

# APPLI: Adaptive Planner Parameter Learning From Interventions

Zizhao Wang<sup>\*1</sup>, Xuesu Xiao<sup>\*2</sup>, Bo Liu<sup>\*2</sup>, Garrett Warnell<sup>2,3</sup>, and Peter Stone<sup>2,4</sup>

**Abstract**—While classical autonomous navigation systems can typically move robots from one point to another safely and in a collision-free manner, these systems may fail or produce suboptimal behavior in certain scenarios. The current practice in such scenarios is to manually re-tune the system’s parameters, e.g. max speed, sampling rate, inflation radius, to optimize performance. This practice requires expert knowledge and may jeopardize performance in the originally good scenarios. Meanwhile, it is relatively easy for a human to identify those failure or suboptimal cases and provide a teleoperated *intervention* to correct the failure or suboptimal behavior. In this work, we seek to learn from those human interventions to improve navigation performance. In particular, we propose *Adaptive Planner Parameter Learning from Interventions* (APPLI), in which multiple sets of navigation parameters are learned during training and applied based on a confidence measure to the underlying navigation system during deployment. In our physical experiments, the robot achieves better performance compared to the planner with static default parameters, and even dynamic parameters learned from a full human demonstration. We also show APPLI’s generalizability in another unseen physical test course, and a suite of 300 simulated navigation environments.

## I. INTRODUCTION

Decades of research has been devoted to developing mobile robot navigation systems that are capable of moving a robot safely from one point to another in obstacle-occupied spaces without collisions. Classical navigation systems, such as Elastic-Bands [1] or Dynamic Window Approach (DWA) [2], have been robustly deployed on mobile robots with verifiable guarantees of safety and explainability and are able to achieve optimal navigation in most cases.

However, in some situations, those classical navigation systems fail or suffer from suboptimal behaviors (Fig. 1). For example, the robot may not be able to find a feasible action in highly-constrained spaces [3], [4], or may drive unnecessarily slowly in open spaces [5]. The current solution to these problems is to manually re-tune the parameters of the underlying navigation system (e.g. max speed, sampling rate, inflation radius) to correct the failure cases or suboptimal behaviors in those places. This re-tuning process not only requires expert knowledge onsite during deployment, but also runs the risk of the re-tuned parameters targeted at the failed




Fig. 1. While classical navigation systems perform well in most places (green), they may fail (red) or suffer from suboptimal behavior (yellow) in others. APPLI utilizes human interventions in these two scenarios (we name them Type A and Type B interventions, respectively) to learn adaptive planner parameters and, based on a confidence measure, uses them during deployment.

or suboptimal comprising performance in the originally good scenarios.

Meanwhile, even a non-expert user (i.e., one who is not familiar with the inner workings of the underlying navigation system) can easily identify the situations where the robot fails or performs suboptimally by watching, and then can intervene by teleoperating the vehicle. In this work, we utilize those human interventions to improve future autonomous navigation in those troublesome places, while maintaining good performance in others.

In particular, we introduce *Adaptive Planner Parameter Learning from Interventions* (APPLI). With a set of teleoperated human interventions, APPLI learns a set of navigation parameters, which are selected dynamically to eliminate failures or suboptimal behaviors during deployment. To assure the learned parameters will not jeopardize navigation performance in other places, the robot only uses the learned parameters when it is confident that they will benefit the current navigation. In our experiments, APPLI learns from interventions in a real-world navigation task. We test APPLI in the same training and another unseen physical environment. More than twenty thousands simulation trials are conducted in unseen environments to test APPLI’s generalizability. Our results show APPLI can improve upon default navigation performance and can generalize well to unseen environments, indicating that interventions are a uniquely valuable form of human interaction for building navigation systems.

## II. RELATED WORK

In this section, we review existing work on adaptive planner parameters, learning from intervention, and uncertainty measurement in deep learning.

\*Equal Contribution

<sup>1</sup>Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, Texas 78712 zizhao.wang@utexas.edu

<sup>2</sup>Department of Computer Science, University of Texas at Austin, Austin, Texas 78712 {xiao, bliu, pstone}@cs.utexas.edu

<sup>3</sup>Computational and Information Sciences Directorate, Army Research Laboratory, Austin, Texas 78712 garrett.a.warnell.civ@mail.mil

<sup>4</sup>Sony AI

Classical navigation methods enjoy safety, explainability, and stable generalization to new environments. However, when facing new environments, they still need a great deal of tuning, which often requires expert robotics knowledge [6], [7]. Prior work has considered automated parameter tuning, e.g., finding trajectory optimization weights [8] for the DWA planner [2], or designing novel systems that can leverage gradient descent to match expert demonstrations [9]. Specifically, Xiao *et al* [10] adopted black-box optimization to automatically map a robot’s local observation to the optimal planner parameters via learning from human demonstration. While this technique can be applied to any parameter-based planner, it is not expected to generalize well in environments not seen in the demonstration. In contrast, APPLI only requires a few, short, local interventions when classical navigation does not perform well, instead of a demonstration of the full trajectory. APPLI also includes confidence estimation over candidate planner parameters during deployment in unseen environments. Notably, in the worst case our method reduces to the planner with default parameters, rather than a poorly chosen parameter set, and therefore enjoys better generalization.

### B. Learning from Intervention

Due to the cost of providing full demonstrations, human intervention is a popular approach to providing minimal guidance for learning. It has been widely used in reinforcement learning [11], [12] and imitation learning [13]–[17]. Learning from intervention essentially focuses the agent on learning from its mistakes, thus improving the data efficiency and reducing the demonstration cost. In this work, we leverage the benefits of learning from intervention to enable robust robot navigation, and further categorize interventions based on the expert’s estimation of the necessity of such interventions.

### C. Measuring Uncertainty in Deep Learning

Recent advances in deep learning have provided a family of tools for measuring the uncertainty in a deep model’s prediction. There are mainly three types of approaches. (1) Bayesian Neural Networks (BNN) represent distributions over network weights and the prediction uncertainty is indirectly inferred via weight uncertainty [18]; (2) Deep Ensemble (DE) uses the outputs from multiple networks, each trained with partial data, as a Monte-Carlo estimator for uncertainty [19]. A specific example is Dropout learning [20]; (3) Other methods train a single network with stationary weights but directly model the predictions in terms of a distribution. Evidential Deep Learning (EDL) [21] is one particular method that models a discrete class of predictions with a Dirichlet distribution. We incorporate EDL into APPLI due to its simplicity and efficiency in terms of both time and space complexity, compared to methods from the other two approaches, which is essential for robot learning.

In this section, we introduce our method, APPLI, which has two novel features: (1) compared with Learning from Demonstration that requires demonstration of the whole task, APPLI only needs a few interventions in challenging scenarios where the default navigation system does not work well; (2) with a confidence measure on candidate parameters learned from interventions, our method knows when to switch back to the default parameters. This confidence measure enables APPLI to generalize well to unseen environments.

### A. Problem Definition

We denote a classical parameterized navigation system as  $G : \mathcal{X} \times \Theta \rightarrow \mathcal{A}$ , where  $\mathcal{X}$  is the state space of the robot (e.g. goal, sensor observations),  $\Theta$  is the parameter space for  $G$  (e.g. max speed, sampling rate, inflation radius), and  $\mathcal{A}$  is the action space (e.g. linear and angular velocities). During deployment, the navigation system repeatedly estimates state  $x$  and takes action  $a$  calculated as  $a = G(x; \theta)$ . Typically, the default parameter set  $\theta$  is tuned by a human designer trying to achieve good performance in most environments. However, being good at everything often means being great at nothing:  $\theta$  usually exhibits suboptimal performance in some situations and may even fail (is unable to find feasible motions, or crashes into obstacles) in particularly challenging ones.

To mitigate this problem, a human can supervise the navigation system’s performance at state  $x$  by observing its action  $a$  and judging whether (s)he should intervene. Here, we consider two types of interventions. A type A intervention is one in which the system performs so poorly that the human *must* intervene (e.g. imminent collision or a signal for help). A type B intervention is one in which a human *might* intervene in order to improve otherwise suboptimal performance (e.g., driving too slowly in an open space). For the  $i^{\text{th}}$  intervention, we assume that the human resets the robot to the position where the failure or suboptimal behavior first occurred and then gives a short teleoperated intervention  $I_i = \{x_t, a_t\}_{t=1}^{T_i}$  of length  $T_i$ , where  $x_{1:T_i}$  is the trajectory starting from the reset state induced by intervention actions  $a_{1:T_i}$ . As this short demonstration shows a cohesive navigation behavior in a specific segment of the environment (open space, narrow corridor, etc), we refer to the segment as a *context*  $c_i$  and denote the space of contexts as  $\mathcal{C}$ . Given  $N$  interventions  $I_{1:N}$ , APPLI finds (1) a mapping  $M : \mathcal{C} \rightarrow \Theta$  that determines the parameter set  $\theta_i$  for each intervention context  $c_i$ , and (2) a parameterized predictor  $B_\phi : \mathcal{X} \rightarrow \mathcal{C}$  that determines to which context (if any) the current state  $x$  belongs.

### B. Parameter Learning

After collecting a set of  $N$  interventions  $I_{1:N}$ , for each  $I_i$ , we learn a set of navigation parameters  $\theta_i$  that can best imitate the demonstrated correction. To find such parameters, we use the same training procedure as in the approach by Xiao *et al* [10], i.e., we use Behavior Cloning to minimize the difference between the actions from the human and those

generated by the navigation system with new parameters  $\theta_i$ . To be specific,

$$\theta_i = \arg \min_{\theta} \sum_{(x,a) \in I_i} \|a - G(x; \theta)\|_{\lambda}, \quad (1)$$

where  $\|d\|_{\lambda} = \sum \lambda_i d_i^2$  is the norm of the action difference with  $\lambda$  weighting the different action dimensions (in our case, linear and angular velocity,  $v$  and  $\omega$ ). The loss in Eqn. (1) is minimized with a black-box optimizer, such as CMA-ES [22]. After identifying parameters in each context, the mapping  $M$  is simply  $M(i) = \theta_i$ .

### C. Confidence-Based Context Prediction

So far, we have described how to learn multiple parameter sets  $\theta_{1:N}$  from human interventions  $I_{1:N}$  in contexts  $c_{1:N}$ . In order to select the correct parameters at deployment time, we must also determine if the current state  $x_t$  falls into any one of the collected intervention contexts  $c_i$ . If such a determination can be made, then we direct the robot to use the parameter set  $\theta_i$  to avoid making the same mistake as before. If it cannot be determined that  $x_t$  belongs to a particular intervention context, then we direct the robot to use the default parameters  $\bar{\theta}$ , as they are optimized for most cases and are expected to generalize better than any parameter set learned for a specific scenario. In our system, the determination above is made using a predictor,  $B_{\phi}$ . To train this predictor, we first use the collected interventions to build a dataset,  $\{\{x_t, c_i\}_{t=1}^{T_i}\}_{i=1}^N$ , and train an intermediate classifier  $f_{\phi}(x)$  with parameter set  $\phi$  using the Evidential Deep Learning method (EDL) [21]. A feature of EDL is that it supplies both a predicted label prediction and a confidence in that prediction  $u_i \in (0, 1]$ , i.e.,

$$f_{\phi}(x_i) = (c_i, u_i). \quad (2)$$

After training  $f_{\phi}$  and during deployment, we can build a confidence-based classifier  $g_{\phi}$  as

$$g_{\phi}(x_i) = c_i \mathbb{1}(u_i \geq \epsilon_u), \quad (3)$$

where  $\epsilon_u$  is the threshold on confidence and  $\mathbb{1}$  is the indicator function. For state  $x_i$ ,  $g_{\phi}$  determine its context from  $N + 1$  contexts ( $N$  intervention contexts and 1 default context). If  $u_i \geq \epsilon_u$ , it suggests the classifier  $f_{\phi}$  is confident and  $g_{\phi}$  predicts  $c_i$ . Otherwise, when  $f_{\phi}$  is unsure about its prediction,  $c_i \mathbb{1}\{u_i \geq \epsilon_u\} = 0$ . In this case,  $g_{\phi}$  believes the current state  $x_i$  is not similar to any intervention context and instead classifies  $x_i$  as the default context. For this default context labeled as  $c_i = 0$ , navigation utilizes the default navigation parameters  $\bar{\theta}$  (i.e., we set  $M(0) = \bar{\theta}$ ).

Then we define our context predictor  $B_{\phi}$  as:

$$B_{\phi}(x_t) = \text{mode}(\{g_{\phi}(x_i)\}_{i=t-w+1}^t). \quad (4)$$

To avoid a context estimation  $c_t$  that changes frequently (e.g. caused by  $g_{\phi}$ 's wrong classifications),  $B_{\phi}$  acts as a mode filter with window length  $w$  and chooses the context  $c_t$  that the majority of classifications agree with over the past  $w$  time steps.

### D. APPLI

Putting together all the components presented above, the entire APPLI pipeline is summarized in Alg. 1. In the training stage, it collects  $N$  interventions from a human supervisor (line 1), and then learns corresponding navigation parameters  $\theta_{1:N}$  i.e., the mapping  $M$  (lines 2-3) and a context predictor  $B_{\phi}$  (lines 5-6). During deployment, we use  $M(B_{\phi}(x_t))$  to select the parameters for the navigation system at time  $t$  (lines 8-10).

---

#### Algorithm 1: APPLI

---

##### 1 Training

**Input:** human interventions  $I_{1:N} = \{\{x_t, a_t\}_{t=1}^{T_i}\}_{i=1}^N$ , navigation system  $G$ , parameter space  $\Theta$ .

##### 2 for $i = 1, \dots, N$ do

3 find parameter  $\theta_i$  for context  $i$  using Eqn. (1).

##### 4 end

5 train the context classifier  $f_{\phi}$  on  $\{\{x_t, c_i\}_{t=1}^{T_i}\}_{i=1}^N$ .

6 build mapping  $M(i) = \theta_i$  and context predictor  $B_{\phi}(x)$ .

##### 7 Deployment

**Input:** navigation system  $G$ , parameter mapping  $M$ , context predictor  $B_{\phi}(x)$ , confidence threshold  $\epsilon_u$ , fallback parameters  $\bar{\theta}$ .

##### 8 for $t = 1, \dots$ do

9 identify the current context  $c_t = B_{\phi}(x_t)$  with confidence threshold  $\epsilon_u$ .

10 navigate with  $G(x_t, M(c_t))$ .

##### 11 end

---

## IV. EXPERIMENTS

In our experiments, we aim to show that APPLI can improve navigation performance by learning from only a few interventions and, with the confidence measurement, that the overall system can generalize well to unseen environments. We apply APPLI on a ClearPath Jackal ground robot in a physical obstacle course. Navigation performance learned through APPLI is then tested both in the same training environment, and also in another unseen physical test course. Furthermore, to investigate generalizability, we test the learned systems on a benchmark suite of 300 unseen simulated navigation environments.

### A. APPLI Implementation

Our Jackal is a differential-drive robot equipped with a Velodyne LiDAR that we use to compute a 720-dimensional planar laser scan with a 270° field of view. The robot uses the Robot Operating System `move_base` navigation stack with Dijkstra's global planner and the default DWA local planner, which works in most situations, but fails or behaves suboptimally in others (see Fig. 1).

During data collection, one of the authors (the *intervener*) follows the robot through the test course and intervenes when necessary, reporting if the intervention is to drive the robot out of a failure case (Type A) or to correct a suboptimal behavior (Type B). The four interventions are shown in Fig. 1: before the two Type A interventions (shown in red), the

default system (DWA with  $\bar{\theta}$ ) fails to plan feasible motions and starts recovery behaviors (rotates in place and moves backward); before the two Type B interventions (shown in yellow), the robot drives unnecessarily slowly in a relatively open space and enters the narrow corridor with unsmooth motions. For every intervention, the intervener stops the robot, drives it back to where they deem the failure or sub-optimal behavior to have begun, and then provides recorded teleoperation  $I$  that avoids the problematic behavior. To compare the performance learned from interventions and learned from a full demonstration, we also collect extra demonstrations for those places where the default planner already works very well (denoted in green in Fig. 1).

This set of interventions comprises the input  $I_{1:N} = \{\{x_t, a_t\}_{t=1}^T\}_{i=1}^N$  to Alg. 1, where  $x_t$  is all the sensory data fed into the `move_base` stack,  $G$ , and  $a_t$  is the linear and angular velocity ( $v$  and  $\omega$ ) from teleoperation. The default and learned parameters are shown in Tab. I, including those learned from Type A and B interventions (A1, A2 and B1, B2), and the extra demonstrations (D1, D2).

TABLE I

Default and Learned Planner Parameters:

*max\_vel\_x* (v), *max\_vel\_theta* (w), *vx\_samples* (s), *vtheta\_samples* (t), *occdist\_scale* (o), *pdist\_scale* (p), *gdist\_scale* (g), *inflation\_radius* (i)

	v	w	s	t	o	p	g	i
DEF.	0.50	1.57	6	20	0.10	0.75	1.00	0.30
A1	0.26	2.00	13	44	0.57	0.76	0.94	0.02
A2	0.22	0.87	13	31	0.30	0.36	0.71	0.30
B1	1.91	1.70	10	47	0.08	0.71	0.35	0.23
B1	0.72	0.73	19	59	0.62	1.00	0.32	0.24
D1	0.37	1.33	9	6	0.95	0.83	0.93	0.01
D2	0.31	1.05	17	20	0.45	0.61	0.22	0.23

### B. Physical Experiments

After the training in Alg. 1 with the collected interventions, we deploy the learned mapping  $M$  and context predictor  $B_\phi$  on the `move_base` navigation stack  $G$ . We use a confidence threshold  $\epsilon_u = 0.8$ .

We first deploy APPLI in the same training physical environment (Fig. 1). We compare the performance of the DWA planner with default parameters, APPLI learned only with Type A interventions, APPLI learned with Type A and Type B interventions, and APPLI learned with a *full demonstration* (which is basically the APPLD framework [10] enhanced by the confidence measure). The motivation for the variation of APPLI learned only with Type A interventions is to study the effect of an unfocused or inexperienced human intervener. In this case, the human would still conduct all Type A interventions, as those mistakes are severe and easy to identify—some robots may even actively ask for help (e.g. by starting recovery behaviors). However, the human may fail to conduct Type B interventions as (s)he is not paying attention, or isn't equipped with the knowledge to

TABLE II  
Traversal Time in Training Environment

Default	Type A	Type A+B	Full Demo
134.0±60.6s	77.4±2.8s	70.6±3.2s	78.0±2.7s




Fig. 2. APPLI Running in an Unseen Physical Environment

identify suboptimal behaviors. For all methods, we run five trials each and report the mean and standard deviation of the traversal time in Tab. II. If the robot gets stuck, we introduce a penalty value of 200 seconds. We also deploy the same sets of variants in an unseen physical environments (Fig. 2) and report the results in Tab. III.

For both training and unseen environments, Type A interventions alone significantly improve upon the default parameters, by correcting all recovery behaviors such as rotating in place or driving backwards, and eliminating all failure cases. Adding Type B interventions further reduces traversal time, since the robot learns to speed up in relatively open spaces and to execute smooth motion when the tightness of the surrounding obstacles changes. All the interventions are able to improve navigation in both training and unseen environments, suggesting APPLI's generalizability. Surprisingly, in both environments, APPLI learned from only Type A and Type B interventions can even outperform APPLI learned from an entire demonstration. One possible reason for this better performance from fewer human interactions is the additional human demonstrations may be suboptimal, especially since they are collected in places where the default navigation system was already deemed to have performed well. For example, in the full demonstration, we find the human intervener is more conservative than the default navigation system and drives slowly in some places. Hence, learning from these suboptimal behaviors introduces suboptimal parameters and consequently worse performance in contexts similar to that intervention.

TABLE III  
Traversal Time in Unseen Environment

Default	Type A	Type A+B	Full Demo
109.2±50.8s	71.0±0.7s	59.0±0.7s	62.0±2.0s

TABLE IV  
Percentage of Simulation Environments that Method 1 is Significantly Worse than Method 2 in Terms of Traversal Time  
(Methods are listed in order of increasing performances. Results mentioned in experiment analysis are bold for better identification)

		Method 2						
		APPLI (A)	DWA	APPLI (A+c)	APPLI (A+B+D+c)	APPLI (A+B+D)	APPLI (A+B+c)	APPLI (A+B)
Method 1	APPLI (A)	0	50	<b>53</b>	62	63	68	66
	DWA	10	0	6	<b>33</b>	40	<b>44</b>	47
	APPLI (A+c)	6	4	0	31	37	45	45
	APPLI (A+B+D+c)	5	7	11	0	25	<b>31</b>	33
	APPLI (A+B+D)	5	7	7	10	0	<b>21</b>	21
	APPLI (A+B+c)	3	3	4	3	5	0	9
	APPLI (A+B)	2	5	5	6	4	6	0




Fig. 3. Navigation Trials in Example Environments with Low, Medium, and High Difficulty Levels

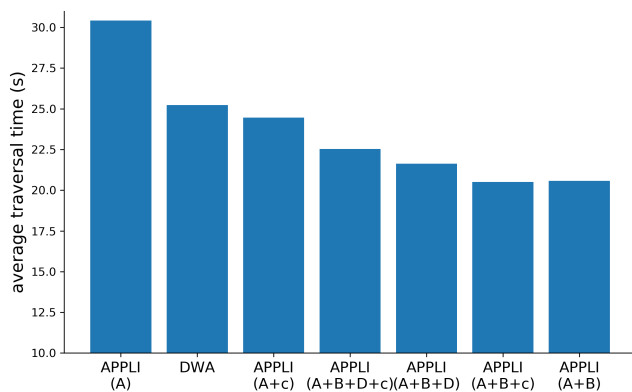


Fig. 4. Average Performance in 300 Simulation Environments under 12 Different Runs

### C. Simulated Experiments

To further test APPLI's generalizability to unseen environments, we test our method with multiple variations and compare them with DWA on the Benchmark for Autonomous Robot Navigation (BARN) dataset [23]. The benchmark dataset consists of 300 simulated navigation environments generated using Cellular Automata, ranging from easy ones with a lot of open spaces to challenging ones where the robot needs to get through dense obstacles. Navigation trials in three example environments with low, medium, and high difficulty levels are shown in Fig. 3. Using the same training data collected from the physical environment shown in Fig. 1, we test the following seven variants:

- APPLI learned from *Type A and B* inventions *with* confidence measure, denoted as APPLI (A+B+c).
- APPLI learned from *Type A and B* inventions *without* confidence measure, i.e., APPLI (A+B).

- APPLI learned from only *Type A* interventions *with* confidence measure, i.e., APPLI (A+c).
- APPLI learned from only *Type A* interventions *without* confidence measure, i.e., APPLI (A).
- APPLI learned from *full demonstration with* confidence measure, i.e., APPLI (A+B+D+c).
- APPLI learned from *full demonstration without* confidence measure, i.e., APPLI (A+B+D).
- the DWA planner with default parameters.

Testing these variations aims at studying the effect of learning from different modes of interventions caused by different degrees of human attention and experience levels, i.e. imperative interventions (A), optional interventions (A + B), and a full demonstration (A + B + D). They also provide an ablation study for the confidence measure in the EDL context classifier  $f_\phi$ : when deployed without the confidence measure, the robot has to choose among the parameters learned from interventions and never uses the default parameters.

For each method in each simulation environment, we measure the traversal time for 12 different runs (the run is terminated after 50s if the robot gets stuck), resulting in 25200 total navigation trials. The average traversal time for each method in all simulation environments are shown in order of increasing performance in Fig. 4. We then conduct a pair-wise t-test for all methods in order to compute the percentage of environments in which one method (denoted as Method 1) is significantly worse ( $p < 0.05$ ) than another (denoted as Method 2). For better illustration, we also reorder the method by their performance and show the pairwise comparison in Tab. IV.

APPLI (A+B+c) and APPLI (A+B+D+c) outperform DWA: they are significantly better in 44% and 33% of environments respectively and significantly worse in only 3% and 7% of environments than DWA. However, APPLI (A+c) is only significantly better than DWA 6% of the time, which suggests that even though type B inventions only correct suboptimal performances, they are crucial for performance improvement. In detail, as the robot learns to go through narrow passages from Type A interventions, the first Type B intervention further teaches the robot to drive fast in safe open spaces, significantly reducing the traversal time for simulation environments whose beginnings and ends

are relatively open. Meanwhile, from the second Type B intervention, the robot learns how to take sharp turns despite constrained surroundings.

In terms of the effect of confidence, APPLI (A) only selects parameters learned from 2 Type A interventions and never uses the default parameters even when they are more appropriate. Removing confidence greatly harms its performance, making it significantly worse than APPLI (A+c) in 53% of environments. However, APPLI (A+B+c) and APPLI (A+B+D+c), which use more interventions or even the full demonstration to train the parameter mapping  $M$  and context predictor  $B_\phi$  are more confident about their predictions most of the time. As a result, removing confidence in the context predictor doesn't result in a significant difference.

Lastly, a counterintuitive, but similar result as in the physical experiments is that compared with APPLI (A+B+D+c) which uses the full demonstration, APPLI (A+B+c) learned from only Type A and B interventions achieves superior performance by being significantly better than APPLI (A+B+D+c) and APPLI (A+B+D) in 31% and 21% of the environments, respectively. However, similar to the discussions about physical experiments, unnecessary human demonstrations are most likely suboptimal. In this sense, APPLI not only reduces the required human interactions from a full demonstration to only a few interventions, but also reduces the chance of performance degradation caused by suboptimal demonstrations.

## V. CONCLUSIONS

In this work, we introduce APPLI, *Adaptive Planner Parameter Learning from Interventions*. In contrast to most existing end-to-end machine learning for navigation approaches, APPLI utilizes existing classical navigation systems and inherits all their benefits, such as safety and explainability. Furthermore, instead of requiring a full expert demonstration or random exploration based on trial-and-error, APPLI only needs a few interventions, where the default underlying navigation system fails or exhibit poor behavior. It also introduces a confidence measure to assure generalizability in unseen environments. We show APPLI's improved performance in training and unseen physical environments. We further test APPLI's generalizability with 25200 simulated navigation trials in 300 unseen environments. While we allow the intervener to start interventions by rewinding the robot navigation to a state before the failure or suboptimal behavior occurs, an interesting direction for future work is to further investigate interventions without "rewinding", i.e. where the intervener takes over from where the robot fails and drives it forward to a good state.

## ACKNOWLEDGMENTS

This work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, The University of Texas at Austin. LARG research is supported in part by grants from the National Science Foundation (CPS-1739964, IIS-1724157, NRI-1925082), the Office of Naval Research (N00014-18-2243), Future of Life

Institute (RFP2-000), Army Research Office (W911NF-19-2-0333), DARPA, Lockheed Martin, General Motors, and Bosch. The views and conclusions contained in this document are those of the authors alone. 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] 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.
- [2] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [3] X. Xiao, B. Liu, G. Warnell, and P. Stone, "Toward agile maneuvers in highly constrained spaces: Learning from hallucination," *arXiv preprint arXiv:2007.14479*, 2020.
- [4] B. Liu, X. Xiao, and P. Stone, "Lifelong navigation," *arXiv preprint arXiv:2007.14486*, 2020.
- [5] X. Xiao, B. Liu, and P. Stone, "Agile robot navigation through hallucinated learning and sober deployment," *arXiv preprint arXiv:2010.08098*, 2020.
- [6] K. Zheng, "Ros navigation tuning guide," *arXiv preprint arXiv:1706.09068*, 2017.
- [7] X. Xiao, J. Dufek, T. Woodbury, and R. Murphy, "Uav assisted usv visual navigation for marine mass casualty incident response," in *2017 IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 6105–6110.
- [8] D. Teso-Fz-Betoño, E. Zulueta, U. Fernandez-Gamiz, A. Saenz-Aguirre, and R. Martinez, "Predictive dynamic window approach development with artificial neural fuzzy inference improvement," *Electronics*, vol. 8, no. 9, p. 935, 2019.
- [9] M. Bhardwaj, B. Boots, and M. Mukadam, "Differentiable gaussian process motion planning," *arXiv preprint arXiv:1907.09591*, 2019.
- [10] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [11] W. Saunders, G. Sastry, A. Stuhlmüller, and O. Evans, "Trial without error: Towards safe reinforcement learning via human intervention," *arXiv preprint arXiv:1707.05173*, 2017.
- [12] B. Prakash, M. Khatwani, N. Waytowich, and T. Mohsenin, "Improving safety in reinforcement learning using model-based architectures and human intervention," *arXiv preprint arXiv:1903.09328*, 2019.
- [13] V. G. Goecks, G. M. Gremillion, V. J. Lawhern, J. Valasek, and N. R. Waytowich, "Efficiently combining human demonstrations and interventions for safe training of autonomous systems in real-time," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 2462–2470.
- [14] M. Kelly, C. Sidrane, K. Driggs-Campbell, and M. J. Kochenderfer, "Hg-dagger: Interactive imitation learning with human experts," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8077–8083.
- [15] J. Zhang and K. Cho, "Query-efficient imitation learning for end-to-end autonomous driving," *arXiv preprint arXiv:1605.06450*, 2016.
- [16] J. Spencer, S. Choudhury, M. Barnes, M. Schmittle, M. Chiang, P. Ramadge, and S. Srinivasa, "Learning from interventions," in *Robotics: Science and Systems (RSS)*, 2020.
- [17] G. Kahn, P. Abbeel, and S. Levine, "Land: Learning to navigate from disengagements," *IEEE Robotics and Automation Letters*, 2021.
- [18] I. Kononenko, "Bayesian neural networks," *Biological Cybernetics*, vol. 61, no. 5, pp. 361–370, 1989.
- [19] K. Lee, Z. Wang, B. Vlahov, H. Brar, and E. A. Theodorou, "Ensemble bayesian decision making with redundant deep perceptual control policies," in *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 2019, pp. 831–837.
- [20] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*, 2016, pp. 1050–1059.

- [21] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," in *Advances in Neural Information Processing Systems*, 2018, pp. 3179–3189.
- [22] N. Hansen, S. D. Müller, 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] D. Perille, A. Truong, X. Xiao, and P. Stone, "Benchmarking metric ground navigation," *arXiv preprint arXiv:2008.13315*, 2020.