

Learning Perceptual Hallucination for Multi-Robot Navigation in Narrow Hallways

Jin-Soo Park¹, Xuesu Xiao^{2,3}, Garrett Warnell^{4,5}, Harel Yedidsion⁴, and Peter Stone^{4,6}

Abstract—While current systems for autonomous robot navigation can produce safe and efficient motion plans in static environments, they usually generate suboptimal behaviors when multiple robots must navigate together in confined spaces. For example, when two robots meet each other in a narrow hallway, they may either turn around to find an alternative route or collide with each other. This paper presents a new approach to navigation that allows two robots to pass each other in a narrow hallway without colliding, stopping, or waiting. Our approach, Perceptual Hallucination for Hallway Passing (PHHP), learns to synthetically generate virtual obstacles (i.e., *perceptual hallucination*) to facilitate passing in narrow hallways by multiple robots that utilize otherwise standard autonomous navigation systems. Our experiments on physical robots in a variety of hallways show improved performance compared to multiple baselines.

I. INTRODUCTION

One of the grand goals of the robotics community is to safely and reliably deploy fully-autonomous mobile robots in common environments over extended periods of time. Indeed, many researchers have moved toward this vision and reported hundreds of hours of unsupervised, collision-free navigation by a single robot [1], [2].

However, long-term deployment of *multiple* autonomous robots in common spaces still remains a difficult task. One reason for this difficulty is that, while conventional navigation systems are good at handling static environments, their performance deteriorates in the presence of dynamic obstacles, e.g., other moving robots. The research community has explored some solutions to this problem, but these solutions typically rely on strict requirements such as a perfectly-controlled space (e.g., a warehouse) or perfect sensing [3], and they cannot guarantee safety in novel environments without employing time-consuming movement schemes such as one robot halting while another moves past [4]. To the best of our knowledge, there are no reports that claim long-term deployment of *multiple* autonomous robots in uncontrolled spaces without human supervision.

Separately, recent work in the navigation community leveraging the concept of *perceptual hallucination* [5]–[7] has demonstrated impressive results in allowing robots to navigate highly constrained spaces successfully. Here, perceptual hallucination refers to the technique of forcing the robot to perceive specific additional virtual obstacles such that

motion plans generated and executed in the presence of these additional obstacles will exhibit certain desired behaviors. One intuitive motivation for such techniques is that the additional obstacles serve as a kind of blinder for the robot by concealing unnecessary (or even distracting) information. To date, however, perceptual hallucination has not been applied in the context of multiple robots or dynamic obstacles.

In this paper, we hypothesize that perceptual hallucination can be used to improve conventional navigation systems in multi-robot and confined settings. In particular, we posit that, by using perceptual hallucination techniques to obscure the presence of moving objects that would otherwise cause these conventional systems to generate suboptimal behavior (e.g., collision or turning around), we can enable multi-robot navigation in confined spaces such as narrow hallways. If true, this would imply that hallucination would allow system designers to solve the multi-robot navigation problem using the same conventional navigation systems that have been thoroughly tested to be stable in static environments.

To investigate our hypothesis, we introduce and evaluate *Perceptual Hallucination for Hallway Passing* (PHHP), a hallucination-based approach to improve a given navigation policy in the setting of multi-robot navigation in narrow hallways. PHHP uses experience gathered in domain-randomized simulation episodes of hallway passing in order to learn the proper size and placement of virtual obstacles so as to enable successful navigation. We investigate the performance and robustness of using PHHP in common hallways with both simulation and real-world experiments, and we find that it can achieve similar performance compared to a leading existing method, Optimal Reciprocal Collision Avoidance (ORCA) [3], while relaxing its assumption of perfect information about the surroundings. We further show that, compared to a rule-based right-lane-following method, PHHP reduces the average delay by 59.41%. Finally, we show that PHHP is robust to the sim-to-real gap, different speeds, detection ranges, and even various hallway shapes and widths.

II. RELATED WORK

The PHHP approach we present is a machine-learning-based solution to the problem of autonomous multi-robot navigation. Therefore, we review here both conventional approaches that have been proposed to solve the multi-robot navigation problem, and also more recent approaches that have specifically incorporated the use of machine learning. We also briefly review recent work on the use of perceptual hallucination in navigation.

¹Department of Electrical and Computer Engineering, The University of Texas at Austin ²Department of Computer Science, George Mason University ³Everyday Robots ⁴Department of Computer Science, The University of Texas at Austin ⁵Army Research Laboratory ⁶Sony AI {js.park@utexas.edu, xiao@gmu.edu, garrett.a.warnell.civ@army.mil, {harel, pstone}@cs.utexas.edu}




Fig. 1: Overview of the proposed Perceptual Hallucination for Hallway Passing approach. (Left) multi-robot hallway passing scenario with an existing navigation system, and (right) how PHHP improves the navigation system with hallucinated sensor readings. PHHP deployment proceeds as follows: (a) two robots try to pass each other in the hallway; (b) one robot detects the other and initiates the hallucination process; (c) the trained PHHP system generates a hallucinated field $\mathcal{H}_\theta$ in the global coordinate system based on the global plan (d) each robot uses its corresponding hallucinated scan l_{θ} and its existing navigation system to handle the rest of the scenario. Note that the hallucination function, $h(l, \mathcal{H})$, computes the depth value of hallucinated readings by taking the *minimum* value between the real and virtual scans.

A. Conventional Approaches

We divide conventional approaches to multi-robot navigation into ones based on coordination between robots and approaches that operate in a fully decentralized manner. Approaches that require tight coordination include the centralized approach deployed by Kiva Systems (now Amazon Robotics) in controlled warehouse spaces [8] and the coordinated task planning [9] or path planning approach [10]. While these approaches can be used successfully, the reliance on coordination limits their use in scenarios without sufficient communication bandwidth and environments with external dynamic obstacles such as third-party robots. Conventional decentralized approaches [11], including ORCA [3], on the other hand, do not require explicit coordination and communication between agents, and have instead focused on modifying single-agent motion planning in an attempt to make them applicable in multi-robot settings. However, these approaches have their own drawbacks, typically exhibiting high sensitivities to errors in dynamic obstacle state estimation, sensor noise, and parameter selection [12], [13].

B. Learning-Based Approaches

Inspired by recent success in machine learning, several in the community have proposed methods that use learning to enable multi-robot navigation [14], [15]. Decentralized end-to-end approaches, i.e., approaches that learn mappings directly from sensor readings to motion commands, have

been shown to be successful in limited settings [16]–[19], but they are typically less successful in novel environments and often suffer from a lack of safety guarantees, e.g., they may not prevent collisions in highly constrained narrow hallways. There have also been some hybrid attempts to combine conventional navigation with machine learning for multi-robot navigation [4], [20], but they have thus far typically resulted in sub-optimal passing behaviors.

C. Perceptual Hallucination

Finally, the concept of perceptual hallucination, which our method also uses, has recently emerged as an effective tool for enabling navigation in highly-constrained spaces [5]–[7] by allowing robots to synthetically modify their own sensor readings for better neural network training, and simplified motion planning during deployment, or both. Despite its success, however, the idea of hallucination has not previously been applied to dynamic scenarios, including the multi-robot scenario that we study here.

III. APPROACH

In this section, we first formulate the multi-robot hallway passing problem. We then describe our solution, Perceptual Hallucination for Hallway Passing (PHHP).

A. Problem Formulation

We consider a hallway passing scenario in which two robots must pass each other in confined hallways that are

barely wide enough to permit collision-free passing. We specifically consider this scenario in four widely observed hallway shapes: I-, L-, T-, and Z-shaped. In this context, let $\vec{p}^1$ and $\vec{p}^2$ denote the two-dimensional global positions of the first and second robot, respectively. We assume the robot is aware of the other robot's approximate position, $\vec{p}^2$, either through sensing or through occasional communication. In IV-E, we evaluate empirically how approximate this awareness can be. Further, let $\mathcal{C}$ denote a set of global points covering the center of the hallway. We assume that each robot is equipped with a two-dimensional LiDAR scanner, and we denote the LiDAR measurements obtained by each robot at time t as $\mathbf{l}_t^1$ and $\mathbf{l}_t^2$. We also assume that both robots use a global planner and a collision-free local navigation system designed for static environments (e.g., ROS `move_base` [21]).

In the context of the above scenario, we seek to investigate whether perceptual hallucination can improve the performance of an existing navigation system to reduce collisions and decrease passing delay. Mathematically, we use h to denote the *hallucination function*, i.e., the sensor reading $\mathbf{l}_{\mathcal{H}} = h(\mathbf{l}, \mathcal{H})$ is modified by transforming a LiDAR scan $\mathbf{l}$ such that it appears as if virtual obstacles specified by an *obstacle field* $\mathcal{H}$ (i.e., the shape and location of hallucinated obstacles) were added to the current environment. To assure safety, $\mathbf{l}_{\mathcal{H}}$ only contains *additional* obstacles, i.e., to compute the depth value at any particular bearing k , the *minimum* value between the real scan, $\mathbf{l}$, and a hallucinated scan corresponding to only obstacles in $\mathcal{H}$, $\mathbf{v}_{\mathcal{H}}$, is sent to the robot. Finally, we only consider here cases in which the hallucinated obstacle field is static.

In order to use perceptual hallucination, we must determine what each robot should use for its hallucinated obstacle field $\mathcal{H}$ to enable better passing. In general, $\mathcal{H}$ could consist of an arbitrary number of obstacles, each with an arbitrary shape. However, to make the problem tractable, we consider only obstacle fields comprised of the union of a set of contiguous circles of some fixed radius. This set of circles is easy to represent while still being able to approximate an arbitrary-shaped obstacle. We denote such obstacle fields as $\mathcal{H}_{\theta}(\mathcal{C})$ parameterized by $\theta = (r, \Delta r, k_{begin}, k_{end})$, where r specifies the radius of each circle, and k_{begin} and k_{end} specify the fraction that divides the distance from the robot to the first and last circle along the hallway center points $\mathcal{C}$ by the detection range of the robot. Between these starting and ending circles, $\mathcal{H}_{\theta}(\mathcal{C})$ contains a new circle every 5cm. Finally, Δr specifies the distance from the center of the circle to the hallway center points, $\mathcal{C}$. PHHP works best with each robot using the same $\mathcal{H}_{\theta}$, but the system can still improve multi-robot navigation ability with different $\mathcal{H}_{\theta}$ as long as $\mathcal{H}_{\theta}$ is valid. See Section IV-F for more detail. An overview of the hallway passing scenario and how perceptual hallucination is applied is illustrated in Figure 1.

In order to quantify which $\mathcal{H}_{\theta}$ is best for hallway passing, we define a hallway-passing cost function. For a given hallway passing episode, we define this cost function to

encourage both fast and safe passing, i.e.,

$$C(\mathcal{H}_{\theta}) = \frac{\text{TTD}_1(\mathcal{H}_{\theta}) + \text{TTD}_2(\mathcal{H}_{\theta})}{2} + c_{coll} * \mathbb{1}_{coll}, \quad (1)$$

where $\text{TTD}_i(\mathcal{H}_{\theta})$ denotes the amount of time (seconds) it takes for robot i to reach its goal using $\mathcal{H}_{\theta}$, and $\mathbb{1}_{coll}$ is a collision indicator function that is 1 if a collision occurred and 0 otherwise. Note that any collision will terminate the episode. In our implementation, we set the collision penalty c_{coll} to 100. With this setup, the problem of finding the best obstacle field to hallucinate for the hallway passing problem becomes one of finding the θ that minimizes this cost, i.e.,

$$\theta^* = \arg \min_{\theta} C(\mathcal{H}_{\theta}). \quad (2)$$

B. Optimal Hallucination

We find θ^* with Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [22], [23], a population-based, black-box optimization method. CMA-ES finds the best (lowest cost) sample within a fixed number of generations. In each generation, a population of N obstacle fields ($\theta_1 \dots \theta_N$) is sampled from a multivariate normal distribution. Each sample, θ_i , is tested in simulation multiple times to evaluate the average cost (Equation 1). The mean and covariance of the next generation is then updated to increase the likelihood of samples with lower costs. Finally, the minimum-cost sample across all generations is returned as θ^* .

To evaluate a particular sample θ when running CMA-ES, we compute $C(\mathcal{H}_{\theta})$ by executing a hallway passing episode with perceptual hallucination in simulation. For each episode, two robots are initialized at each side of an I- or L-shaped hallway and given navigation goals that require them to pass one another. Then each robot begins navigating using its base navigation system. When robots detect one another by communicating, each robot employs PHHP with $\mathcal{H}_{\theta}(\mathcal{C})$ where the hallway center points $\mathcal{C}$ are approximated by the path provided by a global planner that seeks to maximize the margin of the path. The episode ends when both robots have successfully reached their respective goal locations. The amount of time it takes robot i to reach its goal is recorded as TTD_i . Collisions are defined as any time when a robot is in contact with either another robot or a wall.

In order to ensure θ^* is robust to differences between conditions in simulation and those in the real world, we further employ *domain randomization* [24], [25]. That is, we compute the CMA-ES objective for each sample by averaging costs obtained over several simulation episodes, each with randomized starting delay t_i and detection range D_i sampled from uniform random distributions, $\mathcal{U}_{[0, t_{max}]}$ and $\mathcal{U}_{[d_{min}, d_{max}]}$.

The pseudocode of *perceptual hallucination for hallway passing* (PHHP) is given in Algorithm 1.

IV. EXPERIMENTS

In order to evaluate the efficacy of Perceptual Hallucination for Hallway Passing (PHHP), we measure: (1) the performance of PHHP in different hallway shapes, (2) performance of PHHP and alternative approaches with communication

Algorithm 1 Find optimal Hallucination with CMA-ES

Require: $r, \Delta r, k_{begin}, k_{end}$

- 1: CMAES.initialize($r, \Delta r, k_{begin}, k_{end}$)
- 2: $\sigma \leftarrow 0.1$
- 3: best_cost $\leftarrow \infty$
- 4: $\theta^* \leftarrow \text{None}$
- 5: **while** $\sigma \geq \text{threshold}$ **do**
- 6: $\Theta \leftarrow \text{CMAES.generate_samples}()$
- 7: cost $\leftarrow []$
- 8: **for** $k \leftarrow 1$ to N **do**
- 9: $\theta \leftarrow \Theta[k]$
- 10: $t_1, t_2 \leftarrow \mathcal{U}_{[0, t_{max}]}, \mathcal{U}_{[0, t_{max}]}$
- 11: $D_1, D_2 \leftarrow \mathcal{U}_{[d_{min}, d_{max}]}, \mathcal{U}_{[d_{min}, d_{max}]}$
- 12: TTD₁, TTD₂, coll $\leftarrow \text{episode}(\theta, t_1, D_1, t_2, D_2)$
- 13: cost[k] $\leftarrow \frac{\text{TTD}_1 + \text{TTD}_2}{2} + 100 \cdot \mathbb{1}_{\text{coll}}$
- 14: **if** best_cost $\geq \min(\text{cost})$ **then**
- 15: best_cost $\leftarrow \min(\text{cost})$
- 16: $\theta^* \leftarrow \theta$
- 17: **end if**
- 18: **end for**
- 19: CMAES.optimize(Θ , cost)
- 20: $\sigma \leftarrow \text{CMAES.evaluate}()$
- 21: **end while**
- 22: **return** θ^*

noise, and (3) whether PHHP is robust enough to overcome the sim-to-real gap; as well as to deal with different hallway shapes, different characteristics of robots, and heterogeneous obstacle fields.

We compare PHHP to three alternative approaches: a rule-based, right-lane-following baseline; NH_ORCA [3], [13]; and the *halting method* [4].¹ Details of each of these approaches is provided in Section IV-C.

We evaluate the performance of each method using the following metrics:

- Δt : The amount of delay compared to a single robot traversing the same hallway.
- $P_{\text{collision}}$: The probability of collision.
- P_{failure} : The probability that the navigation system fails to generate a valid passing plan, which typically manifests as the robot turning around.

A. Platform

We evaluate PHHP using BWIBots [26], a custom differential-drive robot atop a Segway base. A single BWIBot is 65cm wide and has a maximum linear velocity of 1.0 m/s. The BWIBot is equipped with a front-facing Hokuyo LiDAR sensor with a 170-degree field of view and a maximum range of 20m. For the underlying navigation system, the BWIBot uses the E-Band planner [27] as the local planner, which continually generates a sequence of motion commands that result from planning over a 4m horizon.

¹Referred to as the *adaptive method* in [4].

B. Training

We train PHHP in the widely-used Gazebo [28] simulator since it provides a safe and fast way to collect realistic data. Two types of hallway are used in training: (a) an “I-shaped” straight hallway, and (b) an “L-shaped” hallway corner. Both hallways are 1.6m-wide, a width for which the two 65cm-wide BWIBots can barely pass each other without colliding. Training episodes proceed as described in Section III-B, where the robots spawn at each side of the hallway, 14m apart from one another. As described in Section III, domain randomization is used during simulation training in an effort to ensure that the learned policy works well when deployed in the real world. Specifically, we impose a random starting delay from 0 to 2 seconds for each robot and we randomly set the detection range to a distance between 7 and 9 meters.

We use PHHP to find virtual obstacles for each type of hallway; I-shaped and L-shaped. To accelerate the CMA-ES search, we used $(r, \Delta r, k_{begin}, k_{end}) = (0.5, 0.05, 0.3, 0.6)$ as an initial hypothesis, which, intuitively, represents a virtual obstacle that entirely blocks the left half of the hallway. For each run of CMA-ES, a total of approximately 100 generations occur before the standard deviation of all samples in a generation becomes less than our selected threshold of 0.01. A single generation contains 8 sample configurations, and each configuration is evaluated by the average cost in Equation (1) averaged over 200 domain-randomized episodes. Training takes about 20 hours using a distributed computing cluster with about 150 nodes. The identified configurations of virtual obstacle fields $\mathcal{H}_{\theta^*}$ in two types of hallways are presented in Table I.

TABLE I: Learned configuration of $\mathcal{H}_{\theta^*}$ in various hallways.

train environment	radius	Δr	k_{begin}	k_{end}
I-shape hallway	0.7590	0.7888	0.4845	0.4910
L-shape hallway	0.5122	0.5661	0.4842	0.5001

C. Alternative Approaches

We compare PHHP with three alternative methods; a right-lane-following baseline, NH_ORCA, and the halting method.

The right-lane-following baseline, or simply *baseline* is inspired by the US traffic standard. It is a rule-based algorithm that, upon detection of the other robot, moves the robot into a human-annotated right lane and proceeds in that lane until the two robots pass one another. However, the baseline has two drawbacks; (a) it requires human effort to manually specify the lanes for each hallway in the environment, and (b) since the robot needs to always stay in a narrow lane even when another robot is not present, the average speed of the robot drops significantly.

NH_ORCA [13] is an extension of ORCA [3] that can be applied to differential drive robots by approximating the trajectory of ORCA with differential drive constraints. ORCA finds the optimal collision-avoiding velocity command with minimum deviation from the velocity command that heads directly to the goal. As long as ORCA obtains accurate information about the environment with sufficient frequency




Fig. 2: The performance of PHHP trained in the subscribed environment deployed in various hallways. Each bar consists of 300 test episodes in the given hallway in simulation. Both PHHP with $\mathcal{H}_{\theta_I^*}$ and $\mathcal{H}_{\theta_L^*}$ provide smooth interaction in the I and T-shaped hallway. However, $\mathcal{H}_{\theta_I^*}$ often fails to provide a smooth solution in the L-shaped hallway while $\mathcal{H}_{\theta_L^*}$ does. Note that no collision happened over 1,800 episodes in simulation. Please refer to Table I for the used $\mathcal{H}_{\theta}$.

(i.e., the precise position and velocity of the other robot), it is able to provide collision avoidance behavior with a small delay. However, while ORCA provides excellent performance (we consider it to be an upper bound) in simulation with perfect communication channels, real robots often must rely on noisy communication channels to share position and velocity information, which we will show degrades the performance of ORCA. PHHP, on the other hand, only needs to observe the presence of the other robot once.

Finally, the *halting* method [4] is a system designed for hallway passing in which, when a halting robot detects a potential collision, it immediately moves to the nearest, safe parking spot until the non-halting robot completely passes and then resumes. While this approach can be used to avoid collisions in some hallways settings, the halting behavior itself typically causes the average delay to be high.

D. Testing PHHP in Different Hallway Shapes

We trained PHHP to obtain specialized obstacle fields, $\mathcal{H}_{\theta_I^*}$ and $\mathcal{H}_{\theta_L^*}$, for each I- and L- shaped hallway. Both obstacle fields are tested 300 times in each I-, L-, and T-shaped hallway. Note that any hallway shape not listed in the subscript is a test hallway that was not part of the training scenarios for PHHP. The results can be found in Figure 2, where we can see that the performance of PHHP using $\mathcal{H}_{\theta_L^*}$ outperforms $\mathcal{H}_{\theta_I^*}$ in all three environments, presumably because the wings of the L-shaped hallway can be viewed as an I-shaped hallway with a shorter length. For this reason, we use PHHP with $\mathcal{H}_{\theta_L^*}$ to represent PHHP in further experiments unless otherwise mentioned.

E. Performance in the Presence of Communication Noise

We now seek to investigate the efficacy of PHHP versus alternative approaches. In particular, we are interested in how robust each approach is to communication noise, and so we perform our experiments here in a setting in which an




Fig. 3: The performance analysis of each algorithm in the noisy communication channel; (left) the average delay of each algorithm, and (right) the probability of collision. The average delay of PHHP remains consistently stable, while the delay of NH_ORCA increases rapidly as the noise increases. The right-lane-following baseline and halting methods are also resistant to noise, but the average delay of both algorithms is much higher than PHHP.

artificial noisy channel is imposed using a dropout model. Under this model, messages fail to reach their destination with a specified probability and, when such a failure occurs, each algorithm uses the last received message to perform collision avoidance. Experiments were conducted 30 times in the 1.6m-wide I-shaped hallway in the Gazebo simulator, and the results are shown in Figure 3. PHHP, the baseline, and the halting method show performance that is robust to communication noise since they only require a single detection of the other robot before activating the corresponding collision avoidance behavior. In contrast, the performance of NH_ORCA, which requires continual and flawless detection of the other robot, drops notably as the channel noise increases.

The results indicate that, as expected, the average delay incurred by PHHP becomes lower than that of NH_ORCA when the channel noise exceeds 5%. Additionally, about 10% of NH_ORCA trials ended in collision in high noise settings, while no collisions were reported for PHHP across all noise levels. Taken together, these results suggest that PHHP will fare better than NH_ORCA in real-world deployment settings. Additionally, we note that, while the baseline and halting methods are robust to noise, the raw performance of those methods is far worse than that of PHHP.

F. Robustness Analysis

We investigate the robustness of PHHP in terms of sim-to-real transfer performance in different environments, different characteristics of the robot (i.e., detection range and velocity), and even with different but valid obstacle fields, $\mathcal{H}_{\theta^*}$, used by PHHP. The test setup is as follows. We directly deploy simulation-trained PHHP in the real world and conduct 32 evaluation episodes per condition, each with a specific environment and robot parameters as shown in Table II. “ID” defines the name of each condition, while “Hallway Type” represents the hallway shape and width, where L-shaped hallways are 1.8m wide on one segment and 1.6m wide on the other. Example hallways can be seen in Figure 5. D_i and V_i describe the detection range and velocity of




Fig. 4: Box and whisker plot of real-world experiments with various conditions. Details for each condition can be found in Table II. When the two robots use different conditions, the delay of each robot is plotted separately. Note that PHHP did not incur a single collision or turnaround behavior across all 192 experiments with various conditions.

TABLE II: The configuration used in real world experiments.

ID	Hallway Type	D_1	D_2	V_1	V_2	$\mathcal{H}_1$	$\mathcal{H}_2$
I-shape	1.8m I-shaped	8.0	8.0	1.0	1.0	$\mathcal{H}_L^*$	$\mathcal{H}_L^*$
L-shape	L-shaped	8.0	8.0	1.0	1.0	$\mathcal{H}_L^*$	$\mathcal{H}_L^*$
S-shape	Z-shaped	8.0	8.0	1.0	1.0	$\mathcal{H}_L^*$	$\mathcal{H}_L^*$
DIFF_D	1.8m I-shaped	9.0	7.0	1.0	1.0	$\mathcal{H}_L^*$	$\mathcal{H}_L^*$
DIFF_V	L-shaped	8.0	8.0	1.0	0.6	$\mathcal{H}_L^*$	$\mathcal{H}_L^*$
DIFF_H	1.6m I-shaped	8.0	8.0	1.0	1.0	$\mathcal{H}_L^*$	$\mathcal{H}_I^*$

a robot, respectively, where the subscript i denotes which robot the parameter pertains to. Finally, $\mathcal{H}_i$ represents the specific obstacle field used by PHHP during the experiment as defined in Section IV-D. The robots report their locations to each other over WiFi.

The result of each experiment is shown in Figure 4. Regarding sim-to-real transfer performance, PHHP fares well: the average delay of PHHP in the real I- and L-type corridors is similar to the results obtained in simulation. Interestingly, a Z-shaped hallway only records one second of delay, presumably because the interaction in a Z-shaped hallway happens on the wider side of the hallway. Regarding robustness to different robot characteristics, PHHP resolves the hallway passing problem without collisions or turnarounds, even when robots have significantly different velocities and detection ranges. The results for **DIFF_D** explains why PHHP is robust against noise. PHHP only needs to receive the opponent’s position once every 2 meters or 2 seconds. Finally, regarding robustness to different but valid obstacle fields, the **DIFF_H** results indicate that, as long as individual obstacle fields can resolve the hallway passing problem, a combination of them can also be successful. **Importantly, no collisions or turnarounds were observed** during the entire set of 192 real-world episodes. These results suggest that PHHP is very robust to the sim-to-real gap, environmental changes, different speeds and detection ranges of the robot, and even to using obstacle fields trained in different hallways.




Fig. 5: Photos of real-world test environments; (a) 1.8m I-shaped, (b) 1.6m I-shaped, (c) Z-shaped, and (d) L-shaped.

V. CONCLUSIONS AND FUTURE WORKS

In this paper, we presented Perceptual Hallucination for Hallway Passing (PHHP), a new method that enables multi-robot navigation in constrained spaces. We showed how to find the best obstacle for PHHP to hallucinate for a given environment and navigation policy using CMA-ES. The simulation and real-world deployment results indicate that PHHP achieves comparable performance against ORCA, while removing the assumption that the robot has continuous access to the other robot’s exact position and velocity. Moreover, PHHP outperforms both a right-lane-following baseline and our prior work, the halting method, in terms of delay. Additionally, real-world deployment results experimentally confirm that PHHP, which is trained in simulation, can successfully be deployed in a wide variety of real-world settings, including those in which the size of the hallway is changed, the robots move with different velocities, their perception systems exhibit a different detection range, or the virtual obstacles used by PHHP is different. Despite the successes we presented, PHHP has only been developed and evaluated here with two robots, both using PHHP, and we have only explored using a particular class of hallucinated obstacle fields comprised of a number of circles. Therefore, an important direction for future work is to investigate how to expand PHHP to work with multiple, arbitrarily shaped obstacles in a wider variety of settings.

ACKNOWLEDGMENTS

This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CPS-1739964, IIS-1724157, FAIN-2019844), ONR (N00014-18-2243), ARO (W911NF-19-2-0333), DARPA, GM, Bosch, and UT Austin’s Good Systems grand challenge. Peter Stone serves as the Executive Director of Sony AI America and receives financial compensation for this work. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

REFERENCES

- [1] J. Biswas and M. Veloso, "The 1,000-km challenge: Insights and quantitative and qualitative results," *IEEE Intelligent Systems*, vol. 31, no. 3, pp. 86–96, 2016.
- [2] P. Khandelwal, S. Zhang, J. Sinapov, M. Leonetti, J. Thomason, F. Yang, I. Gori, M. Svetlik, P. Khante, V. Lifschitz, *et al.*, "Bwibots: A platform for bridging the gap between ai and human–robot interaction research," *The International Journal of Robotics Research*, vol. 36, no. 5–7, pp. 635–659, 2017.
- [3] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [4] J.-S. Park, B. Tsang, H. Yedidsion, G. Warnell, D. Kyoung, and P. Stone, "Learning to improve multi-robot hallway navigation," in *4th Conference on Robot Learning (CoRL)*, November 2020.
- [5] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1503–1510, 2021.
- [6] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [7] Z. Wang, X. Xiao, A. J. Nettekoven, K. Umasankar, A. Singh, S. Bommakanti, U. Topcu, and P. Stone, "From agile ground to aerial navigation: Learning from learned hallucination," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021.
- [8] P. R. Wurman, R. D'Andrea, and M. Mountz, "Coordinating hundreds of cooperative, autonomous vehicles in warehouses," *AI magazine*, vol. 29, no. 1, pp. 9–9, 2008.
- [9] Y. Jiang, H. Yedidsion, S. Zhang, G. Sharon, and P. Stone, "Multi-robot planning with conflicts and synergies," *Autonomous Robots*, March 2019.
- [10] F. G. Lopez, J. Abbenseth, C. Henkel, and S. Dörr, "A predictive online path planning and optimization approach for cooperative mobile service robot navigation in industrial applications," in *2017 European Conference on Mobile Robots (ECMR)*. IEEE, 2017, pp. 1–6.
- [11] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, 1998.
- [12] D. Hennes, D. Claes, W. Meeussen, and K. Tuyls, "Multi-robot collision avoidance with localization uncertainty," in *AAMAS*, 2012, pp. 147–154.
- [13] J. Alonso-Mora, A. Breitenmoser, M. Rufli, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed autonomous robotic systems*. Springer, 2013, pp. 203–216.
- [14] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Motion planning and control for mobile robot navigation using machine learning: a survey," *Autonomous Robots*, pp. 1–29, 2022.
- [15] R. Mirsky, X. Xiao, J. Hart, and P. Stone, "Prevention and resolution of conflicts in social navigation—a survey," *arXiv preprint arXiv:2106.12113*, 2021.
- [16] P. Long, W. Liu, and J. Pan, "Deep-learned collision avoidance policy for distributed multiagent navigation," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 656–663, 2017.
- [17] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6252–6259.
- [18] J. Lin, X. Yang, P. Zheng, and H. Cheng, "End-to-end decentralized multi-robot navigation in unknown complex environments via deep reinforcement learning," in *2019 IEEE International Conference on Mechatronics and Automation (ICMA)*. IEEE, 2019, pp. 2493–2500.
- [19] Q. Tan, T. Fan, J. Pan, and D. Manocha, "Deepmnavigate: Deep reinforced multi-robot navigation unifying local & global collision avoidance," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6952–6959.
- [20] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.
- [21] OSRF, "Ros wiki move_base," http://wiki.ros.org/move_base, 2018.
- [22] N. Hansen, S. Muller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary Computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [23] N. Hansen, "The cma evolution strategy: A tutorial," *arXiv preprint arXiv:1604.00772*, 2016.
- [24] N. Jakobi, "Evolutionary robotics and the radical envelope-of-noise hypothesis," *Adaptive Behavior*, vol. 6, no. 2, pp. 325–368, 1997.
- [25] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017.
- [26] P. Khandelwal, S. Zhang, J. Sinapov, M. Leonetti, J. Thomason, F. Yang, I. Gori, M. Svetlik, P. Khante, V. Lifschitz, J. K. Aggarwal, R. Mooney, and P. Stone, "Bwibots: A platform for bridging the gap between ai and human–robot interaction research," *The International Journal of Robotics Research*, 2017. [Online]. Available: <http://www.cs.utexas.edu/users/ai-lab/?khandelwal:ijrr17>
- [27] S. Quinlan and O. Khatib, "Elastic bands: Connecting path planning and control," in *[1993] Proceedings IEEE International Conference on Robotics and Automation*. IEEE, 1993, pp. 802–807.
- [28] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(IEEE Cat. No. 04CH37566), vol. 3. IEEE, 2004, pp. 2149–2154.