
Constraint-Aware Diffusion Guidance for Imitation Learning

Hao Ma*

haomah@ethz.ch

Sabrina Bodmer*

sabodmer@ethz.ch

Andrea Carron*

carrona@ethz.ch

Melanie Zeilinger*[†]

mzeilinger@ethz.ch

Michael Muehlebach[‡]

michaelm@tuebingen.mpg.de

Abstract

We propose Constraint-Aware Diffusion Guidance (CoDiG), a constraint-aware imitation learning framework based on conditional diffusion models. Unlike conventional imitation learning methods, which often fail to generalize to unseen or constrained environments, CoDiG enforces safety and physical feasibility during inference via barrier function guidance. Our method learns from a limited number of expert demonstrations without reward supervision or environment interaction, and is capable of generating safe and feasible trajectories in real time. A warm-start strategy further accelerates sampling by reusing previous outputs. We evaluate CoDiG on a miniature autonomous racing platform in a challenging obstacle avoidance task, demonstrating robust generalization, near time-optimal performance, and 100% success rate in dynamic scenarios. Our results highlight the potential of constraint-aware diffusion models as a data-efficient and deployable solution for safe imitation learning in robotics. Videos are available at: <https://www.youtube.com/watch?v=KNYsTddxOU>.

1 Introduction

Learning policies from expert demonstrations - commonly referred to as imitation learning - has emerged as a compelling alternative to reinforcement learning (RL), particularly in scenarios where reward functions are difficult to define, environment interaction is costly or risky, and safety is critical [Correia and Alexandre, 2024, Argall et al., 2009]. However, conventional imitation learning methods, such as behavior cloning or adversarial imitation, often fail to generalize to unseen environments or enforce essential physical constraints [Ross et al., 2011, Ho and Ermon, 2016, Fu et al., 2018, Lu et al., 2023a]. These limitations significantly hinder their applicability to real-world, safety-critical robotic systems.

Recently, diffusion models have shown promise in offline imitation settings by learning to generate expert-like policies through denoising-based generative processes for robotics [Wolf et al., 2025, Chi et al., 2024, Urain et al., 2023]. While powerful, many of these approaches are trained purely on data without explicitly enforcing constraints, which can lead to collisions or dynamic infeasibility, particularly when encountering out-of-distribution scenarios [Kondo et al., 2024, Palo et al., 2025]. Additionally, existing methods often rely heavily on large-scale offline datasets to promote generalization, which can limit their adaptability in unseen environments [Lee et al., 2025, Liang et al., 2023].

*Department of Mechanical and Process Engineering, ETH Zürich, Switzerland.

[†]Shared last author.

[‡]Max Planck Institute for Intelligent Systems, Tübingen, Germany.

Addressing these challenges is crucial for enabling the safe and reliable deployment of diffusion models in safety-critical imitation learning in robotic applications.

To overcome these limitations, we propose Constraint-Aware Diffusion Guidance (CoDiG), a general-purpose, data-efficient imitation learning framework that integrates diffusion-based trajectory generation with real-time constraint enforcement. Our method learns from a limited set of expert demonstrations, without requiring environment interaction or reward functions. CoDiG integrates a barrier function directly into the reverse diffusion process, steering the sampling away from unsafe or dynamically infeasible regions without relying on external classifiers or auxiliary models. To further accelerate sampling and enhance its stability, CoDiG employs a warm-start strategy by initializing the diffusion process near feasible solutions. By augmenting the score updates with barrier gradients during inference, CoDiG enforces safety constraints, enabling efficient and reliable deployment in safety-critical environments. Our main contributions are summarized as follows:

- We introduce Constraint-Aware Diffusion Guidance (CoDiG), a general-purpose and data-efficient diffusion-based imitation learning framework that learns from a small number of expert demonstrations and enforces safety constraints *at inference time*, enabling generalization to unseen and time-varying environments without reward supervision or online interaction.
- We propose a warm-start strategy that significantly accelerates the inference process, achieving real-time performance suitable for high-frequency, safety-critical control, while ensuring smooth transitions between trajectories generated at successive iterations.
- We evaluate CoDiG on a real-world autonomous racing platform, demonstrating that CoDiG can safely imitate expert behavior (near time-optimal) and perform obstacle avoidance in dynamic and constrained environments with high reliability.

2 Related work

Our method provides a constraint-aware imitation learning framework that is grounded in diffusion-based trajectory generation. Unlike traditional reinforcement learning approaches that rely on online exploration and reward feedback, our method learns purely from offline demonstrations and enforces constraints directly during inference. This aligns with recent trends in safe offline control and imitation learning.

In the following, we focus on how recent diffusion-based approaches in robotics have attempted to incorporate constraints in order to ensure safety and feasibility across a variety of tasks. Some approaches enforce constraints during training: Bastek et al. [2025] integrate physical laws into the training objective to ensure physically consistent outputs; Giannone et al. [2023] align sampling trajectories with constrained optimization paths; and Power et al. [2023] separately train on different constraints and combine them at inference. Others address constraints during inference: Carvalho et al. [2023] condition the sampling process on goal-reaching and obstacle avoidance; Christopher et al. [2024], Xiao et al. [2023] enforce feasibility through projection steps, albeit with significant computational overhead; Römer et al. [2024] incorporate model-based projections directly into the backward diffusion process to enforce constraint satisfaction during trajectory generation, avoiding the need for external post-sampling corrections; and Yu et al. [2024] generate local collision-free motions through conditional sampling. Several methods handle constraints in both training and inference phases, such as Ajay et al. [2023] for decision-making, Gong et al. [2025] with trajectory-level diffusion, and Botteghi et al. [2023], which train safe priors and apply runtime filtering. Among these, Yu et al. [2024] primarily handle inference-time constraints, while Botteghi et al. [2023] combine both stages. Overall, incorporation during training time promotes inherent feasibility, while inference-time methods offer flexibility at the cost of higher computational complexity during inference.

Compared to prior inference-time approaches, our CoDiG framework handles constraints by augmenting score updates with lightweight barrier gradients during sampling, without relying on projections, auxiliary models, or simulators, unlike classifier-guided [Dhariwal and Nichol, 2021] and energy-guided diffusion [Lu et al., 2023b] that learn auxiliary networks at each step. This provides efficient, continuous guidance toward feasible trajectories while preserving the generative flexibility of diffusion models with a time-dependent weight γ_t that ramps up during denoising. Warm-start initialization further accelerates convergence and enhances sampling stability, enabling real-time deployment (2.5 Hz on hardware). Unlike previous works mainly evaluated in simulation or in quasi-static environments, we demonstrate CoDiG on a real-world autonomous racing platform, where strict dynamic feasibility

and rapid obstacle avoidance are critical. While Sheebaelhamd et al. [2025] have also suggested autoregressive architectures as an alternative to diffusion-based generation, it is unclear whether constrained-aware generation via barrier functions is also effective with these architectures. These aspects highlight the unique contributions of CoDiG in enabling efficient, reliable, and real-time constraint handling within generative robotic planning. See Appendix A for a compact side-by-side summary of training- vs. inference-time strategies (Table 1).

3 Generative imitation via diffusion

From an imitation learning perspective, diffusion models amount to learning a generative model over expert demonstrations conditioned on the environment context (e.g., racetrack, obstacles), without access to rewards or interaction. Recent advances in score-based generative modeling formulate diffusion processes via stochastic differential equations (SDEs), offering a continuous-time view of forward noise injection and reverse denoising [Song et al., 2021]. Since our work builds on this foundation, we briefly review score-based generative modeling and introduce the notation used throughout the paper.

3.1 Preliminaries

Let $x_0 \in \mathbb{R}^d$ denote a noise-free data sample drawn from the data distribution $p_0(x)$. A score-based generative model defines a continuous-time diffusion process $\{x_t\}_{t \in [0, T]}$, where t denotes the diffusion time, such that x_T becomes approximately Gaussian. It is important to note that throughout this paper, we encounter two notions of “time”: here, t refers to the artificial diffusion time governing the processes, while later, τ will denote the physical time in real-world dynamical systems.

Diffusion Process. The forward diffusion process gradually perturbs the data by solving the following SDE:

$$dx_t = f(x_t, t) dt + g(t) dw_t, \quad t \in [0, T], \quad x_0 \sim p_0,$$

where $x_t \in \mathbb{R}^d$ is the perturbed data at time t , $f : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}^d$ is the drift term, $g : [0, T] \rightarrow \mathbb{R}$ is the scalar-valued diffusion term, and $w_t \in \mathbb{R}^d$ denotes a standard Wiener process.

A common instantiation of the diffusion process is the *Ornstein–Uhlenbeck (OU) process* [Øksendal, 1995], in which the drift pulls x_t toward a mean $\mu \in \mathbb{R}^d$:

$$dx_t = \beta(t)(\mu - x_t) dt + g(t) dw_t, \quad t \in [0, T], \quad (1)$$

where $\beta(t)$ is a positive scalar-valued function controlling the drift strength. In this case, the OU process admits a closed-form solution for the mean and variance of the marginal distribution of x_t . Specifically, letting

$$\bar{\beta}_t := \exp\left(-\int_0^t \beta(\nu) d\nu\right),$$

then the marginal distribution of x_t is Gaussian:

$$p_t(x_t | x_0) = \mathcal{N}\left(x_t; \mu - (\mu - x_0)\bar{\beta}_t, \frac{g(t)^2}{2\beta(t)}(\mathbb{I} - \bar{\beta}_t^2\mathbb{I})\right), \quad (2)$$

where $\mathbb{I} \in \mathbb{R}^{d \times d}$ is the identity matrix.

Sampling Process. To generate new data, one samples from the reverse-time SDE corresponding to the forward process. Under suitable regularity conditions, this reverse SDE takes the form [Anderson, 1982]:

$$dx_t = \left[f(x_t, t) - g(t)^2 \nabla_x \log p_t(x_t)\right] dt + g(t) d\tilde{w}_t, \quad t \in [0, T], \quad x_T \sim p_{x_T}, \quad (3)$$

where \tilde{w}_t is a standard Wiener process running backward in time, and $\nabla_x \log p_t(x_t)$ is the score function of the marginal distribution.

In practice, the score function is unknown and approximated by a neural network $s_\theta(x_t, t)$ trained using denoising score matching. The training objective minimizes the expected squared error between the predicted score and the true score:

$$\mathbb{E}_{t \sim \mathcal{U}[0, T]} \mathbb{E}_{x_0 \sim p_0(x)} \mathbb{E}_{x_t \sim p_t(x_t | x_0)} \left[|s_\theta(x_t, t) - \nabla_x \log p_t(x_t | x_0)|^2 \right],$$

where $|\cdot|$ denotes the ℓ_2 -norm, and $\mathcal{U}[0, T]$ the uniform distribution with support $[0, T]$.

3.2 Constraint-aware diffusion guidance

To ensure the generated trajectories are not only expert-like but also physically feasible and safe, we inject constraint-awareness directly into the generative process. Specifically, we augment the score updates with barrier gradients. Firstly, we specify the functional forms of the drift and diffusion terms in (1) for concreteness and clarity. It is important to emphasize that the proposed framework does not rely on these specific choices - the following definitions are adopted purely for illustrative purposes and to remain consistent with the experimental setup described later.

We let $\mu = 0$, and define the drift term and the diffusion term as

$$f(x_t, t) = -\beta(t) x_t, \quad g(t) = \sqrt{2\beta(t)}, \quad t \in [0, T],$$

which yields the so-called variance preserving SDE [Song et al., 2021], where $g(t)^2 = 2\beta(t)$ holds for all $t \in [0, T]$ such that the marginal variance of x_t is preserved over time. This specific choice ensures that the forward process remains stable and tractable for training and sampling, while still allowing for an expressive and well-defined reverse-time generative process. Under this formulation, when the terminal time T is sufficiently large, the forward diffusion process described by (1) converges to a standard Gaussian distribution. As analyzed in Song et al. [2021], the term $\sqrt{2\beta(t)}$ should grow with time, requiring $\beta(t)$ to be strictly increasing.

For simplicity and numerical stability, we normalize the diffusion process to $t \in [0, 1]$. To ensure convergence to a standard Gaussian, the diffusion term $\sqrt{2\beta(t)}$ must grow rapidly within this interval. In our implementation, we model $\beta(t)$ as a quadratic function, $\beta(t) = r_1 t^2 + r_0$, with parameters detailed in Appendix D.3. In this case, (3) can be reformulated as:

$$dx_t = [-\beta(t) x_t - 2\beta(t) \nabla_x \log p_t(x_t)] dt + \sqrt{2\beta(t)} d\tilde{w}_t, \quad t \in [0, 1], \quad x_1 \sim p_{x_1}. \quad (4)$$

Next, we consider the marginal distribution $p_t(x_t)$, which represents the probability distribution of a sample at an intermediate time step, in the absence of constraints. Before incorporating constraints into this distribution, we first introduce the following definitions. Let $c : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^k$ denote a time-dependent constraint function, encoding the safety or feasibility requirements of the system. We define the feasible region at time $\tau \geq 0$ as

$$\mathcal{C}_\tau := \{x \in \mathbb{R}^d \mid c(x, \tau) \leq 0\},$$

where the inequality is interpreted element-wise. Naturally, when constraints are introduced, the distribution of interest becomes the conditional distribution:

$$p_t(x_t \mid \mathcal{C}_\tau), \quad t \in [0, 1], \quad \tau \geq 0.$$

These constraints may encode different forms of feasibility or safety requirements, depending on the task setting. For example, in autonomous racing, \mathcal{C}_τ refers to the obstacle-free, drivable region of a racing track. While in diffusion-based control policies, \mathcal{C}_τ must account for system dynamics, as the use of visual feedback necessitates control that respects the underlying physical constraints of the system. Here, we use the time subscript τ to emphasize that such constraints can be time-varying, which is often the case in dynamic or interactive environments. For simplicity, and without loss of clarity, we will omit this subscript in the following when no confusion arises.

Several existing methods attempt to directly model the marginal distribution $p_t(x_t \mid \mathcal{C})$ by injecting the constraint representation into the diffusion model architecture [Ho and Salimans, 2022]. While effective in big-data domains such as image synthesis, these approaches face significant limitations in the context of robotics: (i) Learning $p_t(x_t \mid \mathcal{C})$ from scratch requires many expert demonstrations that satisfy \mathcal{C} , which are often expensive or impractical to collect in robotics. (ii) Since \mathcal{C} is often

time-varying and task-specific, models trained on a fixed distribution may fail to generalize to unseen or dynamic constraints at test time.

To overcome these limitations, we leverage the known structure of the constraint \mathcal{C} during sampling to dynamically guide the generation process. We propose an alternative formulation of the constrained distribution $p_t(x_t | \mathcal{C})$, which does not require learning the conditional model directly from data:

$$p_t(x_t | \mathcal{C}) = p_t(x_t) \frac{e^{-\gamma_t V(x_t; \mathcal{C})}}{Z_t},$$

where $Z_t := \int_{\mathbb{R}^d} p_t(x) e^{-\gamma_t V(x; \mathcal{C})} dx$ is a normalization constant. The barrier function $V : \mathbb{R}^d \rightarrow \mathbb{R}_+$ assigns large values to infeasible data points, while remaining close to zero within the feasible region. Intuitively, applying the barrier function pushes the probability of infeasible data points toward zero. Importantly, the barrier function is *derived from explicit task constraints* (e.g., obstacle clearance and near time-optimality) rather than a heuristic penalty, and we use a *time-dependent weight* γ_t that ramps up across denoising steps to enforce constraints more strongly as samples approach the data manifold. As a result, the constrained distribution focuses its support almost entirely on the feasible region. We substitute the above formula into the score function and get

$$\nabla_x \log p_t(x_t | \mathcal{C}) = \nabla_x \log p_t(x_t) - \gamma_t \nabla_x V(x_t; \mathcal{C}), \quad (5)$$

where the normalization constant Z_t vanishes when taking the gradient of the log-probability, and hence does not affect the reverse-time dynamics. By substituting the right-hand side of (5) into (4), we obtain the modified reverse SDE that incorporates constraint information:

$$dx_t = \beta(t) [-x_t - (1 + \eta) (\nabla_x \log p_t(x_t) - \gamma_t \nabla_x V(x_t; \mathcal{C}))] dt + \eta \sqrt{2\beta(t)} d\tilde{w}_t, \quad (6)$$

where a constant $\eta \in [0, 1]$ is introduced to accelerate convergence and enhance the stability of the sampling process [Song and Ermon, 2020]. We observe that the first term on the right-hand side of (5) corresponds exactly to the unconstrained score function defined in (4). This term can still be approximated by the neural network $s_\theta(x_t, t)$ trained without considering any constraints. Crucially, the effect of the constraint appears only during the denoising process, in an explicit gradient-based form - as an additive term derived from the constraint potential (e.g., a barrier function). This formulation significantly reduces the need for large amounts of constraint-compliant training data, as the constraint is not encoded in the network itself but *instead injected at inference time*. Moreover, because the constraint enters the reverse SDE as a differentiable time-varying potential, the framework can naturally accommodate dynamic, time-varying constraints.

It is important to note, as pointed out by Bastek et al. [2025], that applying constraints to data that is close to pure noise in diffusion models is not meaningful. Therefore, during the denoising process - i.e., as t decreases from one to zero - we gradually increase the value of γ_t starting from zero at $t = 1$. This progressive scheduling is crucial for ensuring the stability of the denoising process. The specific design of γ_t is detailed in Appendix B.

4 Constraint-aware imitation in autonomous racing

We now demonstrate how a policy trained via constraint-aware imitation learning can be deployed in a real-world robotic task. We choose autonomous racing with obstacle avoidance as a representative challenge to evaluate constraint satisfaction, physical feasibility, and generalization. While this example serves to ground our discussion, the barrier function design and constraint-handling mechanisms are task-agnostic. Thus, our framework is not limited to autonomous racing but serves as a general-purpose solution for safety-critical robotics applications. For details on expert demonstration, dataset construction, diffusion model architecture, and training procedures, please refer to Appendix D.

4.1 Constraint-aware barrier function

For the considered application, the barrier function, which is instantiated from task constraints (safety and near time-optimality) rather than tuned heuristics, is designed as follows:

$$V(\hat{y}, \hat{\phi}; \mathcal{C}) = \underbrace{\sum_{k=0}^{N-1} \alpha \mathbb{1}\{\hat{y}_k \notin \mathcal{C}_k\}}_{\text{first part}} + \underbrace{\frac{\epsilon}{2} |\hat{y}_k - \hat{y}_{\text{nominal},k}|^2 + \frac{\epsilon}{2} |\hat{\phi}_k - \hat{\phi}_{\text{nominal},k}|^2}_{\text{second part}}, \quad (7)$$

where the symbol $\mathbb{1}\{\cdot\}$ denotes the indicator function⁴, and the subscript $(\cdot)_{\text{nominal}}$ refers to the time-optimal solution computed offline in the absence of obstacles, which serves as a reliable reference. Here, N represents the number of discrete points obtained by uniformly sampling along the track center line. In our setting, \hat{y} denotes the lateral displacement and $\hat{\phi}$ represents the yaw angle in the Frenet coordinate system (see Appendix D.1). The feasible region \mathcal{C}_k is also defined in the Frenet frame, capturing the obstacle-free area at each sampled position along the track.

In (7), the first part accounts for time-varying obstacles by guiding the sampling process toward safe, obstacle-free regions. The second part addresses the missing curvature information of the track during training, which is intentionally discarded when transforming trajectories into the local Frenet coordinate system. Beyond promoting near time-optimality without requiring global geometric knowledge of the track, this part also facilitates the generation of dynamically feasible trajectories. The local representation ensures that the resulting motions adhere more closely to the physical and kinematic constraints of the system. The positive constants α and ϵ are tunable hyperparameters that balance the importance of the two components. Specifically, α modulates the influence of physical safety constraints, while ϵ regulates the adherence to nominal time-optimality.

It is worth noting that the design of the barrier function is not unique and can be tailored to the specific task. While such customization may require a small amount of tuning effort, it is negligible compared to the cost of collecting expert demonstrations, especially in robotics domains where data is expensive. This makes our framework both flexible and data-efficient.

4.2 Constraint-aware inference

We train the diffusion model as described in Appendix D.2 for 500 epochs. During inference, we applied the Euler-Maruyama discretization (8), which corresponds to the discretized version of (6). The denoising process proceeds from $t = 1.0$ to $t = 0.0$ in 1000 steps, gradually transforming samples from noise to data. The results are shown in Fig. 1. Fig. 1a illustrates the denoising process without using the barrier function, while Fig. 1b shows the effect of the proposed constraint-aware guided generation. Each row depicts intermediate generation results at $t = 1\text{ s}$, 0.591 s , 0.002 s , from left to right. The black points and arrows indicate the evolving trajectories in the z - y plane. See Appendix B for the concrete values of the hyperparameters used during inference.

As shown in Fig. 1a, a diffusion model trained purely on expert demonstrations (i.e., standard imitation learning without constraints) learns to generate trajectories that broadly follow the racetrack and respect basic physical layout. However, due to limited number of expert demonstrations, such models struggle to consistently avoid obstacles, especially in areas not well-covered by the training data. This behavior highlights a key limitation of imitation learning methods that rely solely on demonstration data: they lack explicit mechanisms to enforce safety or task-specific constraints.

After incorporating the barrier function into the inference process, the model is able to systematically guide samples away from unsafe regions while still preserving expert-like behavior. As seen in Fig. 1b, the guided trajectories successfully avoid all obstacles and produce dynamically feasible motion. Moreover, the model converges faster during the denoising process due to stronger guidance, resulting in cleaner and more structured trajectories earlier in the sampling process.

Interestingly, beyond simply avoiding obstacles, the guided trajectories also exhibit characteristics of the expert’s behavior - in particular, they closely track the nominal time-optimal solution even in the presence of obstacles. This is achieved via the second part of the barrier function, which softly penalizes deviation from a precomputed time-optimal baseline without any obstacles. As a result, the generated motion achieves a balance between safety and efficiency, a hallmark of high-quality imitation in safety-critical settings.

4.3 Near time-optimality

To quantitatively evaluate the near time-optimality of the generated trajectories, we compare them with offline-computed time-optimal solutions [Verschuere et al., 2016] across multiple obstacle scenarios. Fig. 2 presents several representative obstacle configurations extracted from a real-world experiment. In each scenario, the red trajectory denotes the real-time obstacle-avoidance path generated by the CoDiG framework, while the black trajectory represents the time-optimal path computed offline under

⁴One possible differentiable approximation is provided in Appendix E.

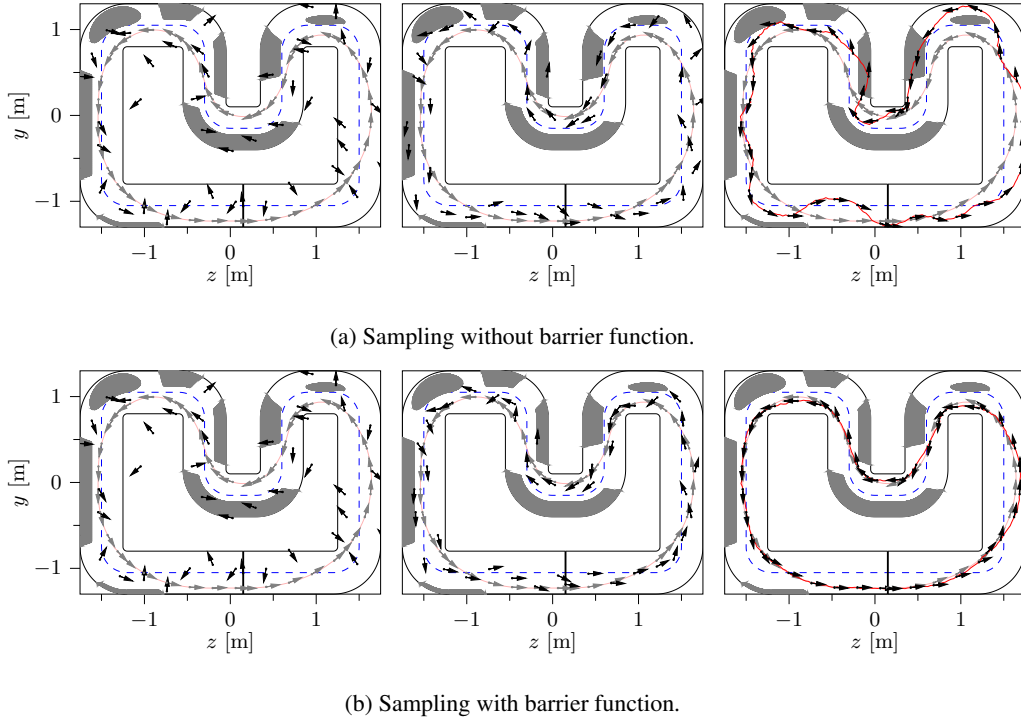


Figure 1: Intermediate denoising results during sampling at three representative time steps $t = 1$ s, 0.591 s, 0.002 s, from left to right. (a) Sampling process without the barrier function. (b) Sampling process with the proposed barrier function. Black dots and arrows represent the generated trajectory points and their heading directions in the global frame.

the same obstacle layout. As illustrated in Fig. 2, CoDiG produces trajectories that closely match the expert’s behavior while also avoiding obstacles in real time. Minor deviations typically arise in two cases: (i) conservative safety margins near obstacles, and (ii) slight smoothing near sharp corners for dynamic feasibility. Nonetheless, the overall match confirms that the model has successfully internalized the expert’s motion style and planning objectives, achieving near time-optimal behavior even under dynamic constraints.

5 Real-world experiments

We evaluate CoDiG in experiments on a real-world miniature autonomous racing platform [Bodmer et al., 2024, Carron et al., 2023]. For more details on the experimental platform, the obstacle setup, and a flowchart illustrating how the CoDiG framework is deployed to achieve real-time obstacle avoidance, please refer to Appendix G.

5.1 Warm-starting

Real-time obstacle avoidance requires not only safe trajectories but also fast replanning. As shown in Sec. 4.2, our diffusion model with a barrier function produces high-quality trajectories after 1000 denoising steps, but this results in a low sampling frequency of 0.25 Hz, which is insufficient for real-time racing.

While various acceleration techniques exist [Song et al., 2022, Palo et al., 2025, Zhang and Chen, 2023], we propose a warm-start strategy tailored to robotic control. Unlike standard diffusion generation, which samples each trajectory from pure noise, our proposed warm-start technology perturbs the previous output with small noise and reuses it as the next input. This maintains temporal consistency, reduces trajectory variance, and improves control stability [Morari and Lee, 1999]. By




Figure 2: Comparison between trajectories generated in real time by CoDiG (red) and offline-computed time-optimal trajectories (black) under various obstacle configurations.

promoting smooth transitions between consecutive trajectories, warm-starting significantly lowers the number of denoising steps required and enhances real-time feasibility. A detailed analysis and comparison of results with and without warm-starting are provided in Appendix F.

5.2 Experimental results

Through the integration of our warm-start technique, we achieve a sampling frequency of 2.5 Hz on a computer equipped with an NVIDIA RTX 4090 GPU. While there is still potential for further acceleration, this performance already satisfies the real-time requirements of obstacle avoidance in racing scenarios.

We successfully deployed CoDiG on our real-world autonomous racing platform for real-time trajectory planning. A tracking model predictive control (TMPC) [Limon et al., 2008, Soloperto et al., 2023] is employed to follow the planned trajectories. Notably, the TMPC operates without any knowledge of obstacles, relying solely on the reference trajectories for control. Fig. 3 illustrates two representative obstacle avoidance maneuvers during autonomous driving. In both Fig. 3a and Fig. 3b, the red lines represent the trajectories planned by CoDiG, the gray circles denote static obstacles, while the black circles indicate dynamic obstacles. The black dashed lines show the predicted trajectories generated by the TMPC as it attempts to follow the red reference trajectory. Each sequence from left to right captures a complete avoidance cycle: (i) Obstacle Encroachment: An obstacle intrudes into a previously feasible trajectory, making it infeasible. (ii) Replanning: The planner detects the encroachment and generates a new collision-free trajectory. (iii) Successful Avoidance: The vehicle safely bypasses the obstacle.

As shown in the figures, the predicted trajectories from the TMPC closely align with the reference trajectories generated by the diffusion model. This highlights that the planned trajectories are closely aligned with physical feasibility, enabled by the barrier function, which is crucial for effective tracking

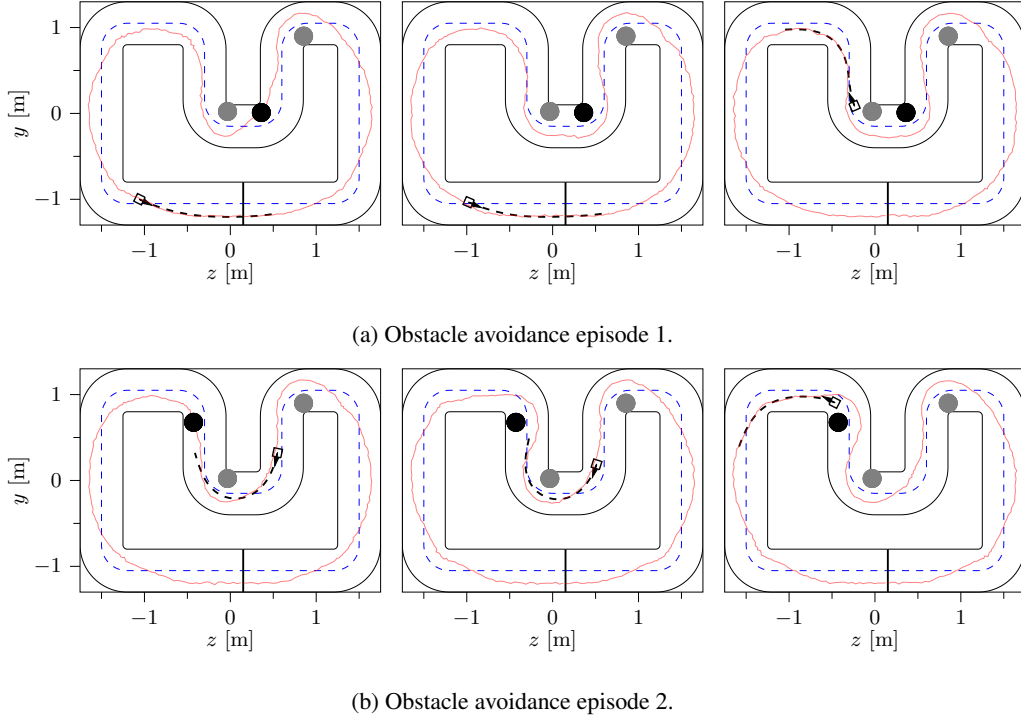


Figure 3: Real-world demonstration of real-time obstacle avoidance using CoDiG. Red lines represent the planned trajectories generated by the CoDiG diffusion planner. Gray circles indicate static obstacles, and black circles represent dynamic obstacles. Black dashed lines show the predicted trajectory from the TMPC while following the reference plan. Each episode illustrates a complete avoidance cycle: obstacle encroachment, real-time replanning, and successful passage.

performance. Additionally, even in the presence of obstacles, the generated trajectories maintain near time-optimality, indicating that the planner does not overly sacrifice efficiency for safety.

Finally, thanks to the warm-start strategy, significant replanning is only triggered when the obstacle actually interferes with the current path. In static conditions, consecutive trajectories remain almost unchanged, ensuring system stability. We conducted five experimental trials, each consisting of 15 racing laps, across ten different obstacle configurations. The framework achieved a 100 % success rate in obstacle avoidance, demonstrating its robustness and reliability in diverse scenarios.

6 Conclusion

In this work, we introduce CoDiG, a constraint-aware diffusion-based imitation learning framework for robotics. CoDiG learns offline from a small number of expert demonstrations, without requiring environment interaction or reward supervision. At inference time, it integrates physical and safety constraints via barrier-guided score updates, enabling the generation of feasible and safe trajectories in dynamic and constrained environments.

We demonstrate CoDiG on autonomous racing with dynamic obstacles, achieving robust real-world performance with reliable obstacle avoidance and precise tracking at 2.5 Hz. The resulting behavior closely mimics time-optimal expert trajectories while ensuring safety, making CoDiG a promising direction for deploying scalable and safe imitation learning in safety-critical domains such as autonomous racing.

Acknowledgments

We thank Matteo Facchino for providing code related to time-optimal control solvers. We also gratefully acknowledge Jan-Hendrik Bastek for the insightful discussions on constraint handling in diffusion models. We thank the German Research Foundation and the Max-Planck ETH Center for Learning Systems for the support.

References

- André Correia and Luís A Alexandre. A Survey of Demonstration Learning. *Robotics and Autonomous Systems*, 182(0):104812–104841, 2024.
- Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A Survey of Robot Learning from Demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.
- Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. In *International Conference on Artificial Intelligence and Statistics*, pages 627–635, 2011.
- Jonathan Ho and Stefano Ermon. Generative Adversarial Imitation Learning. In *Advances in Neural Information Processing Systems*, pages 1–9, 2016.
- Justin Fu, Katie Luo, and Sergey Levine. Learning Robust Rewards with Adversarial Inverse Reinforcement Learning. *arXiv:1710.11248*, pages 1–15, 2018.
- Yiren Lu, Justin Fu, George Tucker, Xinlei Pan, Eli Bronstein, Rebecca Roelofs, Benjamin Sapp, Brandyn White, Aleksandra Faust, Shimon Whiteson, Dragomir Anguelov, and Sergey Levine. Imitation Is Not Enough: Robustifying Imitation with Reinforcement Learning for Challenging Driving Scenarios. In *International Conference on Intelligent Robots and Systems*, pages 7553–7560, 2023a.
- Rosa Wolf, Yitian Shi, Sheng Liu, and Rania Rayyes. Diffusion Models for Robotic Manipulation: A Survey. *arXiv:2504.08438*, pages 1–28, 2025.
- Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 0(0):1–21, 2024.
- Julen Urain, Niklas Funk, Jan Peters, and Georgia Chalvatzaki. SE(3)-DiffusionFields: Learning smooth cost functions for joint grasp and motion optimization through diffusion. In *International Conference on Robotics and Automation*, pages 5923–5930, 2023.
- Kota Kondo, Andrea Tagliabue, Xiaoyi Cai, Claudius Tewari, Olivia Garcia, Marcos Espitia-Alvarez, and Jonathan P. How. CGD: Constraint-Guided Diffusion Policies for UAV Trajectory Planning. *arXiv:2405.01758*, pages 1–8, 2024.
- Norman Di Palo, Leonard Hasenclever, Jan Humplik, and Arunkumar Byravan. Diffusion Augmented Agents: A Framework for Efficient Exploration and Transfer Learning. In *Conference on Lifelong Learning Agents*, pages 268–284, 2025.
- Kin Man Lee, Sean Ye, Qingyu Xiao, Zixuan Wu, Zulfiqar Zaidi, David B. D’Ambrosio, Pannag R. Sanketi, and Matthew Gombolay. Learning Diverse Robot Striking Motions with Diffusion Models and Kinetically Constrained Gradient Guidance. *arXiv:2409.15528*, pages 1–8, 2025.
- Zhixuan Liang, Yao Mu, Mingyu Ding, Fei Ni, Masayoshi Tomizuka, and Ping Luo. AdaptDiffuser: Diffusion Models as Adaptive Self-evolving Planners. In *International Conference on Machine Learning*, pages 20725–20745, 2023.
- Jan-Hendrik Bastek, WaiChing Sun, and Dennis M. Kochmann. Physics-Informed Diffusion Models. *arXiv:2403.14404*, pages 1–26, 2025.
- Giorgio Giannone, Akash Srivastava, Ole Winther, and Faez Ahmed. Aligning Optimization Trajectories with Diffusion Models for Constrained Design Generation. In *Advances in Neural Information Processing Systems*, pages 51830–51861, 2023.

- Thomas Power, Rana Soltani-Zarrin, Soshi Iba, and Dmitry Berenson. Sampling Constrained Trajectories Using Composable Diffusion Models. In *International Conference on Intelligent Robots and Systems*, pages 1–5, 2023.
- João Carvalho, An T. Le, Mark Baierl, Dorothea Koert, and Jan Peters. Motion Planning Diffusion: Learning and Planning of Robot Motions with Diffusion Models. In *International Conference on Intelligent Robots and Systems*, pages 1916–1923, 2023.
- Jacob K. Christopher, Stephen Baek, and Ferdinando Fioretto. Constrained Synthesis with Projected Diffusion Models. *arXiv:2402.03559*, pages 1–20, 2024.
- Wei Xiao, Tsun-Hsuan Wang, Chuang Gan, and Daniela Rus. SafeDiffuser: Safe Planning with Diffusion Probabilistic Models. *arXiv:2306.00148*, pages 1–19, 2023.
- Ralf Römer, Alexander von Rohr, and Angela P. Schoellig. Diffusion Predictive Control with Constraints. *arXiv:2412.09342*, pages 1–14, 2024.
- Wenhao Yu, Jie Peng, Huanyu Yang, Junrui Zhang, Yifan Duan, Jianmin Ji, and Yanyong Zhang. LDP: A Local Diffusion Planner for Efficient Robot Navigation and Collision Avoidance. In *International Conference on Intelligent Robots and Systems*, pages 5466–5472, 2024.
- Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is Conditional Generative Modeling All You Need for Decision-Making? *arXiv:2211.15657*, pages 1–24, 2023.
- Ze Gong, Akshat Kumar, and Pradeep Varakantham. Offline Safe Reinforcement Learning Using Trajectory Classification. In *Conference on Artificial Intelligence*, pages 16880–16887, 2025.
- Nicolò Botteghi, Federico Califano, Mannes Poel, and Christoph Brune. Trajectory Generation, Control, and Safety with Denoising Diffusion Probabilistic Models. *arXiv:2306.15512*, pages 1–18, 2023.
- Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*, pages 8780–8794, 2021.
- Cheng Lu, Huayu Chen, Jianfei Chen, Hang Su, Chongxuan Li, and Jun Zhu. Contrastive Energy Prediction for Exact Energy-Guided Diffusion Sampling in Offline RL. In *International Conference on Machine Learning*, pages 22825–22855, 2023b.
- Ziyad Sheebaelhamd, Michael Tschannen, Michael Muehlebach, and Claire Vernade. Quantization-Free Autoregressive Action Transformer. *arXiv:2503.14259*, pages 1–15, 2025.
- Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Machine Learning*, pages 10362–10383, 2021.
- Bernt Øksendal. *Stochastic Differential Equations*. Springer Berlin Heidelberg, 1995.
- Brian D.O. Anderson. Reverse-Time Diffusion Equation Models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. *arXiv:2207.12598*, pages 1–14, 2022.
- Yang Song and Stefano Ermon. Improved Techniques for Training Score-Based Generative Models. In *Advances in Neural Information Processing Systems*, pages 12438–12448, 2020.
- Robin Verschueren, Mario Zanon, Rien Quirynen, and Moritz Diehl. Time-optimal Race Car Driving using an Online Exact Hessian based Nonlinear MPC Algorithm. In *European Control Conference*, pages 141–147, 2016.
- Sabrina Bodmer, Lukas Vogel, Simon Muntwiler, Alexander Hansson, Tobias Bodewig, Jonas Wahlen, Melanie N. Zeilinger, and Andrea Carron. Optimization-Based System Identification and Moving Horizon Estimation Using Low-Cost Sensors for a Miniature Car-Like Robot. *arXiv:2404.08362*, pages 1–11, 2024.

- Andrea Carron, Sabrina Bodmer, Lukas Vogel, René Zurbrügg, David Helm, Rahel Rickenbach, Simon Muntwiler, Jerome Sieber, and Melanie N. Zeilinger. Chronos and CRS: Design of a miniature car-like robot and a software framework for single and multi-agent robotics and control. In *International Conference on Robotics and Automation*, pages 1371–1378, 2023.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. *arXiv:2010.02502*, pages 1–22, 2022.
- Qinsheng Zhang and Yongxin Chen. Fast Sampling of Diffusion Models with Exponential Integrator. *arXiv:2204.13902*, pages 1–33, 2023.
- Manfred Morari and Jay H. Lee. Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(4-5):667–682, 1999.
- Daniel Limon, Ignacio Alvarado Aldea, Teodoro Alamo, and Eduardo F. Camacho. MPC for tracking piecewise constant references for constrained linear systems. *Automatica*, 44(9):2382–2387, 2008.
- Raffaele Soloperto, Johannes Köhler, and Frank Allgöwer. A Nonlinear MPC Scheme for Output Tracking Without Terminal Ingredients. *Transactions on Automatic Control*, 68(4):2368–2375, 2023.
- Timothy Sauer. Numerical Solution of Stochastic Differential Equations in Finance. In *Handbook of Computational Finance*, pages 529–550. Springer, 2012.
- Hugh C. Crenshaw and Leah Edelstein-Keshet. Orientation by Helical Motion—II. Changing the Direction of the Axis of Motion. *Bulletin of Mathematical Biology*, 55(1):213–230, 1993.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. *arXiv:2112.10752*, pages 1–45, 2022.
- James Hensman, Nicolas Durrande, and Arno Solin. Variational Fourier Features for Gaussian Processes. *Journal of Machine Learning Research*, 18(151):1–52, 2018.

A Comparison of Constraint-Enforcement Strategies

This section provides a concise, side-by-side comparison of how constraints are handled across methods in Table 1: whether constraints are enforced during training, inference, or both; the guidance source (e.g., classifier, energy, projection, or physics-based barrier); whether learned auxiliary models or projection steps are required; and whether any real-time evidence is reported.

B Hyperparameters

During inference, the hyperparameters are set as follows:

$$\eta = 0.1, \quad \alpha = 0.4, \quad \epsilon = 16.0.$$

In practice, we set α in the range 10–15, whereas ϵ is chosen at a smaller magnitude, typically 0.1–0.5. We also observe that our method is not particularly sensitive to these hyperparameters across environments: the same settings generalize well on both training and test scenarios.

In particular, the time-varying weight $\gamma(t)$ is assigned non-uniform values according to the following scheme:

$$\gamma(t) = \frac{\hbar_1}{1 + \exp(-\hbar_2(\hbar_3 - t))}, \quad t \in [0, 1],$$

where $\hbar_1 = 1.0$, $\hbar_2 = 50.0$, and $\hbar_3 = 0.7$. The choice of the function γ is not unique. The guiding principle is to introduce the constraint progressively during denoising so that its gradient increasingly shapes the samples as they approach the data manifold. Any monotone schedule that follows this principle (e.g., linear or logistic ramp-up) works in practice.

C Discrete-time integration

Assuming a denoising process over $M \in \mathbb{N}_{++}$ steps, we partition the interval $[0, 1]$ non-uniformly as follows:

$$t_k = \left(1 - \frac{k}{M}\right)^p, \quad k = 0, \dots, M,$$

where $p = 2.2$ in our case. Starting from an initial sample x_0 drawn from a standard Gaussian distribution, we perform denoising according to the following discrete Euler-Maruyama [Sauer, 2012] update scheme:

$$\begin{aligned} \bar{x}_{k+1} &= x_k + \beta(t_k) [-x_k - (1 + \eta)(s_\theta(x_k, t_k) - \gamma_{t_k} \nabla_x V(x_k; \mathcal{C}))] \Delta t_k, \\ x_{k+1} &= \bar{x}_{k+1} + \eta \sqrt{2\beta(t_k)} \sqrt{|\Delta t_k|} \sigma_k, \quad k = 0, \dots, M-1, \end{aligned} \tag{8}$$

where $\Delta t_k = t_{k+1} - t_k$ denoting the step size between successive time points. The noise term $\sigma_k \in \mathbb{R}^d$ is sampled from a standard Gaussian distribution. Here, \bar{x} denotes the mean estimate at each step, while x denotes the noisy sample.

D Data and model pipeline

In this section, we describe the pipeline used to train our diffusion model for obstacle avoidance in racing scenarios. We begin by presenting our data collection process, where expert demonstrations are gathered to reflect optimal driving behaviors in the presence of obstacles. Then, we introduce a data augmentation strategy that diversifies the training distribution while preserving expert intent. Next, we detail the architecture of our proposed diffusion model, which is adapted to handle time-conditioned inputs and spatial constraints relevant to the racing task. Finally, we present the training results of the diffusion model under various input configurations, demonstrating how different modalities affect the training performance.

D.1 Expert dataset construction

Even on a miniature autonomous racing platform, collecting expert demonstrations via manual teleoperation is highly challenging and time-consuming. Therefore, we generate expert data by

Method	Stage	Guidance source	Aux. model	Projection	RT evidence
CoDiG (ours)	Inference	Barrier gradient (physics & safety)	✗	✗	✓ (2.5 Hz)
Dhariwal and Nichol [2021]	Inference	Classifier $\nabla \log p(y x)$	✓	✗	–
Lu et al. [2023b]	Inference	Energy gradient	✓	✗	–
Christopher et al. [2024]	Inference	Projection operator	○	✓	–
Xiao et al. [2023]	Inference	Model-based constraints	○	○	–
Römer et al. [2024]	Inference	Local collision-free conditioning	○	✗	–
Yu et al. [2024]	Inference				
Bastek et al. [2025]	Train	Physics-informed training loss	✗	✗	–
Giannone et al. [2023]	Train	Align to constrained opt. paths	✗	✗	–
Power et al. [2023]	Train	Multi-constraint training	○	✗	–
Ajay et al. [2023]	Both	Constrained decision diffusion	○	○	–
Gong et al. [2025]	Both	Trajectory-level diffusion w/ constraints	○	○	–
Botteghi et al. [2023]	Both	Safe priors + runtime filtering	○	○	–

Table 1: Comparison across enforcement phases. “Aux.” indicates learned auxiliary networks; “Projection” indicates projection steps; RT = real-time. “✓” means the item is required, “✗” means it is not needed, “○” denotes that it is optional, and “–” means that no result is reported.

solving a time-optimal control problem [Verschuere et al., 2016], including car states and control inputs.

To collect expert data, we randomly place obstacles on the track and solve the aforementioned time-optimal control problem to obtain optimal driving trajectories with continuous looping and corresponding control inputs. An example is shown in Fig. 4a, where the gray regions indicate obstacles. The red curve shows the trajectory in the z - y plane, and the black rectangles and arrows illustrate the approximate shape and orientation of the vehicle, respectively, reflecting the fact that the vehicle is not treated as a point mass to account for the system dynamics.

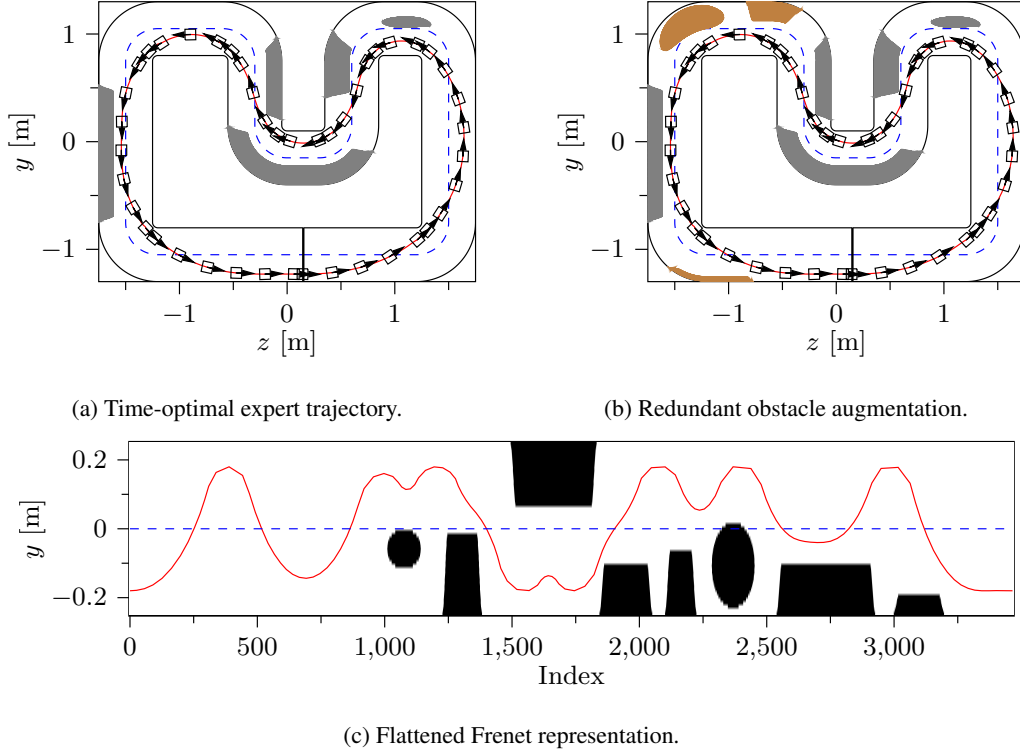


Figure 4: (a) A time-optimal trajectory (red line) computed for a given obstacle configuration (gray regions) on the racing track. Black rectangles and arrows indicate the approximate vehicle shape and heading. (b) Redundant obstacles (brown regions) added in areas that do not affect the trajectory, providing data augmentation without solving additional optimal control problems. (c) A flattened track representation in the local Frenet coordinate system, visualizing both the trajectory (red line) and obstacles (black regions).

As previously mentioned, collecting expert data is expensive. Solving a single time-optimal control problem takes around 10 minutes on average. To address this limitation, we propose a method for dataset augmentation. We observe that once the time-optimal solution is obtained for a given map (with a specific obstacle configuration), adding extra obstacles within the safe region that do not interfere with the trajectory will not alter the time-optimal solution. These redundant obstacles - illustrated as brown regions in Fig. 4b - can be arbitrarily placed without affecting the outcome. Based on this observation, we first collect 100 trajectories by solving time-optimal problems with randomly placed obstacles, which takes approximately 16 hours in total. We then expand this dataset to 10 000 trajectories by adding random redundant obstacles in safe regions, using 80 % of them when training the diffusion model.

During training, we only use the pose information - namely y and yaw angle ϕ - which are transformed into a local Frenet coordinate system [Crenshaw and Edelstein-Keshet, 1993]. This yields the local variables \hat{y} and $\hat{\phi}$, representing the lateral displacements and heading relative to the reference path. Together with the obstacle representation, this results in a flattened map as shown in Fig. 4c. In

this map, the presence of obstacles naturally induces an obstacle-free region, denoted by \mathcal{C} , which is already defined in the local Frenet frame. For notational simplicity, we omit the explicit time index τ , but we emphasize that \mathcal{C} is inherently time-varying, reflecting the dynamic nature of the environment. The set \mathcal{C} provides a time-varying constraint in the planning process and is considered in the definition of our constraint-aware barrier function.

By performing this transformation, we deliberately discard information about the global curvature of the track. This enhances the generalization capability of the trained diffusion model, enabling the model to handle arbitrary (even moving) obstacles. However, this also means that the generated trajectories may not inherently account for curvature constraints, an issue we address using a barrier function in the denoising process, which is detailed in Sec. 4.2.

D.2 Diffusion model architecture

As illustrated in Fig. 5, we adopt a time-conditioned U-Net architecture as the backbone of our diffusion model [Rombach et al., 2022]. The network follows a classic encoder-decoder structure, augmented with time and conditional information to support trajectory generation in dynamic environments.




Figure 5: Architecture of the proposed time-conditioned score-based generative model. The U-Net backbone extracts multi-scale features through a sequence of convolutional and deconvolutional layers, with temporal embeddings injected via dense layers. Spatial transformer modules enable conditional attention guided by task-specific context. Skip connections ensure spatial consistency across scales.

The input is a single-channel spatial-temporal representation of the trajectory, and the output preserves the same spatial resolution. Temporal conditioning is achieved via Gaussian Fourier features [Hensman et al., 2018], which embed the diffusion time step into a high-dimensional representation. This embedding is injected at every resolution level to inform the network of the denoising progress.

The encoder consists of a sequence of down-sampling convolutional blocks, each followed by time embedding fusion and group normalization. To enhance spatial reasoning and enable conditional generation, spatial transformer modules are inserted at deeper layers, where they incorporate context information - such as a reference track - encoded via a lightweight convolutional neural network.

The decoder mirrors the encoder with up-sampling blocks and skip connections, allowing the network to reconstruct high-resolution outputs by fusing low-level and high-level features. Each decoding layer is also conditioned on time to ensure consistency with the diffusion process.

This architecture is designed to be data-efficient, modular, and generalizable. It supports plug-and-play conditional guidance and is easily extendable to other tasks in robotics beyond the case study of autonomous racing.

D.3 Training of the network

We experiment with different input modalities for the diffusion model. Specifically, we considered: (i) the lateral displacement \hat{y} after transforming into the Frenet coordinate system; (ii) both the lateral

displacement \hat{y} and the yaw angle $\hat{\phi}$ in the Frenet frame; (iii) the states including \hat{x} , \hat{y} , $\hat{\phi}$ along with their corresponding velocities \hat{v}_x , \hat{v}_y , $\hat{\omega}$ in the Frenet frame. For each input configuration, we train the model for 500 epochs and explore different values of r_1 and r_0 in constructing the noise schedule $\beta(t) = r_1 t^2 + r_0$ for $t \in [0, 1]$. The training results are shown in Fig. 6, where the three plots from left to right correspond to the aforementioned three input configurations, respectively.




Figure 6: Training performance of the diffusion model under different input configurations and noise schedules. From left to right, the three plots correspond to using (i) lateral displacement \hat{y} in the Frenet frame only, (ii) lateral displacement \hat{y} and yaw angle $\hat{\phi}$ in the Frenet frame, and (iii) the states \hat{x} , \hat{y} , $\hat{\phi}$ along with velocities \hat{v}_x , \hat{v}_y , $\hat{\omega}$ as model inputs. Each setting was trained for 500 epochs while varying the parameters r_1 and r_0 in the noise schedule.

Our experiments show that varying r_1 has negligible impact on the final training performance. In contrast, increasing r_0 generally improves training outcomes, suggesting that larger initial noise levels may facilitate better learning. However, due to the limited size of our training dataset, excessively large values of r_0 can lead to overfitting risks. Additionally, we observe that as the input dimensionality increases, the training performance degrades, likely due to the increased complexity of the data distribution and the limited model capacity under fixed training resources. Based on these observations, we choose to use only the lateral displacement \hat{y} and the yaw angle $\hat{\phi}$ in the Frenet frame as inputs in our final framework, setting $r_1 = 100.0$ and $r_0 = 30.0$.

It is important to emphasize that although we adopt a simplified input representation in this work, our approach remains general and can naturally extend to handle higher-dimensional or multimodal inputs. This flexibility paves the way toward directly modeling control inputs using diffusion models in future work.

E Differentiable Approximation of the Indicator

To facilitate reproducibility, we detail a specific approximation of the indicator in (7), while noting that alternative formulations are possible. We approximate the indicator $\mathbb{1}_{\{\hat{y}_k \notin \mathcal{C}_k\}}$ with a *piecewise-linear function* defined on the signed distance to the nearest obstacle center. Let $\text{dist}(\hat{y}_k)$ be the signed distance that is *positive inside obstacles*, zero on the boundary, and negative outside. We use the normalized linear map

$$\hat{\mathbb{1}}_{\{\hat{y}_k \notin \mathcal{C}_k\}} = \min \left(1, \max \left(0, \frac{\text{dist}(\hat{y}_k)}{\rho(\mathcal{C}_k)} \right) \right)$$

so that points on the obstacle edge map to zero and (approximately) the obstacle center maps to one. Here ρ is a normalization scale corresponding to the obstacle half-width (for disks, the radius; for general shapes, the inradius or a fixed calibration constant). In practice, we compute $\text{dist}(\cdot)$ from the obstacle *binary mask* using a (Euclidean) distance transform; outside obstacles $\text{dist}(\cdot) \leq 0$ hence $\hat{\mathbb{1}} = 0$, while inside obstacles it increases linearly with the interior distance and saturates at one.

F Warm-start evaluation

In this work, we incorporate a warm-starting strategy to accelerate the sampling process, thereby enabling real-time obstacle avoidance. This section presents a quantitative analysis of the effects introduced by this partial diffusion strategy on trajectory generation performance.

Fig. 7 illustrates the reference trajectories generated with and without the application of the warm start technique under an identical obstacle configuration, sampled at consistent time instances. In the figures, gray circles denote static obstacles, while black circles denote dynamic obstacles. The nine subfigures are arranged sequentially from left to right and top to bottom. In each subfigure, the black solid line represents the trajectory obtained using the standard diffusion model, which initiates from standard Gaussian noise and progresses through 500 denoising steps. In contrast, the red solid line corresponds to the trajectory generated with the warm start method, which undergoes 50 denoising steps of partial noised initial trajectory.

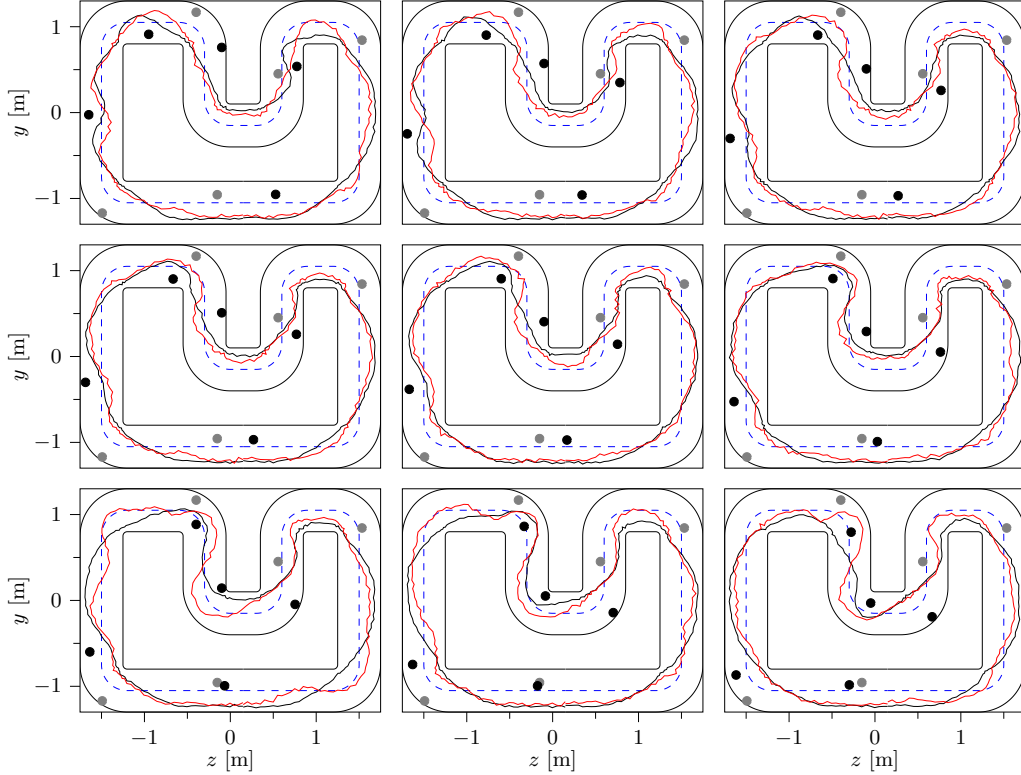



Figure 7: Comparison of reference trajectories generated with and without the warm start technique under an identical obstacle configuration. Gray circles denote static obstacles, and black circles denote static obstacles. The black solid lines represent trajectories produced by the standard diffusion model after 500 denoising steps starting from standard Gaussian noise. The red solid lines represent trajectories generated using the warm start approach, where 50 denoising steps are performed. The warm start method accelerates the sampling process while maintaining successful obstacle avoidance, albeit with slightly coarser trajectory profiles and more conservative motion planning behavior.

As evidenced by the results, both approaches successfully achieve obstacle avoidance at all time steps, demonstrating their respective effectiveness. Nevertheless, the trajectories generated via the warm start technique exhibit a coarser structure, primarily due to the incomplete denoising process inherent to partial diffusion. Furthermore, from the perspective of physical feasibility, the trajectories derived from the standard diffusion model better adhere to realistic vehicle dynamics. Specifically, the final subfigure demonstrates that the warm start method tends to converge to a local solution and favors a more conservative path - remaining closer to the previous time point - to avoid obstacle.


Despite this conservatism, the warm start approach proves crucial, as it reduces the sampling time by approximately a factor of three, thereby making real-time obstacle avoidance feasible. Moreover, the conservative behavior introduced by warm start contributes positively to the overall system stability.

G Experimental platform and the CoDiG framework

Fig. 8 illustrates the experimental platform used to evaluate the performance of the CoDiG framework for real-time obstacle avoidance in autonomous racing. The platform consists of a down-scaled race track (Fig. 8a), a custom-built autonomous car (Fig.8b), and a motion capture system (not shown in the figure). This setup enables agile maneuvering and real-time control in dynamic, safety-critical scenarios such as obstacle avoidance. It provides a reproducible environment to evaluate our approach under realistic conditions.



(a) Down-scaled race track.



(b) Custom-built autonomous car.



(c) Obstacle configuration.

Figure 8: Experimental platform used to evaluate the performance of the CoDiG framework for real-time obstacle avoidance in autonomous racing. The setup includes (a) a down-scaled race track, (b) a custom-built autonomous vehicle, and (c) an obstacle configuration that simulates a challenging and realistic racing scenario.

Additionally, Fig. 8c depicts the obstacle configuration used during the experiments. The vehicle positioned at the starting line is the one under our control, responsible for executing the obstacle avoidance task. Yellow boxes represent static obstacles, while the remaining vehicles serve as either dynamic or static obstacles. This setup faithfully simulates a complex and challenging racing

environment, emphasizing the effectiveness and robustness of our framework under realistic, high-difficulty conditions.

The flowchart illustrating how the CoDiG framework enables real-time obstacle avoidance for autonomous racing on the experimental platform is shown in Fig. 9. The core component of the CoDiG framework is a trained diffusion planner module, which generates a safe reference trajectory y_{ref} capable of avoiding all obstacles. This is achieved by incorporating map and obstacle information, and guiding the sampling process via gradients provided by a constraint-aware guidance mechanism.

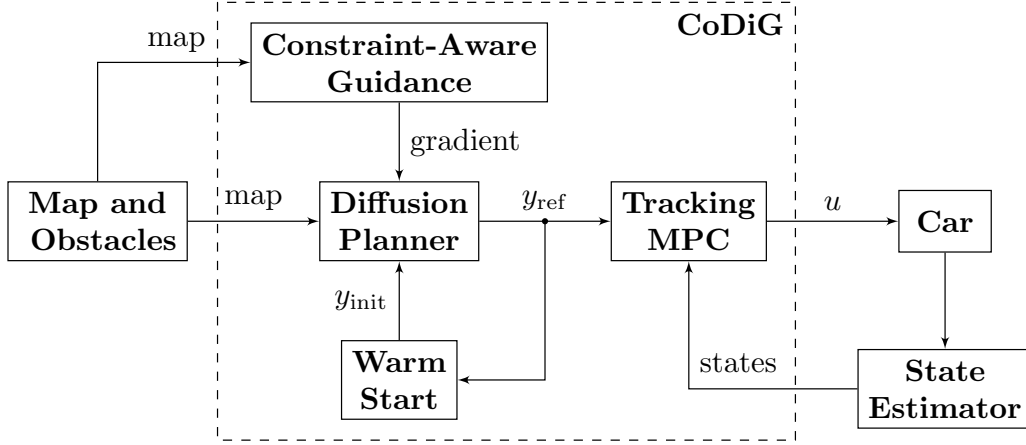


Figure 9: Flowchart of the proposed CoDiG framework for real-time obstacle avoidance in autonomous racing. The framework integrates a diffusion-based trajectory planner, a constraint-aware guidance module that guides the denoising process, a warm start strategy to accelerate sampling, and a tracking MPC controller. All modules operate within the experimental platform described in Fig. 8.

To improve sampling efficiency, the reference trajectory generated at the current time point is further used to construct the initial input y_{init} for the diffusion process at the next time step, via a warm start strategy. This replaces the conventional use of standard Gaussian noise as the initial condition, thereby accelerating the trajectory generation process.

Subsequently, a tracking MPC module computes the control input u required to follow the reference trajectory y_{ref} , based on the current vehicle state estimated by a state estimator module. Finally, the control input u is applied to the vehicle to execute real-time obstacle avoidance.