

Flow Matching-based Trajectory Generation for Intelligent and Reliable Motion Control in End-to-End Autonomous Driving

Yeryeong Cho[‡], Soohyun Park[†], and Joongheon Kim[‡]

[‡]Korea University [†]Sookmyung Women's University

E-mails: {joyena0909, joongheon}@korea.ac.kr, soohyun.park@sookmyung.ac.kr,

Abstract—End-to-end autonomous driving requires robust trajectory planning to ensure both safety and adaptability under diverse conditions. Traditional reinforcement learning (RL) and imitation learning (IL) approaches encounter challenges such as unstable training, poor generalization, and high sample complexity. For these reasons, flow matching (FM) in generative modeling is proposed as a promising alternative for action planning. Therefore, this paper proposes the feasibility of FM-based planning using a simplified environment as a controlled abstraction of road networks. The evaluation demonstrates that FM achieves rapid and stable convergence and generates near-optimal trajectories. Furthermore, the proposed algorithm produces smoother paths than the RL method. These results highlight the potential of FM to ensure both optimality and adaptability in data-driven generative models. Therefore, this paper establishes FM as a foundation for advancing end-to-end autonomous driving systems.

Index Terms—Flow Matching (FM), Generative Model, End-to-End Autonomous Driving

I. INTRODUCTION

End-to-end autonomous driving requires reliable action planning that generates safe and efficient trajectories. Traditional reinforcement learning (RL) and imitation learning (IL) often suffer from instability and limited generalization to unseen environments [1], [2]. Flow matching (FM), a recent generative modeling technique, provides a promising alternative by effectively capturing continuous trajectory distributions and enabling smooth action generation [3]. Therefore, this paper presents an initial validation of FM-based action planning using a simplified environment. The environment is designed as an abstracted representation of road networks, where an agent must navigate from a start position to a goal while avoiding obstacles. This environment offers a controlled testbed for evaluating the feasibility of FM in action planning before extending it to full-scale autonomous driving systems. This paper presents an initial exploration of FM-based action planning for autonomous driving. The key contributions are summarized as follows:

- Introduce FM as a generative framework for action planning in autonomous driving.
- Comparative analysis with established planning methods under controlled experimental conditions.

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT)(RS-2025-00561377). (Corresponding author: Joongheon Kim)




Fig. 1: Proposed FM architecture.

- Suggest future direction to enhance end-to-end autonomous driving through improved generalization and trajectory modeling.

Consequently, this paper establishes FM as not only a theoretical framework but also a practical foundation for advancing scalable and reliable end-to-end autonomous driving systems.

II. RELATED WORKS

Research on autonomous driving has widely explored learning-based approaches for end-to-end control. IL algorithms directly map sensory inputs to driving actions [4]. Although they provide efficiency in training, they typically fail to handle distributional shifts and exhibit poor generalization in diverse or rare driving conditions [5]. RL methods, in contrast, enable adaptability through trial-and-error interaction with the environment [6]. However, their reliance on extensive exploration results in high sample complexity, and training stability remains a persistent challenge [7]. These limitations restrict their deployment in real-world safety-critical systems [8].

Generative modeling has recently gained traction as a promising alternative for sequential decision-making [9]. Diffusion-based methods demonstrate strong capacity in modeling complex data distributions and have been applied to trajectory prediction and motion generation tasks [10]. FM, as a specific class of generative modeling, offers distinct advantages by learning vector fields that transform noise distributions into structured trajectories [11], as shown in Fig. 1. This architecture provides stable training dynamics and efficient trajectory sam-

pling, reducing computational overhead compared to traditional diffusion processes [12].

The integration of such generative techniques into planning enables the generation of smooth, robust, and dynamically feasible action sequences [13]. This stands in contrast to deterministic classical methods, which guarantee optimality under simplified assumptions but lack flexibility in uncertain or dynamic conditions [14]. FM thus represents a new direction that unifies the strengths of data-driven learning with principled generative modeling, establishing a foundation for more reliable end-to-end autonomous driving [15].

III. ALGORITHM

The proposed algorithm formulates end-to-end autonomous driving as an action planning problem that can be effectively addressed through FM. Unlike classical optimization-based planners that rely on deterministic search, FM provides a generative framework capable of learning continuous distributions of feasible trajectories. In this architecture, a velocity field is learned to progressively transform random noise samples into structured trajectories. As a result, the algorithm ensures that the generated paths are smooth and dynamically consistent. It also remains robust against environmental variations that typically degrade the performance of conventional methods. The training process follows the rectified flow formulation. For each demonstration collected from expert references, a target trajectory x_1 is defined as the ground truth sequence of waypoints. A Gaussian noise vector $x_0 \sim \mathcal{N}(0, I)$ is sampled to serve as the initialization. These two elements are combined through linear interpolation to create an intermediate state:

$$x_t = (1 - t)x_0 + tx_1, \quad t \in [0, 1]. \quad (1)$$

The ideal transport velocity guiding this transformation is

$$u_t = x_1 - x_0. \quad (2)$$

The model then is trained to predict this $v_\theta(x_t, t)$, minimizing the mean squared error, shown as

$$\mathcal{L}(\theta) = \|v_\theta(x_t, t) - u_t\|^2. \quad (3)$$

Through this process, the model acquires a continuous mapping from noise to structured trajectories.

In inference time, a trajectory is then generated by sampling $x_0 \sim \mathcal{N}(0, I)$ and repeatedly applying the learned velocity field. The integration is performed via a discretized update rule, defined as

$$x_{t+\Delta t} = x_t + v_\theta(x_t, t) \Delta t. \quad (4)$$

This recursive update gradually transports the noisy initialization toward a valid trajectory. As t approaches 1, the process converges to a coherent and dynamically consistent path. Because each x_0 can yield a distinct yet feasible path, this architecture guarantees not only diversity but also reliability.




Fig. 2: Loss of Proposed FM Algorithm for 1000 Epochs.




Fig. 3: Comparison of the Generated Path.

Furthermore, the refinement process enforces convergence toward dynamically consistent and smooth trajectories. The generated plans can be directly mapped into driving environments, forming the basis for safe, adaptable, and scalable end-to-end autonomous driving.

IV. PERFORMANCE EVALUATION

A. Evaluation Setup

The evaluation is conducted in a two-dimensional grid-based environment, which provides a simplified abstraction of navigation scenarios. Each experiment requires an agent to move from a start point to a designated goal while avoiding static obstacles. The environment offers a controlled setup in which trajectories are modeled as discrete spatial steps and resampled into continuous paths. It enables systematic assessment of the ability to generate feasible and smooth plans. This abstraction validates the fundamental capability to model trajectory distributions without the added complexity of real driving environments. Furthermore, its simplicity reduces computational overhead to ensure feasibility and stability. This is essential for isolating the action planning component and demonstrating its effectiveness before real-world deployment.

B. Evaluation Result

Fig. 2 shows the training loss of the proposed FM algorithm. The loss exhibits a steep decrease at the beginning, which indicates that the model rapidly captures the fundamental structure of the transport mapping from noise to valid trajectories. The reduction rate then becomes gradual, and the loss converges toward a stable value. This convergence pattern confirms the numerical robustness of the FM framework and its stability over long iterations. The observed stability implies that the learned velocity field is well-aligned with the target dynamics, thereby ensuring reliable generation of trajectories in subsequent inference stages. Furthermore, Fig. 3 shows a qualitative comparison of paths generated by different planning approaches. The trajectory produced by FM aligns closely with the reference solution obtained by the A* algorithm. This proximity demonstrates that FM can capture near-optimal solutions while simultaneously producing smoother trajectories. In contrast, the RL method tends to generate irregular and less reliable paths, which occasionally deviate from the intended direction. Such deviations demonstrate the limitations of RL in constrained navigation environments. In summary, the evaluation results demonstrate that FM-based action planning offers several advantages: (i) rapid and stable convergence, (ii) reliable trajectory generation that remains close to optimal references, and (iii) smooth path structures that outperform benchmarks. Therefore, this paper demonstrates that FM provides a promising foundation for end-to-end autonomous driving, as it combines the optimality of classical planning with the robustness and adaptability of generative modeling.

V. FUTURE WORKS AND CONCLUSION

The evaluation results demonstrate the effectiveness of FM in stable, near-optimal, and smooth trajectories. However, it is necessary to extend the framework to address the complexities of real-world autonomous driving. Therefore, future research should focus on validation and extension to a multi-agent system. Validation in realistic simulators is crucial for scalability under continuous states and dynamic agents. Such validation enables assessment of scalability and robustness under complex interactions and variable driving scenarios. Beyond single-agent navigation, the integration of FM into multi-agent contexts represents a critical next step. Autonomous driving naturally involves continuous cooperation among vehicles, pedestrians, and infrastructure systems. Embedding interaction-awareness into the flow fields will allow cooperative planning, dynamic adaptation at intersections, and conflict-free decision-making in densely populated environments. In conclusion, FM represents a promising foundation for end-to-end autonomous driving. The demonstrated balance between trajectory feasibility and smoothness highlights the potential of FM for both controlled testbeds and practical driving systems.

REFERENCES

- [1] W. Liu, W. Hu, W. Jing, L. Lei, L. Gao, and Y. Liu, "Learning to model diverse driving behaviors in highly interactive autonomous driving scenarios with multiagent reinforcement learning," *IEEE Systems Journal*, vol. 19, no. 1, pp. 317–326, February 2025.
- [2] W. J. Yun, M. Shin, S. Jung, S. Kwon, and J. Kim, "Parallelized and randomized adversarial imitation learning for safety-critical self-driving vehicles," *Journal of Communications and Networks*, vol. 24, no. 6, pp. 710–721, April 2022.
- [3] Y. Lipman, R. T. Q. Chen, H. Ben-Hamu, M. Nickel, and M. Le, "Flow matching for generative modeling," in *Proc. International Conference on Learning Representations (ICLR)*, Kigali, Rwanda, May 2023.
- [4] L. Cultrera, F. Becattini, L. Seidenari, P. Pala, and A. D. Bimbo, "Addressing limitations of state-aware imitation learning for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 1, pp. 2946–2955, November 2024.
- [5] S. Teng, L. Chen, Y. Ai, Y. Zhou, Z. Xuanyuan, and X. Hu, "Hierarchical interpretable imitation learning for end-to-end autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 1, pp. 673–683, 2023.
- [6] S. Govinda, B. Brik, and S. Harous, "A survey on deep reinforcement learning applications in autonomous systems: Applications, open challenges, and future directions," *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 7, pp. 11 088–11 113, May 2025.
- [7] W. Huang, H. Liu, Z. Huang, and C. Lv, "Safety-aware human-in-the-loop reinforcement learning with shared control for autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 11, pp. 16 181–16 192, July 2024.
- [8] J. Chen, S. E. Li, and M. Tomizuka, "Interpretable end-to-end urban autonomous driving with latent deep reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 6, pp. 5068–5078, February 2022.
- [9] L. H. Meftah, A. Cherif, and R. Braham, "Improving autonomous vehicles maneuverability and collision avoidance in adverse weather conditions using generative adversarial networks," *IEEE Access*, vol. 12, pp. 89 679–89 690, June 2024.
- [10] E. Sun, "CondDG: Enhancing autonomous navigation in adverse weather with conditional diffusion synthetic data," in *Proc. International Conference on Control and Robotics (ICCR)*, Yokohama, Japan, December 2024, pp. 63–68.
- [11] W. Seo, W. Park, and M. Kim, "Lightweight optical flow estimation using 1D matching," *IEEE Access*, vol. 12, pp. 45 560–45 571, March 2024.
- [12] Q. Rouxel, A. Ferrari, S. Ivaldi, and J.-B. Mouret, "Flow matching imitation learning for multi-support manipulation," in *Proc. International Conference on Humanoid Robots (Humanoids)*, Nancy, France, November 2024, pp. 528–535.
- [13] Y. Fu, Q. Yan, L. Wang, K. Li, and R. Liao, "MoFlow: One-step flow matching for human trajectory forecasting via implicit maximum likelihood estimation based distillation," in *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, June 2025, pp. 17 282–17 293.
- [14] Y. Jin, Z. Sun, N. Li, K. Xu, H. Jiang, N. Zhuang, Q. Huang, Y. Song, Y. Mu, and Z. Lin, "Pyramidal flow matching for efficient video generative modeling," in *Proc. International Conference on Learning Representations (ICLR)*, Singapore, April 2025.
- [15] Z. Xing, X. Zhang, Y. Hu, B. Jiang, T. He, Q. Zhang, X. Long, and W. Yin, "GoalFlow: Goal-driven flow matching for multimodal trajectories generation in end-to-end autonomous driving," in *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, June 2025, pp. 1602–1611.