

APPL: Adaptive Planner Parameter Learning

Xuesu Xiao*, Zizhao Wang, Zifan Xu, Bo Liu, Garrett Warnell, Gauraang Dhamankar, Anirudh Nair, Peter Stone

Department of Computer Science, The University of Texas at Austin, Austin, TX 78712, United States of America

ARTICLE INFO

Article history:

Received 2 January 2022

Received in revised form 20 March 2022

Accepted 20 April 2022

Available online 29 April 2022

Keywords:

Mobile robot navigation

Machine learning

Motion planning

ABSTRACT

While current autonomous navigation systems allow robots to successfully drive themselves from one point to another in specific environments, they typically require extensive manual parameter re-tuning by human robotics experts in order to function in new environments. Furthermore, even for just one complex environment, a single set of fine-tuned parameters may not work well in different regions of that environment. These problems prohibit reliable mobile robot deployment by non-expert users. As a remedy, we propose *Adaptive Planner Parameter Learning* (APPL), a machine learning framework that can leverage non-expert human interaction via several modalities – including teleoperated demonstrations, corrective interventions, and evaluative feedback – and also unsupervised reinforcement learning to learn a *parameter policy* that can dynamically adjust the parameters of classical navigation systems in response to changes in the environment. APPL inherits safety and explainability from classical navigation systems while also enjoying the benefits of machine learning, i.e., the ability to adapt and improve from experience. We present a suite of individual APPL methods and also a unifying cycle-of-learning scheme that combines all the proposed methods in a framework that can improve navigation performance through continual, iterative human interaction and simulation training.

© 2022 Elsevier B.V. All rights reserved.

1. Introduction

While engineered systems developed over the past several decades have enabled autonomous mobile robot navigation in certain scenarios, it is typically the case that the parameters of these systems need to be adjusted for each new deployment environment. This process is referred to as “parameter-tuning”, i.e., the process through which a human with expert knowledge of the underlying system and the role of each parameter – along with trial-and-error, experience, and intuition – manually finds the best set of system parameters for a new environment. This traditional setup (1) precludes non-expert users from optimizing the default robot system; (2) assumes a *single* set of parameters will work well for all regions of a specific environment; and (3) completely abandons previously fine-tuned parameters when using a new set of parameters.

In contrast, many humans – even those with little to no robotics experience – can easily teleoperate mobile robots in new environments [1], and even more humans are able to provide simple evaluative feedback (e.g., a numerical assessment of performance) on the robot’s behavior. These *in situ* end-user interactions can provide a rich source of knowledge regarding how the system should behave in deployment environments. We

hypothesize that systems that can capture and leverage these interactions will be able to quickly adapt to their deployment environments, i.e., exhibit robust and reliable navigation system performance.

In this article, we investigate this hypothesis by developing and studying a novel family of algorithms called Adaptive Planner Parameter Learning (APPL). In particular, we present an algorithm that can learn from human demonstrations (APPLD), an algorithm that can learn from human interventions (APPLI), and an algorithm that can learn from human-generated evaluative feedback (APPLE). Additionally, we present an algorithm in this family that can, if available, leverage simulation experience in an unsupervised fashion using reinforcement learning (APPLR). Importantly, the framework under which we develop APPL does not replace existing systems for autonomous navigation, but rather *augments* them by providing a new learned module that adjusts only certain hyperparameters of these systems (e.g., obstacle inflation radius, sampling rate, cost function coefficients, etc.). By adopting this approach, APPL systems inherit the advantages of existing navigation systems (e.g., safety and explainability), while also enjoying the benefits of machine learning methods (e.g., adaptivity and improvement from experience). Moreover, the APPL framework explicitly allows for the possibility of adjusting these hyperparameters “on-the-fly”, i.e., dynamically changing parameters at every time step to achieve robust navigation. Finally, we also present a unifying vision for how our APPL algorithms can be

* Corresponding author.

E-mail address: xiao@cs.utexas.edu (X. Xiao).

used in concert to enable a cycle-of-learning framework for continual system improvement through iterative interaction with simulation and end users.

We demonstrate the efficacy of the proposed algorithms in a series of experiments in which a small ground robot, accompanied by a human, must autonomously navigate in a number of complex, constrained environments. Our results show that leveraging human interaction does indeed allow the robot to quickly adapt to each new environment, learning to overcome failures and other suboptimal behavior to the point where the human is no longer necessary. Moreover, we show that combining the proposed approaches in the unified framework that can also leverage unsupervised simulated training in a series of system iterations exhibits increasing performance with each new deployment.

The remainder of this article is organized as follows. We begin by reviewing related work in terms of learning for navigation and learning from humans in Section 2. As one of the two novel contributions of this article, we formalize the *Adaptive Planner Parameter Learning* (APPL) framework in Section 3, which is an overarching paradigm that encompasses all the following more specialized methods. Specifically, in Sections 4, 5, 6, and 7, we present Adaptive Planner Parameter Learning from Demonstration [2], from Interventions [3], from Evaluative Feedback [4], and from Reinforcement [5], respectively. Note that APPLD and APPLR were introduced in our previous Robotics and Automation Letters articles [2,4], so we only briefly summarize these two methods. For APPLI and APPLR, which have only appeared in conference papers [3,5], we re-formalize these two methods under the overarching APPL framework and present detailed descriptions. As the second novel contribution of this article, we introduce a cycle-of-learning scheme in Section 8, which combines all four individual APPL methods in the unified framework, leveraging demonstration, interventions, evaluative feedback, and reinforcement in different deployment scenarios with different human users, and achieves cyclic and continual improvement in navigation performance. The experiments pertaining to this cycle-of-learning scheme are also new in this article. All experiment videos can be found at <https://www.cs.utexas.edu/~xiao/Research/APPL/APPL.html>.

2. Related work

This section reviews state-of-the-art approaches that have applied machine learning to the problem of mobile robot navigation and related work in terms of using different human interaction modalities to assist machine learning.

2.1. Learning for navigation

Autonomous mobile robot navigation has been a topic of interest to the robotics community for decades [6,7]. For example, the DWA planner [7] samples feasible motion commands within a dynamic window, evaluates each sample using a forward prediction model, and selects the best sample based on a cost function that considers distance to obstacles, deviation from the global path, and progress toward the goal. While providing verifiable guarantees, these classical approaches still require extensive knowledge from robotics expert onsite during deployment, e.g., through in-situ parameter tuning, to adapt to different navigation scenarios [8,9]. For example, max linear and angular velocities and their sampling rates determine if the DWA planner can find good motion commands with limited onboard computation, while its different optimization weights affect the final navigation behavior. All experiments reported in this article use DWA as the underlying motion planning algorithm. Furthermore,

those classical systems generally do not learn with increasing navigation experience [10].

With the recent advances in machine learning research, data-driven techniques have started being applied to the mobile robot navigation problem. Xiao, et al. [11] presented a survey on using machine learning for motion control in mobile robot navigation: in addition to some works in the emerging social [12] and terrain-based [13] navigation, most existing learning approaches for navigation adopt an end-to-end learning paradigm to address the classical collision-free navigation problem and are capable of generating navigational behaviors [14,15]. However, those end-to-end approaches still cannot outperform classical navigation methods or are not even compared to them. More importantly, end-to-end learning is extremely data-hungry, usually requiring millions of training data or training steps, and forgoes verifiable guarantees such as safety and explainability. On the other hand, other work which targeted a specific navigational component [16–18] has achieved superior performance when being compared to their classical counterparts. Therefore, it is promising to use machine learning at the subsystem or component level and to combine it with the structure of classical approaches [11].

APPL uses machine learning at the parameter level and devises extra learning components that interact with classical navigation systems. Therefore, APPL inherits all the benefits of classical approaches, while enjoying the adaptivity and flexibility of learning methods.

2.2. Learning from humans

Using interactions with humans is an effective approach to facilitate learning in general, e.g., Learning from Demonstration or Imitation Learning [19]. Of particular interest to the robotics community is learning through interactions with non-expert users, which relaxes the requirement for expert roboticists. For mobile robot navigation, it is relatively easy for most non-expert users to provide a teleoperated demonstration [20]. If a robot can successfully navigate in most scenarios, it is only necessary to teach the robot where it makes mistakes. Learning from Intervention can therefore be applied only at those trouble-some situations [21]. For non-expert users who are not able to take control of the robot, learning from evaluative feedback [22] provides another interaction modality that teaches the robot with a simple scalar feedback value. Most aforementioned approaches of learning from humans have been proposed in the learning community and applied to test domains such as Atari games or simulations, and have yet been used to tackle physical mobile robot navigation problems. We posit that this gap can be caused by the requirement for extensive training data and training time, which is prohibitive for physical mobile robots navigating in the real world.

APPL utilizes all these human interaction modalities (demonstration, interventions, and evaluative feedback), and combines them with Reinforcement Learning (RL) in a Cycle-of-Learning scheme [23] for autonomous mobile robot navigation. To alleviate the heavy dependency on huge amounts of high quality training data from the real world, APPL learns a *parameter policy* that interacts with an underlying navigation system, in contrast to replacing it by learning an end-to-end motion policy.

2.3. Self-supervised learning

Researchers have also investigated self-supervised learning techniques for robotics. One such example is reinforcement learning from trial and error without access to expert or non-expert humans [24]. These self-supervised approaches work very well given access to a high-fidelity simulator [25,26] (e.g., in Atari

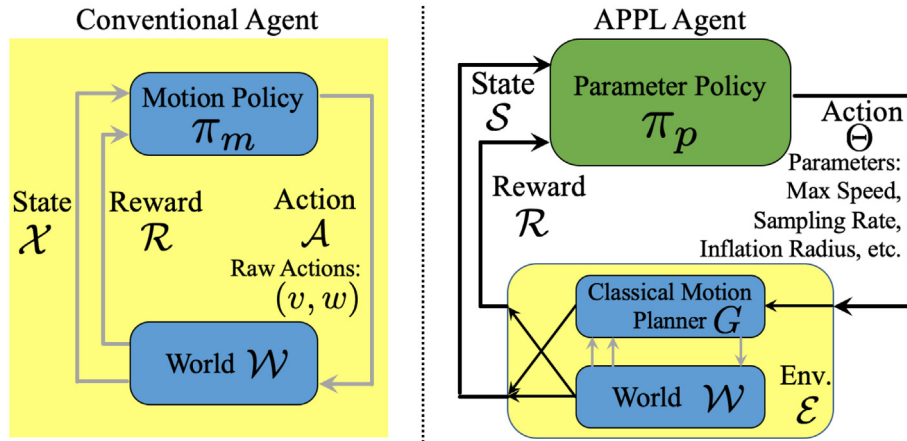


Fig. 1. Contrast between Conventional Agent vs. APPL Agent.

games) and a well-defined reward function to specify the desired behavior to be learned [27]. Although such self-learning approaches remove the burden of having a human in the loop, they require an extensive amount of training data, which can be impractical to collect in the real world. Even with specific techniques to improve sample efficiency and allow real-world training, random exploration in the physical world is often very risky. For example, the BADGR [28] system failed catastrophically (e.g., flipping over) during its exploration phase and required manual intervention to reset the robot. Another approach, VOILA [29], learns a visual navigation policy in the real world in a self-supervised manner with the guidance of the visual observation of another robot executing the same navigation task. But such methods do not easily generalize to other environments and other navigation tasks. Self-supervised model learning [13] has also been investigated to learn kinodynamic models for real-world unstructured off-road terrain. But the self-supervised learning happens in an open-space, so that the probability of the robot colliding with any obstacles is minimized.

APPL also leverages self-supervised learning in simulation using reinforcement learning, in addition to the non-expert human interactions collected from the wide spectrum of modalities in the real world. Such a Cycle-of-Learning scheme allows robots to continually improve their navigation behavior using both human interactions and self-supervision.

3. Adaptive planner parameter learning

In this section, we formalize the *Adaptive Planner Parameter Learning* (APPL) problem, which serves as the foundation for all the following APPL methods.

3.1. Underlying navigation planner

APPL assumes that a mobile robot has an underlying navigation planner $G : \mathcal{X} \times \Theta \rightarrow \mathcal{A}$, where \mathcal{X} is the planner's state space (e.g., robot odometry, sensory inputs, navigation goal), Θ is the space of free parameters for G (e.g., inflation radius, sampling rate, planner optimization coefficients), and \mathcal{A} is the planner's action space (e.g., linear and angular velocity v ω). Using G and a particular set of parameters θ , the robot performs navigation by repeatedly estimating its state x and applying action $a = G(x; \theta) = G_\theta(x)$. Traditional learning methods for navigation replace classical navigation planner G with a differentiable neural network, and use gradient descent to find thousands (or millions) of neural network weights. In contrast, APPL works within the framework of classical planner G , but treats it as a black box, e.g., it does not need to be differentiable, and APPL does not need to understand what each component of θ does.

3.2. Meta-environment

APPL works in the context of a *meta-environment* \mathcal{E} composed of both the underlying navigation world \mathcal{W} (the physical, obstacle-occupied world) and the given classical planner G with adjustable parameters $\theta \in \Theta$ and state input $x \in \mathcal{X}$. An APPL agent interacts with this meta-environment \mathcal{E} through a meta-state $s_t \in \mathcal{S}$, which includes both the planner's current state $x_t \in \mathcal{X}$ and previous parameters $\theta_{t-1} \in \Theta$, i.e., $\mathcal{S} = \mathcal{X} \times \Theta$. Instead of the raw action $a \in \mathcal{A}$, APPL's meta-action takes the form of $\theta \in \Theta$, which is the current parameters to be used by G .

3.3. Parameter policy

We formulate the APPL problem as a meta Markov Decision Process (MDP) in this meta-environment, i.e., a tuple $(\mathcal{S}, \Theta, \mathcal{T}, \gamma, R)$. Note this meta-MDP differs from the conventional MDP defined by traditional learning-based motion planners, whose states are x and actions are a (Fig. 1 left). Based on a meta-state s_t , an APPL agent takes action θ_t so that the state advances to s_{t+1} based on the transition function $s_{t+1} \sim \mathcal{T}(\cdot | s_t, \theta_t)$, and then receives a reward $r_t = R(s_t, \theta_t)$ (Fig. 1 right). In general, the objective of an APPL agent is to learn a *parameter policy* $\pi : \mathcal{S} \rightarrow \Theta$ that can be used to select actions (as parameters θ) that maximize the expected cumulative reward over time,

$$\max_{\pi} J^{\pi} = \mathbb{E}_{s_0, \theta_t \sim \pi(s_t), s_{t+1} \sim \mathcal{T}(s_t, \theta_t)} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right]. \quad (1)$$

Within the meta-environment \mathcal{E} , with the selected θ , $G_\theta(x)$ produces navigation action a to interact with the physical world \mathcal{W} . On top of this general form of parameter policy based on a general reward function, we impose additional structures on the policy and on the reward for each individual APPL method. For example, APPLR adopts this general notion of parameter policy and optimizes a reward function; an APPLE policy maximizes the expected evaluative feedback from human users as an approximation of the underlying reward; APPLI and APPLD simplify the general parameter policy in terms of a pre-trained context predictor to select appropriate parameters from a parameter library.

3.4. Parameter library

A specific type of APPL parameter policy $\pi : \mathcal{S} \rightarrow \Theta$ can be instantiated by imposing two intermediate mappings, i.e. a parameterized context predictor $B_\phi : \mathcal{S} \rightarrow \mathcal{C}$, and a one-to-one

mapping $M : \mathcal{C} \rightarrow \Theta$, where \mathcal{C} is the space of (finite) contexts \mathcal{C} (a relatively cohesive region in the physical world \mathcal{W}). Therefore, $\theta = \pi(s) = M(B_\phi(s))$. Each predefined context is associated with a set of static parameters. These parameter sets comprise a *parameter library* \mathcal{L} . APPLD and APPLI impose such structure on the general parameter policy, and selecting an appropriate parameter set which minimizes a Behavior Cloning loss approximates maximizing the underlying reward in Eq. (1).

The high-level algorithm of APPL is shown in Alg. 1. The subroutine $\pi = \text{LearnParameterPolicy}(\mathcal{I}, \Theta, G)$ is instantiated differently in each APPL method, as specified in Sections 4–7.

Algorithm 1 APPL

```

1: // Training
2: Input: human interaction  $\mathcal{I}$ , space of possible parameters  $\Theta$ ,
   and navigation stack  $G$ .
3:  $\pi = \text{LearnParameterPolicy}(\mathcal{I}, \Theta, G)$ .
4: // Deployment
5: Input: navigation stack  $G$ , parameter policy  $\pi$ .
6: for  $t = 1 : T$  do
7:   construct meta-state  $s_t$  from  $x_t$  and  $\theta_{t-1}$ .
8:    $\theta_t = \pi(s_t)$ .
9:   Navigate with  $G_{\theta_t}(x_t)$ .
10: end for

```

4. APPL from demonstration (APPLD)

In this section, we briefly summarize *Adaptive Planner Parameter Learning from Demonstration* (APPLD) and re-formalize it based on the APPL formulation in Section 3. APPLD [2] utilizes a context predictor and a *parameter library* learned from human demonstration. A teleoperated demonstration with planner state x and demonstrated action a is collected, which is first segmented into different *contexts* using Bayesian change point detection. Within each context, we use behavior cloning to find a set of planner parameters θ to match the planner's output with the human demonstration (Fig. 2). For full details and experiment results, please refer to our Robotics and Automation Letters (RAL) article [2].

4.1. Learning APPLD policy

A human demonstration of successful navigation is recorded as time series data $\mathcal{I} = \mathcal{D} = \{x_i^D, a_i^D, t_i^D\}_{i=1}^N$, where N is the length of the series, and x_i^D and a_i^D represent the planner state and demonstrated action at time t_i^D (Line 2 Alg. 1). As mentioned in Section 3, we impose additional structures on the APPLD policy, i.e. $\pi_D : \mathcal{S} \rightarrow \Theta$ is instantiated by a parameterized context predictor $B_\phi : \mathcal{S} \rightarrow \mathcal{C}$, and a one-to-one mapping $M : \mathcal{C} \rightarrow \Theta$ (Line 3 Alg. 1). Given M and B_ϕ , our system then performs navigation by selecting actions according to $G_\theta(x) = G_{M(B_\phi(x))}(x)$ (Lines 8–9 Alg. 1).

4.1.1. Demonstration segmentation

The key of learning π_D is to construct the space of *contexts* \mathcal{C} , each of which corresponds to a cohesive navigation region where a single set of parameters suffices. For each specific context, we apply black-box optimization to learn a parameter set. Generally speaking, any changepoint detection method can be used to solve this segmentation problem [30]. A changepoint detection algorithm A_{segment} can automatically detect how many changepoints exist in \mathcal{D} and where those changepoints are within the demonstration. Denote the number of changepoints found by A_{segment} as $K - 1$ and the changepoints as $\tau_1, \tau_2, \dots, \tau_{K-1}$ with $\tau_0 = 0$ and $\tau_K = N + 1$, the demonstration \mathcal{D} is then segmented into K pieces $\{\mathcal{D}_k = \{x_i^D, a_i^D, t_i^D \mid \tau_{k-1} \leq i < \tau_k\}\}_{k=1}^K$.

4.1.2. Parameter learning

For each of the segmented contexts, $\mathcal{D}_k = \{x_i^D, a_i^D, t_i^D \mid \tau_{k-1} \leq i < \tau_k\}$, we employ *behavioral cloning* (BC) [31] to learn a suitable set of parameters θ_k^* . More specifically, BC seeks θ_k^* which minimizes the difference between the demonstrated actions and the actions that G_{θ_k} would produce on $\{x_i^D\}$:

$$\theta_k^* = \operatorname{argmin}_{\theta} \sum_{(x,a) \in \mathcal{D}_k} \|a - G_\theta(x)\|_H, \quad (2)$$

where $\|v\|_H = v^T H v$ is the induced norm by a diagonal matrix H with positive real entries, which is used for weighting each dimension of the action. We solve Eq. (2) with a black-box optimization method $A_{\text{black-box}}$. Having found each θ_k^* , the *parameter library* \mathcal{L} is formed and the mapping M is simply $M(k) = \theta_k^*$. This fully parallelizable optimization takes approximately eight hours in our experiments on a single Dell XPS laptop (Intel Core i9-9980HK) using 16 parallel threads, but this time could be significantly reduced with more computational resources and engineering effort.

4.1.3. Online context prediction

The context predictor B_ϕ is learned with a supervised dataset $\{x_i^D, c_i\}_{i=1}^N$, where $c_i = k$ if $x_i^D \in \mathcal{D}_k$. To classify which segment x_i^D comes from, we learn a parameterized function $f_\phi(x)$ via supervised learning:

$$\phi^* = \operatorname{argmax}_{\phi} \sum_{i=1}^N \log \frac{\exp(f_\phi(x_i^D)[c_i])}{\sum_{c=1}^K \exp(f_\phi(x_i^D)[c])}. \quad (3)$$

Our context predictor B_ϕ is then defined as:

$$B_\phi(x_t) = \operatorname{mode} \left\{ \operatorname{argmax}_c f_\phi(x_i)[c], \quad t-p < i \leq t \right\}. \quad (4)$$

In other words, B_ϕ acts as a mode filter on the context predicted by f_ϕ over a sliding window of length p .

The *LearnParameterPolicy* subroutine in Alg. 1 for APPLD is shown in Alg. 2, where the above three stages are applied sequentially to learn a *parameter library* $\mathcal{L} = \{\theta_k^*\}_{k=1}^K$ (hence the mapping M) and a *context predictor* B_ϕ , both of which constitute the APPLD policy π_D . During deployment, Eq. (4) is applied online to pick the right set of parameters for G , π_D , as a special case of π , does not consider θ_{t-1} as part of s_t (only x_t , line 7 Alg. 1), but consults a history of p steps of x_t for smoothness.

Algorithm 2 *LearnParameterPolicy* (APPLD)

```

1: Input:  $\mathcal{I} = \mathcal{D} = \{x_i^D, a_i^D, t_i^D\}_{i=1}^N, \Theta, G$ .
2: Call  $A_{\text{segment}}$  on  $\mathcal{D}$  to detect changepoints  $\tau_1, \dots, \tau_{K-1}$  with
    $\tau_0 = 0$  and  $\tau_K = N + 1$ .
3: Segment  $\mathcal{D}$  into  $\{\mathcal{D}_k = \{x_i^D, a_i^D, t_i^D \mid \tau_{k-1} \leq i < \tau_k\}\}_{k=1}^K$ .
4: Train a classifier  $f_\phi$  on  $\{x_i^D, c_i\}_{i=1}^N$ , where  $c_i = k$  if  $x_i^D \in \mathcal{D}_k$ .
5: for  $k = 1 : K$  do
6:   Call  $A_{\text{black-box}}$  with objective defined in Eq. (2) on  $\mathcal{D}_k$  to find
   parameters  $\theta_k^*$  for context  $k$ .
7: end for
8: Form the map  $M(k) = \theta_k^*, \forall 1 \leq k \leq K$  and context predictor
    $B_\phi(x)$ .

```

4.2. Experiments

We implement APPLD on a physical ClearPath Jackal robot to experimentally validate that using the learned parameter library and context predictor from a teleoperated demonstration can achieve better navigation performance in complex environments compared to that obtained by the underlying navigation system

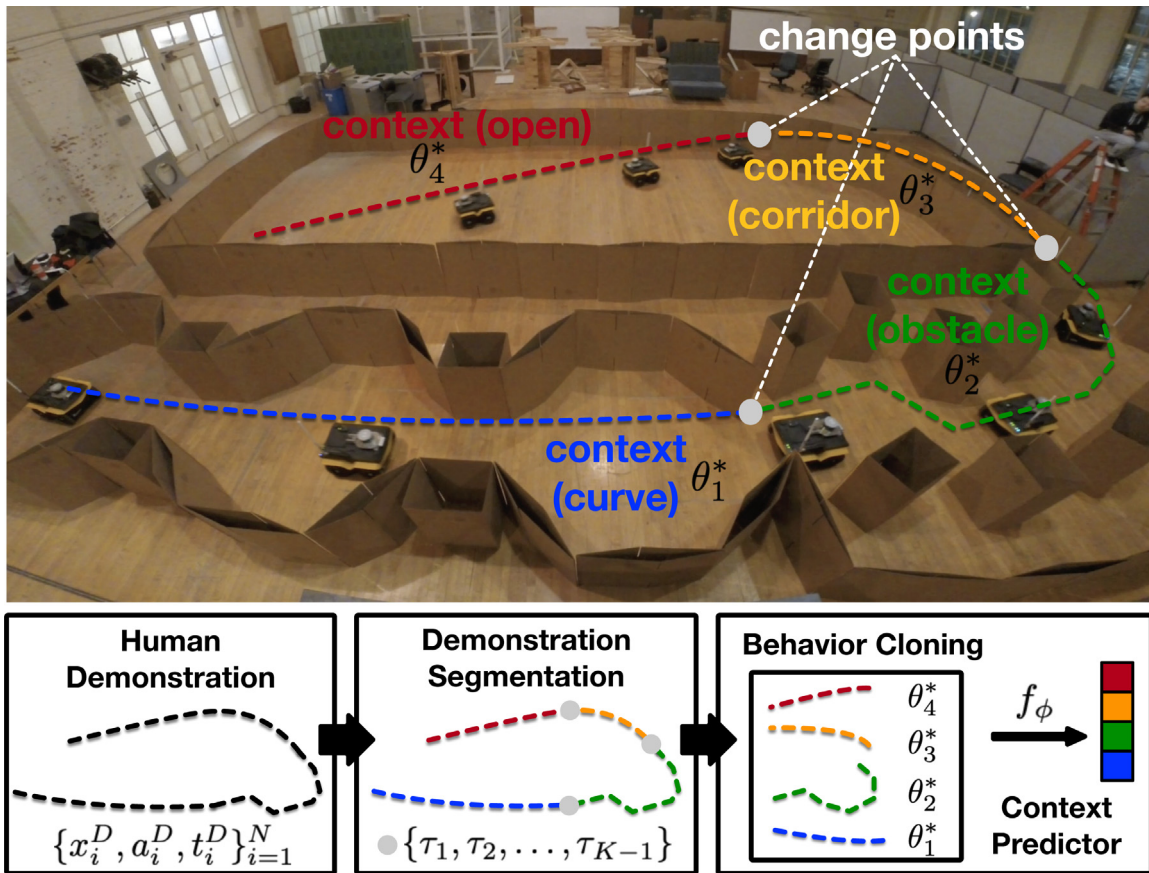


Fig. 2. APPLD: a human demonstration is segmented into different contexts, for each of which, a set of parameters θ_k^* is learned via Behavior Cloning. During deployment, proper parameters are selected by an online context predictor.

using (a) its default parameters from the robot platform manufacturer, and (b) parameters we found using behavior cloning but without context. The four-wheeled, differential-drive, unmanned ground vehicle with a top speed of 2.0 m/s is tasked to move through a custom-built maze as quickly as possible (Fig. 2). A Velodyne LiDAR provides 3D point cloud data, which is transformed into 2D laser scan for 2D navigation. The Jackal runs Robot Operating System (ros) onboard, and APPLD is applied to the local planner, DWA [7], in the commonly-used move_base navigation stack. Other parts of the navigation stack, e.g. global planning with Dijkstra’s algorithm, remain intact.

We use CHAMP as $A_{segment}$ (Line 2 Alg. 2), a state-of-the-art Bayesian segmentation algorithm [32]. We find each θ_k^* using CMA-ES [33] as our black-box optimizer (Line 6 Alg. 2). θ includes DWA’s max_vel_x (v), max_vel_theta (w), $vx_samples$ (s), $vtheta_samples$ (t), $occdist_scale$ (o), $pdist_scale$ (p), $gdist_scale$ (g) and costmap’s $inflation_radius$ (i).

In our experiments, APPLD achieves superior performance in terms of fastest traversal time compared to the default parameters, parameters learned without context, and even the demonstrator. For all the experiments in this paper, we use traversal time as the comparison metric since most suboptimal navigation behavior cause stop-and-go motions, induce recovery behaviors, cause the robot to get stuck, or collide with obstacles (termination) - each of which will result in a higher traversal time. For full details about the experimental setup and results, along with experimental results on a different robot with a different navigation system in a different environment, please refer to our RAL article [2].

5. APPL from interventions (APPLI)

APPLD assumes human demonstration is a good approximation for optimal navigation behavior and default parameters do not work well in most places, which are not always the case. In *Adaptive Planner Parameter Learning from Interventions* (APPLI) [3], instead of providing a full demonstration of the entire navigation task (such as APPLD), non-expert users can easily identify failure or suboptimal cases by watching and then provide a few teleoperated *interventions* to correct the failure or suboptimal behaviors (Fig. 3). APPLI utilizes this human interaction modality and learns sets of planner parameters specifically for the scenarios where failure and suboptimal behaviors take place, to create a *parameter library* including the default parameters. During deployment, APPLI applies those learned parameters, only when it is confident that they will benefit the current navigation based on a confidence measure of the context predictor, to the underlying navigation system in those troublesome places, while maintaining good performance in others by switching back to the default parameters. This confidence measure also enables APPLI to generalize well to unseen environments.

5.1. Learning APPLI policy

APPLI assumes a default parameter set $\bar{\theta}$ 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: $\bar{\theta}$ 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.




Fig. 3. While classical navigation systems perform well in most places (green), they may fail (red) or suffer from suboptimal behavior (yellow) in others. APPLI learns from human interventions in these two scenarios (Type A and Type B). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

5.1.1. Interventions as contexts

To mitigate this problem, a human can supervise the navigation system's performance at state x by observing its action a and judging whether she 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 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 naturally shows a cohesive navigation behavior in a specific segment of the environment (open space, narrow corridor, etc.), it automatically forms a *context* $c_i \in \mathcal{C}$. Using human interaction composed of N interventions $\mathcal{I} = I_{1:N}$ instead of the full demonstration sequence $\mathcal{D} = \{x_i^D, a_i^D, t_i^D\}_{i=1}^N$ (Section 4), APPLI finds the mapping $M : \mathcal{C} \rightarrow \Theta$ that determines the parameter set θ_i for each intervention context c_i , and the parameterized context predictor $B_\phi : \mathcal{X} \rightarrow \mathcal{C}$ that determines to which context (if any) the current state x belongs. After collecting a set of N interventions $I_{1:N}$, for each I_i , we learn a set of navigation parameters θ_i that can best imitate the demonstrated correction with the same Behavior Cloning loss (Eq. (2)) with a black-box optimizer, such as CMA-ES [33]. After identifying parameters in each context, the mapping M is simply $M(i) = \theta_i$.

5.1.2. Confidence-based context prediction

In contrast to APPLD, APPLI does not learn parameters for places where the original navigation performance is good. Therefore, it determines if the current state x_t falls into any one of the collected intervention contexts c_i . If so, APPLI directs the robot to use the parameter set θ_i to avoid making the same mistake as before. If not, APPLI directs 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 new context predictor, B_ϕ , with confidence measure. To train this predictor, we build a similar dataset as in APPLD, $\{\{x_t, c_i\}_{t=1}^{T_i}\}_{i=1}^N$, and train an intermediate classifier $f_\phi(x)$ with parameter set ϕ using the Evidential Deep Learning method (EDL) [34]. A feature of EDL is that it supplies both a predicted label and a confidence measure, $u_i \in (0, 1]$, i.e.,

$$f_\phi(x_i) = (c_i, u_i). \quad (5)$$

After training f_ϕ and during deployment, we build a confidence-based classifier g_ϕ as

$$g_\phi(x_i) = c_i \mathbb{1}(u_i \geq \epsilon_u), \quad (6)$$

where ϵ_u is a threshold on confidence and $\mathbb{1}$ is the indicator function. For state x_i , g_ϕ determines its context from $N + 1$ contexts (N intervention contexts and one default context). If $u_i \geq \epsilon_u$, it suggests the classifier f_ϕ is confident and g_ϕ predicts c_i . Otherwise, when f_ϕ is unsure about its prediction, $c_i \mathbb{1}(u_i \geq \epsilon_u) = 0$. In this case, g_ϕ 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 $M(0) = \bar{\theta}$). The new confidence-based context predictor B_ϕ is then defined as:

$$B_\phi(x_t) = \text{mode} \left\{ g_\phi(x_i), t - p < i \leq t \right\}. \quad (7)$$

Again, B_ϕ acts as a mode filter and chooses the context c_i that the majority agrees with over the past p time steps.

The *LearningParameterPolicy* subroutine in Alg. 1 for APPLI is shown in Alg. 3. APPLI learns a confidence-based classifier (Line 2) and navigation parameters $\theta_{1:N}$ for each context (Lines 3–5). During deployment, we use $\theta_t = \pi_I(x_t) = M(B_\phi(x_t))$ to select the parameters for the navigation system at time t . Similar to π_D , π_I does not consider θ_{t-1} as part of s_t (only x_t , line 7 Alg. 1), but consults a history of p steps of x_t for smoothness.

Algorithm 3 LearningParameterPolicy (APPLI)

- 1: **Input:** $\mathcal{I} = I_{1:N} = \{\{x_t, a_t\}_{t=1}^{T_i}\}_{i=1}^N$, Θ , G .
 - 2: Train a confidence-based classifier f_ϕ on $\{\{x_t, c_i\}_{t=1}^{T_i}\}_{i=1}^N$.
 - 3: **for** $i = 1, \dots, N$ **do**
 - 4: Find parameter θ_i for context i using Eq. (2) on I_i .
 - 5: **end for**
 - 6: Form the map $M(i) = \theta_i, \forall 1 \leq i \leq N$, and confidence-based context predictor $B_\phi(x)$.
-

5.2. 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 the same ClearPath Jackal ground robot with the same setup as in APPLD (Section 4.2) in a physical obstacle course.

Table 1
APPLI traversal time.

	Default	Type A	Type A + B	Full Demo
Training environment	134.0 ± 60.6 s	77.4 ± 2.8 s	70.6 ± 3.2 s	78.0 ± 2.7 s
Unseen environment	109.2 ± 50.8 s	71 ± 0.7 s	59 ± 0.7 s	62.0 ± 2.0 s




Fig. 4. APPLI running in an unseen physical environment.

Navigation performance learned through APPLI is then tested both in the same training environment, and also in another unseen physical test course, which is qualitatively similar to the training environment (i.e., similar contexts were created in a different ordering). Furthermore, to investigate generalizability, we test the learned systems on a benchmark suite of 300 unseen simulated navigation environments [35].

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. 3: 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 suboptimal behavior to have begun, and then provides recorded teleoperation I that avoids the problematic behavior. It takes less than an hour to learn a set of parameters for each intervention. Training the EDL-based context predictor takes only a few minutes using the same computational infrastructure as specified in Section 4.1.1. 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 well (shown in green in Fig. 3).

5.2.1. Physical experiments

After training, we deploy π_I with learned mapping M and context predictor B_ϕ on the `move_base` navigation stack G with an empirically chosen threshold value $\epsilon_u = 0.8$.

We first deploy APPLI in the same training physical environment (Fig. 3). 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 APPLD enhanced by manual context segmentation and 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 she is not paying attention, or is not equipped with the knowledge to identify suboptimal behaviors. For each method, we run five trials and report the mean and standard deviation of the traversal time in Table 1 1st row. If the robot gets stuck, we introduce a penalty value of 200 s. We also deploy the same sets of variants in an unseen physical environments (Fig. 4 and Table 1 2nd row).

For both the training and unseen environments, Type A interventions alone significantly improve upon the default parameters ($p = 0.016$), 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 ($p = 5 \times 10^{-8}$), 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 ($p_{\text{env.}} = 0.11$), 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 ($p = 1.5 \times 10^{-4}$). 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.

5.2.2. Simulated experiments

To further test APPLI's generalizability to unseen environments, we test our method and compare it with two baselines on the Benchmark for Autonomous Robot Navigation (BARN) dataset [35]. The benchmark dataset consists of 300 simulated navigation environments randomly 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 (see examples in Fig. 5). Using the same training data collected from the physical environment shown in Fig. 3, we test the following seven variants: (1) APPLI (A + B + c): APPLI



Fig. 5. Example navigation environments in the BARN dataset.




Fig. 6. Normalized performance in 300 simulation environments from 12 runs each. Error bars: standard deviation.

Table 2

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						
		(A)	DWA	(A + c)	(A + B + D + c)	(A + B + D)	(A + B + c)	(A + B)
Method 1	I	0	50	53	62	63	68	66
		II	10	0	6	33	40	44
	III	6	4	0	31	37	45	45
	IV	5	7	11	0	25	31	33
	V	5	7	7	10	0	21	21
	VI	3	3	4	3	5	0	9
	VII	2	5	5	6	4	6	0

learned from *Type A and B* inventions with confidence measure, (2) APPLI (A + B): APPLI learned from *Type A and B* inventions without confidence measure, (3) APPLI (A + c): APPLI learned from only *Type A* interventions with confidence measure, (4) APPLI (A): APPLI learned from only *Type A* interventions without confidence measure, (5) APPLI (A + B + D + c): APPLI learned from *full demonstration* with confidence measure, (6) APPLI (A + B + D): APPLI learned from *full demonstration* without confidence measure, (7) 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_ϕ : 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, resulting in 25 200 total navigation trials. The average traversal time for each method in all simulation environments is normalized by DWA and is shown

in order of increasing performance in Fig. 6. We conduct a pairwise 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 performance and show the pairwise comparisons in Table 2.

APPLI (A + B + c) and APPLI (A + B + D + c) outperform DWA: their average traversal time is shorter than that of DWA by 8% and 17% respectively, and they are significantly better in 44% and 33% of environments respectively and significantly worse in only 3% and 7% of environments than DWA. However, for APPLI (A + c), its traversal time is only 1% better than DWA (significantly better 6% of the time), which suggests that even though type B inventions only correct suboptimal performances, they are crucial for performance improvement.

In terms of the effect of confidence, APPLI (A) only selects parameters learned from 2 *Type A* inventions and never uses the default parameters even when they are more appropriate. Removing confidence greatly harms its performance, making its traversal time even longer than DWA by 31% (significantly worse in 53% of environments). However, APPLI (A + B + c) and APPLI (A + B + D + c), which use more interventions or even the full



Fig. 7. For non-expert users who are unable or unwilling to take control of the robot, evaluative feedback, e.g. *good job* (green thumbs up) or *bad job* (red thumbs down), is a more accessible human interaction modality, but still valuable for improving navigation systems during deployment.

demonstration to train the parameter mapping M and context predictor B_ϕ , are more confident about their predictions most of the time. As a result, removing confidence in the context predictor does not result in a significant difference.

Lastly, a counterintuitive, but similar result as in the physical experiments is that APPLI (A + B + c) learned from only Type A and B interventions achieves shorter traversal time than APPLI (A + B + D + c) and APPLI (A + B + D) by 9.2% and 6.1% (significantly better in 31% and 21% of the environments), respectively. 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.

6. APPL from evaluative feedback (APPLE)

APPLD and APPLI require non-expert users to take full control of the moving robot with a joystick, which some non-expert users may not feel comfortable with, e.g., due to perceived risk of human error and causing collisions. For these users, we introduce *Adaptive Planner Parameter Learning from Evaluative Feedback* (APPLE) [4], where they only need to observe the robot navigating and provide real-time positive or negative assessments of the observed navigation behavior through *evaluative feedback*. This more-accessible modality provides a new interaction channel for a larger community of non-expert users with mobile robots (Fig. 7). In this section, we briefly summarize APPLE and reformalize it based on the APPL formulation in Section 3. For full details and experiment results, please refer to our RAL article [4].

6.1. Learning APPLE policy

APPLE learns a parameter policy from human evaluative feedback in order to select the appropriate parameter set θ for the current deployment scenario. The parameter set can be selected from either a discrete parameter set library or from a continuous full parameter space. To be specific, a human observes the underlying planner taking action a at state x , and then provides evaluative feedback e . The human can provide either discrete (e.g., “good job”/“bad job”) or continuous (e.g., a score ranging in $[0, 1]$) evaluative feedback. With the evaluative feedback from the human, APPLE finds (1) a parameterized predictor $F_\phi : \mathcal{X} \times \Theta \rightarrow \mathcal{E}$ that predicts human evaluative feedback for each state-parameter pair (x, θ) , and (2) a parameterized parameter policy $\pi_\psi : \mathcal{X} \rightarrow \Theta$ that determines the appropriate planner parameter set θ for the current state x .

6.1.1. Discrete parameter policy

A discrete parameter policy is designed for the situations where K candidate parameter sets (e.g., the default set or sets tuned for specific environments like narrow corridors, open spaces, etc.) are already available. These K candidate parameter sets comprise a parameter library $\mathcal{L} = \{\theta^i\}_{i=1}^K$ (superscript i denotes the index in the library). In this case, a discrete APPLE policy learns to select the most appropriate of these parameters given the state x using the provided evaluative feedback e .

In our discrete APPLE policy, the feedback predictor F_ϕ is parameterized in a way similar to the value network in DQN [36]: the input is the state x , while the output is K predicted feedback values $\{\hat{e}^i\}_{i=1}^K$. Each \hat{e}^i is a prediction of the evaluative feedback a human user would give if the planner were using the respective parameter set $\theta^i \in \mathcal{L}$ at state x . During training, a feedback dataset for supervised learning, $\mathcal{F} = \{x_j, \theta_j, e_j\}_{j=1}^N$ ($\theta_j \in \mathcal{L}$, subscript j denotes the time step), is built using the evaluative feedback collected so far. We use supervised learning which minimizes the difference between predicted feedback and the label to learn F_ϕ :

$$\phi^* = \underset{\phi}{\operatorname{argmin}} \mathbb{E}_{(x_j, \theta_j, e_j) \sim \mathcal{F}} \ell(F_\phi(x_j, \theta_j), e_j) \quad (8)$$

where $\ell(\cdot, \cdot)$ is the binary cross entropy loss if the feedback e_j is discrete (e.g., $e_j = 1$ for “good job”, and $e_j = 0$ for “bad job”), or mean squared error given continuous feedback.

With F_ϕ^* , the discrete parameter policy $\pi(\cdot|x)$ simply chooses the parameter set that maximizes the expected human feedback:

$$\pi(\cdot|x) = \underset{\theta \in \mathcal{L}}{\operatorname{argmax}} F_\phi^*(x, \theta). \quad (9)$$

Note ψ is omitted since only ϕ^* is needed for π_ψ to select the appropriate θ .

6.1.2. Continuous parameter policy

APPLE can also learn to select appropriate parameters from continuous parameter spaces (e.g., deciding the max speed from $[0.1, 2]$ m/s). For continuous APPLE policy, the parameter policy π_ψ and the feedback predictor F_ϕ are parameterized in the actor-critic style [37]. Similar to the discrete APPLE policy, we also train F_ϕ by minimizing the difference between predicted and collected feedback using the collected dataset $\mathcal{F} = \{x_j, \theta_j, e_j\}_{j=1}^N$ (Eq. (8)). The parameter policy π_ψ is trained to not only choose the action that maximizes expected feedback, but also to maximize the entropy of policy $\mathcal{H}(\pi_\psi(\cdot|x))$ at state x . We use the same entropy regularization as Soft Actor Critic (SAC) [37], so that π_ψ favors more stochastic policies for better exploration during training:

$$\psi^* = \underset{\psi}{\operatorname{argmin}} \mathbb{E}_{x_j \in \mathcal{F}} \left[-F_\phi(x_j, \tilde{\theta}_j) + \alpha \log \pi_\psi(\tilde{\theta}_j|x_j) \right], \quad (10)$$

$\tilde{\theta}_j \sim \pi_\psi(\cdot|x_j)$

Table 3
Traversal time in training and unseen environment.

	Default	APPLI	APPLE (disc.)
Training	143.1 ± 20.0 s	79.8 ± 8.1 s	75.2 ± 4.1 s
Unseen	150.5 ± 24.0 s	86.4 ± 1.1 s	83.9 ± 4.6 s

where α is the temperature controlling the importance of the entropy bonus and is automatically tuned as in SAC [37].

The *LearningParameterPolicy* subroutine in Alg. 1 for APPLE to learn π_E is simply learning F_{ϕ^*} with Eq. (8) and learning π_{ψ^*} with SAC [37] for the discrete and continuous parameter policy, respectively.

6.2. Experiments

In our experiments, we find that APPLE can improve navigation performance by learning from evaluative feedback, in contrast to a teleoperated demonstration or a few corrective interventions, both of which require the non-expert user to take control of the moving robot. We also show APPLE's generalizability to unseen environments. To be specific, we implement APPLE on the same ClearPath Jackal ground robot in BARN [35], and in two physical obstacle courses. We show that APPLE can achieve significant improvement compared to the default parameters, and even slightly better performance (not statistically significant) than APPLI learned from a similar obstacle course (Table 3). For full details about the experimental setup and results, please refer to our RAL article [4].

7. APPL from reinforcement (APPLR)

Adaptive Planner Parameter Learning from Reinforcement (APPLR) [5] adopts the general notion of *parameter policy* (Section 3 Fig. 1). One disadvantage of learning planner parameters from different human interaction modalities is that the learner's performance is limited by the human's (most likely sub-optimal) teleoperated demonstration, corrective interventions, and evaluative feedback. With reinforcement learning in a wide variety of simulation environments, APPLR does *not* need access to any human interaction, and learns to make planner parameter decisions in such a way that allows the system to take suboptimal actions at one state in order to perform even better in the future.

7.1. Learning APPLR policy

APPLR uses an existing RL algorithm and a reward function to learn a parameter policy π_R .

7.1.1. Reward function

In general, we encourage three types of behaviors: (1) behaviors that lead to the global goal faster; (2) behaviors that make more local progress; and (3) behaviors that avoid collisions and danger. Correspondingly, the designed reward function can be summarized as

$$R_t(s_t, a_t, s_{t+1}) = c_f R_f + c_p R_p + c_c R_c. \quad (11)$$

Here, c_f , c_p , c_c are coefficients for the three types of reward functions R_f , R_p , R_c . Specifically, $R_f(s_t, a_t) = \mathbb{1}(s_t \text{ is terminal}) - 1$ applies a -1 penalty to every step before reaching the global goal. To encourage the local progress of the robot, we add a dense shaping reward R_p . Assume $\beta = (\beta_x, \beta_y) \in \mathbb{R}^2$ is the global goal and at time t , the absolute coordinates of the robot are $p_t = (p_t^x, p_t^y)$, we define

$$R_p = \frac{(p_{t+1} - p_t) \cdot (\beta - p_t)}{|\beta - p_t|}, \quad (12)$$

In other words, R_p denotes the robot's local progress ($p_{t+1} - p_t$) projected on the direction toward the global goal ($\beta - p_t$). Finally, a penalty for the robot colliding with or coming too close to obstacles is defined as $R_c = -1/d(p_{t+1})$, where $d(p_{t+1})$ is a distance function measuring how close the robot is to obstacles based on sensor observations (e.g., using LiDAR).

7.1.2. Reinforcement learning algorithm

For continuous actions and high sample efficiency, we use the Twin Delayed Deep Deterministic policy gradient algorithm (TD3) [38], an actor-critic algorithm that keeps an estimate for both the policy and two state-action value functions.

7.2. Experiments

Again, we use the same ClearPath Jackal and the same setup to validate that APPLR can enable adaptive autonomous navigation *without* access to expert tuning or human demonstrations and is generalizable to many deployment environments, both in simulation and in the real-world. The results of APPLR are compared with those obtained by the underlying navigation system using its default parameters from the robot platform manufacturer. For the physical experiments, we also compare to parameters learned from APPLD.

7.2.1. Training

x_t is composed of 720-dimensional laser scan (capped to 2 m) and a relative goal angle, computed from a local goal taken from the global path 1 m away from the robot. The meta-state s_t is composed of x_t and θ_{t-1} . APPLR learns a parameter policy π_R to select the same DWA parameters θ_t as other APPL methods. π_R is trained in simulation using 250 randomly selected training environments from the BARN dataset [35]. In each of the environments, the robot aims to navigate from a fixed start to a fixed goal location in a safe and fast manner. For the reward function, while R_f penalizes each time step before reaching the global goal, we simplified R_p by replacing it with its projection along the y -axis (longitudinal direction) of all BARN environments, because the traversal paths from start to goal in all BARN environments are along the positive direction of the y -axis. The distance function in $d(p_{t+1})$ in R_c is the minimal value among the 720 laser beams. π_R produces a new set of planner parameters every two seconds.

This simulated navigation task is implemented in a Singularity container, which enables easy parallelization on a computer cluster. TD3 [38] is implemented to learn the parameter policy π_R in simulation. The policy network and the two target Q-networks are represented as multilayer perceptrons with three 512-unit hidden layers. The policy is learned under the distributed architecture Gorila [39]. The acting processes are distributed over 500 CPUs with each CPU running one individual navigation task. On average, two actors work on a given navigation task. A single central learner periodically collects samples from the actors' local buffers and supplies the updated policy to each actor. Gaussian exploration noise with 0.5 standard deviation is applied to the actions at the beginning of the training. Afterward the standard deviation linearly decays at a rate of 0.125 per million steps and stays at 0.02 after 4 million steps. The entire training process takes about 6 h and requires 5 million transition samples in total.

7.2.2. Simulated experiments

After training, we deploy the learned parameter policy π_R on both the 250 training and 50 test environments. We also use traversal time to evaluate the performance of the policy. A maximum traversal time of 50 s is used, and the failing trials are set to be 50 s plus a 20 s penalty. We average over the




Fig. 8. Traversal time difference between APPLR and DWA for the training environments (top) and the test environments (bottom). The bars represent the environments ordered by traversal time difference. The colored bars indicate the environments that show statistically significant difference. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 4

Average traversal time of APPLR and DWA.

	APPLR	DWA	Improvement	P-value
Training	20.8 ± 1.8	27.1 ± 1.3	6.3 (23.2%)	8.2×10^{-3}
Test	22.4 ± 1.0	26.3 ± 2.1	3.9 (14.8%)	5.3×10^{-2}

Table 5

Number and percentage of all environments in which one method is better compared to the other.

	APPLR better	DWA better
Training	106 (42.4%)	12 (4.8%)
Test	20 (40%)	4 (8%)

traversal time of 40 trials for DWA and APPLR in each environment. We also perform t-tests for each pair of APPLR and DWA performance to check statistical significance. The results over the 250 training environments and 50 test environments are shown in Fig. 8. For the majority of the environments (green bars), APPLR shows statistically significantly better navigation performance. Table 4 shows the average traversal time of APPLR and DWA, and relative improvement (averaged over three training runs). APPLR yields an improvement of 23.2% in the training environments, and 14.8% in the test environments. In addition, Table 5 compares the number of environments that show significant improvement and deterioration. In both training and test set, APPLR achieves statistically significantly better navigation performance in over 40% of environments than DWA does, while DWA is only better in 4.8% and 8% of environments in the training and test set, respectively.

7.2.3. Physical experiments

To validate the sim-to-real transfer of APPLR, we also test the learned parameter policy π_R on a physical Jackal robot. The learned policy is deployed in a real-world obstacle course (Fig. 9). This physical environment is different from any of the navigation environments in BARN. Therefore, both generalizability and sim-to-real transfer of APPLR can be tested with this unseen real-world environment. Note that the use of LiDAR input also reduces the difference between simulation and the real-world.

Given this target environment, we further collect a teleoperated demonstration provided by one of the authors and learn a

Table 6

Traversal time in physical experiments.

DWA	APPLD	APPLR
72.8 ± 10.1 s ^a	43.2 ± 4.1 s	34.4 ± 4.8 s

^aDenotes one additional failure trial.

parameter tuning policy based on the notion of navigational context (APPLD). The author aims at driving the robot to traverse the entire obstacle course in a safe and fast manner. APPLD identifies three contexts using the human demonstration and learns three sets of navigation parameters.

We compare the performance of APPLR with that of APPLD and the DWA planner using a set of hand-tuned default parameters. For each trial, the robot navigates from the fixed start point to a fixed goal point. Each trial is repeated five times and we report the mean and standard deviation in Table 6. We also conduct t-tests for APPLR vs. DWA and APPLR vs. APPLD and find the p-values to be 2.5×10^{-4} and 1.4×10^{-2} respectively, showing APPLR's statistically significant improvement. In all DWA trials, the robot gets stuck in many places, especially where the surrounding obstacles are very tight. It has to engage in many recovery behaviors, i.e. rotating in place or driving backwards, to "get unstuck". Furthermore, in relatively open space, the robot drives unnecessarily slowly. All these behaviors contribute to the large traversal time and high variance (plus an additional failure trial). Unlike many simulation environments in BARN [35], where obstacles are generated by cellular automata and therefore very cluttered, the relatively open space in the physical environment (Fig. 9) allows faster speed and gives APPLR a greater advantage. Surprisingly, APPLR even achieves better navigation performance than APPLD, which has access to a human demonstration in the same environment. One of the reasons we observe this result in the physical experiments is that the human demonstrator is relatively conservative in some places; the parameters learned by APPLD are upper-bounded by this suboptimal human performance. Another reason is that APPLR is given the flexibility to continually change parameters, and RL is able to utilize the sequential aspect of the parameter selection problem, in contrast to APPLD's three sets of static parameters.




Fig. 9. APPLR physical experiments.



Fig. 10. One learning cycle. White arrows: Learning in the real-world. Black arrow: Learning in simulation.

8. Cycle-of-learning

The four individual APPL methods are combined to form a cycle-of-learning scheme [23], in which a mobile robot can interact with different users in different deployment environments and continually improve its navigation in a cyclic fashion. We present our cycle-of-learning algorithm in Alg. 4. We provide detailed explanation of Alg. 4 using one full cycle (Fig. 10) as an example, which starts from APPLR, goes through APPLD, APPLI, and APPLE, and comes back to APPLR. To distinguish the start and end of the cycle, we name them APPLR1 and APPLR2.

We assume a mobile robot is given an underlying navigation system G (e.g., DWA planner [7]) and a space of possible parameters Θ (e.g., DWA parameters). We also assume there is a set of simulation environments, where the robot can learn through trial and error, e.g., the BARN dataset [35] (line 1 Alg. 4). Before any real-world deployment, we first train the APPLR1 policy π_R in the BARN dataset \mathcal{E} (line 3 Alg. 4). Then the cycle-of-learning starts (lines 5–28 Alg. 4). With only one parameter policy π_R in the policy set Π (line 7 Alg. 4), the first deployment environment in our example cycle is a relatively open space, where default DWA slowly drives at the default 0.5 m/s max velocity and APPLR1 is able to accelerate to 2.0 m/s (line 9 Alg. 4). Therefore user 0 is satisfied with the navigation performance and does not choose to interact with the robot (skipping lines 10–15 Alg. 4). Moving on to the second environment (line 9 Alg. 4), APPLR1 produces unsmooth motion at the narrow gap (line 10 Alg. 4). User 1 provides a full demonstration for Environment 2 (line 11 Alg. 4), from which APPLD learns two contexts and corresponding parameters (line 12 Alg. 4). The navigation therefore improves (line 13 Alg. 4). When deployed in the third environment (line 9 Alg. 4), APPLD gets stuck multiple times in the narrow corridor (line 10 Alg. 4), where User 2 intervenes (line 11 Alg. 4) and uses APPLI to learn an extra context (line 12 Alg. 4). In Environment 4 (line 9 Alg. 4), APPLI suffers from suboptimal behaviors in the obstacle field (line 10 Alg. 4). User 3 provides evaluative feedback (line 11 Alg. 4) and uses APPLE to improve the context predictor of APPLI (line 12 Alg. 4). The performance of different APPL variants in the four environments is shown in Table 7 (averaged over 5 trials). The method which directly utilizes human interaction in the particular environment achieves the best results in the corresponding environment (bold).

We further perform t-tests for the best two methods in each of the four environments and report the corresponding p-values

Table 7
Cycle-of-learning performance in Env. 1–4.

Env.	DWA	APPLR1	APPLD	APPLI	APPLE	APPLR2
1	9.8 ± 0.5 s	4.4 ± 0.3 s				
2	30.2 ± 1.4 s	18.6 ± 3.6 s	12.5 ± 0.4 s			
3	56.9 ± 2.5 s	57.5 ± 5.8 s	38.1 ± 2.3 s	37.0 ± 2.0 s		
4	109.2 ± 6.5 s	103.4 ± 12.8 s	90.6 ± 6.0 s	88.8 ± 2.9 s	76.4 ± 3.5 s	81.6 ± 5.6 s

Table 8
P-values of the best two methods for Env. 1–4.

	Env. 1 R1 vs. DWA	Env. 2 D vs. R1	Env. 3 I vs. D	Env. 4 E vs. I
p-value	4.5×10^{-8}	5.5×10^{-3}	0.45	2.9×10^{-4}

in Table 8. In Environments 1, 2, and 4, APPLR1, APPLD, and APPLE are statistically significantly better than the default DWA, APPLR1, and APPLI, respectively. In Environment 3, APPLI does not achieve statistically significant improvement over APPLD, possibly because the extremely difficult extra narrow corridor in Environment 3 also causes trouble for the parameters learned by APPLI.

After deploying in Environment 4 with APPLE, we finish all physical deployment in this cycle (line 16 Alg. 4) and then train APPLR2 on the BARN dataset (lines 18–27 Alg. 4). For each policy learned from different human interactions in the current policy set Π (line 19 Alg. 4), we identify the BARN environments where the APPL variants perform better than APPLR1 (line 21 Alg. 4), and use the learned APPLD, APPLI, or APPLE policies for exploration (line 22 Alg. 4)¹ with a probability which decreases from 1 to 0 in those environments. APPLR2 achieves significantly better performance in these environments due to the better exploration policy learned within the cycle from different human interaction modalities. We further test APPLR2’s performance against APPLR1’s in all BARN environments and do not observe significant deterioration in other environments. In fact, APPLR2’s average traversal time in BARN improves from APPLR1’s 24.7 s to 24.0 s. We also deploy APPLR2 in Environment 4 in Fig. 10. Although the training experience in BARN slightly increases traversal time compared to

¹ Other techniques for combining policies may be possible (e.g., [40]).

APPLE, APPLR2 significantly outperforms APPLR1 by incorporating all human interactions, including the learned context predictor and parameters, into a stand-alone parameter policy. We hypothesize that APPLR2's performance degradation compared to APPLE in Environment 4 arises because while APPLE's feedback specifically targets improving performance in Environment 4, APPLR2 has to consider performance in the simulated BARN environments as well, which may compromise the performance in Environment 4.

Algorithm 4 CYCLE-OF-LEARNING

```

1: Input: navigation stack  $G$ , space of possible parameters  $\Theta$ ,
   and simulation environments  $\mathcal{E}$ .
2: // Initialization
3: Train APPLR policy  $\pi_R$  in  $\mathcal{E}$  to select  $\theta \in \Theta$  for  $G$ .
4: // Cycle-of-Learning
5: while True do
6:   // Physical Deployment
7:   Initialize policy set  $\Pi = \{\pi_R\}$ .
8:   for each deployment do
9:     Deploy a policy  $\pi \in \Pi$  (e.g., the latest one).
10:    if unsatisfactory navigation performance then
11:      User provides interaction  $\mathcal{I}$  (Demonstration, Interventions, or Evaluative Feedback).
12:       $\pi = \text{LearnParameterPolicy}(\mathcal{I}, \Theta, G)$  (APPLD, APPLI, or APPLE, based on the chosen modality).
13:      Deploy with  $\pi$ .
14:       $\Pi = \Pi \cup \pi$ .
15:    end if
16:  end for
17:  // Simulated Training
18:  Initialize training dataset  $\mathcal{D} = \emptyset$ .
19:  for each  $\pi \in \Pi \setminus \pi_R$  do
20:    for each  $e_i \in \mathcal{E}$  do
21:      if  $\pi$  performs better than  $\pi_R$  in  $e_i$  then
22:        Explore with  $\pi$  in  $e_i$  to gather training data  $d$ .
23:         $\mathcal{D} = \mathcal{D} \cup d$ .
24:      end if
25:    end for
26:  end for
27:  Update  $\pi_R$  with  $\mathcal{D}$  for next cycle of deployments.
28: end while

```

9. Discussion

The APPL paradigm integrated with the Cycle-of-Learning scheme opens up at least three dimensions for autonomous mobile robots deployed in the real-world co-existing with non-expert human users. First, building upon classical navigation planners, the APPL agent interfaces with these systems through their hyper-parameters. Therefore, benefits of these classical systems can be inherited when being deployed in the real world, e.g., safety can still be provided by the provable guarantees of the classical systems and the behavior of the APPL agent is still explainable through the designed role of each learned hyper-parameter. While enjoying these benefits, APPL simultaneously empowers mobile robots with adaptivity to a variety of scenarios in the wild, which is usually an advantage of purely learning-based methods.

Second, APPL allows mobile robots to learn from a variety of non-expert human users through multi-model interactions, ranging from teleoperated demonstration, corrective interventions, and evaluative feedback. Relaxing the assumption of access to expert roboticists when facing new deployment scenarios or suboptimal navigation behavior in existing environments, many

non-expert human users, or even bystanders, can also contribute to the improvement of autonomous navigation performance in the real world.

Third, the Cycle-of-Learning scheme further removes the traditionally myopic focus on in-situ adjustment or improvement in one single deployment scenario with a particular human interaction, and aims toward a human-robot ecosystem where with non-expert humans' help robots are able to continually improve during real-world deployment throughout their lifetime. Robots are then expected to require less and less frequent interactions with human users but achieve better and better performance in their future deployments.

Such a learning paradigm that leverages the agent's own exploration experiences and interactions with other agents, human or artificial, has been considered by the learning community for many years. Lin [41] proposed to combine reinforcement learning with teaching frameworks to speed up reinforcement learning in solving complicated learning tasks three decades ago. At that time, the complicated learning task only involved navigating in a discrete 2D maze environment, similar to an Atari game. A decade later, Smart and Kaelbling [42] proposed a similar paradigm and moved closer toward real-world robotics applications. Similar to APPL's underlying classical planner, their method utilized a supplied control policy (either actual control code, or a human directly controlling the robot with a joystick) to kick start the first learning phase, where the learning agent only passively watches the supplied policy generating states, actions, and rewards and learns a value function. In the second phase, the learning agent takes control of the robot using the learned value function from the first phase. Such a paradigm assumes that after the first phase, the learning agent is fully capable of reliably executing any task in the workspace, which, however, is not always true. During real-world robot deployment which is the focus of APPL, encountering unseen scenarios is inevitable, and the learning agent may find it difficult to generalize well to such scenarios. On the other hand, a classical scripted system or a human may be able to address these scenarios in a more trustworthy way.

The APPL paradigm with the Cycle-of-Learning scheme is very reminiscent of the approach proposed by Smart and Kaelbling [42], i.e., leveraging classical non-learning-based methods and/or humans to improve learning, but with a focus on addressing the generalizability issue in real-world unseen scenarios mentioned above. We posit that a robot interacting with the real-world with an end-to-end learning system will inevitably encounter situations which it cannot handle, therefore APPL relies on classical systems throughout the entire robot deployment period to assure safety and explainability, in contrast to only during the initial phase. The learning only happens at the hyper-parameter level to fine-tune the underlying classical planner to address different deployment scenarios.

Furthermore, while Lin's work [41] proposed to use teaching to accelerate reinforcement learning of complex tasks, the learning overhead of training has been largely alleviated by the development of faster computation hardware, e.g., GPUs, and better function approximators, e.g., deep neural networks, in the last three decades. Leveraging the reduced learning cost, APPL can train a parameter policy in a variety of simulation environments, i.e., the BARN dataset, with a highly parallelizable, containerized, distributed training system [43], in order to cover as many real-world environments as possible in simulation. During real-world deployment, when out-of-distribution scenarios are inevitably encountered, non-expert users can help the robots to adapt to these unseen scenarios using selected interaction modalities. Although it is beyond the scope of this paper, another computational factor is onboard computation [44,45] for the multi-model learning during deployment. Due to limited onboard resources, such learning may not take place on the robot,

and may need to be uploaded to offboard computation resources to assure the parameter policy can be updated quickly and transferred back to the robot during a particular deployment. How to allocate both onboard and offboard computation and how to balance the trade-off between high computation power and communication latency remains an interesting future research direction.

10. Conclusions

We present APPL, a learning paradigm that leverages different human interaction modalities, including demonstration, interventions, evaluative feedback, and unsupervised reinforcement learning, to dynamically fine-tune classical navigation planner parameters to achieve better navigation performance. In addition to the individual APPL methods and experiments, we also introduce a cycle-of-learning scheme, in which different human interaction modalities can be utilized to continually improve future navigation performance in a cyclic fashion. One interesting direction for future work is to investigate other methods to incorporate human interaction to RL training in simulation, e.g., using inverse reinforcement learning to infer an underlying reward function, or creating high fidelity simulations similar to real-world deployment environments.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: The University of Texas at Austin, US Army Research Laboratory, George Mason University.

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, NRI-1925082), ONR, United States of America (N00014-18-2243), FLI (RFP2-000), ARO, United States of America (W911NF19-2-0333), DARPA, United States of America, Lockheed Martin, United States of America, GM, United States of America, and Bosch. 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.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.robot.2022.104132>.

References

- [1] I. Farkhatdinov, J.-H. Ryu, J. Poduraev, A user study of command strategies for mobile robot teleoperation, *Intell. Serv. Robot.* 2 (2) (2009) 95–104.
- [2] X. Xiao, B. Liu, G. Warnell, J. Fink, P. Stone, APPLD: Adaptive planner parameter learning from demonstration, *IEEE Robot. Autom. Lett.* 5 (3) (2020) 4541–4547.
- [3] Z. Wang, X. Xiao, B. Liu, G. Warnell, P. Stone, APPLI: Adaptive planner parameter learning from interventions, in: 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2021, pp. 6079–6085.
- [4] Z. Wang, X. Xiao, G. Warnell, P. Stone, APPL: Adaptive planner parameter learning from evaluative feedback, *IEEE Robot. Autom. Lett.* 6 (4) (2021) 7744–7749.
- [5] Z. Xu, G. Dharamkar, A. Nair, X. Xiao, G. Warnell, B. Liu, Z. Wang, P. Stone, APPLR: Adaptive planner parameter learning from reinforcement, in: 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2021, pp. 6086–6092.
- [6] S. Quinlan, O. Khatib, Elastic bands: Connecting path planning and control, in: [1993] Proceedings IEEE International Conference on Robotics and Automation, IEEE, 1993, pp. 802–807.
- [7] D. Fox, W. Burgard, S. Thrun, The dynamic window approach to collision avoidance, *IEEE Robot. Autom. Mag.* 4 (1) (1997) 23–33.
- [8] K. Zheng, Ros navigation tuning guide, 2017, arXiv preprint arXiv:1706.09068.
- [9] X. Xiao, J. Dufek, T. Woodbury, R. Murphy, UAV assisted USV visual navigation for marine mass casualty incident response, in: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2017, pp. 6105–6110.
- [10] B. Liu, X. Xiao, P. Stone, A lifelong learning approach to mobile robot navigation, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 1090–1096.
- [11] X. Xiao, B. Liu, G. Warnell, P. Stone, Motion planning and control for mobile robot navigation using machine learning: a survey, *Auton. Robots* (2022) 1–29.
- [12] J. Liang, U. Patel, A.J. Sathyamoorthy, D. Manocha, CrowdSteer: Realtime smooth and collision-free robot navigation in dense crowd scenarios trained using high-fidelity simulation, 2020, arXiv preprint arXiv:2004.03089.
- [13] X. Xiao, J. Biswas, P. Stone, Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain, *IEEE Robot. Autom. Lett.* 6 (3) (2021) 6054–6060.
- [14] S. Thrun, An approach to learning mobile robot navigation, *Robot. Auton. Syst.* 15 (4) (1995) 301–319.
- [15] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, C. Cadena, From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots, in: 2017 IEEE International Conference on Robotics and Automation (Icra), IEEE, 2017, pp. 1527–1533.
- [16] X. Xiao, B. Liu, P. Stone, Agile robot navigation through hallucinated learning and sober deployment, in: 2021 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2021, pp. 7316–7322.
- [17] X. Xiao, B. Liu, G. Warnell, P. Stone, Toward agile maneuvers in highly constrained spaces: Learning from hallucination, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 1503–1510.
- [18] Z. Wang, X. Xiao, A.J. Nettekoven, K. Umasankar, A. Singh, S. Bommakanti, U. Topcu, 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, pp. 148–153.
- [19] B.D. Argall, S. Chernova, M. Veloso, B. Browning, A survey of robot learning from demonstration, *Robot. Auton. Syst.* 57 (5) (2009) 469–483.
- [20] S. Siva, M. Wigness, J. Rogers, H. Zhang, Robot adaptation to unstructured terrains by joint representation and apprenticeship learning, in: *Robotics: Science and Systems (RSS)*, 2019.
- [21] V.G. Goecks, G.M. Gremillion, V.J. Lawhern, J. Valasek, 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.
- [22] W.B. Knox, P. Stone, C. Breazeal, Training a robot via human feedback: A case study, in: *International Conference on Social Robotics*, Springer, 2013, pp. 460–470.
- [23] N.R. Waytowich, V.G. Goecks, V.J. Lawhern, Cycle-of-learning for autonomous systems from human interaction, in: *AI-HRI Symposium, AAAI Fall Symposium Series*, 2018.
- [24] R.S. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.
- [25] H.-T.L. Chiang, A. Faust, M. Fiser, A. Francis, Learning navigation behaviors end-to-end with actor, *IEEE Robot. Autom. Lett.* 4 (2) (2019) 2007–2014.
- [26] Z. Xu, X. Xiao, G. Warnell, A. Nair, P. Stone, Machine learning methods for local motion planning: A study of end-to-end vs. parameter learning, in: 2021 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), IEEE, 2021, pp. 217–222.
- [27] W.B. Knox, A. Allievi, H. Banzhaf, F. Schmitt, P. Stone, Reward (mis) design for autonomous driving, 2021, arXiv preprint arXiv:2104.13906.
- [28] G. Kahn, P. Abbeel, S. Levine, Badgr: An autonomous self-supervised learning-based navigation system, *IEEE Robot. Autom. Lett.* 6 (2) (2021) 1312–1319.
- [29] H. Karnan, G. Warnell, X. Xiao, P. Stone, Voila: Visual-observation-only imitation learning for autonomous navigation, in: 2022 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2022.
- [30] S. Aminikhanghahi, D.J. Cook, A survey of methods for time series change point detection, *Knowl. Inf. Syst.* 51 (2) (2017) 339–367.
- [31] D.A. Pomerleau, Alvin: An autonomous land vehicle in a neural network, in: *Advances in Neural Information Processing Systems*, 1989, pp. 305–313.

- [32] S. Niekum, S. Osentoski, C.G. Atkeson, A.G. Barto, Online bayesian change-point detection for articulated motion models, in: 2015 IEEE International Conference on Robotics and Automation (ICRA), IEEE, 2015, pp. 1468–1475.
- [33] N. Hansen, S.D. Müller, P. Koumoutsakos, Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES), *Evol. Comput.* 11 (1) (2003) 1–18.
- [34] M. Sensoy, L. Kaplan, M. Kandemir, Evidential deep learning to quantify classification uncertainty, in: *Advances in Neural Information Processing Systems*, 2018, pp. 3179–3189.
- [35] D. Perille, A. Truong, X. Xiao, P. Stone, Benchmarking metric ground navigation, in: 2020 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), IEEE, 2020, pp. 116–121.
- [36] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing atari with deep reinforcement learning, 2013, arXiv preprint [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- [37] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., Soft actor-critic algorithms and applications, 2018, arXiv preprint [arXiv:1812.05905](https://arxiv.org/abs/1812.05905).
- [38] S. Fujimoto, H. van Hoof, D. Meger, Addressing function approximation error in actor-critic methods, 2018, [arXiv:1802.09477](https://arxiv.org/abs/1802.09477).
- [39] A. Nair, P. Srinivasan, S. Blackwell, C. Alciçek, R. Fearon, A.D. Maria, V. Panneershelvam, M. Suleyman, C. Beattie, S. Petersen, S. Legg, V. Mnih, K. Kavukcuoglu, D. Silver, Massively parallel methods for deep reinforcement learning, 2015, [arXiv:1507.04296](https://arxiv.org/abs/1507.04296).
- [40] V.G. Goecks, G.M. Gremillion, V.J. Lawhern, J. Valasek, N.R. Waytowich, Integrating behavior cloning and reinforcement learning for improved performance in dense and sparse reward environments, in: *International Conference on Autonomous Agents and Multi Agent Systems*, 2020.
- [41] L.-J. Lin, Self-improving reactive agents based on reinforcement learning, planning and teaching, *Mach. Learn.* 8 (3) (1992) 293–321.
- [42] W.D. Smart, L.P. Kaelbling, Effective reinforcement learning for mobile robots, in: *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, Vol. 4, IEEE, 2002, pp. 3404–3410.
- [43] A. Nair, Z. Xu, G. Dhamankar, X. Xiao, Using parallelized containers for reinforcement learning on large computer clusters, 2021.
- [44] R. Dromnelle, E. Renaudo, G. Pourcel, R. Chatila, B. Girard, M. Khamassi, How to reduce computation time while sparing performance during robot navigation? A neuro-inspired architecture for autonomous shifting between model-based and model-free learning, in: *Conference on Biomimetic and Biohybrid Systems*, Springer, 2020, pp. 68–79.
- [45] R. Dromnelle, B. Girard, E. Renaudo, R. Chatila, M. Khamassi, Coping with the variability in humans reward during simulated human-robot interactions through the coordination of multiple learning strategies, in: 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), IEEE, 2020, pp. 612–617.



Xuesu Xiao is an incoming Assistant Professor in the Department of Computer Science at George Mason University starting Fall 2022. Currently, he is a roboticist on The Everyday Robot Project at X, The Moonshot Factory, and a research affiliate in the Department of Computer Science at The University of Texas at Austin. Dr. Xiao's research focuses on field robotics, motion planning, and machine learning. He develops highly capable and intelligent mobile robots that are robustly deployable in the real world with minimal human supervision. Dr. Xiao received his Ph.D. in Computer Science from Texas A&M University in 2019, Master of Science in Mechanical Engineering from Carnegie Mellon University in 2015, and dual Bachelor of Engineering in Mechatronics Engineering from Tongji University and FH Aachen University of Applied Sciences in 2013.