsymbol{x}_t) \triangleq \sup_{\boldsymbol{u}_{[t_l,t]}} \left\{ \int_{t_l}^{t} L(s, \boldsymbol{x}_s, \boldsymbol{u}_s)\, ds + \phi(t_l, \boldsymbol{x}_{t_l}) \right\}, \tag{11}$$

to define the best accumulated reward achievable by time t. Bellman's principle of optimality (Bellman, 1957) shows that $\hat{V}$ must satisfy the reverse-time Hamilton–Jacobi–Bellman (HJB) equation[1]

$$\hat{V}_t(t, \boldsymbol{x}_t) = \sup_{\boldsymbol{u}_t \in U(t)} \left\{ L(t, \boldsymbol{x}_t, \boldsymbol{u}_t) - \left\langle \hat{V}_{\boldsymbol{x}_t}(t, \boldsymbol{x}_t),\, \boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{u}_t) \right\rangle \right\}, \quad \forall\, t \in (t_l, T], \tag{12}$$

with boundary condition $\hat{V}(t_l, \boldsymbol{x}_{t_l}) = \phi(t_l, \boldsymbol{x}_{t_l})$. Throughout, we take the termination time $t_l \in (0, T)$, which keeps $\beta_s^2 > 0$ (and hence the relevant densities well-defined) over the whole horizon. Additionally, we require that L, $\hat{V}$, and $\boldsymbol{b}$ be regular[2] and C^1 differentiable in t, $\boldsymbol{x}_t$, and $\boldsymbol{u}_t$.

We now choose L and $\boldsymbol{b}$. For the running reward, we penalise the divergence of the control $\boldsymbol{u}_s$ from the *unconditional* sample, multiplied by some time-dependent cost scale $r(s) > 0$, alongside an additional, as-yet-unspecified term h capturing further path costs/rewards,

$$L(s, \boldsymbol{x}_s, \boldsymbol{u}_s) \triangleq -r(s)\, \|\boldsymbol{u}_s - \hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)\|_2^2 + h(s, \boldsymbol{x}_s, \boldsymbol{u}_s). \tag{13}$$

The choice of $\hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)$ as the anchor for the ℓ_2 penalty is deliberate: by Tweedie's formula (5), it is the unconditional posterior mean of $\boldsymbol{x}_s$, i.e. the most plausible denoising target available at that noise level. Penalising the ℓ_2 deviation of $\boldsymbol{u}_s$ from this quantity therefore encodes a natural prior: a candidate target should only be preferred over the unconditional estimate if there is sufficient reward for doing so. Framing guidance this way turns the question of *why* CFG works into a question about h and $\boldsymbol{b}$: is there a choice of h and $\boldsymbol{b}$ whose optimal control recovers the CFG update of (8)?

To connect this to diffusion sampling, we take the dynamics $\boldsymbol{b}$ to be the probability-flow ODE of (2), written in Tweedie form via (4) and (5), with the denoising target replaced by the control $\boldsymbol{u}$ itself,

$$\dot{\boldsymbol{x}}_s = \boldsymbol{b}(s, \boldsymbol{x}_s, \boldsymbol{u}_s) \triangleq f(s)\boldsymbol{x}_s - g(s)^2\big(\alpha_s \boldsymbol{u}_s - \boldsymbol{x}_s\big)/(2\beta_s^2). \tag{14}$$

Setting $\boldsymbol{u}_s = \hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)$ at every step recovers ordinary unconditional sampling, while any other choice of control defines an equally valid, but different, flow-ODE trajectory. What remains is to specify h, and here we make a significant assumption: we posit that the value function is equal to the quantity CFG attempts to maximise at each timestep:

$$\hat{V}(t, \boldsymbol{x}_t) \triangleq \phi(t, \boldsymbol{x}_t) = \log\left(p(\boldsymbol{x}_t \mid \xi)/p(\boldsymbol{x}_t)\right). \tag{15}$$

We argue that this is a reasonable assertion: CFG is already effective at conditioning the diffusion trajectory–so choosing a value function whose gradients align with the CFG score incentivises movement in that direction, and this assumption immediately satisfies the boundary condition of (12). We now ask what constraints (what h) has ϕ as its value function under (12).

Substituting (14) and (15) into the reverse-time HJB equation (12) and using the Kolmogorov forward equation to derive $\hat{V}_t$ allows us to solve for a unique h. We show in Theorem E.11 that $\hat{V} = \phi$ solves

[1] We state the reverse-time HJB equation in Lemma E.9, and provide a rigorous HJB derivation in Appendix E.2 based on Bellman (1957); Liberzon (2012); Berkovitz & Medhin (2012).

[2] See Problem E.5 for full conditions, and note that many are immediately satisfied due to the differentiability of the trained denoiser and Gaussian transition kernels.

(12) for every $s \in [t_l, T]$ when

$$h(s, \boldsymbol{x}_s) = \underbrace{\frac{1}{2} g(s)^2 \big\langle \nabla_{\boldsymbol{x}_s} \phi(s, \boldsymbol{x}_s), \nabla_{\boldsymbol{x}_s} \log p(\boldsymbol{x}_s \mid \xi) \big\rangle}_{\text{conditional score alignment}} + \underbrace{r(s) \, \|\Upsilon(s, \boldsymbol{x}_s)\|_2^2}_{\text{CFG score magnitude}}$$

$$\underbrace{-\frac{1}{2} \frac{g(s)^2 \alpha_s}{\beta_s^2} \big\langle \nabla_{\boldsymbol{x}_s} \phi(s, \boldsymbol{x}_s), \Upsilon(s, \boldsymbol{x}_s) \big\rangle}_{\text{CFG score regularisation}} + \underbrace{\frac{1}{2} g(s)^2 \Delta_{\boldsymbol{x}_s} \phi(s, \boldsymbol{x}_s)}_{\text{early-convergence penalty}}, \quad \text{where} \tag{16}$$

$$\Upsilon(s, \boldsymbol{x}_s) \triangleq \text{proj}_{U(s) - \hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)} \left\{ \frac{1}{4} \frac{g(s)^2 \alpha_s}{r(s) \beta_s^2} \nabla_{\boldsymbol{x}_s} \phi(s, \boldsymbol{x}_s) \right\}. \tag{17}$$

The corresponding optimal control is then given by

$$\boldsymbol{u}_t^\star \triangleq \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y) + \Upsilon(t, \boldsymbol{x}_t). \tag{18}$$

CFG-like guidance can then be recovered by reversing Tweedie's formula (5), and using (4) to solve for the score imposed by the optimal control:

$$\nabla_{\boldsymbol{x}_t} \log \bar{p}(\boldsymbol{x}_t \mid \xi) = -\frac{\boldsymbol{x}_t - \alpha_t \boldsymbol{u}_t^\star}{\beta_t^2} = -\frac{1}{\beta_t} \Big[\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \Omega_Y) - \frac{\alpha_t}{\beta_t} \Upsilon \Big]. \tag{19}$$

Traditional CFG, scheduled variants, and Kingma's SNR. CFG and other scheduled guidance variants with a positive guidance-weight function $\omega(t) > 0$ can be recovered by setting

$$r(t) = \frac{g(t)^2 \alpha_t^2}{\beta_t^4} \cdot \frac{1}{4\omega(t)}, \quad U(t) = \mathbb{R}^m \; \forall \, t \in (0, T], \tag{20}$$

and substituting (17) and (20) into (19). *Traditional CFG* (Ho & Salimans, 2022) follows when $\omega(t) = w > 0$, and we cover CFG++ (Chung et al., 2025) and IntCFG (Kynkäänniemi et al., 2024) in Appendix B. Note that for CFG, CFG++, and IntCFG, the $r(t)$ of (20) is proportional to $-d\,\text{SNR}(t)/dt = g(t)^2 \alpha_t^2 / \beta_t^4$, with $\text{SNR}(t) \triangleq \alpha_t^2 / \beta_t^2$. Let $\lambda_t \triangleq \log \text{SNR}(t)$ (Kingma's log-SNR), then Kingma & Gao (2023) show $-d\lambda_t/dt$ is the weighting that converts the uniform noise-matching ELBO objective into an equivalent image-matching objective, and $\|\boldsymbol{\epsilon} - \hat{\boldsymbol{\epsilon}}\|_2^2 = \text{SNR}(t) \|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0\|_2^2$. Thus, minimising $\mathbb{E}\|\boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0\|_2^2$ weighted by $g(t)^2 \alpha_t^2 / \beta_t^4$ is the same training objective as the (unweighted) noise-prediction ELBO. In other words, CFG, CFG++, and IntCFG are optimal controls with an ℓ_2 penalty that is specified with the same scale as the denoising training objective.

We now ask: *what do the constraints of (16) imply about the paths constructed by CFG-like guidance?* The *conditional score alignment* term in (16) rewards paths that align the conditional and CFG scores, which is reasonable, given we seek trajectories that converge to high-density conditional regions. This term is scaled by $g(s)^2$, which decreases as $s \to 0$ under most diffusion parameterisations (Karras et al., 2022), suggesting that the conditional alignment matters most in early-to-mid sampling times. The CFG score projection, Υ, can be thought of as the *maximum allowable CFG movement* given the current Tweedie estimate, $\hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)$, at time s. Assuming that $\hat{\boldsymbol{x}}_0 \in U(s)$ (the control set should contain all Tweedie estimates), then Υ is either $\mathbf{0}$ or aligned with the CFG score, $\nabla_{\boldsymbol{x}_s} \phi$, in the sense that $\langle \nabla_{\boldsymbol{x}_s} \phi(s, \boldsymbol{x}_s), \Upsilon(s, \boldsymbol{x}_s) \rangle \geq 0$ (by Bertsekas, 2009, Prop. 1.1.9), and thus, the *CFG score regularisation* term is ≤ 0. This term is negligible at early sampling times, $\alpha_s \to 0$ as $s \to T$, but grows as $\beta_s \to 0$ as $s \to 0$, and the only way to avoid this penalty is to find a local optimum ($\nabla \phi = 0$) or enter a region that no longer has viable movement in the direction of the CFG score, $\Upsilon = 0$. Thus, we can view this term as a penalty for not finding a local maximum, or not maximising allowable growth of ϕ. The *CFG score magnitude* term blindly rewards large CFG score magnitudes regardless of sign; however, this is entirely negated by the ℓ_2 penalty of (13) at the optimal control, $\boldsymbol{u}_t^\star$. The *early-convergence penalty* term is negative at a mode of ϕ ($\Delta_{\boldsymbol{x}_s} \phi < 0$) and is scaled by $g(s)^2$ (which shrinks as $s \to 0$), so this term penalises trajectories that sit near a sharp maximum at early sampling times, discouraging them from settling into a mode too early (reducing mode collapse).

Thus, CFG-like guidance is an optimal control specified in the same loss scale as the denoising training objective, and it finds paths that: (i) are encouraged to align the conditional and CFG scores at early-to-mid sampling times; (ii) are encouraged to converge to a local maximum at late sampling times; and (iii) are penalised for nearing local maximums at early sampling times.

5 CURVATURE-AWARE CFG (CACFG)

Section 4 shows that CFG-style guidance solves an optimal control problem, but says nothing about whether the resulting control is a *safe* one to take. Compensating for weak class adherence with a large guidance scale ω (or small $r(s)$) may move the sampling trajectory into regions with few training samples, and thus, require the model to extrapolate. This is a known failure mode of neural networks–adversarial regions populate extrapolated sections of the learned manifold, and they cause networks to behave erratically (Szegedy et al., 2014; Goodfellow et al., 2015; Madry et al., 2018).

Existing corrections to CFG address this by assuming the underlying image manifold has some known geometry (Chung et al., 2022; 2025). Such hypotheses are reasonable when $\boldsymbol{x}_0$ is itself a natural image, but for the latent diffusion models used in practice (Rombach et al., 2022), $\boldsymbol{x}_0$ is instead a VAE latent code, and the manifold it occupies is whatever geometry that particular VAE has learned. *A manifold hypothesis tuned to one model's VAE carries no guarantee of transferring to another's.*

Algorithm 1 Curvature-Aware Guidance

Require: diffusion model: $\boldsymbol{\epsilon}_{\theta^\star}$; target $y \in \Omega_Y$; sampling steps with $t_0 = t_l \approx 0 < \cdots < t_n = T$; noise schedules α_t, β_t; guidance weight schedule $\omega(t)$; control set $U(t)$; upper radius r_u; interpolation parameter γ.

$\boldsymbol{x}_{t_n} \sim \mathcal{N}(0, I)$
for $i = n, \dots, 1$ **do**
  $\boldsymbol{\epsilon}_\Omega \leftarrow \boldsymbol{\epsilon}_{\theta^\star}(t_i, \boldsymbol{x}_{t_i}, \Omega_Y)$
  $\boldsymbol{\epsilon}_\xi \leftarrow \boldsymbol{\epsilon}_{\theta^\star}(t_i, \boldsymbol{x}_{t_i}, \xi)$
  $\tilde{\boldsymbol{\epsilon}} \leftarrow \boldsymbol{\epsilon}_\Omega + \omega(t_i)(\boldsymbol{\epsilon}_\xi - \boldsymbol{\epsilon}_\Omega)$
  $\hat{\boldsymbol{x}}_0^{\text{cfg}} \leftarrow (\boldsymbol{x}_{t_i} - \beta_{t_i}\tilde{\boldsymbol{\epsilon}})/\alpha_{t_i}$
  $\rho_i^2 \leftarrow (1-\gamma)\dfrac{\alpha_{t_i}^2}{\alpha_{t_i}^2 + \beta_{t_i}^2}\, r_u^2 + \gamma\, r_u^2$
  $\boldsymbol{u}_i^\star \leftarrow \text{proj}_{U_\gamma(t_i)}(\hat{\boldsymbol{x}}_0^{\text{cfg}})$ ▷ clip to $\|\cdot\|_2 \le \rho_i$
  $\boldsymbol{\epsilon}_i^\star \leftarrow (\boldsymbol{x}_{t_i} - \alpha_{t_i}\boldsymbol{u}_i^\star)/\beta_{t_i}$
  $\boldsymbol{x}_{t_{i-1}} \leftarrow$ Reverse flow-ode step with $\boldsymbol{\epsilon}_i^\star$, $\boldsymbol{x}_{t_i}$
end for
return $\boldsymbol{x}_{t_0}$

We instead propose a correction built on a property the VAE is trained to have: the regularisation term of the VAE objective in (9) penalises the encoder distribution for departing from an isotropic standard normal, $q_\phi(\boldsymbol{x}_0 \mid \boldsymbol{z}) \approx \mathcal{N}(\mathbf{0}, I^{(m\times m)})$ (Kingma & Welling, 2014). This term rewards encodings that are close to the origin, and we hypothesise that valid encodings become more sparse as distance from the origin increases (Appendix C.2), which means that these regions are likely adversarial from the perspective of the diffusion denoiser. We propose restricting the control set from Section 4 to latents contained within a hypersphere with radius r_u. That is, we consider a situation where $\|\boldsymbol{X}_0\|_2^2 \le r_u^2$, and then by Lemma E.4, the magnitude of Tweedie's estimate is also upper bounded by r_u[3]: $\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 \le r_u^2$. This limits the ability of large control inputs to direct the sample into these hypothesised adversarial regions.

Unfortunately, restricting the generative domain in this way would require knowledge of the learned VAE latent space. To circumvent this we again leverage the Gaussian regularisation of the VAE loss and assume $\boldsymbol{X}_0 \sim \mathcal{N}(\mathbf{0}, I^{(m\times m)})$, under which $\|\boldsymbol{X}_0\|_2^2$ follows a chi-squared distribution with m degrees of freedom. We then define a confidence level p and a corresponding radius r_u such that:

$$P\big(\|\boldsymbol{X}_0\|_2^2 \le r_u^2\big) = p. \tag{21}$$

Empirically, we find that the Tweedie estimates (i.e. the unclipped control input) of standard CFG sampling routinely violate the 95.0% bound imposed by this assumption (Figure 1). Furthermore, under $\boldsymbol{X}_0 \sim \mathcal{N}(\mathbf{0}, I^{(m\times m)})$ and the transition kernel of Theorem E.1, $\boldsymbol{X}_0$ and $\boldsymbol{x}_t$ are jointly Gaussian, so the Tweedie estimate $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)$ is itself Gaussian, and thus,

$$P\left(\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 \le \frac{\alpha_t^2}{\alpha_t^2 + \beta_t^2}\, r_u^2\right) = p \quad \forall\, t \in [t_l, T]. \tag{22}$$

That is, for $t > 0$ the Tweedie estimate is expected to lie *closer* to the origin than the raw bound r_u.

A curvature-aware control set. The two bounds above represent the two extremes of the Gaussian hypothesis: (21) assumes only the raw radius r_u, while (22) assumes the hypothesis holds exactly. We interpolate between the two with a single tuneable parameter $\gamma \in [0, 1]$, and take the interpolated radius to define a time-varying *convex* control set within the optimal control framework of Section 4,

$$U_\gamma(t) \triangleq \left\{\boldsymbol{u}_t \in \mathbb{R}^m : \|\boldsymbol{u}_t\|_2^2 \le \frac{\alpha_t^2}{\alpha_t^2 + \beta_t^2}\, r_u^2 + \gamma\left(r_u^2 - \frac{\alpha_t^2}{\alpha_t^2 + \beta_t^2}\, r_u^2\right)\right\} \quad \forall\, t \in [t_l, T]. \tag{23}$$

At $\gamma = 0$ we assume the full Gaussian hypothesis, while $\gamma = 1$ requires only the upper bound to hold. We substitute $U_\gamma(t)$ into (20), and refer to this guidance scheme as curvature-aware CFG (CACFG).

[3]There exists no measure on the truncated distribution that would permit $\|\hat{x}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 > r_u^2$.

6 CACFG EXPERIMENTS

Diffusion models. We evaluate CACFG using two latent diffusion models of varying sizes: Stable Diffusion 1.5 (SD1.5) (Rombach et al., 2022) and Stable Diffusion XL (SDXL) (Podell et al., 2024). Both are text-conditioned latent diffusion models that run the diffusion of (1) on a VAE latent space trained using the same objective of (9). SD1.5 and SDXL have differing generative resolutions of 512×512 and 1024×1024, which have substantially different compute costs. Thus, for SD1.5 experiments we generate 10,000 samples for each guidance-dataset-strength setting, while for SDXL we generate 2,000. We run both models on a NVIDIA RTX4090, and the curvature-aware projection does not incur substantive compute overhead–all models run at their originally reported speeds.

Datasets. We generate images conditioned on two datasets: ImageNet (Russakovsky et al., 2015), which provides class-conditional generation across 1000 object categories, and COCO (Lin et al., 2014), which provides caption-conditioned generation–a more sophisticated generative task.

Guidance schedules. We compare three guidance-weight schedules $\omega(t)$: the constant-weight CFG of (8) (Ho & Salimans, 2022), and the two scheduled alternatives CFG++ (Chung et al., 2025) and IntCFG (Kynkäänniemi et al., 2024), whose exact forms are given in Appendix B. Each schedule is evaluated both on its own and combined with the curvature-aware control set of (23), giving the CACFG, CACFG++, and IntCACFG guidance schemes respectively. Across all curvature-aware variants we fix the confidence level of (21) at $p = 0.95$, which corresponds to $r_u = 129.16$ and $r_u = 257.16$ for SD1.5 and SDXL, respectively. Based on initial ablations (Appendix C.3), we use the blending parameter $\gamma = 0.5$ and $\gamma = 1.0$ for SD1.5 and SDXL, respectively. Additionally, we use the DDIM ODE solver (Song et al., 2021a) for SD1.5 and SDXL, as it is the most common deterministic solver.

Guidance weights. We report the guidance weight of CFG and compute an approximate equivalent weight under each other schedule. That is, we find schedule parameters, θ, such that $\int_0^T \omega_{\text{CFG}} dt = \int_0^T \omega(t;\theta) dt$. This compares schedules with equal approximate total path perturbations. For SD1.5 we use CFG weights $\omega_{\text{CFG}} = 5, 7.5, 10, 15$, and for SDXL we use $\omega_{\text{CFG}} = 10, 15, 20, 25$. This discrepancy arises because SDXL is significantly more resistant to excessive CFG guidance, making comparisons at lower guidance strengths less meaningful.

Metrics. We report nine metrics measuring image quality across varying axes. Fréchet Inception Distance (FID) (Heusel et al., 2018) measures the similarity of real and generated image distributions. Density and Coverage (Naeem et al., 2020) report the fidelity and diversity of the generated images by comparing manifolds fitted to real and generated images. NIQE (Mittal et al., 2013), BRISQUE (Mittal et al., 2012), and TReS (Golestaneh et al., 2022) are no-reference perceptual quality metrics, scoring an image's naturalness based on its pixel statistics. CLIP score (Radford et al., 2021) measures prompt alignment via the cosine similarity between a generated image's CLIP embedding and its conditioning. Human Preference Score v2 (HPSv2) (Wu et al., 2023) is a model trained to mimic human rankings of image quality. We note that HPSv2 and CLIP score are known to artificially improve at large CFG strengths despite poorer image quality and increasing saturation (Xie et al., 2026). We therefore additionally report Saturation (Sadat et al., 2025), computed as the mean of the saturation channel after converting each image from RGB to HSV, to measure this effect.

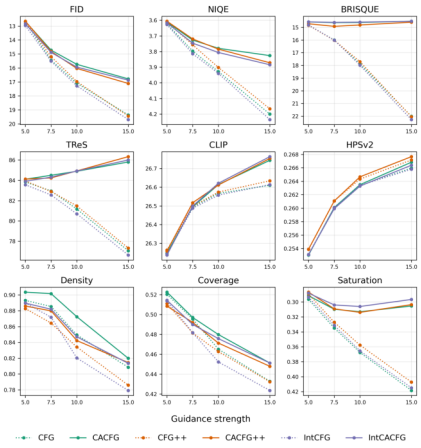


Figure 2: Performance of CFG, CFG++, and IntCFG with (solid line) and without (dotted line) our curvature-aware correction using SD1.5 on COCO. Some metric axes have been flipped so that the improvement direction is vertically up the page. Zoom in for easier viewing.

6.1 RESULTS

Distribution quality and diversity. When using SD1.5 we find FID, Density, and Coverage are improved or equivalent under the curvature-aware (CA) correction for all three guidance schedules and both datasets, with the gap widening substantially at higher guidance strengths (Figures 2 and 6).

Since Density and Coverage improve alongside FID rather than trading off against it, the corrected samples are more accurately contained within the real-image manifold and also cover a wider region of it, indicating an improvement in both quality and diversity.

The same pattern holds for SDXL (Figure 7); however, Density and Coverage also depend on guidance strength in a way they do not for SD1.5. At the lowest guidance strength tested, we find both Density and Coverage are marginally degraded by the CA correction, before favouring the correction from moderate guidance strengths onward, alongside a FID advantage that similarly grows across the same range. We hypothesise two explanations for this phenomenon. First, SDXL is a higher-capacity model and is more robust to the manifold departures, thus, the CA control set has less to correct at low guidance, and its benefit only becomes apparent once uncorrected guidance has sufficiently corrupted the sampling trajectory. Second, at low guidance strengths the more restrictive control constraint, $U_\gamma(t)$, limits the sampler's ability to search the latent space, a cost that is outweighed once the guidance strength (effective search budget) is high enough. Thus, across both SD1.5 and SDXL, the CA correction provides superior conditional generation at moderate-to-high guidance strengths, with a less severe quality-diversity tradeoff than that exhibited by unconstrained guidance methods. These findings are supported qualitatively by the samples in Figure 3 and Appendix D.2.

Perceptual quality. NIQE, BRISQUE, and TReS all favour the curvature-aware variants at every guidance strength for both SD1.5 and SDXL, with the largest gains at high guidance. As these are no-reference measures of image pixel statistics, this indicates that curvature-aware guidance produces images with more natural low-level statistics, independently of any comparison to the real data distribution, and that this benefit is consistent across model scale, guidance schedule, and diffusion model.



Figure 3: Comparison of CFG (Ho & Salimans, 2022) and CACFG (ours) at guidance strength $\omega = 15$ on the COCO (Lin et al., 2014) dataset. CACFG better preserves fine detail: the middle image has a correctly numbered clock face, while the rightmost image reconstructs Van Gogh's *The Starry Night*.

Prompt alignment. CLIP score improves only marginally under ImageNet but more substantially under COCO, for both SD1.5 and SDXL. We hypothesise that ImageNet's single-object class prompts are simple enough that standard guidance can satisfy them at these guidance strengths, leaving no room for the CA correction to help, whereas COCO's complex caption prompts are harder to satisfy: by keeping the sampling trajectory closer to the latent image manifold, curvature-aware guidance allows higher guidance strengths to be spent productively toward matching the full caption, rather than being wasted on manifold-departing artefacts.

Human-preference metrics and saturation. HPSv2 is the one metric where the curvature-aware variants are outperformed by their uncorrected counterparts, an effect that is mild for SD1.5 but more pronounced for SDXL. However, the Saturation values of both Figures 6 and 7 show that this coincides with increasing saturation levels, while curvature-aware guidance keeps saturation roughly constant. Given that Xie et al. (2026) have shown HPSv2 and CLIP scores artificially increase alongside saturation, and that every other image-quality metric favours the curvature-aware correction for both models, we interpret the higher HPSv2 of uncorrected guidance as an instance of this failure mode rather than a quality advantage.

7 Conclusion

We show via the HJB equation that CFG and other guidance schedules can be derived as the optimal control for a continuous-time optimisation problem. Existing CFG corrections assume the underlying image manifold has some known geometry, but this is not guaranteed for latent diffusion models, whose latent image manifold is governed by what the VAE has learned. We use the Gaussian regularisation term of the VAE loss to infer that most image encodings lie within a hypersphere of estimable radius. We then propose a curvature-aware CFG correction that constrains the control set of our optimisation problem to a time-varying hypersphere whose radius is informed by the estimated radius. We find that at high guidance scales this generates images with substantially better quality, as measured by FID, NIQE, BRISQUE, TReS, Density, and Coverage. Thus, our constrained guidance provides superior conditional generation with a less severe quality-diversity tradeoff than that exhibited by unconstrained guidance methods, and functionally no additional compute cost.

REFERENCES

Brian D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.

Richard Bellman. *Dynamic programming*. Princeton University Press, Princeton, NJ, 1957.

Leonard D. Berkovitz and Negash G. Medhin. *Nonlinear Optimal Control Theory*. Chapman and Hall/CRC, Boca Raton, FL, 2012. ISBN 978-1-4665-6026-0.

Dimitri P. Bertsekas. *Convex Optimization Theory*. Athena Scientific, 2009.

Arwen Bradley and Preetum Nakkiran. Classifier-free guidance is a predictor-corrector. *Trans. Mach. Learn. Res.*, 2025, 2025.

Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35, NeurIPS*, 2022.

Hyungjin Chung, Suhyeon Lee, and Jong Chul Ye. Decomposed diffusion sampler for accelerating large-scale inverse problems. In *12th International Conference on Learning Representations, ICLR*, 2024.

Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. CFG++: manifold-constrained classifier free guidance for diffusion models. In *13th International Conference on Learning Representations, ICLR*, 2025.

Max Collins, Jordan Vice, Tim French, and Ajmal Mian. NatADiff: adversarial boundary guidance for natural adversarial diffusion. In *14th International Conference on Learning Representations, ICLR*, 2026.

G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989.

Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems 34, NeurIPS*, 2021.

Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011. ISSN 01621459.

S. Alireza Golestaneh, Saba Dadsetan, and Kris M. Kitani. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *IEEE/CVF Winter Conference on Applications of Computer Vision, WACV*, pp. 3989–3999. IEEE, 2022.

Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *3rd International Conference on Learning Representations, ICLR*, 2015.

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, 2018.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems 33, NeurIPS*, 2020.

Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Netw.*, 4(2): 251–257, 1991.

Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6:695–709, 2005.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems 35, NeurIPS*, 2022.

Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. Guiding a diffusion model with a bad version of itself. In *Advances in Neural Information Processing Systems 37, NeurIPS*, 2024.

Diederik P. Kingma and Ruiqi Gao. Understanding diffusion objectives as the ELBO with simple data augmentation. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36, NeurIPS*, 2023.

Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR*, 2014.

Tuomas Kynkäänniemi, Miika Aittala, Tero Karras, Samuli Laine, Timo Aila, and Jaakko Lehtinen. Applying guidance in a limited interval improves sample and distribution quality in diffusion models. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37, NeurIPS*, 2024.

Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press, Princeton, NJ, 2012. ISBN 978-0-691-15187-8.

Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *European Conference on Computer Vision, ECCV*, volume 8693 of *Lecture Notes in Computer Science*, pp. 740–755, 2014.

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.

Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Trans. Image Process.*, 21(12):4695–4708, 2012.

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal Process. Lett.*, 20(3):209–212, 2013.

Badr Moufad, Yazid Janati, Alain Durmus, Ahmed Ghorbel, Eric Moulines, and Jimmy Olsson. Conditional diffusion models with classifier-free Gibbs-like guidance, 2025.

Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *Proceedings of the 37th International Conference on Machine Learning, ICML*, Proceedings of Machine Learning Research. PMLR, 2020.

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 *12th International Conference on Learning Representations, ICLR*, 2024.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE/CVF, 2022.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet large scale visual recognition challenge. *Int. J. Comput. Vis.*, 115(3):211–252, 2015.

Seyedmorteza Sadat, Otmar Hilliges, and Romann M. Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *13th International Conference on Learning Representations, ICLR*, 2025.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *9th International Conference on Learning Representations, ICLR*, 2021a.

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*, 2021b.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In *2nd International Conference on Learning Representations, ICLR*, 2014.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23:1661–1674, 2011.

Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.

Dian Xie, Shitong Shao, Lichen Bai, Zikai Zhou, Bojun Cheng, Shuo Yang, Jun Wu, and Zeke Xie. Guidance matters: Rethinking the evaluation pitfall for text-to-image generation. In *14th International Conference on Learning Representations, ICLR*, 2026.

Appendix

Table of Contents

A Statements **14**
- A.1 Limitations . . . 14
- A.2 Ethics and broader impacts . . . 14
- A.3 Reproducibility . . . 14
- A.4 AI use . . . 14

B Comparison Methods CFG Form **14**

C Ablations **15**
- C.1 SD1.5 theoretical Tweedie upper bound . . . 15
- C.2 True VAE latent distribution . . . 17
- C.3 Curvature-aware upper bound blending parameter . . . 17

D Curvature-Aware Experiments **20**
- D.1 Metric plots . . . 20
- D.2 Additional samples . . . 22
- D.3 Quantitative results . . . 30
- D.4 Comparison with adaptive projected guidance (APG) . . . 32

E HJB Derivation **34**
- E.1 Necessary background results . . . 34
- E.2 HJB CFG formulation . . . 37

A STATEMENTS

A.1 LIMITATIONS

We note three limitations of our work, each of which we view as a promising direction for future work rather than a fundamental shortcoming of the proposed optimal control framework. First, the control set $U_\gamma(t)$ we optimise over is time-varying but not state-varying; extending it to depend jointly on sampling time and the current state $\boldsymbol{x}_t$ would more directly incentivise training a model to perform the image-manifold projection. Second, our optimal control perspective does not currently explain guidance techniques such as APG (Sadat et al., 2025), which act by manipulating the guidance gradient itself, rather than by reweighting it through the guidance schedule. Finally, our formulation is posed for deterministic sampling processes, and does not immediately extend to stochastic samplers. However, these extensions are unnecessary for understanding and improving traditional classifier-free guidance on deterministic samplers, which was the primary objective of this work.

A.2 ETHICS AND BROADER IMPACTS

We adhere to the ICLR code of ethics. We acknowledge that this work is concerned with improving the quality and controllability of conditional image generation, a capability that can raise legitimate ethical concerns, such as the generation of deep fakes or offensive content. However, we do not identify any major ethical concerns arising directly from our contribution: CACFG refines an existing, widely used guidance mechanism (Ho & Salimans, 2022) rather than introducing a new generative capability, and all experiments in this work are conducted using publicly available, pretrained diffusion models. We do not use private or sensitive data, and all datasets and models used in this work are publicly released and broadly studied. By making our code publicly available, we aim to support transparency and reproducibility, and we encourage the responsible use of the guidance techniques we study.

A.3 REPRODUCIBILITY

We ensure reproducibility by providing a detailed description of our algorithm in the main text (Algorithm 1). Additionally, our code and the configuration files necessary to replicate our experiments are made publicly available. Comparison methods are implemented using default parameters or with parameters selected based on the respective authors' tuning algorithms.

A.4 AI USE

We used large language models as a writing and verification aid during the preparation of this paper. Specifically, AI assistance was used to check and typeset proofs that we had already derived by hand, to edit and improve the clarity of rough prose drafts, and to assist with LaTeX formatting of figures. AI assistance was not used to write experiment code, nor to generate the ideas, methods, or theoretical contributions presented in this work; all such content originates from the authors.

B COMPARISON METHODS CFG FORM

Both CFG++ (Chung et al., 2025) and Interval-CFG (IntCFG) (Kynkäänniemi et al., 2024) propose guidance scheduling schemes to modulate the strength of CFG-like guidance across diffusion sampling paths. Similarly to Section 4, we let the guidance scale of (8) vary with time, $\omega(t)$. We consider a discretisation $0 = t_0 < t_1 < \cdots < t_k = T$ of the canonical image diffusion of Definition E.1, with sampling proceeding from t_k down to t_0, i.e. each step moves $\boldsymbol{x}_{t_i}$ to the less-noisy $\boldsymbol{x}_{t_{i-1}}$. The guidance schedules of CFG++ and IntCFG are as follows.

CFG++. Chung et al. (2025) use a manifold-constrained gradient to force the diffusion sampling trajectory to remain on the theorised linear image manifold (Chung et al., 2022; 2024). Under common diffusion scheduling schemes this culminates in a time-varying CFG guidance strength, $\omega(t)$. Given the discretisation above and using a DDIM diffusion schedule, Chung et al.'s CFG++

guidance scale is given by

$$\omega(t_i) = w \cdot \frac{\alpha_{t_{i-1}}\beta_{t_i}}{\alpha_{t_{i-1}}\beta_{t_i} - \alpha_{t_i}\beta_{t_{i-1}}}, \tag{24}$$

where $w > 0$ is a hyperparameter that governs the strength of the CFG++ guidance. Using Euler derivative approximations: $\dot{\alpha}_{t_i} \approx \frac{\alpha_{t_i} - \alpha_{t_{i-1}}}{\Delta t} \implies \alpha_{t_{i-1}} \approx \alpha_{t_i} - \Delta t \cdot \dot{\alpha}_{t_i}$ and $\dot{\beta}_{t_i} \approx \frac{\beta_{t_i} - \beta_{t_{i-1}}}{\Delta t} \implies \beta_{t_{i-1}} \approx \beta_{t_i} - \Delta t \cdot \dot{\beta}_{t_i}$, then (24) can be written as

$$\omega(t) \approx w \cdot \left[\frac{1}{\Delta t} \cdot \frac{\alpha_t\beta_t}{\alpha_t\dot{\beta}_t - \dot{\alpha}_t\beta_t} - \frac{\dot{\alpha}_t\beta_t}{\alpha_t\dot{\beta}_t - \dot{\alpha}_t\beta_t}\right], \tag{25}$$

where Δt functions as a parameter describing how coarse the diffusion discretisation is. Notably, from (25), it is clear that $\omega(t)$ diverges as $\Delta t \to 0$. This suggests that CFG++ is sensitive to the number of sampling steps, with an exploding guidance scale as the granularity of the ODE solver increases.

Interval CFG (IntCFG). Kynkäänniemi et al. (2024) hypothesise that guidance at high noise levels reduces sample diversity, and is redundant at low noise levels, where the trajectory has already converged onto the correct mode. Thus, IntCFG confines guidance to an interval of intermediate noise levels. In our notation this is the piecewise-constant schedule

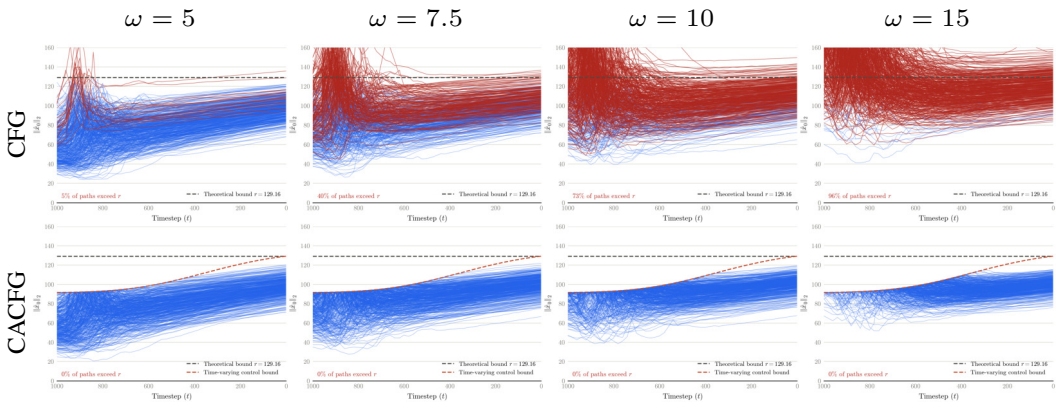
$$\omega(t) = \begin{cases} w, & \frac{\beta_t}{\alpha_t} \in (\sigma_{\text{lo}}, \sigma_{\text{hi}}], \\ 1, & \text{otherwise}, \end{cases} \tag{26}$$

for hyperparameters $0 < \sigma_{\text{lo}} < \sigma_{\text{hi}}$ and guidance weight $w > 1$, selected by grid search against a downstream metric (e.g. FID) rather than derived. Note that technically IntCFG results in a discontinuous $\omega(t)$ scale, which conflicts with HJB assumptions made in Section 4. To circumvent this, a smooth sigmoid transition could be used at each cutoff point σ_{lo} and σ_{hi} to ensure continuity. However, in this paper we use the standard IntCFG proposed by Kynkäänniemi et al. (2024) and specified in (26).

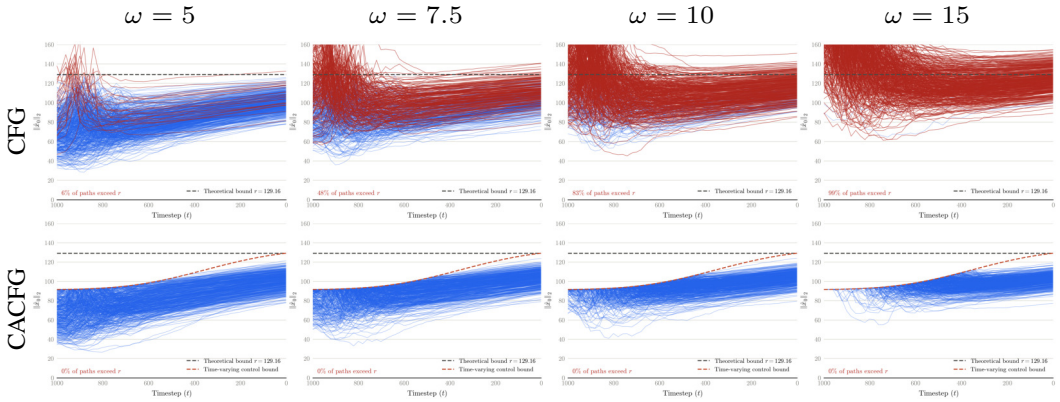
C Ablations

C.1 SD1.5 theoretical Tweedie upper bound

Figure 4 extends Figure 1 to the full range of guidance strengths tested in the main paper, showing Tweedie's magnitude ($\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2$) of SD1.5 CFG and CACFG sampling trajectories across $\omega \in \{5, 7.5, 10, 15\}$ on (a) ImageNet (Russakovsky et al., 2015) and (b) COCO (Lin et al., 2014). Notably, it highlights that increasing guidance strength dramatically increases the number of paths that violate the theoretical 95% upper limit, $r_u = 129.16$.



(a) ImageNet



(b) COCO

Figure 4: Tweedie's magnitude ($\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2$) of SD1.5 CFG and CACFG (ours) sampling trajectories at guidance strengths $\omega \in \{5, 7.5, 10, 15\}$ on (a) ImageNet (Russakovsky et al., 2015) and (b) COCO (Lin et al., 2014). Red paths violate the theoretical 95% upper magnitude bound $r_u = 129.16$ (black dashed line); the violating fraction is annotated in each panel. The proportion of violating paths grows sharply with guidance strength for CFG, and this growth is suppressed by our curvature-aware correction, with its time-varying upper bound (red dashed line). We recommend zooming in for easier viewing.

C.2 True VAE latent distribution

Figure 5 shows the empirical distribution of SD1.5 VAE encodings ($\|\boldsymbol{x}_0\|_2$) of the ImageNet validation set, alongside the theoretical chi density implied by the isotropic Gaussian hypothesis proposed in Section 5. The two do not agree because the VAE is trained with a regularisation term that encourages encodings towards the theoretical density, but does not enforce their adherence. In practice we find that the high-dimensional theoretical density appears as a narrow spike, whereas the true distribution of encodings is visibly broader and right-skewed, with a tail extending up to an empirical maximum of 211.19. Despite this, $r_u = 129.16$ falls just past the mode of the empirical distribution, where the bulk of the density has already accumulated, so the bound continues to capture a large proportion of real encodings even though the theoretical density does not capture the distribution's true shape.

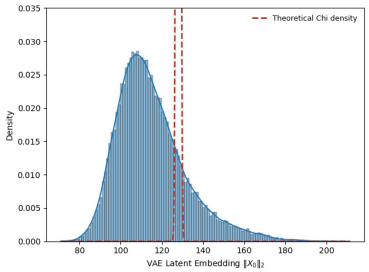


Figure 5: Empirical distribution of SD1.5 VAE encodings ($\|\boldsymbol{x}_0\|_2$) of the ImageNet validation set, i.e. the distribution of $\|\boldsymbol{x}_0\|_2$. Red line denotes the theoretical chi density implied by the Gaussian hypothesis in Section 5.

To test whether this mismatch matters in practice, we compare CACFG against *CACFG**, an otherwise identical scheme that instead sets $r_u = 211.19$, i.e. the maximum SD1.5 VAE encoding magnitude observed on the ImageNet validation set. From Table 1 we see that CACFG* is worse than CACFG on BRISQUE, TReS, and Saturation at every guidance strength tested. FID, NIQE, Density, and Coverage follow the same pattern from strength 7.5 onward. CLIP favours CACFG* at strengths 5 and 7.5, ties at strength 10, and favours CACFG at 15. HPSv2 favours CACFG* across all guidance strengths. However, as we discuss in the main paper, Xie et al. (2026) tie HPSv2 to oversaturation, and CACFG*'s saturation exceeds CACFG's throughout, and all other metrics favour CACFG at high guidance strengths. Thus, we interpret the higher HPSv2 of CACFG* as an instance of the saturation-induced failure mode rather than a quality advantage.

We hypothesise that CACFG outperforms CACFG* because the tighter, theoretically-derived radius encloses a denser region of the encoding manifold: Figure 5 shows valid encodings are more sparse with distance from the origin, so the smaller radius keeps sampling trajectories within the well-populated region the denoiser was trained on, rather than permitting convergence to low-confidence regions (Szegedy et al., 2014; Goodfellow et al., 2015; Madry et al., 2018).

C.3 Curvature-aware upper bound blending parameter

The blending parameter γ interpolates between the two extremes of the Gaussian hypothesis underlying $U_\gamma(t)$: at $\gamma = 0$ the control set assumes the hypothesis holds exactly, while at $\gamma = 1$ it retains only the raw upper bound r_u (23). We sweep $\gamma \in \{0.0, 0.1, \ldots, 1.0\}$ for CACFG on ImageNet, for both SD1.5 (Table 2) and SDXL (Table 3). A consistent pattern emerges across both models: loosening the control set steadily improves prompt alignment and human-preference scores (CLIP, HPSv2) and steadily increases saturation, while the remaining fidelity metrics (FID, NIQE, BRISQUE, Density, Coverage) improve sharply over the first few steps of the sweep and then plateau. No single value of γ is simultaneously optimal across every metric, so selecting γ requires an explicit choice about how to balance prompt alignment against the residual fidelity cost of moving away from $\gamma = 0.0$; we describe that choice, and its cost, for each model below.

For SD1.5, we select $\gamma = 0.5$. At this value, six of the nine metrics (NIQE, BRISQUE, Density, Coverage, CLIP, and HPSv2) sit within 1% of the best value achieved anywhere in the sweep, and FID (14.97) is close to its sweep-best (14.53) and far below the $\gamma = 0.0$ value (16.95). Beyond

Table 1: Image quality for SD1.5 samples generated for ImageNet using CFG (Ho & Salimans, 2022) and CACFG with varying upper bound limits, r_u. CACFG denotes the $r_u = 129.16$ bound used in the main paper, while CACFG* is the $r_u = 211.19$ bound derived from the maximum SD1.5 VAE encoding of the ImageNet validation dataset. The best result for each dataset and guidance strength is **bold**.

Strength	Guidance	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
5	CFG	**13.67**	4.06	14.98	83.24	24.851	25.076	**1.128**	0.453	0.299
	CACFG	13.70	**4.05**	**14.83**	**83.47**	24.849	25.070	1.124	**0.454**	**0.295**
	CACFG*	**13.67**	**4.05**	14.98	83.27	**24.856**	**25.081**	1.123	**0.454**	0.299
7.5	CFG	15.12	4.27	15.47	82.80	**25.251**	25.680	1.139	0.443	0.338
	CACFG	**14.97**	**4.22**	**14.59**	**83.80**	25.243	25.653	**1.140**	**0.445**	**0.316**
	CACFG*	15.11	4.27	15.33	82.85	25.249	**25.687**	1.132	0.444	0.337
10	CFG	16.22	4.44	17.13	81.41	25.391	25.967	1.084	0.424	0.371
	CACFG	**15.73**	**4.32**	**14.66**	**84.10**	**25.395**	25.933	**1.119**	**0.433**	**0.321**
	CACFG*	16.16	4.40	16.57	81.79	**25.395**	**25.979**	1.094	0.426	0.365
15	CFG	17.78	4.75	21.95	78.14	25.491	26.214	0.994	0.391	0.424
	CACFG	**16.52**	**4.42**	**14.80**	**84.60**	**25.540**	26.175	**1.063**	**0.413**	**0.317**
	CACFG*	17.55	4.57	19.06	80.02	25.505	**26.264**	1.023	0.400	0.397

$\gamma = 0.5$, prompt-alignment and preference gains diminish sharply: CLIP and HPSv2 each move by less than 0.15% between $\gamma = 0.5$ and $\gamma = 1.0$. Over the same range, the two metrics that continue to degrade past $\gamma = 0.5$ worsen further: TReS, which peaks early at $\gamma = 0.1$ and declines steadily thereafter, falling below even its $\gamma = 0.0$ value; and saturation, which increases monotonically with γ throughout the sweep. We therefore treat $\gamma = 0.5$ as close to the point of diminishing returns for SD1.5: it captures nearly all of the available prompt-alignment gain without paying the additional TReS and saturation cost of continuing to $\gamma = 1.0$.

For SDXL, we instead select $\gamma = 1.0$. Unlike SD1.5, fidelity metrics for SDXL plateau from roughly $\gamma = 0.3$–0.4 onward and remain within a small margin of their sweep-best all the way to $\gamma = 1.0$. FID, TReS, Density, and Coverage are each within 2% of their respective optima at $\gamma = 1.0$, while CLIP and HPSv2 are both maximised at $\gamma = 1.0$. This is consistent with our hypothesis that SDXL, as a higher-capacity model, is more robust to the manifold departure the curvature-aware correction addresses, and so tolerates a looser control set with comparatively little fidelity cost. The corresponding cost is a mild increase in BRISQUE (2.7% above its sweep-best) and saturation, which reaches its highest value in the sweep at $\gamma = 1.0$; we accept this cost given the consistent, and in several cases maximal, gains elsewhere.

Table 2: Image quality for SD1.5 (Rombach et al., 2022) samples generated for ImageNet (Russakovsky et al., 2015) dataset using CACFG across varying curvature-aware blending parameter values.

γ	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
0.0	16.95	5.44	23.59	85.15	24.835	23.844	0.938	0.395	0.194
0.1	14.53	4.56	17.36	86.28	25.159	25.069	1.111	0.444	0.236
0.2	14.54	4.32	15.28	85.65	25.199	25.403	1.142	0.446	0.266
0.3	14.72	4.24	14.64	84.96	25.233	25.550	1.140	0.448	0.289
0.4	14.89	4.22	14.50	84.29	25.242	25.621	1.140	0.447	0.305
0.5	14.97	4.22	14.59	83.80	25.243	25.653	1.140	0.445	0.316
0.6	15.03	4.23	14.75	83.48	25.244	25.669	1.145	0.447	0.323
0.7	15.11	4.24	14.93	83.28	25.244	25.672	1.138	0.445	0.328
0.8	15.10	4.25	15.04	83.12	25.247	25.673	1.137	0.444	0.331
0.9	15.14	4.26	15.14	83.04	25.255	25.681	1.137	0.445	0.333
1.0	15.09	4.26	15.20	82.97	25.255	25.690	1.133	0.445	0.335

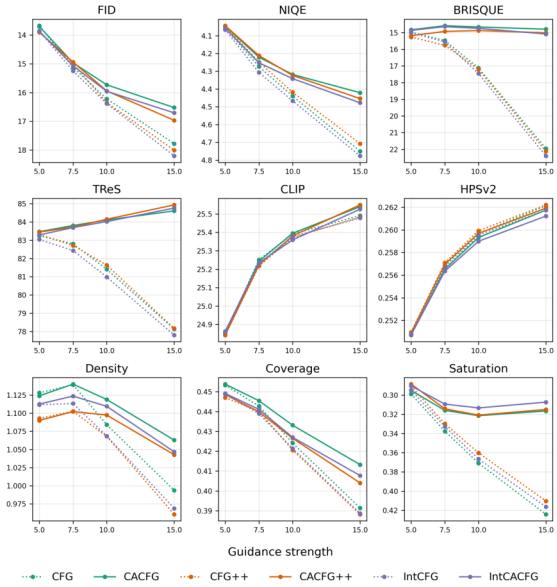
Table 3: Image quality for SDXL (Podell et al., 2024) samples generated for ImageNet (Russakovsky et al., 2015) dataset using CACFG across varying curvature-aware blending parameter values.

γ	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
0.0	35.84	4.96	26.93	78.06	24.751	23.531	0.863	0.143	0.139
0.1	29.85	4.48	26.04	81.82	25.213	25.409	1.073	0.174	0.165
0.2	29.22	4.42	26.05	83.37	25.274	25.924	1.133	0.183	0.182
0.3	29.09	4.43	26.21	83.96	25.291	26.114	1.147	0.184	0.193
0.4	29.02	4.44	26.45	84.22	25.322	26.234	1.117	0.181	0.201
0.5	29.02	4.47	26.53	84.31	25.361	26.294	1.140	0.184	0.206
0.6	29.06	4.48	26.59	84.38	25.342	26.319	1.157	0.187	0.209
0.7	29.09	4.48	26.70	84.41	25.350	26.339	1.159	0.186	0.211
0.8	29.10	4.49	26.66	84.43	25.356	26.353	1.159	0.185	0.213
0.9	29.10	4.50	26.76	84.40	25.358	26.370	1.159	0.187	0.214
1.0	29.12	4.49	26.75	84.39	25.369	26.380	1.158	0.186	0.215

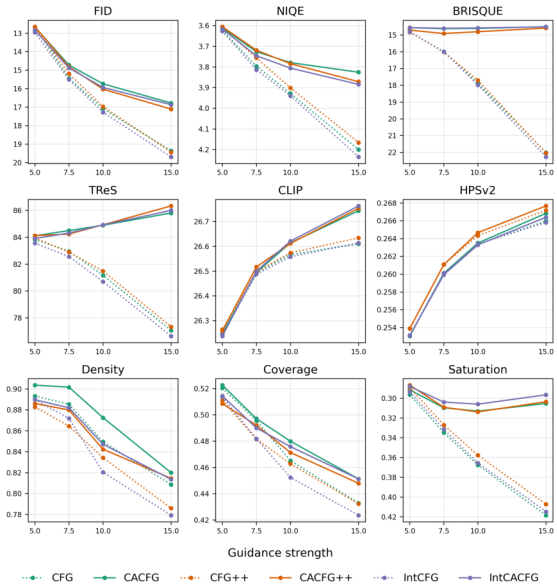
D CURVATURE-AWARE EXPERIMENTS

D.1 METRIC PLOTS

Here we provide the metric plots of SD1.5 (Rombach et al., 2022) and SDXL (Podell et al., 2024) diffusion models on the ImageNet (Russakovsky et al., 2015) and COCO (Lin et al., 2014) generation tasks. We discuss these results in Section 6.1.

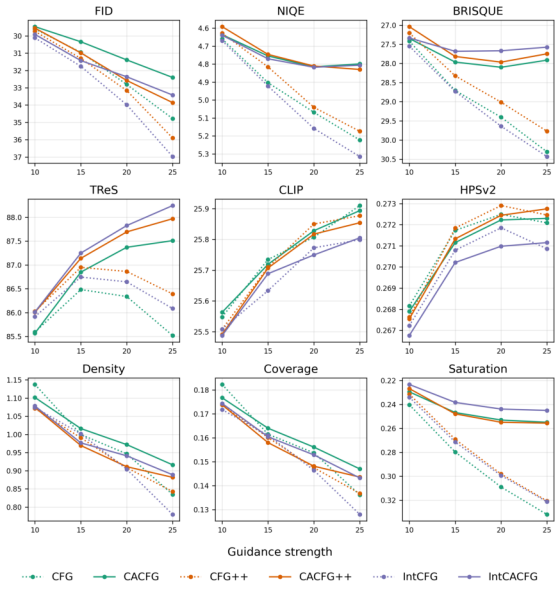


(a) ImageNet

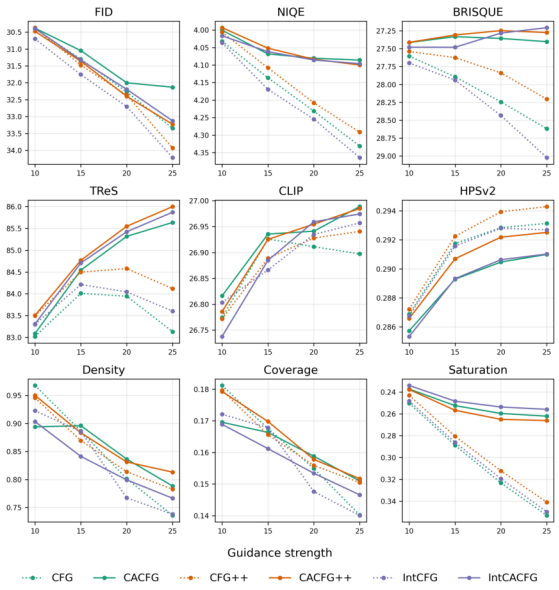


(b) COCO

Figure 6: Performance of CFG, CFG++, and IntCFG with (solid line) and without (dotted line) our curvature-aware correction using SD1.5 on (a) ImageNet and (b) COCO. Some metric axes have been flipped so that the improvement direction is vertically up the page. Note at larger guidance strengths FID, NIQE, BRISQUE, TReS, Density, Coverage, and CLIP are all improved by our curvature-aware correction. We recommend zooming in for easier viewing.



(a) ImageNet



(b) COCO

Figure 7: Performance of CFG, CFG++, and IntCFG with (solid line) and without (dotted line) our curvature-aware correction using SDXL on (a) ImageNet and (b) COCO. Some metric axes have been flipped so that the improvement direction is vertically up the page. Note at larger guidance strengths FID, NIQE, BRISQUE, TReS, Density, and Coverage are all improved by our curvature-aware correction. We recommend zooming in for easier viewing.

D.2 Additional samples

Here we present additional curvature-aware diffusion samples across varying guidance methods, guidance strengths, and models.

D.2.1 SD1.5



(a) CFG vs CACFG (ours)



(b) CFG++ vs CACFG++ (ours)



(c) IntCFG vs IntCACFG (ours)

Figure 8: CFG (Ho & Salimans, 2022), CFG++ (Chung et al., 2025), and IntCFG (Kynkäänniemi et al., 2024) visualised against their curvature-aware variants (ours): CACFG, CACFG++, IntCACFG. Samples are generated using SD1.5 (Rombach et al., 2022) targeting the ImageNet (Russakovsky et al., 2015) dataset with a guidance strength of $\omega = 15$.

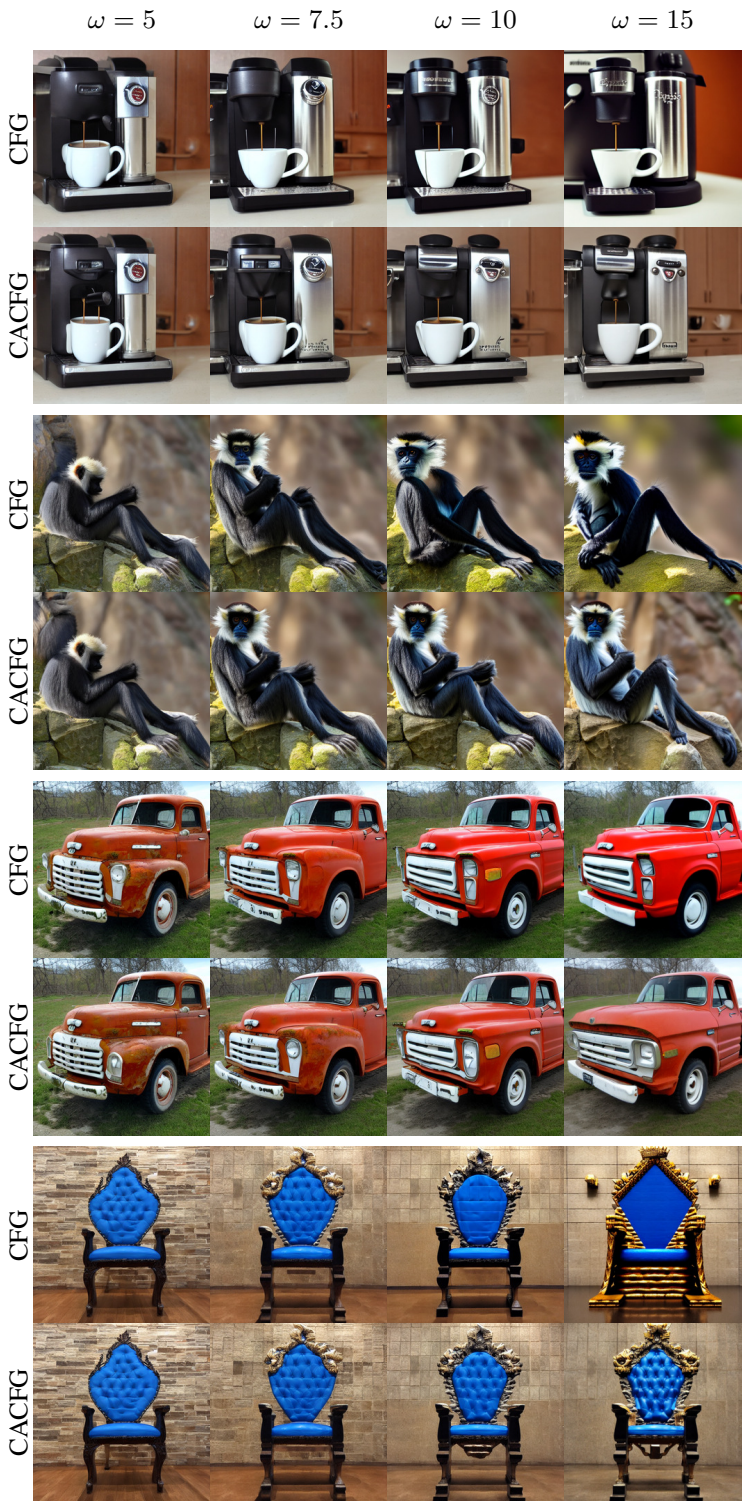
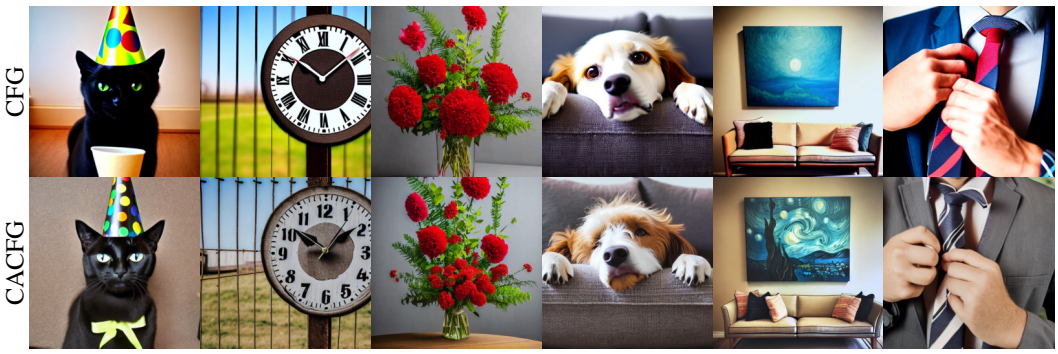
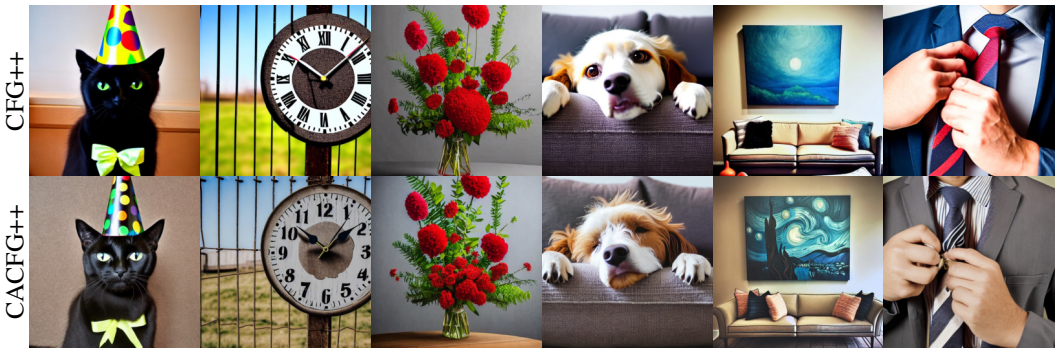


Figure 9: CFG (Ho & Salimans, 2022) and CACFG (ours) samples generated using SD1.5 (Rombach et al., 2022) targeting the ImageNet (Russakovsky et al., 2015) dataset. Guidance scale varies across $\omega = 5, 7.5, 10, 15$.



(a) CFG vs CACFG (ours)



(b) CFG++ vs CACFG++ (ours)



(c) IntCFG vs IntCACFG (ours)

Figure 10: CFG (Ho & Salimans, 2022), CFG++ (Chung et al., 2025), and IntCFG (Kynkäänniemi et al., 2024) visualised against their curvature-aware variants (ours): CACFG, CACFG++, IntCACFG. Samples are generated using SD1.5 (Rombach et al., 2022) targeting the COCO (Lin et al., 2014) dataset with a guidance strength of $\omega = 15$.

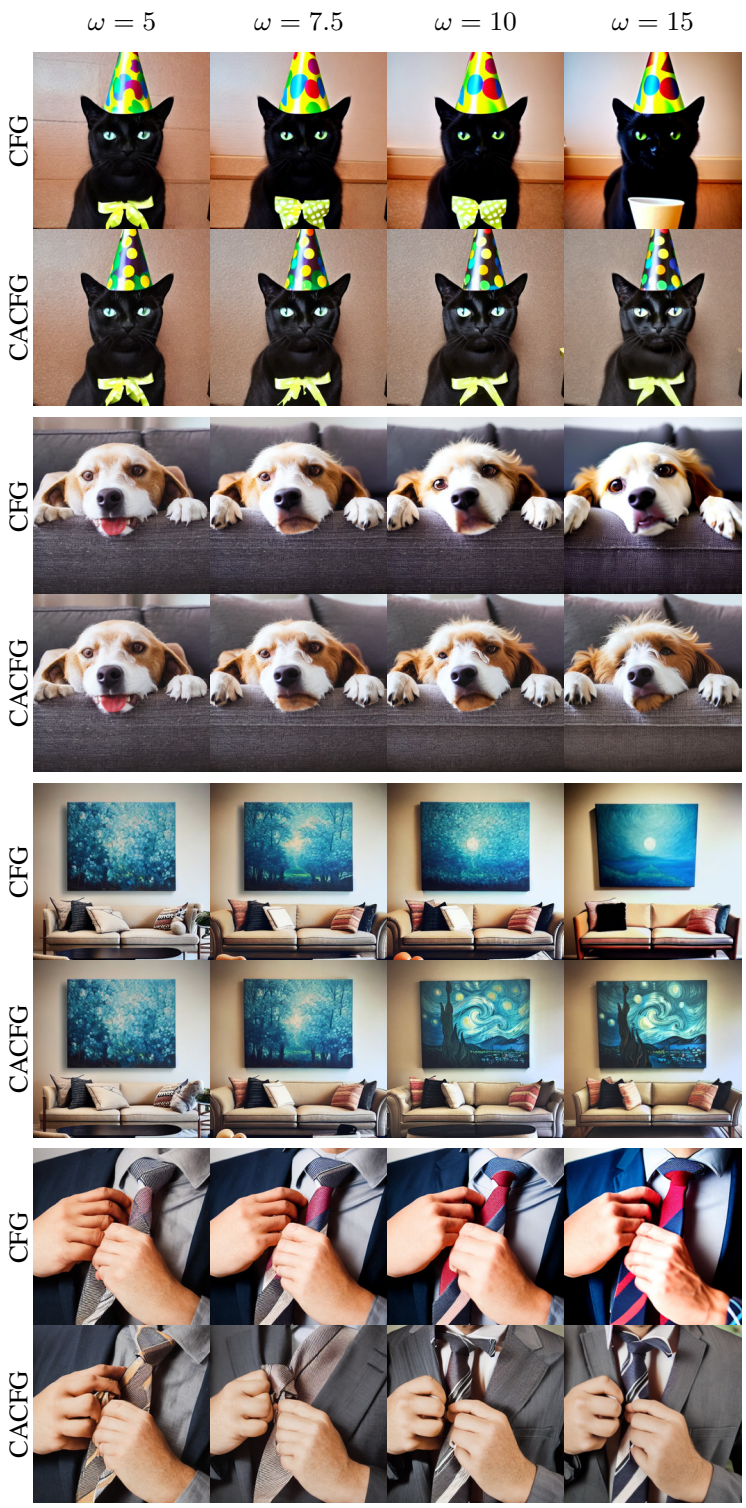


Figure 11: CFG (Ho & Salimans, 2022) and CACFG (ours) samples generated using SD1.5 (Rombach et al., 2022) targeting the COCO (Lin et al., 2014) dataset. Guidance scale varies across $\omega = 5, 7.5, 10, 15$.

D.2.2 SDXL



(a) CFG vs CACFG (ours)



(b) CFG++ vs CACFG++ (ours)



(c) IntCFG vs IntCACFG (ours)

Figure 12: CFG (Ho & Salimans, 2022), CFG++ (Chung et al., 2025), and IntCFG (Kynkäänniemi et al., 2024) visualised against their curvature-aware variants (ours): CACFG, CACFG++, IntCACFG. Samples are generated using SDXL (Podell et al., 2024) targeting the ImageNet (Russakovsky et al., 2015) dataset with a guidance strength of $\omega = 25$.



Figure 13: CFG (Ho & Salimans, 2022) and CACFG (ours) samples generated using SDXL (Podell et al., 2024) targeting the ImageNet (Russakovsky et al., 2015) dataset. Guidance scale varies across $\omega = 10, 15, 20, 25$.



(a) CFG vs CACFG (ours)



(b) CFG++ vs CACFG++ (ours)



(c) IntCFG vs IntCACFG (ours)

Figure 14: CFG (Ho & Salimans, 2022), CFG++ (Chung et al., 2025), and IntCFG (Kynkäänniemi et al., 2024) visualised against their curvature-aware variants (ours): CACFG, CACFG++, IntCACFG. Samples are generated using SDXL (Podell et al., 2024) targeting the COCO (Lin et al., 2014) dataset with a guidance strength of $\omega = 25$.

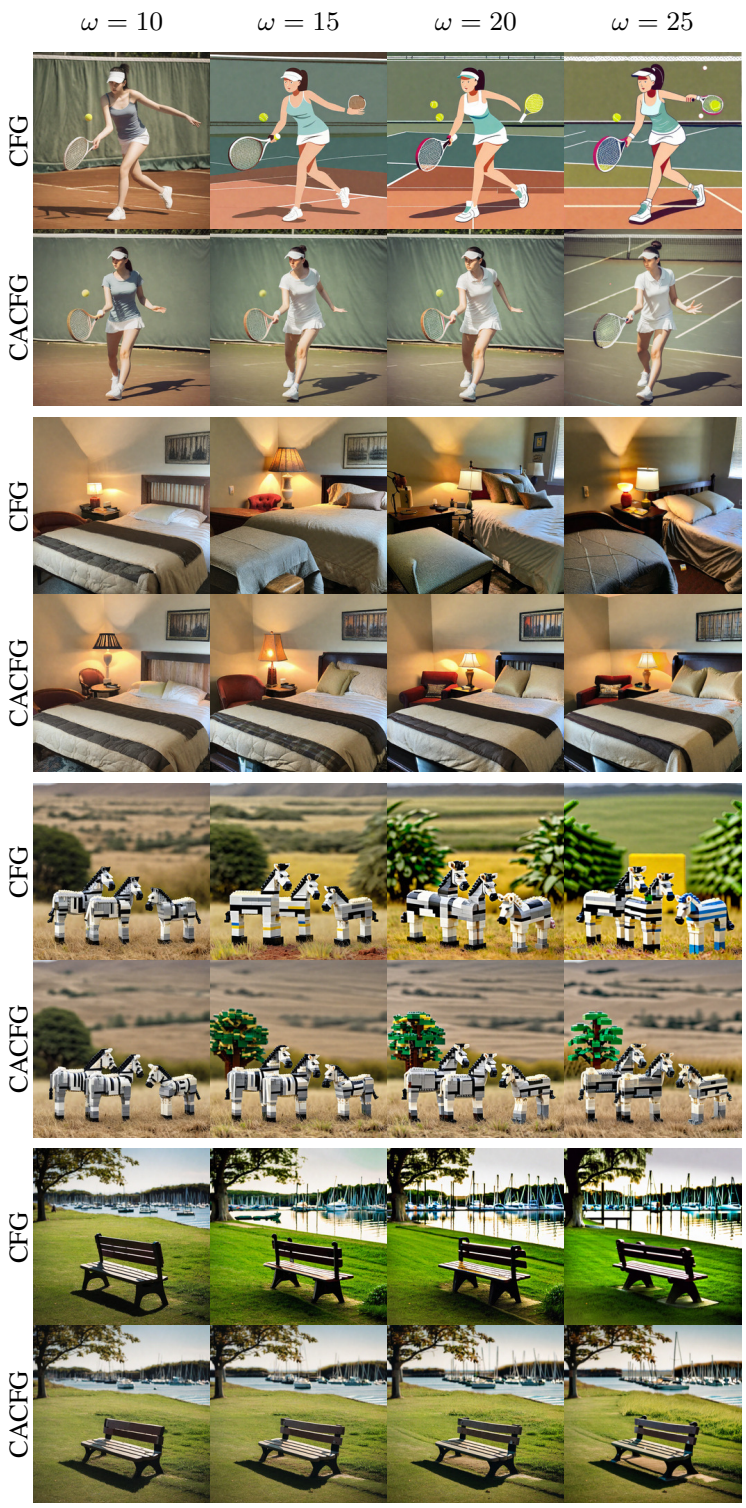


Figure 15: CFG (Ho & Salimans, 2022) and CACFG (ours) samples generated using SDXL (Podell et al., 2024) targeting the COCO (Lin et al., 2014) dataset. Guidance scale varies across $\omega = 10, 15, 20, 25$.

D.3 QUANTITATIVE RESULTS

Here we provide the numeric values that the Figures 6 and 7 were created from.

Table 4: Image quality for SD1.5 samples generated for ImageNet and COCO datasets across a range of guidance strengths and schedules. Best result for each dataset and guidance strength is **bold**, and the best value between curvature-aware and non-curvature-aware guidance is given in light green, i.e., colour denotes whether CA or non-CA methods were favoured.

Dataset	Strength	Guidance	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
ImageNet	5	CFG	**13.67**	4.06	14.98	83.24	24.851	25.076	**1.128**	0.453	0.299
		CFG++	13.88	4.05	15.27	83.31	24.842	**25.096**	1.093	0.447	0.291
		IntCFG	13.86	4.07	15.00	83.04	24.859	25.079	1.112	0.449	0.295
		CACFG	13.70	4.05	**14.83**	**83.47**	24.849	25.070	1.124	**0.454**	0.295
		CACFG++	13.90	**4.04**	15.19	83.44	24.843	25.090	1.090	0.449	**0.289**
		IntCACFG	13.86	4.06	14.87	83.29	**24.862**	25.072	1.113	0.449	0.290
	7.5	CFG	15.12	4.27	15.47	82.80	**25.251**	25.680	1.139	0.443	0.338
		CFG++	14.99	4.25	15.76	82.71	25.218	**25.709**	1.103	0.440	0.330
		IntCFG	15.25	4.31	15.57	82.44	25.231	25.658	1.113	0.439	0.333
		CACFG	14.97	4.22	**14.59**	**83.80**	25.243	25.653	**1.140**	**0.445**	0.316
		CACFG++	**14.94**	**4.21**	14.93	83.72	25.223	25.696	1.102	0.440	0.314
		IntCACFG	15.07	4.25	14.64	83.69	25.235	25.635	1.123	0.441	**0.309**
	10	CFG	16.22	4.44	17.13	81.41	25.391	25.967	1.084	0.424	0.371
		CFG++	16.35	4.42	17.21	81.63	25.372	**25.993**	1.068	0.421	0.360
		IntCFG	16.38	4.47	17.46	80.98	25.376	25.954	1.069	0.421	0.367
		CACFG	**15.73**	**4.32**	**14.66**	84.10	**25.395**	25.933	**1.119**	**0.433**	0.321
		CACFG++	15.94	**4.32**	14.88	**84.14**	25.381	25.970	1.097	0.427	0.321
		IntCACFG	15.96	4.34	14.75	84.02	25.359	25.900	1.109	0.427	**0.313**
	15	CFG	17.78	4.75	21.95	78.14	25.491	26.214	0.994	0.391	0.424
		CFG++	18.01	4.71	22.12	78.17	25.483	**26.221**	0.961	0.388	0.410
		IntCFG	18.21	4.78	22.39	77.81	25.478	26.175	0.969	0.389	0.416
		CACFG	**16.52**	**4.42**	**14.80**	84.60	25.540	26.175	**1.063**	**0.413**	0.317
		CACFG++	16.97	4.45	15.01	**84.93**	**25.549**	26.192	1.043	0.404	0.315
		IntCACFG	16.71	4.48	15.09	84.76	25.526	26.122	1.047	0.408	**0.307**
COCO	5	CFG	12.69	3.62	14.82	83.84	26.245	25.310	0.893	0.520	0.296
		CFG++	12.68	**3.61**	14.83	83.93	**26.264**	25.386	0.883	0.511	0.289
		IntCFG	12.96	3.63	14.81	83.56	26.243	25.312	0.890	0.514	0.294
		CACFG	12.67	**3.61**	14.57	84.11	26.247	25.303	**0.904**	**0.523**	0.292
		CACFG++	**12.65**	**3.61**	14.71	**84.12**	26.258	**25.389**	0.886	0.509	**0.287**
		IntCACFG	12.84	3.62	**14.56**	83.91	26.237	25.310	0.890	0.514	0.288
	7.5	CFG	15.44	3.80	16.00	82.96	26.495	26.011	0.885	0.496	0.335
		CFG++	15.19	3.76	16.02	82.90	26.498	**26.109**	0.864	0.482	0.327
		IntCFG	15.51	3.81	16.01	82.56	26.488	26.009	0.872	0.482	0.331
		CACFG	**14.73**	3.73	14.63	**84.50**	26.495	26.010	**0.902**	**0.497**	0.310
		CACFG++	14.82	**3.72**	14.92	84.24	**26.517**	26.107	0.880	0.492	0.309
		IntCACFG	14.89	3.75	**14.60**	84.32	26.499	25.993	0.882	0.490	**0.304**
	10	CFG	17.09	3.93	17.88	81.17	26.567	26.335	0.850	0.465	0.368
		CFG++	16.98	3.90	17.70	81.48	26.574	26.432	0.834	0.463	0.358
		IntCFG	17.28	3.94	17.98	80.70	26.558	26.338	0.820	0.452	0.366
		CACFG	**15.74**	**3.78**	14.61	84.90	26.614	26.349	**0.872**	**0.480**	0.313
		CACFG++	16.03	**3.78**	14.81	**84.93**	26.612	**26.466**	0.842	0.471	0.314
		IntCACFG	15.94	3.81	**14.58**	84.92	**26.621**	26.328	0.847	0.476	**0.306**
	15	CFG	19.36	4.20	22.01	77.05	26.610	26.597	0.809	0.433	0.419
		CFG++	19.44	4.17	22.05	77.33	26.634	26.717	0.786	0.432	0.407
		IntCFG	19.70	4.24	22.27	76.64	26.614	26.580	0.779	0.423	0.415
		CACFG	**16.78**	**3.83**	**14.51**	85.81	26.744	26.683	**0.820**	**0.451**	0.305
		CACFG++	17.11	3.87	14.59	**86.34**	26.754	**26.766**	0.814	0.448	0.304
		IntCACFG	16.87	3.88	14.52	86.00	**26.763**	26.639	0.814	**0.451**	**0.297**

Table 5: Image quality for SDXL samples generated for ImageNet and COCO datasets across a range of guidance strengths and schedules. Best result for each dataset and guidance strength is **bold**, and the best value between curvature-aware and non-curvature-aware guidance is given in light green, i.e., colour denotes whether CA or non-CA methods were favoured.

Dataset	Strength	Guidance	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
ImageNet	10	CFG	29.60	4.66	27.42	85.59	25.548	**26.815**	**1.137**	**0.182**	0.240
		CFG++	29.74	4.63	27.21	**86.03**	25.508	26.762	1.071	0.174	0.231
		IntCFG	30.12	4.67	27.55	85.91	25.508	26.722	1.074	0.172	0.234
		CACFG	**29.47**	4.64	27.35	85.57	**25.563**	26.789	1.102	0.177	0.229
		CACFG++	29.55	**4.59**	**27.04**	86.01	25.491	26.754	1.074	0.174	0.227
		IntCACFG	29.91	4.64	27.33	86.01	25.488	26.675	1.078	0.174	**0.223**
	15	CFG	30.95	4.90	28.71	86.48	**25.735**	27.173	1.000	0.162	0.280
		CFG++	31.30	4.82	28.32	86.95	25.710	**27.184**	0.990	0.160	0.269
		IntCFG	31.76	4.92	28.73	86.75	25.633	27.079	1.001	0.161	0.271
		CACFG	**30.34**	**4.75**	27.97	86.85	25.719	27.116	**1.016**	**0.164**	0.247
		CACFG++	30.98	**4.75**	27.82	87.14	25.707	27.133	0.969	0.158	0.248
		IntCACFG	31.43	4.77	**27.68**	**87.25**	25.689	27.021	0.977	0.160	**0.238**
	20	CFG	32.81	5.07	29.40	86.34	25.807	27.249	0.947	0.154	0.309
		CFG++	33.16	5.04	29.01	86.86	**25.850**	**27.291**	0.908	0.147	0.298
		IntCFG	33.98	5.16	29.64	86.65	25.773	27.186	0.904	0.146	0.299
		CACFG	**31.39**	4.82	28.10	87.37	25.828	27.223	**0.972**	**0.156**	0.253
		CACFG++	32.58	**4.81**	27.97	87.69	25.817	27.245	0.911	0.148	0.255
		IntCACFG	32.36	4.82	**27.67**	**87.83**	25.750	27.098	0.941	0.153	**0.244**
	25	CFG	34.78	5.22	30.31	85.52	**25.910**	27.209	0.835	0.136	0.332
		CFG++	35.90	5.17	29.77	86.39	25.877	27.246	0.843	0.137	0.321
		IntCFG	36.98	5.32	30.43	86.08	25.799	27.086	0.780	0.128	0.321
		CACFG	**32.40**	**4.80**	27.91	87.51	25.894	27.230	**0.916**	**0.147**	0.255
		CACFG++	33.86	4.83	27.75	87.97	25.854	**27.276**	0.882	0.144	0.256
		IntCACFG	33.42	4.81	**27.58**	**88.25**	25.805	27.115	0.889	0.143	**0.245**
COCO	10	CFG	30.40	4.03	27.60	83.02	26.774	28.688	**0.968**	**0.181**	0.250
		CFG++	**30.37**	4.01	27.54	83.49	26.771	**28.722**	0.946	0.180	0.243
		IntCFG	30.70	4.04	27.70	83.30	26.803	28.675	0.923	0.172	0.248
		CACFG	30.38	4.00	**27.41**	83.09	**26.816**	28.573	0.894	0.170	0.237
		CACFG++	30.47	**3.99**	**27.41**	**83.50**	26.786	28.658	0.950	0.179	0.238
		IntCACFG	30.39	4.02	27.48	83.31	26.737	28.533	0.903	0.169	**0.234**
	15	CFG	31.30	4.14	27.89	84.01	26.926	29.175	0.887	0.168	0.289
		CFG++	31.48	4.11	27.62	84.50	26.889	**29.225**	0.869	0.166	0.281
		IntCFG	31.76	4.17	27.94	84.21	26.866	29.158	0.886	0.168	0.286
		CACFG	**31.05**	4.07	27.33	84.54	**26.936**	28.929	**0.896**	0.166	0.252
		CACFG++	31.36	**4.05**	**27.31**	**84.77**	26.925	29.069	0.883	**0.170**	0.257
		IntCACFG	31.33	4.06	27.48	84.71	26.885	28.933	0.841	0.161	**0.248**
	20	CFG	32.27	4.23	28.24	83.94	26.911	29.284	0.801	0.155	0.323
		CFG++	32.37	4.21	27.84	84.58	26.928	**29.393**	0.814	0.156	0.312
		IntCFG	32.70	4.25	28.43	84.04	26.935	29.280	0.767	0.148	0.319
		CACFG	**32.00**	**4.08**	27.36	85.32	26.941	29.047	**0.836**	**0.159**	0.260
		CACFG++	32.40	**4.08**	**27.25**	**85.55**	26.955	29.219	0.831	0.158	0.265
		IntCACFG	32.19	4.09	27.28	85.42	**26.959**	29.063	0.799	0.153	**0.254**
	25	CFG	33.35	4.33	28.62	83.13	26.897	29.313	0.736	0.140	0.353
		CFG++	33.93	4.29	28.20	84.12	26.941	**29.429**	0.783	0.150	0.341
		IntCFG	34.22	4.36	29.02	83.60	26.957	29.269	0.738	0.140	0.350
		CACFG	**32.13**	**4.09**	27.40	85.64	**26.989**	29.099	0.788	0.151	0.262
		CACFG++	33.23	4.10	27.27	**86.00**	26.985	29.251	**0.813**	**0.152**	0.266
		IntCACFG	33.13	4.10	**27.21**	85.87	26.975	29.102	0.767	0.147	**0.256**

D.4 Comparison with adaptive projected guidance (APG)

Adaptive Projected Guidance (APG) (Sadat et al., 2025) decomposes the CFG update direction, $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi) - \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y)$, into components parallel and orthogonal to the conditional prediction. They show that manipulating the parallel and orthogonal gradient components allows the saturation and image quality to be controlled more effectively. In this section we compare APG against CFG and CACFG; however, unlike CFG++ (Chung et al., 2025) and IntCFG (Kynkäänniemi et al., 2024), there is no form of $\omega(t)$ that converts the CFG-like guidance in Section 4 to APG. This prevents us from comparing methods at equivalent path costs as we do in the main paper. To address this, we instead compare methods across a range of strengths. We report image quality of samples generated using CFG (Ho & Salimans, 2022), CACFG, and APG (Sadat et al., 2025) across a range of guidance strengths for the ImageNet (Russakovsky et al., 2015) and COCO (Lin et al., 2014) datasets, and SD1.5 (Rombach et al., 2022) and SDXL (Podell et al., 2024) diffusion models.

Counting only best-in-category and second-best-in-category results, we find CACFG is equal to or better than all other guidance methods; however, this overlooks certain patterns of performance (see Tables 6 and 7). We find APG has lower saturation across all regimes, which is reasonable given that it was designed specifically to address oversaturation. By comparison, CACFG's strengths are more nuanced: it achieves superior no-reference quality (NIQE, BRISQUE) across models and datasets. However, on SD1.5 it most consistently attains the best or second-best Density/Coverage scores, whereas on SDXL, CACFG's advantage shifts to TReS, with FID differences between all three methods narrowing to within hundredths. This suggests CACFG's benefit is not tied to a fixed metric, but rather it maintains a strong performance across metrics as model scale changes. Taken together, no method is necessarily favoured, but CACFG is more theoretically grounded due to its optimal control formulation, and can be applied across guidance schedules, while APG requires tuning to a specific diffusion model.

Table 6: Image quality for SD1.5 (Rombach et al., 2022) samples generated for ImageNet (Russakovsky et al., 2015) and COCO (Lin et al., 2014) datasets using CFG (Ho & Salimans, 2022), CACFG, and APG (Sadat et al., 2025). Best result for each dataset are given in **bold green** and second best in light green.

Dataset	Guidance	Strength	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
ImageNet	CFG	5	**13.67**	4.06	14.98	83.24	24.851	25.076	1.128	0.453	0.299
		7.5	15.12	4.27	15.47	82.80	25.251	25.680	1.139	0.443	0.338
		10	16.22	4.44	17.13	81.41	25.391	25.967	1.084	0.424	0.371
		15	17.78	4.75	21.95	78.14	25.491	**26.214**	0.994	0.391	0.424
	APG	5	14.07	4.19	17.07	84.85	24.240	23.826	1.009	0.427	**0.230**
		7.5	14.06	4.51	17.04	85.55	24.806	24.567	1.091	0.443	0.233
		10	14.77	4.72	16.94	86.05	25.105	24.965	1.123	0.443	0.237
		15	16.12	4.82	15.87	**86.17**	25.396	25.481	1.132	0.436	0.252
	CACFG	5	13.70	**4.05**	14.83	83.47	24.849	25.070	1.124	**0.454**	0.295
		7.5	14.97	4.22	**14.59**	83.80	25.243	25.653	**1.140**	0.445	0.316
		10	15.73	4.32	14.66	84.10	25.395	25.933	1.119	0.433	0.321
		15	16.52	4.42	14.80	84.60	**25.540**	26.175	1.063	0.413	0.317
COCO	CFG	5	12.69	3.62	14.82	83.84	26.245	25.310	0.893	0.520	0.296
		7.5	15.44	3.80	16.00	82.96	26.495	26.011	0.885	0.496	0.335
		10	17.09	3.93	17.88	81.17	26.567	26.335	0.850	0.465	0.368
		15	19.36	4.20	22.01	77.05	26.610	26.597	0.809	0.433	0.419
	APG	5	**11.56**	3.88	16.02	85.88	25.755	23.653	0.820	0.504	**0.222**
		7.5	12.43	4.12	15.98	86.63	26.210	24.528	0.859	0.503	0.224
		10	13.68	4.29	15.95	87.01	26.451	25.030	0.870	0.495	0.226
		15	15.47	4.32	15.13	**87.35**	26.648	25.649	0.838	0.467	0.241
	CACFG	5	12.67	**3.61**	14.57	84.11	26.247	25.303	**0.904**	**0.523**	0.292
		7.5	14.73	3.73	14.63	84.50	26.495	26.010	0.902	0.497	0.310
		10	15.74	3.78	14.61	84.90	26.614	26.349	0.872	0.480	0.313
		15	16.78	3.83	**14.51**	85.81	**26.744**	**26.683**	0.820	0.451	0.305

Table 7: Image quality for SDXL (Podell et al., 2024) samples generated for ImageNet (Russakovsky et al., 2015) and COCO (Lin et al., 2014) datasets using CFG (Ho & Salimans, 2022), CACFG, and APG (Sadat et al., 2025). Best result for each dataset are given in **bold green** and second best in light green.

Dataset	Guidance	Strength	FID ↓	NIQE ↓	BRISQUE ↓	TReS ↑	CLIP ↑	HPSv2 ↑	Density ↑	Coverage ↑	Saturation ↓
ImageNet	CFG	10	29.60	4.66	27.42	85.59	25.548	26.815	**1.137**	**0.182**	0.240
		15	30.95	4.90	28.71	86.48	25.735	27.173	1.000	0.162	0.280
		20	32.81	5.07	29.40	86.34	25.807	**27.249**	0.947	0.154	0.309
		25	34.78	5.22	30.31	85.52	**25.910**	27.209	0.835	0.136	0.332
	APG	10	29.70	**4.57**	**26.92**	81.76	25.256	25.528	1.107	0.178	**0.160**
		15	**29.46**	4.70	27.36	83.78	25.612	26.136	1.058	0.172	0.168
		20	30.51	4.76	27.68	85.41	25.706	26.521	1.017	0.165	0.179
		25	31.19	4.80	27.86	86.37	25.818	26.756	1.009	0.165	0.193
	CACFG	10	29.47	4.64	27.35	85.57	25.563	26.789	1.102	0.177	0.229
		15	30.34	4.75	27.97	86.85	25.719	27.116	1.016	0.164	0.247
		20	31.39	4.82	28.10	87.37	25.828	27.223	0.972	0.156	0.253
		25	32.40	4.80	27.91	**87.51**	25.894	27.230	0.916	0.147	0.255
COCO	CFG	10	30.40	4.03	27.60	83.02	26.774	28.688	**0.968**	**0.181**	0.250
		15	31.30	4.14	27.89	84.01	26.926	29.175	0.887	0.168	0.289
		20	32.27	4.23	28.24	83.94	26.911	29.284	0.801	0.155	0.323
		25	33.35	4.33	28.62	83.13	26.897	**29.313**	0.736	0.140	0.353
	APG	10	30.40	4.15	28.63	78.51	26.659	26.580	0.892	0.170	**0.166**
		15	31.08	4.19	28.43	80.84	26.880	27.391	0.870	0.163	0.175
		20	31.01	4.15	27.94	82.62	26.984	27.950	0.933	0.172	0.187
		25	31.39	4.13	27.63	83.94	**26.990**	28.362	0.855	0.161	0.202
	CACFG	10	**30.38**	**4.00**	27.41	83.09	26.816	28.573	0.894	0.170	0.237
		15	31.05	4.07	**27.33**	84.54	26.936	28.929	0.896	0.166	0.252
		20	32.00	4.08	27.36	85.32	26.941	29.047	0.836	0.159	0.260
		25	32.13	4.09	27.40	**85.64**	26.989	29.099	0.788	0.151	0.262

E HJB DERIVATION

In this section we present the derivation of the HJB formulation of CFG-like diffusion guidance, and the modifications necessary to re-derive classic CFG.

E.1 NECESSARY BACKGROUND RESULTS

Canonical image diffusion: We first fix the notation for the forward diffusion process used throughout the remainder of this appendix, following Collins et al. (2026).

Definition E.1 (Canonical Image Diffusion)**.** Let $\boldsymbol{x}_t \in \mathbb{R}^m$ for $t \in [0, T]$, let $f : \mathbb{R} \to \mathbb{R}$ and $g : \mathbb{R} \to \mathbb{R}$ be continuous functions of t, and let $\cdot\, d\boldsymbol{B}_t$ denote an Itô integral with respect to the standard m-dimensional Brownian motion process $\boldsymbol{B}_t$. We define the *canonical image diffusion* to be the forward process

$$d\boldsymbol{x}_t = f(t)\boldsymbol{x}_t\, dt + g(t)I^{(m\times m)} d\boldsymbol{B}_t, \quad \forall\, t \in [0, T]. \tag{27}$$

Let $\boldsymbol{X}_0$ be an $\mathbb{R}^m$-valued random variable representing the initial (image) state, and let Y denote auxiliary conditioning information available for $\boldsymbol{X}_0$ (e.g. a class label), taking values in a countable set Ω_Y with associated sigma-algebra $\mathcal{Y}$ (we take $\mathcal{Y}$ to contain every singleton $\{y\}\ \forall\ y \in \Omega_Y$). We take $\boldsymbol{X}_0$ and Y to be jointly distributed, with joint law $P(\boldsymbol{X}_0, Y)$, a probability measure on $(\mathbb{R}^m \times \Omega_Y,\ \mathcal{B}(\mathbb{R}^m) \otimes \mathcal{Y})$; the marginal laws $P(\boldsymbol{X}_0)$ and $P(Y \in \xi)$ for $\xi \in \mathcal{Y}$, follow in the usual way. For $\xi \in \mathcal{Y}$ with $P(Y \in \xi) > 0$, we write $P(\boldsymbol{X}_0 \mid \xi)$ for the conditional law of $\boldsymbol{X}_0$ given the event $\{Y \in \xi\}$, i.e.

$$P\left(\boldsymbol{X}_0 |\ \xi\right) \triangleq \frac{P(\boldsymbol{X}_0,\ Y \in \xi)}{P(Y \in \xi)}, \tag{28}$$

a probability measure on $(\mathbb{R}^m, \mathcal{B}(\mathbb{R}^m))$. Two special cases of ξ recur throughout: taking $\xi = \Omega_Y$ recovers the **unconditional** case, $P(\boldsymbol{X}_0 \mid \Omega_Y) = P(\boldsymbol{X}_0)$, since $\{Y \in \Omega_Y\}$ is the sure event; and for a specific label $y \in \Omega_Y$ with $P(Y = y) > 0$, taking $\xi = \{y\}$ gives the **class-conditional** law $P(\boldsymbol{X}_0 \mid \{y\})$, which we abbreviate $P(\boldsymbol{X}_0 \mid y)$. For a measurable function $h : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}$ and for $\xi \in \mathcal{Y}$ with $P(Y \in \xi) > 0$, we understand $\int h(\boldsymbol{x}_t, \boldsymbol{x}_0)\, dP(\boldsymbol{x}_0 \mid \xi)$ to denote integration of $h(\boldsymbol{x}_t, \cdot)$ against $P(\boldsymbol{X}_0 \mid \xi)$, without requiring this measure to admit a density (this accommodates empirical measures over a finite dataset).

Conditional forward distribution: The following theorem describes the conditional forward distribution of the canonical image diffusion when conditioned on the state at an arbitrary time τ. This is a well-known result, and we restate the formulation by Collins et al. (2026); we omit the proof here and refer the reader there (the derivation solves (27) via its Stratonovich representation and converts back to the equivalent Itô form).

Theorem E.1 (Conditional Forward Distribution for Denoising Diffusion, Collins et al., 2026)**.** *Under the canonical image diffusion of Definition E.1, define for* $0 \le \tau \le t \le T$

$$\alpha(\tau, t) \triangleq \exp\left(\int_\tau^t f(u)\, du\right), \qquad \beta(\tau, t)^2 \triangleq \alpha(\tau, t)^2 \int_\tau^t \frac{g(u)^2}{\alpha(\tau, u)^2}\, du. \tag{29}$$

Then the forward process in (27) admits the conditional forward distribution

$$\boldsymbol{X}_t \mid \boldsymbol{X}_\tau = \boldsymbol{x}_\tau \sim \mathcal{N}\left(\alpha(\tau, t)\boldsymbol{x}_\tau,\ \beta(\tau, t)^2 I^{(m\times m)}\right) \quad \forall\, t > \tau. \tag{30}$$

With a slight abuse of notation we write $\alpha_t \triangleq \alpha(0, t)$ *and* $\beta_t \triangleq \beta(0, t)$*, so that*

$$\boldsymbol{X}_t \mid \boldsymbol{X}_0 = \boldsymbol{x}_0 \sim \mathcal{N}\left(\alpha_t \boldsymbol{x}_0,\ \beta_t^2 I^{(m\times m)}\right). \tag{31}$$

Score-model link: The following Score-Model Link theorem is based on Karras et al.'s 2022 argument, extended by Collins et al. (2026) by conditioning on the auxiliary random variable Y of Definition E.1. For additional treatments see (Karras et al., 2022; Hyvärinen, 2005; Vincent, 2011). We restate it here under the canonical image diffusion of Definition E.1, replacing the empirical-density formulation of the original statement with integration against the measure $P(\boldsymbol{X}_0 \mid \xi)$ directly; this is equivalent to the version in (Collins et al., 2026), and we omit the proof, referring the reader there.

Theorem E.2 (Score-Model Link, Collins et al., 2026; Karras et al., 2022)**.** *Let $\boldsymbol{x}_t$ evolve according to the canonical image diffusion of Definition E.1, with initial law $P(\boldsymbol{X}_0 \mid \xi)$ for some $\xi \in \mathcal{Y}$.*

Define

$$\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi; \theta) = \frac{\boldsymbol{x}_t - \beta_t \boldsymbol{\epsilon}_\theta(t, \boldsymbol{x}_t, \xi)}{\alpha_t}, \tag{32}$$

where α_t and β_t are as in Theorem E.1, and $\boldsymbol{\epsilon}_\theta : \mathbb{R}^m \times \mathbb{R} \times \mathcal{Y} \to \mathbb{R}^m$ is a model parameterised by $\theta \in \Theta$ with sufficient capacity such that the Universal Approximation Theorem (Cybenko, 1989; Hornik, 1991) holds for all $\xi \in \mathcal{Y}$. Then if $\beta_t^2 > 0 \; \forall \, t \in (0, T]$, the following statements are true:

1.

$$\nabla_{\boldsymbol{x}_t} \log(p(\boldsymbol{x}_t \mid \xi)) = -\frac{1}{\beta_t} \boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi) \quad \forall \, t \in (0, T], \; \boldsymbol{x}_t \in \mathbb{R}^m, \; \xi \in \mathcal{Y}; \tag{33}$$

2.

$$\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi) = \frac{\displaystyle\int_{\mathbb{R}^m} \frac{1}{\beta_t} \big(\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_0\big) p(\boldsymbol{x}_t | \boldsymbol{x}_0) \, dP(\boldsymbol{x}_0 \mid \xi)}{p(\boldsymbol{x}_t \mid \xi)} \quad \forall \, t \in (0, T], \tag{34}$$

$\boldsymbol{x}_t \in \mathbb{R}^m, \; \xi \in \mathcal{Y}$; where

$$\theta^\star \triangleq \arg\min_\theta \, \mathbb{E}_{\boldsymbol{x}_0 \sim P(\boldsymbol{x}_0 | \xi), \; \boldsymbol{x}_t \sim p(\boldsymbol{x}_t | \boldsymbol{x}_0), \; t \sim p(t)} \left[\| \boldsymbol{x}_0 - \hat{\boldsymbol{x}}_0 \|_2^2 \right], \tag{35}$$

and $p(t) \triangleq \frac{1}{T}$.

(36)

By the shorthand of Definition E.1, taking $\xi = \{y\}$ for a specific label $y \in \Omega_Y$ gives the class-conditional score-model $\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, y)$, and taking $\xi = \Omega_Y$ gives the unconditional score-model $\boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \Omega_Y)$; both are the objects we later manipulate in the HJB derivation.

Tweedie's formula: Tweedie's formula (Efron, 2011) relates the posterior mean of $\boldsymbol{X}_0$ given a noisy observation $\boldsymbol{x}_t$ to the score of the marginal density of $\boldsymbol{X}_t$. We state and prove the specialisation of this result to the canonical image diffusion of Definition E.1, conditioned on the event ξ of Definition E.1, as we require both the conditioning and the proof mechanics directly in the upper-bounding lemma that follows.

Given the initial law $P(\boldsymbol{X}_0 \mid \xi)$ of Definition E.1, which we do not assume admits a density, define the conditional marginal density of $\boldsymbol{x}_t$ by

$$p(\boldsymbol{x}_t \mid \xi) \triangleq \int_{\mathbb{R}^m} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0) \, dP(\boldsymbol{x}_0 \mid \xi), \tag{37}$$

where $p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$ is the Gaussian transition density of Theorem E.1. As this transition density is, for $\beta_t^2 > 0$, positive-everywhere and integrates the mass of $P(\boldsymbol{X}_0 \mid \xi)$ against a bounded, smooth kernel, then (37) defines a genuine, positive-everywhere probability density on $\mathbb{R}^m$, regardless of whether $P(\boldsymbol{X}_0 \mid \xi)$ itself admits one.

Theorem E.3 (Tweedie's Formula, adapted from Efron, 2011)**.** *Under the canonical image diffusion of Definition E.1, with $\beta_t^2 > 0 \; \forall \, t \in (0, T]$, and let $p(\boldsymbol{x}_t \mid \xi)$ be as in (37) for $\xi \in \mathcal{Y}$. Define*

$$\boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \xi) \triangleq -\beta_t \nabla_{\boldsymbol{x}_t} \log(p(\boldsymbol{x}_t \mid \xi)). \tag{38}$$

Then

$$\mathbb{E}_{\boldsymbol{X}_0 | \boldsymbol{x}_t, \, \xi} \left[\boldsymbol{X}_0 \right] = \frac{\boldsymbol{x}_t - \beta_t \boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \xi)}{\alpha_t} \quad \forall \, \boldsymbol{x}_t \in \mathbb{R}^m, \; t \in (0, T], \; \xi \in \mathcal{Y},$$

where $\mathbb{E}_{\boldsymbol{X}_0 | \boldsymbol{x}_t, \, \xi}$ denotes expectation with respect to the posterior measure $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \, \xi)$ constructed in the proof (writing $\boldsymbol{x}_t$ for the conditioning event $\{\boldsymbol{X}_t = \boldsymbol{x}_t\}$, analogous to the ξ shorthand of Definition E.1).

Proof. Differentiating (37) under the integral sign (justified since $\beta_t^2 > 0$ implies $p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$ and $\nabla_{\boldsymbol{x}_t} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$ are uniformly bounded in $\boldsymbol{x}_0$, so dominated convergence permits the interchange),

$$\begin{aligned}\nabla_{\boldsymbol{x}_t} \log(p(\boldsymbol{x}_t \mid \xi)) &= \nabla_{\boldsymbol{x}_t} \log\left(\int_{\mathbb{R}^m} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)\, dP(\boldsymbol{x}_0 \mid \xi)\right) \\ &= \frac{1}{p(\boldsymbol{x}_t \mid \xi)} \int_{\mathbb{R}^m} \nabla_{\boldsymbol{x}_t} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)\, dP(\boldsymbol{x}_0 \mid \xi).\end{aligned}$$

By Theorem E.1, $p(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathcal{N}\left(\alpha_t \boldsymbol{x}_0,\ \beta_t^2 I^{(m\times m)}\right) \implies \nabla_{\boldsymbol{x}_t} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = -\frac{1}{\beta_t^2}(\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_0) p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)$, thus,

$$\nabla_{\boldsymbol{x}_t} \log(p(\boldsymbol{x}_t \mid \xi)) = \frac{1}{p(\boldsymbol{x}_t \mid \xi)} \int_{\mathbb{R}^m} \left[-\frac{1}{\beta_t^2}(\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_0)\right] p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)\, dP(\boldsymbol{x}_0 \mid \xi). \tag{39}$$

Define the posterior measure $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi)$ on $(\mathbb{R}^m, \mathcal{B}(\mathbb{R}^m))$ by

$$dP(\boldsymbol{x}_0 \mid \boldsymbol{x}_t,\, \xi) \triangleq \frac{p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)}{p(\boldsymbol{x}_t \mid \xi)}\, dP(\boldsymbol{x}_0 \mid \xi). \tag{40}$$

As $p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)/p(\boldsymbol{x}_t \mid \xi) \geq 0$ and $\int_{\mathbb{R}^m} p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)/p(\boldsymbol{x}_t \mid \xi)\, dP(\boldsymbol{x}_0 \mid \xi) = p(\boldsymbol{x}_t \mid \xi)/p(\boldsymbol{x}_t \mid \xi) = 1$ by (37), (40) is a valid Radon-Nikodym reweighting of $P(\boldsymbol{X}_0 \mid \xi)$ and hence defines a genuine probability measure; this is just Bayes' theorem expressed without requiring $P(\boldsymbol{X}_0 \mid \xi)$ to admit a density. In particular, $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi) \ll P(\boldsymbol{X}_0 \mid \xi)$: *absolute continuity*, i.e., any set assigned zero probability by $P(\boldsymbol{X}_0 \mid \xi)$ is also assigned zero probability by the posterior, since the posterior is obtained by integrating a finite, nonnegative Radon-Nikodym derivative against $P(\boldsymbol{X}_0 \mid \xi)$. Substituting (40) into (39),

$$\begin{aligned}\nabla_{\boldsymbol{x}_t} \log(p(\boldsymbol{x}_t \mid \xi)) &= \int_{\mathbb{R}^m} \left[-\frac{1}{\beta_t^2}(\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_0)\right] dP(\boldsymbol{x}_0 \mid \boldsymbol{x}_t,\, \xi) \\ &= -\frac{1}{\beta_t^2}\boldsymbol{x}_t + \frac{\alpha_t}{\beta_t^2} \int_{\mathbb{R}^m} \boldsymbol{x}_0\, dP(\boldsymbol{x}_0 \mid \boldsymbol{x}_t,\, \xi) \\ &= -\frac{1}{\beta_t^2}\boldsymbol{x}_t + \frac{\alpha_t}{\beta_t^2}\, \mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi}\left[\boldsymbol{X}_0\right]. \end{aligned} \tag{41}$$

Multiplying (41) by $-\beta_t$ and applying (38),

$$\boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \xi) = \frac{\boldsymbol{x}_t - \alpha_t\, \mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi}\left[\boldsymbol{X}_0\right]}{\beta_t} \implies \mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi}\left[\boldsymbol{X}_0\right] = \frac{\boldsymbol{x}_t - \beta_t \boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \xi)}{\alpha_t},$$

which proves the claim. □

Comparing (38) with Theorem E.2, we note that $\boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \xi) = \boldsymbol{\epsilon}_{\theta^\star}(t, \boldsymbol{x}_t, \xi)$ for every $\xi \in \mathcal{Y}$; that is, the population noise-prediction function of Tweedie's formula is the optimal score-model of Theorem E.2, and by the shorthand of Definition E.1 this holds equally for the class-conditional case ($\xi = \{y\}$, written $\boldsymbol{\epsilon}(t, \boldsymbol{x}_t, y)$) and the unconditional case ($\xi = \Omega_Y$, written $\boldsymbol{\epsilon}(t, \boldsymbol{x}_t, \Omega_Y)$).

Boundedness of the Tweedie estimate: An important property of Tweedie's estimate is that it inherits some bounding properties of the underlying data distribution. The following lemma shows that if the data lies almost surely within a closed ball of radius r, then so does the Tweedie estimate, for every $\boldsymbol{x}_t$, t, and conditioning event ξ.

Lemma E.4 (Boundedness of the Tweedie Estimate)**.** *Under the canonical image diffusion of Definition E.1, with $\beta_t^2 > 0\ \forall\, t \in (0, T]$, and suppose there exists $r > 0$ and $\xi \in \mathcal{Y}$ such that*

$$P\left(\|\boldsymbol{X}_0\|_2^2 \leq r^2 \;\middle|\; \xi\right) = 1.$$

Then the Tweedie estimator $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi) \triangleq \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi; \theta^\star) = \mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t,\, \xi}\left[\boldsymbol{X}_0\right]$ of Theorem E.3 satisfies

$$\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 \leq r^2 \quad \forall\, \boldsymbol{x}_t \in \mathbb{R}^m,\ t \in (0, T].$$

Proof. Fix $\boldsymbol{x}_t \in \mathbb{R}^m$ and $t \in (0, T]$, and let $A \triangleq \left\{\boldsymbol{x}_0 \in \mathbb{R}^m : \|\boldsymbol{x}_0\|_2^2 > r^2\right\}$. By hypothesis, $P(\boldsymbol{X}_0 \in A \mid \xi) = 0$. By the absolute continuity $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi) \ll P(\boldsymbol{X}_0 \mid \xi)$ established in the proof of Theorem E.3 (via (40)), it follows that

$$P\left(\boldsymbol{X}_0 \in A \mid \boldsymbol{x}_t, \xi\right) = \int_A \frac{p(\boldsymbol{x}_t \mid \boldsymbol{x}_0)}{p(\boldsymbol{x}_t \mid \xi)}\, dP(\boldsymbol{x}_0 \mid \xi) = 0,$$

since the integrand is finite and nonnegative and is being integrated over the $P(\boldsymbol{X}_0 \mid \xi)$-null set. Hence $\|\boldsymbol{X}_0\|_2^2 \le r^2$ almost surely under the posterior measure $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi)$.

Integrating this almost-sure bound against the probability measure $P(\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi)$ gives

$$\mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi}\left[\|\boldsymbol{X}_0\|_2^2\right] \le r^2. \tag{42}$$

As $\|\cdot\|_2^2$ is convex on $\mathbb{R}^m$, Jensen's inequality gives

$$\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 = \left\|\mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi}\left[\boldsymbol{X}_0\right]\right\|_2^2 \le \mathbb{E}_{\boldsymbol{X}_0 \mid \boldsymbol{x}_t, \xi}\left[\|\boldsymbol{X}_0\|_2^2\right]. \tag{43}$$

Combining (42) and (43),

$$\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 \le r^2.$$

As $\boldsymbol{x}_t \in \mathbb{R}^m$ and $t \in (0, T]$ were arbitrary, the claim holds for all $\boldsymbol{x}_t$ and t. □

Since conditioning cannot introduce probability mass on a set the unconditional law assigns zero mass to, taking $\xi = \Omega_Y$ in the hypothesis of Lemma E.4, i.e. an unconditional bound $P(\|\boldsymbol{X}_0\|_2^2 \le r^2) = 1$, then $P(\|\boldsymbol{X}_0\|_2^2 \le r^2 \mid \xi) = 1$, and hence $\|\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)\|_2^2 \le r^2$, for *every* $\xi \in \mathcal{Y}$ with $P(Y \in \xi) > 0$– in particular, for every class-conditional label $y \in \Omega_Y$. Additionally, note that we are taking $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi)$ to refer to the optimal $\hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \xi; \theta^\star)$ for brevity.

E.2 HJB CFG FORMULATION

Optimal control problem and HJB equation: We first fix our notation for the optimal control problem underlying the HJB equation, following the general formulation of Berkovitz & Medhin (2012) (§2.3, pp. 19–21) and Liberzon (2012) (§5.1.2, pp. 158–159), generalising the control constraint set to vary with time as in Berkovitz & Medhin (2012) (§6.2, p. 150).

Problem E.5 (Optimal Control Problem)**.** *Fix $t_l < T$. Let $\boldsymbol{x}_t \in \mathbb{R}^m$ for $t \in [t_l, T]$, let $\boldsymbol{u}_t \in U(t) \subseteq \mathbb{R}^p$ be a control taking values in a nonempty, time-varying control set $U(t)$, and let $\boldsymbol{b} : \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}^m$. Let $L : \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^p \to \mathbb{R}$ be a running cost and $K : \mathbb{R}^m \to \mathbb{R}$ a terminal cost, and let both be sufficiently regular such that a unique trajectory exists for every admissible control.*[4] *For $(t, \boldsymbol{x}_t) \in [t_l, T] \times \mathbb{R}^m$,* ***minimise the functional***

$$J(t, \boldsymbol{x}_t, \boldsymbol{u}_t) \triangleq \int_t^T L(s, \boldsymbol{x}_s, \boldsymbol{u}_s)\, ds + K(\boldsymbol{x}_T) \tag{44}$$

subject to the state equation

$$\dot{\boldsymbol{x}}_s = \frac{d\boldsymbol{x}_s}{ds} = \boldsymbol{b}(s, \boldsymbol{x}_s, \boldsymbol{u}_s), \quad \boldsymbol{x}_t \text{ given}, \tag{45}$$

and control constraint $\boldsymbol{u}_s \in U(s)$, $s \in [t, T]$.

Definition E.2 (Value Function)**.** The **value function** of Problem E.5 is

$$V(t, \boldsymbol{x}_t) \triangleq \inf_{\boldsymbol{u}_{[t,T]}} J(t, \boldsymbol{x}_t, \boldsymbol{u}_t), \tag{46}$$

where the infimum ranges over controls with $\boldsymbol{u}_s \in U(s)$, $s \in [t, T]$, and $\boldsymbol{x}_s$ solves (45).

[4] Specifically: joint continuity, locally Lipschitz in $\boldsymbol{x}$ with a t-dependent, square-integrable modulus, uniformly over $\boldsymbol{u}_t$, and of at most linear growth in $|\boldsymbol{x}|$ (Berkovitz & Medhin, 2012, Assumption 12.2.1(i)–(iv), pp. 338–339). Parts (v)–(vi) of that assumption, concerning the terminal manifold and cost, are immediate from our choice of a fixed terminal time T; part (vii), concerning existence of an optimal control, is not required for the verification Lemma E.7 below.

Hamilton–Jacobi–Bellman equation: The following theorem states that the value function of Definition E.2 satisfies the associated Hamilton–Jacobi–Bellman (HJB) equation. This is a well-known result, following from the dynamic programming principle; we omit the proof here and refer the reader to Bellman, 1957, Berkovitz & Medhin, 2012 (§6.2, pp. 150–154), and Liberzon, 2012 (§5.1.3, pp. 161–162).

Theorem E.6 (HJB Equation, after Bellman, 1957; Berkovitz & Medhin, 2012; Liberzon, 2012)**.** *The value function V of Definition E.2 satisfies*

$$-V_t(t,\boldsymbol{x}_t) = \inf_{\boldsymbol{u}_t \in U(t)} \Big\{ L(t,\boldsymbol{x}_t,\boldsymbol{u}_t) + \big\langle V_{\boldsymbol{x}_t}(t,\boldsymbol{x}_t),\, \boldsymbol{b}(t,\boldsymbol{x}_t,\boldsymbol{u}_t) \big\rangle \Big\}, \quad \forall\, (t,\boldsymbol{x}_t) \in [t_l,T)\times\mathbb{R}^m, \tag{47}$$

with boundary condition $V(T,\boldsymbol{x}_T) = K(\boldsymbol{x}_T)$.

Sufficient condition for optimality: The following verification lemma gives conditions under which a C^1 solution of the HJB equation coincides with the value function, and a control satisfying the associated stationarity condition is optimal. We state the classical smooth specialisation of the argument; see Bellman, 1957, Berkovitz & Medhin, 2012 (Thm. 12.7.1, pp. 359–360, there stated more generally for Lipschitz continuous value functions via lower Dini derivatives), and Liberzon, 2012 (§5.1.4, pp. 165–166, the classical C^1 statement we specialise to here) for the proof.

Lemma E.7 (Sufficient Condition for Optimality, after Bellman, 1957; Berkovitz & Medhin, 2012; Liberzon, 2012)**.** *Suppose $\hat{V} \in C^1([t_l,T]\times\mathbb{R}^m)$ satisfies the HJB equation (47) and boundary condition of Theorem E.6. Suppose there exists a control $\hat{\boldsymbol{u}}$ with $\hat{\boldsymbol{u}}_t \in U(t)$ and trajectory $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{x}}_{t_l}$ given, such that*

$$L(t,\hat{\boldsymbol{x}}_t,\hat{\boldsymbol{u}}_t) + \big\langle \hat{V}_{\boldsymbol{x}_t}(t,\hat{\boldsymbol{x}}_t), \boldsymbol{b}(t,\hat{\boldsymbol{x}}_t,\hat{\boldsymbol{u}}_t)\big\rangle = \min_{\boldsymbol{u}_t\in U(t)} \Big\{ L(t,\hat{\boldsymbol{x}}_t,\boldsymbol{u}_t) + \big\langle \hat{V}_{\boldsymbol{x}_t}(t,\hat{\boldsymbol{x}}_t), \boldsymbol{b}(t,\hat{\boldsymbol{x}}_t,\boldsymbol{u}_t)\big\rangle \Big\}$$

for every $t\in[t_l,T]$. Then $\hat{V}(t_l,\boldsymbol{x}_{t_l}) = V(t_l,\boldsymbol{x}_{t_l})$ and $\hat{\boldsymbol{u}}$ is an optimal control.

Reverse-time optimal control problem: We now consider the analogous problem in which the objective is maximised and running reward accumulated *since* a fixed reference time t_l, rather than cost-to-go accrued *until* a fixed terminal time. This is the formulation our subsequent CFG derivation is stated on directly.

Problem E.8 (Reverse-Time Optimal Control Problem)**.** *Fix $t_l < T$. Let $\boldsymbol{x}_t \in \mathbb{R}^m$ for $t\in[t_l,T]$, let $\boldsymbol{u}_t \in U(t) \subseteq \mathbb{R}^p$, $\boldsymbol{b}$, L be as in Problem E.5 (with the same regularity), and let $K:\mathbb{R}\times\mathbb{R}^m\to\mathbb{R}$ be an initial reward. For $(t,\boldsymbol{x}_t)\in[t_l,T]\times\mathbb{R}^m$,* ***maximise the functional***

$$J(t,\boldsymbol{x}_t,\boldsymbol{u}_t) \triangleq \int_{t_l}^{t} L(s,\boldsymbol{x}_s,\boldsymbol{u}_s)\,ds + K(t_l,\boldsymbol{x}_{t_l}) \tag{48}$$

subject to the state equation

$$\dot{\boldsymbol{x}}_s = \frac{d\boldsymbol{x}_s}{ds} = \boldsymbol{b}(s,\boldsymbol{x}_s,\boldsymbol{u}_s), \quad \boldsymbol{x}_t \text{ given}, \tag{49}$$

and control constraint $\boldsymbol{u}_s \in U(s)$, $s\in[t_l,t]$.

Definition E.3 (Reverse-Time Value Function)**.** The **value function** of Problem E.8 is

$$\hat{V}(t,\boldsymbol{x}_t) \triangleq \sup_{\boldsymbol{u}_{[t_l,t]}} J(t,\boldsymbol{x}_t,\boldsymbol{u}_t), \tag{50}$$

where the supremum ranges over controls with $\boldsymbol{u}_s\in U(s)$, $s\in[t_l,t]$, and $\boldsymbol{x}_s$ solves (49).

Lemma E.9 (Reverse-Time HJB Equation)**.** *The value function $\hat{V}$ of Definition E.3 satisfies*

$$\hat{V}_t(t,\boldsymbol{x}_t) = \sup_{\boldsymbol{u}_t\in U(t)} \Big\{ L(t,\boldsymbol{x}_t,\boldsymbol{u}_t) - \big\langle \hat{V}_{\boldsymbol{x}_t}(t,\boldsymbol{x}_t),\, \boldsymbol{b}(t,\boldsymbol{x}_t,\boldsymbol{u}_t)\big\rangle \Big\}, \quad \forall\, t\in(t_l,T], \tag{51}$$

with boundary condition $\hat{V}(t_l,\boldsymbol{x}_{t_l}) = K(t_l,\boldsymbol{x}_{t_l})$.

Proof. Define the change of variables at the level of the free (evaluation) variable,

$$t \triangleq T + t_l - \tau, \qquad \tilde{\boldsymbol{x}}_\tau \triangleq \boldsymbol{x}_{T+t_l-\tau}, \qquad \tilde{\boldsymbol{u}}_\tau \triangleq \boldsymbol{u}_{T+t_l-\tau},$$

and apply the same substitution to the dummy (integration) variable, writing ψ for its image: $s \triangleq T + t_l - \psi$, so that $\tilde{\boldsymbol{x}}_\psi = \boldsymbol{x}_s$ and $\tilde{\boldsymbol{u}}_\psi = \boldsymbol{u}_s$ are the same trajectory and control functions evaluated at ψ rather than τ. This maps $[t_l, T]$ to itself, exchanging its endpoints: $t = t_l \iff \tau = T$, $t = T \iff \tau = t_l$, and likewise $s = t_l \iff \psi = T$, $s = T \iff \psi = t_l$.

State equation. Writing $s = T + t_l - \psi$, the chain rule gives

$$\frac{d\tilde{\boldsymbol{x}}_\psi}{d\psi} = \frac{d\boldsymbol{x}_s}{ds} \cdot \frac{ds}{d\psi} = -\boldsymbol{b}(s, \boldsymbol{x}_s, \boldsymbol{u}_s) = -\boldsymbol{b}(T + t_l - \psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi) \triangleq \tilde{\boldsymbol{b}}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi).$$

Functional. Since $s \mapsto \psi$ reverses orientation on $[t_l, T]$, substituting $s = T + t_l - \psi$, $ds = -d\psi$ into (48) and flipping the limits of integration ($s = t_l \mapsto \psi = T$, $s = t \mapsto \psi = T + t_l - t = \tau$),

$$\int_{t_l}^{t} L(s, \boldsymbol{x}_s, \boldsymbol{u}_s)\, ds = \int_{\tau}^{T} L(T + t_l - \psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)\, d\psi \triangleq \int_{\tau}^{T} \tilde{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)\, d\psi,$$

where $\tilde{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi) \triangleq L(T + t_l - \psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)$. The bounds of integration now match Problem E.5's: the free variable τ as the lower limit, the fixed time T as the upper limit. Using $\boldsymbol{x}_{t_l} = \tilde{\boldsymbol{x}}_T$,

$$J(t, \boldsymbol{x}_t, \boldsymbol{u}_t) = \int_{\tau}^{T} \tilde{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)\, d\psi + K(t_l, \tilde{\boldsymbol{x}}_T).$$

As $\boldsymbol{u}_s \in U(s) \iff \tilde{\boldsymbol{u}}_\psi \in U(T + t_l - \psi) \triangleq \tilde{U}(\psi)$, taking the supremum over $\boldsymbol{u}_{[t_l, t]}$ is equivalent to taking it over $\tilde{\boldsymbol{u}}_{[\tau, T]}$, so, writing $\tilde{V}(\tau, \tilde{\boldsymbol{x}}_\tau) \triangleq \hat{V}(t, \boldsymbol{x}_t)$,

$$\tilde{V}(\tau, \tilde{\boldsymbol{x}}_\tau) = \sup_{\tilde{\boldsymbol{u}}_{[\tau,T]}} \left[\int_{\tau}^{T} \tilde{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)\, d\psi + K(t_l, \tilde{\boldsymbol{x}}_T) \right].$$

Min-max flip. Define $\bar{V} \triangleq -\tilde{V}$, $\bar{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi) \triangleq -\tilde{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)$, and $\bar{K}(\tilde{\boldsymbol{x}}_T) \triangleq -K(t_l, \tilde{\boldsymbol{x}}_T)$. Then $\sup\{\cdot\} = -\inf\{-\cdot\}$ gives

$$\bar{V}(\tau, \tilde{\boldsymbol{x}}_\tau) = \inf_{\tilde{\boldsymbol{u}}_{[\tau,T]}} \left[\int_{\tau}^{T} \bar{L}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)\, d\psi + \bar{K}(\tilde{\boldsymbol{x}}_T) \right],$$

subject to $\dot{\tilde{\boldsymbol{x}}}_\psi = \tilde{\boldsymbol{b}}(\psi, \tilde{\boldsymbol{x}}_\psi, \tilde{\boldsymbol{u}}_\psi)$ on $[\tau, T]$. This is precisely Problem E.5 with t replaced by τ (and its dummy variable s correspondingly by ψ), and $L, K, \boldsymbol{b}, U$ replaced by $\bar{L}, \bar{K}, \tilde{\boldsymbol{b}}, \tilde{U}$. By Theorem E.6, $\bar{V}$ therefore satisfies (47) under this substitution:

$$-\bar{V}_\tau(\tau, \tilde{\boldsymbol{x}}_\tau) = \inf_{\boldsymbol{v} \in \tilde{U}(\tau)} \left\{ \bar{L}(\tau, \tilde{\boldsymbol{x}}_\tau, \boldsymbol{v}) + \left\langle \bar{V}_{\tilde{\boldsymbol{x}}_\tau}(\tau, \tilde{\boldsymbol{x}}_\tau), \tilde{\boldsymbol{b}}(\tau, \tilde{\boldsymbol{x}}_\tau, \boldsymbol{v}) \right\rangle \right\}, \quad \tau \in [t_l, T).$$

Back-substitution. Since $\bar{V}(\tau, \tilde{\boldsymbol{x}}_\tau) = -\hat{V}(t, \boldsymbol{x}_t)$ with $t = T + t_l - \tau$, the chain rule gives $\bar{V}_\tau(\tau, \tilde{\boldsymbol{x}}_\tau) = \hat{V}_t(t, \boldsymbol{x}_t)$ and $\bar{V}_{\tilde{\boldsymbol{x}}_\tau}(\tau, \tilde{\boldsymbol{x}}_\tau) = -\hat{V}_{\boldsymbol{x}_t}(t, \boldsymbol{x}_t)$. Substituting these together with $\bar{L}(\tau, \tilde{\boldsymbol{x}}_\tau, \boldsymbol{v}) = -L(t, \boldsymbol{x}_t, \boldsymbol{v})$, $\tilde{\boldsymbol{b}}(\tau, \tilde{\boldsymbol{x}}_\tau, \boldsymbol{v}) = -\boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{v})$, and $\tilde{U}(\tau) = U(t)$,

$$-\hat{V}_t(t, \boldsymbol{x}_t) = \inf_{\boldsymbol{v} \in U(t)} \left\{ -L(t, \boldsymbol{x}_t, \boldsymbol{v}) + \left\langle \hat{V}_{\boldsymbol{x}_t}(t, \boldsymbol{x}_t), \boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{v}) \right\rangle \right\}, \quad t \in (t_l, T],$$

using $\langle -\boldsymbol{a}, -\boldsymbol{b} \rangle = \langle \boldsymbol{a}, \boldsymbol{b} \rangle$. Multiplying by -1 and using $-\inf\{-\cdot\} = \sup\{\cdot\}$ gives (51). The boundary condition follows from $\hat{V}(t_l, \boldsymbol{x}_{t_l}) = -\bar{V}(T, \tilde{\boldsymbol{x}}_T) = -\bar{K}(\tilde{\boldsymbol{x}}_T) = K(t_l, \boldsymbol{x}_{t_l})$. □

Sufficient condition for optimality (reverse-time): The following lemma is the mirror image of Lemma E.7, obtained by transporting it through the same change of variables used to prove Lemma E.9. We first record the single-point statement that falls out of a direct transport, then extend it to hold on the whole interval.

Lemma E.10 (Reverse-Time Sufficient Condition for Optimality). *Suppose $\check{V} \in C^1([t_l, T] \times \mathbb{R}^m)$ satisfies the reverse-time HJB equation (51) and boundary condition of Lemma E.9. Suppose there exists a control $\check{\boldsymbol{u}}$ with $\check{\boldsymbol{u}}_t \in U(t)$ and trajectory $\check{\boldsymbol{x}}$, $\check{\boldsymbol{x}}_T$ given, such that*

$$L(t, \check{\boldsymbol{x}}_t, \check{\boldsymbol{u}}_t) - \left\langle \check{V}_{\boldsymbol{x}_t}(t, \check{\boldsymbol{x}}_t), \boldsymbol{b}(t, \check{\boldsymbol{x}}_t, \check{\boldsymbol{u}}_t) \right\rangle = \max_{\boldsymbol{u}_t \in U(t)} \left\{ L(t, \check{\boldsymbol{x}}_t, \boldsymbol{u}_t) - \left\langle \check{V}_{\boldsymbol{x}_t}(t, \check{\boldsymbol{x}}_t), \boldsymbol{b}(t, \check{\boldsymbol{x}}_t, \boldsymbol{u}_t) \right\rangle \right\}$$

for every $t \in [t_l, T]$. Then $\check{V}(T, \boldsymbol{x}_T) = \hat{V}(T, \boldsymbol{x}_T)$ and $\check{\boldsymbol{u}}$ is an optimal control.

Proof. Apply the change of variables from the proof of Lemma E.9, at the level of the free variable alone: $t \triangleq T + t_l - \tau$, $\tilde{\boldsymbol{x}}_\tau \triangleq \boldsymbol{x}_{T+t_l-\tau}$, $\tilde{\boldsymbol{u}}_\tau \triangleq \boldsymbol{u}_{T+t_l-\tau}$ (no dummy variable is needed, as the hypothesis to be transported is a pointwise condition rather than an integral). As shown there, the reverse-time HJB equation and boundary condition transport $\check{V}$ into a candidate $\check{\tilde{V}} \triangleq -\check{V}(T + t_l - \tau, \tilde{\boldsymbol{x}}_\tau)$ satisfying the (forward) HJB equation (47) and boundary condition of Theorem E.6, for the problem with $(L, K, \boldsymbol{b}, U)$ replaced by $(\bar{L}, \bar{K}, \tilde{\boldsymbol{b}}, \tilde{U})$. Since $t \mapsto \tau$ is a bijection of $[t_l, T]$ onto itself, "for every $t \in [t_l, T]$" transports directly to "for every $\tau \in [t_l, T]$"; applying the same substitution to the stationarity hypothesis above (using $\langle -\boldsymbol{a}, -\boldsymbol{b}\rangle = \langle \boldsymbol{a}, \boldsymbol{b}\rangle$ and $\max\{-f\} = -\min\{f\}$, as in that proof) shows that $\check{\tilde{\boldsymbol{u}}}_\tau \triangleq \check{\boldsymbol{u}}_{T+t_l-\tau}$ and $\check{\tilde{\boldsymbol{x}}}_\tau \triangleq \check{\boldsymbol{x}}_{T+t_l-\tau}$ satisfy

$$\bar{L}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau, \check{\tilde{\boldsymbol{u}}}_\tau) + \big\langle \check{\tilde{V}}_{\tilde{\boldsymbol{x}}_\tau}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau), \tilde{\boldsymbol{b}}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau, \check{\tilde{\boldsymbol{u}}}_\tau)\big\rangle = \min_{\boldsymbol{v} \in \tilde{U}(\tau)} \Big\{ \bar{L}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau, \boldsymbol{v}) + \big\langle \check{\tilde{V}}_{\tilde{\boldsymbol{x}}_\tau}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau), \tilde{\boldsymbol{b}}(\tau, \check{\tilde{\boldsymbol{x}}}_\tau, \boldsymbol{v})\big\rangle \Big\}$$

for every $\tau \in [t_l, T]$–precisely the hypothesis of Lemma E.7, with $\check{\tilde{\boldsymbol{x}}}_{t_l} = \check{\boldsymbol{x}}_T$ given. Lemma E.7 therefore gives $\check{\tilde{V}}(t_l, \check{\tilde{\boldsymbol{x}}}_{t_l}) = \bar{V}(t_l, \check{\tilde{\boldsymbol{x}}}_{t_l})$ and that $\check{\tilde{\boldsymbol{u}}}$ is optimal for the transported problem. Translating back via $\tau = t_l \iff t = T$ and $\check{\tilde{V}}(t_l, \check{\tilde{\boldsymbol{x}}}_{t_l}) = -\check{V}(T, \boldsymbol{x}_T)$, $\bar{V}(t_l, \tilde{\boldsymbol{x}}_{t_l}) = -\hat{V}(T, \boldsymbol{x}_T)$ gives $\check{V}(T, \boldsymbol{x}_T) = \hat{V}(T, \boldsymbol{x}_T)$, and optimality of $\check{\tilde{\boldsymbol{u}}}$ for the transported problem translates back to optimality of $\tilde{\boldsymbol{u}}$ for Problem E.8. □

CFG-style guidance as optimal control: We now show that a guidance scheme built from Tweedie's estimate and the class-conditional score-ratio solves the reverse-time optimal control problem of Definition E.3, for an entire family of running rewards indexed by a free schedule $r(s)$.

Theorem E.11 (CFG-Style Guidance Solves the Reverse-Time HJB Equation)**.** *Fix $\xi \in \mathcal{Y}$ and $t_l \in (0, T)$ (so that $\beta_s^2 > 0$ throughout $[t_l, T]$), let $r : [t_l, T] \to (0, \infty)$ be continuous, and suppose $U(s) \subseteq \mathbb{R}^m$ is nonempty, closed, and convex for every $s \in [t_l, T]$.*[5] *Define*

$$\phi(s, \boldsymbol{x}_s) \triangleq \log\left(\frac{p(\boldsymbol{x}_s \mid \xi)}{p(\boldsymbol{x}_s)}\right), \tag{52}$$

and consider Problem E.8 with dynamics

$$\dot{\boldsymbol{x}}_s = \boldsymbol{b}(s, \boldsymbol{x}_s, \boldsymbol{u}_s) \triangleq f(s)\boldsymbol{x}_s - \frac{1}{2}\frac{g(s)^2}{\beta_s^2}\big(\alpha_s \boldsymbol{u}_s - \boldsymbol{x}_s\big), \tag{53}$$

running reward

$$\begin{aligned} L(s, \boldsymbol{x}_s, \boldsymbol{u}_s) \triangleq\ & -r(s)\,\|\boldsymbol{u}_s - \hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)\|_2^2 + \frac{1}{2}g(s)^2\big\langle \nabla_{\boldsymbol{x}_s}\phi(s, \boldsymbol{x}_s), \nabla_{\boldsymbol{x}_s}\log p(\boldsymbol{x}_s \mid \xi)\big\rangle \\ & + r(s)\,\|\Upsilon(s, \boldsymbol{x}_s)\|_2^2 - \frac{1}{2}\frac{g(s)^2\alpha_s}{\beta_s^2}\big\langle \nabla_{\boldsymbol{x}_s}\phi(s, \boldsymbol{x}_s), \Upsilon(s, \boldsymbol{x}_s)\big\rangle \\ & + \frac{1}{2}g(s)^2\Delta_{\boldsymbol{x}_s}\phi(s, \boldsymbol{x}_s), \end{aligned} \tag{54}$$

and initial reward $K(t_l, \boldsymbol{x}_{t_l}) \triangleq \phi(t_l, \boldsymbol{x}_{t_l})$, where

$$\Upsilon(s, \boldsymbol{x}_s) \triangleq \mathit{proj}_{U(s) - \hat{\boldsymbol{x}}_0(s, \boldsymbol{x}_s, \Omega_Y)}\left\{\frac{1}{4}\frac{g(s)^2\alpha_s}{r(s)\beta_s^2}\nabla_{\boldsymbol{x}_s}\phi(s, \boldsymbol{x}_s)\right\}. \tag{55}$$

Then the value function of this problem is $\hat{V}(t, \boldsymbol{x}_t) = \phi(t, \boldsymbol{x}_t)$ for every $(t, \boldsymbol{x}_t) \in [t_l, T] \times \mathbb{R}^m$, and the control

$$\boldsymbol{u}_t^\star \triangleq \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y) + \Upsilon(t, \boldsymbol{x}_t) \tag{56}$$

is optimal.

[5] When $U(s)$ is bounded, the regularity of Problem E.5 (joint continuity, local Lipschitz in $\boldsymbol{x}$, linear growth) is understood to hold on the set; when $U(s)$ is unbounded, as in Corollary E.12, the regularity of Problem E.5 (joint continuity, local Lipschitz in $\boldsymbol{x}$, linear growth) is understood to hold on every compact subset of $\mathbb{R}^m$ rather than globally, as Assumption 12.2.1(iii) of Berkovitz & Medhin (2012) requires.

Proof. We verify the hypotheses of Lemma E.10 for $\check{V} \triangleq \phi$ and $\check{\boldsymbol{u}} \triangleq \boldsymbol{u}^\star$.

Regularity. Since $t_l > 0$, $\beta_t^2 > 0$ for every $t \in [t_l, T]$, so by the argument preceding Theorem E.3 both $p(\cdot \mid \xi)$ and $p(\cdot)$ are positive-everywhere and $C^{1,\infty}$ on $[t_l, T] \times \mathbb{R}^m$; hence $\phi \in C^{1,\infty}([t_l, T] \times \mathbb{R}^m)$.

Stationarity. Fix $(t, \boldsymbol{x}_t)$ and write $\hat{\boldsymbol{x}}_0 \triangleq \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y)$, $\nabla\phi \triangleq \nabla_{\boldsymbol{x}_t}\phi(t, \boldsymbol{x}_t)$ for brevity. As $\boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{u}_t)$ is affine in $\boldsymbol{u}_t$ and the only $\boldsymbol{u}_t$-dependence of $L(t, \boldsymbol{x}_t, \boldsymbol{u}_t)$ is through $-r(t)\|\boldsymbol{u}_t - \hat{\boldsymbol{x}}_0\|_2^2$ (every remaining term of (54) depends on $(t, \boldsymbol{x}_t)$ alone), the map

$$H(\boldsymbol{u}_t) \triangleq L(t, \boldsymbol{x}_t, \boldsymbol{u}_t) - \langle \nabla\phi, \boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{u}_t)\rangle$$

is, as a function of $\boldsymbol{u}_t$, a strictly concave quadratic (since $r(t) > 0$):

$$H(\boldsymbol{u}_t) = -r(t)\|\boldsymbol{u}_t - \hat{\boldsymbol{x}}_0\|_2^2 + \frac{1}{2}\frac{g(t)^2\alpha_t}{\beta_t^2}\langle \nabla\phi, \boldsymbol{u}_t\rangle + C(t, \boldsymbol{x}_t)$$

for some $\boldsymbol{u}_t$-independent $C(t, \boldsymbol{x}_t)$. Completing the square, H is maximised over $\mathbb{R}^m$ uniquely at $\boldsymbol{u}_0 \triangleq \hat{\boldsymbol{x}}_0 + \frac{1}{4}\frac{g(t)^2\alpha_t}{r(t)\beta_t^2}\nabla\phi$. As $U(t)$ is nonempty, closed, and convex, the Hilbert-space projection theorem gives that H restricted to $U(t)$ is maximised uniquely at

$$\mathrm{proj}_{U(t)}\{\boldsymbol{u}_0\} = \hat{\boldsymbol{x}}_0 + \mathrm{proj}_{U(t)-\hat{\boldsymbol{x}}_0}\{\boldsymbol{u}_0 - \hat{\boldsymbol{x}}_0\} = \hat{\boldsymbol{x}}_0 + \Upsilon(t, \boldsymbol{x}_t) = \boldsymbol{u}_t^\star.$$

HJB equation. It remains to show $H(\boldsymbol{u}_t^\star) = \phi_t(t, \boldsymbol{x}_t)$, i.e. that (51) holds at $\check{V} = \phi$, $\boldsymbol{u}_t = \boldsymbol{u}_t^\star$. Writing $A \triangleq \log p(\boldsymbol{x}_t \mid \xi)$, $B \triangleq \log p(\boldsymbol{x}_t)$ (so $\phi = A - B$), the Kolmogorov forward equation for the canonical image diffusion of Definition E.1 gives, for either log-density,

$$\partial_t \log p = -f(t)\left(m + \langle \boldsymbol{x}_t, \nabla \log p\rangle\right) + \frac{1}{2}g(t)^2\left(\|\nabla \log p\|_2^2 + \Delta \log p\right),$$

and subtracting the two copies (for A and B; the $-f(t)m$ terms cancel, and $\|\nabla A\|_2^2 - \|\nabla B\|_2^2 = \langle \nabla A - \nabla B, \nabla A + \nabla B\rangle = \langle \nabla\phi, \nabla A\rangle + \langle \nabla\phi, \nabla B\rangle$ since $\nabla A - \nabla B = \nabla\phi$) gives

$$\begin{aligned}\phi_t(t, \boldsymbol{x}_t) = &- f(t)\langle \boldsymbol{x}_t, \nabla\phi\rangle + \frac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t \mid \xi)\rangle + \frac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t)\rangle \\ &+ \frac{1}{2}g(t)^2\Delta\phi. \end{aligned} \tag{57}$$

On the other hand, substituting $\boldsymbol{u}_t^\star = \hat{\boldsymbol{x}}_0 + \Upsilon$ into $\boldsymbol{b}$ and using Tweedie's formula (Theorem E.3) in the form $\alpha_t\hat{\boldsymbol{x}}_0 = \beta_t^2\nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t) + \boldsymbol{x}_t$ (which follows from combining Theorem E.3 with (33) at $\xi = \Omega_Y$), the $\mp\frac{1}{2}\frac{g(t)^2}{\beta_t^2}\boldsymbol{x}_t$ terms cancel and

$$\langle \nabla\phi, \boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{u}_t^\star)\rangle = f(t)\langle \boldsymbol{x}_t, \nabla\phi\rangle - \frac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t)\rangle - \frac{1}{2}\frac{g(t)^2\alpha_t}{\beta_t^2}\langle \nabla\phi, \Upsilon\rangle.$$

Substituting $\boldsymbol{u}_t = \boldsymbol{u}_t^\star$ into (54) (so $\boldsymbol{u}_t^\star - \hat{\boldsymbol{x}}_0 = \Upsilon$, and the $\mp r(t)\|\Upsilon\|_2^2$ terms cancel) and subtracting the expression above,

$$\begin{aligned} H(\boldsymbol{u}_t^\star) &= L(t, \boldsymbol{x}_t, \boldsymbol{u}_t^\star) - \langle \nabla\phi, \boldsymbol{b}(t, \boldsymbol{x}_t, \boldsymbol{u}_t^\star)\rangle \\ &= \tfrac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t \mid \xi)\rangle - \tfrac{1}{2}\tfrac{g(t)^2\alpha_t}{\beta_t^2}\langle \nabla\phi, \Upsilon\rangle + \tfrac{1}{2}g(t)^2\Delta\phi \\ &\quad - f(t)\langle \boldsymbol{x}_t, \nabla\phi\rangle + \tfrac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t)\rangle + \tfrac{1}{2}\tfrac{g(t)^2\alpha_t}{\beta_t^2}\langle \nabla\phi, \Upsilon\rangle \\ &= -f(t)\langle \boldsymbol{x}_t, \nabla\phi\rangle + \tfrac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t \mid \xi)\rangle + \tfrac{1}{2}g(t)^2\langle \nabla\phi, \nabla_{\boldsymbol{x}_t}\log p(\boldsymbol{x}_t)\rangle \\ &\quad + \tfrac{1}{2}g(t)^2\Delta\phi, \end{aligned}$$

where the $\mp\frac{1}{2}\frac{g(t)^2\alpha_t}{\beta_t^2}\langle \nabla\phi, \Upsilon\rangle$ terms cancel. This is precisely (57), so $H(\boldsymbol{u}_t^\star) = \phi_t(t, \boldsymbol{x}_t)$, and since $\boldsymbol{u}_t^\star$ maximises H over $U(t)$, (51) holds at $\check{V} = \phi$ for every $t \in [t_l, T]$.

Boundary condition. $\check{V}(t_l, \boldsymbol{x}_{t_l}) = \phi(t_l, \boldsymbol{x}_{t_l}) = K(t_l, \boldsymbol{x}_{t_l})$ by definition of K.

The hypotheses of Lemma E.10 hold on $[t_l, T]$ with $\check{V} = \phi$, $\check{\boldsymbol{u}} = \boldsymbol{u}^\star$. Since $\hat{V}(t, \boldsymbol{x}_t)$ depends, by Definition E.3, only on $L, \boldsymbol{b}, K$ restricted to $[t_l, t]$, these hypotheses (inherited by restriction) give $\phi(t, \boldsymbol{x}_t) = \hat{V}(t, \boldsymbol{x}_t)$ for every $t \in (t_l, T]$, not merely $t = T$; equality at $t = t_l$ follows directly from the boundary condition. Hence $\phi = \hat{V}$ on $[t_l, T] \times \mathbb{R}^m$, and $\boldsymbol{u}^\star$ is optimal. □

Corollary E.12 (Unconstrained CFG-Style Guidance)**.** *Under the hypotheses of Theorem E.11, if* $U(s) = \mathbb{R}^m$ *for every* $s \in [t_l, T]$*, then* $\Upsilon(t, \boldsymbol{x}_t) = \frac{1}{4}\frac{g(t)^2\alpha_t}{r(t)\beta_t^2}\nabla_{\boldsymbol{x}_t}\phi(t, \boldsymbol{x}_t)$*, and the optimal control (56) reduces to*

$$\boldsymbol{u}_t^\star = \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y) + \frac{1}{4}\frac{g(t)^2\alpha_t}{r(t)\beta_t^2}\nabla_{\boldsymbol{x}_t}\log\left(\frac{p(\boldsymbol{x}_t \mid \xi)}{p(\boldsymbol{x}_t)}\right).$$

Proof. $U(t) - \hat{\boldsymbol{x}}_0(t, \boldsymbol{x}_t, \Omega_Y) = \mathbb{R}^m$, and projection onto $\mathbb{R}^m$ is the identity map. □