Diffusion Forcing Planner: History-Annealed Planning with Time-Dependent Guidance for Autonomous Driving

Supplementary Material

1. Supplementary Methods

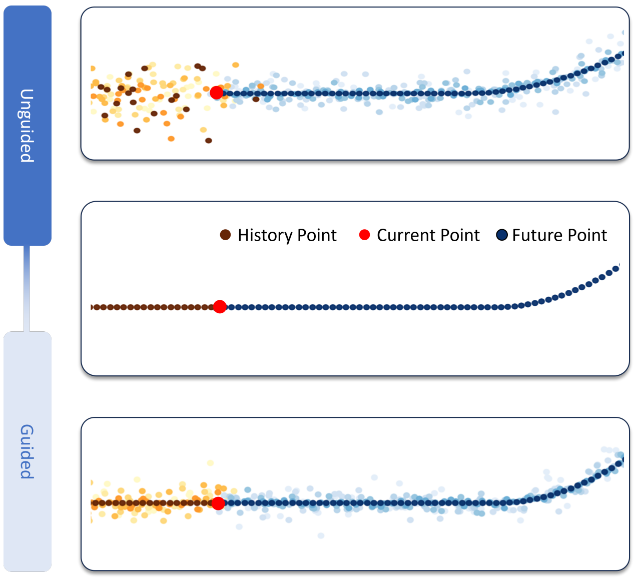

Figure 1. **Annealed history CFG.** Lighter points indicate earlier diffusion timesteps.

Algorithms 1 and 2 are presented in the Method section as pseudocode, and together they provide the complete pipelines of the two components of DFP: Algorithm 1 — Training with Diffusion Forcing, and Algorithm 2 — History-annealed CFG Inference. Figure1 visualizes the annealed history CFG procedure.

2. Additional Experiments on Flow Matching and Mix-of-Transformer

While our main approach utilizes a Diffusion-based decoder to ensure stability and history controllability, we recognize the potential of the Flow Matching paradigm for achieving straighter generation trajectories. In this exploratory study, we substituted the Diffusion with the Flow Matching [4]. Additionally, to enhance the model's capacity in capturing diverse multi-modal interactions, we upgrade our backbone to a Mixture-of-Transformer structure[1, 3, 4]. This design leverages multiple expert transformers to specialize in varied traffic contexts. As shown in Table 2 and Table 1, our new baseline outperforms the original DFP framework, and achieves state-of-the-art closed-loop performance on nuPlan non-reactive benchmarks.

Algorithm 1 Diffusion-Forcing Training with Chunkwise Noising-as-Masking

Require: scene contexts C; GT trajectory $x_0 = [x_0^{-H}, \dots, x_0^0, \dots, x_0^F]$, total length $S = H + 1 + F$; chunk length L, the number of chunks N (N_H history, 1 current, N_F future); SDE marginals (α_t, σ_t); DiT f_θ; weights $\lambda_{\text{hist}}, \lambda_{\text{futr}}$

Ensure: parameter update of θ

Split x_0 into N index sets $\{\mathcal{I}_b\}_{b=1}^N$, then each chunk $x_0^{(b)} = \{x_{0,i}\}_{i \in \mathcal{I}_b}$

Sample per-block noise levels t_b, $b = 1, \dots, N$; set current chunk $t_0 \leftarrow 0$

for $b = 1$ to N **do**

  Sample $\varepsilon^{(b)} \sim \mathcal{N}(0, I)$

  Noised chunk: $x_{t_b}^{(b)} \leftarrow \alpha_{t_b} x_0^{(b)} + \sigma_{t_b} \varepsilon^{(b)}$

end for

Pack tokens $X_t = \text{concat}\big(\{x_{t_b}^{(b)}\}_{b=1}^N\big), \quad t = \{t_b\}_{b=1}^N$

Predict per-block $\{\hat{x}_{t_b}^{(b)}\}_{b=1}^N \leftarrow f_\theta\big(X_t, t, C\big)$

Reconstruction losses:

$L_{\text{hist}} \leftarrow \frac{1}{N_H} \sum_{b=1}^{N_H} \sum_{i \in I_b} \big\|\hat{x}_{0,i}^{(b)} - x_{0,i}^{(b)}\big\|_2^2$

$L_{\text{futr}} \leftarrow \frac{1}{N_F} \sum_{b=N_H+2}^{N} \sum_{i \in I_b} \big\|\hat{x}_{0,i}^{(b)} - x_{0,i}^{(b)}\big\|_2^2$

Total loss: $L \leftarrow \lambda_{\text{hist}} L_{\text{hist}} + \lambda_{\text{futr}} L_{\text{futr}}$

Update $\theta \leftarrow \theta - \eta \nabla_\theta L$

3. Ablation Study on Temporal Stability: Impact of History Guidance

To isolate the contribution of the proposed History-Annealed-Guidance mechanism on temporal stability, we conduct and ablation study comparing the full DFP model (with history guidance) against a variant which predicts history and future trajectories jointly (without history guidance). Figure 2 visualizes the jerk and acceleration over a 15-second segment. The results highlight a clear difference in the quality of the generated trajectories: the model with history guidance generates significantly smoother yaw rate and longitudinal acceleration profiles, indicating reduced jerk and improved comfort.

Figure 3 report the distributions of frame-level quantities. Compared to DP, DFP reduces the mean ego jerk by **18.9%** and the standard deviation by **17.0%**, indicating not only smoother trajectories on average, but also significantly reduced variability across consecutive planning frames. Notably, on ego yaw acceleration, DFP maintains performance on par with the baseline, aligning with our comfort expec-

Table 1. **Case study on scenario-specific metrics.** Evaluation is conducted on the nuPlan `Val14` benchmark in the non-reactive setting. *: We reproduce the baselines[5] on the nuPlan[2], and may differ slightly from the original report, likely due to differences in the randomly sampled 1M training subset. †: Motivated by the design of Flow Planner [4], we incorporate Flow Matching and a Mixture-of-Transformer architecture into our Diffusion Forcing Planner; **Score** denotes the overall score for the scenario; **Collision/TTC/Drivable/Comfort/Progress** denote the respective per-scenario metrics. All per-scenario metrics are Boolean indicators, reported as the fraction of cases satisfying the condition (higher is better).

Scenario type	Planner	Score	Collision	TTC	Drivable	Comfort	Progress
All	PlanTF	84.27	94.01	89.00	95.89	93.02	98.75
	Diffusion Planner*	87.80	95.53	92.13	97.32	91.86	**100.00**
	DFP (ours)	90.33	96.60	91.86	92.49	96.69	99.91
	DFP (ours) †	**92.68**	**98.03**	**93.20**	**98.12**	**98.48**	99.82
Low magnitude speed	PlanTF	85.00	96.00	92.00	95.00	89.00	95.00
	Diffusion Planner*	86.51	97.00	97.00	95.00	94.00	96.00
	DFP (ours)	**91.08**	98.00	97.00	97.00	96.00	**99.00**
	DFP (ours) †	90.74	**98.00**	**97.00**	**97.00**	**100.00**	98.00
High magnitude speed	PlanTF	89.39	94.95	94.95	95.96	94.95	98.99
	Diffusion Planner*	84.50	95.96	95.96	96.97	60.61	100.00
	DFP (ours)	94.95	98.99	97.98	97.98	96.97	100.00
	DFP (ours) †	**97.40**	**100.0**	**98.99**	**98.99**	**98.99**	**100.00**
Starting left turn	PlanTF	80.42	93.00	85.00	98.00	81.00	98.00
	Diffusion Planner*	82.38	93.00	90.00	97.00	90.00	99.00
	DFP (ours)	86.74	93.00	88.00	98.00	94.00	100.00
	DFP (ours) †	**93.15**	**97.00**	**93.00**	**100.00**	**99.00**	**100.00**
Starting right turn	PlanTF	70.13	91.84	81.63	90.82	91.84	94.90
	Diffusion Planner*	78.16	89.80	81.63	94.90	94.90	100.00
	DFP (ours)	79.38	94.90	82.65	88.78	95.92	100.00
	DFP (ours) †	**86.61**	**96.94**	**83.67**	**95.92**	**97.96**	**100.00**

Table 2. **Main results on the nuPlan benchmarks.** Overall score (%) in non-reactive (NR) and reactive (R) closed-loop evaluations on *Val14*, *Test14*, and *Test14-hard*. No post-processing; raw model outputs are evaluated. *: We reproduce the baselines[5] on the nuPlan[2], and may differ slightly from the original report, likely due to differences in the randomly sampled 1M training subset. †: Motivated by the design of Flow Planner [4], we incorporate Flow Matching and a Mixture-of-Transformer architecture into our Diffusion Forcing Planner.

Type	Planner	Val14		Test14		Test14-hard	
		NR	R	NR	R	NR	R
Expert	Log-replay	93.53	80.32	85.96	68.80	94.03	75.86
Learning-based	PDM-Open*	53.53	54.24	52.81	57.23	33.51	35.83
	UrbanDriver	68.57	64.11	51.83	67.15	50.40	49.95
	GameFormer	13.32	8.69	11.36	9.31	7.08	6.69
	PlanTF	84.27	76.95	85.62	79.58	69.70	61.61
	PLUTO	88.89	78.11	89.90	78.62	70.03	59.74
	CoPlanner	89.48	79.00	90.31	78.81	76.82	64.47
	Diffusion Planner	89.87	82.80	89.19	82.93	75.99	69.22
	Flow Planner	90.43	**83.31**	89.88	82.93	76.47	**70.42**
	Diffusion Planner *	87.87	77.48	90.01	79.61	74.26	61.25
	DFP (ours)	90.33	79.97	**90.69**	81.96	76.91	63.56
	DFP (ours) †	**92.68**	81.30	90.62	**83.59**	**79.43**	67.94

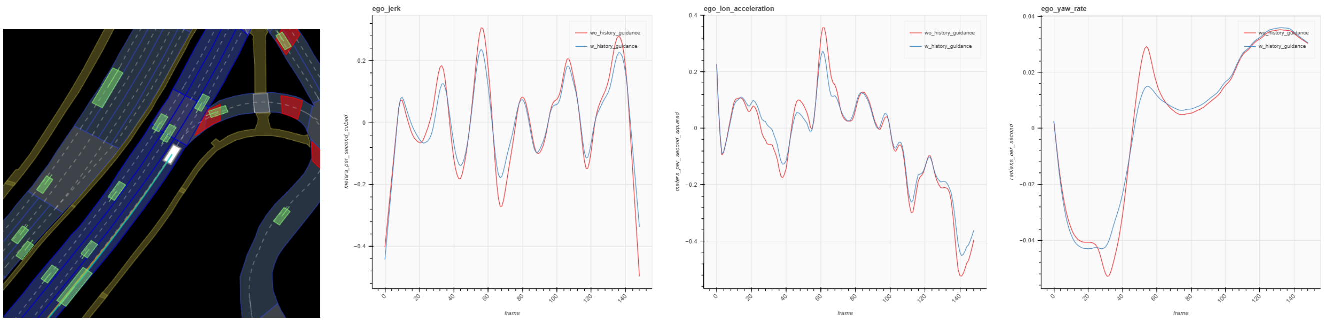

(a) nuPlan scene (b) Ego jerk (c) Ego longitudinal acceleration (d) Ego yaw rate

Figure 2. nuPlan scene visualization and ego kinematic profiles. From left to right: scene view, jerk, longitudinal acceleration, and yaw rate.

Algorithm 2 Annealed-History Classifier-Free Guidance (CFG) Inference

Require: scene context C; clean history chunks X_{history}; current chunk $x_0^{(0)}$; full length $S{=}H{+}1{+}F$; chunk length L; the number of chunks $N{=}N_H{+}1{+}N_F$; diffusion time t_s; guidance weight w; history anneal exponent β; SDE (α_t, σ_t); DiT f_θ

Ensure: predicted future trajectory $\widehat{x}_{\text{future}}{=}[x_0^1, x_0^2, \dots, x_0^F]$

Initialize future chunks with noise: $x_{t_s}^{(b)} \sim \mathcal{N}(0, I)$, for $b{=}1, 2, \dots, N_F$

Form $X_{t_s} \leftarrow \left[x_0^{(0)},\, x_{t_s}^{(1)},\, \dots,\, x_{t_s}^{(N_F)}\right]$

for $t = t_s,\, \dots,\, 0$ **do**

Unguided branch

Sample history noise chunks $\varepsilon \sim \mathcal{N}(0, I)$

$X \leftarrow \left[\varepsilon,\, X_{t_s}\right]$

$\mathbf{t} \leftarrow \Big[\, \underbrace{1, \dots, 1}_{N_H},\, 0,\, \underbrace{t_s, \dots, t_s}_{N_F} \Big]$

$\widehat{X}_{0,unguided} \leftarrow f_\theta\big(X,\, \mathbf{t},\, C\big)$

Guided branch

Sample $\varepsilon \sim \mathcal{N}(0, I)$; $t_{\text{hist}} \leftarrow (t_s)^\beta$

$X_{\text{guidance}} \leftarrow \alpha_{t_{\text{hist}}} X_{\text{history}} + \sigma_{t_{\text{hist}}}\, \varepsilon$

$X \leftarrow \left[X_{\text{guidance}},\, X_{t_s}\right]$

$\mathbf{t} \leftarrow \Big[\, \underbrace{t_{\text{hist}}, \dots, t_{\text{hist}}}_{N_H},\, 0,\, \underbrace{t_s, \dots, t_s}_{N_F} \Big]$

$\widehat{X}_{0,guided} \leftarrow f_\theta\big(X,\, \mathbf{t},\, C\big)$

CFG fusion

$\widehat{X}_0 \leftarrow \widehat{X}_{0,unguided} + w\big(\widehat{X}_{0,guided} - \widehat{X}_{0,unguided}\big)$

Hard constraint & step update

Reset current chunk to $x_0^{(0)}$

end for

Concatenate all future chunks to compose $\widehat{x}_{\text{future}}$

tations. The full CFG inference currently runs at 7.7 FPS (129.79ms), while unguided branch only runs at 14.8 FPS (67.71 ms).

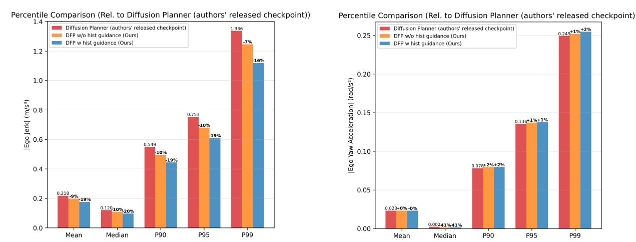

Figure 3. Frame-level comparison.

References

[1] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024. 1

[2] Holger Caesar, Juraj Kabzan, Kok Seang Tan, Whye Kit Fong, Eric Wolff, Alex Lang, Luke Fletcher, Oscar Beijbom, and Sammy Omari. Nuplan: A closed-loop ml-based planning benchmark for autonomous vehicles. In *CVPR ADP3 workshop*, 2021. 2

[3] Weixin Liang, LILI YU, Liang Luo, Srini Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen tau Yih, Luke Zettlemoyer, and Xi Victoria Lin. Mixture-of-transformers: A sparse and scalable architecture for multi-modal foundation models. *Transactions on Machine Learning Research*, 2025. 1

[4] Tianyi Tan, Yinan Zheng, Ruiming Liang, Zexu Wang, Kexin Zheng, Jinliang Zheng, Jianxiong Li, Xianyuan Zhan, and Jingjing Liu. Flow matching-based autonomous driving planning with advanced interactive behavior modeling. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. 1, 2

[5] Yinan Zheng, Ruiming Liang, Kexin ZHENG, Jinliang Zheng, Liyuan Mao, Jianxiong Li, Weihao Gu, Rui Ai, Shengbo Eben Li, Xianyuan Zhan, and Jingjing Liu.

Diffusion-based planning for autonomous driving with flexible guidance. In *The Thirteenth International Conference on Learning Representations*, 2025. 2