

Competence-Aware Path Planning Via Introspective Perception

Sadegh Rabiee[✉], Connor Basich, *Graduate Student Member, IEEE*, Kyle Hollins Wray, *Member, IEEE*, Shlomo Zilberstein, and Joydeep Biswas[✉], *Member, IEEE*

Abstract—Robots deployed in the real world over extended periods of time need to reason about unexpected failures, learn to predict them, and to proactively take actions to avoid future failures. Existing approaches for competence-aware planning are either model-based, requiring explicit enumeration of known failure sources, or purely statistical, using state- and location-specific failure statistics to infer competence. We instead propose a structured model-free approach to competence-aware planning by reasoning about plan execution failures due to errors in perception, without requiring a priori enumeration of failure sources or requiring location-specific failure statistics. We introduce *competence-aware path planning via introspective perception (CPIP)*, a Bayesian framework to iteratively learn and exploit task-level competence in novel deployment environments. CPIP factorizes the competence-aware planning problem into two components. First, perception errors are learned in a model-free and location-agnostic setting via *introspective perception* prior to deployment in novel environments. Second, during actual deployments, the prediction of task-level failures is learned in a context-aware setting. Experiments in a simulation show that the proposed CPIP approach outperforms the frequentist baseline in multiple mobile robot tasks, and is further validated via real robot experiments in environments with perceptually challenging obstacles and terrain.

Index Terms—Motion and path planning, visual learning, failure detection and recovery.

I. INTRODUCTION

AS ROBOTS become increasingly available, they are deployed for tasks where autonomous navigation in uncontrolled environments is crucial to success, such as package delivery, warehouse automation, and home service settings. Deploying robots over extended periods of time and in such

Manuscript received September 9, 2021; accepted January 4, 2022. Date of publication January 25, 2022; date of current version February 7, 2022. This letter was recommended for publication by Associate Editor L. Ott and Editor C. Cadena-Lerma upon evaluation of the reviewers comments. This work was supported in part by NSF under Grants CAREER-2046955 and IIS-1954778 and in part by DARPA under Grant HR001120C0031. The views and conclusions contained in this document are those of the authors only. (*Corresponding author: Sadegh Rabiee.*)

Sadegh Rabiee and Joydeep Biswas are with the Department of Computer Science, The University of Texas at Austin, Austin, TX 78712 USA (e-mail: sad.rabiee@gmail.com; joydeepb@cs.utexas.edu).

Connor Basich and Shlomo Zilberstein are with the College of Information and Computer Sciences, The University of Massachusetts Amherst, Amherst, MA 01003 USA (e-mail: cbasich@cs.umass.edu; shlomo@cs.umass.edu).

Kyle Hollins Wray is with the Alliance Innovation Lab Silicon Valley, Santa Clara, CA 95051 USA (e-mail: kyle.hollins.wray@gmail.com).

This letter has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2022.3145517>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2022.3145517

open world setting requires addressing failures originating from real-world uncertainty and imperfect perception. Continuous operator monitoring, while effective, is cumbersome and thus not scalable to many robots or large environments. We are thus interested in developing *competence-aware* agents capable of assessing the probability of successfully completing a given task. Such agents would learn from failures and leverage the acquired knowledge when planning to improve their robustness and reliability. Previous efforts towards competence-aware path planning and motion planning either rely solely on statistical analysis of logged instances of failures in the configuration space of the robot and do not benefit from the sensing information collected by the robot [1], or are application specific and designed to reduce the probability of failure for a specific perception module such as visual SLAM [2]. While there has been progress on introspective perception to enable perception algorithms to learn to predict their sources of errors [3], [4], the outputs of such algorithms have not yet been exploited in robot planning.

We present competence-aware path planning via introspective perception (CPIP), a general framework that bridges the gap between path planning and introspective perception and allows the robot to iteratively learn and exploit task-level competence in novel deployment environments. CPIP models the path planning problem as a Stochastic Shortest Path (SSP) problem and builds a model that represents both the topological map of the environment as well as the competence of the robot in traversing each part of the map autonomously. CPIP leverages introspective perception to predict the task-level competence of the robot in novel deployment environments and employs a Bayesian approach to update its estimate of the robot competence online and during the deployment. CPIP then uses this information to plan paths that reduce the risk of failures.

Our experimental results demonstrate that CPIP converges to the optimal planning policy in novel deployment environments while reducing the frequency of navigation failures by more than 80% compared to the state-of-the-art competence-aware path planning algorithms that do not leverage introspective perception.

II. RELATED WORK

The idea of integrating perception with planning and control was introduced by pioneering works on active perception that suggested performance of perception can be improved by

selecting control strategies that depend on the current state of perception data interpretation as well as the goal of the task [5], [6]. Researchers have applied this idea to various levels of control ranging from active vergence control for a stereo pair of cameras [7] to object manipulation given the next best view for surface reconstruction of unknown objects [8].

One line of work predicts and avoids degradation of perception performance given features extracted from the raw sensory data. Costante *et al.* [2] propose a perception-aware path planning for MAVs that maximizes the information gain from image matching while solving for dense V-SLAM. Sadat *et al.* [9] and Deng *et al.* [10] follow a similar approach and use an RRT* planner where the cost of a path is defined as a linear combination of the length of the path and the predicted density of image features along the path to reduce localization errors. In these works, estimates of competence for perception are obtained via hand crafted metrics and the path planner cost function is designed to specifically address the reliability of V-SLAM and is not generalizable to arbitrary perception tasks.

There exists a body of work on risk-aware path planning, where it is assumed that the autonomous agent has accurate models of the uncertainty of perception. Jasour *et al.* [11] use the *a priori* known parametric probability distributions for obstacle locations and leverage chance constrained optimization to plan paths that have collision probabilities below a user-specified threshold. Schirmer *et al.* [12] use an offline-built localization uncertainty map of the environment to do risk-aware path planning. Barbosa *et al.* [13] and similarly Chung *et al.* [14] relax the full-observability assumption and do path planning in partially known environments, yet they assume that the agent has an observation model with a known noise process that is used to update its belief over the state of the world. However, in practice estimates of uncertainty of perception algorithms, as obtained by methods such as the Cramer-Rao lower bound, are often overconfident and inaccurate. An alternative approach is to directly model task-level failures as a function of the state of the world as acquired by perception. Saxena *et al.* [15] learn to predict task-level failures that are due to errors in perception from the raw sensory data; however, predicted failures are used to trigger an enumerated set of recovery actions rather than proactively generating plans that reduce the probability of failures. Similarly Gurau *et al.* [16] leverage image data and location specific features to do reactive planning by selecting between different levels of autonomy at any point in time.

A different line of work on competence-aware path planning that has a more holistic view of failures includes keeping track of all of the robot failures regardless of the perception algorithm that is the cause, and then leveraging this information to proactively generate plans with reduced risk of failures. Lacerda *et al.* [17] aggregate the failure instances of a service mobile robot while navigating the environment to model the probability of success for traversing each edge of a topological map using an MDP and generate navigation policies that prefer paths with high success probabilities. Krajník [18] use a spectral model to learn mid to long-term environmental changes assuming they have a periodic nature and exploit it to improve robot navigation and localization by predicting such changes. Vintr *et*

al. [19] use a similar approach to learn a spatio-temporal model for predicting presence of humans in the robot's deployment environment at different times of the day. Since these methods are based on statistical analysis of the frequency of navigation failures, they require ample experience and several samplings from any location in the map, in order to achieve an accurate estimate of the robot's competence in navigating that specific location. Moreover, due to using location specific features of the environment for estimating the robot competence, these estimates cannot be generalized to novel deployment environments. Basich *et al.* [20] further expand the concept of competence to the optimal level of autonomy and define a stochastic model for solving the path planning problem, where the generated plans consist of a path and the optimal level of autonomy for each segment of the path. In order to learn to predict the probability of failure at each level of autonomy this work requires a curated list of environmental features that are potentially correlated with robot failures.

In this work we leverage machine learned models capable of predicting errors of individual perception modules, to obtain an accurate estimate of the robot's competence at successfully navigating throughout an environment. CPIP uses the estimate of competence to plan reliable and short-duration paths. Our work is similar to [17], [20] in that it reasons about the competence of the robot at successfully performing navigation tasks at a topological map level; however, it removes the need for an enumerated list of perception related features by automatically learning to extract such features from the raw sensory data. Furthermore, CPIP significantly reduces the frequency of failures experienced in new environments by exploiting the generalizable learned perception features instead of merely relying on statistical analysis of the location of previous navigation failures.

III. CPIP DEFINITION

CPIP is a framework for integrating path planning with introspective perception in life-long learning settings. It is defined as a tuple $\langle \mathcal{M}, \mathcal{I}, \mathcal{H} \rangle$, where \mathcal{M} is a stochastic planning model, $\mathcal{I} = \{\mathbb{I}_k\}_{k=1}^N$ is a set of introspective perception modules, and \mathcal{H} is a task-level competence predictor. CPIP leverages introspective perception and the competence predictor model to predict the probability of task-level failures given the raw sensory data at every time step and uses these estimates to update the planning model iteratively during robot deployments, hence learns policies that reduce the probability of failures. In Section IV, we introduce the planning model \mathcal{M} and explain how it incorporates the probability of autonomous navigation failure in path planning. We then explain introspective perception \mathcal{I} and the competence predictor model \mathcal{H} and how they are used to structure the problem of learning to predict instances of navigation failures in Section V.

IV. COMPETENCE-AWARE PLANNING

The CPIP planning model \mathcal{M} uses a representation of the environment that includes both the connectivity of a set of sparse locations on the map as well as the probability of successful traversal between each two connected neighboring locations. In

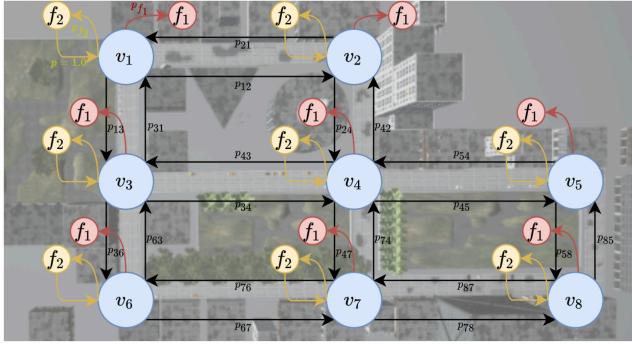


Fig. 1. Planning SSP for an example environment.

this section, we explain this model and how it is actively updated during deployments.

A. Planning Model Description

The input to our problem is a topological map of the environment in the form of a directed graph $G = \langle N, E \rangle$ composed of a set of nodes, N , and a set of edges, E . Each node represents a location, and each edge e is defined by a tuple $\langle n_e^i, n_e^j, t_e, p_e \rangle$. Here, n_e^i is the starting vertex, n_e^j is the ending vertex, t_e is the expected traversal time for the edge e , and p_e is the probability of successfully traversing the edge.

Given the topological map, G , we model the planning problem as a Stochastic Shortest Path (SSP) problem, a formal decision-making model for reasoning in stochastic environments where the objective is to find the least-cost path from a start state to a goal state. An SSP is a tuple $\langle S, A, T, C, s_0, G \rangle$ where S is a finite set of states, A is a finite set of actions, $T : S \times A \times S \rightarrow [0, 1]$ represents the probability of reaching state $s' \in S$ after performing action $a \in A$ in state $s \in S$, $C : S \times A \rightarrow \mathbb{R}^+$ represents the expected immediate cost of performing action $a \in A$ in state $s \in S$, $s_0 \in S$ is an initial state, and $G \subset S$ is a finite (possibly singleton) set of goal states such that $T(s_g, a, s_g) = 1 \wedge C(s_g, a) = 0 \forall a \in A, s_g \in G$.

A solution to an SSP is a policy $\pi : S \rightarrow A$ that indicates that action $\pi(s) \in A$ should be taken in state $s \in S$. A policy π induces the value function $V^\pi : S \rightarrow \mathbb{R}$ that represents the expected cumulative cost $V^\pi(s)$ of reaching s_g from state s following the policy π . An optimal policy π^* minimizes the expected cumulative cost $V^*(s_0)$ from the initial state s_0 .

In our problem, $S = N \times \tilde{S}$ is a finite set of states comprised of the map nodes N and a finite set of failure states \tilde{S} and $A = E \cup \tilde{A}$ is a finite set of actions comprised of the directed edges in the graph and a finite set of recovery actions \tilde{A} . $T(s, a, s')$ is determined by the probability of successfully traversing the edge $e = (s, s')$, p_e , which is zero if the action a does not correspond to the edge e . In a failure state s , $T(s, a, s') = 0$ if $a \notin \tilde{A}$ and $s \neq s'$. $C(s, a)$ is set to t_e if $a \in E$, and the expected recovery cost for s otherwise. Fig. 1 illustrates the planning SSP for an example urban environment. During robot deployments, the transition function is updated to reflect the latest belief over the probability of navigation failures in traversing each edge on the map, or equivalently the probability of successful traversals. Next, we explain the method for updating the transition function.

B. Updating the Failure Belief During Deployment

CPIP builds an SSP model to represent the topological map of the environment as described in Section IV-A. CPIP updates the aforementioned SSP model structure during deployments as it collects more observational data from the environment, altering the underlying transition function such that the resultant model represents not just the map but the competence of the robot in traversing it. In order to achieve that, the occurrence of a failure of class f_i at edge e is assumed to be a random variable from the categorical distribution $F_e \sim \text{Cat}(p_{1:L})$. The belief over this variable is defined as $B_t(F_e = f_i) = p(F_e = f_i|z_{1:t})$, where the subscript t indicates the t^{th} traversal of the edge and z_t is the observation made by the robot during that traversal. Applying the Bayes rule yields

$$\begin{aligned} \text{bel}_t(f_{i,e}) &= \frac{p(z_t|f_{i,e}, z_{1:t-1}) p(f_{i,e}|z_{1:t-1})}{p(z_t|z_{1:t-1})} \\ &= \frac{p(f_{i,e}|z_t) p(z_t) \text{bel}_{t-1}(f_{i,e})}{p(f_{i,e}) p(z_t|z_{1:t-1})} . \end{aligned} \quad (1)$$

Defining the negation of $f_{i,e}$ as $p(\neg f_{i,e}) = 1 - p(f_{i,e}) = \sum_{j \neq i} p(f_{j,e})$ the belief can be implemented as the log odds ratio

$$\begin{aligned} l_t(f_{i,e}) &= \log \left(\frac{\text{bel}_t(f_{i,e})}{\text{bel}_t(\neg f_{i,e})} \right) \\ &= \log \left(\frac{p(f_{i,e}|z_t)}{1 - p(f_{i,e}|z_t)} \right) + l_{t-1}(f_{i,e}) - l_0(f_{i,e}) , \end{aligned} \quad (2)$$

where $l_0(f_{i,e}) = \log \left(\frac{p(f_{i,e})}{1-p(f_{i,e})} \right)$ is the prior in log odds form. Before the first deployment of the robot in a new environment, $p(f_{i,e}) = \epsilon$ and $p_e = p(f_{L,e}) = 1 - \sum_{i \neq L} p(f_{i,e})$ for every $e \in E$. Then upon each traversal of an edge, the above relation is used to update the transition function of the planning SSP model such that $T_t(s, e, \tilde{s}_t) = \text{bel}_t(f_{i,e}) = 1 - \frac{1}{1+\exp(l_t(f_{i,e}))}$. The main term that needs to be computed for updating the belief in Eq. 2 after each traversal is $p(f_{i,e}|z_t)$, which is known as the inverse observation likelihood and in CPIP it is implemented by two different functions, each handling one of the two different types of observations z_t : 1) Occurrence of failures of class f_i which is indicated via intervention signals issued either by a human or a supervisory sensing unit and is denoted by $s_{t,i}$; 2) Sensory input that the robot continuously acquires such as RGB images captured by cameras on the robot, which is denoted by \mathbf{I}_t . For the former, the inverse observation likelihood is implemented as

$$p(f_{i,e}|z_t = s_{t,j}) = \begin{cases} \delta & i = j \\ \frac{1-\delta}{L-1} & i \neq j \end{cases} \quad (3)$$

where δ is a constant coefficient. The inverse observation likelihood function for the latter type of observations, however, is machine learned and is one of the key components of this work that allows CPIP to reach an accurate estimate of $\text{bel}_t(f_{i,e})$ without requiring the robot to experience costly failures. CPIP structures the problem of learning $p(f_{i,e}|z_t = \mathbf{I}_t)$ such that it can be achieved with a small number of failure examples for training

data. Introspective perception is leveraged to extract features associated with errors in perception from the high dimensional raw sensory data. These features are then used to learn to predict the probability of different classes of failures of navigation. By learning this likelihood function, the robot will learn to better navigate its environment, proactively avoiding paths that are known to lead to failure cases, and reactively adjusting its policy upon encountering novel situations that may lead to failures. In the following section we describe the different parts of this learning problem.

V. FAILURE PREDICTION VIA INTROSPECTIVE PERCEPTION

In order to predict failures of navigation given the sensory data, we need to approximate the function $p(f_{i,e}|z_t = \mathbf{I}_t)$. End-to-end learning of this function is intractable because it requires a great amount of training data, yet catastrophic failures in robotics when executing tasks such as autonomous navigation do not happen frequently. The scarcity of these examples makes it challenging to learn a classifier that predicts the probability of task execution failure directly from the raw sensory data. Without enough training data and without abstracting the acquired high dimensional sensory data, the learned classifier is bound to overfit to the training data. We instead propose to factorize $p(f_{i,e}|z_t) = \int_{\phi} p(f_{i,e}|\phi)p(\phi|z_t)$ where ϕ are the features extracted from observations by introspective perception — a model-free approach to predicting arbitrary errors of perception.

A. Introspective Perception

Early works on introspective perception [21], [22] defined a perception algorithm to be introspective if it is capable of predicting the occurrence of errors in its output given the current state of the robot. Follow-up works [3], [4] extended this definition and required such perception algorithms to predict the probability of perception error conditioned on the region in the raw sensory data that the output is dependent upon, e.g. an image patch in the captured image by an RGB camera where the estimated depth of the scene is erroneous. This is obtained by means of an introspection function that is trained on empirical data.

In CPIP, the robot is assumed to be equipped with one or more introspective perception modules; each module has a learned function $\mathbb{I}_k : Z \rightarrow \mathbb{R}^n$, which extracts features $\phi \in \mathbb{R}^n$ from the raw sensory data Z that encode information about sources of perception errors. The outputs of all introspective perception modules are fed to a navigation competence predictor $h : \mathbb{R}^{n \times K} \rightarrow [0, 1]^L$, which learns to estimate the likelihood of each of the different classes of failure $f_{1:L}$ given a set of sources of perception errors, i.e. $P(f_l|\phi_{1:K})$ such that $\sum_l P(f_l|\phi_{1:K}) \leq 1$. The inverse observation likelihood function in Eq. 2 is then estimated as the composition of the above two functions, i.e. $p(f_{i,e}|z_t = \mathbf{I}_t) = h(\mathbb{I}_{1:k}(\mathbf{I}_t))_{[i]}$. It should be noted that although CPIP assumes a constant set of failure classes, the number of distinct sources of failures, which is often much larger than the number of different classes of failure, are not enumerated *a priori*. Each failure class corresponds to a different severity level and hence a different failure recovery cost that is considered in

planning. There exist, however, a large number of failure sources that lead to failures with the same severity level. For example, a high-severity class of failure in robot navigation are collisions, for which there exist numerous failure sources including false negatives in obstacle detection due to texture-less surfaces or small object sizes, terrain type mis-classification, dynamic obstacles, occlusions, etc. Furthermore, while the distinct sources of failure differ between environments, the classes of failure are specific to the objective of the domain (e.g. navigation, manipulation, etc.), irrespective of the environment, and are hence comparatively easy to enumerate.

In this paper we implement introspective perception for a block matching-based stereo depth estimator [23] using the same convolutional neural network architecture as that used in [3] for the introspection function. The training data is collected autonomously using a depth camera as supervisory sensing, which is only occasionally available and provides oracular information about the true depth of the scene.

B. Competence Predictor Model

We implement the navigation competence predictor model $h : \mathbb{R}^{n \times K} \rightarrow [0, 1]^L$ as an ensemble of two deep neural networks. The input to the model is a list of image patches $I_i \in \mathbb{R}^n$ extracted from the same input image \mathbf{I} and the output is the probability of each class of failure. The architecture as shown in Fig. 2 consists of two sub-models that are trained independently.

The `global_info` network is a convolutional neural network (CNN) that operates simultaneously on all input image patches arranged on a blank image in their original pixel coordinates. The input to this network is equivalent to the input image masked at all regions except for those predicted by introspective perception to lead to errors. The `global_info` CNN captures task-contextual and spatial information from the current frame related to competence. By masking out parts of the full image deemed to be unrelated to perception failures we are able to ensure that the `global_info` CNN does not overfit to specific environments.

The `local_info` network is a CNN that is fed as input individual image patches. The output of this branch is the probability of each class of failure for each single image patch. This network learns correlations between navigation failures and image features that lead to perception errors. The goal of this branch is to locally pinpoint the potential source of navigation failures in the image space, when a class of failure is predicted by the `global_info` network.

The last stage of the model is a temporal filtering of the output of each of the two networks. Failure class probabilities that are produced by the `global_info` network are passed through a mean filter to output $P_f \in [0, 1]^L$. Moreover, image patches that are predicted by the `local_info` network to lead to navigation failures are tracked in the full image over consecutive frames to form a set of active tracklets Λ_i for each class of failure f_i . The output of the model is obtained via strict consensus on the output

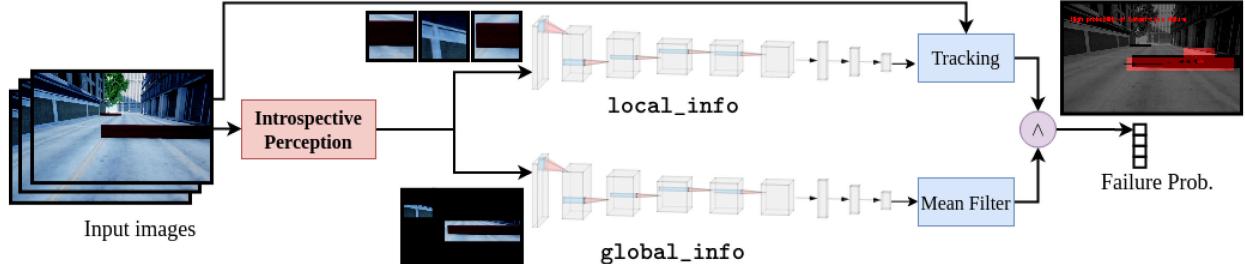


Fig. 2. Navigation competence predictor model architecture.

such that

$$P(f_i) = \begin{cases} P_{f_i} & \text{if } \Lambda_i \neq \emptyset \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

In other words, the predicted probability of each class of failure provided by the `global_info` network is only accepted if the `local_info` network also supports that by detecting at least one potential cause for the same class of failure in the image space. During deployment, if $\exists j \mid P(f_j) > \epsilon$, i.e. there exists consensus between the two branches of the network on the existence of any class of failure, the output of the competence predictor model will be used to update the belief in (2).

VI. IMPLEMENTATION DETAILS

In this section we provide implementation details for our application of CPIP, i.e. path planning for an unmanned ground vehicle (UGV) that uses a stereo vision-based depth estimator [23] for obstacle avoidance.

A. Autonomy Stack

Our navigation software consists of global path planning on the navigation graph and local path planning to follow the planned path while avoiding dynamic obstacles that the robot does not know about *a priori*. CPIP's planning model does the global path planning and we use a trajectory roll-out local path planner. The 3D reconstruction of the environment by stereo-vision is processed and any points with their height larger than 15 cm and less than the height of the robot are detected as obstacles. All obstacle points coordinates are projected to the ground plane, converted to a 2D laser-scan format, and used by the local path planner to select the least cost trajectory from a set of sampled trajectories, such that the robot keeps a large clearance from obstacles and makes progress towards the next way-point on the global plan.

We implement introspective perception for the depth estimator and train a CNN to predict depth estimation errors similar to prior work [3]. The network is composed of 5 convolution layers followed by 3 fully connected layers. The input to the network are the image patches of size 70×70 pixels extracted from the 512×512 pixel images captured by the left stereo pair. The output is the probability of depth estimation error and all image patches predicted to lead to perception error with a probability of > 0.5 are passed to the competence predictor model. We use a similar CNN architecture as that used for introspective depth estimation for both sub-models of the competence predictor with

the only difference being the number of nodes in the output layer of the network. The full pipeline of navigation competence predictor which consists of introspective perception and the competence predictor model runs at 5 Hz on a laptop with Intel Core i9-9880H and GeForce RTX 2080 Max-Q.

B. Training of CPIP

CPIP has two learned components, i.e. introspective perception and the competence predictor model and they are trained sequentially. The training data is extracted from logs of robot deployments in the training environment. The logs include data collected by the primary sensors, i.e. RGB images captured by the stereo cameras, as well as data collected by supervisory sensors, i.e. the Orbbec Astra depth camera that is only occasionally available. Furthermore, intervention signals issued by a human operator upon occurrence of navigation failures are also recorded.

The deployment logs are processed offline. First, introspective perception is trained with data that is autonomously labeled using the supervisory sensing. Then, the training data for the competence predictor model is prepared by passing the raw sensory data through the introspective perception module and labeling the output image patches as associated with one of the classes of navigation failures if they fall within a fixed time window preceding the occurrence of such failures. Each of the two sub-models of the competence predictor model explained in §V-B are then trained using a cross-entropy loss.

VII. EXPERIMENTAL RESULTS

In this section: 1) We evaluate CPIP on how well it predicts sources of robot failures. 2) We compare CPIP against baseline global path planners in terms of their task completion success rate and their task completion time. 3) We evaluate the importance of introspective perception in CPIP's performance and generalizability via ablation studies.

A. Experimental Setup

1) *Simulation:* In order to evaluate CPIP and compare it against SOTA extensively, we use AirSim [12], a photo-realistic simulation environment, where robot failures are not expensive and the robot can easily be reset upon occurrence of navigation failures. A simulated car is equipped with a stereo pair of RGB cameras as well as a depth camera that provides ground truth depth readings for every pixel in the camera frame. We use



Fig. 3. Training and test environments in the real-world experiments.

two separate urban environments for training and testing. The environments are populated with obstacles of different shapes, textures, and colors.

2) *Real-robot maze*: We also evaluate CPIP on a real robot. We use a Clearpath Husky equipped with a stereo pair of RGB cameras, an Orbbec Astra depth camera, and a Velodyne VLP-16 3D Lidar. We use different indoor sites for training and testing of CPIP. Each environment has different types of terrain such as tiles and carpet, and is populated with obstacles of different shapes, textures, and surface materials. The test environment is a maze constructed in an area of size 60 m^2 .

3) *Real-robot large-scale*: In order to test CPIP extensively and in more natural settings, we also conduct a large scale experiment, where we deploy the robot on the entire floor of a building. This test environment has an area of larger than 400 m^2 and the robot traverses more than 1.5 km during the deployment. Fig. 3 shows the training and both the large scale and maze test environments.

B. FAILURE PREDICTION ACCURACY

In order to evaluate the accuracy of CPIP in predicting failures of navigation, we have the autonomous agent traverse each of the edges of the navigation graph in the test simulation environment 50 times and run the captured images by the robot camera through the CPIP’s introspective perception module and the competence predictor model to predict instances of navigation failure. In this paper, we implement CPIP with two classes of failures. 1) Catastrophic failures, where the robot ends up in a state that precludes completion of the task and is not recoverable with human intervention. Examples of this class include collisions and the robot getting stuck off-road. 2) Non-catastrophic failures, where the robot will not be able to complete its task unless intervention is provided by a human operator or a supervisory sensor. The robot getting stuck due to false detection of obstacles or because of localization errors are examples of this type of failure. Fig. 4 illustrates the predicted and actual navigation failures in a confusion matrix. CPIP correctly predicts occurrence of navigation failures more than 70% of the time for both types of failures. Prediction errors mostly correspond to cases, where the source of failure is significantly different looking from the examples available in the training data.

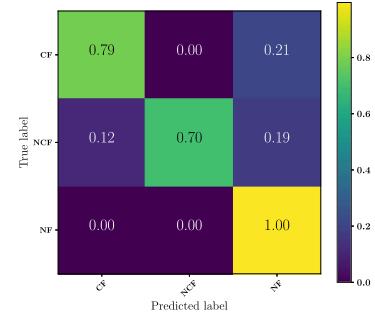


Fig. 4. Prediction results of the competence predictor model in the previously unseen test environment for the three classes of catastrophic failures (CF), non-catastrophic failures (NCF), and no failures (NF).

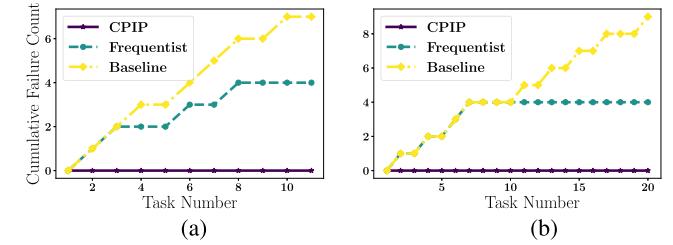


Fig. 5. Comparison of cumulative failure count (a) in the real-robot maze experiment, and (b) the real-robot large scale experiment for this work (CPIP), SOTA (frequentist), and the baseline with no competence-aware planning.

C. NAVIGATION SUCCESS RATE AND PLAN OPTIMALITY

We test the end-to-end system in predicting navigation failures and proactively planning paths that reduce the probability of failures, by deploying the robot in a previously unseen test environments. The robot is commanded to complete randomly generated navigation tasks that consist of a starting pose and a target pose. We conduct this experiment in all three settings explained in Section VII-A, i.e. simulation, real-robot maze, and the large-scale real-robot deployment.

We compare CPIP with a baseline path planner that does not reason about the competence of the robot as well as a state-of-the-art approach for competence-aware path planning — called the Frequentist approach — that relies on keeping track of the frequency of past failures in traversing each of the edges of the navigation graph [17]. Fig. 5 compares the cumulative failure count for all three methods throughout the real-robot experiments. With the Frequentist approach, the robot learns to avoid regions of the environment, where it cannot navigate reliably as it experiences navigation failures. However, CPIP enables the robot to predict and avoid most of these failures, leading to the least number of experienced failures.

We also evaluate the optimality of the planned paths by comparing the task completion time for all the methods under test with an oracular path planner that is given the true probability of navigation failures for each edge of the navigation graph. The ground truth failure probabilities are obtained by having the agent traverse each edge of the navigation graph numerous times and logging the frequency of each type of failure. Fig. 6 shows the completion duration for each task in the simulation experiment. The duration values are normalized by the task completion duration when the oracular path planner is used. The

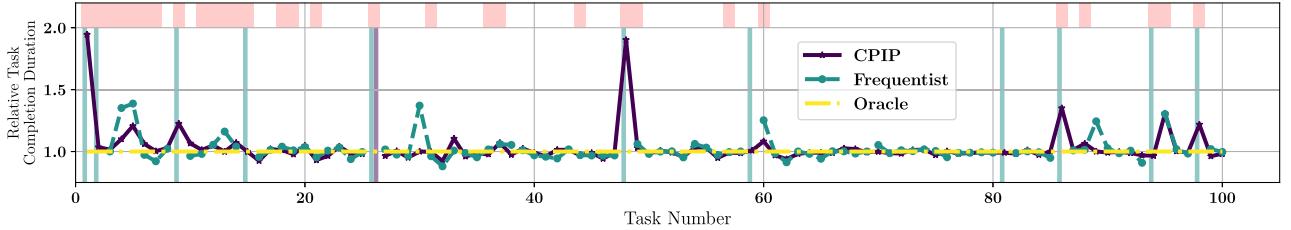


Fig. 6. Task completion duration w.r.t. an oracle planner that is provided with the true probability of failure throughout the environment ahead of deployment. Vertical bars visualize the incomplete tasks for each method annotated by color. Highlighted red regions in the top band demonstrate tasks, during which the robot encounters previously unseen parts of the environment. Best viewed in color.

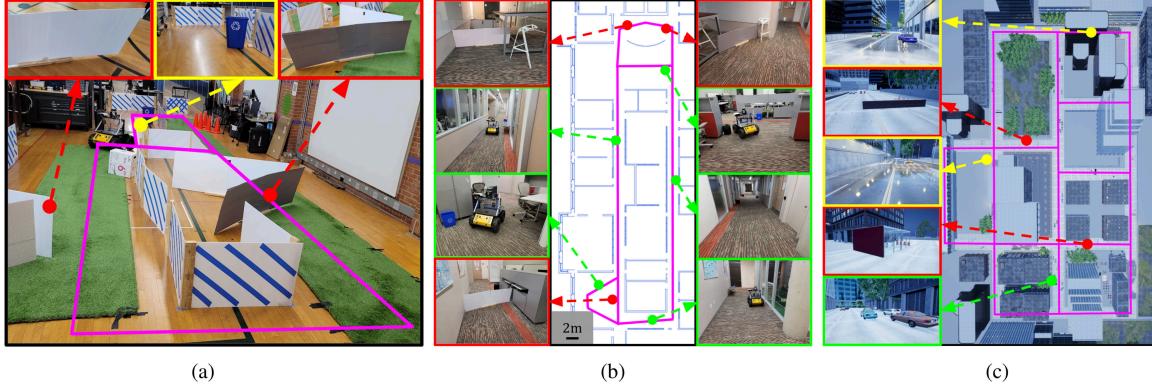


Fig. 7. Test environments in (a) the real robot maze experiment, (b) large-scale real robot experiment, and (c) the simulation experiment. Regions of the environments highlighted in red cause catastrophic failures, regions highlighted in yellow illustrate sources of non-catastrophic failures, and areas annotated with green, show areas where the robot can successfully operate autonomously.

TABLE I
TASK COMPLETION AND FAILURE PREVENTION RATE

	# Tasks	TCR (%)	Relative TCD*		# Avoided Failures		
			Mean	Std	CF	NCF	
Real Robot Maze	CPIP	11	100	1.04	0.08	5 (100%)	3 (100%)
	Frequentist	11	73	1.22	0.43	3 (60%)	1 (33%)
Real Robot Large-Scale	CPIP	20	100	0.98	0.04	9 (100%)	-
	Frequentist	20	80	0.99	0.04	5 (55%)	-
Simulation	CPIP	100	97	1.00	0.05	14 (93%)	61 (97%)
	Frequentist	100	83	1.02	0.09	9 (60%)	52 (83%)

*Task completion duration statistics are only calculated for tasks that were completed by both algorithms.

figure also illustrates instances of task completion failures for both CPIP and the Frequentist method. Such instances include occurrence of catastrophic failures or occurrence of consecutive non-catastrophic failures such that the robot cannot recover from a stuck state by re-planning. CPIP task completion duration is similar to that of the oracular path planner except for tasks where the robot visits a previously unseen part of the environment and has to re-plan upon prediction of a source of navigation failure. An example of such re-planning can be seen around task number 50 in Fig. 6. Table I summarizes the task completion rate (TCR), relative task completion duration (TCD), and the number of avoided navigation failures by CPIP and the Frequentist method for both simulation and real-robot experiments. The TCD statistics are computed only for tasks that were completed by both CPIP and the Frequentist approach. CPIP achieves a significantly higher TCR across all experiments; moreover, CPIP either performs similarly or outperforms the Frequentist

approach in terms of relative TCD. The reduced task completion duration achieved by CPIP is due to proactively predicting and avoiding non-catastrophic failures, e.g. getting stuck behind falsely detected obstacles. The Frequentist approach would experience these failures and although it might complete the task by replanning, it suffers a longer task completion duration due to performing recovery actions. CPIP outperforms the Frequentist in terms of the relative TCD in the real-robot maze environment, where the distance traveled by the robot in each task is shorter compared to the other experiments, hence the relative task completion delay caused by performing recovery actions by the Frequentist approach is significantly larger. Moreover, in the real-robot large-scale experiment the relative TCD is very similar for CPIP and the Frequentist since there exist no sources of non-catastrophic failures in the environment. Fig. 7 illustrates snapshots of the test environments and highlights the different sources of navigation failures encountered by the robot, which includes different types of texture-less obstacles as well as reflective surfaces.

D. ABLATION STUDY

In order to evaluate the importance of introspective perception in the pipeline of CPIP, we conduct an ablation study. We train a classifier that instead of leveraging the extracted information by introspective perception, directly receives the raw captured RGB images as input and outputs the probability of each class of failure occurring in a specified time window in the future. We use a convolutional neural network with the AlexNet architecture

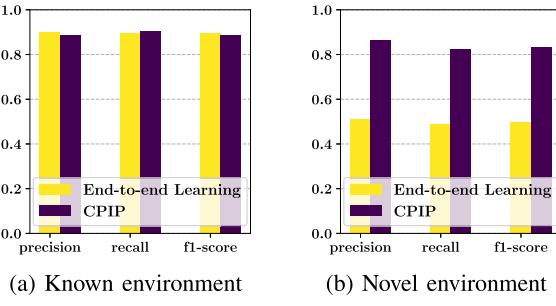


Fig. 8. Results of the navigation failure prediction for CPIP vs. an end-to-end classifier that does not use introspective perception (a) in a previously seen environment and (b) in a novel environment.

similar to that used in prior work [22] for predicting failures of perception.

We train the classifier on the same simulation dataset used for training CPIP and we compare the performance of both methods in predicting navigation failures both in a previously unseen environment—the same test dataset described in Section VII-B—as well as in a new set of deployments of the agent in the training environment. Fig. 8 shows the average precision, recall, and f1-score metrics over all classes, i.e. two classes of failures and a no-failure class, for both CPIP and the end-to-end classifier. While both methods perform similarly good in a previously seen environment, CPIP significantly outperforms the alternative classifier in the novel environment. Leveraging the extracted features by introspective perception simplifies the learning task and allows CPIP to achieve better generalizability given the same amount of training data. This is specifically a benefit for task-level failure prediction, where the volume of training data is limited due to the costly nature of acquiring data from examples of robot failures.

VIII. CONCLUSION

In this paper, we introduced CPIP, a framework for integrating introspective perception with path planning in order to learn to reduce robot navigation failures in the deployment environment and with limited amount of training data. We empirically demonstrated that by leveraging introspective perception CPIP can learn a navigation competence predictor model that generalizes to novel environments and results in significantly reduced frequency of navigation failures. CPIP currently addresses the problem of robot global path planning on a coarse navigation map of the environment. As future directions, the CPIP framework can be extended to support competence-aware local motion planning as well as high-level task planning for mobile robots.

REFERENCES

- [1] N. Hawes *et al.*, “The strands project: Long-term autonomy in everyday environments,” *IEEE Robot. Automat. Mag.*, vol. 24, no. 3, pp. 146–156, Sep. 2017.
- [2] G. Costante, J. Delmerico, M. Werlberger, P. Valigi, D. Scaramuzza, “Exploiting photometric information for planning under uncertainty,” in *Robotics research*, Springer, pp. 107–124, 2018.
- [3] S. Rabiee and J. Biswas, “IVOA: Introspective vision for obstacle avoidance,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 1230–1235.
- [4] S. Rabiee and J. Biswas, “IV-SLAM: Introspective vision for simultaneous localization and mapping,” in *Proc. Conf. Robot Learn.*, 2020, pp. 1100–1109.
- [5] R. Bajcsy, “Active perception,” *Proc. IEEE*, vol. 76, no. 8, pp. 966–1005, Aug. 1988.
- [6] J. Aloimonos, I. Weiss, and A. Bandyopadhyay, “Active vision,” *Int. J. Comput. Vis.*, vol. 1, pp. 333–356, 1988.
- [7] E. Krotkov, “Focusing,” *Int. J. Comput. Vis.*, vol. 1, no. 3, pp. 223–237, 1988.
- [8] M. Krainin, B. Curless, and D. Fox, “Autonomous generation of complete 3 d object models using next best view manipulation planning,” in *Proc. Int. Conf. Robot. Automat.*, 2011, pp. 5031–5037.
- [9] S. A. Sadat, K. Chutskoff, D. Jungic, J. Wawerla, and R. Vaughan, “Feature-rich path planning for robust navigation of mavs with monoslam,” in *Proc. Int. Conf. Robot. Automat.*, 2014, pp. 3870–3875.
- [10] X. Deng, Z. Zhang, A. Sintov, J. Huang, and T. Bretl, “Feature-constrained active visual SLAM for mobile robot navigation,” in *Proc. Int. Conf. Robot. Automat.*, 2018, pp. 7233–7238.
- [11] A. M. Jasour and B. C. Williams, “Risk contours map for risk bounded motion planning under perception uncertainties,” in *Robotics: Sci. Syst.*, 2019, doi: 10.15607/RSS.2019.XV.056.
- [12] S. Shah, D. Dey, C. Lovett, and A. Kapoor, “Airsim: High-fidelity visual and physical simulation for autonomous vehicles,” in *Proc. Field Service Robot.*, 2018, pp. 621–635.
- [13] F. S. Barbosa, B. Lacerda, P. Duckworth, J. Tumova, and N. Hawes, “Risk-aware motion planning in partially known environments,” in *Proc. IEEE Conf. Decis. Control*, 2021.
- [14] J. J. Chung, A. J. Smith, R. Skeele, and G. A. Hollinger, “Risk-aware graph search with dynamic edge cost discovery,” *Int. J. Robot. Res.*, vol. 38, no. 2–3, pp. 182–195, 2019.
- [15] D. M. Saxena, V. Kurtz, and M. Hebert, “Learning robust failure response for autonomous vision based flight,” in *Proc. Int. Conf. Robot. Automat.*, 2017, pp. 5824–5829.
- [16] C. Gurau, D. Rao, C. H. Tong, and I. Posner, “Learn from experience: Probabilistic prediction of perception performance to avoid failure,” *Int. J. Robot. Res.*, vol. 37, no. 9, pp. 981–995, 2018.
- [17] B. Lacerda, D. Parker, and N. Hawes, “Optimal and dynamic planning for markov decision processes with co-safe ltl specifications,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 1511–1516.
- [18] T. Krajná, J. P. Fentanes, J. M. Santos, and T. Duckett, “Fremen: Frequency map enhancement for long-term mobile robot autonomy in changing environments,” *IEEE Trans. Robot.*, vol. 33, no. 4, pp. 964–977, Aug. 2017.
- [19] T. Vintr, Z. Yan, T. Duckett, and T. Krajná, “Spatio-temporal representation for long-term anticipation of human presence in service robotics,” in *Proc. Int. Conf. Robot. Automat.*, 2019, pp. 2620–2626.
- [20] C. Basich, J. Svegliato, K. H. Wray, S. Witwicki, J. Biswas, and S. Zilberman, “Learning to optimize autonomy in competence-aware systems,” in *Proc. 19th Int. Conf. Auton. Agents MultiAgent Syst.*, 2020, pp. 123–131.
- [21] P. Zhang, J. Wang, A. Farhadi, M. Hebert, and D. Parikh, “Predicting failures of vision systems,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 3566–3573.
- [22] S. Daffry, S. Zeng, J. A. Bagnell, and M. Hebert, “Introspective perception: Learning to predict failures in vision systems,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1743–1750.
- [23] K. Pulli, A. Baksheev, K. Konyakov, and V. Eruhimov, “Real-time computer vision with opencv,” *Commun. ACM*, vol. 55, no. 6, pp. 61–69, 2012.