TRAJCFG: TOKEN-WISE CLASSIFIER-FREE GUIDANCE VIA DENOISING TRAJECTORY FEEDBACK

Anonymous authors
Paper under double-blind review

ABSTRACT

Classifier-free guidance (CFG) typically applies a shared scale to all latent tokens, although their velocity predictions evolve differently during sampling. We introduce Trajectory-Adaptive Classifier-Free Guidance (TrajCFG), a training-free method for flow-matching models that adapts token-wise guidance using historical changes in guided velocity direction. TrajCFG maps this feedback to bounded scales, assigning stronger guidance to larger directional changes and weaker guidance to more consistent directions. It combines adjacent-step and interval-normalized longer-span feedback, reusing standard CFG predictions without additional model evaluations. Controlled diagnostics and ablations examine the roles of directional feedback, token-wise allocation, and history-span fusion. Experiments on text-to-video, image-to-video, text-to-audio-video, text-to-image, and image-to-3D generation show improvements over base samplers across multiple backbones. For text-to-video generation, TrajCFG improves VBench Total scores by 0.30–0.75 points over the strongest evaluated baseline on each of three backbones. These results support historical directional feedback as a broadly applicable signal for fine-grained guidance allocation.

1 INTRODUCTION

Flow matching (Lipman et al., 2023) provides a common framework for generative modeling by learning velocity fields that transport noise into data. Its applications span image, video, audiovisual, and 3D generation. Across these tasks, classifier-free guidance (CFG) (Ho & Salimans, 2022) steers generation toward desired conditions by combining target- and reference-conditioned predictions. Standard CFG typically applies a shared guidance scale to all latent sites at a given sampling step. Such uniform guidance does not account for differences in how local predictions evolve along the sampling trajectory, motivating adaptation to local sampling behavior.

Recent work has improved CFG by adapting when and how guidance is applied. Guidance schedules vary its strength across noise levels (Kynkäänniemi et al., 2024; Wang et al., 2024), while spatially adaptive and velocity-based methods determine the scale from local guidance energy or within-step velocity alignment (Li et al., 2026; Luo et al., 2026). Other approaches modify the guided update through optimized scaling or feedback corrections (Fan et al., 2025; Wang et al., 2026). Inference-time refinement also uses a pretrained generator's own predictions to improve samples through additional sampling operations (Jang et al., 2026). These developments raise a natural question: *can directional changes in guided velocities already computed during sampling provide feedback for adapting CFG at each latent site?*

Our starting point is that guided velocity history records how the sampler revises its local predictions. Under Euler sampling, changes in guided velocity correspond to revisions of the sampler's clean-endpoint estimate. Separating changes in direction from changes in magnitude provides a normalized directional statistic. Video visualizations show that historical velocity angles can exhibit content-related spatial structure. A controlled local reconstruction diagnostic further shows that reducing guidance tends to benefit sites with small historical angles, whereas increasing guidance tends to benefit those with larger angles. Together, these observations motivate guidance allocation based on historical directional changes. We evaluate the effectiveness of this principle across modalities through generation experiments.

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.

We introduce **Trajectory-Adaptive Classifier-Free Guidance (TrajCFG)**, a training-free method for token-wise guidance adaptation in flow matching models. As illustrated in Figure 1, TrajCFG compares guided velocities from the two most recent completed steps at each latent site and maps their cosine similarity to a bounded guidance scale. Larger directional changes receive stronger guidance, while more consistent directions receive weaker guidance. We use these scales to combine the current target- and reference-conditioned predictions. The guided velocities are then supplied to the solver and retained for subsequent steps, forming a feedback loop in which past sampling behavior informs future conditioning strength. We also incorporate interval-normalized longer-span feedback to complement the adjacent-step signal. Because the rule uses only local guided velocity histories and existing CFG predictions, it applies across modalities with different latent representations. TrajCFG requires neither auxiliary networks nor extra model evaluations and leaves the pretrained model and sampling solver unchanged.

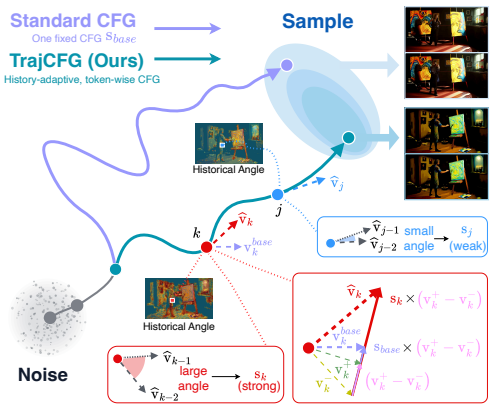

Figure 1: **Overview of TrajCFG.** Standard CFG uses a fixed scale (purple), while TrajCFG (teal) adapts each token's scale using guided velocities from the two most recent completed steps. Larger historical angles yield stronger guidance; smaller angles yield weaker guidance. Predictions are reused without extra model evaluations. For clarity, the illustration shows adjacent-step feedback. Trajectories are schematic.

We evaluate TrajCFG across text-to-video, image-to-video, text-to-audio-video, text-to-image, and image-to-3D generation. Experiments show improvements in generation quality and conditional alignment across diverse model families. Controlled diagnostics and ablation studies examine the roles of token-wise allocation, historical directional feedback, and the combination of history spans. Together, these results support TrajCFG as a broadly applicable approach to adaptive CFG in flow matching models.

2 Related Work

Scheduling and adaptive classifier-free guidance. Classifier-free guidance (CFG) combines target- and reference-conditioned predictions to control generation (Ho & Salimans, 2022). Guidance schedules and limited-interval guidance vary conditioning strength with noise level (Wang et al., 2024; Kynkäänniemi et al., 2024), while Adaptive Guidance and Step AG selectively omit guidance evaluations (Castillo et al., 2025; Zhang & Wan, 2025). C^2FG and stage-wise analyses motivate time-dependent scales through score discrepancies and sampling dynamics, respectively (Gao et al., 2026; Jin et al., 2026). Online-feedback CFG selects scales using trained latent-space evaluators (Papalampidi et al., 2026), whereas PathGuide uses on-policy transport alignment for online scale adaptation or offline schedule fitting (Nevo & Hazan, 2026). TrajCFG extends adaptation to token-wise scales driven by cross-step changes in guided velocity direction.

Geometry-, spatial-, and history-aware guidance. CFG++, APG, and TCFG modify guidance through manifold constraints, projection and momentum, and score-component filtering, respectively (Chung et al., 2025; Sadat et al., 2024; Kwon et al., 2025). CFG-Zero combines within-step projection with suppression of initial updates (Fan et al., 2025), while VAGS adjusts a sample-wise scale using sampling progress and conditional–unconditional velocity alignment (Luo et al., 2026). For spatial adaptation, S-CFG uses attention-based semantic regions (Shen et al., 2024), whereas SAMG derives point-wise scales from current-step guidance difference energy (Li et al., 2026). History informs APG's momentum, CFG-Ctrl's discrepancy-based feedback correction (Wang et al., 2026), and HiGS's additional direction from current and aggregated past denoised predictions (Sadat et al., 2026). TrajCFG instead measures cross-step angular changes of the *guided* velocity at each fixed latent site to set bounded token-wise CFG scales for the next step. The feedback is com-

puted from guided velocities at completed sampling steps and modulates the current target–reference prediction difference without introducing a separate update direction.

Inference-time refinement and sampling interventions. Beyond CFG adaptation, FreeInit refines low-frequency initial noise (Wu et al., 2024), while Restart and Zigzag Sampling introduce additional noising–denoising or inversion operations (Xu et al., 2023; Bai et al., 2025). VideoGuide uses additional teacher or self-guided steps (Lee et al., 2025), and Self-Refining Video Sampling performs fixed-noise predict-and-perturb refinement (Jang et al., 2026). FlowMo optimizes latents with a motion objective (Shaulov et al., 2026), whereas SpecLoR spectrally corrects a clean-latent lookahead (Zhang et al., 2026). PAG and STG construct weak predictions through attention perturbation and layer skipping (Ahn et al., 2024; Hyung et al., 2025); STAS intervenes on sparse massive network activations (Cheng et al., 2026). TrajCFG uses a short history of velocities already computed during standard CFG sampling, requiring no auxiliary evaluator, extra model evaluations, latent optimization, or model-internal changes.

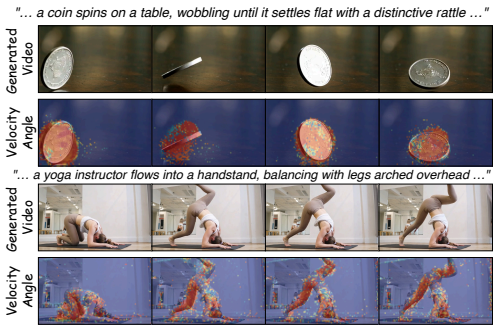

Figure 2: **Content-related structure in historical angles.** Generated frames and historical velocity-angle maps at sampling step $k = 5$ from Wan2.2-14B on Dynamic-Bench. Warmer colors indicate larger angular changes.

3 PRELIMINARIES

Flow Matching. Flow matching (Lipman et al., 2023) learns a velocity field that continuously transforms noise into data. Under an affine path formulation, the interpolation between a clean data latent $\mathbf{x}$ and noise $\boldsymbol{\epsilon}$ is

$$\mathbf{z}_\tau = (1-\tau)\mathbf{x} + \tau\boldsymbol{\epsilon}, \qquad \tau \in [0,1], \tag{1}$$

with path velocity $\boldsymbol{\epsilon} - \mathbf{x}$. A model $f_\theta(\mathbf{z}_\tau, \tau; y)$ learns the velocity field conditioned on y. During generation, a numerical solver integrates the learned field from $\tau = 1$ to $\tau = 0$, starting from noise.

Classifier-Free Guidance. At sampling step k, given the latent $\mathbf{z}_k$ and time τ_k, a pretrained model predicts velocities under target conditioning y^+ and reference conditioning y^-:

$$\mathbf{v}_k^+ = f_\theta(\mathbf{z}_k, \tau_k; y^+), \qquad \mathbf{v}_k^- = f_\theta(\mathbf{z}_k, \tau_k; y^-). \tag{2}$$

The reference condition follows the backbone's configuration, such as a null condition or a negative prompt. Standard CFG combines these predictions using a baseline guidance scale s_{base} shared across latent sites:

$$\widehat{\mathbf{v}}_k = \mathbf{v}_k^- + s_{\text{base}}(\mathbf{v}_k^+ - \mathbf{v}_k^-). \tag{3}$$

The sampling solver then updates the latent using the guided velocity $\widehat{\mathbf{v}}_k$. We use *latent site* to denote a position in the latent representation and *token* to denote its channel vector; a token-wise guidance scale is shared across that vector's channels.

4 METHOD

TrajCFG adapts the guidance scale at each latent site using historical changes in guided velocity direction. We first motivate this feedback signal, then define the token-wise allocation rule and its extension to longer history spans.

4.1 MOTIVATION: DIRECTIONAL CHANGES AS GUIDANCE FEEDBACK

Standard CFG in Eq. (3) applies a shared guidance scale s_{base} across all latent sites, even though their velocity predictions may evolve differently during sampling. Prior work shows that guidance

utility varies across noise levels (Kynkäänniemi et al., 2024; Wang et al., 2024) and that local prediction statistics can inform adaptive guidance (Li et al., 2026). These findings motivate us to investigate whether the directional evolution of guided velocities provides useful feedback for allocating guidance across latent sites.

Figure 2 illustrates historical velocity angles in two video examples. Larger angles concentrate around the foreground subjects, while much of the background exhibits smaller angles. This content-related structure suggests that guided velocity history captures differences in local sampling behavior. The statistic itself depends only on velocity histories at individual latent sites and can be computed across modalities. We next examine what directional changes represent and whether they inform local guidance benefits.

Interpreting directional changes. At a fixed latent site, with its index omitted, the affine path formulation gives the clean-endpoint estimate $\mathbf{D}_k = \mathbf{z}_k - \tau_k \widehat{\mathbf{v}}_k$. Under an Euler update,

$$\mathbf{D}_{k+1} - \mathbf{D}_k = -\tau_{k+1}\big(\widehat{\mathbf{v}}_{k+1} - \widehat{\mathbf{v}}_k\big), \tag{4}$$

so changes in guided velocity correspond to revisions of the sampler's clean-endpoint estimate, rather than directly measuring its error against a ground-truth target. For successive nonzero velocities, let $n_k = \|\widehat{\mathbf{v}}_k\|_2$ and define their angle $\theta_k \in [0, \pi]$ by

$$\cos\theta_k = \frac{\langle \widehat{\mathbf{v}}_{k+1}, \widehat{\mathbf{v}}_k \rangle}{n_{k+1} n_k}.$$

Their squared difference admits the decomposition

$$\|\widehat{\mathbf{v}}_{k+1} - \widehat{\mathbf{v}}_k\|_2^2 = (n_{k+1} - n_k)^2 + 2n_{k+1}n_k(1 - \cos\theta_k). \tag{5}$$

The two terms separate changes in magnitude from magnitude-weighted changes in direction. Normalizing the directional term by $2n_{k+1}n_k$ yields $1 - \cos\theta_k$, which increases with the angle and is invariant to independent positive rescaling of the two velocities. We therefore consider this statistic as a measure of directional revision over the observed sampling interval. Appendix A.1 provides the derivation, the extension to general solvers, and an analysis of interval dependence. This interpretation motivates an empirical question: do sites with different degrees of historical directional change benefit from different guidance scales?

Guidance response across directional changes. We evaluate Wan2.2-5B on 90 DAVIS (Pont-Tuset et al., 2017) samples at all sampling timesteps. For each video, noise seed, and step k, we construct the current and two preceding states by adding noise to ground-truth clean latents along the affine path, without sampler rollouts. We group sites by the angle between baseline-guided velocities at the constructed states for steps $k-2$ and $k-1$. The five groups in Figure 3 are 0–3°, 3–6°, 6–9°, 9–12°, and $> 12°$. Across candidate CFG scales, we fix these groups, the current noisy latent, and the target- and reference-conditioned predictions. We compute channel-averaged clean-endpoint MSE against the ground-truth latent and aggregate it within each group across evaluation instances.

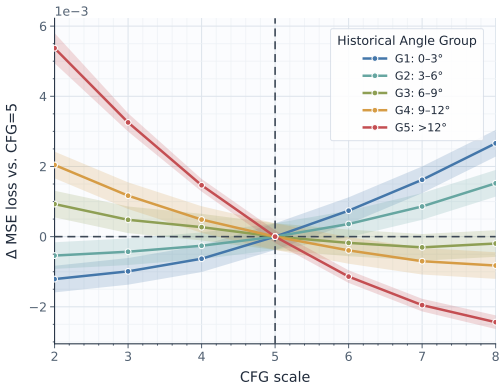

Figure 3: **CFG response by historical angle group.** Groups are fixed across candidate CFG scales. Clean-latent MSE differences are relative to CFG $= 5$ (negative is better); shading indicates standard deviations.

Figure 3 shows the resulting group-wise MSE differences relative to CFG $= 5$. Among the groups shown, reducing guidance below the baseline tends to lower error for sites with small historical angles, whereas increasing guidance tends to lower error for those with large historical angles. These controlled comparisons suggest that historical directional changes carry information about local guidance benefits.

From observation to policy. These observations motivate a simple allocation principle: assign weaker guidance to sites with more consistent historical directions and stronger guidance to those

with larger directional revisions. Its use in generation rests on the hypothesis that the association observed in the controlled reconstruction diagnostic remains informative along actual sampling trajectories. The following subsection implements this principle through a bounded mapping from historical directional feedback to token-wise guidance scales. We evaluate its effectiveness across generative modalities and examine its components through ablation studies.

4.2 Token-Wise Guidance from Trajectory Feedback

Following the allocation principle above, TrajCFG uses historical directional changes to determine the guidance scale at each latent site. Let $\mathbf{z}_k \in \mathbb{R}^{P\times C}$ denote the latent at sampling step k, where P is the number of sites and C is the number of channels; we omit the batch dimension. Each site corresponds to a token indexed by p, whose velocity is a C-dimensional vector. Depending on the backbone, the site axis may represent spatial, temporal, spatiotemporal, or generic sequence positions. TrajCFG operates on this flattened site axis without assuming a modality-specific latent geometry.

Historical directional feedback. At step $k \geq 2$, we compare the guided velocities at each site from the two most recent completed steps:

$$q_{k,p} = \cos_\delta\big(\widehat{\mathbf{v}}_{k-2,p}, \widehat{\mathbf{v}}_{k-1,p}\big), \tag{6}$$

where $\widehat{\mathbf{v}}_{j,p}$ is the guided velocity actually supplied to the solver at step j. We use the numerically stabilized cosine similarity

$$\cos_\delta(\mathbf{a},\mathbf{b}) = \mathrm{clip}_{[-1,1]}\left(\frac{\langle \mathbf{a},\mathbf{b}\rangle}{\max(\|\mathbf{a}\|_2,\delta)\max(\|\mathbf{b}\|_2,\delta)}\right), \tag{7}$$

with a norm floor $\delta > 0$. When both velocity norms exceed δ, $q_{k,p}$ is the cosine of their historical angle: higher similarity indicates more consistent directions, while lower similarity indicates larger directional changes. At a fixed site, this historical angle corresponds to θ_{k-2} in the preceding directional decomposition. Because both velocities come from completed steps, the feedback is available before selecting the current guidance scale.

Mapping feedback to guidance. We map historical directional change to guidance strength using a clipped linear function:

$$g(q) = s_{\min} + (s_{\max} - s_{\min})\,\mathrm{clip}_{[0,1]}\left(\frac{1-q}{1-\rho_{\mathrm{lo}}}\right), \tag{8}$$

where $0 \leq s_{\min} < s_{\max}$ specifies the guidance range and $-1 \leq \rho_{\mathrm{lo}} < 1$ determines when the scale reaches $s_{\max}$. At $q = 1$, the scale equals $s_{\min}$. As similarity decreases, guidance increases linearly with the directional-change statistic $1 - q$, reaching $s_{\max}$ at $q = \rho_{\mathrm{lo}}$ and remaining capped below this threshold. This mapping provides a simple bounded realization of the empirically motivated allocation principle.

Applying token-wise guidance. Let K denote the number of sampling updates, indexed from 0 to $K-1$, and let $\mathcal{A} \subseteq \{2,\dots,K-1\}$ be the set of steps at which adaptation is enabled. With adjacent-step feedback, the scale applied to token p is

$$s_{k,p} = \begin{cases} g(q_{k,p}), & k \in \mathcal{A}, \\ s_{\mathrm{base}}, & k \notin \mathcal{A}. \end{cases} \tag{9}$$

We then combine the current target- and reference-conditioned predictions as

$$\widehat{\mathbf{v}}_{k,p} = \mathbf{v}^-_{k,p} + s_{k,p}\big(\mathbf{v}^+_{k,p} - \mathbf{v}^-_{k,p}\big). \tag{10}$$

Each scale is shared across the channels of its token. The resulting guided velocity is passed to the original solver to update the latent and is retained as history for subsequent steps. The first two steps use s_{base} to initialize this history. As sampling proceeds, the feedback therefore reflects the trajectory produced by the guidance scales actually applied.

Longer-span feedback. We extend the adjacent-step rule with feedback from more distant historical predictions. Let $s^{(h)}_{k,p}$ denote the guidance scale obtained by comparing $\widehat{\mathbf{v}}_{k-1-h,p}$ with $\widehat{\mathbf{v}}_{k-1,p}$,

normalizing the angle to the adjacent-step interval, and applying the same guidance mapping. Appendix A.1.3 defines the interval normalization. At adaptive steps $k \in \mathcal{A}$, we replace $g(q_{k,p})$ in Eq. (9) with the fused guidance scale

$$s_{k,p} = \sum_{h \in \mathcal{H}} w_h s_{k,p}^{(h)}, \qquad w_h \geq 0, \quad \sum_{h \in \mathcal{H}} w_h = 1, \tag{11}$$

where $\mathcal{H}$ contains the available history spans. We assign decreasing weights to longer spans, prioritizing recent directional changes that more directly reflect the current sampling state. More distant history provides complementary context with a smaller contribution. For $k \geq 3$, we use adjacent-step ($h = 1$) and stride-2 ($h = 2$) feedback with weights $w_1 = 0.8$ and $w_2 = 0.2$. The additional span compares $\widehat{\mathbf{v}}_{k-3,p}$ with $\widehat{\mathbf{v}}_{k-1,p}$; at $k = 2$, we use $\mathcal{H} = \{1\}$ and $w_1 = 1$. Steps outside $\mathcal{A}$ retain s_{base}. This implementation retains three historical velocity tensors without extra model evaluations. Algorithm 1 in Appendix A.2 summarizes the complete sampling procedure.

5 Experiments

We compare TrajCFG with inference-time enhancement methods for text-to-video and image-to-video generation, ablate its guidance allocation rule, and evaluate three additional generation tasks.

5.1 Benchmarks

For T2V, we use VBench (Huang et al., 2024) and Dynamic-Bench (Jang et al., 2026); for I2V, we use VBench-I2V (Huang et al., 2025) and PAI-Bench-G (Zhou et al., 2026). We follow the official protocols for VBench and VBench-I2V. Dynamic-Bench and PAI-Bench-G use the supported dimensions under the VBench custom evaluation protocol; these scores are distinct from the official aggregate scores.

5.2 Backbones and Baselines

For T2V, we evaluate Wan2.1-1.3B, Wan2.2-5B, and Wan2.2-14B (Wan et al., 2025); for I2V, we use Cosmos-Predict2.5-2B (Ali et al., 2025) and both Wan2.2 backbones. We compare with the original sampling configuration (*Base*) and four training-free methods: CFG-Zero (Fan et al., 2025), CFG-Ctrl (Wang et al., 2026), PAG (Ahn et al., 2024), and Self-Refining Video Sampling (SRVS) (Jang et al., 2026). TrajCFG is labeled *Ours* in figures; guidance hyperparameters and step counts appear in Appendix Table 7.

5.3 Quantitative Comparison

Table 1: **Text-to-video generation results on VBench and Dynamic-Bench.** VBench reports the official Quality (Qual.), Semantic (Sem.), and Total scores. For Dynamic-Bench, Mot., Over., and Avg. denote Motion Smoothness, Overall Consistency, and the average over evaluated VBench dimensions under the custom protocol. Higher scores are better. Runtime is measured relative to Base ($1.0\times$). Best and second-best scores in each metric column are bold and underlined, respectively.

Method	Wan2.1-1.3B						Wan2.2-5B						Wan2.2-14B						Time
	VBench			Dynamic-Bench			VBench			Dynamic-Bench			VBench			Dynamic-Bench			
	Qual.	Sem.	Total	Mot.	Over.	Avg.	Qual.	Sem.	Total	Mot.	Over.	Avg.	Qual.	Sem.	Total	Mot.	Over.	Avg.	
Base	83.03	72.77	80.98	96.40	22.36	66.18	82.88	74.52	81.21	96.58	24.00	67.45	82.93	73.05	80.96	98.14	23.47	69.14	1.0×
CFG-Zero	82.69	<u>73.95</u>	80.94	97.13	23.81	66.64	83.79	74.84	82.00	96.68	24.19	67.94	83.32	75.27	81.71	98.02	25.18	69.41	1.0×
CFG-Ctrl	83.65	72.68	81.45	97.28	**24.18**	66.90	<u>84.12</u>	<u>76.75</u>	<u>82.64</u>	96.95	24.09	67.42	83.09	76.11	81.70	<u>98.38</u>	24.79	69.50	1.0×
PAG	83.24	73.92	81.38	97.15	23.89	67.25	83.37	76.14	81.93	96.91	24.11	67.63	83.07	73.31	81.12	98.11	24.90	69.46	1.5×
SRVS	<u>83.69</u>	73.10	<u>81.57</u>	<u>97.38</u>	24.15	<u>67.48</u>	83.93	76.63	82.47	96.87	<u>24.30</u>	<u>68.11</u>	<u>84.24</u>	<u>76.59</u>	<u>82.71</u>	98.26	<u>25.24</u>	<u>69.86</u>	1.5×
TrajCFG	**84.06**	**75.37**	**82.32**	**97.53**	<u>24.17</u>	**67.66**	**84.82**	**77.45**	**83.35**	**97.02**	**24.33**	**68.30**	**84.41**	**77.39**	**83.01**	**98.43**	**25.31**	**70.27**	1.0×

Text-to-video generation. Table 1 reports the T2V results. TrajCFG achieves the highest official VBench Quality, Semantic, and Total scores on all three backbones, improving Total over the strongest competitor by 0.30–0.75 points. It also obtains the highest Dynamic-Bench average

Table 2: **Image-to-video generation results on VBench-I2V and PAI-Bench-G.** Img. (Image Alignment) denotes the mean of I2V Subject and I2V Background; Mot. denotes Motion Smoothness; and Qual. denotes the mean of Aesthetic Quality and Imaging Quality. Higher scores are better. Runtime is measured relative to Base (1.0×). Best and second-best scores in each metric column are bold and underlined, respectively.

Method	Cosmos-Predict2.5-2B						Wan2.2-5B						Wan2.2-14B						Time
	VBench-I2V			PAI-Bench-G			VBench-I2V			PAI-Bench-G			VBench-I2V			PAI-Bench-G			
	Img.	Mot.	Qual.	Img.	Mot.	Qual.	Img.	Mot.	Qual.	Img.	Mot.	Qual.	Img.	Mot.	Qual.	Img.	Mot.	Qual.	
Base	98.46	99.09	66.87	96.76	99.09	61.32	98.74	99.02	67.49	97.14	98.72	59.56	98.51	98.77	66.87	96.69	98.50	58.86	1.0×
CFG-Zero	98.53	**99.14**	67.07	96.83	99.13	61.61	98.25	98.92	67.17	96.86	98.72	60.00	97.69	98.69	65.65	96.44	98.48	58.52	1.0×
CFG-Ctrl	98.13	98.97	66.88	95.85	98.37	60.93	98.37	99.10	66.37	97.23	98.97	59.81	98.48	98.79	66.63	96.42	98.52	58.18	1.0×
PAG	98.46	99.10	66.87	96.71	99.08	61.62	98.71	99.02	67.37	96.95	98.69	60.26	98.20	98.76	66.85	96.66	98.53	58.69	1.5×
SRVS	98.12	99.11	66.88	96.03	99.10	**61.65**	98.81	99.09	67.92	96.75	98.50	**60.39**	98.57	98.79	66.84	96.73	98.61	58.86	1.5×
TrajCFG	**98.55**	99.13	**67.16**	**96.89**	**99.14**	61.54	**98.97**	**99.32**	**68.15**	**97.54**	**98.99**	59.69	**98.63**	**98.90**	**67.02**	**96.94**	**98.65**	**58.90**	1.0×

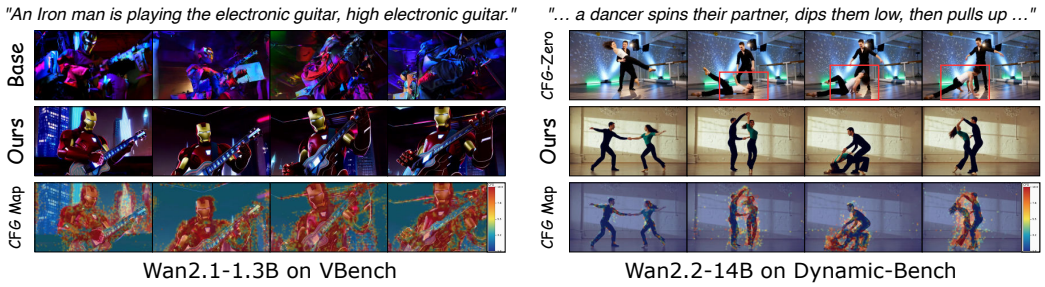

(a) **Text-to-video (T2V) generation.**

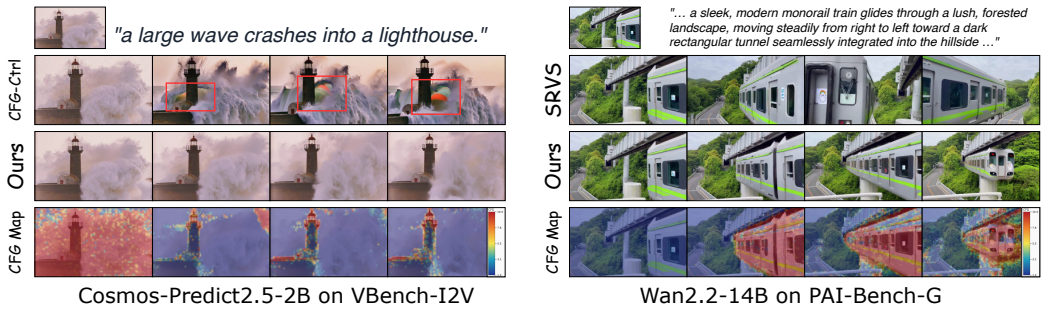

(b) **Image-to-video (I2V) generation.**

Figure 4: **Qualitative comparisons for video generation.** (a) TrajCFG is compared with Base on Wan2.1-1.3B/VBench (left) and CFG-Zero on Wan2.2-14B/Dynamic-Bench (right). (b) TrajCFG is compared with CFG-Ctrl on Cosmos-Predict2.5-2B/VBench-I2V (left) and SRVS on Wan2.2-14B/PAI-Bench-G (right); input images appear beside the prompts. Guidance-scale rows visualize TrajCFG's token-wise scales at sampling step $k = 5$, with warmer colors indicating stronger guidance. Red boxes highlight artifacts in the competing-method sequences.

and Motion Smoothness on every backbone, with average-score gains of 0.18–0.41 points over the strongest competitor.

Image-to-video generation. TrajCFG improves every metric summarized in Table 2 over Base. It achieves the highest VBench-I2V Image Alignment and Quality on all backbones and the highest PAI-Bench-G Image Alignment and Motion Smoothness. PAI-Bench-G Quality gains are smaller and not uniformly best against all competitors, while the clearest advantages remain reference-image preservation and motion smoothness. Complete per-dimension video results are provided in the appendix.

5.4 Qualitative Comparison

Text-to-video generation. In Figure 4a, Base fails to preserve Iron Man's characteristic appearance, while CFG-Zero abruptly changes the woman's orientation from front-facing to back-facing after she contacts the floor. TrajCFG better preserves character appearance and continuity of body orientation. Its guidance maps complement the historical angle maps in Figure 2.

Image-to-video generation. In Figure 4b, CFG-Ctrl introduces unnatural iridescent colors in the seawater, while the SRVS monorail leaves its track. TrajCFG reduces these color artifacts and keeps the train aligned with the track. Further T2V/I2V cases appear in Appendix Figures 5 and 6.

5.5 Ablation Study

Table 3 reports ablations of adaptation granularity, feedback–allocation correspondence, and history-span fusion on VBench using Wan2.1-1.3B and Wan2.2-5B.

Table 3: **Ablation study of TrajCFG for text-to-video generation on VBench.** We report the official Quality, Semantic, and Total scores (higher is better). Fusion weights (w_1, w_2) correspond to adjacent-step ($h = 1$) and stride-2 ($h = 2$) guidance estimates, respectively; TrajCFG (Full) uses $(0.8, 0.2)$. Best and second-best scores in each metric column are bold and underlined, respectively.

	Variant	Wan2.1-1.3B			Wan2.2-5B		
		Quality	**Semantic**	**Total**	**Quality**	**Semantic**	**Total**
	TrajCFG (Full)	**84.06**	**75.37**	**82.32**	<u>84.82</u>	**77.45**	**83.35**
Q1	Time-only	<u>83.94</u>	73.17	81.78	84.62	<u>76.94</u>	83.08
	Sample-wise	83.83	74.02	81.87	84.48	76.34	82.85
Q2	Shuffled allocation	83.77	73.67	81.75	84.43	76.73	82.89
	Reversed allocation	83.70	72.96	81.55	84.55	75.82	82.81
Q3	Adjacent-only $(1, 0)$	83.81	<u>74.41</u>	<u>81.93</u>	**84.85**	76.93	<u>83.27</u>
	Stride-2-only $(0, 1)$	83.86	73.39	81.77	84.65	75.64	82.85

Q1: Is token-wise guidance adaptation beneficial? We compare TrajCFG with *Time-only*, a fixed schedule averaged over a calibration set, and *Sample-wise*, which averages each online scale map across tokens. TrajCFG improves every aggregate score; its Total-score gains over these variants are 0.45–0.54 points on Wan2.1-1.3B and 0.27–0.50 points on Wan2.2-5B, supporting adaptation at individual latent sites.

Q2: Does feedback-aligned guidance allocation matter? Using the same initial noise, we replay cached scale maps while either shuffling token assignments or reversing their strength order; scales are not recomputed along the altered trajectories. Both variants underperform the original allocation on every aggregate score, reducing Total by 0.46–0.77 points. Thus, performance depends on matching guidance to each token's feedback, not only on the per-step distribution of scales.

Q3: Does combining history spans improve performance? Fusion $(w_1, w_2) = (0.8, 0.2)$ outperforms *Adjacent-only* and *Stride-2-only* in Semantic and Total scores on both backbones, improving Total over Adjacent-only by 0.39 and 0.08 points. Adjacent-only has a slightly higher Quality score on Wan2.2-5B. These results suggest that longer-span feedback modestly complements the primary adjacent-step signal.

5.6 Evaluation across Generative Modalities

Using each backbone's original configuration as Base, we apply the same historical feedback rule across three latent representations; qualitative results appear in Appendix E, Figures 7–9.

Text-to-audio-video generation. We evaluate TrajCFG for text-to-audio-video (T2AV) generation using LTX-2.3 (HaCohen et al., 2026) and Ovi (Low et al., 2025) on VABench (Hua et al., 2026) and T2AV-Compass (Cao et al., 2026). TrajCFG improves all metrics summarized in Table 4 on both backbones, although some additional metrics decline in the full evaluation (Appendix F). For

Table 4: **Text-to-audio-video results.** Audio Aes., T-A, A-V, VT, AA, and DS denote Audio Aesthetics, text–audio alignment, audio–visual alignment, Video Technological, Audio Attribute, and DeSync. Lower is better for Desync and DS; best scores within each backbone are bold.

Backbone	Method	VABench					T2AV-Compass				
		Visual Realism↑	Audio Aes.↑	T-A↑	A-V↑	Desync↓	VT↑	AA↑	T-A↑	A-V↑	DS↓
LTX-2.3	Base	4.615	3.345	0.3118	0.2490	0.5911	7.295	6.709	0.1722	0.2469	0.7140
	TrajCFG	**4.638**	**3.351**	**0.3202**	**0.2567**	**0.5685**	**7.510**	**6.738**	**0.1735**	**0.2481**	**0.7016**
Ovi	Base	4.522	3.282	0.3288	0.1564	0.7659	6.714	6.676	0.0897	0.1596	0.9624
	TrajCFG	**4.547**	**3.364**	**0.3316**	**0.1654**	**0.7080**	**7.199**	**6.733**	**0.0922**	**0.1622**	**0.9356**

Table 5: **Text-to-image results.** Higher is better; best scores within each backbone are bold, including ties.

Backbone	Method	GenEval							DPG-Bench					
		Single Obj.↑	Two Obj.↑	Counting↑	Colors↑	Position↑	Attr. Bind.↑	Overall↑	Global↑	Entity↑	Attribute↑	Relation↑	Other↑	Overall↑
Qwen-Image	Base	0.9875	0.9318	0.6094	0.8324	**0.4950**	0.5825	0.7412	89.39	90.84	89.48	91.76	90.63	87.11
	TrajCFG	**0.9906**	**0.9394**	**0.6188**	**0.8617**	**0.4950**	**0.6050**	**0.7517**	**91.03**	**93.20**	**90.40**	**93.21**	**94.35**	**87.19**
SD3-Medium	Base	0.9969	0.8611	0.6125	0.8351	0.3125	0.5650	0.7021	86.98	84.24	86.26	87.30	85.19	83.56
	TrajCFG	**1.0000**	**0.8813**	**0.6344**	**0.8484**	**0.3225**	**0.5975**	**0.7140**	**88.35**	**86.37**	**86.55**	**89.38**	**86.15**	**83.99**

Table 6: **Image-to-3D results on GSO.** Metrics follow Hunyuan3D 2.0; -T and -I denote mesh–text and mesh–image similarity. Arrows indicate direction; best scores are bold.

Backbone	Method	Shape				Texture			
		ULIP-T↑	ULIP-I↑	Uni3D-T↑	Uni3D-I↑	CMMD↓	FID_{CLIP} ↓	CLIP-score↑	LPIPS↓
Hunyuan3D 2.0	Base	0.1753	0.2277	0.2190	0.2714	0.0571	0.0606	0.2503	0.1015
	TrajCFG	**0.1762**	**0.2290**	**0.2198**	**0.2719**	**0.0532**	**0.0581**	**0.2505**	**0.1008**

Ovi, T2AV-Compass VT increases by 0.485, while VABench Desync decreases by 0.0579, showing complementary gains in visual quality and audio–visual synchronization.

Text-to-image generation. We evaluate TrajCFG for text-to-image (T2I) generation using Qwen-Image (Wu et al., 2025) and SD3-Medium (Esser et al., 2024) on GenEval (Ghosh et al., 2023) and DPG-Bench (Hu et al., 2024). As shown in Table 5, TrajCFG improves or matches Base on every metric, with consistent gains in counting, color, and attribute binding on GenEval and more modest Overall gains on DPG-Bench. GenEval Overall increases by 0.0105 and 0.0119 for Qwen-Image and SD3-Medium, respectively; the corresponding DPG-Bench Overall gains are 0.08 and 0.43 points.

Image-to-3D generation. We evaluate TrajCFG for image-to-3D generation using Hunyuan3D 2.0 (Zhao et al., 2025) on GSO (Downs et al., 2022). Table 6 shows improvements over Base on all eight shape and texture metrics, including lower CMMD and FID_{CLIP}. In particular, CMMD and FID_{CLIP} decrease by 0.0039 and 0.0025, while all four shape similarities also improve. Together, the three evaluations support applicability across different latent representations.

6 Conclusion

We introduced TrajCFG, a training-free method that adapts token-wise CFG scales using historical changes in guided velocity direction. By combining adjacent-step and interval-normalized longer-span feedback, TrajCFG allocates guidance from sampling history without additional model evaluations. Experiments show improvements across video backbones and applicability to audiovisual, image, and 3D generation. These results support historical directional feedback as a useful signal for local guidance allocation across latent representations.

Historical directional feedback tracks revisions of the sampler's predictions rather than directly estimating generation error. Its effectiveness beyond the evaluated models and sampling configurations remains to be established, particularly for very low-step sampling and alternative generative parameterizations. Future work could investigate self-calibrating guidance rules that reduce manual parameter selection without additional training or model evaluations.

REFERENCES

Donghoon Ahn, Hyoungwon Cho, Jaewon Min, Wooseok Jang, Jungwoo Kim, SeonHwa Kim, Hyun Hee Park, Kyong Hwan Jin, and Seungryong Kim. Self-rectifying diffusion sampling with perturbed-attention guidance. In *ECCV*, 2024.

Arslan Ali, Junjie Bai, Maciej Bala, Yogesh Balaji, Aaron Blakeman, Tiffany Cai, Jiaxin Cao, Tianshi Cao, Elizabeth Cha, Yu-Wei Chao, et al. World simulation with video foundation models for physical ai. *arXiv preprint arXiv:2511.00062*, 2025.

Lichen Bai, Shitong Shao, Zipeng Qi, Zhiqiang Xu, Haoyi Xiong, Zeke Xie, et al. Zigzag diffusion sampling: Diffusion models can self-improve via self-reflection. In *ICLR*, 2025.

Zhe Cao, Tao Wang, Jiaming Wang, Yanghai Wang, Yuanxing Zhang, Jiahao Wang, Jialu Chen, Miao Deng, Yubin Guo, Chenxi Liao, et al. T2av-compass: Towards unified evaluation for text-to-audio-video generation. In *ICML*, 2026.

Angela Castillo, Jonas Kohler, Juan C Pérez, Juan Pablo Pérez, Albert Pumarola, Bernard Ghanem, Pablo Arbeláez, and Ali Thabet. Adaptive guidance: Training-free acceleration of conditional diffusion models. In *AAAI*, 2025.

Xianhang Cheng, Yujian Zheng, Zhenyu Xie, Tingting Liao, and Hao Li. Steering video diffusion transformers with massive activations. *arXiv preprint arXiv:2603.17825*, 2026.

Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. Cfg++: Manifold-constrained classifier free guidance for diffusion models. In *ICLR*, 2025.

Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *ICRA*, 2022.

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 *ICML*, 2024.

Weichen Fan, Amber Yijia Zheng, Raymond A Yeh, and Ziwei Liu. Cfg-zero*: Improved classifier-free guidance for flow matching models. *arXiv preprint arXiv:2503.18886*, 2025.

Jiayang Gao, Tianyi Zheng, Jiayang Zou, Fengxiang Yang, Shice Liu, Luyao Fan, Zheyu Zhang, Hao Zhang, Jinwei Chen, Peng-Tao Jiang, et al. C2fg: Control classifier-free guidance via score discrepancy analysis. *arXiv preprint arXiv:2603.08155*, 2026.

Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. In *NeurIPS*, 2023.

Yoav HaCohen, Benny Brazowski, Nisan Chiprut, Yaki Bitterman, Andrew Kvochko, Avishai Berkowitz, Daniel Shalem, Daphna Lifschitz, Dudu Moshe, Eitan Porat, et al. Ltx-2: Efficient joint audio-visual foundation model. *arXiv preprint arXiv:2601.03233*, 2026.

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

Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024.

Daili Hua, Xizhi Wang, Bohan Zeng, Xinyi Huang, Hao Liang, Junbo Niu, Xinlong Chen, Quanqing Xu, and Wentao Zhang. Vabench: A comprehensive benchmark for audio-video generation. In *CVPR*, 2026.

Ziqi Huang, Yinan He, Jiashuo Yu, Fan Zhang, Chenyang Si, Yuming Jiang, Yuanhan Zhang, Tianxing Wu, Qingyang Jin, Nattapol Chanpaisit, et al. Vbench: Comprehensive benchmark suite for video generative models. In *CVPR*, 2024.

Ziqi Huang, Fan Zhang, Xiaojie Xu, Yinan He, Jiashuo Yu, Ziyue Dong, Qianli Ma, Nattapol Chanpaisit, Chenyang Si, Yuming Jiang, et al. Vbench++: Comprehensive and versatile benchmark suite for video generative models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Junha Hyung, Kinam Kim, Susung Hong, Min-Jung Kim, and Jaegul Choo. Spatiotemporal skip guidance for enhanced video diffusion sampling. In *CVPR*, 2025.

Sangwon Jang, Taekyung Ki, Jaehyeong Jo, Saining Xie, Jaehong Yoon, and Sung Ju Hwang. Self-refining video sampling. In *ICML*, 2026.

Cheng Jin, Qitan Shi, and Yuantao Gu. Stage-wise dynamics of classifier-free guidance in diffusion models. In *ICLR*, 2026.

Mingi Kwon, Jaeseok Jeong, Yi Ting Hsiao, Youngjung Uh, et al. Tcfg: Tangential damping classifier-free guidance. In *CVPR*, 2025.

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 *NeurIPS*, 2024.

Dohun Lee, Bryan Sangwoo Kim, Geon Yeong Park, and Jong Chul Ye. Videoguide: Improving video diffusion models without training through a teacher's guide. In *CVPR*, 2025.

Haosen Li, Wenshuo Chen, Lei Wang, Shaofeng Liang, Bowen Tian, Soning Lai, and Yutao Yue. Delta score matters! spatial adaptive multi guidance in diffusion models. *arXiv preprint arXiv:2604.26503*, 2026.

Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. In *ICLR*, 2023.

Chetwin Low, Weimin Wang, and Calder Katyal. Ovi: Twin backbone cross-modal fusion for audio-video generation. *arXiv preprint arXiv:2510.01284*, 2025.

Yan Luo, Ahmadou Aidara, Jingyi Lu, Jeremy Moebel, Kai Han, and Mengyu Wang. Vags: Velocity adaptive guidance scale for image editing and generation. *arXiv preprint arXiv:2605.15661*, 2026.

Avishag Nevo and Tamir Hazan. Pathguide: Dynamic classifier-free guidance via on-policy transport alignment. *arXiv preprint arXiv:2608.29107*, 2026.

Pinelopi Papalampidi, Olivia Wiles, Ira Ktena, Aleksandar Shtedritski, Emanuele Bugliarello, Ivana Kajić, Isabela Albuquerque, and Aida Nematzadeh. Dynamic classifier-free diffusion guidance via online feedback. In *ICLR*, 2026.

Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.

Seyedmorteza Sadat, Otmar Hilliges, and Romann M Weber. Eliminating oversaturation and artifacts of high guidance scales in diffusion models. In *ICLR*, 2024.

Seyedmorteza Sadat, Farnood Salehi, and Romann Weber. Higs: History-guided sampling for plug-and-play enhancement of diffusion models. In *International Conference on Learning Representations*, volume 2026, pp. 63432–63462, 2026.

Ariel Shaulov, Itay Hazan, Lior Wolf, and Hila Chefer. Flowmo: Variance-based flow guidance for coherent motion in video generation. In *NeurIPS*, 2026.

Dazhong Shen, Guanglu Song, Zeyue Xue, Fu-Yun Wang, and Yu Liu. Rethinking the spatial inconsistency in classifier-free diffusion guidance. In *CVPR*, 2024.

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.

Hanyang Wang, Yiyang Liu, Jiawei Chi, Fangfu Liu, Ran Xue, and Yueqi Duan. Cfg-ctrl: Control-based classifier-free diffusion guidance. In *CVPR*, 2026.

Xi Wang, Nicolas Dufour, Nefeli Andreou, Marie-Paule Cani, Victoria Fernández Abrevaya, David Picard, and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *arXiv preprint arXiv:2404.13040*, 2024.

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.

Tianxing Wu, Chenyang Si, Yuming Jiang, Ziqi Huang, and Ziwei Liu. Freeinit: Bridging initialization gap in video diffusion models. In *ECCV*, 2024.

Yilun Xu, Mingyang Deng, Xiang Cheng, Yonglong Tian, Ziming Liu, and Tommi Jaakkola. Restart sampling for improving generative processes. In *NeurIPS*, 2023.

Huixuan Zhang and Xiaojun Wan. How much to guide: Revisiting adaptive guidance in classifier-free guidance text-to-vision diffusion models. In *MM Asia*, 2025.

Xu Zhang, Yu Lu, Ruijie Quan, Zhaozheng Chen, Bohan Wang, and Yi Yang. Speclor: Spectral lookahead rectification for motion-coherent text-to-video generation. *arXiv preprint arXiv:2606.11969*, 2026.

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.

Fengzhe Zhou, Jiannan Huang, Jialuo Li, Deva Ramanan, and Humphrey Shi. Pai-bench: A comprehensive benchmark for physical ai. In *CVPR*, 2026.

AI USE STATEMENT

Generative AI tools were used solely for language editing and polishing of the manuscript. They were not used to develop the research ideas, method, experimental design, code, data, results, or conclusions. The authors reviewed and verified all AI-assisted edits and take full responsibility for the content of the paper.

ETHICS STATEMENT

This work uses publicly available generative models and benchmarks and does not involve human-subject studies or the collection of private or personally identifiable data. TrajCFG inherits the potential societal risks of the underlying generative models, including the creation of misleading or harmful content. We encourage its use in accordance with the licenses and responsible-use policies of the corresponding models and datasets.

REPRODUCIBILITY STATEMENT

All experiments use publicly available backbones and benchmarks. The method and evaluation protocols are specified in Sections 4 and 5, respectively. Appendix A.2 provides the complete sampling procedure (Algorithm 1), and Appendix B reports task- and backbone-specific settings (Table 7), including the base CFG scale, $s_{\max}$, $s_{\min}$, ρ_{lo}, and the default number of denoising steps. Appendix F further provides the complete per-dimension results underlying the main-paper tables. Together, these materials provide the information needed to reproduce the reported experiments.

A Additional Method Details

A.1 Endpoint Revision and Directional Decomposition

This section derives the identities used to interpret directional feedback and explains the interval normalization used for longer history spans. Unless otherwise stated, all quantities refer to a fixed latent site, whose index is omitted.

A.1.1 Endpoint Revision and Solver Dependence

For the affine path $\mathbf{z}_\tau = (1-\tau)\mathbf{x}+\tau\boldsymbol{\epsilon}$, the path velocity is $\boldsymbol{\epsilon}-\mathbf{x}$, so $\mathbf{x} = \mathbf{z}_\tau - \tau(\boldsymbol{\epsilon}-\mathbf{x})$. Substituting the guided prediction for this velocity defines the sampler's clean-endpoint estimate:

$$\mathbf{D}_k = \mathbf{z}_k - \tau_k \widehat{\mathbf{v}}_k. \tag{12}$$

Let $\Delta\tau_k = \tau_{k+1} - \tau_k$. Under the Euler update $\mathbf{z}_{k+1} = \mathbf{z}_k + \Delta\tau_k \widehat{\mathbf{v}}_k$,

$$\begin{aligned} \mathbf{D}_{k+1} - \mathbf{D}_k &= \mathbf{z}_{k+1} - \mathbf{z}_k - \tau_{k+1}\widehat{\mathbf{v}}_{k+1} + \tau_k \widehat{\mathbf{v}}_k \\ &= \Delta\tau_k \widehat{\mathbf{v}}_k - \tau_{k+1}\widehat{\mathbf{v}}_{k+1} + \tau_k \widehat{\mathbf{v}}_k \\ &= -\tau_{k+1}\big(\widehat{\mathbf{v}}_{k+1} - \widehat{\mathbf{v}}_k\big). \end{aligned} \tag{13}$$

Thus, under Euler sampling, changes in guided velocity correspond directly to revisions of the clean-endpoint estimate.

For a general solver, define its displacement relative to the Euler update as

$$\mathbf{e}_k^{\mathrm{sol}} = \mathbf{z}_{k+1} - \mathbf{z}_k - \Delta\tau_k \widehat{\mathbf{v}}_k. \tag{14}$$

The same calculation gives

$$\mathbf{D}_{k+1} - \mathbf{D}_k = -\tau_{k+1}\big(\widehat{\mathbf{v}}_{k+1} - \widehat{\mathbf{v}}_k\big) + \mathbf{e}_k^{\mathrm{sol}}. \tag{15}$$

This identity also covers multistep solvers, with the effects of their internal history included in $\mathbf{e}_k^{\mathrm{sol}}$. The residual is zero for Euler updates and measures departure from the Euler displacement; it is not a numerical error estimate with respect to the exact ODE solution. Both identities concern revisions of the sampler's own endpoint estimate, rather than its error against a ground-truth target.

A.1.2 Magnitude–Direction Decomposition

For successive nonzero guided velocities, define

$$n_k = \|\widehat{\mathbf{v}}_k\|_2, \qquad c_k = \frac{\langle \widehat{\mathbf{v}}_{k+1}, \widehat{\mathbf{v}}_k \rangle}{n_{k+1} n_k} = \cos\theta_k. \tag{16}$$

Here c_k compares steps k and $k+1$. The historical feedback used at step k in the main text instead compares steps $k-2$ and $k-1$, and equals c_{k-2} when the norm floor is inactive. Expanding the squared velocity difference yields

$$\begin{aligned} \|\widehat{\mathbf{v}}_{k+1} - \widehat{\mathbf{v}}_k\|_2^2 &= n_{k+1}^2 + n_k^2 - 2n_{k+1}n_k c_k \\ &= (n_{k+1} - n_k)^2 + 2n_{k+1}n_k(1 - c_k). \end{aligned} \tag{17}$$

The first term measures changes in magnitude, and the second measures magnitude-weighted changes in direction. Dividing the second term by $2n_{k+1}n_k$ isolates the directional statistic $1 - c_k$. Equivalently, with unit directions $\mathbf{u}_k = \widehat{\mathbf{v}}_k / n_k$,

$$1 - c_k = \frac{1}{2}\|\mathbf{u}_{k+1} - \mathbf{u}_k\|_2^2. \tag{18}$$

This statistic is invariant to independent positive rescaling of the velocities. It isolates their directional change, while the full velocity difference and endpoint revision also depend on the velocity magnitudes and the sampling time.

A.1.3 Sampling Intervals and Longer-Span Normalization

Suppose the unit direction of a locally nonzero velocity admits a C^2 curve $\mathbf{u}(\tau)$. For sufficiently small $\Delta > 0$, Taylor expansion gives

$$\mathbf{u}(\tau - \Delta) = \mathbf{u}(\tau) - \Delta\dot{\mathbf{u}}(\tau) + O(\Delta^2). \tag{19}$$

Using the unit-direction identity above,

$$1 - \langle \mathbf{u}(\tau), \mathbf{u}(\tau - \Delta)\rangle = \frac{\Delta^2}{2}\|\dot{\mathbf{u}}(\tau)\|_2^2 + O(\Delta^3). \tag{20}$$

Directional change therefore depends on both the local directional rate and the observation interval. For small angles, $1 - \cos\theta \approx \theta^2/2$, motivating angle rescaling in proportion to the interval length when comparing nearby history spans.

For span h, let $\Delta_k^{(h)} = |\tau_{k-1} - \tau_{k-1-h}|$ and $q_{k,p}^{(h)} = \cos_\delta(\widehat{\mathbf{v}}_{k-1-h,p}, \widehat{\mathbf{v}}_{k-1,p})$. We normalize its angle to the adjacent-step interval and apply the same guidance mapping:

$$\widetilde{q}_{k,p}^{(h)} = \cos\left(\frac{\Delta_k^{(1)}}{\Delta_k^{(h)}} \arccos q_{k,p}^{(h)}\right), \qquad s_{k,p}^{(h)} = g(\widetilde{q}_{k,p}^{(h)}). \tag{21}$$

For $h = 1$, this recovers the adjacent-step feedback. The rescaling accounts for interval length under a local smoothness approximation; signals from different spans can still differ because the direction evolves over their respective intervals.

Scope of interpretation. The exact directional identities and rescaling invariance apply when the velocity norms are nonzero and the numerical norm floor is inactive. The feedback measures directional change over an observed interval, rather than step-size-independent curvature. These relationships explain the signal and its normalization; the allocation of stronger guidance to larger directional changes is motivated by the local response diagnostic and evaluated through generation experiments and ablations.

A.2 Sampling Algorithm

Algorithm 1 presents the sampling procedure with adjacent-step and longer-span feedback. The first two steps use the baseline guidance scale to initialize the velocity history. At subsequent adaptive steps, guidance is determined from previously computed guided velocities. Adjacent-step feedback is used alone at $k = 2$; from $k \geq 3$, it is combined with interval-normalized longer-span feedback using weights 0.8 and 0.2.

All cosine similarities are computed token-wise, and the angle normalization and guidance mapping are applied element-wise. The operator $\odot$ broadcasts each token's guidance scale across its channels. $\mathrm{SolverStep}$ denotes the original solver update, including any internal solver state. We assume a strictly decreasing sampling schedule $\tau_0 > \cdots > \tau_K$.

The guidance estimates are fused after applying g to each span's feedback. At non-adaptive steps, the baseline scale is used. Every step updates the cache with the guided velocity actually passed to the solver, so subsequent feedback reflects the realized sampling trajectory.

Algorithm 1 TrajCFG with Adjacent-Step and Longer-Span Feedback

Require: Initial latent $\mathbf{z}_0$; schedule $\{\tau_k\}_{k=0}^{K}$
Require: Model f_θ; target/reference conditions y^+, y^-
Require: Base scale s_{base}; adaptive steps $\mathcal{A} \subseteq \{2, \dots, K-1\}$
Require: Guidance mapping g from Eq. (8); stabilized cosine $\cos_\delta$ from Eq. (7)
Ensure: Final latent $\mathbf{z}_K$

Initialize an empty velocity cache $\mathcal{V}$
for $k = 0, \dots, K-1$ **do**
  $\mathbf{s}_k \leftarrow s_{\text{base}} \mathbf{1}_P$
  if $k \in \mathcal{A}$ **then**
    $\mathbf{q}^{\text{short}} \leftarrow \cos_\delta\big(\mathcal{V}[k-2], \mathcal{V}[k-1]\big)$
    $\mathbf{s}_k \leftarrow g(\mathbf{q}^{\text{short}})$
    if $k \geq 3$ **then**
      $\mathbf{q}^{\text{long}} \leftarrow \cos_\delta\big(\mathcal{V}[k-3], \mathcal{V}[k-1]\big)$
      $\Delta_{\text{short}} \leftarrow |\tau_{k-1} - \tau_{k-2}|$
      $\Delta_{\text{long}} \leftarrow |\tau_{k-1} - \tau_{k-3}|$
      $\widetilde{\mathbf{q}}^{\text{long}} \leftarrow \cos\Big(\frac{\Delta_{\text{short}}}{\Delta_{\text{long}}} \arccos \mathbf{q}^{\text{long}}\Big)$
      $\mathbf{s}_k \leftarrow 0.8\, \mathbf{s}_k + 0.2\, g\big(\widetilde{\mathbf{q}}^{\text{long}}\big)$
    end if
  end if
  $\mathbf{v}_k^+ \leftarrow f_\theta(\mathbf{z}_k, \tau_k; y^+)$
  $\mathbf{v}_k^- \leftarrow f_\theta(\mathbf{z}_k, \tau_k; y^-)$
  $\widehat{\mathbf{v}}_k \leftarrow \mathbf{v}_k^- + \mathbf{s}_k \odot (\mathbf{v}_k^+ - \mathbf{v}_k^-)$
  $\mathbf{z}_{k+1} \leftarrow \text{SolverStep}(\mathbf{z}_k, \widehat{\mathbf{v}}_k, \tau_k, \tau_{k+1})$
  $\mathcal{V}[k] \leftarrow \widehat{\mathbf{v}}_k$
  Retain only the latest three entries in $\mathcal{V}$
end for
return $\mathbf{z}_K$

B HYPERPARAMETER SETTINGS

Table 7 summarizes the task- and backbone-specific settings. Here, s_{base} is the baseline CFG scale, $s_{\min}$ and $s_{\max}$ define the adaptive guidance range, ρ_{lo} is the lower similarity threshold, and the final column gives the default number of denoising steps.

Table 7: Hyperparameter settings for each task and backbone.

Task	Backbone	s_{base}	$s_{\min}$	$s_{\max}$	ρ_{lo}	Steps
T2V	Wan2.1-1.3B	6.0	5.0	10.0	0.95	40
	Wan2.2-5B	5.0	3.0	10.0	0.95	40
	Wan2.2-14B	3.5	1.0	10.0	0.90	40
I2V	Cosmos-Predict2.5-2B	5.0	4.0	9.0	0.90	35
	Wan2.2-5B	5.0	2.0	10.0	0.95	40
	Wan2.2-14B	2.5	1.0	7.0	0.95	40
T2AV	LTX-2.3	5.0	4.0	9.0	0.90	30
	Ovi	3.5	3.5	10.0	0.95	50
T2I	Qwen-Image	5.0	4.0	9.0	0.95	30
	SD3-Medium	7.0	6.0	9.0	0.95	28
I2-3D	Hunyuan3D 2.0 (Shape)	7.0	5.0	9.0	0.95	50
	Hunyuan3D 2.0 (Texture)	2.0	2.0	3.0	0.95	50

C ADDITIONAL TEXT-TO-VIDEO QUALITATIVE COMPARISONS

Figure 5 provides six additional comparisons spanning appearance, physical consistency, temporal stability, and action semantics. From upper left to lower right, Base retains color despite a black-and-white prompt; produces water splashes in distant dry regions; fails to maintain the snowy setting and depicts the motorcycle moving backward; causes the cat's paw to intersect the yarn ball; abruptly darkens the guitar scene; and turns a single slicing action into two tomato slices. TrajCFG better follows the requested appearance while improving spatial localization, object interaction, temporal consistency, and action–result correspondence.

D ADDITIONAL IMAGE-TO-VIDEO QUALITATIVE COMPARISONS

Figure 6 provides six comparisons of physical interaction, material dynamics, structural stability, object continuity, camera control, and manipulation. From upper left to lower right, Base lets the beer spill past the woman's mouth instead of being drunk; makes the honey behave like a solid blob rather than a viscous liquid that spreads; introduces artifacts and structural distortions; causes a cup to appear spontaneously in the woman's hand; fails to execute the requested zoom-out; and lifts the supporting dish together with the bowl. TrajCFG yields more coherent liquid behavior, geometry, object transfer, camera motion, and fine-grained physical interactions.

E CROSS-MODAL QUALITATIVE COMPARISONS

Figure 7 illustrates complementary visual and audio–visual failures in Base. In the upper-left example, squeezing the bottle produces visibly distorted fingers. In the upper-right example, the subject moves his mouth, yet the accompanying audio contains no speech. In the lower-left example, Base omits the prompted kitten meow. Finally, in the lower-right example, Base depicts the hand contacting a doorbell while generating the sound of knocking on wood, creating a cross-modal content mismatch. TrajCFG reduces the hand artifact and more faithfully aligns the visible actions and prompted events with the generated audio.

Figure 8 illustrates improvements in prompt adherence across both backbones. Representative Base failures include an implausible pair of scissors, an incorrect dog–teddy-bear relation, swapped colors between the toilet and suitcase, three circles instead of two, an extra baseball, two bats instead of

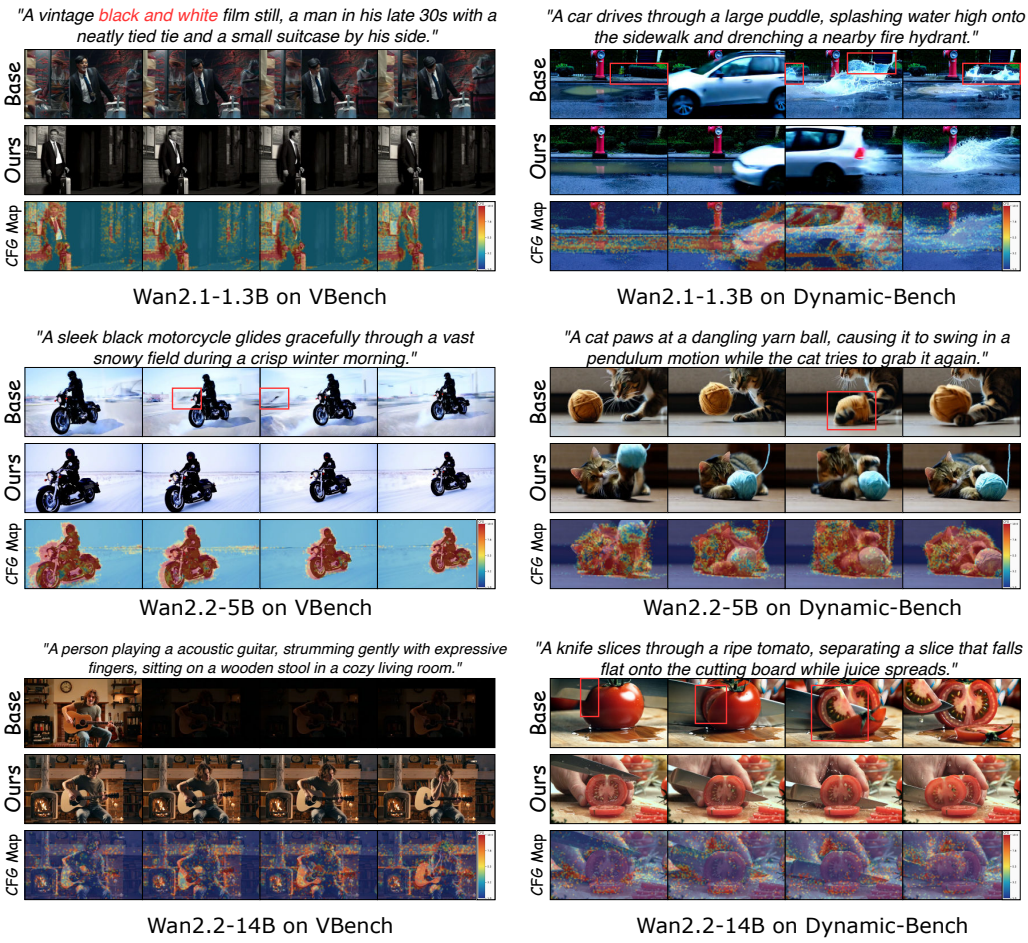

Figure 5: **Additional qualitative comparisons for text-to-video generation.** Results are shown for Wan2.1-1.3B, Wan2.2-5B, and Wan2.2-14B on VBench and Dynamic-Bench. Each example presents selected frames from Base and TrajCFG together with the token-wise guidance map produced by TrajCFG; red boxes mark representative Base failures.

three, and a person in place of one of the two requested rabbits. TrajCFG better preserves the requested object identities, counts, spatial relations, colors, and attribute bindings while maintaining comparable visual quality.

Figure 9 shows that TrajCFG better preserves the light-colored shoe sole and stacked-piece colors, while reducing apparent texture leakage from the toy kitchen's green front panel.

F Complete Quantitative Results

This section reports the complete per-dimension results underlying the summarized tables in the main paper. Tables 8 and 10 present text-to-video (T2V) results on VBench and Dynamic-Bench, respectively, while Table 9 reports the complete VBench ablation results. Tables 11 and 12 present image-to-video (I2V) results on VBench-I2V and PAI-Bench-G, respectively. Tables 13 and 14 present text-to-audio-video (T2AV) results on VABench and T2AV-Compass, respectively. Unless otherwise specified, higher values indicate better performance. For the T2V and I2V tables, the best and second-best results within each backbone are **boldfaced** and underlined, respectively; TrajCFG rows and the full TrajCFG ablation rows are shaded in gray. For the T2AV tables, only the best results are shown in **bold**.

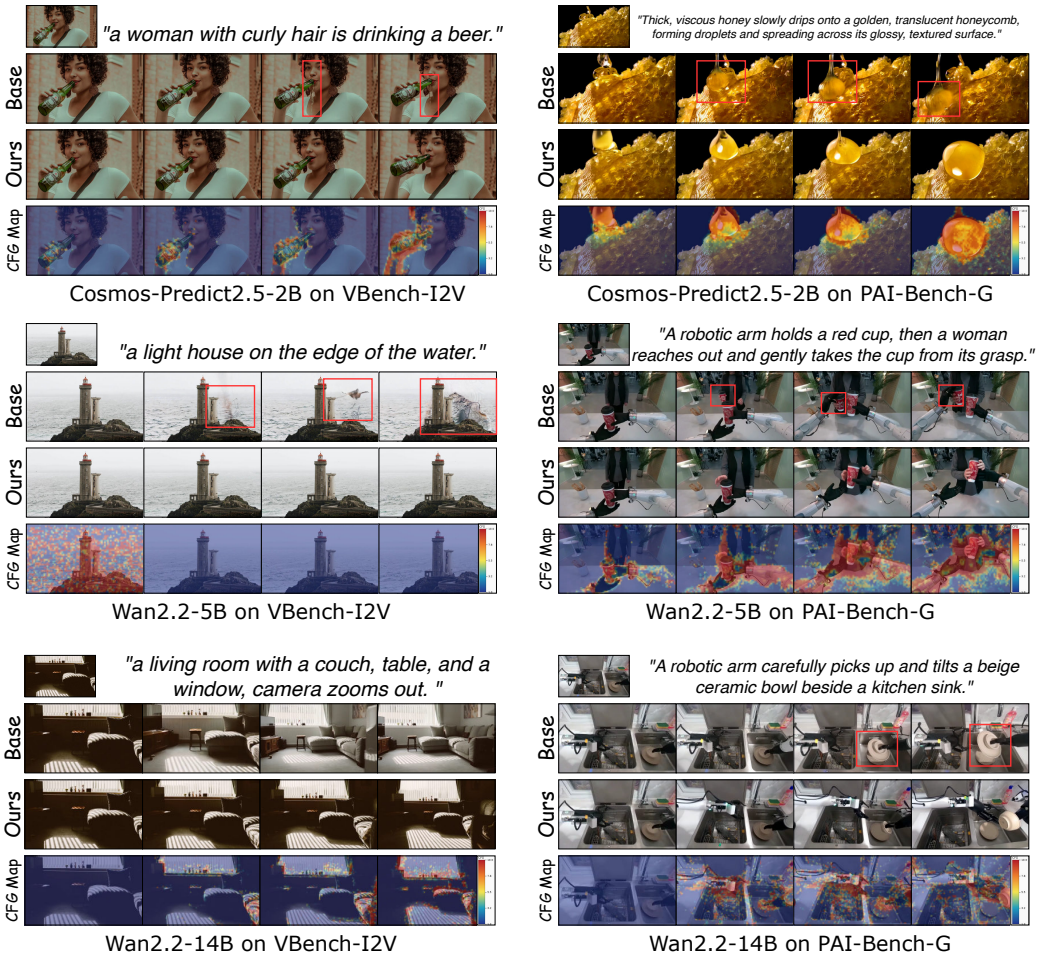

Figure 6: **Additional qualitative comparisons for image-to-video generation.** The three rows show Cosmos-Predict2.5-2B, Wan2.2-5B, and Wan2.2-14B, while the left and right columns use VBench-I2V and PAI-Bench-G, respectively. Each example presents the input image, selected frames from Base and TrajCFG, and the token-wise guidance map produced by TrajCFG; red boxes mark representative Base failures.

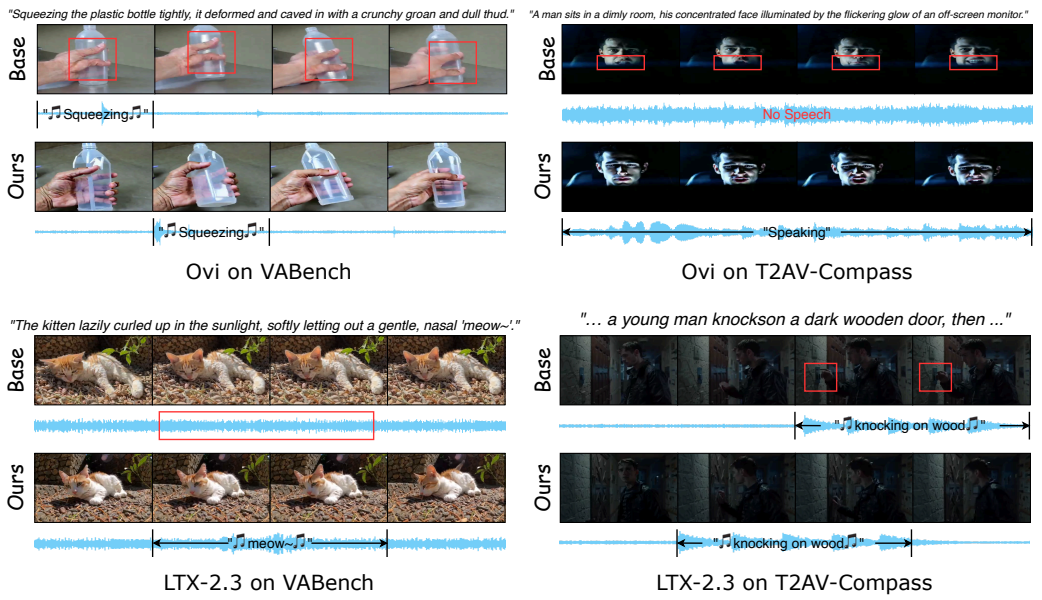

Figure 7: **Qualitative comparisons for text-to-audio-video generation.** The four examples show Ovi on VABench (upper left), Ovi on T2AV-Compass (upper right), LTX-2.3 on VABench (lower left), and LTX-2.3 on T2AV-Compass (lower right). Each example pairs selected video frames with the corresponding audio waveform. Arrows mark sound-event intervals, while red boxes highlight spatial or temporal discrepancies in Base.

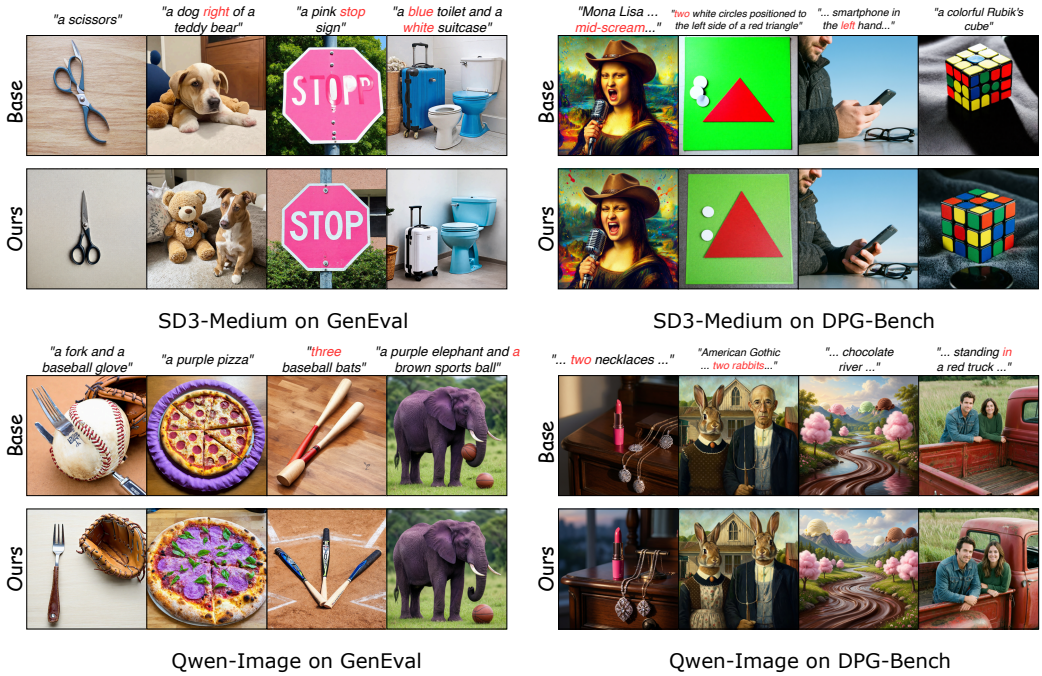

Figure 8: **Qualitative comparisons for text-to-image generation.** The four groups show SD3-Medium on GenEval (upper left), SD3-Medium on DPG-Bench (upper right), Qwen-Image on GenEval (lower left), and Qwen-Image on DPG-Bench (lower right). Within each group, Base and TrajCFG are shown in the top and bottom rows, respectively. Red text highlights selected prompt requirements.

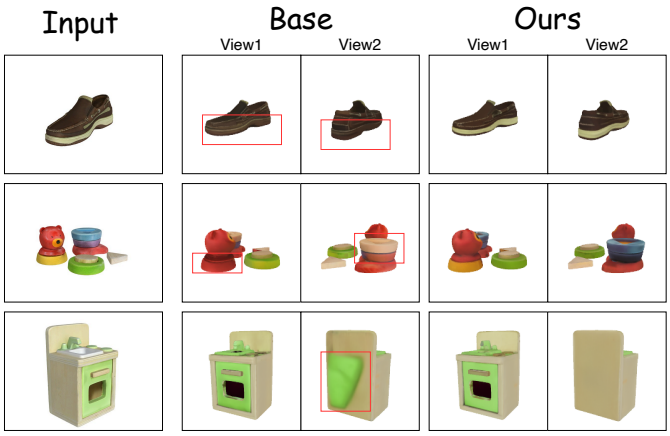

Figure 9: **Qualitative comparisons for image-to-3D generation.** Input images and two rendered views from Base and TrajCFG. Red boxes highlight appearance discrepancies in Base.

Table 8: Complete per-dimension results for text-to-video (T2V) generation on VBench.

Backbone	Method	Aesthetic Quality	Appearance Style	Background Consistency	Color	Dynamic Degree	Human Action	Imaging Quality	Motion Smoothness	Multiple Objects	Object Class
Wan2.1-1.3B	Base	64.11	22.04	96.70	84.71	65.28	91.00	62.92	98.00	71.19	87.42
	CFG-Zero	65.26	22.16	95.61	**87.67**	68.90	93.00	64.91	97.41	69.74	86.31
	CFG-Ctrl	64.49	22.12	96.77	86.69	68.06	93.00	64.80	98.09	68.90	84.97
	PAG	64.76	22.24	96.58	85.62	66.67	93.00	64.71	97.89	70.50	85.76
	SRVS	64.87	22.37	96.64	82.77	66.67	92.00	**64.94**	**98.19**	71.49	88.92
	TrajCFG	**65.48**	**22.45**	**96.90**	86.82	**70.83**	**93.75**	64.71	98.12	**75.68**	**89.14**
Wan2.2-5B	Base	65.17	22.34	96.23	90.05	55.56	**95.00**	63.89	98.44	65.78	79.17
	CFG-Zero	66.17	22.40	96.98	87.20	62.50	93.00	66.60	98.38	67.61	86.71
	CFG-Ctrl	66.20	22.38	97.35	88.38	62.50	94.00	65.99	98.33	71.80	87.18
	PAG	65.49	22.22	97.09	86.15	56.94	**95.00**	66.74	98.38	71.42	**88.61**
	SRVS	**66.22**	21.84	97.41	88.86	56.94	94.00	66.94	98.44	72.26	88.21
	TrajCFG	65.49	**22.43**	**97.53**	**91.36**	**68.06**	**95.00**	**67.26**	**98.45**	**73.86**	85.68
Wan2.2-14B	Base	66.37	21.56	95.26	77.83	72.22	90.00	66.75	97.31	76.30	87.66
	CFG-Zero	67.12	21.78	95.09	82.04	69.44	92.00	67.36	97.69	74.54	86.08
	CFG-Ctrl	66.58	22.61	95.37	79.83	65.28	94.00	66.91	97.67	**78.58**	88.29
	PAG	66.97	20.92	95.39	81.40	69.44	94.00	66.37	97.28	68.29	88.45
	SRVS	67.21	22.40	95.42	80.34	77.78	**95.24**	**67.70**	97.71	77.52	89.64
	TrajCFG	**67.26**	**22.66**	**95.43**	**83.73**	**79.17**	95.00	67.57	**97.73**	78.20	**90.03**

Backbone	Method	Overall Consistency	Scene	Spatial Relationship	Subject Consistency	Temporal Flickering	Temporal Style	Quality Score	Semantic Score	Total
Wan2.1-1.3B	Base	25.30	38.52	65.42	93.08	99.47	22.48	83.03	72.77	80.98
	CFG-Zero	25.35	39.10	71.52	91.01	99.03	22.77	82.69	73.95	80.94
	CFG-Ctrl	25.30	37.79	66.02	93.10	99.46	22.43	83.65	72.68	81.45
	PAG	25.49	38.84	72.58	92.39	99.33	22.83	83.24	73.92	81.38
	SRVS	25.48	39.17	64.97	**93.28**	**99.49**	22.51	83.69	73.10	81.57
	TrajCFG	**25.61**	**39.64**	**72.79**	93.13	**99.49**	**22.93**	**84.06**	**75.37**	**82.32**
Wan2.2-5B	Base	24.82	43.60	77.12	93.86	99.50	23.35	82.88	74.52	81.21
	CFG-Zero	25.65	40.26	**77.24**	92.38	99.38	23.29	83.79	74.84	82.00
	CFG-Ctrl	25.75	48.69	76.98	94.12	99.51	23.34	84.12	76.75	82.64
	PAG	24.91	48.18	75.19	92.53	99.48	23.33	83.37	76.14	81.93
	SRVS	25.78	47.82	77.06	**94.46**	99.40	23.25	83.93	76.63	82.47
	TrajCFG	**25.84**	**49.49**	77.19	**94.46**	**99.57**	**23.41**	**84.82**	**77.45**	**83.35**
Wan2.2-14B	Base	24.15	37.94	75.69	89.42	98.88	22.59	82.93	73.05	80.96
	CFG-Zero	25.88	45.64	77.73	91.01	98.74	22.65	83.32	75.27	81.71
	CFG-Ctrl	26.10	44.58	75.69	**91.50**	99.01	**23.15**	83.09	76.11	81.70
	PAG	25.67	42.95	72.78	91.18	98.87	21.46	83.07	73.31	81.12
	SRVS	25.98	46.22	77.50	91.04	99.05	22.97	84.24	76.59	82.71
	TrajCFG	**26.21**	**46.95**	**78.44**	91.32	**99.09**	22.82	**84.41**	**77.39**	**83.01**

Table 9: Complete per-dimension ablation results for text-to-video (T2V) generation on VBench.

Backbone	Variant	Aesthetic Quality	Appearance Style	Background Consistency	Color	Dynamic Degree	Human Action	Imaging Quality	Motion Smoothness	Multiple Objects	Object Class
Wan2.1-1.3B	**TrajCFG (Full)**	**65.48**	**22.45**	96.90	86.82	**70.83**	93.75	64.71	98.12	**75.68**	**89.14**
	Time-only	64.49	22.08	97.14	85.06	69.44	**95.00**	**65.02**	98.16	72.26	86.55
	Sample-wise	64.17	22.13	97.02	86.77	69.44	**95.00**	64.90	98.15	71.04	87.10
	Shuffled allocation	64.12	22.17	96.99	84.68	69.44	94.00	64.63	98.15	72.26	87.74
	Reversed allocation	64.01	22.28	96.98	86.67	69.44	94.00	64.32	98.16	69.59	85.21
	Adjacent-only $(1, 0)$	64.79	22.31	**97.18**	87.17	68.06	94.00	64.91	98.10	75.08	87.34
	Stride-2-only $(0, 1)$	64.71	22.22	97.05	**87.57**	69.44	94.00	64.40	**98.18**	73.09	86.16
Wan2.2-5B	**TrajCFG (Full)**	**65.49**	**22.43**	**97.53**	91.36	**68.06**	**95.00**	67.26	98.45	**73.86**	85.68
	Time-only	65.35	22.37	97.43	91.32	66.67	94.00	67.19	98.40	72.26	85.71
	Sample-wise	65.47	22.22	97.52	91.30	65.28	93.00	67.13	98.28	69.36	85.69
	Shuffled allocation	65.18	22.23	97.49	90.22	66.67	**95.00**	66.88	98.18	71.49	85.48
	Reversed allocation	65.23	22.22	97.47	86.77	66.67	92.00	66.59	98.39	71.11	85.18
	Adjacent-only $(1, 0)$	65.44	22.31	97.41	**93.26**	**68.06**	94.00	**67.56**	**98.50**	71.95	85.33
	Stride-2-only $(0, 1)$	65.15	22.21	97.51	87.81	66.67	93.00	66.89	98.46	72.03	**85.92**

Backbone	Variant	Overall Consistency	Scene	Spatial Relationship	Subject Consistency	Temporal Flickering	Temporal Style	Quality Score	Semantic Score	Total
Wan2.1-1.3B	**TrajCFG (Full)**	**25.61**	39.64	**72.79**	93.13	99.49	**22.93**	**84.06**	**75.37**	**82.32**
	Time-only	25.45	38.37	64.18	**93.19**	**99.50**	22.42	83.94	73.17	81.78
	Sample-wise	25.47	**40.55**	67.62	93.16	**99.50**	22.52	83.83	74.02	81.87
	Shuffled allocation	25.46	39.83	66.47	93.10	**99.50**	22.56	83.77	73.67	81.75
	Reversed allocation	25.44	39.31	64.03	93.10	**99.50**	22.38	83.70	72.96	81.55
	Adjacent-only $(1, 0)$	25.41	39.03	68.69	93.07	99.49	22.58	83.81	74.41	81.93
	Stride-2-only $(0, 1)$	25.44	37.79	63.98	**93.19**	99.49	22.59	83.86	73.39	81.77
Wan2.2-5B	**TrajCFG (Full)**	**25.84**	49.49	**77.19**	94.46	99.57	23.41	84.82	**77.45**	**83.35**
	Time-only	25.74	50.03	74.94	94.33	**99.59**	23.42	84.62	76.94	83.08
	Sample-wise	25.67	49.78	74.51	94.35	**99.59**	23.44	84.48	76.34	82.85
	Shuffled allocation	25.75	**50.09**	74.56	94.29	**99.59**	23.44	84.43	76.73	82.89
	Reversed allocation	25.71	49.65	74.61	94.57	**99.59**	23.27	84.55	75.82	82.81
	Adjacent-only $(1, 0)$	25.71	48.84	75.24	94.46	99.56	**23.47**	**84.85**	76.93	83.27
	Stride-2-only $(0, 1)$	25.74	47.67	71.48	**94.67**	**99.59**	23.36	84.65	75.64	82.85

Table 10: Complete per-dimension results for text-to-video (T2V) generation on Dynamic-Bench.

Backbone	Method	Subject Consistency	Background Consistency	Motion Smoothness	Aesthetic Quality	Imaging Quality	Overall Consistency
Wan2.1-1.3B	**Base**	85.54	91.42	96.40	45.17	56.21	22.36
	CFG-Zero	85.75	91.38	97.13	46.39	55.40	23.81
	CFG-Ctrl	86.80	92.47	97.28	44.40	56.27	**24.18**
	PAG	87.27	92.55	97.15	46.23	56.40	23.89
	SRVS	87.27	92.65	97.38	47.07	56.39	24.15
	TrajCFG	**87.83**	**92.70**	**97.53**	**47.31**	**56.42**	24.17
Wan2.2-5B	**Base**	86.50	92.84	96.58	45.82	58.95	24.00
	CFG-Zero	87.78	93.23	96.68	45.42	60.35	24.19
	CFG-Ctrl	87.73	93.16	96.95	46.40	56.16	24.09
	PAG	87.82	93.40	96.91	43.23	60.29	24.11
	SRVS	88.12	93.39	96.87	46.13	59.88	24.30
	TrajCFG	**88.18**	**93.44**	**97.02**	**46.41**	**60.39**	**24.33**
Wan2.2-14B	**Base**	89.29	92.59	98.14	50.62	60.73	23.47
	CFG-Zero	88.97	92.44	98.02	50.69	61.14	25.18
	CFG-Ctrl	90.13	92.71	98.38	50.88	60.12	24.79
	PAG	89.17	92.66	98.11	50.98	60.97	24.90
	SRVS	89.93	92.82	98.26	50.63	62.29	25.24
	TrajCFG	**90.74**	**93.12**	**98.43**	**51.24**	**62.80**	**25.31**

Table 11: Complete per-dimension results for image-to-video (I2V) generation on VBench-I2V.

Backbone	Method	I2V Subject	I2V Background	Subject Consistency	Background Consistency	Motion Smoothness	Aesthetic Quality	Imaging Quality	Overall Consistency
Cosmos-Predict2.5-2B	Base	98.18	98.74	96.00	96.18	99.09	62.23	71.51	23.48
	CFG-Zero	98.26	98.79	96.39	96.42	**99.14**	62.46	71.67	23.50
	CFG-Ctrl	97.83	98.43	95.37	95.53	98.97	62.21	71.55	23.36
	PAG	98.18	98.74	95.95	96.16	99.10	62.25	71.49	23.44
	SRVS	97.77	98.47	94.68	96.07	99.11	62.27	71.48	23.42
	TrajCFG	**98.29**	**98.81**	**96.51**	**96.49**	99.13	**62.59**	**71.73**	**23.53**
Wan2.2-5B	Base	98.51	98.97	97.62	97.48	99.02	63.68	71.30	22.91
	CFG-Zero	97.95	98.55	96.22	96.54	98.92	62.46	71.88	23.16
	CFG-Ctrl	98.04	98.69	96.33	96.24	99.10	61.05	71.68	23.19
	PAG	98.47	98.94	97.50	97.37	99.02	63.54	71.20	23.03
	SRVS	98.60	99.02	97.63	97.65	99.09	63.91	**71.93**	23.22
	TrajCFG	**98.78**	**99.16**	**98.33**	**98.04**	**99.32**	**64.36**	**71.93**	**23.24**
Wan2.2-14B	Base	98.28	98.74	94.74	95.04	98.77	61.79	71.94	23.29
	CFG-Zero	97.35	98.03	97.07	96.48	98.69	60.24	71.06	23.32
	CFG-Ctrl	98.24	98.72	96.59	96.02	98.79	61.28	71.98	23.21
	PAG	97.90	98.49	96.98	96.42	98.76	61.78	71.91	23.24
	SRVS	98.36	98.78	97.10	96.62	98.79	61.79	71.89	23.29
	TrajCFG	**98.42**	**98.84**	**97.55**	**96.84**	**98.90**	**62.02**	**72.02**	**23.36**

Table 12: Complete per-dimension results for image-to-video (I2V) generation on PAI-Bench-G.

Backbone	Method	I2V Subject	I2V Background	Subject Consistency	Background Consistency	Motion Smoothness	Aesthetic Quality	Imaging Quality	Overall Consistency
Cosmos-Predict2.5-2B	Base	96.27	97.25	92.43	94.17	99.09	52.46	70.18	21.18
	CFG-Zero	96.33	97.32	92.70	94.38	99.13	52.83	70.38	**21.19**
	CFG-Ctrl	95.33	96.36	91.23	93.54	98.37	51.59	70.27	20.94
	PAG	96.21	97.21	92.14	94.08	99.08	**52.87**	70.36	21.17
	SRVS	95.40	96.65	91.67	93.71	99.10	52.86	70.43	21.15
	TrajCFG	**96.39**	**97.38**	**92.77**	**94.44**	**99.14**	52.62	**70.46**	21.15
Wan2.2-5B	Base	96.47	97.80	92.72	94.57	98.72	51.24	67.88	21.26
	CFG-Zero	96.17	97.55	92.64	94.47	98.72	51.58	68.41	21.27
	CFG-Ctrl	96.53	97.92	92.80	94.21	98.97	50.81	68.80	21.06
	PAG	96.27	97.63	92.70	94.46	98.69	51.65	68.86	21.33
	SRVS	96.11	97.38	92.10	94.39	98.50	**51.82**	**68.96**	21.33
	TrajCFG	**96.95**	**98.13**	**94.01**	**95.23**	**98.99**	51.45	67.93	**21.35**
Wan2.2-14B	Base	96.12	97.26	91.81	92.32	98.50	51.11	66.60	21.36
	CFG-Zero	95.84	97.04	91.05	91.81	98.48	51.09	65.95	21.41
	CFG-Ctrl	95.80	97.04	91.89	92.02	98.52	50.51	65.84	21.30
	PAG	96.08	97.24	91.72	92.30	98.53	50.95	66.42	21.43
	SRVS	96.17	97.29	91.71	92.26	98.61	51.09	66.62	21.43
	TrajCFG	**96.37**	**97.51**	**92.77**	**92.72**	**98.65**	**51.13**	**66.67**	**21.47**

Table 13: Complete per-dimension results for text-to-audio-video (T2AV) generation on VABench. Lower is better for Desync.

Backbone	Method	Speech Q&N	Audio Aes.	T–V Align.	T–A Align.	A–V Align.	Lip–Sync	Desync↓	Alignment	Expressiveness	Visual Realism	Audio Realism	Audio QA	Visual QA
LTX-2.3	Base	2.698	3.345	0.1928	0.3118	0.2490	**2.752**	0.5911	4.470	4.299	4.615	4.032	0.7049	0.6616
	TrajCFG	**2.763**	**3.351**	**0.1963**	**0.3202**	**0.2567**	2.555	**0.5685**	**4.479**	**4.321**	**4.638**	**4.066**	**0.7178**	**0.6730**
Ovi	Base	3.181	3.282	0.1968	0.3288	0.1564	4.420	0.7659	**4.189**	4.138	4.522	4.063	0.5251	0.5948
	TrajCFG	**3.463**	**3.364**	**0.1977**	**0.3316**	**0.1654**	**4.780**	**0.7080**	4.179	**4.141**	**4.547**	**4.074**	**0.5257**	**0.5953**

Table 14: Complete per-dimension results for text-to-audio-video (T2AV) generation on T2AV-Compass. Lower is better for DS.

Backbone	Method	Video Quality		Audio Quality		Cross-modal Alignment				
		VT↑	VA↑	AA↑	SQ↑	A–V↑	T–A↑	T–V↑	DS↓	LS↑
LTX-2.3	Base	7.295	4.921	6.709	1.562	0.2469	0.1722	**0.0432**	0.7140	2.484
	TrajCFG	**7.510**	**4.924**	**6.738**	**1.578**	**0.2481**	**0.1735**	0.0423	**0.7016**	**2.820**
Ovi	Base	6.714	4.516	6.676	1.566	0.1596	0.0897	0.0297	0.9624	1.939
	TrajCFG	**7.199**	**4.555**	**6.733**	**1.665**	**0.1622**	**0.0922**	**0.0316**	**0.9356**	**2.449**