

CFG-Ctrl: Control-Based Classifier-Free Diffusion Guidance

Hanyang Wang*, Yiyang Liu*, Jiawei Chi, Fangfu Liu, Ran Xue, Yueqi Duan[†]
Tsinghua University

Abstract

Classifier-Free Guidance (CFG) has emerged as a central approach for enhancing semantic alignment in flow-based diffusion models. In this paper, we explore a unified framework called **CFG-Ctrl**, which reinterprets CFG as a control applied to the first-order continuous-time generative flow, using the conditional-unconditional discrepancy as an error signal to adjust the velocity field. From this perspective, we summarize vanilla CFG as a proportional controller (P-control) with fixed gain, and typical follow-up variants develop extended control-law designs derived from it. However, existing methods mainly rely on linear control, inherently leading to instability, overshooting, and degraded semantic fidelity especially on large guidance scales. To address this, we introduce Sliding Mode Control CFG (**SMC-CFG**), which enforces the generative flow toward a rapidly convergent sliding manifold. Specifically, we define an exponential sliding mode surface over the semantic prediction error and introduce a switching control term to establish nonlinear feedback-guided correction. Moreover, we provide a Lyapunov stability analysis to theoretically support finite-time convergence. Experiments across text-to-image generation models including Stable Diffusion 3.5, Flux, and Qwen-Image demonstrate that SMC-CFG outperforms standard CFG in semantic alignment and enhances robustness across a wide range of guidance scales. Project Page: <https://hanyang-21.github.io/CFG-Ctrl>.

1. Introduction

Diffusion models [16, 44, 45] have recently achieved state-of-the-art performance in high-fidelity image synthesis across diverse domains [34, 36]. Building on the similar probabilistic formulation, flow matching [24, 31] provides a more straightforward alternative by directly estimating deterministic velocity fields, realizing stable training and faster sampling than diffusion [9, 11]. These flow-based methods have demonstrated strong capability across text-to-image [8, 20], text-to-video [18, 46, 54], and other visual generation applications [21, 47, 60].

*Equal contribution. [†]Corresponding author.

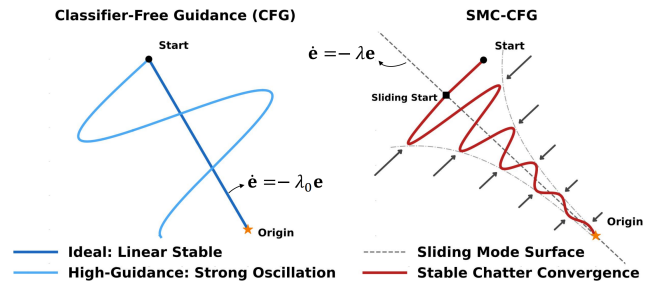


Figure 1. **Phase diagram in the e - \dot{e} plane.** We schematically illustrate the convergence patterns of CFG and the proposed SMC-CFG. *Left:* CFG’s ideal linear convergence trajectory and the strong oscillatory divergence under high guidance scales. *Right:* the proposed SMC-CFG, through a switching-forcing mechanism, drives the system states toward the sliding mode surface governed by parameter λ , achieving robust and rapid convergence.

A key technique widely adopted in diffusion models is Classifier-Free Guidance (CFG) [15], which enhances semantic alignment between the generated sample and the input condition. Previous studies commonly interpret CFG as a linear extrapolation between unconditional and conditional predictions within deterministic diffusion flows [40]. While this perspective offers an intuitive interpretation, the resulting linear extrapolation can distort the generative trajectory from the learned data manifold, leading to oversaturated colors, warped structures, and strong sensitivity to the guidance scale [5]. Several improved methods have been proposed to alleviate these issues, including linear recomposition [52], orthogonal decomposition [9, 38], and dynamic weighting schedules [40, 48].

We observe that the discrepancy between the conditional and unconditional velocity predictions gradually diminishes in diffusion flow progress, effectively serving as a natural error signal. This observation motivates us to reinterpret CFG not as a static extrapolation rule, but as a form of feedback control applied to the latent generative flow. Based on this observation, we explore a unified theoretical framework called **CFG-Ctrl** for Classifier-Free Guidance in flow matching diffusion. Under this CFG-Ctrl paradigm, the standard CFG corresponds to a proportional controller (P-control) that amplifies the semantic error with a fixed gain and feeds it back into the system, while existing CFG

variants can be regarded as alternative designs of feedback control laws. However, most of these methods rely on approximately linear control laws for feedback, which cannot ensure stable convergence when the underlying generative dynamics become highly nonlinear—particularly as model capacity increases or the guidance scale becomes large as shown in Figure 1 (left).

To address this, we further propose Sliding Mode Control CFG (**SMC-CFG**), a control-based guidance mechanism that directs the flow trajectory onto a rapidly converging sliding mode surface. This design draws on the proven success of Sliding Mode Control (SMC) [7, 58] in stabilizing nonlinear dynamical systems. As shown in Figure 1 (right), our approach constructs a sliding mode surface over the semantic prediction error, corresponding to the gray dashed line in the figure. We also introduce a switching control term that enforces nonlinear, feedback-driven corrective force, which are represented by the arrows at both sides of the convergence curve. This design adaptively regulates the evolution of the flow trajectory and preserves stability even under strong guidance. To theoretically substantiate convergence, we provide a Lyapunov stability analysis based on the principle of monotonically decreasing energy, demonstrating that SMC-CFG supports finite-time convergence toward the desired semantic manifold. Extensive experiments on three state-of-the-art text-to-image (T2I) models show that SMC-CFG consistently improves semantic fidelity, reduces visual artifacts, and maintains robustness across multiple semantic and perceptual metrics. Our contributions are summarized as follows:

- We explore CFG-Ctrl, a novel theoretical framework for Classifier-Free Guidance in flow matching models grounded in control theory, unifying the systematic interpretation of diverse guidance strategies.
- We propose SMC-CFG, a sliding-mode-based nonlinear feedback controller for flow models, and prove finite-time convergence under Lyapunov stability analysis.
- Extensive experiments across multiple diffusion backbones demonstrate that SMC-CFG achieves superior semantic fidelity, visual coherence, and robustness, particularly under high guidance scales.

2. Related Work

Diffusion and Flow Matching. Diffusion models [16, 44, 45] have garnered significant attention in recent years as a class of generative models that iteratively transform simple distributions into more complex ones, ultimately generating high-quality samples. Early diffusion models define a forward diffusion process, where noise is gradually added to data samples, typically starting from a simple prior such as an isotropic Gaussian. The reverse process is then learned by training a neural network to estimate the score function of the data distribution [36, 45], enabling the model to pro-

gressively recover the original data. More recently, flow matching [24, 31] has been proposed to model the transformation process via a learned velocity field, which simplifies the generative formulation and leads to better empirical performance. This paradigm has been widely adopted in large-scale foundation models across multiple domains, including image generation [20, 49, 50], video generation [12, 26, 28], and 3D content generation [29, 30, 56], demonstrating its scalability and strong performance advantages.

Guidance in Diffusion. Guidance techniques play a crucial role across a wide range of visual tasks [27, 55]. In diffusion-based generative models, guidance emerges as a core mechanism for conditional generation. Early approaches such as Classifier Guidance (CG) [6] improve sample quality by leveraging an external classifier to steer the denoising process toward desired semantic targets, but require training a separate noise-aware classifier and are difficult to scale to complex or multimodal conditioning signals. To address these limitations, Classifier-Free Guidance (CFG) [15] was introduced, enabling conditional generation [25, 37, 39] without an auxiliary classifier. By jointly training the diffusion model with and without conditioning inputs, CFG allows flexible control at inference time through a simple interpolation between conditional and unconditional predictions. Subsequent works [5, 19, 22, 61] explore adaptive guidance strategies for CFG by dynamically adjusting the guidance scale [48, 52] or refining the guidance direction [38] to mitigate oversaturation, thereby improving fidelity and reducing artifacts. Building upon these, recent studies [9, 40] have extended the CFG to flow matching models. For example, CFG-Zero* [9] introduces an optimized guidance scale to correct velocity estimation, while Rectified-CFG++ [40] proposes an adaptive predictor-corrector scheme that integrates the deterministic efficiency of rectified flows. These methods demonstrate that guidance remains a powerful and extensible mechanism for controllable and efficient generative modeling.

Control Theory. Control theory provides a foundational framework for designing systems that can regulate their behavior to achieve desired objectives. Its principles are fundamental to ensuring the performance, safety, and efficiency of complex systems across aerospace [3, 4], robotics [43, 57], and industrial process control [33, 42]. Among various approaches, Proportional-Integral-Derivative (PID) control [1] remains one of the most widely adopted strategies due to its simplicity and broad applicability, effectively balancing responsiveness, stability, and steady-state accuracy using feedback errors. Beyond PID, advanced paradigms address more complex challenges: Model Predictive Control (MPC) optimizes future actions based on a system model [10], while Adaptive Control adjusts parameters online to manage uncertainties [2]. Furthermore, robust control strategies guarantee stability against defined model

inaccuracies. Sliding Mode Control (SMC) [7], as a prominent example of robust control, introduces a discontinuous law that forces the system trajectory onto a predefined manifold, ensuring exceptional resilience to disturbances. These diverse control strategies have inspired recent efforts to integrate feedback-based and stability-driven principles into learning-based and generative modeling frameworks.

3. Method

3.1. Preliminaries

Classifier-Free Guidance (CFG) [15] introduces guidance by linearly interpolating between the conditional and unconditional velocity fields. Let $\mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset)$ denote the *unconditional* velocity, obtained by dropping the condition \mathbf{c} during training. The guided velocity is computed as:

$$\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{c}) = \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset) + w \cdot (\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset)), \quad (1)$$

where $w \geq 1$ is the guidance weight. Rearranging yields:

$$\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{c}) = (1 - w) \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset) + w \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}). \quad (2)$$

When $w = 1$, the model reduces to the standard conditional predictor. Increasing $w > 1$ amplifies the conditional component, improving semantic alignment at the cost of reduced sample diversity.

Weight-Scheduler [48] introduces a time-varying guidance weight $w(t)$ in place of the fixed weight w in standard CFG. The guided velocity becomes:

$$\hat{\mathbf{v}}_\theta(\mathbf{x}_t, t, \mathbf{c}) = \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset) + w(t) \cdot (\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset)). \quad (3)$$

Here, the scheduler $w(t)$ is a monotonically increasing function of the denoising step to avoid overshooting the guidance in the initial stages.

Adaptive Projected Guidance (APG) [38] mitigates oversaturation by down-weighting the component of the guidance direction that is parallel to the conditional prediction. The standard CFG update direction $\Delta \mathbf{v}_t = \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset)$ is decomposed into parallel and orthogonal components:

$$\Delta \mathbf{v}_t^\parallel = \frac{\langle \Delta \mathbf{v}_t, \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) \rangle}{\|\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c})\|^2} \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}), \quad \Delta \mathbf{v}_t^\perp = \Delta \mathbf{v}_t - \Delta \mathbf{v}_t^\parallel. \quad (4)$$

APG reduces oversaturation by down-weighting the parallel term. The guided velocity then becomes:

$$\mathbf{v}_{\text{APG}}(\mathbf{x}_t, t, \mathbf{c}) = \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset) + w \cdot (\Delta \mathbf{v}_t^\perp + \eta \Delta \mathbf{v}_t^\parallel), \quad \eta \leq 1. \quad (5)$$

APG chooses $\eta < 1$ to suppress oversaturation while preserving the quality-enhancing orthogonal component.

3.2. Motivation

Classifier-Free Guidance (CFG) has demonstrated remarkable empirical success across numerous diffusion-based

generative models and related applications. In prior formulations, CFG can be viewed as a linear extrapolation within deterministic diffusion flows [40], as expressed in Eq. (1). We denote the guidance term as

$$\mathbf{e}(t) = \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset). \quad (6)$$

Ideally, during the denoising process from time step T to 0, CFG continuously injects conditional information into the trajectory \mathbf{x}_t . This mechanism progressively enriches the semantic content encoded in \mathbf{x}_t as the timestep decreases. In the final stages of denoising, when most semantic information has already been embedded in \mathbf{x}_t , the conditional and unconditional predictions tend to converge, *i.e.*, $\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) \approx \mathbf{v}_\theta(\mathbf{x}_t, t, \emptyset)$, such that both \mathbf{e} and its temporal derivative $\dot{\mathbf{e}}$ approach zero. This ideal behavior can be viewed geometrically as a guidance process evolving on the $(\mathbf{e}, \dot{\mathbf{e}})$ plane, aiming to drive the system state toward the equilibrium point $(0, 0)$. The most direct and stable convergence path under such a setting corresponds to the first-order linear system:

$$\dot{\mathbf{e}} = -\lambda_0 \cdot \mathbf{e}, \quad \lambda_0 = -\frac{\dot{\mathbf{e}}(T) \cdot \mathbf{e}(T)}{\|\mathbf{e}(T)\|^2}, \quad (\mathbf{e}(T) \neq \mathbf{0}). \quad (7)$$

From the viewpoint of differential-equation, it implies a rapid and stable exponential convergence.

In practice, however, the assumption of collinear $(\mathbf{e}, \dot{\mathbf{e}})$ relationship rarely holds, especially when model capacity increases and the CFG guidance scale is enlarged. The resulting system becomes highly nonlinear, and the standard CFG formulation may exhibit oscillatory or divergent behavior, as illustrated in Fig. 1 (left). Such instability often manifests as color distortion, loss of fine details, or inconsistent textures in generated images [5, 38, 61].

Motivated by the effectiveness of control methods in stabilizing oscillatory behavior and ensuring convergence in dynamical systems, we revisit CFG from a control-theoretic perspective. Rather than treating CFG as a static extrapolation method, we propose to view CFG as a feedback control process that actively regulates the evolution of $\mathbf{e}(t)$, driving it toward the equilibrium in a principled, rate-aware manner.

3.3. Theoretical Formulation of CFG-Ctrl

In this section, we introduce CFG-Ctrl, a unified theoretical framework for CFG in flow matching models, which systematically interprets diverse guidance strategies. We first model the flow matching sampling process as a continuous-time controlled dynamical system. Let $\mathbf{x}_t \in \mathcal{X} \subseteq \mathbb{R}^d$ denote the latent state at time $t \in [0, T]$, whose evolution follows the control-affine ordinary differential equation

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_\theta(\mathbf{x}_t, t) + \mathbf{G}(\mathbf{x}_t, t)\mathbf{u}_t, \quad (8)$$

Table 1. **Typical CFG variants under CFG-Ctrl formulation.** We summarize the key components of various methods under the control formulation, along with their corresponding types of control interpretations.

Method	Gain K_t	Operator Π_t	Error $\mathbf{e}(t)$	Control Interpretation
CFG [15]	w	I	$\Delta \mathbf{v}_\theta(t)$	Proportional control
Weight Scheduler [48]	$w(t)$	I	$\Delta \mathbf{v}_\theta(t)$	Time-varying gain scheduling
APG [38]	$w [I \ \eta I]$	$\begin{bmatrix} I - P_t \\ P_t \end{bmatrix}, P_t = \frac{\mathbf{v}_\theta(\mathbf{c})\mathbf{v}_\theta(\mathbf{c})^\top}{ \mathbf{v}_\theta(\mathbf{c}) ^2}$	$\Delta \mathbf{v}_\theta(t)$	Projection-based Feedback Control
CFG-Zero* [9]	$[wI \ \frac{s_t}{1-s_t}I], s_t = \frac{\mathbf{v}_\theta(\mathbf{c})^\top \mathbf{v}_\theta(\varnothing)}{ \mathbf{v}_\theta(\varnothing) ^2}$	$\begin{bmatrix} I - P_t \\ P_t \end{bmatrix}, P_t = \frac{\mathbf{v}_\theta(\varnothing)\mathbf{v}_\theta(\varnothing)^\top}{ \mathbf{v}_\theta(\varnothing) ^2}$	$\Delta \mathbf{v}_\theta(t)$	Projection-based Feedback Control
Rectified-CFG++ [40]	$[I \ \alpha(t)I], \alpha(t) = \lambda_{max}(1-t)^\gamma$	I	$\begin{bmatrix} \Delta \mathbf{v}_\theta(t) \\ \Delta \mathbf{v}_\theta(t - \frac{\Delta t}{2}) \end{bmatrix}$	Model Predictive Control
SMC-CFG	w	I	$\Delta \mathbf{v}_\theta(t) - k \cdot \text{sign}(s_t)$	Sliding Mode Control

with the initial condition $\mathbf{x}_0 \sim \mathcal{N}(0, \mathbf{I})$. In Eq. (8), $\mathbf{v}_\theta : \mathcal{X} \times [0, T] \rightarrow \mathcal{X}$ is the pre-trained velocity field, $\mathbf{G} : \mathcal{X} \times [0, T] \rightarrow \mathbb{R}^{d \times m}$ is the input mapping matrix, and $\mathbf{u}_t \in \mathcal{U} \subseteq \mathbb{R}^m$ is the guidance control input.

Guidance mechanisms act directly in the latent coordinates without cross-space transformations; hence we set $\mathbf{G}(\mathbf{x}_t, t) = \mathbf{I}_d$ (full actuation, $m = d$), reducing Eq. (8) to the additive-velocity form

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_\theta(\mathbf{x}_t, t) + \mathbf{u}_t. \quad (9)$$

To better analyze guidance mechanisms, we propose to formulate the control signal \mathbf{u}_t using a general state-feedback law, decomposing it into two key components:

$$\mathbf{u}_t = K_t \Pi_t(\mathbf{e}(t)). \quad (10)$$

Here, $\mathbf{e}(t)$, defined in Eq. (6), is regarded as the semantic error of the system. We term K_t the guidance schedule, as it schedules the guidance strength, and Π_t the direction operator, as it shapes the correction direction (*e.g.*, normalization or projection).

Under the CFG-Ctrl formulation, we can interpret the standard CFG as a specific, simple instance of this general control law. The standard CFG update Eq. (1) modifies the closed-loop dynamics as:

$$\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + w (\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing)), \quad (11)$$

where w is the guidance scale. This specific form is recovered when the guidance schedule K_t is a constant scalar and the direction operator Π_t is identity:

$$K_t = w, \quad \Pi_t = I. \quad (12)$$

Substituting these into the closed-loop dynamics $\frac{d\mathbf{x}_t}{dt} = \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + K_t \Pi_t(\mathbf{e}(t))$ in Eq. (9) and (10) yields:

$$\begin{aligned} \frac{d\mathbf{x}_t}{dt} &= \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + K_t \Pi_t(\mathbf{e}(t)) \\ &= \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + w (\mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing)), \end{aligned} \quad (13)$$

which recovers the standard CFG update in Eq. (11). Thus, CFG is mathematically equivalent to a proportional state-feedback controller (P-control) acting on the semantic alignment error $\mathbf{e}(t)$. The guidance scale w , serving as the constant guidance schedule, directly plays the role of the proportional gain.

This state-feedback perspective, decomposing guidance into the guidance schedule K_t and the direction operator Π_t , provides a structured framework to understand existing CFG advancements. Many follow-up typical CFG variants can be reinterpreted as specific control laws for modulating either the strength via K_t or the direction via Π_t .

Guidance Schedule. We next focus on the guidance schedule component K_t . Recall that standard CFG applies a *constant* guidance schedule w to the semantic feedback signal.

However, K_t does not need to be fixed. A prominent example of a dynamic guidance schedule is guidance weight scheduling. Recent work [48] has shown that replacing the constant gain w with a time-varying schedule $w(t)$ leads to substantial improvements in sample quality and semantic consistency. Under our formulation, this corresponds to choosing a time-dependent guidance schedule while the direction operator remains identity:

$$K_t = w(t), \quad \Pi_t = I, \quad (14)$$

with the same semantic error signal $\mathbf{e}(t)$ in Eq. (6). The resulting closed-loop dynamics are shown in Eq. (3).

This reveals that the weight-scheduler approach is a time-varying proportional feedback controller, also known in control theory as gain-scheduled control. The key difference from standard CFG is that the guidance schedule K_t is no longer fixed.

From a control perspective, the motivation for a dynamic guidance schedule is clear: in early stages of sampling, the state \mathbf{x}_t is dominated by noise, so applying strong correction (a large K_t) may amplify noise rather than semantic alignment. A smaller gain $w(t)$ is therefore desirable at high noise levels. As the sample becomes more structured, the feedback signal becomes more semantically meaningful, and the gain can be safely increased.

Direction Operator. The direction operator Π_t can be combined with a more advanced guidance schedule K_t . Whereas weight-schedulers [48] use a scalar K_t and an identity Π_t , Adaptive Projected Guidance (APG) [38] demonstrates a case where K_t becomes a matrix gain, working in conjunction with a structured Π_t . Within the first-order state-feedback framework, APG can be written as:

$$K_t = w [I \ \eta I], \quad \Pi_t = \begin{bmatrix} I - P_t \\ P_t \end{bmatrix}, \quad P_t = \frac{\mathbf{v}_\theta(\mathbf{c})\mathbf{v}_\theta(\mathbf{c})^\top}{|\mathbf{v}_\theta(\mathbf{c})|^2}, \quad (15)$$

where P_t is an orthogonal projection onto the conditional direction $\mathbf{v}_\theta(\mathbf{c}) = \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c})$. Applying the direction operator Π_t first decomposes the semantic error $\mathbf{e}(t)$ into orthogonal and parallel components:

$$\begin{bmatrix} \Delta \mathbf{v}_t^\perp \\ \Delta \mathbf{v}_t^\parallel \end{bmatrix} = \Pi_t(\mathbf{e}(t)) = \begin{bmatrix} I - P_t \\ P_t \end{bmatrix} \mathbf{e}(t). \quad (16)$$

The guidance schedule K_t , now a structured matrix, applies the global CFG scaling w while introducing an additional factor η specifically on the parallel component, yielding:

$$\begin{aligned} \frac{d\mathbf{x}_t}{dt} &= \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + K_t \Pi_t(\mathbf{e}(t)) \\ &= \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing) + w \left(\Delta \mathbf{v}_t^\perp + \eta \Delta \mathbf{v}_t^\parallel \right). \end{aligned} \quad (17)$$

APG therefore reshapes the feedback signal: instead of uniformly amplifying the semantic correction (as in scalar K_t), the matrix-based guidance schedule K_t selectively enhances the parallel component aligned with the conditional direction. In control-theoretic terms, APG is a projection-based feedback controller. This design improves semantic alignment without the instability of simply increasing the proportional gain w , since it adjusts *how strongly* the guidance acts on different components (via K_t) rather than just *how strongly overall*. We show more control interpretations of various CFG methods in Table 1. For notation list and theoretical details, please refer to supplementary material.

3.4. Sliding Mode Control CFG

Existing CFG methods primarily rely on linear feedback, such as linear combinations or orthogonal projections of the conditional and unconditional velocity estimates. However, the ODE flow is inherently a highly nonlinear dynamical system, particularly when the model capacity becomes large and guidance scale is high. In such regimes, linear guidance tends to amplify nonlinear distortions, leading to oversaturated textures and semantic inconsistency.

To address these issues, we reinterpret CFG under the first-order state-feedback control framework introduced in Sec. 3.3. Under this perspective, we propose Sliding Mode Control CFG (SMC-CFG), which introduces a nonlinear

Algorithm 1 SMC-CFG

```

1: Input: Velocity model  $\mathbf{v}_\theta(\cdot, t, \mathbf{c})$ , input condition  $\mathbf{c}$ ,
   guidance scale  $w$ , SMC parameters  $\lambda, k$ .
2:  $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$ 
3: for  $t = T$  to 1 do
4:    $\mathbf{v}_t(\mathbf{c}) \leftarrow \mathbf{v}_\theta(\mathbf{x}_t, t, \mathbf{c})$       # Conditional prediction
5:    $\mathbf{v}_t(\varnothing) \leftarrow \mathbf{v}_\theta(\mathbf{x}_t, t, \varnothing)$     # Unconditional prediction
6:    $\mathbf{e}(t) \leftarrow \mathbf{v}_t(\mathbf{c}) - \mathbf{v}_t(\varnothing)$ 
7:   if  $\mathbf{e}(t+1)$  is None then
8:      $\mathbf{e}(t+1) \leftarrow \mathbf{e}(t)$ 
9:   end if
10:   $\mathbf{s}_t \leftarrow (\mathbf{e}(t) - \mathbf{e}(t+1)) + \lambda \cdot \mathbf{e}(t+1)$     # Sliding
   surface
11:   $\Delta \mathbf{e} \leftarrow -k \cdot \text{sign}(\mathbf{s}_t)$                 # Switching control
12:   $\mathbf{e}(t) \leftarrow \mathbf{e}(t) + \Delta \mathbf{e}$                 # SMC guidance update
13:   $\hat{\mathbf{v}}_t \leftarrow \mathbf{v}_t(\varnothing) + w \cdot \mathbf{e}(t)$ 
14:   $\hat{\mathbf{x}}_{t-1} \leftarrow \text{ODEUpdate}(\mathbf{x}_t, \hat{\mathbf{v}}_t, t)$ 
15: end for
16: Return  $\mathbf{x}_0$ 

```

sliding surface that continuously corrects the semantic deviation while constraining the diffusion trajectory to evolve toward a stable low-energy semantic manifold.

For the semantic error $\mathbf{e}(t)$ defined in Eq. (6), the ideal target behavior is that $(\mathbf{e}(t), \dot{\mathbf{e}}(t))$ decays directly toward origin, as shown in Figure 1 (left):

$$\dot{\mathbf{e}}(t) = -\lambda_0 \mathbf{e}(t), \quad \lambda > 0. \quad (18)$$

Here λ_0 denotes the slope of the ideal line, with its value typically determined by the initial state $(\mathbf{e}(T), \dot{\mathbf{e}}(T))$. The ODE solution $\mathbf{e}(t) = \mathbf{e}(T)\exp(-\lambda t)$ also ensures smooth, monotonic exponential convergence.

However, the diffusion dynamics cannot ensure the ideal process of \mathbf{e} ; thus we define the sliding mode surface:

$$\mathbf{s}(t) = \dot{\mathbf{e}}(t) + \lambda \mathbf{e}(t), \quad (19)$$

where λ is an adjustable shape parameter of the sliding mode surface, and the surface implicitly encodes the target error dynamics in Eq. (18). The manifold $\mathbf{s}(t) = \mathbf{0}$, as illustrated by the dashed line in Figure 1 (right), defines the desired semantic equilibrium flow. We adopt the Lyapunov function [32] V to measure the deviation of the system from the sliding manifold. For stable convergence, the system energy must monotonically decrease over time:

$$V(\mathbf{s}) = \frac{1}{2} \|\mathbf{s}\|^2, \quad \dot{V} = \mathbf{s}^\top \dot{\mathbf{s}} < 0. \quad (20)$$

We derive $\dot{\mathbf{s}}$ from the semantic guidance error in Eq. (6). Using the chain rule, its time derivative is

$$\begin{aligned} \dot{\mathbf{e}}(t) &= \frac{\partial \mathbf{v}_\theta(\mathbf{x}, t, \mathbf{c})}{\partial \mathbf{x}} \dot{\mathbf{x}} - \frac{\partial \mathbf{v}_\theta(\mathbf{x}, t, \varnothing)}{\partial \mathbf{x}} \dot{\mathbf{x}} \\ &\quad + \frac{\partial \mathbf{v}_\theta(\mathbf{x}, t, \mathbf{c})}{\partial t} - \frac{\partial \mathbf{v}_\theta(\mathbf{x}, t, \varnothing)}{\partial t}. \end{aligned} \quad (21)$$

Table 2. **Quantitative evaluation of CFG methods.** Lower (\downarrow) FID and higher (\uparrow) CLIP, Aesthetic, ImageReward, PickScore, HPSv2, HPSv2.1 and MPS scores indicate better performance. Note that Qwen-Image preserves natural image statistics, yielding the lowest FID.

Guidance	FID \downarrow	CLIP \uparrow	Aesthetic \uparrow	ImageReward \uparrow	PickScore \uparrow	HPSv2 \uparrow	HPSv2.1 \uparrow	MPS \uparrow
SD3.5 [8]	41.725	0.3399	5.4256	0.3591	0.2124	0.2710	0.2372	6.5554
w/ CFG [15]	21.421	0.3681	5.5875	0.8889	0.2190	0.2930	0.2842	7.2476
w/ CFG-Zero* [9]	20.317	0.3691	5.6124	0.9312	0.2195	0.2942	0.2862	7.0430
w/ Rect-CFG++ [40]	20.550	0.3655	5.5663	0.7097	0.2173	0.2887	0.2748	6.7854
w/ SMC-CFG	20.044	0.3694	5.5790	0.9486	0.2211	0.2945	0.2875	7.5719
Flux-dev [20]	52.598	0.3272	5.4568	0.2572	0.2137	0.2650	0.2280	5.9592
w/ CFG [15]	27.323	0.3692	5.5397	0.8749	0.2228	0.2917	0.2828	7.8387
w/ CFG-Zero* [9]	26.901	0.3742	5.7053	1.0300	0.2262	0.2987	0.2992	8.1573
w/ Rect-CFG++ [40]	27.219	0.3728	5.6909	1.0075	0.2252	0.2974	0.2963	7.9746
w/ SMC-CFG	26.398	0.3743	5.7342	1.0558	0.2268	0.2986	0.3021	8.2307
Qwen-Image [49]	24.894	0.3626	5.4081	0.5742	0.2157	0.2815	0.2613	6.7152
w/ CFG [15]	35.431	0.3815	5.5995	1.1063	0.2260	0.2996	0.3038	8.1852
w/ CFG-Zero* [9]	35.391	0.3822	5.6598	1.1941	0.2279	0.3019	0.3092	8.3739
w/ Rect-CFG++ [40]	34.371	0.3834	5.6007	1.1727	0.2276	0.3017	0.3068	8.1026
w/ SMC-CFG	33.371	0.3856	5.6289	1.2035	0.2275	0.3026	0.3105	8.4320

We introduce a sliding-mode correction term $\Delta \mathbf{e}(t)$, giving the full control $\mathbf{u}(t) = w(\mathbf{e}(t) + \Delta \mathbf{e}(t))$, which modulates the semantic guidance by directly shaping the error dynamics rather than altering the model prediction.

Substituting the controlled state dynamics in Eq. (8) into the time derivative of the semantic error yields:

$$\dot{\mathbf{e}}(t) = \Phi(t, \mathbf{x}) + w \left(\frac{\partial v_\theta(\mathbf{x}, t, \mathbf{e})}{\partial \mathbf{x}} - \frac{\partial v_\theta(\mathbf{x}, t, \emptyset)}{\partial \mathbf{x}} \right) \mathbf{G}(\mathbf{x}, t) \Delta \mathbf{e}(t), \quad (22)$$

where Φ absorbs terms independent of $\Delta \mathbf{e}(t)$. Differentiating and substituting the sliding surface definition in Eq. (19), we obtain:

$$\dot{\mathbf{s}}(t) = \Phi_s(t, \mathbf{x}) + \Gamma_s(t) \Delta \mathbf{e}(t), \quad (23)$$

where Γ_s denotes the coefficient matrix multiplying $\Delta \mathbf{e}$, and $\Phi_s(t, \mathbf{x})$ represents all remaining terms. We assume that Γ_s has minimum singular value lower-bounded (*i.e.*, $\sigma_{\min}(\Gamma_s) \geq b_{\min} > 0$), and Φ_s is bounded (*i.e.*, $\|\Phi_s\| \leq \delta$, $\exists \delta > 0$), which are standard in sliding mode control.

Substituting into the Lyapunov derivative yields:

$$\dot{V} = \mathbf{s}^\top \Phi(t, \mathbf{e}) + \mathbf{s}^\top \Gamma(t) \Delta \mathbf{e}(t). \quad (24)$$

We apply the classical switching control law:

$$\Delta \mathbf{e}(t) = -\mathbf{K} \cdot \text{sign}(\mathbf{s}(t)), \quad (25)$$

where $\mathbf{K} = k\mathbf{I}$ is a positive diagonal gain matrix. Since $\sigma_{\min}(\Gamma(t)\mathbf{K}) \geq k b_{\min}$, we obtain:

$$\dot{V} \leq \|\mathbf{s}\| \delta - k b_{\min} \|\mathbf{s}\| = -(k b_{\min} - \delta) \|\mathbf{s}\|. \quad (26)$$

Therefore, choosing k such that $k b_{\min} > \delta$ ensures:

$$\dot{V} = \mathbf{s}^\top \dot{\mathbf{s}} \leq -\eta \|\mathbf{s}\|, \quad \eta = k b_{\min} - \delta > 0. \quad (27)$$

Dividing both sides by $\|\mathbf{s}\| > 0$ yields the scalar differential inequality

$$\frac{d}{dt} \|\mathbf{s}(t)\| \leq -\eta.$$

Integrating from 0 to t gives

$$\|\mathbf{s}(t)\| \leq \|\mathbf{s}(0)\| - \eta t,$$

which supports finite-time convergence of $\mathbf{s}(t)$ to zero. In particular:

$$\|\mathbf{s}(t)\| = 0 \quad \text{for some } t \leq \frac{\|\mathbf{s}(0)\|}{\eta}. \quad (28)$$

We present the entire method in Algorithm 1. The proposed sliding mode surface and switching control law enforce stable semantic guidance by ensuring that the diffusion trajectory converges to the desired manifold, depicted by the red curve in Figure 1, eliminating oscillations and improving consistency during guided sampling.

4. Experiments

4.1. Experimental Setups

Datasets and Baselines. We conduct experiments on a subset of the MS-COCO [23] dataset, comprising 5,000 image-text pairs. To demonstrate the generality of our method across diverse model scales, we evaluate it on several state-of-the-art flow-based T2I models, including Stable Diffusion 3.5 (SD3.5) [8], Flux-dev [20], and Qwen-Image [49] with 8B, 12B, and 20B parameters, respectively. In addition to comparing against the standard Classifier-Free Guidance baseline, we include two recent guidance variants designed

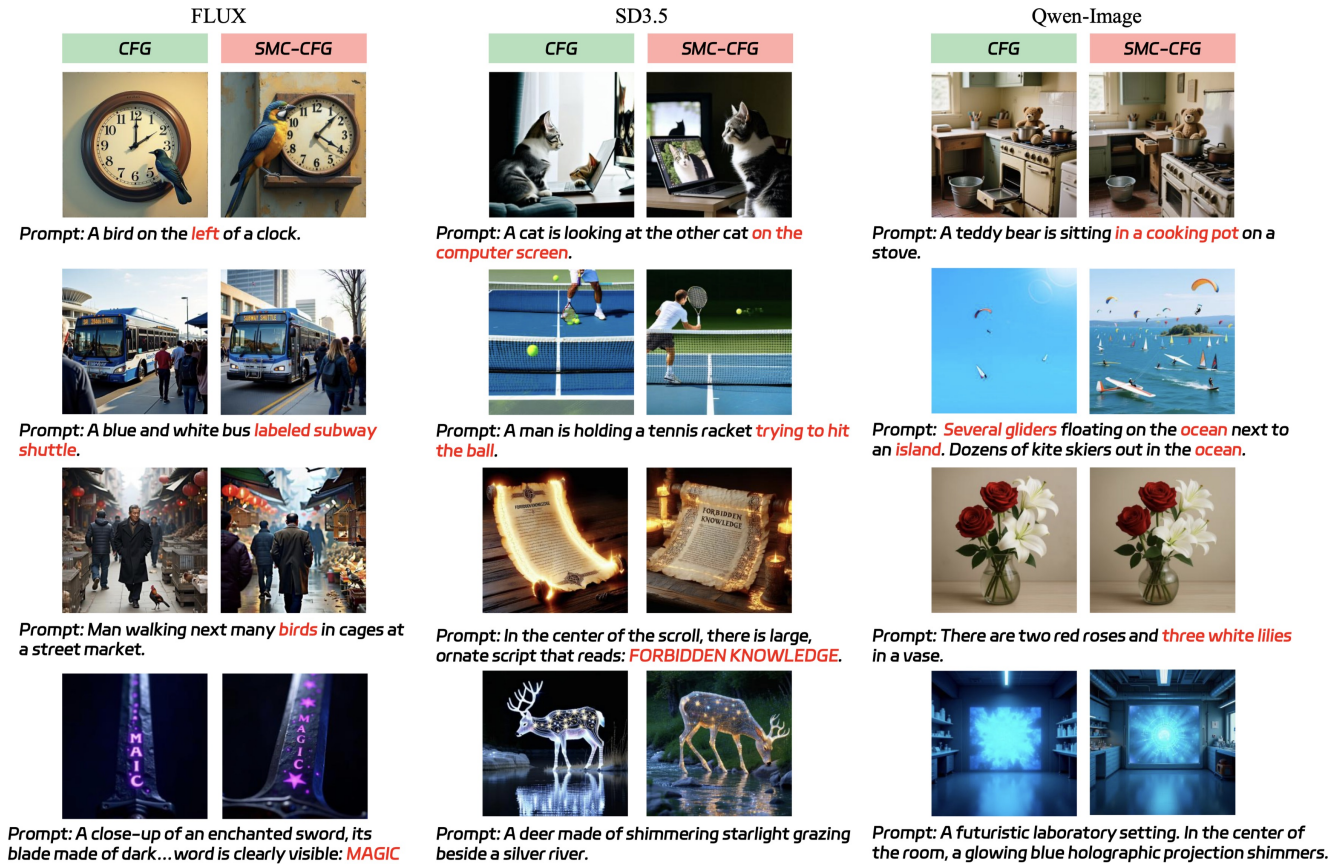


Figure 2. **Qualitative results across different T2I models.** We provide visual comparisons between CFG and our SMC-CFG across various models. SMC-CFG exhibits better performance in positional relationships, text generation, and detailed object representation.

specifically for flow-matching generative models: CFG-zero* [9] and Rectified-CFG++ [40]. We implement both methods on all evaluated backbones following the official paper and open-source code to ensure a fair and consistent comparison. For more comprehensive experiments on additional T2I benchmark and diffusion model, please refer to our supplementary material.

Evaluation Metrics. To assess image quality and visual realism, we report the Fréchet Inception Distance (FID) [14]. To measure the alignment between generated images and input text prompts, we use the CLIP Score [13, 35], which quantifies semantic consistency in the joint vision–language embedding space. In addition to these core metrics, we further provide a comprehensive evaluation of aesthetic quality and human preference, including Aesthetic Score [41], ImageReward [53], PickScore [17], HPSv2 [51], HPSv2.1 [51], and MPS [59]. Together, these metrics offer a holistic perspective on both the fidelity and human-perceived appeal of the generated content.

Implementation Details. All experiments are conducted on a single NVIDIA A100 GPU (40GB). We implement the proposed method SMC-CFG on three representative pre-trained T2I diffusion models: SD3.5 [8], Flux-dev [20], and

Qwen-Image [49]. For all models, we adopt their default CFG scales provided in the official implementations. In our SMC-CFG framework, the two hyperparameters λ and K are kept fixed within each model and shared across all datasets and experimental conditions to ensure fair comparison. See supplementary material for more implementation details and complete hyperparameter configurations.

4.2. Text-to-Image Generation

In this section, we evaluate the effectiveness of our proposed SMC-CFG in text-to-image generation. Experiments are conducted on the MS-COCO [23] dataset using three state-of-the-art flow-based models. To ensure a fair and up-to-date comparison, we implement two recent CFG variants (CFG-zero* [9] and Rectified-CFG++ [40]) designed for flow-matching models as baselines.

Quantitative Evaluation. Table 2 reports the quantitative results of SMC-CFG compared with the standard CFG and baselines across different T2I models. Our method consistently achieves lower FID scores, indicating the generated images exhibit improved visual quality and realism. Meanwhile, the higher CLIP Scores demonstrate stronger semantic alignment between the generated images and the input



Figure 3. **Qualitative comparison with baseline methods.** For challenging scenarios including relative positions, clothing styles, and human actions, baseline methods produce irrational outputs, while SMC-CFG preserves robust text consistency.

prompts. Furthermore, SMC-CFG attains superior scores on ImageReward, HPSv2.1, and MPS scores, signifying that the generated images are more aligned with human aesthetic and preference judgments. For additional metrics, our method also achieves comparable or better results than the baselines, demonstrating strong overall generation quality.

Qualitative Evaluation. We further present qualitative comparisons to illustrate the improvements achieved by SMC-CFG. As shown in Figure 2, across different model backbones, our method produces images with sharper details, more coherent object structures, and more faithful adherence to the textual descriptions compared to standard CFG. This demonstrates that our approach is consistently effective and model-agnostic. In addition, Figure 3 highlights results on more challenging prompts involving complex compositions, fine-grained semantics, or stylistic attributes. Compared with recent flow-matching-based guidance variants, SMC-CFG generates images that better preserve semantic correctness and maintain aesthetic quality, without introducing over-smoothing or mode collapse.

Table 3. **Ablation study on hyperparameter λ and k .** We conduct ablation across various hyperparameter settings in four metrics: FID, CLIP, Aesthetic (Aesth), and ImageReward (ImgRwd), respectively measuring generation quality, semantic alignment, aesthetic level, and human preference.

λ	k	FID ↓	CLIP ↑	Aesth ↑	ImgRwd ↑
3	0.1	26.193	0.3698	5.7064	1.0174
4	0.1	26.006	0.3701	5.7098	1.0219
5	0.1	25.951	0.3709	5.7128	1.0248
6	0.1	26.143	0.3703	5.7071	1.0228
5	0.1	25.951	0.3709	5.7128	1.0248
5	0.4	26.143	0.3719	5.7218	1.0504
5	0.7	26.416	0.3739	5.7175	1.0406
5	1.0	26.281	0.3741	5.7054	1.0453



Figure 4. Visual comparison between CFG (top) and SMC-CFG (bottom) across different CFG scales.

4.3. Ablation Studies and Analysis

Ablation on Hyperparameters. To gain a deeper understanding of the roles of hyperparameters in SMC-CFG, we perform an ablation study on their distinct impacts. The top of Table 3 illustrates how λ shapes the sliding mode surface. Extreme values (too low or too high) distort this manifold, impairing guidance stability and diminishing output fidelity. On the bottom of Table 3, we explore influence of k (with fixed λ), which governs the force toward the sliding mode surface. Modest k yields slow convergence and meanwhile alleviates the distortions introduced by CFG, thereby weakening text-image alignment (*i.e.*, lower CLIP scores) but preserves realism (*i.e.*, lower FID). In contrast, excessive k causes abrupt pulls, triggering erratic sampling or vibrations. Though boosting semantic match, such outputs suffer from reduced aesthetic appeal and poor human-preference ratings. Overall, suitable hyperparameters strike a trade-off between perceptual excellence and textual fidelity.

Guidance Scale. We analyze how guidance scales influence the generation performance of SMC-CFG in Figure 4. At large guidance scales, CFG improves semantic alignment at the cost of significant degradation in image quality and realism. In contrast, SMC-CFG exhibits more stable performance across a wide range of guidance scales, maximizing the capability of CFG while avoiding significant reductions in image quality and aesthetic appeal.

5. Conclusion

We explore a unified framework called CFG-Ctrl, reinterpreting CFG as a feedback control in flow matching models and analyzing its nonlinear behaviors under high guidance scales. From this perspective, we further propose SMC-CFG, a nonlinear control-based guidance mechanism that introduces a switching control term to enforce fast and stable convergence along the sliding mode surface. Extensive experiments demonstrate that SMC-CFG consistently improves semantic alignment and visual fidelity while maintaining robustness across diverse guidance scales. Ablation studies also reveal how its hyperparameters affect stability and perception. We believe this control-theoretic perspective provides a promising direction for more effective and robust guidance in future large-scale generative models.

Acknowledgements. This work was supported in part by the Beijing Natural Science Foundation under Grant L252011, by the National Natural Science Foundation of China under Grant 62576185, and by the Young Elite Scientist Sponsorship Program by CAST under Grant YESS20240544.

References

- [1] K.J. Åström and T. Hägglund. *PID Controllers*. International Society for Measurement and Control, 1995. 2
- [2] K. J. Åström. Adaptive Control. In *Mathematical System Theory: The Influence of R. E. Kalman*, pages 437–450. Springer, Berlin, Heidelberg, 1991. 2
- [3] Randal W Beard and Timothy W McLain. *Small unmanned aircraft: Theory and practice*. Princeton university press, 2012. 2
- [4] Arthur Earl Bryson. *Applied optimal control: optimization, estimation and control*. Routledge, 2018. 2
- [5] Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. Cfg++: Manifold-constrained classifier free guidance for diffusion models. *arXiv preprint arXiv:2406.08070*, 2024. 1, 2, 3
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2
- [7] Christopher Edwards and Sarah K Spurgeon. *Sliding mode control: theory and applications*. CRC press, 1998. 2, 3
- [8] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 1, 6, 7
- [9] Weichen Fan, Amber Yijia Zheng, Raymond A Yeh, and Zhiwei Liu. Cfg-zero*: Improved classifier-free guidance for flow matching models. *arXiv preprint arXiv:2503.18886*, 2025. 1, 2, 4, 6, 7
- [10] Carlos E. García, David M. Prett, and Manfred Morari. Model predictive control: Theory and practice—A survey. *Automatica*, 25(3):335–348, 1989. 2
- [11] Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018. 1
- [12] Yoav HaCohen, Nisan Chiprut, Benny Brazowski, Daniel Shalem, Dudu Moshe, Eitan Richardson, Eran Levin, Guy Shiran, Nir Zabari, Ori Gordon, et al. Ltx-video: Realtime video latent diffusion. *arXiv preprint arXiv:2501.00103*, 2024. 2
- [13] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pages 7514–7528, 2021. 7
- [14] 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. *Advances in neural information processing systems*, 30, 2017. 7
- [15] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 1, 2, 3, 4, 6
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2
- [17] Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Pick-a-pic: An open dataset of user preferences for text-to-image generation. *Advances in neural information processing systems*, 36:36652–36663, 2023. 7
- [18] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. 1
- [19] 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. *Advances in Neural Information Processing Systems*, 37:122458–122483, 2024. 2
- [20] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 1, 2, 6, 7
- [21] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 1
- [22] Shanchuan Lin, Bingchen Liu, Jiashi Li, and Xiao Yang. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 5404–5411, 2024. 2
- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6, 7
- [24] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 1, 2
- [25] Fangfu Liu, Hanyang Wang, Weiliang Chen, Haowen Sun, and Yueqi Duan. Make-your-3d: Fast and consistent subject-driven 3d content generation. In *European Conference on Computer Vision*, pages 389–406. Springer, 2024. 2
- [26] Fangfu Liu, Hanyang Wang, Shunyu Yao, Shengjun Zhang, Jie Zhou, and Yueqi Duan. Physics3d: Learning physical properties of 3d gaussians via video diffusion. *arXiv preprint arXiv:2406.04338*, 2024. 2
- [27] Fangfu Liu, Hao Li, Jiawei Chi, Hanyang Wang, Minghui Yang, Fudong Wang, and Yueqi Duan. Langscene-x: Reconstruct generalizable 3d language-embedded scenes with

- trimap video diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 29010–29020, 2025. 2
- [28] Fangfu Liu, Hanyang Wang, Yimo Cai, Kaiyan Zhang, Xiaohang Zhan, and Yueqi Duan. Video-t1: Test-time scaling for video generation. *arXiv preprint arXiv:2503.18942*, 2025. 2
- [29] Fangfu Liu, Junliang Ye, Yikai Wang, Hanyang Wang, Zhengyi Wang, Jun Zhu, and Yueqi Duan. Dreamreward-x: Boosting high-quality 3d generation with human preference alignment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2
- [30] Fangfu Liu, Wenqiang Sun, Hanyang Wang, Yikai Wang, Haowen Sun, Junliang Ye, Jun Zhang, and Yueqi Duan. Reconx: Reconstruct any scene from sparse views with video diffusion model. *IEEE Transactions on Image Processing*, 2026. 2
- [31] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 1, 2
- [32] Aleksandr Mikhailovich Lyapunov. The general problem of the stability of motion. *International journal of control*, 55(3):531–534, 1992. 5
- [33] David Q Mayne, James B Rawlings, Christopher V Rao, and Pierre OM Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000. 2
- [34] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 1
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 7
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 2
- [37] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22500–22510, 2023. 2
- [38] Seyedmorteza Sadat, Otmar Hilliges, and Romann M Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2024. 1, 2, 3, 4, 5
- [39] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 2
- [40] Shreshth Saini, Shashank Gupta, and Alan C Bovik. Rectified-cfg++ for flow based models. *arXiv preprint arXiv:2510.07631*, 2025. 1, 2, 3, 4, 6, 7
- [41] Christoph Schuhmann. LAION-Aesthetics. <https://laion.ai/blog/laion-aesthetics/>, 2022. Accessed: 2023-11-10. 7
- [42] Dale E Seborg, Thomas F Edgar, Duncan A Mellichamp, and Francis J Doyle III. *Process dynamics and control*. John Wiley & Sons, 2016. 2
- [43] Bruno Siciliano, Lorenzo Sciacivco, Luigi Villani, and Giuseppe Oriolo. *Robotics: modelling, planning and control*. Springer, 2009. 2
- [44] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 1, 2
- [45] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 1, 2
- [46] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 1
- [47] Hanyang Wang, Fangfu Liu, Jiawei Chi, and Yueqi Duan. Videoscene: Distilling video diffusion model to generate 3d scenes in one step. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16475–16485. IEEE, 2025. 1
- [48] Xi Wang, Nicolas Dufour, Nefeli Andreou, Marie-Paule Cani, Victoria Fernández Abrevaya, David Picard, and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *Transactions on Machine Learning Research Journal*, 2024. 1, 2, 3, 4, 5
- [49] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 2, 6, 7
- [50] Kailu Wu, Fangfu Liu, Zhihan Cai, Runjie Yan, Hanyang Wang, Yating Hu, Yueqi Duan, and Kaisheng Ma. Unique3d: High-quality and efficient 3d mesh generation from a single image. *Advances in Neural Information Processing Systems*, 37:125116–125141, 2024. 2
- [51] 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. 7
- [52] Mengfei Xia, Nan Xue, Yujun Shen, Ran Yi, Tieliang Gong, and Yong-Jin Liu. Rectified diffusion guidance for conditional generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13371–13380, 2025. 1, 2
- [53] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023. 7

- [54] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. [1](#)
- [55] Runmao Yao, Yi Du, Zhuoqun Chen, Haoze Zheng, and Chen Wang. Airroom: Objects matter in room reidentification. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1385–1394, 2025. [2](#)
- [56] Runmao Yao, Junsheng Zhou, Zhen Dong, and Yu-Shen Liu. Anchoredream: Zero-shot 360 $\{\deg\}$ indoor scene generation from a single view via geometric grounding. *arXiv preprint arXiv:2601.16532*, 2026. [2](#)
- [57] Tsuneo Yoshikawa. *Foundations of robotics: analysis and control*. MIT press, 1990. [2](#)
- [58] Meysar Zeinali and Leila Notash. Adaptive sliding mode control with uncertainty estimator for robot manipulators. *Mechanism and Machine Theory*, 45(1):80–90, 2010. [2](#)
- [59] Sixian Zhang, Bohan Wang, Junqiang Wu, Yan Li, Tingting Gao, Di Zhang, and Zhongyuan Wang. Learning multi-dimensional human preference for text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8018–8027, 2024. [7](#)
- [60] Zibo Zhao, Zeqiang Lai, Qingxiang Lin, Yunfei Zhao, Haolin Liu, Shuhui Yang, Yifei Feng, Mingxin Yang, Sheng Zhang, Xianghui Yang, et al. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202*, 2025. [1](#)
- [61] Candi Zheng and Yuan Lan. Characteristic guidance: Non-linear correction for diffusion model at large guidance scale. In *International Conference on Machine Learning*, pages 61386–61412. PMLR, 2024. [2](#), [3](#)