

Adaptive and Explainable Deployment of Navigation Skills via Hierarchical Deep Reinforcement Learning

Kywoon Lee^{*1}, Seongun Kim^{*2} and Jaesik Choi²

Abstract—For robotic vehicles to navigate robustly and safely in unseen environments, it is crucial to decide the most suitable navigation policy. However, most existing deep reinforcement learning based navigation policies are trained with a hand-engineered curriculum and reward function which are difficult to be deployed in a wide range of real-world scenarios. In this paper, we propose a framework to learn a family of low-level navigation policies and a high-level policy for deploying them. The main idea is that, instead of learning a single navigation policy with a fixed reward function, we simultaneously learn a family of policies that exhibit different behaviors with a wide range of reward functions. We then train the high-level policy which adaptively deploys the most suitable navigation skill. We evaluate our approach in simulation and the real world and demonstrate that our method can learn diverse navigation skills and adaptively deploy them. We also illustrate that our proposed hierarchical learning framework presents explainability by providing semantics for the behavior of an autonomous agent.

I. INTRODUCTION

Autonomous navigation of mobile robots has gained much interest due to a wide variety of important applications in industry. These include assistive robots [1], a last-mile delivery [2], a guidance at airports [3], and warehouse navigation [4], among others. To navigate a robot to a desired goal in complex, cluttered, and highly dynamic environments with multiple static and dynamic obstacles, a reliable and robust navigation policy is essential.

Traditional approaches to address the navigation problem typically consist of a series of modules, each of which is specifically designed for solving a particular sub-task of navigation problems such as human detection, prediction of future trajectories of humans, and path planning [5], [6]. However, these approaches require extensive computational efforts and rely on manually engineered parameters, which limit the ability of mobile robots to operate in previously unseen environments, or across different robotic platforms, as the modular system suffers from a lack of generalization and sub-optimal performance [7]. In this regard, deep reinforcement learning (DRL) approaches have recently been proposed as learning based solutions for autonomous navigation, which directly map raw sensor observations to controls and have shown remarkable results compared to

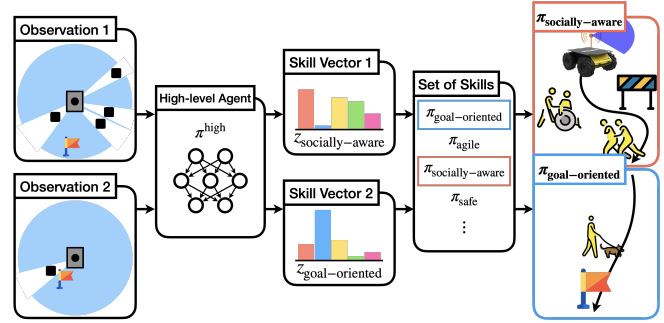


Fig. 1. An overview of the proposed hierarchical framework of navigation skills. A high-level policy invokes low-level navigation skills from a raw sensory observation. A low-level policy is adopted from a continuous skill vector among the infinite number of skills, which then drives a robot while understanding the context.

traditional methods and robustness to sensor noise [8], [9], [10], [11]. However, existing DRL-based approaches require hand-engineered reward functions that can be exceedingly time consuming. Moreover, in challenging scenarios, DRL-based approaches with a fixed reward function easily get stuck in local optima which makes complex situations including navigating across narrow corridors or dense crowds and turning corners, unsolvable.

Another potential problem of existing DRL-based robot navigation is that a policy represented by a deep neural network often lacks transparency and cannot provide explanations on decision-making reasons. Recent works investigate DRL methods which are interpretable and provide decision-making reasons [12], [13]. However, they are only applicable in video games or in simple simulation environments. Another line of research provides explanations in real-world robots by highlighting input features that a deep network most refers to when making a decision [14]. Nevertheless, this approach has the limitation that it is not directly relevant to increasing the performance of the policy model.

To address the aforementioned problems, we develop a framework of learning to adaptively deploy navigation skills that are explainable with hierarchical reinforcement learning (HRL).¹ Specifically, our approach can be divided into two phases. The first phase trains a family of low-level navigation policies, each of which is optimized for a particular skill vector in a continuous representation. This skill vector is associated with a corresponding reward function. For example, a skill vector that imposes an additional penalty for the robot to be more cautious with obstacles can exhibit socially-

¹Source codes are available at <https://github.com/leekwoon/hrl-nav>

^{*}Equal contribution

¹ Department of Computer Science and Engineering, Ulsan National Institute of Science and Technology, Ulsan, Republic of Korea. leekwoon@unist.ac.kr.

² Graduate School of Artificial Intelligence, Korea Advanced Institute of Science and Technology, 291 Daehak-ro, Daejeon 34141, Republic of Korea. {seongun, jaesik.choi}@kaist.ac.kr Correspondence to: Jaesik Choi

aware navigation behavior, while a skill vector encouraging the robot getting closer to its goal may result in more aggressive behavior. In the second phase, we train the high-level policy which deploys the most suitable navigation skill for every time step. The main contributions of this work can be summarized as follows:

- Proposal of a hierarchical reinforcement learning approach which learns diverse navigation skills and deploys them.
- Extensive evaluation of our approach on various scenarios including a real-world robot navigation which demonstrates effectiveness and explainability of our approach.

II. RELATED WORK

A. Reinforcement Learning for Navigation

Recent advances in deep reinforcement learning have enabled the solving of complex navigation tasks from raw sensory measurements. Long et al. [8] propose a DRL-based multi-agent navigation framework to train collision avoidance policy with proximal policy optimization (PPO) [15] using 2D LiDAR observation. This DRL-based collision avoidance policy is further integrated into a hybrid control framework [16] and a conventional global planner from robot operating system (ROS) [17]. A similar approach is proposed in [10], however, additional ego-safety and social-safety rewards are used to consider human-awareness. Apart from using 2D LiDAR observation, navigation that learns from RGB images has also been studied [18], [19].

B. Reinforcement Learning with Sparse Rewards

Existing approaches which handle sparse rewards in reinforcement learning (RL) can be divided into two categories, curriculum reinforcement learning and reward shaping. In the context of curriculum reinforcement learning, the implicit curriculum uses goal relabeling techniques which randomly sample goals from failed trajectories [20], [21]. The explicit curriculum considers the difficulty of goals during training, e.g., by generating the goal further from the start in navigation or increasing the number of obstacles [22], [8]. The reward shaping technique modifies the reward signal by learning a parameterized dense reward function [23]. Sibling rivalry [24] uses a model-free, dynamic reward shaping method that preserves optimal policies on sparse-rewards tasks. On the other hand, AutoRL [25] uses large-scale hyperparameter optimization to shape the reward.

C. Hybrid Control framework for Navigation

To further improve the robustness and effectiveness of navigation policy by leveraging the strength of multiple local planners, a hybrid control framework has been widely adopted which uses a high-level switching controller to manage a set of low-level control rules. For example, Jin et al. [26] propose hand-designed switching rules to combine goal-navigation and obstacle avoidance, Shucker et al. [27] introduce switching rules to handle challenging cases where collisions or noises cannot be handled appropriately, and

Fan et al. [16] classify the scenarios into three cases by considering a robot's sensor measurement and using both a PID and DRL-based controller. Most similar to our work, Kastner et al. [28] propose a hierarchical navigation system integrating model-based and learning-based local planners by training an agent which decides between multiple local planners. However, the approach uses only one DRL-based local planner with the traditional model-based planners. As a result, navigation systems following this approach are often strongly dependent on deployed model-based planners. Instead, we train a family of policies that exhibit different behaviors and focus on how to adaptively deploy them.

III. PROBLEM FORMULATION

HRL decomposes a general RL into a hierarchy of multiple sub-problems which themselves are RL problems [29]. Higher-level problems abstract an original RL problem and adopt which lower-level problems to solve, while lower-level problems are defined to solve the original problem given an abstraction from the higher-level problems.

In this work, we decompose a problem of learning navigation skills into a hierarchy of two sub-problems, a high-level problem and a low-level problem. The high-level problem is to find a macro policy which understands the context of a navigation environment and adopts which navigation skills to invoke. The macro policy observes a current state and demonstrates the context that the agent encounters. The low-level problem is to discover a set of navigation skills that drives a mobile robot to reach a target position. Each low-level navigation skill issues a command by giving primitive actions such as motor velocities, conditioned on the context distinguished by the macro policy.

We formalize each sub-problem as a Markov decision process (MDP), in particular, a goal-conditioned MDP [30] for learning general-purpose navigation skills. A high-/low-level problem is modeled as an MDP $\langle \mathcal{S}^{\text{high/low}}, \mathcal{A}^{\text{high/low}}, \mathcal{G}, P^{\text{high/low}}, R_g^{\text{high/low}}, \gamma \rangle$, where \mathcal{S} is the set of states, \mathcal{G} is the set of goals, \mathcal{A} is the set of actions, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, +\infty)$ is the transition probability, $R_g : \mathcal{S} \times \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}$ is the goal-conditioned reward function, $\gamma \in [0, 1]$ is the discount factor, and each superscript high and low represents a notation for the high-level and the low-level problem respectively.

To build the hierarchy, an action from the high-level agent $a^{\text{high}} \in \mathcal{A}^{\text{high}}$ is observed as a state by the low-level agent, where this high-level action implies the context of the environment. In addition to a^{high} , the low-level agent observes the same state as the high-level agent, which results in $s^{\text{low}} = (s^{\text{high}}, a^{\text{high}}) \in \mathcal{S}^{\text{low}}$. While interacting with the environment and observing the state s^{low} , each low-level agent receives a reward from a reward function $R_g^{\text{low}}(s^{\text{low}}, a^{\text{low}}, a^{\text{high}})$ that consists of multiple reward terms which we will describe in more detail in the following section. To effectively train each low-level agent to present discriminative primitive actions conditioned on a^{high} , they are trained with different reward functions by weighting each reward term in R_g^{low} by a^{high} .

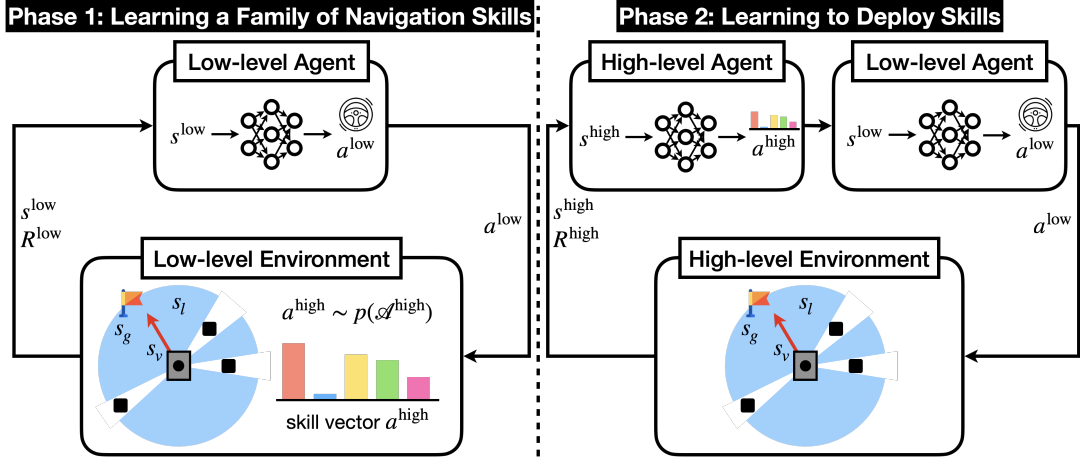


Fig. 2. A two-phase training procedure of the proposed hierarchical framework. A family of navigation skills is trained in the first phase, while a high-level policy is trained in the second phase. To train the infinite number of navigation skills in the first phase, a single policy network which is conditioned on a continuous skill vector is utilized. The skill vector a^{high} is sampled from a fixed distribution per episode to construct s^{low} along with a scan observation. The high-level agent is then trained by taking actions from the low-level skill deployed by the high-level agent.

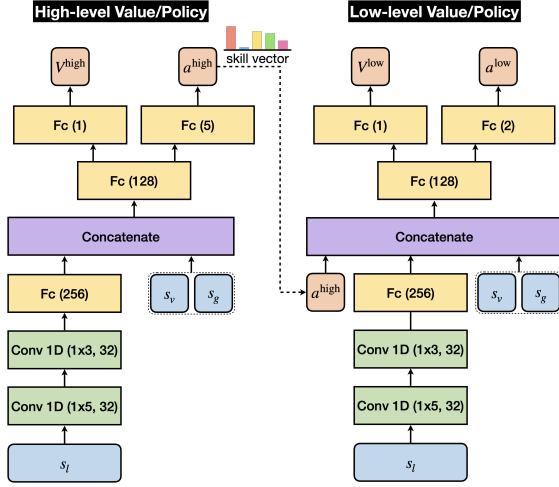


Fig. 3. Actor-critic architecture used for learning a family of navigation skills (left) and learning to deploy skills (right). Conv denotes a convolutional layer and Fc denotes a fully-connected layer.

Under this problem definition, the optimal low-level policies and the high-level policy in the hierarchy are achieved by an off-the-shelf RL algorithm which maximizes a corresponding value function formulated as follows:

$$V^{\text{low}}(s^{\text{low}}, g | \pi^{\text{high}}) \triangleq \mathbb{E}_{\substack{a_t^{\text{low}} \sim \pi^{\text{low}}(a_t^{\text{low}} | s_t^{\text{low}}, g), \\ a_t^{\text{high}} \sim \pi^{\text{high}}(a_t^{\text{high}} | s_t^{\text{high}}, g), \\ s_{t+1}^{\text{low}} \sim P(s_{t+1}^{\text{low}} | s_t^{\text{low}}, a_t^{\text{low}})}} \left[\sum_{t=0}^{\infty} \gamma^t R_g^{\text{low}}(s_t^{\text{low}}, a_t^{\text{low}}, a_t^{\text{high}}) \middle| s_0^{\text{low}} = s^{\text{low}} \right], \quad (1)$$

$$V^{\text{high}}(s^{\text{high}}, g | \pi^{\text{low}}) \triangleq \mathbb{E}_{\substack{a_t^{\text{high}} \sim \pi^{\text{high}}(a_t^{\text{high}} | s_t^{\text{high}}, g), \\ a_t^{\text{low}} \sim \pi^{\text{low}}(a_t^{\text{low}} | s_t^{\text{low}}, g), \\ s_{t+1}^{\text{high}} \sim P(s_{t+1}^{\text{high}} | s_t^{\text{high}}, a_t^{\text{high}})}} \left[\sum_{t=0}^{\infty} \gamma^t R_g^{\text{high}}(s_t^{\text{high}}, a_t^{\text{high}}, a_t^{\text{low}}) \middle| s_0^{\text{high}} = s^{\text{high}} \right]. \quad (2)$$

IV. APPROACH

In the hierarchy, the high-level agent and the low-level agent interact with each other. To train agents in the hierarchy, we decompose a training procedure into two phases. In this section, we describe how each component of MDP is defined and how the agent is trained for each training phase.

A. Learning a Family of Navigation Skills

In the first phase, we learn a family of low-level navigation policies, each optimized for a particular skill vector in continuous representation. We present a detailed setup of phase 1 as follows:

1) *State space*: A state s_t^{low} at time step t consists of four components: the raw 2D LiDAR measurements $s_t^l \in \mathbb{R}^{512}$ casting 512 rays over 360° with up to 6m range, the linear and angular velocity of the robot $s_t^v \in \mathbb{R}^2$, the relative goal state $s_t^g \in \mathbb{R}^2$ represented in the polar coordinate, and the high-level action $a_t^{\text{high}} \in \mathbb{R}^5$. We call a_t^{high} a *skill vector* since it induces a specific behavior by weighting a number of reward terms.

2) *Action space*: Considering nonholonomic kinematics constraints, an action $a_t^{\text{low}} \in \mathbb{R}^2$ is composed of the linear velocity $v_t \in [0, 0.5]$ and angular velocity $w_t \in [-0.64, 0.64]$ which is normalized into $[-1, 1]$ in the neural network.

3) *Reward function*: Our goal is to learn a family of navigation policies with a diverse set of behavioral characteristics. Thus, we set the reward function parameterized by the skill vector a^{high} as follows:

$$R_t^{\text{low}}(s_t^{\text{low}}, a_t^{\text{low}}, a_t^{\text{high}}) = r_{\text{success}} + a^{\text{high}} [r_{\text{collision}} \ r_{\text{progress}} \ r_v \ r_w \ r_{\text{safety}}]^\top.$$

The first term, r_{success} , is the success reward, which is 1 if the goal is achieved and 0 otherwise. The collision reward, $r_{\text{collision}}$ is -1 when the robot collides with obstacles and 0 otherwise. The progress reward, r_{progress} is the difference between the previous distance to the goal and the current

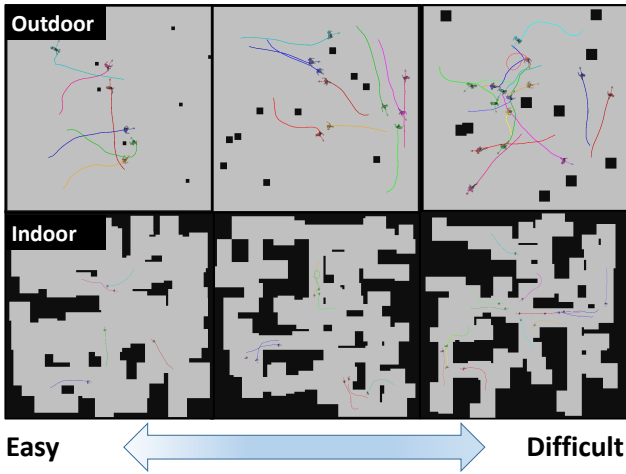


Fig. 4. Training environments used to train the policies, with both randomized maps and obstacles. We randomize training environments so that the agent learns diverse skills. To control the difficulty, we manipulate the number of humans, obstacles, and corridor and the width of corridor.

distance to the goal. The driving reward, r_v is the positive linear speed which encourages the robot to drive as fast as possible. The turning reward, r_w is the negative angular speed for smooth driving. Finally, inspired by the social-safety zone in [10], r_{safety} is the safety reward that gives a penalty with value $1 - \frac{d_t}{r+0.5}$ when the robot enters the safety zone, where d_t is the distance to the closest obstacle and r is the robot radius.

4) *Network architecture*: The neural network architecture used for learning a family of navigation skills is illustrated in the right of Fig. 3. We follow the architecture in the prior work [8], except that the policy and the value function are additionally parameterized by the skill vector a^{high} which is concatenated with the output of the third layer, the current velocity s_v , and the relative goal state s_g .

5) *Training procedure*: As each individual policy is trained under the similar structure of reward function, we use a single deep neural network to train the entire family of low-level policies simultaneously. Inspired by a universal policy [31], we train a policy $\pi^{\text{low}} : (s_t^{\text{low}}, a_t^{\text{high}}) \rightarrow a_t^{\text{low}}$ that is conditioned on both state of the robot s_t^{low} and the skill vector a_t^{high} . In contrast to existing approaches which learn a single policy, at the beginning of the episode, we sample a skill vector from predefined distribution and fix it during the rollout to train corresponding policy $\pi_{a_t^{\text{high}}}^{\text{low}} : s_t^{\text{low}} \rightarrow a_t^{\text{low}}$, which we call a skill.

B. Learning to Deploy Skills

The second phase of our approach is to learn a high-level policy to adaptively deploy the navigation skills learned in phase 1. The following provides a detailed formulation.

1) *State space*: A state s_t^{high} composed of three components, s_t^l , s_t^v and s_t^g , which are used in the first phase.

2) *Action space*: An action $a_t^{\text{high}} \in \mathbb{R}^5$ is a skill vector that decides the learned low-level behavior characteristic.

3) *Reward function*: The objective of the second phase is to minimize the time to reach the desired goal and we use

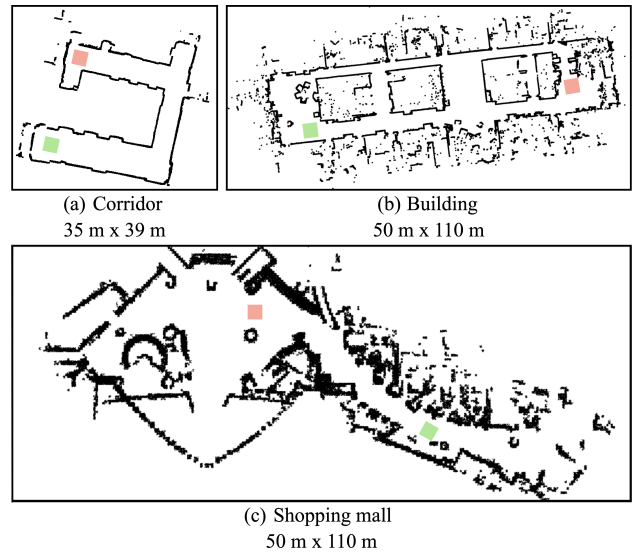


Fig. 5. Unseen environments used to evaluate the policies. The robot navigates from the green square to the goal is located on the red square.

the following sparse reward function:

$$R_t^{\text{high}}(s_t^{\text{high}}, a_t^{\text{high}}, a_t^{\text{low}}) = \begin{cases} 0 & \text{if reach the goal} \\ -1 & \text{otherwise,} \end{cases}$$

where the agent gets the reward of 0 when the goal is achieved, and -1 otherwise.

4) *Network architecture*: We parameterize each of the policies and value functions with the architecture from the prior work [8], except that the policy predicts a particular skill vector that will be used by the low-level policy. We further feed the output of high-level policy as an input to the low-level policy to infer the command velocity of the robot during a rollout as shown in Fig. 3.

5) *Training procedure*: After finishing training the set of navigation skills, we can further perform hierarchical control by training the high-level policy. During a rollout, the high-level policy predicts the skill vector which decides the behavior characteristic. This skill vector is observed as an additional input to the low-level policy which outputs the command velocity of a robot. Note that, in the second phase, gradients flow only through the high-level policy, not the low-level policy. Due to a problem with the sparse reward setting where a positive reward is only provided when the goal is achieved, we use the hindsight experience replay (HER) technique [20] which provides an effective way to handle the sparse reward challenge by revisiting previous states in the replay buffer and storing additional trajectories with hindsight goals for generating reward signals.

V. EXPERIMENTS AND RESULTS

A. Setup

We train a policy using an OpenAI-gym-compatible simulator that we specially design to integrate it into the robot operating system (ROS) with the goal of open-sourcing it for easy comparison of various approaches including state-of-the-art learning-based approaches and conventional ones




Fig. 6. Quantitative results in three unseen environments with 100 episodes. We evaluate the proposed method by comparing four metrics: success rate, collision rate, time to reach the goal, and path length. Our hierarchical method (blue bar) presents comparative performance in all evaluation environments.




Fig. 7. Results of the qualitative analysis. We compare the average value of the skill vector with standard deviation among three evaluation environments. The high-level agent deploys slow but careful skills in the first environment, while it deploys smooth but fast skills in the other environments.

(see Fig. 4). To create a scenario with various difficulty levels, we adopt the map generation method from [28] to generate indoor (50m×50m) and outdoor (20m×20m) maps. We randomize the noise on sensors, and static and dynamic obstacles where we control the motion of the dynamic obstacles using a learning-based collision avoidance method that extends the prior work [8] with A^* global planner to make it follow the sub-goal generated from the global path for more realistic long-range behavior. We terminate the simulation when reaching the goal, colliding with an obstacle, or exceeding 1000 time steps.

To train the policies, we use soft actor-critic (SAC) [32] with an automatically tuned entropy coefficient, employing the Adam optimizer [33] with a learning rate $1e-4$ for both the actor and the critic, a target smoothing coefficient of $5e-4$, and a replay buffer size of 1M.

We compare our approach with a traditional model-based approach dynamic window approach (DWA), two DRL-based approaches LONG [8] and GRING [17], and four agents trained with a fixed skill vector which we term $\pi_{\text{goal-oriented}}$, $\pi_{\text{socially-aware}}$, π_{safe} , and π_{mild} where each policy is trained with the skill vector of $[0, 0.2, 0.1, 0, 0]$, $[0.1, 0, 0, 0.1, 0.2]$, $[0.2, 0, 0, 0.1, 0.1]$, and $[0.1, 0.1, 0.1, 0.2, 0]$ respectively.

B. Quantitative Analysis

We evaluate the trained policies in three unseen large environments (see Fig. 5). These scenarios include (a) a corridor with 4 dynamic obstacles, (b) a building with 9 dynamic obstacles, and (c) a shopping mall with 8 dynamic obstacles which are designed to evaluate the ability of the agent including global planning, collision avoidance, and safe navigation, among other real-world situations. We validate the effectiveness of our method with four metrics: success

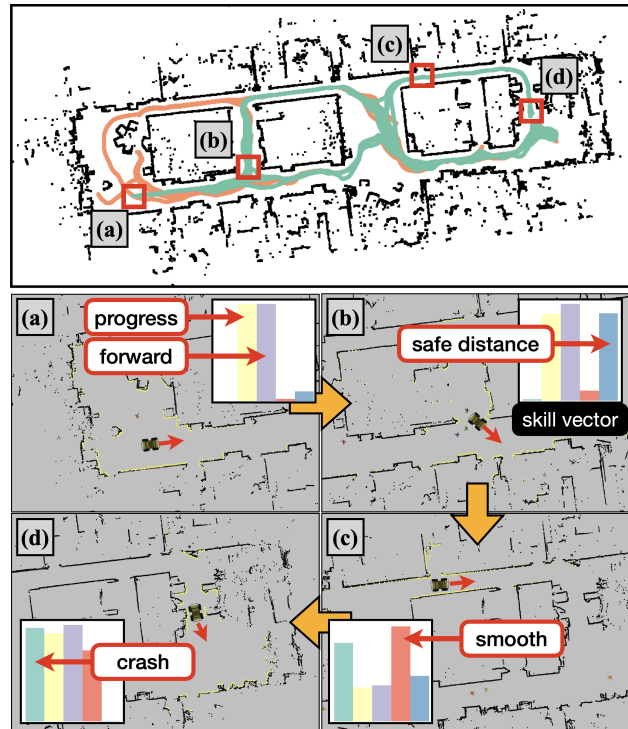


Fig. 8. Trajectories of the our approach in the unseen building environment (top) and skill vectors in particular situations (bottom).

rate, collision rate, arriving time, and path length. Fig. 6 shows the evaluation results with 100 episodes in each scenario. To solely evaluate the performance of learned navigation policy, all approaches do not use a global planner that is frequently used to handle a long-range navigation challenge. A traditional model-based approach DWA fails in dense crowds scenarios. Further, in all evaluation environments, our approach shows comparative performance to other baselines including DRL-based agents and agents trained with the fixed skill vector, except for goal-oriented policy. LONG and GRING present comparative results in densely crowded environments where a local navigation policy can solve but performance drops significantly for a situation where a global navigation policy is required. An interesting observation is that $\pi_{\text{goal-oriented}}$ performs well in long-range and dynamic environments without crash or safety rewards typically used in prior works. We assume that this is due to the goal-oriented reward, which encourages exploration

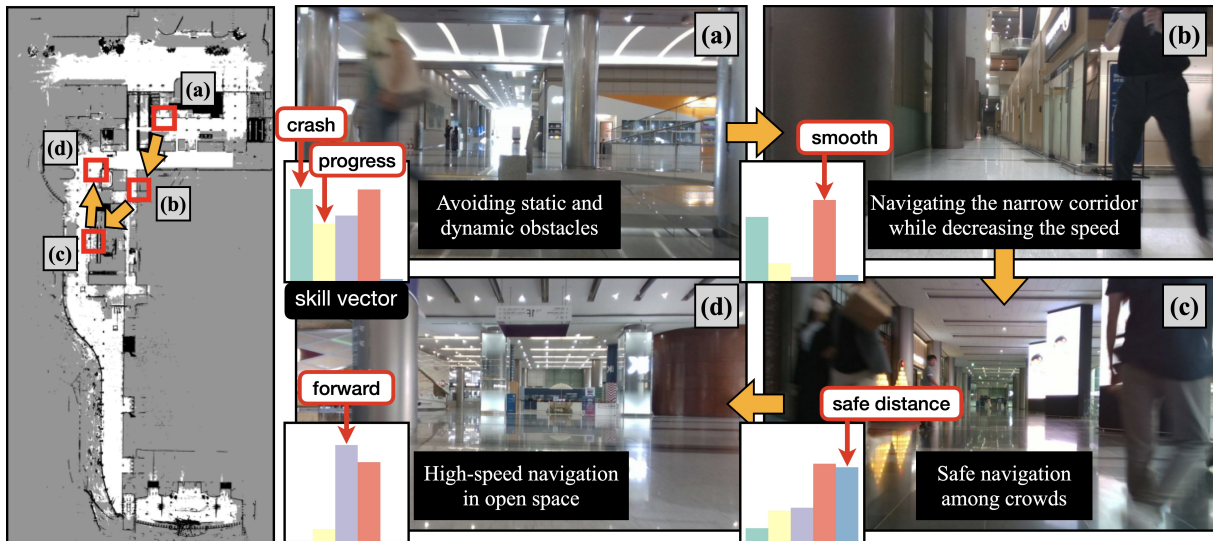


Fig. 9. Demonstration of our approach in unseen real scenarios. The scenarios are navigating in (a) cluttered environment, (b) narrow corridor, (c) dense crowd and (d) open space. The skill vectors are predicted by transferred high-level policy.

and goal achievement. It would be an appealing research direction to study the role of reward terms rigorously.

C. Qualitative Analysis

To qualitatively analyze the effectiveness of the proposed approach, we interpret the average skill vector along the trajectory for each unseen evaluation environment in Fig. 7. In the corridor environment, there is a long wall. To successfully reach the goal, the robot needs to detour a long range. To this end, the high-level agent deploys skills that present slow but careful motor commands with relatively high crash skill value. On the other hand, in the other two environments, the building and the shopping mall, there are little static obstacles on the straight line from the initial position of the robot to the target position. In these environments, the skill vector has large progress and forward values so that the robot presents agile movement. In addition, the high-level agent demonstrates smooth motor command through the smooth skill value. The average discomfort skill value in the last two crowded environments is relatively large, which validates the effectiveness of our approach.

The whole trajectories of our approach in the building environment are shown in Fig. 8 where green lines represent trajectories of successes. The skill vectors demonstrate the interpretable deployment of learned navigation skills. At the beginning where there are few obstacles around the robot as shown in (a) of Fig. 8, progress and forward skill values are huge for fast navigation towards the goal. In the middle of the navigation shown in (b), though progress and forward skill values are still large, the discomfort value increases when the robot is surrounded by humans, which would lead the robot to avoid collision with them. When the robot enters the narrow corridor in (c), the smooth value increases to prevent oscillatory behaviors. Around the goal which is close to obstacles in (d), crash is high to safely arrive the goal. The skill vector presents the same tendency across the other

evaluation scenarios. We refer the reader to the videos in supplementary material for further information.

D. Real-World Experiment

To demonstrate the validity of the proposed approach, we collect the expert trajectory with our robot which has two 270° LiDARs that are merged together to cover 360° . We then evaluate how well the high-level policy understands the context in a real-world shopping mall. Fig. 9 shows the skill vector generated by high-level policy for different scenarios. The high-level policy encourages the robot to keep a safe distance from crowds and to decrease the speed in narrow corridors. On the other hand, the goal-oriented navigation skill with high speed is preferred in open space.

VI. CONCLUSION

In this paper, we propose a hierarchical learning framework to train reinforcement learning based navigation policies that can present explainability by providing semantics of a behavior of an autonomous agent and can reduce the effort to design hand-engineered reward functions. The evaluation results on unseen environments including real-world cases demonstrate the adaptive and explainable deployment of learning navigation skills.

ACKNOWLEDGMENT

This work was supported by the Industry Core Technology Development Project, 20005062, Development of Artificial Intelligence Robot Autonomous Navigation Technology for Agile Movement in Crowded Space, funded by the Ministry of Trade, Industry & Energy (MOTIE, Republic of Korea) and was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00984, Development of Artificial Intelligence Technology for Personalized Plug-and-Play Explanation and Verification of Explanation, No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

REFERENCES

- [1] D. Feil-Seifer and M. J. Matarić, "Socially assistive robotics," *IEEE Robotics & Automation Magazine*, vol. 18, no. 1, pp. 24–31, 2011.
- [2] T. Hoffmann and G. Prause, "On the regulatory framework for last-mile delivery robots," *Machines*, vol. 6, no. 3, p. 33, 2018.
- [3] S. Kayukawa, K. Higuchi, J. Guerreiro, S. Morishima, Y. Sato, K. Kitani, and C. Asakawa, "Bbep: A sonic collision avoidance system for blind travellers and nearby pedestrians," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–12.
- [4] W. Xiao-Long, W. Chun-Fu, L. Guo-Dong, and C. Qing-Xie, "A robot navigation method based on rfid and qr code in the warehouse," in *2017 Chinese Automation Congress (CAC)*. IEEE, 2017, pp. 7837–7840.
- [5] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.
- [6] L. Liu, D. Dugas, G. Cesari, R. Siegwart, and R. Dubé, "Robot navigation in crowded environments using deep reinforcement learning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5671–5677.
- [7] A. J. Sathymoorthy, U. Patel, T. Guan, and D. Manocha, "Frozone: Freezing-free, pedestrian-friendly navigation in human crowds," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4352–4359, 2020.
- [8] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6252–6259.
- [9] T. Fan, X. Cheng, J. Pan, P. Long, W. Liu, R. Yang, and D. Manocha, "Getting robots unfrozen and unlost in dense pedestrian crowds," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1178–1185, 2019.
- [10] J. Jin, N. M. Nguyen, N. Sakib, D. Graves, H. Yao, and M. Jagersand, "Mapless navigation among dynamics with social-safety-awareness: a reinforcement learning approach from 2d laser scans," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 6979–6985.
- [11] D. Dugas, J. Nieto, R. Siegwart, and J. J. Chung, "Navrep: Unsupervised representations for reinforcement learning of robot navigation in dynamic human environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7829–7835.
- [12] D. Lee, H. Tang, J. O. Zhang, H. Xu, T. Darrell, and P. Abbeel, "Modular architecture for starcraft ii with deep reinforcement learning," in *Fourteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2018.
- [13] D. Trivedi, J. Zhang, S.-H. Sun, and J. J. Lim, "Learning to synthesize programs as interpretable and generalizable policies," *Advances in neural information processing systems*, vol. 34, pp. 25 146–25 163, 2021.
- [14] S. Kim and J. Choi, "Explaining the decisions of deep policy networks for robotic manipulations," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 2663–2669.
- [15] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [16] T. Fan, P. Long, W. Liu, and J. Pan, "Fully distributed multi-robot collision avoidance via deep reinforcement learning for safe and efficient navigation in complex scenarios," *arXiv preprint arXiv:1808.03841*, 2018.
- [17] R. Guldenring, M. Görner, N. Hendrich, N. J. Jacobsen, and J. Zhang, "Learning local planners for human-aware navigation in indoor environments," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6053–6060.
- [18] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3357–3364.
- [19] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, and J. Malik, "Cognitive mapping and planning for visual navigation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2616–2625.
- [20] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, O. Pieter Abbeel, and W. Zaremba, "Hindsight experience replay," *Advances in neural information processing systems*, vol. 30, 2017.
- [21] M. Fang, C. Zhou, B. Shi, B. Gong, J. Xu, and T. Zhang, "Dher: Hindsight experience replay for dynamic goals," in *International Conference on Learning Representations*, 2018.
- [22] C. Florensa, D. Held, X. Geng, and P. Abbeel, "Automatic goal generation for reinforcement learning agents," in *International conference on machine learning*. PMLR, 2018, pp. 1515–1528.
- [23] A. Y. Ng, D. Harada, and S. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *Icml*, vol. 99, 1999, pp. 278–287.
- [24] A. Trott, S. Zheng, C. Xiong, and R. Socher, "Keeping your distance: Solving sparse reward tasks using self-balancing shaped rewards," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [25] H.-T. L. Chiang, A. Faust, M. Fiser, and A. Francis, "Learning navigation behaviors end-to-end with autorl," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2007–2014, 2019.
- [26] J. Jin, Y. Kim, S. Wee, D. Lee, and N. Gans, "A stable switched-system approach to collision-free wheeled mobile robot navigation," *Journal of Intelligent & Robotic Systems*, vol. 86, no. 3, pp. 599–616, 2017.
- [27] B. Shucker, T. Murphey, and J. K. Bennett, "Switching rules for decentralized control with simple control laws," in *2007 American Control Conference*. IEEE, 2007, pp. 1485–1492.
- [28] L. Kastner, J. Cox, T. Buiyan, and J. Lambrecht, "All-in-one: A drl-based control switch combining state-of-the-art navigation planners," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 2861–2867.
- [29] B. Hengst, *Hierarchical Reinforcement Learning*. Boston, MA: Springer US, 2010, pp. 495–502.
- [30] T. Schaul, D. Horgan, K. Gregor, and D. Silver, "Universal value function approximators," in *International Conference on Machine Learning*. PMLR, 2015, pp. 1312–1320.
- [31] W. Yu, J. Tan, C. K. Liu, and G. Turk, "Preparing for the unknown: Learning a universal policy with online system identification," *arXiv preprint arXiv:1702.02453*, 2017.
- [32] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1861–1870.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.