

SaferPath: Hierarchical Visual Navigation with Learned Guidance and Safety-Constrained Control

Lingjie Zhang, Zeyu Jiang, Changhao Chen*

Abstract—Visual navigation is a core capability for mobile robots, yet end-to-end learning-based methods often struggle with generalization and safety in unseen, cluttered, or narrow environments. These limitations are especially pronounced in dense indoor settings, where collisions are likely and end-to-end models frequently fail. To address this, we propose SaferPath, a hierarchical visual navigation framework that leverages learned guidance from existing end-to-end models and refines it through a safety-constrained optimization-control module. SaferPath transforms visual observations into a traversable-area map and refines guidance trajectories using Model Predictive Stein Variational Evolution Strategy (MP-SVES), efficiently generating safe trajectories in only a few iterations. The refined trajectories are tracked by an MPC controller, ensuring robust navigation in complex environments. Extensive experiments in scenarios with unseen obstacles, dense unstructured spaces, and narrow corridors demonstrate that SaferPath consistently improves success rates and reduces collisions, outperforming representative baselines such as ViNT and NoMaD, and enabling safe navigation in challenging real-world settings.

I. INTRODUCTION

Navigation is a fundamental capability of mobile robots, enabling autonomous operation and the execution of complex tasks in dynamic and uncertain environments. Due to limited prior knowledge and the inherent unpredictability of the real world, enhancing navigation in previously unseen environments remains one of the central challenges in robotics [1], [2]. Traditional navigation approaches typically rely on pre-built maps (e.g., occupancy grids [3] or topological graphs [4]) to plan paths. These methods often require complex sensor fusion, such as integrating LiDAR or depth cameras, and planning algorithms such as A* [5] or RRT [6].

With the rapid advancement of embodied AI, end-to-end visual navigation has attracted growing attention. These methods leverage deep neural networks (DNNs) to directly map visual inputs to actions, reducing reliance on explicit environmental modeling and planning [1]. However, their performance depends heavily on large-scale, diverse training datasets for robust generalization. Collecting and annotating such datasets is costly, and distribution mismatches between training and deployment environments can lead to significant performance degradation. Furthermore, real-world scenarios

Lingjie Zhang, Zeyu Jiang and Changhao Chen are with PEAK-Lab, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, 511453, China (zljacob@gmail.com, zjiang122@connect.hkust-gz.edu.cn, changhaochen@hkust-gz.edu.cn)

This work was supported by National Natural Science Foundation of China (NSFC) under the Grant Number 62573370 and Key Area Project of Education Department of Guangdong Province (No. 2025ZDZX3051).

* Corresponding Author.

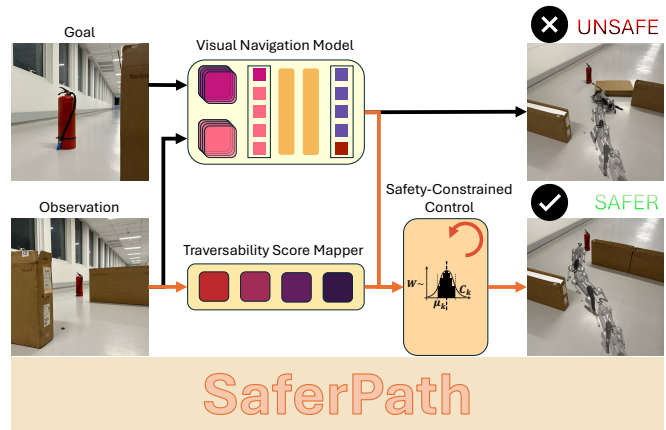


Fig. 1: SaferPath is a hierarchical navigation framework that uses an end-to-end visual navigation model for guidance and integrates safety-constrained control module, significantly improving obstacle avoidance and narrow-passage traversal for safe navigation in challenging environments.

often include unseen obstacles, unstructured layouts, and narrow passages—for example, indoor navigation for domestic robots—where collisions are likely and safety is critical. Diffusion-based navigation models [7] are particularly prone to performance degradation in these settings, as confirmed in our experiments.

This raises two fundamental questions:

- Q1.** How can the safety and robustness of end-to-end visual navigation models be improved without retraining when deployed in diverse and dynamic environments?
- Q2.** Since training data often lack dense and narrow scenarios, can the proposed approach enhance navigation performance and safety in such challenging environments?

To address these challenges, we propose SaferPath, a novel hierarchical visual navigation framework with learned guidance and safety-constrained control. SaferPath is designed as a flexible modular layer above end-to-end navigation models. Instead of directly executing the trajectories predicted by an end-to-end model, we use them only as guidance trajectories. Specifically, visual observations are first transformed into a traversable-area map. The learned guidance trajectory is then refined by formulating the problem as a model predictive optimization task. To reduce the high computation requirements of nonlinear optimization, we introduce the Model Predictive Stein Variational Evolution Strategy (MP-SVES), which recasts trajectory refinement as an efficient variational inference process. MP-SVES generates safety-constrained solutions within only a few iterations, thereby avoiding the

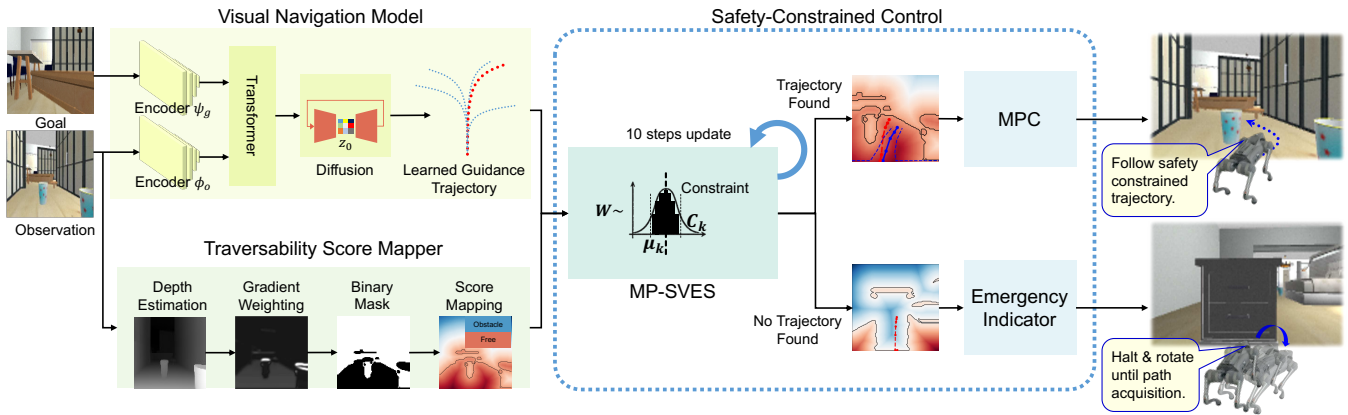


Fig. 2: SaferPath Pipeline overview. SaferPath can be integrated with various end-to-end visual navigation models. The Traversability Score Mapper processes RGB observations to generate a score map \mathcal{S} , which, together with the learned guidance trajectory, is passed to the MP-SVES module. If optimization fails, the Emergency Indicator stops the robot and rotates it until a traversable trajectory is found; otherwise, the MPC controller executes trajectory tracking.

heavy computational cost of conventional nonlinear solvers. The refined, safety-constrained trajectory is subsequently tracked by an MPC controller, ensuring robust and precise navigation in complex environments.

Across diverse scenarios, SaferPath consistently outperforms strong baselines, maintaining high success rates and low collision counts even in cluttered or narrow environments. Compared to the representative visual navigation model NoMaD [7], SaferPath improves performance by 40% in robust navigation under unseen obstacles and over 50% in exploration within dense unstructured environments, while successfully navigating narrow corridors where all baselines fail.

In summary, our main contributions are as follows.

- We propose SaferPath, a hierarchical visual navigation framework that flexibly builds on end-to-end models by leveraging learned guidance while incorporating safety-constrained control, thereby enhancing navigation accuracy and robustness in complex unknown environments.
- We introduce Model Predictive Stein Variational Evolution Strategy (MP-SVES), a novel optimization algorithm that efficiently solves nonconvex nonlinear problems and generates feasible, safety-constrained trajectories within only a few iterations.
- We conduct extensive experiments across diverse scenarios—including navigation with unseen obstacles, exploration in dense unstructured environments, and obstacle avoidance in narrow indoor passages—demonstrating that SaferPath significantly outperforms representative baselines such as ViNT and NoMaD.

II. RELATED WORK

A. Embodied Navigation

With the development of embodied intelligence, learning-based methods, in contrast to geometry-based approaches, emphasize the development of reliable navigation strategies by learning from data-driven navigation experiences[8].

Typically, learning-based navigation strategies are formulated as sequence prediction tasks, where current and historical observations are exploited to generate future behaviors under a generative modeling paradigm[9], [10], [11]. Several studies have explored end-to-end navigation frameworks that integrate heterogeneous sensing modalities, including LiDAR, depth perception, and proprioceptive measurements, to enhance long-range global navigation performance [12], [13], [14]. ViNT[15] and other recent works[7], [16], [17] propose purely vision-based, end-to-end, map-free robot navigation approaches, where models directly process sequences of visual observations to infer future motions or navigation actions, enabling autonomous navigation in complex or previously unseen environments.

B. Learning-Enhanced Model Predictive Control

Model Predictive Control (MPC) uses known system dynamics to optimize control inputs for predicting and regulating future states. Its strengths in real-time optimization and constraint handling make it ideal for motion planning in embodied intelligent robots[18]. However, MPC is less effective for long-horizon predictions due to increasing computational complexity and accumulation of model inaccuracies over time.

Recent works have integrated model-based control approaches with generative models, combining the flexibility of these models with the real-time performance and robustness of MPC to enhance control accuracy and efficiency in robotic systems. Huang et al. introduced the Subgoal Diffuser, a diffusion-based framework that generates adaptive, coarse-to-fine subgoals to guide MPC, enabling robust planning for long-horizon robotic manipulation tasks[19]. In SICNav-Diffusion, the diffusion model is integrated with the MPC framework to facilitate safe trajectory prediction in crowd navigation environments[20]. Other researchers[21] also try to combine MPC with a generative model. In this work, we integrate end-to-end visual navigation models with MPC-based prediction under safety constraints, addressing the limitations of purely end-to-end approaches and achieving more

reliable and robust navigation in real-world environments.

III. METHOD

Prior end-to-end navigation models often improve performance via sophisticated designs or larger datasets, yet still generalize poorly to unseen environments. To address this, we propose a hierarchical navigation framework where an end-to-end visual navigation model providing learned guidance trajectory based on current and past observations while a separate optimization and control module enforces safety and dynamic constraints. This design reduces reliance on the end-to-end model, requiring it only to provide coarse directional guidance for more reliable navigation. Figure 2 illustrates the overall hierarchical navigation framework, with further details provided in the following sections.

A. Preliminaries

SaferPath builds upon end-to-end visual navigation models. These models take the robot’s current and historical RGB observations as input, optionally conditioned on a target image or target position. The inputs are first processed by a visual encoder to produce latent representations, which are then passed through a downstream neural network to predict the trajectory. Formally, the model can be expressed as:

$$f_\theta : \mathbb{R}^{L \times H \times W \times C} \times \mathcal{G} \rightarrow \mathbf{X}_t, \quad \mathbf{X}_t = f_\theta(\mathbf{I}_t, g), \quad (1)$$

where $\mathbf{I}_t = \{I_{t-L+1}, \dots, I_t\}$, $\mathbf{I}_t \in \mathbb{R}^{L \times H \times W \times C}$ denotes the current and historical observations, where L is the history length, H and W denote the image height and width, and C is the number of channels. $g \in \mathcal{G}$ is the goal image, which is typically one of a sequence of images organized as a topological map along the navigation path. $\mathbf{X}_t = \{\mathbf{x}_0, \dots, \mathbf{x}_{t+T}\}$ is the predicted trajectory with $\mathbf{X}_t \in \mathbb{R}^d$ representing the robot’s position trajectory over time.

For instance, we use NoMaD [7] as a base end-to-end model. NoMaD employs an EfficientNet-B0 encoder [22] to encode both the observation sequence \mathbf{I}_t and the goal image g , producing tokenized latent representations z_t . These tokens are then processed by Transformer layers [23] to generate a contextual vector c_t , which a diffusion model [24] uses to predict the trajectory \mathbf{X}_t . Additionally, NoMaD supports a goal-mask input, enabling exploration without an explicit goal.

B. Traversability Score Mapper

The Traversability Score Mapper (TSM) module constructs a safety-aware score map from RGB images to delineate navigable regions for trajectory planning. First, depth estimation is performed using the DepthAnythingV2 model [25], producing a normalized depth map in the range $[0, 1]$. Gaussian smoothing reduces noise, followed by vertical gradient computation via a Sobel operator to highlight depth discontinuities associated with obstacles. To emphasize regions near the robot, a spatially varying weight that increases from top to bottom of the image is applied. The weighted gradient is thresholded to generate a preliminary binary map, labeling obstacles as 0 and traversable areas

Algorithm 1: Trajectory Optimization via MP-SVES

Input: Learned guidance trajectory $\mathbf{X}_{\text{ref},t}$, score map \mathcal{S}

Output: Safety constrained trajectory $\mathbf{X}_{\text{opt},t}$

1 Parameters: Particle number P , sample number m , iteration number N , kernel weight γ , step size η

2 Initialize μ_p, \mathbf{C}_p for $p = 1, \dots, P$

3 for each timestep t do

4 for $n = 1$ to N do

5 for $p = 1$ to P do

6 $\epsilon_{p,i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, for $i = 1, \dots, m$

7 $\mathbf{W}_{p,i} \leftarrow \mu_p + \mathbf{L}_p \epsilon_{p,i}$, $\mathbf{C}_p = \mathbf{L}_p \mathbf{L}_p^\top$

8 $J_{p,i} \leftarrow J(\mathbf{W}_{p,i}, \mathcal{S}, \mathbf{X}_{\text{ref}})$

9 $\omega_{p,i} \leftarrow \frac{\exp(-J_{p,i})}{\sum_{j=1}^m \exp(-J_{p,j})}$

10 $\Delta \bar{\mathbf{W}}_p \leftarrow \sum_{i=1}^m \omega_{p,i} (\mathbf{W}_{p,i} - \mu_p)$

11 $\mathbf{g}_p^{\text{ker}} \leftarrow \sum_{q=1, q \neq p}^P \nabla_{\mu_p} \text{ker}(\mu_p, \mu_q)$

12 $\mu_p \leftarrow \mu_p + \eta \cdot (\Delta \bar{\mathbf{W}}_p + \gamma \mathbf{g}_p^{\text{ker}})$

13 $\mathbf{C}_p \leftarrow \mathbf{C}_p + \eta \cdot \sum_{i=1}^m \omega_{p,i} \epsilon_{p,i} \epsilon_{p,i}^\top$

14 end

15 end

16 $\mathbf{W}_t^* \leftarrow \arg \min_{\mathbf{W}_{p,i}} J_{p,i}$ s.t. constraints

17 $\mathbf{X}_{\text{opt},t} \leftarrow \mathbf{W}_t^* \cdot \mathbf{X}_{\text{ref},t}$

18 end

as 1, with adjustments near the bottom rows to prevent misclassification.

This binary map is refined using horizontal and vertical gradient analysis to detect boundaries and segment obstacle versus free-space regions. A Euclidean distance transform then computes pixel-wise distances to these regions, yielding a smooth score map. In this map, costs are defined as: positive near obstacles (increasing with distance from boundaries), negative in free-space (decreasing with distance), and zero at boundary transitions. Finally, the map is normalized to $[-1, 1]$, producing a continuous score map \mathcal{S} that encapsulates safe navigation regions for trajectory planning.

C. Safety-Constrained Control

We employ the score map \mathcal{S} from TSM as a safety constraint for trajectory planning. Since \mathcal{S} is nonconvex, nonlinear, and highly variable, solving the optimal control problem directly using classic methods such as sequential quadratic programming (SQP)[26] or interior point methods (IPM)[27] is computationally expensive and unsuitable for real-time operation. To overcome this, we propose **Model Predictive Stein Variational Evolution Strategy (MP-SVES)**, which reformulates trajectory optimization as a model predictive task. MP-SVES combines sampling-based methods [28], [29] with Stein variational strategies [30], [31] to accelerate convergence. The safety constrained trajectory is then executed using an MPC control module.

MP-SVES Optimization: In the first step, the learned guidance trajectory $\mathbf{X}_{\text{ref}} = \left\{ \left[x_{\text{ref}}^{(k)}, y_{\text{ref}}^{(k)} \right]^\top \right\}_{k=0}^{K-1}$ from end-

to-end visual navigation model is passed to the MP-SVES optimization module. The goal of this module is to optimize the learned guidance trajectory by adjusting its parameters while ensuring that the safety constrained trajectory remains as close as possible to learned guidance trajectory, while also satisfying safety constraints defined by the score map \mathcal{S} . The optimization is formulated by applying an anisotropic scaling to the learned guidance trajectory points:

$$\mathbf{X}_{\text{opt}}^{(k)} = \begin{bmatrix} x_{\text{opt}}^{(k)} \\ y_{\text{opt}}^{(k)} \end{bmatrix} = \begin{bmatrix} w_x & 0 \\ 0 & w_y \end{bmatrix} \begin{bmatrix} x_{\text{ref}}^{(k)} \\ y_{\text{ref}}^{(k)} \end{bmatrix}, \quad k = 0, \dots, K-1, \quad (2)$$

where w_x and w_y are optimization parameters controlling the scaling of learned guidance trajectory in the x and y dimensions. This scaling-based approach allows the optimization to preserve key geometric characteristics of the original trajectory, such as general shape and heading, while providing flexibility to adjust its size and improve compliance with safety constraints.

The MP-SVES optimization algorithm formulates trajectory optimization as a model predictive problem, where the goal is to identify the optimal scaling parameters $\mathbf{W} = \text{diag}(w_x, w_y)$ that minimize the following objective function subject to system dynamics and feasibility constraints:

$$\begin{aligned} \min_{\mathbf{W}} J(\mathbf{W}, \mathcal{S}, \mathbf{X}_{\text{ref}}) &= \sum_{k=0}^{K-1} \|\mathbf{X}_{\text{opt}}^{(k)} - \mathbf{X}_{\text{ref}}^{(k)}\|_Q^2 \\ &+ \lambda_s \sum_{k=0}^{K-1} \mathcal{S}(\mathcal{X}_{\text{tube}}^{(k)}) + \lambda_r \|\mathbf{W} - \mathbf{I}\|_2^2 \\ \text{s. t. } \quad \mathbf{X}_{\text{opt, init}} &= \mathbf{X}_{\text{ref}} \\ \mathbf{X}_{\text{opt}}^{(k)} &= \mathbf{W} \cdot \mathbf{X}_{\text{ref}}^{(k)}, \quad k = 0, \dots, K-1 \\ \mathcal{S}(\mathcal{X}_{\text{tube}}^{(k)}) &\leq \delta, \quad k = 0, \dots, K-1 \\ \text{where } \mathcal{X}_{\text{tube}}^{(k)} &= \left\{ \mathbf{X}_{\text{opt}}^{(k)} + \alpha \mathbf{n}^{(k)} \mid \alpha \in \left[\frac{-d}{2}, \frac{d}{2} \right] \right\}, \end{aligned} \quad (3)$$

K is the prediction horizon at each timestep. Q , λ_s , and λ_r are tuning parameters controlling tracking accuracy, safety, and regularization. d denotes the robot width. $\mathbf{n}^{(k)}$ denotes the unit normal vector orthogonal to the trajectory at timestep k .

This formulation seeks scaling parameters that keep safety constrained trajectory close to the learned guidance trajectory while satisfying the safety constraints. The MP-SVES iterative procedure used to update \mathbf{W} and obtain safety constrained trajectory is summarized in Algorithm 1, where $\boldsymbol{\mu}_p$ and \mathbf{C}_p represent the mean and covariance matrix of particle p . The repulsion gradient, $\nabla_{\boldsymbol{\mu}_p} \text{ker}(\boldsymbol{\mu}_p, \boldsymbol{\mu}_q)$, is the gradient of the kernel function, typically the RBF kernel $\text{ker}(\boldsymbol{\mu}_p, \boldsymbol{\mu}_q) = \exp\left(-\frac{\|\boldsymbol{\mu}_p - \boldsymbol{\mu}_q\|^2}{2\sigma^2}\right)$, generating repulsive forces to prevent mode collapse and maintain diversity. \mathbf{W}_t^* denotes the optimal scaling parameters that satisfy safety constraints.

In practice, MP-SVES algorithm requires only about 10 iterations to obtain a feasible safety constrained trajectory, making it significantly more efficient than gradient descent methods that directly optimize over the score map. When

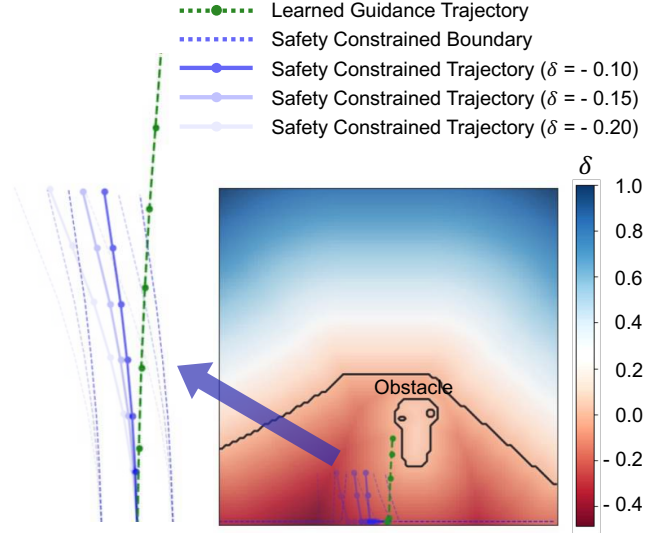


Fig. 3: Safety constrained trajectories under different safety thresholds, with the robot maintaining varying safe distances from obstacles by adjusting threshold δ

facing obstacles, we vary the safety threshold to show how safety constrained trajectories change: Smaller (more negative) thresholds lead to stricter safety constraints, causing safety constrained trajectories to deviate further from obstacles (Fig. 3). In cases where obstacles are too close and no feasible solution can be found, the situation is classified as dangerous. In such cases, the Emergency Indicator commands the robot to stop immediately and rotate toward the direction with a larger free-space region until a feasible safety constrained trajectory is obtained.

Model Predictive Control: Once a safety constrained trajectory has been obtained, it is passed to the MPC module. The purpose of this module is to compute the optimal control inputs $\mathbf{U}_t = (v_t, \omega_t)$ at each time step t , ensuring that the robot's state $\mathbf{X}_t = (x_t, y_t, \theta_t)$ closely follows safety constrained trajectory \mathbf{X}_{opt} , while respecting system dynamics, physical constraints. The MPC problem is formulated as:

$$\begin{aligned} \min_{\mathbf{U}_t^{(0)}, \dots, \mathbf{U}_t^{(K-1)}} & \left\{ \sum_{k=0}^{K-1} \left(\|\mathbf{X}_t^{(k)} - \mathbf{X}_{\text{opt}, t}^{(k)}\|_Q^2 + \|\mathbf{U}_t^{(k)}\|_R^2 \right) \right. \\ & \left. + \underbrace{\|\mathbf{X}_t^{(K)} - \mathbf{X}_{\text{opt}, t}^{(K)}\|_{Q_f}^2}_{\text{terminal cost}} \right\} \end{aligned} \quad (4)$$

$$\begin{aligned} \text{s. t. } \quad \mathbf{X}_t^{(0)} &= \mathbf{X}_{\text{opt}, t}^{(0)}, \\ \mathbf{X}_t^{(k+1)} &= f(\mathbf{X}_t^{(k)}, \mathbf{U}_t^{(k)}), \quad k = 0, \dots, K-1, \\ \mathbf{X}_t^{(k)} &\in \mathcal{X}, \quad \mathbf{U}_t^{(k)} \in \mathcal{U}, \quad k = 0, \dots, K-1, \end{aligned}$$

where $f(\cdot)$ denotes the robot's kinematic model, $\mathbf{U}_t^{(k)} = (v_t^{(k)}, \omega_t^{(k)})$ is the control input at time $t+k$, \mathcal{X} and \mathcal{U} are the feasible state and control input constraint sets.

The objective of the MPC is to penalize deviations from safety constrained trajectory \mathbf{X}_{opt} and excessive control inputs, while ensuring that the trajectory remains within the

dynamic constraints.

The robot’s kinematic model $f(\cdot)$ is defined as:

$$\begin{aligned} x_{t+1} &= x_t + v_t \cos(\theta_t) \Delta t, \\ y_{t+1} &= y_t + v_t \sin(\theta_t) \Delta t, \\ \theta_{t+1} &= \theta_t + \omega_t \Delta t, \end{aligned} \quad (5)$$

where Δt is the time step.

Complete Process The entire process is summarized as follows:

- **MP-SVES Optimization** refines learned guidance trajectory, ensuring it adheres to safety constraints while remaining as close as possible to learned guidance trajectory.
- **MPC Control** computes the optimal control inputs based on safety constrained trajectory, ensuring the robot follows the trajectory while respecting dynamic and physical constraints.

By combining MP-SVES optimization and MPC control, this approach enables robust trajectory tracking in dynamic and narrow environments.

IV. EXPERIMENTS

In this section, we conducted comprehensive experiments to evaluate the effectiveness of our method. Section IV-A describes the experimental setup. Section IV-B evaluates robustness when facing previously unseen obstacles. Section IV-C assesses obstacle avoidance performance in unguided exploration within dense environments. Section IV-D tests navigation in narrow corridors, requiring precise obstacle avoidance and path planning. Section IV-E validates our method in real-world scenarios, demonstrating safe and reliable navigation under practical conditions. Finally, Section IV-F analyzes the contribution of each module through ablation study.

A. Experimental Setup

Baselines: We compare our method against the following pure RGB visual navigation baselines:

- **General Navigation Model (GNM):** A general-purpose visual navigation policy trained on diverse datasets collected across multiple environments and robot platforms [16].
- **Visual Navigation Transformer (ViNT):** A transformer-based foundation model for visual navigation that generates subgoals via diffusion-based planning [15].
- **Navigation with Goal Masked Diffusion (NoMaD):** A diffusion-based navigation policy that handles both goal-directed navigation and exploration in unknown environments [7].

Datasets: To ensure fair comparison, all methods are trained on the same set of navigation datasets:

- **SACSoN/HuRoN:** 70+ hours indoor trajectories with human interactions[32].
- **SCAND:** 9 hours of trajectories from Spot and Jackal robots in diverse social navigation scenarios [33].

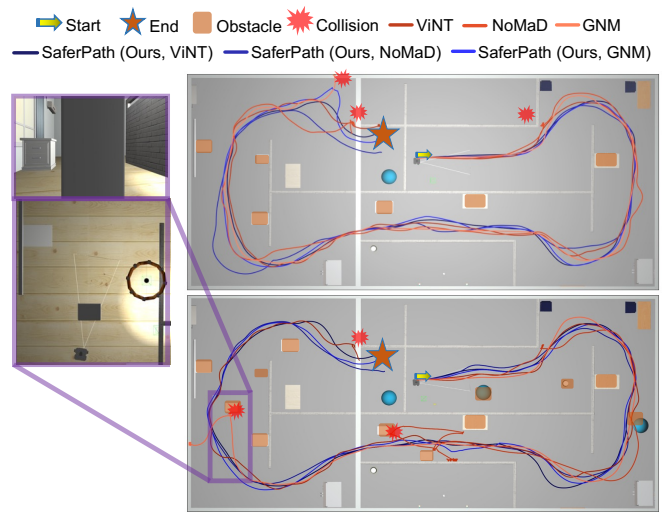


Fig. 4: The figure illustrates the scene from Robust Navigation under Unseen Obstacles experiment. The scene above (Scene A) introduces obstacles along the sides of the road in the original environment, while the scene below (Scene B) includes all the obstacles from Scene A and adds additional obstacles in the middle of the road. The figures display some examples of successful trajectories of our method and examples of collision scenarios of the baseline methods.

- **GoStanford2:** 16 hours of Turtlebot2 trajectories collected across multiple campus buildings [34].
- **RECON:** 18 months of navigation data supporting cross-domain tasks in varied environments [35].

Training: All methods are trained on RGB images and trajectory sequences from the datasets above, using a single NVIDIA RTX 4090 GPU, batch size 256, for 20 epochs.

Evaluation Metrics: We evaluate navigation performance using *success rate* and *average number of collisions per trial*. All simulation experiments and ablation study are repeated 50 times to ensure statistical reliability.

Scenario Configuration: All evaluation scenarios are unseen during training. The robot’s full trajectory is recorded within a fixed time horizon. A trial is considered successful if the robot reaches the designated goal region within this limit; otherwise, it is marked as a failure. Collisions are counted whenever the robot briefly contacts an obstacle or wall before reaching the goal.

B. Robust Navigation under Unseen Obstacles

In this experiment, we evaluate the robustness of our framework when the robot encounters obstacles that are absent from the learned guidance. Specifically, we use guidance generated in the original environment without these obstacles, and during deployment we introduce additional, previously unseen obstacles along the navigation path while keeping the guidance unchanged. This setup contains dynamic environmental changes and tests the model’s ability to handle unforeseen perturbations.

The experimental design is illustrated in Fig. 4, where obstacles are randomly placed along indoor corridors to interfere with the planned navigation trajectory. In Scene A, obstacles are positioned along the sides of the corridors,

TABLE I: Performance comparison of Robust Navigation under Unseen Obstacles experiment. Our method outperforms all baseline methods in the same scene. When more obstacles are introduced into the scene, our method experiences less performance drop compared to the baseline methods.

Method	Scene A		Scene B	
	Success (%)	Coll.	Success (%)	Coll.
GNM	68	0.70	26	1.22
ViNT	94	0.22	80	0.80
NoMaD	56	1.46	20	2.34
SaferPath(Ours, GNM)	82	0.08	74	0.06
SaferPath(Ours, ViNT)	94	0	94	0.08
SaferPath(Ours, NoMaD)	98	0.02	94	0.06

TABLE II: Performance comparison of Exploration in Dense Unstructured Environments experiment. Our method demonstrates significantly better obstacle avoidance capability compared to the baseline method.

Method	Success (%)	Coll.
NoMaD	38	2.66
SaferPath(Ours, NoMaD)	96	0.04

causing minimal occlusion of the robot’s view. In Scene B, additional obstacles are placed in the middle of the corridors, creating significant occlusion and posing greater challenges to navigation. As shown in Fig. 4, these occlusions can temporarily restrict the robot’s field of view, leading to partial or complete loss of guidance information and potentially unsafe trajectory predictions.

Table I summarizes the results. SaferPath(Ours, GNM/ViNT/NoMaD) denotes our framework built on these end-to-end visual navigation models. In Scene A, with only a few additional obstacles, all baseline methods achieve moderate success rates. In Scene B, where occlusions are more severe, baseline performance drops significantly: GNM and NoMaD show larger declines, while ViNT remains relatively more robust. NoMaD also exhibits higher collision counts, likely due to its reliance on goal masks during training, which can lead to unsafe trajectories when guidance is unavailable; in contrast, GNM and ViNT produce more conservative behaviors under such conditions. When our module is integrated, all enhanced versions maintain high success rates and low collision counts across both Scene A and Scene B, with only minor performance reduction under heavy occlusion. These results demonstrate that our module effectively mitigates guidance unavailability, enabling safe and reliable navigation in dynamically changing environments.

C. Exploration in Dense Unstructured Environments

In this experiment, we evaluate the framework’s ability to explore challenging unstructured environments with dense obstacles in the absence of explicit goal guidance. The method leverages its goal mask mechanism to enable exploration without external inputs. A dense obstacle field is constructed, and the robot is initialized within this region (see Fig. 5). The task is considered successful if the robot exits the obstacle field within a predefined time limit. Performance is assessed using two metrics: success rate and average number of collisions.

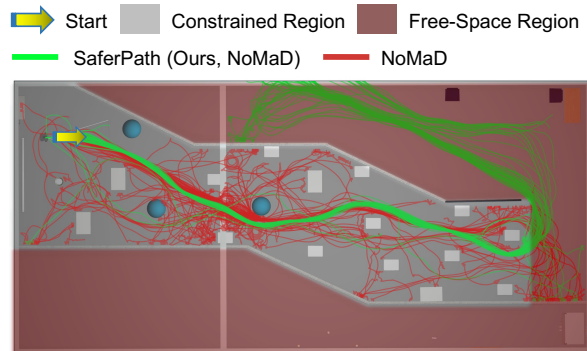


Fig. 5: In Exploration in Dense Unstructured Environments experiment, the trajectory of SaferPath (Ours, NoMaD) (green) effectively avoids dense obstacles and reaches the free-space region. In contrast, the NoMaD trajectory (red) is more chaotic and results in significantly more collisions.

TABLE III: Performance comparison in Obstacle Avoidance in Narrow Corridors experiment. Our method significantly outperforms the baseline methods, with all baseline methods failing in this experiment.

Method	Success (%)	Coll.
GNM	0	4.36
NoMaD	0	4.10
ViNT	0	4.04
SaferPath(Ours, GNM)	76	0.42
SaferPath(Ours, NoMaD)	54	0.52
SaferPath(Ours, ViNT)	80	0.16

Table II reports the results. In dense obstacle fields without guidance, our method demonstrates effective obstacle avoidance and reliable navigation. Compared to NoMaD, it improves the success rate by over 50% while also significantly reducing collisions, underscoring its robustness and safety in unguided exploration scenarios.

D. Obstacle Avoidance in Narrow Corridors

In this experiment, we evaluate navigation performance in narrow and unstructured indoor environments (see Fig. 7). This setting imposes stringent requirements on accurate obstacle avoidance and precise path planning through confined corridors, where even small deviations can result in wall collisions. The irregular layout of the passageways further challenges the real-time performance and robustness of different methods.

Table III summarizes the results. All baseline methods fail in this demanding navigation task, unable to maintain safe trajectories in such constrained spaces. In contrast, the three enhanced versions equipped with our module achieve significantly higher success rates and lower collision counts. Taking ViNT as an example, we visualize the trajectories from all trials for both the original and enhanced versions. The results clearly show that our module enables the robot to remain within safe passage areas, while the original ViNT frequently produces unsafe trajectories that lead to collisions and task failure.

E. Real-World Navigation on a Quadruped Robot

We conducted real-world experiments to validate the effectiveness of our method in practical scenarios. The exper-

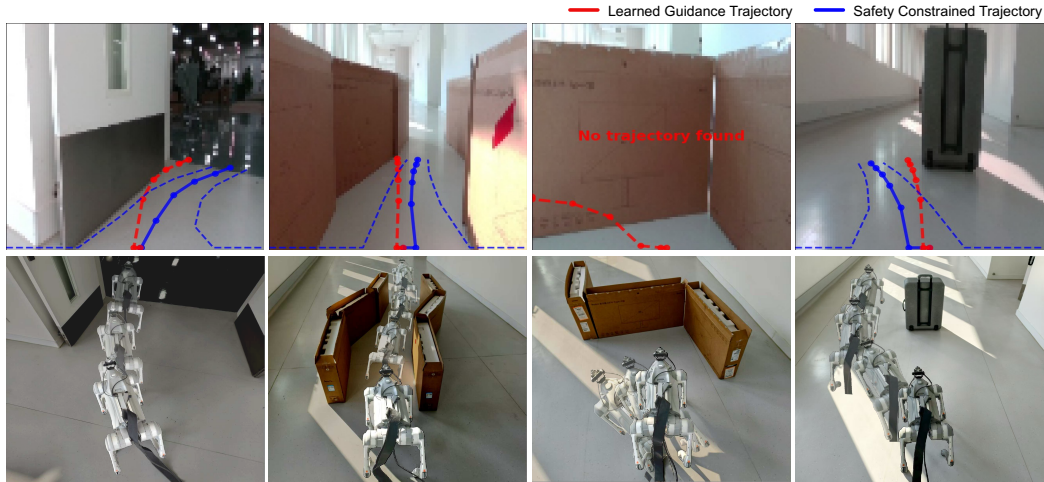


Fig. 6: Our SaferPath deployed on the Unitree Go2 successfully avoids both unseen and dynamic obstacles, navigates narrow spaces, and consistently remains within safe areas. When no traversable trajectory is available, the Emergency Indicator stops the robot and reorients it toward a safe direction.




Fig. 7: The figure illustrates partial trajectories from Obstacle Avoidance in Narrow Corridors experiment, where the trajectory of SaferPath (Ours, ViNT) (green) is clearly smoother and safer, with significantly fewer collisions compared to ViNT (red). This demonstrates the superior obstacle avoidance capability and robustness of SaferPath in complex environments over ViNT.

iments were performed on a Unitree Go2 quadruped robot equipped with an Intel D435i camera, using only its RGB functionality for perception. For navigation, we deployed SaferPath (ours) alongside NoMaD for comparison. Figure 6 shows representative scenarios where four baseline methods are prone to collisions. In these cases, the reference trajectories often bring the robot dangerously close to obstacles. By contrast, after integrating SaferPath, the robot is able to (i) maintain a safe distance from side walls, (ii) navigate narrow passages without collisions, (iii) react promptly to emergency situations by stopping and rerouting to a feasible path, and (iv) avoid obstacles positioned in the center of its trajectory. These results demonstrate that SaferPath enables robust and safe navigation in real-world environments, effectively handling narrow passages and dynamic obstacles.

TABLE IV: The ablation study reveals that all modules are essential, with MP-SVES being the most critical for system stability and performance.

Method	Success (%)	Coll.
SaferPath(Ours w/o TSM, ViNT)	60	1.34
SaferPath(Ours w/o MP-SVES, ViNT)	0	2.28
SaferPath(Ours w/o Em, ViNT)	70	1.14
SaferPath(Ours, ViNT)	80	0.16

F. Ablation Study

To analyze the contribution of individual components, we conduct an ablation study by selectively removing the Traversability Score Mapper (TSM), MP-SVES, and Emergency Indicator modules:

- *w/o TSM*: Remove the Traversability Score Mapper, eliminating the score-map safety constraint.
- *w/o MP-SVES*: Remove the MP-SVES module, placing all constraints into the MPC module and solving with a conventional nonlinear solver IPOPT[36].
- *w/o Em*: Remove the Emergency Indicator module, such that when a safety-constrained trajectory is infeasible, the learned guidance trajectory is executed directly.

We repeat the obstacle avoidance in complex environments experiment using ViNT as the underlying end-to-end visual navigation model. Table IV summarizes the results. Removing any module degrades performance. In particular, eliminating MP-SVES causes the success rate to drop to 0%, underscoring its critical role in ensuring system stability. Without the Traversability Score Mapper, the success rate decreases by 20% and collisions increase significantly, highlighting its importance for obstacle avoidance. Removing the Emergency Indicator module also reduces success rates and increases collisions, showing its necessity in extreme cases.

These results directly address the research questions posed in the Introduction. For **Q1**, the proposed module significantly improves the safety and robustness of end-to-end visual navigation: in environments with additional obstacles, augmented models maintain high success rates and low colli-

sion counts, demonstrating that feasible and safe trajectories are preserved. For **Q2**, in narrow and cluttered environments that are underrepresented in training data, the module enables efficient and collision-free navigation, yielding substantial performance gains. Consistent results across both simulation and real-world experiments confirm the effectiveness and generality of our approach.

V. CONCLUSION

We propose SaferPath, a hierarchical visual navigation framework that leverages learned guidance from end-to-end visual navigation models while ensuring safe and feasible trajectory execution. By effectively utilizing the directional information provided by these models, SaferPath employs the MP-SVES algorithm to optimize trajectories, generating safety constrained paths in just a few iterations while handling both safe zone navigation and emergency obstacle avoidance. Extensive experiments demonstrate that our approach outperforms state-of-the-art end-to-end visual models by over 40%, and is capable of safe navigation even in complex and narrow environments.

Despite these improvements, the method is limited by the restricted field of view of RGB observations, which may lead to collisions when obstacles are outside the camera's view. Future work will focus on incorporating richer sensory modalities and predictive scene understanding to further improve robustness in cluttered and dynamic settings.

REFERENCES

- [1] Y. Liu, L. Liu, Y. Zheng, Y. Liu, F. Dang, N. Li, and K. Ma, "Embodied navigation," *Science China Information Sciences*, vol. 68, no. 4, p. 141101, 2025.
- [2] T. Zhang, X. Hu, J. Xiao, and G. Zhang, "A survey of visual navigation: From geometry to embodied ai," *Engineering Applications of Artificial Intelligence*, vol. 114, p. 105036, 2022.
- [3] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, 1989.
- [4] S. Thrun and A. Bücken, "Integrating grid-based and topological maps for mobile robot navigation," in *Proceedings of the national conference on artificial intelligence*, 1996, pp. 944–951.
- [5] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [6] S. M. LaValle, "Rapidly-exploring random trees : a new tool for path planning," *The annual research report*, 1998.
- [7] A. Sridhar, D. Shah, C. Glossop, and S. Levine, "NoMaD: Goal Masked Diffusion Policies for Navigation and Exploration," *arXiv preprint*, 2023.
- [8] H. Liu, D. Guo, and A. Cangelosi, "Embodied intelligence: A synergy of morphology, action, perception and learning," *ACM Computing Surveys*, Feb. 2025.
- [9] W. Zheng, R. Song, X. Guo, C. Zhang, and L. Chen, "Genad: Generative end-to-end autonomous driving," *arXiv preprint arXiv:2402.11502*, 2024.
- [10] Y. Hu, J. Yang, L. Chen, K. Li, C. Sima, X. Zhu, S. Chai, S. Du, T. Lin, W. Wang, L. Lu, X. Jia, Q. Liu, J. Dai, Y. Qiao, and H. Li, "Planning-oriented autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [11] G. Kahn, P. Abbeel, and S. Levine, "Badgr: An autonomous self-supervised learning-based navigation system," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.
- [12] J. Liang, P. Gao, X. Xiao, A. J. Sathiamoorthy, M. Elnoor, M. C. Lin, and D. Manocha, "Mtg: Mapless trajectory generator with traversability coverage for outdoor navigation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 2396–2402.
- [13] J. Liang, A. Payandeh, D. Song, X. Xiao, and D. Manocha, "Dtg: Diffusion-based trajectory generation for mapless global navigation," 2024.
- [14] A. Saviolo and G. Loianno, "Nova: Navigation via object-centric visual autonomy for high-speed target tracking in unstructured gps-denied environments," *arXiv preprint arXiv:2506.18689*, 2025.
- [15] D. Shah, A. Sridhar, N. Dashora, K. Stachowicz, K. Black, N. Hirose, and S. Levine, "ViNT: A foundation model for visual navigation," in *7th Annual Conference on Robot Learning*, 2023.
- [16] D. Shah, A. Sridhar, A. Bhorkar, N. Hirose, and S. Levine, "GNM: A General Navigation Model to Drive Any Robot," in *International Conference on Robotics and Automation (ICRA)*, 2023.
- [17] Y. Zeng, H. Ren, S. Wang, J. Huang, and H. Cheng, "Navidiffuser: Cost-guided diffusion model for visual navigation," *arXiv preprint arXiv:2504.10003*, 2025.
- [18] R. Grandia, F. Farshidian, A. Dosovitskiy, R. Ranftl, and M. Hutter, "Frequency-aware model predictive control," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1517–1524, 2019.
- [19] Z. Huang, Y. Lin, F. Yang, and D. Berenson, "Subgoal diffuser: Coarse-to-fine subgoal generation to guide model predictive control for robot manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 16489–16495.
- [20] S. Samavi, A. Lem, F. Sato, S. Chen, Q. Gu, K. Yano, A. P. Schoellig, and F. Shkurti, "Sicnav-diffusion: Safe and interactive crowd navigation with diffusion trajectory predictions," *IEEE Robotics and Automation Letters*, vol. 10, no. 9, pp. 8738–8745, 2025.
- [21] H. Xue, C. Pan, Z. Yi, G. Qu, and G. Shi, "Full-order sampling-based mpc for torque-level locomotion control via diffusion-style annealing," 2024.
- [22] M. Tan and Q. V. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," 2020.
- [23] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023.
- [24] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," 2020.
- [25] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," in *CVPR*, 2024.
- [26] J. F. Bonnans, J. C. Gilbert, C. Lemaréchal, and C. A. Sagastizábal, *Numerical Optimization: Theoretical and Practical Aspects*, second revised edition of translation of 1997 french edition ed., ser. Universitext. Berlin: Springer-Verlag, 2006.
- [27] S. J. Wright, *Primal-Dual Interior-Point Methods*. Society for Industrial and Applied Mathematics, 1997.
- [28] E. Theodorou, J. Buchli, and S. Schaal, "A generalized path integral control approach to reinforcement learning," *Journal of Machine Learning Research*, vol. 11, no. 104, pp. 3137–3181, 2010.
- [29] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1433–1440.
- [30] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," 2019.
- [31] C. V. Braun, R. T. Lange, and M. Toussaint, "Stein variational evolution strategies," *ArXiv*, vol. abs/2410.10390, 2024.
- [32] N. Hirose, D. Shah, A. Sridhar, and S. Levine, "Sacson: Scalable autonomous control for social navigation," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 49–56, 2024.
- [33] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, "Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation," *IEEE Robotics and Automation Letters*, 2022.
- [34] N. Hirose, A. Sadeghian, M. Vázquez, P. Goebel, and S. Savarese, "Gonet: A semi-supervised deep learning approach for traversability estimation," pp. 3044–3051, 2018.
- [35] D. Shah, B. Eysenbach, N. Rhinehart, and S. Levine, "Rapid exploration for open-world navigation with latent goal models," in *5th Annual Conference on Robot Learning*, 2021.
- [36] F. E. Curtis, O. Schenk, and A. Wächter, "An interior-point algorithm for large-scale nonlinear optimization with inexact step computations," *SIAM Journal on Scientific Computing*, vol. 32, no. 6, pp. 3447–3475, 2010.