


# ViPlanner: Visual Semantic Imperative Learning for Local Navigation

Pascal Roth<sup>1</sup>, Julian Nubert<sup>1</sup>, Fan Yang<sup>1</sup>, Mayank Mittal<sup>1,2</sup>, and Marco Hutter<sup>1</sup>



**Fig. 1:** Quadrupedal navigation in large scale urban environments requires semantic understanding to successfully follow side- and crosswalks. Four local planning events (A - D) along the autonomously traversed path are shown. The planned path of the proposed semantic imperative planner is projected into the semantic images (middle row), whereas the estimated path of the purely geometric iPlanner [1] is overlaid onto the depth image (bottom row). The traversed path is shown in an environment reconstruction generated by [2].

**Abstract—**Real-time path planning in outdoor environments still challenges modern robotic systems due to differences in terrain traversability, diverse obstacles, and the necessity for fast decision-making. Established approaches have primarily focused on geometric navigation solutions, which work well for structured geometric obstacles but have limitations regarding the semantic interpretation of different terrain types and their affordances. Moreover, these methods fail to identify traversable geometric occurrences, such as stairs. To overcome these issues, we introduce ViPlanner, a learned local path planning approach that generates local plans based on geometric and semantic information. The system is trained using the Imperative Learning paradigm, for which the network weights are optimized end-to-end based on the planning task objective. This optimization uses a differentiable formulation of a semantic costmap, which enables the planner to distinguish between the traversability of different terrains and accurately identify obstacles. The semantic information is represented in 30 classes using an RGB colorspace that can effectively encode the multiple levels of traversability. We show that the planner can adapt to diverse real-world environments without requiring any real-world training. In fact, the planner is trained purely in simulation, enabling a highly scalable training data generation. Experimental results demonstrate resistance to noise, zero-shot sim-to-real transfer, and a decrease of 38.02% in terms of traversability cost compared to purely geometric-based approaches. Code and models are made publicly available: <https://github.com/leggedrobotics/viplanner>.

This work is supported in part by the Max Planck ETH Center for Learning Systems, the EU Horizon 2020 programme grant agreement No.852044, 10107045 and 101016970, the EU Horizon Europe Framework Programme grant agreement No. 101070405 and 101070596, the NCCR digital fabrication and robotics, and the SNSF project No.188596.

<sup>1</sup>All authors are with the Robotic Systems Lab, ETH Zürich, 8092 Zürich, Switzerland. Contact: {rothpa, nubertj, fanyang1, mmittal, mahutter}@ethz.ch.

<sup>2</sup>The author is with NVIDIA.

## I. INTRODUCTION

Path planning is a fundamental problem in robotics. Significant progress has been made for mobile navigation in environments where pre-built high-definition maps are available [3]. In contrast, several challenges still exist for planning in unknown environments fully relying on onboard sensors due to sensor noise, dynamic or moving objects, and diverse scenes [1]. At present, most works address these challenges using purely geometric navigation solutions [1,4]–[8]. While showing good performance in structured or unpopulated (known) environments, such as indoors [9] or underground environments [10], navigation becomes much harder once robots enter unstructured outdoor environments. On the one hand, upcoming systems must be able to distinguish terrain of varying traversability with the same geometric appearance (e.g., mud vs. concrete), while on the other hand, seemingly geometric obstacles (such as steps or stairs) must be interpreted correctly. Including the semantic domain in the traversability estimation has the potential to improve the assessment in semi-structured environments [11]–[13].

For unknown environments, current path planning algorithms are either designed as end-to-end learned or modular approaches. In the latter, a perception module typically estimates the traversability while the path is generated by a sampling or optimization-based method [4]–[6]. While these methods can generalize well, they require collecting a vast amount of difficult-to-acquire real-world data for the traversability assessment and searching a path in this map, which can introduce large latencies. On the other hand, in end-to-end learned solutions, the path is directly predicted from sensor measurements, which reduces latencies.

These methods are trained either through imitation learning (IL) with expert demonstrations [7,14], with reinforcement learning (RL) [8,15] or very recently via imperative learning (ImpL) [1]. While IL historically suffers from low generalization capabilities due to the limited availability of demonstrations, RL and ImpL can be trained entirely in simulation or with a mix of simulated and real-world data. ImpL employs an offline Bi-Level Optimization (BLO) over a predefined differentiable cost (map) to generate smooth (path) predictions. It enhances training efficiency compared to RL and has been shown to outperform previous methods [1]. Nevertheless, the existing ImpL planner, called iPlanner [1], is restricted to the geometric domain and requires diverse real-world data to be applied safely.

In this work, we present ViPlanner, an end-to-end learned, multi-domain local planner that uses the ImpL paradigm by building up on iPlanner. Our core contributions are:

- 1) The development of a semantically-aware local planner using an unsupervised Imperative Learning approach.
- 2) The achievement of zero-shot transfer from simulation to real-world by combining and integrating a pre-trained semantic segmentation network and geometric input during end-to-end training.
- 3) Evaluations and benchmarks of the proposed method against the geometric-based approach [1] in both simulated and real-world settings using the quadrupedal robot ANYmal [16].
- 4) The released open-source code featuring an efficient, scalable pipeline for data generation and planner evaluation, applicable to indoor and outdoor environments, using high-fidelity simulation [17] based on *NVIDIA Omniverse*.

## II. RELATED WORK

Local path planning has been extensively explored over the past two decades. Traditionally, a modular approach consisting of *i*) a traversability estimation module, and *ii*) a path-searching algorithm is applied [18]. Early works assumed a mostly observable environment with given traversability estimation and tackled the local path planning problem with optimization-based [19], sampling-based [20,21], as well as heuristic/primitive-based methods [22,23]. Over the past few years, the techniques have gotten more advanced and can master more complex environments [5] or robot configurations [24], recently often combining sampling-based and optimization-based methods in one approach [24,25]. However, the conceptual separation in *i*) and *ii*) has mostly remained unchanged. Accelerated by the advent of deep learning, recent advances in literature have explored data-driven approaches, from learning-based traversability assessment [26,27], up to fully end-to-end learned approaches [1,8].

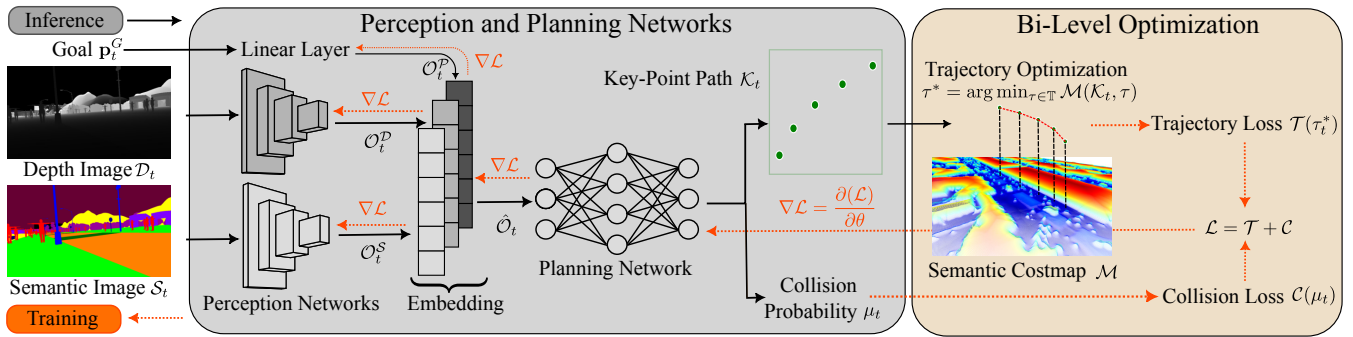
*a) Modular Geometric Approaches:* Classical approaches analyze the environment primarily through geometric information in the form of point clouds [4,28] or meshes [29] and determine traversability based on metrics such as occupancy or stepping difficulty [30]. A sampling or optimization-based path-searching module then determines the

path by employing traditional techniques such as PRM\* [5], visibility graphs [4] or first learning-based approaches [6]. However, only using geometric measurements results in the failure to reason about paths over different terrains with similar geometry or predict paths that could traverse geometric obstacles, like stairs, making the planner inadequate for outdoor environments. Instead, this work focuses on a multi-domain approach for enhanced environmental understanding.

*b) Modular Semantic Approaches:* Semantics as additional domains within the traversability assessment allow for enhanced reasoning about the environment where each semantic class is assigned a traversability cost. By annotating a geometric point cloud with semantic labels, [11,12] created a dense traversability map and enabled autonomous off-road navigation. Moreover, [31]–[33] uses semantic information to improve reactivity and the planner’s safety. The resulting composition of three modules introduces large latencies in the systems and limits the application to mostly static environments. While we also include semantics in our approach, we fuse the domains in the latent space and generate paths end-to-end to minimize latencies. Moreover, our learned solution goes beyond pure reactivity by exposing the planner to different behaviors during training while not relying on a globally build map.

*c) Imitation Learning and Self-Supervised Learning:* Another line of research aims to learn local path planning from expert demonstrations [7,34,35]. For autonomous driving, demonstrations fused with high-definition maps in an end-to-end learned setting allow for applications in environments shared with humans [3,14]. In unstructured environments, the traversability estimation module can learn from demonstrations in a self-supervised manner [36]. In these cases, only weak, platform-dependent supervision is required [26,27]. However, such methods suffer from poor generalization to unfamiliar environments and labor-intensive data generation [37]. Additionally, the optimality is limited to the sub-optimality of the expert. On the contrary, we train our method entirely in simulation by simultaneously optimizing the network weights and the task objective to improve generalization while not being bound by sub-optimality.

*d) RL and Imperative Learning:* While previous approaches often relied on extensive expert knowledge and real-world data, Reinforcement Learning (RL) offers a structured approach to end-to-end path planning learning, performed in simulation [8]. Domain randomization and noised measurements are used to transfer the planner to the real world. However, low sampling efficiency and sparse rewards make training RL models time-intensive, especially for dense input data [38]. Most recently, Yang et al. [1] introduced iPlanner, which uses the Imperative Learning paradigm [39] to address path planning as an offline Bi-Level Optimization. While this concept improves convergence speed compared to RL, iPlanner requires real-world data to bridge the reality gap. Our method builds on top of [1], which demonstrated advantages over previous demonstration, modular, and RL-based methods in semi-structured environments, and overcomes its main limitation by going beyond the geometric domain.



**Fig. 3:** Overview of the integral components of the proposed approach. The perception and planning networks take a depth image, a semantic image, and the desired goal position as input and estimate a coarse plan together with a collision probability. The network weights and final path are jointly optimized as part of a Bi-Level Optimization scheme.

### III. PROBLEM FORMULATION

We define the environment within which the robot operates as  $\mathcal{Q} \subset \mathbb{R}^3$ . Here,  $\mathcal{Q}$ , consists of non-traversable geometric and semantic obstacles represented by the subset  $\mathcal{Q}_{\text{obs}} \subset \mathcal{Q}$ , and the traversable space  $\mathcal{Q}_{\text{trav}} = \mathcal{Q} \setminus \mathcal{Q}_{\text{obs}}$  where the robot is safe to walk. The traversable area  $\mathcal{Q}_{\text{trav}}$  is divided into  $N$  subsets  $\mathcal{Q}_i$  with finite motion costs  $c_i \in [0, \text{inf}), \forall i \in \{1, \dots, N\}$ . Consequently, the robot workspace exhibits a fine-grained non-binary separation, fit for representing complex real-world environments, with the complete traversable space  $\mathcal{Q}_{\text{trav}}$  defined as  $\bigcup_{i=1}^N \mathcal{Q}_i$ . The traversability cost of each path  $\mathcal{P}$  is defined as its cost integral  $\mathcal{T}_{\mathcal{P}}^{\mathcal{T}} = \int_{\mathcal{P}} c(x, y) dp$ , with  $c(x, y) \in \{c_1, \dots, c_N\}$ , depending on the location. The goal cost is defined as the overall length of the path  $\mathcal{T}_{\mathcal{P}}^{\mathcal{G}} = \int_{\mathcal{P}} dp$ .

*Objective:* Navigation tasks are ubiquitous in robotics and are generally described as finding a safe, fast, and collision-free path from a start to a goal position in  $\mathcal{Q}$ . For this work, we are interested in online motion planning using cheap onboard sensing only. In our case, this means that given a depth image observation  $\mathcal{D}_t \in \mathbb{R}^{H_D \times W_D}$ , a semantic image observation  $\mathcal{S}_t \in \mathbb{R}^{H_S \times W_S}$ , and an intended goal position  $\mathbf{p}_t^G \in \mathcal{Q}_{\text{trav}}$  at timestep  $t$  (compare Fig. 3), the final goal is to estimate a trajectory  $\tau_t = \Phi(\mathcal{D}_t, \mathcal{S}_t, \mathbf{p}_t^G, \theta)$  that guides the robot from its current position  $\mathbf{p}_t^R$  to the goal  $\mathbf{p}_t^G$ , while minimizing the combined traversability and goal cost  $\mathcal{T}_{\tau}^{\mathcal{T}} + \mathcal{T}_{\tau}^{\mathcal{G}}$ . Here,  $\Phi$  refers to the neural network approximator with weights  $\theta$ . Moreover, the collision risk with the environment should be minimized to reduce safety hazards and increase reliability. It is important to note that this path planning problem must be solved only with partial observations; the two input image streams.

### IV. METHODOLOGY

The proposed pipeline integrates two stages, as visible in Fig. 3. The first stage consists of the perception and planning networks that encode and concatenate semantic, depth, and goal inputs (Sec. IV-A) and predict a sparse key-point-based path  $\mathcal{K}$  towards the goal, along with the collision confidence probability  $\mu$  of the generated  $\tau$  (Sec. IV-B). The second stage is the BLO process, including the metric-based trajectory optimizer (TO) and network updates. The TO process optimizes the path regarding the semantic costmap

Classes	Cost	Spectrum
sidewalk, crosswalk, floor, stairs	$c_{free}$	Green
gravel, sand, snow	$c_{mid_1}$	Light Green
terrain (grass, dirt)	$c_{mid_2}$	Yellow
road	$c_{mid_3}$	Orange
person, animal, vehicle, trains, motorcycle, bicycle	$c_{obs}$	Red
building, wall, fence, bridge, tunnel, furniture, tree, water surface	$c_{obs}$	Purple
pole, traffic sign, traffic light, bench	$c_{obs}$	Blue
sky, ceiling, unknown	$c_{obs}$	Black

**TABLE I:** RGB-encoded semantic colorspace for navigation. While each class has its color, in this table, multiple similar classes are grouped by spectrum. Each group is associated with a motion cost  $c$ , with  $c_{free}$  the lowest and  $c_{obs}$  the largest.


(Sec. IV-C). The TO formulation is introduced in [1] and will not be discussed in detail in this work for brevity. The task-level cost function (Sec. IV-D) to be minimized by both the embedding and planning networks is denoted as  $\mathcal{L}$ . It consists of trajectory ( $\mathcal{T}$ ) and collision probability ( $\mathcal{C}$ ) costs.

#### A. Semantic Encoding

The presented work utilizes a semantic space of 30 classes typically encountered during navigation challenges. In contrast to other approaches, such as one-hot encoding, our method encodes the traversability directly in RGB colorspace, ensuring that classes with similar traversability are grouped. Generally, traversable areas are more within the green, while obstacles are more within the red and blue colorspace. During training, the deployed neural network then learns to utilize the class characteristics as supplementary information for planning. Table I provides an overview of all classes with their corresponding color.

#### B. Perception and Planning Networks

1) *Perception Networks:* For every time stamp  $t$ , the two perception networks receive a depth and semantic image, respectively. At their core, both networks consist of a ResNet-18 [40] architecture and are trained from scratch without weight sharing. During inference, a separate semantic segmentation network generates  $\mathcal{S}_t$  from raw RGB input. We transform both measurements into the same camera frame and estimate embeddings  $\mathcal{O}^D \in \mathbb{R}^{C_I \times M}$  and  $\mathcal{O}^S \in \mathbb{R}^{C_I \times M}$ .



**Fig. 4:** Example training environment for the urban CARLA dataset [41] with its semantic reconstruction, and created geometric and semantic costmaps.

2) *Combined Feature Embedding:* To retrieve the target efficiently, the commanded goal position  $\mathbf{p}_t^R$  is mapped by a linear layer to a higher dimensional embedding  $\mathcal{O}_t^P \in \mathbb{R}^{C_G \times M}$ , with  $C_G \geq 3$ . We then concatenate the embeddings  $\mathcal{O}_t^P$  and  $\mathcal{O}_t^S$  with this goal position embedding, resulting in the combined embedding  $\hat{\mathcal{O}}_t \in \mathbb{R}^{(2C_I + C_G) \times M}$ . This embedding is essential, as it constitutes the input to the planning network.

3) *Planning Network:* Our lightweight planning network consists of convolutional layers (CNN) and a multilayer perceptron (MLP). Building on [1], it includes two distinct heads: the *path planning-* and the *collision probability* head. The former predicts a sparse set of key points  $\mathcal{K}_t \in \mathbb{R}^{n_k \times 3}$  that are the core input to the trajectory optimization. The latter estimates the risk of collision  $\mu$  with obstacles for each trajectory and acts as a supplementary safety measure. This estimate is necessary, as a naive increase of the loss for obstacle violations leads to overly conservative policies. Instead, the collision head allows for more flexibility, which is crucial in scenarios where the system is trapped in local minima. Only trajectories with a collision probability of less than  $\delta_\mu = 0.5$  are executed during online inference.

### C. Semantic Costmap

The creation of the semantic costmap  $\mathcal{M}$  and its usage during the trajectory and network optimization is one of the core components of this work. Note, however, that this costmap is only needed during training time. During inference, the condensed policy estimates the resulting paths from the incoming image stream.

To obtain the costmap, we create a 2D grid with a set resolution of the size of the environment. Each cell is assigned a class label depending on the mesh at the corresponding location. Fig. 4 shows the semantically annotated environment as a point-cloud reconstruction. Cost factors, as given in Tab. I, are assigned to each cell based on the class. We require the costmap to be differentiable and smooth to enable successful (network) optimization. The necessary smoothing is carried out by first applying a Gaussian filter to remove possible classification errors without affecting small obstacles and second by using a signed-distance gradient value toward the closest class boundaries, to reduce the impact of areas with

constant loss. For the area with the smallest loss, this operation is inverted with a gradient pointing towards the center, guiding the robot to the center of elongated parts of the map (such as hallways). This step is crucial to accelerate the training and push the path towards the lowest loss areas. Moreover, it enables us to effectively train in outdoor environments with large spaces of constant loss values. Lastly, a second Gaussian Filter is applied to smooth the boundaries between areas of different classes. An example of a resulting costmap is shown in Fig. 4.

### D. Training Loss

As introduced in [1], the ImpL training loss  $\mathcal{L}(\tau_t)$  includes a trajectory loss term,  $\mathcal{T}_t(\tau_t)$  and a collision loss term,  $\mathcal{C}_t(\tau_t, \mu_t)$ . We directly adopt the definition of  $\mathcal{C}(\tau, \mu)$ ,

$$\mathcal{C} = \begin{cases} \text{BCELoss}(\mu, 0.0) & \mathbf{p}_i^K \in \mathcal{Q}_{tra} \forall \mathbf{p}_i^K \in \tau \\ \text{BCELoss}(\mu, 1.0) & \text{otherwise,} \end{cases} \quad (1)$$

with the Binary-Cross-Entropy Loss (BCELoss) and the center points  $\mathbf{p}_i^K$  of the path  $\tau$ . Our implementation of  $\mathcal{T}$  differs from the original formulation in multiple ways. The resulting trajectory loss term is formulated as follows:

$$\mathcal{T}(\tau) = \alpha \mathcal{T}^T + \beta \mathcal{T}^G + \gamma \mathcal{T}^M + \delta \mathcal{T}^H. \quad (2)$$

Here,  $\alpha, \beta, \gamma$  and  $\delta$  are for loss scaling, and  $\mathcal{T}^M$  is the same motion loss term as used in [1]. Our work introduces new formulations for the obstacle, now called *traversability* ( $\mathcal{T}^T$ ) and goal ( $\mathcal{T}^G$ ) terms, and extends the formulation with an additional height loss  $\mathcal{T}^H$ .

1) *Traversability Loss:* The new traversability loss takes the physical size of the robot into account by evaluating the cost not only at the center points  $\mathbf{p}_i^K$ , but also at points perpendicular to the path at a distance  $w^R$ , corresponding to the robot width. The resulting loss is

$$\mathcal{T}^T(\tau) = \frac{1}{3n} \sum_{i=1}^n \tilde{m}(\mathbf{p}_i^K) + \tilde{m}(\mathbf{p}_i^K + w^R \cdot \mathbf{n}_i^K) + \tilde{m}(\mathbf{p}_i^K - w^R \cdot \mathbf{n}_i^K), \quad (3)$$

where  $\tilde{m}(\cdot)$  is the bi-linearly interpolated value of the costmap  $\mathcal{M}$  at the corresponding point.

2) *Goal Loss:* For the goal loss, we introduce a log scaling to limit the influence of goal points at large distances. The loss can be expressed as


$$\mathcal{T}^G(\tau) = \log(\|\mathbf{p}_n^K - \mathbf{p}_i^G\|_2 + 1.0), \quad (4)$$

with  $\mathbf{p}_n^K$  denoting the last point of the path.

3) *Height Loss:* The additional, newly introduced height loss regularizes the path to maintain the base height of the robot  $h^R$ , avoiding any maneuvers that might circumvent obstacles by going above or below them. The corresponding loss is given as

$$\mathcal{T}^H(\tau_t) = \frac{1}{n} \sum_{i=1}^n |z_{p,i}^K - \tilde{h}(\mathbf{p}_i^K) - h^R|, \quad (5)$$

where  $\tilde{h}(\cdot)$  is the bilinearly interpolated value of the heightmap  $\mathcal{H}$  at the given point.



**Fig. 5:** Qualitative comparison between the proposed ViPlanner (orange) and the purely geometric iPlanner [1] (blue). Five trials of both planners are shown for the manually selected waypoints (red). The street (a) and the yellow working areas (b) are successfully avoided.

## V. IMPLEMENTATION

*Simulation Environments:* The proposed planner has been fully trained with *NVIDIA Omniverse* and evaluated using the legged robot ANYmal [16] and a pre-trained RL locomotion policy provided in the *Orbit* Framework [17]. To allow for successful navigation in semi-structured environments, we picked realistic indoor scenes from the Matterport3D dataset [42] and relevant outdoor scenes as released in CARLA [41], compare Fig. 5. We developed new plugins to load both datasets into *Omniverse*, to benefit from its scalability and realistic physics engine. This allows us to simulate the full robot model, including physical contacts, instead of a perfect point approximation for evaluation, as done in [1]. All code, including the dataset processing and necessary plugins, are made publicly available and will help develop and test future applications.

*Training Set Creation:* For the successful deployment of our method and scaling to larger amounts of data, a flexible and automatic sampling procedure for capturing diverse (semantic) image viewpoints within the environments is required. In contrast to the works in [1,43], where a similar problem has been solved through random-pose sampling around manually defined paths, in this work, the sensor model is spawned fully automatically in simulation according to the so-called *Halton* sequence [44] at robot-accessible locations. The corresponding images at each sampled point are then rendered at angles that maximize the coverage of the traversable space. Moreover, goal points are placed randomly at the sampled camera centers from before. Their reachability is identified by constructing a graph between all centers and removing connections that pass through areas on the costmap  $\mathcal{M}$  higher than a certain threshold. For successful obstacle avoidance learning, the goal should preferably be within the robot’s field of view (FoV). Thus, our data generation pipeline allows us to define the desired ratio of the relative goal point locations w.r.t. the FoV. Overall, we generated approx. 80k start-goal pairs from eleven Matterport (avg.  $36 \times 33\text{m}$ ), one CARLA ( $400 \times 400\text{m}$ ) and three warehouse (avg.  $25 \times 35\text{m}$ ) environments.

## VI. EXPERIMENTS

*Model Training:* The training process is managed using the SGD optimizer, a learning rate scheduler, and an early-stopping strategy. The training procedure is executed on an *NVIDIA RTX 3090* for around six hours. The semantic cost factors, compare Tab. I, are in the range of zero to two.


*Experiment Setup:* We demonstrate the effectiveness of our model, data generation, and training strategy by comparing it in various simulation environments to [1]. In addition, we employ our method on the legged robot ANYmal [16] to showcase the zero-shot transfer to the real world. The planner runs on a *Nvidia Jetson Orin AGX* on the robot and uses the latest semantic image generated by the state-of-the-art *Mask2Former* segmentation model [45] and the current depth image as input. The semantic segmentation model has been fine-tuned on a small set of images collected in Zurich, Switzerland, where we conducted our outdoor experiments. By asynchronously reusing the semantic images generated at approximated 3Hz, an independent planner frequency of 10Hz can be achieved.

### A. Simulation Experiments

We test our proposed planner in three diverse simulation environments: i) the *CARLA* urban dataset, ii) the *NVIDIA* warehouse environment, iii) the indoor *Matterport3d* dataset (Fig. 5), and compare it against the previous geometric iPlanner [1]. Both methods are trained from scratch with the data generated in our proposed pipeline. Three navigation tasks are shown in Fig. 5, for which the predictions of the two planners are provided. Intuitively, the proposed planner successfully uses both semantic and depth information to traverse low-cost areas, such as crosswalks in the urban environment, and avoid geometric and semantic obstacles, such as shelves or yellow working areas (treated as class unknown) in the warehouse environment. In contrast, the geometric planner perceives no difference between the terrains, resulting in straight motions

Scene	Method	Geom. Loss	Sem. Loss	Goal Reached
Matterport3D	iPlanner	$1.77 \pm 0.379$	$1.29 \pm 0.376$	55.37%
	Ours (Geom)	$1.76 \pm 0.543$	$1.27 \pm 0.538$	71.05%
	Ours (Sem)	<b><math>1.70 \pm 0.412</math></b>	<b><math>1.19 \pm 0.410</math></b>	<b>71.60%</b>
CARLA	iPlanner	$1.15 \pm 0.653$	$1.17 \pm 0.669$	79.59%
	Ours (Geom)	$1.03 \pm 0.579$	$1.06 \pm 0.708$	<b>95.91%</b>
	Ours (Sem)	<b><math>1.00 \pm 0.602</math></b>	<b><math>0.69 \pm 0.530</math></b>	87.75%
Warehouse	iPlanner	$0.90 \pm 0.544$	$1.793 \pm 1.274$	86.56%
	Ours (Geom)	$0.82 \pm 0.436$	$1.641 \pm 1.283$	<b>88.72%</b>
	Ours (Sem)	<b><math>0.71 \pm 0.396</math></b>	<b><math>0.96 \pm 1.041</math></b>	85.58%

**TABLE II:** Comparison between the geometric iPlanner [1] against our method trained with geometric data only (geom) and with geometric and semantic information (sem). Geom. and Sem. Loss metrics refer to the traversability loss of the trajectories evaluated on the geometric or semantic cost map respectively. Adding the semantics leads to recognizing geometrically invisible obstacles and interrupting more paths due to collision probabilities larger than  $\delta_{\mu}$ .



**Fig. 6:** Real-world experiment in the presence of geometric occurrences such as stairs. Four planning events (A - D) along the path with corresponding semantic and depth images are shown with the estimated paths of the proposed planner and of iPlanner. iPlanner’s first two predictions are marked completely in red (indicating a high estimated collision probability of  $\sim 0.98 > \delta_\mu$ ), causing the robot to stop.

towards the goal position.

Next, we evaluate our proposed method quantitatively by running experiments on a larger scale. We train our planner once with access only to the geometric domain and once with the fusion of depth and semantics to highlight the effects of the new loss and costmap design. Table II shows the results of 500 random paths with unique start-goal configurations in each of the three simulated environments. We report the rate of paths that reach the goal up to a threshold distance of  $0.5m$  and the traversability loss based on geometric and semantic costmaps for these paths. When comparing iPlanner to ViPlanner trained only with the geometric domain, a constant decrease over both costmaps is evident, together with an average increase of 16% in reached goal points for the Matterport3D and Carla environment. Due to the spacious layout of the warehouse environment, paths are easier to reach, which results in a less significant decrease. This loss improvement can be attributed to the usage of the robot width in Equation (3), compared to the point approximation of the robot used in iPlanner and the different smoothing process of the costmap. Our planner succeeds in using the semantic information, evident in an average semantic loss decrease of 38.02% in the CARLA and warehouse environment, where not all obstacles are detectable through geometry, e.g., the roads. In the same environments, a more conservative behavior of the proposed method can be observed, as the planner is aware of high-cost areas in the semantic domain, such as a road or a working area. Due to the unawareness of the geometric planners, they are not bound by these additional constraints and proceed to the goal, resulting in a higher percentage of "Goal Reached".

### B. Real-World Experiments

We demonstrate zero-shot sim-to-real capabilities in two experiments, contrasting previous work such as [1,8], where real data must be mixed into the training process. Our planner, trained solely on simulated data, has been employed in various scenarios despite diverse lighting conditions, new obstacle shapes, unknown scene compositions, and sensor noise.

The first experiment is shown in Fig. 1, for which the planner is tasked to navigate in an urban outdoor environment. As the goal is to navigate safely, the planner correctly commands the robot to cross the street at the crosswalk location and proceed further on the sidewalk. With the provided sensor measurements, it recognizes lower cost areas and demonstrates adjusted paths that consider the geometrically invisible constraints. Larger distances can be traversed with a small number of manually selected waypoints. The second experiment investigates the planners’s ability to handle seemingly geometric obstacles that, in reality, can be overcome by a legged robot, such as stairs. Fig. 6 showcases parts of the experiment, highlighting the ability to recognize and navigate stairs correctly. A comparison to iPlanner shows that a purely geometric approach struggles in such situations, and its collision estimates ( $> \delta_\mu$ ) will prevent the execution of the predictions, ultimately failing to reach the goal position.

## VII. CONCLUSIONS & FUTURE WORK

In this work, we presented a semantic-aware end-to-end trained local planner for deployment in semi-structured environments, fully trained on simulated data and transferable to the real world. With the presented loss and costmap design, up to 16% more goals can be reached compared to previous geometric approaches. Moreover, by fusing geometric and semantic information, our planner decreases the semantic traversability loss by up to 42%. The developed simulation plugins, the scalable data generation pipeline, and the planner models are open-sourced to expedite forthcoming research.

For future work, we will investigate how to remove hand-crafted loss values to make the costmap more generic and work with open-vocabulary representations. In addition, we aim to determine the sensitivity against semantic segmentation errors. Also, we will incorporate memory to enhance consistency and prevent the planner from "forgetting" obstacles.

### ACKNOWLEDGEMENT

The authors thank Turcan Tuna and David Höller for their support during experiments and the scientific discussions.

## REFERENCES

- [1] F. Yang, *et al.*, “iPlanner: Imperative path planning,” in *Robotics: Science and Systems Conference (RSS)*. Robotics: Science and Systems Foundation, 2023.
- [2] E. Jelavic, *et al.*, “Open3d slam: Point cloud based mapping and localization for education,” in *Robotic Perception and Mapping: Emerging Techniques, ICRA 2022 Workshop*. ETH Zurich, Robotic Systems Lab, 2022, p. 24.
- [3] Z. Bao, *et al.*, “A review of high-definition map creation methods for autonomous driving,” *Engineering Applications of Artificial Intelligence*, vol. 122, p. 106125, 2023.
- [4] C. Cao, *et al.*, “Autonomous exploration development environment and the planning algorithms,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 8921–8928.
- [5] L. Wellhausen and M. Hutter, “Rough terrain navigation for legged robots using reachability planning and template learning,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 6914–6921.
- [6] B. Yang, *et al.*, “Real-time optimal navigation planning using learned motion costs,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 9283–9289.
- [7] A. Loquercio, *et al.*, “Learning high-speed flight in the wild,” *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021.
- [8] D. Hoeller, *et al.*, “Learning a state representation and navigation in cluttered and dynamic environments,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5081–5088, 2021.
- [9] N. El-Sheimy and Y. Li, “Indoor navigation: State of the art and future trends,” *Satellite Navigation*, vol. 2, no. 1, pp. 1–23, 2021.
- [10] M. Tranzatto, *et al.*, “Cerberus in the darpa subterranean challenge,” *Science Robotics*, vol. 7, no. 66, p. eabp9742, 2022.
- [11] D. Maturana, *et al.*, “Real-time semantic mapping for autonomous off-road navigation,” in *Field and Service Robotics*, M. Hutter and R. Siegwart, Eds. Springer International Publishing, 2018, vol. 5, pp. 335–350, series Title: Springer Proceedings in Advanced Robotics.
- [12] A. Shaban, *et al.*, “Semantic terrain classification for off-road autonomous driving,” in *Proceedings of the 5th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Faust, *et al.*, Eds., vol. 164. PMLR, 08–11 Nov 2022, pp. 619–629.
- [13] M. Mueller, *et al.*, “Driving policy transfer via modularity and abstraction,” in *Proceedings of The 2nd Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, A. Billard, *et al.*, Eds., vol. 87. PMLR, 29–31 Oct 2018, pp. 1–15.
- [14] S. Hecker, *et al.*, “Learning accurate and human-like driving using semantic maps and attention,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 2346–2353.
- [15] L. Liu, *et al.*, “Robot navigation in crowded environments using deep reinforcement learning,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5671–5677.
- [16] M. Hutter, *et al.*, “Anymal-toward legged robots for harsh environments,” *Advanced Robotics*, vol. 31, no. 17, pp. 918–931, 2017.
- [17] M. Mittal, *et al.*, “Orbit: A unified simulation framework for interactive robot learning environments,” *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023.
- [18] B. Paden, *et al.*, “A survey of motion planning and control techniques for self-driving urban vehicles,” *IEEE Transactions on intelligent vehicles*, vol. 1, no. 1, pp. 33–55, 2016.
- [19] N. Ratliff, *et al.*, “Chomp: Gradient optimization techniques for efficient motion planning,” in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 489–494.
- [20] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [21] L. Kavraki, *et al.*, “Probabilistic roadmaps for path planning in high-dimensional configuration spaces,” *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [22] S. Schaal, *et al.*, “Control, planning, learning, and imitation with dynamic movement primitives,” in *Workshop on Bilateral Paradigms on Humans and Humanoids: IEEE International Conference on Intelligent Robots and Systems (IROS 2003)*, 2003, pp. 1–21.
- [23] M. Dharmadhikari, *et al.*, “Motion primitives-based path planning for fast and agile exploration using aerial robots,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 179–185.
- [24] E. Jelavic, *et al.*, “Lstp: Long short-term motion planning for legged and legged-wheeled systems,” *IEEE Transactions on Robotics*, 2023.
- [25] J.-P. Sleiman, *et al.*, “Versatile multicontact planning and control for legged loco-manipulation,” *Science Robotics*, vol. 8, no. 81, p. eadg5014, 2023.
- [26] J. Frey, *et al.*, “Fast traversability estimation for wild visual navigation,” in *Robotics: Science and Systems Conference (RSS)*. Robotics: Science and Systems Foundation, 2023.
- [27] L. Wellhausen, *et al.*, “Where should i walk? predicting terrain properties from images via self-supervised learning,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1509–1516, 2019, conference Name: IEEE Robotics and Automation Letters.
- [28] J. Frey, *et al.*, “Locomotion policy guided traversability learning using volumetric representations of complex environments,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 5722–5729.
- [29] N. Hudson, *et al.*, “Heterogeneous ground and air platforms, homogeneous sensing: Team CSIRO data61’s approach to the DARPA subterranean challenge,” *Field Robotics*, vol. 2, no. 1, pp. 595–636, Mar. 2022.
- [30] D. D. Fan, *et al.*, “Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation,” in *Robotics: Science and Systems Conference (RSS)*. Robotics: Science and Systems Foundation, 2021.
- [31] L. Bartolomei, *et al.*, “Perception-aware path planning for uavs using semantic segmentation,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 5808–5815.
- [32] V. Vasilopoulos, *et al.*, “Reactive semantic planning in unexplored semantic environments using deep perceptual feedback,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4455–4462, 2020.
- [33] X. Cai, *et al.*, “Risk-aware off-road navigation via a learned speed distribution map,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 2931–2937.
- [34] M. Pfeiffer, *et al.*, “From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots,” in *2017 IEEE international conference on robotics and automation (icra)*. IEEE, 2017, pp. 1527–1533.
- [35] T. Zhang, *et al.*, “Learning deep control policies for autonomous aerial vehicles with mpc-guided policy search,” in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 528–535.
- [36] M. Caron, *et al.*, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
- [37] M. Bansal, *et al.*, “ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst,” in *Robotics: Science and Systems XV*. Robotics: Science and Systems Foundation, 2019.
- [38] E. Wijmans, *et al.*, “DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames,” in *International Conference on Learning Representations (ICLR)*, 2020.
- [39] T. Fu, *et al.*, “islam: Imperative slam,” *arXiv preprint arXiv:2306.07894*, 2023.
- [40] K. He, *et al.*, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [41] A. Dosovitskiy, *et al.*, “Carla: An open urban driving simulator,” in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [42] A. Chang, *et al.*, “Matterport3d: Learning from rgb-d data in indoor environments,” *International Conference on 3D Vision (3DV)*, 2017.
- [43] J. Nubert, *et al.*, “Learning-based localizability estimation for robust lidar localization,” in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 17–24.
- [44] J. H. Halton, “On the efficiency of certain quasi-random sequences of points in evaluating multi-dimensional integrals,” *Numerische Mathematik*, vol. 2, no. 1, pp. 84–90, Dec. 1960.
- [45] B. Cheng, *et al.*, “Masked-attention mask transformer for universal image segmentation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 1290–1299.