

Efficient Exploration in Constrained Environments with Goal-Oriented Reference Path

Ota, Kei; Sasaki, Yoko; Jha, Devesh K.; Yoshiyasu, Yusuke; Kanezaki, Asako

TR2020-141 November 04, 2020

Abstract

In this paper, we consider the problem of building learning agents that can efficiently learn to navigate in constrained environments. The main goal is to design agents that can efficiently learn to understand and generalize to different environments using high-dimensional inputs (a 2D map), while following feasible paths that avoid obstacles in obstacle-cluttered environment. To achieve this, we make use of traditional path planning algorithms, supervised learning, and reinforcement learning algorithms in a synergistic way. The key idea is to decouple the navigation problem into planning and control, the former of which is achieved by supervised learning whereas the latter is done by reinforcement learning. Specifically, we train a deep convolutional network that can predict collision-free paths based on a map of the environment—this is then used by an reinforcement learning algorithm to learn to closely follow the path. This allows the trained agent to achieve good generalization while learning faster. We test our proposed method in the recently proposed Safety Gym suite that allows testing of safety-constraints during training of learning agents. We compare our proposed method with existing work and show that our method consistently improves the sample efficiency and generalization capability to novel environments.

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)

Efficient Exploration in Constrained Environments with Goal-Oriented Reference Path

Kei Ota^{1,2}, Yoko Sasaki², Devesh K. Jha³, Yusuke Yoshiyasu², and Asako Kanezaki²

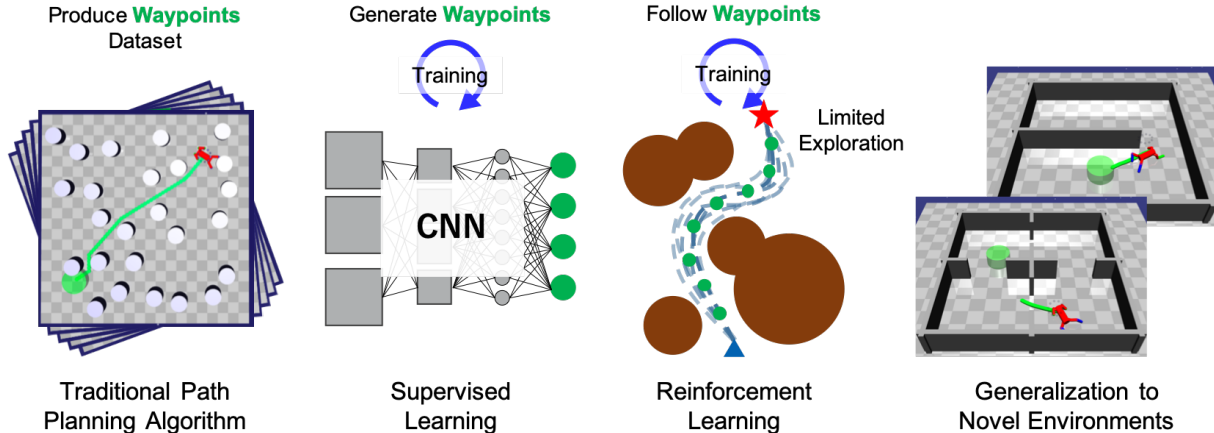


Fig. 1: We decouple planning and control by combining traditional path planning algorithms, supervised learning (SL) and reinforcement learning (RL) algorithms in a synergistic way. First, a path planning algorithm such as A* search generates a set of pairs of aerial images of environments and waypoints, and then a convolutional neural network (CNN) learns the general rule of planning by SL and generates waypoints for the agent to follow. The agent learns to follow arbitrary waypoints while using path-conditioned RL and thus, resulting in efficient exploration. We show that our trained agent can achieve good sample efficiency as well as generalization to novel environments.

Abstract—In this paper, we consider the problem of building learning agents that can efficiently learn to navigate in constrained environments. The main goal is to design agents that can efficiently learn to understand and generalize to different environments using high-dimensional inputs (a 2D map), while following feasible paths that avoid obstacles in obstacle-cluttered environment. To achieve this, we make use of traditional path planning algorithms, supervised learning, and reinforcement learning algorithms in a synergistic way. The key idea is to decouple the navigation problem into planning and control, the former of which is achieved by supervised learning whereas the latter is done by reinforcement learning. Specifically, we train a deep convolutional network that can predict collision-free paths based on a map of the environment—this is then used by an reinforcement learning algorithm to learn to closely follow the path. This allows the trained agent to achieve good generalization while learning faster. We test our proposed method in the recently proposed *Safety Gym* suite that allows testing of safety-constraints during training of learning agents. We compare our proposed method with existing work and show that our method consistently improves the sample efficiency and generalization capability to novel environments.

I. INTRODUCTION

Designing RL agents that can learn complex, safe behavior in constrained environments for navigation has been getting a lot of attention recently [1]. This has been mainly driven by the race to achieve artificial general intelligence where AI agents can achieve human-like performance [2]. It is desirable that the RL agents can generalize to novel environments while achieving optimal performance. Motivated by this problem, in this paper, we present a sample-efficient way of designing agents that can learn to navigate in constrained environments while generalizing to different environments by combining classical planning algorithms, supervised learning and reinforcement learning (see Fig. 1 for pictorial explanation of the proposed method).

Motion planning is central to robotics and thus has been extensively studied [3], [4], [5]. Most of these techniques explicitly consider the geometry of obstacles as well as the robot and find feasible paths ensuring collision avoidance while guaranteeing performance. However, most of these methods handle only kinematic and geometric planning and do not necessarily provide control trajectories that the robot can follow the planned path [5]. Furthermore, these methods may have difficulties when the planning space becomes larger or the workspace of the robot changes dynamically.

Deep Learning has been to achieve super-human performance in the fields of image recognition [6], [7] as well

¹Kei Ota is with Information Technology R&D Center, Mitsubishi Electric Corporation, Japan. Ota.Kei@ds.MitsubishiElectric.co.jp

²Kei Ota, Yoko Sasaki, Yusuke Yoshiyasu and Asako Kanezaki are with National Institute of Advanced Industrial Science and Technology (AIST), Japan. y-sasaki, yusuke-yoshiyasu, kanezaki.asako@aist.go.jp

³Devesh K. Jha is with Mitsubishi Electric Research Labs, Cambridge, MA, USA. jha@merl.com

as when used with RL algorithms for learning games [8]. Deep RL is built on top of recent progresses of Deep Neural Networks (DNNs), which enables it to work even in a large state space such as 2D images, 3D point clouds, etc., by exploiting the high representation capability of DNNs. Also, it enables an agent to acquire complex behaviours considering physical constraints by neatly designing reward functions. Deep RL, however, needs huge amounts of interaction with environments, and could easily overfit to perform optimally in a particular environment. Deep RL has been recently applied to real-time motion planning tasks [9], [10].

Our goal is to efficiently train AI agents that can learn to navigate constrained environments while ensuring safety of the robot. Our proposed method is explained pictorially in Fig. 1. We make use of traditional path planning algorithms, supervised learning (SL) and RL algorithms. We use these three components in a synergistic way to train agents efficiently which can learn to generate feasible trajectories while generalizing to novel environments. Our main idea is to decouple planning and control for an agent that navigates in obstacle-cluttered environment. We use SL to generate waypoints (or subgoals) and the agent then uses RL to learn to follow any sequence of waypoints by using path-conditioned RL. This allows us to generalize to novel environments as the SL can generate waypoints for arbitrary environments after training and the agent can follow any arbitrary trajectory after training with RL. Furthermore, the goal-oriented reference path encourages the RL agent to explore limited regions and thus helps improve sample efficiency. We test our proposed method in the recently proposed *Safety Gym* [1] suite that allows testing of safety-constraints during training of learning agents. We compare our proposed method with existing work and show that our method consistently improves the sample efficiency and generalization capability.

II. RELATED WORK

Deep RL has recently seen tremendous success in learning complex policies for several complex sequential tasks and has attracted huge attention by often achieving superhuman performance [8], [11]. While this shows that deep RL can allow agents to learn complex policies [12], [13], it is also clear that these algorithms require impractical amount of data to learn behavior which are oftentimes very intuitively simple. This has driven a lot of research to allow agents to learn efficiently by allowing better exploration or improving the convergence rates of RL algorithms. Our work falls in the category of allowing better exploration strategies for the agents during learning. Under this category, our work draws similarity with existing work on RL with reference trajectory.

The combination of reference paths and RL has been widely researched [14], [15], [16], [17]. In [14], Probabilistic Roadmaps (PRM) [18] is used to find reference paths, and RL is used to point-to-point navigation as a local planner for PRM. In contrast, [16] uses RL to learn both point-to-point and path-following navigation, which however requires a huge training time because it involves hyper-parameter tuning to improve performance. In [15], a model-based

RL agent is learned for autonomous robotic assembly by exploiting the prior knowledge in the form of CAD data, using a Guided Policy Search (GPS) [19] approach to learn a trajectory-tracking controller. However, the GPS algorithm still produces very local policies and cannot generalize to changing environments. The closest work to ours is [17], which learns an RL agent that optimizes trajectory for 6-DoF manipulator arm. They produce a reference path by Rapidly-exploring Random Trees (RRT) [5] and those points are then used to compute informative reward to reach goal. The important difference is that [17] assumes obstacles and desired location do not change during the whole training/test phase. Our approach, however, can deal with changes in environments by conditioning an RL agent with reference path which is also generated based on an observation of current environment.

Similar to our work, SL is used to learn a path generator in [20]. The navigation is however limited to small environments because it uses first-person-view images captured by a single on-board camera. In order to deal with a navigation task in a large environment, we use aerial view maps as in [21], [22], [23]. Another related work [24], [25] proposes neural network-based motion planning network. They learn to produce feasible paths directly from a point cloud measurement, and shows the high generalization capability. However, their approach considers only path generation, which is similar to the SL part of our work, and requires to track the generated path using some kind of lower-controller, and this could be difficult for a robot that has high dimensional state and action space. Our approach solves the full problem as the waypoints generator is responsible for generalizing to unseen environments, and RL can generate primitive actions that realizes to track the path even in a high dimensional dynamical system.

III. BACKGROUND

In this section, we provide some background for algorithms and methods used in our paper. We describe a standard RL formulation and a state representation [22] which are used in our proposed method to achieve efficiency and generalization capability.

A. Reinforcement Learning

We consider the settings of an RL agent interacting with a stochastic environment to learn a policy that produces an optimal action. The environment is modeled as a Markov decision process (MDP) which consists of a set of states \mathcal{S} , a set of actions \mathcal{A} , a distribution of initial states $p(s_0)$, a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, transition probabilities $p(s_{t+1}|s_t, a_t) : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, and a discount factor $\gamma \in [0, 1]$.

The goal of the RL agent is to acquire the policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the expected sum of rewards in an episode. An episode starts with an initial observation s_0 sampled from $p(s_0)$.

B. GOSELO

We borrow the generalization capability from GOSELO (Goal-directed Obstacle and SELF-Location), which provides

a representation of an image with multiple channels for the input to a CNN [22]. GOSELO is a combination of two bird’s-eye view maps: a map of obstacles observed by the agent’s laser scanner and a map of self-location history from the beginning to the current state. These maps are transformed in the manner as first, rotate and translate maps so that the goal G is located directly above the current location P and the center point M is at the center of the entire image. Next, the rotated image is cropped as a square of size $(L+4) \times (L+4)$ centered at M , where L denotes the number of the pixels on the line GP . The rotated image is then also cropped squares of size $\alpha L \times \alpha L$ and merge them as additional channels. Here, we use two additional channels with $\alpha = (2, 4)$ which is experimentally chosen considering the field size of the environment we use. Finally, those images are concatenated as six channels (i.e., an image of size $W \times H \times 6$) as the input of a CNN, which consists of three channels from an obstacle map and the three channels from a self-location history. The motivation of GOSELO is to make the input environment representation irrelevant to the specific locations of start and goal as well as the shape of obstacles, so as to achieve a *general* navigation policy.

IV. PROPOSED METHOD

In this section, we present details of the proposed method. Specifically, we describe our method of training an agent using SL followed by RL. Figure 2 shows a schematic of the proposed method used to train agents in the paper.

A. Waypoints Generator

Using high-dimensional inputs (images) for training RL agents generally requires huge number of samples for learning meaningful policies as they require bigger networks and thus, more data for convergence [26]. However, for navigation tasks, generating an optimal path (i.e., the sequence of waypoints) is easy by using traditional path planning algorithms, and imitating those optimal paths using DNNs with SL fashion is also easy. We use this idea to first train a module for the agent that can predict a rough path using a map of the environment.

In order to train such an optimal waypoints generator, we first generate a dataset that consists of pairs of GOSELO-style inputs o_k and optimal paths z_k . The optimal paths are generated by A* algorithm [27] from a random start, goal, and obstacles location. Then, the waypoints generator learns to produce the waypoints by minimizing the following objective:

$$L_{\text{waypoints}} = \frac{1}{K} \sum_{k=1}^K \|g(o_k) - z_k\|_2^2. \quad (1)$$

The waypoints generator, g in Eq. (1), consists of a convolutional neural network (CNN) that takes a GOSELO-style input image and generates a short horizon path that consists of 10 waypoints, each of which has 2 dimensional relative position to the current location of the agent. Thus, the SL module provides the agent a rough, collision-free path that it can follow to reach the desired goal.

B. Path Conditioned Reinforcement Learning

We consider the standard RL problem described in Sec. III with the waypoints $z \in \mathcal{Z}$ that lead an RL agent to the goal location. Figure 2 shows the architecture of the proposed approach. The low-level controller (an RL agent) then takes the sequence of the waypoints and internal robot state as an input, and generates low-level action to closely track the waypoints. The waypoints z can be considered as an intermediate representation of a state, because it is acquired by learning a map of GOSELO-style image to the waypoints. We propose to exploit this information to improve the performance of RL agents by combining it with a conventional reward function as $r : \mathcal{S} \times \mathcal{A} \times \mathcal{Z} \rightarrow \mathbb{R}$, and inputs to an RL agent as $o : \mathcal{S} \times \mathcal{Z} \rightarrow \mathcal{O}$, where \mathcal{O} is the extended state space for the agent. Therefore, the reward function can be written as

$$r(s_t, a_t, z_t) = f(s_t, a_t) + h(s_t, a_t, z_t), \quad (2)$$

where $f(s_t, a_t)$ is the reward that originates from an original RL setting. We obtain $h(s_t, a_t, z_t)$ using the waypoints. We specifically define the reward function $h(s_t, a_t, z_t)$ as:

$$h(s_t, a_t, z_t) = w_1 d_{\text{path}} + w_2 n_{\text{progress}}, \quad (3)$$

where d_{path} is the distance to the reference path and n_{progress} is the progress along the path. The first term limits exploration area by penalizing the distance to the waypoints, whereas the second term encourages the agent to move along the waypoints towards the goal location.

In order to calculate d_{path} and n_{progress} , we divide the waypoints and agent’s path at regular intervals, as shown in Fig. 3. By dividing the path, we obtain the sub-sampled vertices $\mathbf{p}_0, \mathbf{p}_1, \dots, \mathbf{p}_{N_p-1}$ for the waypoints, and $\mathbf{x}_t^0, \mathbf{x}_t^1, \dots, \mathbf{x}_t^{N_t-1}$ for the agent’s path, where N_p and N_t are the numbers of vertices in each divided path. We can then define the distance to the given path as $d_{\text{path}} = \max_i D(\mathbf{x}_t^i)$, where $D(\mathbf{x})$ is the distance to the path calculated as $\min_i \|\mathbf{x} - \mathbf{p}_i\|$. We can also observe the progress along the path as $n_{\text{progress}} = \text{NNI}(\mathbf{x}_{t+1}) - \text{NNI}(\mathbf{x}_t)$, where $\text{NNI}(\mathbf{x})$ is the vertex index of the nearest neighbor to \mathbf{x} , i.e., $\text{NNI}(\mathbf{x}) = \arg \min_i \|\mathbf{x} - \mathbf{p}_i\|$. In the case of Fig. 3, $n_{\text{progress}} = 10 - 2 = 8$.

This reward function is then used by an RL algorithm to train a policy to follow the sequence. An important point to note here is that the RL agent learns a path-conditional policy, i.e., a path-tracking policy. This allows flexibility in achieving generalization as the RL is not trained specifically for a goal. As a result, as long as the waypoints generator provides a feasible path to the RL module, the agent can follow it to reach the desired goal.

C. Online Replanning

As we discussed in Sec. IV-A, the waypoints generator outputs only 10 waypoints every-time it is queried. However, since the waypoints are not generated considering the dynamics of the robot, the robot doesn’t exactly follow these waypoints and might start deviating from it after a short

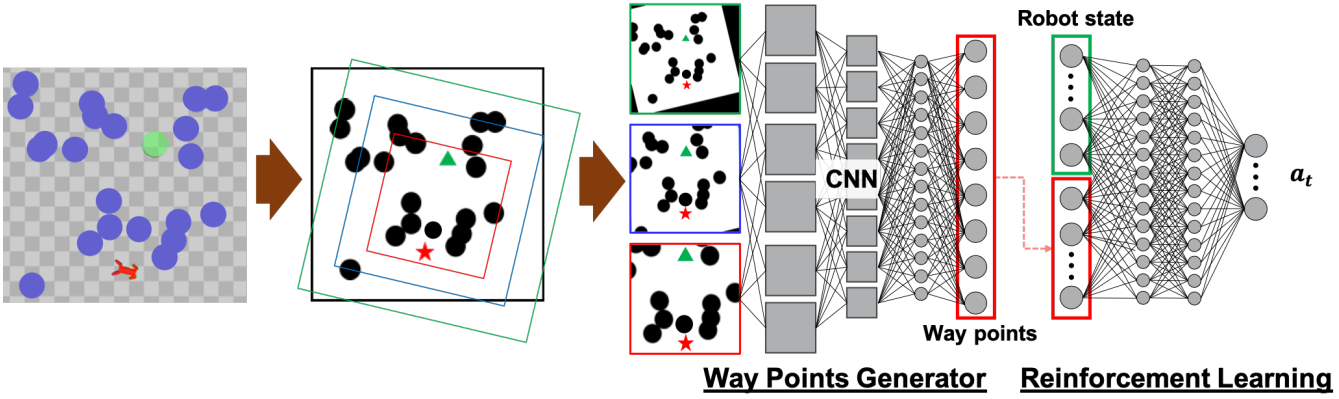


Fig. 2: Architecture of the proposed method. The aerial view of a field is translated into a binary map and then converted to a GOSELO style image. Waypoints generator takes the image and outputs the waypoints that roughly guides the agent to the desired goal. The input to the RL agent is concatenation of the waypoints and an internal state of each robot, and outputs a low-level action to control the robot.

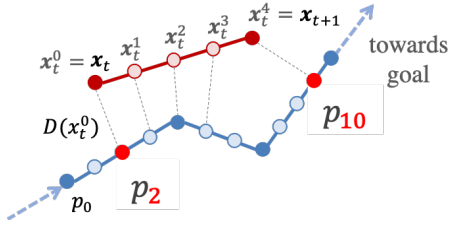


Fig. 3: Path division for calculating rewards. The blue line is the part of waypoints that leads to a goal location and the red line is the agent’s path. x_t is the position of the agent at time step t and p_j is the index of the divided waypoints. The dashed lines indicate the correspondence to the closest neighbor.

horizon. Consequently, we keep updating the waypoints after a certain time window. This is done by updating the state of the robot (and the environment in the presence of moving obstacles) in the GOSELO-style image and then generate new waypoints. The robot then tries to follow the new sequence of waypoints. This is then repeated till the robot reaches the desired goal. Thus the agent needs to decide when it should update its plan. Ideally, updating its plan after every step will result in optimal behavior but may be computationally heavy as computation of waypoints requires the forward propagation of relatively computationally heavy CNN architecture. On the other hand if the update frequency is too small, the agent might not be able to move efficiently. Therefore, the update frequency should be defined by a trade-off between optimal behaviour and the computational cost. To solve this trade-off, we propose a rule-based online replanning strategy, where the agent asks for new waypoints depending on whether 1) the distance of the agent and the closest way point d_{path} is larger than a threshold d_{update} , or 2) the closest index of waypoints $\text{NNI}(x_t)$ is larger than a half of the number of waypoints, where d_{path} and $\text{NNI}(x_t)$ are defined in Sec. IV-B. The threshold in condition (1) is a hyperparameter that

needs to be tuned.

V. EXPERIMENTAL SETTINGS

This section briefly describes the experimental settings including the environments and tasks used in the paper to validate the proposed method, for more details refer to Sec. VIII-A.

A. Environment

We evaluate our method on custom environment, which is based on Safety Gym [1]. Here we describe the details of the environment.

1) *States*: The states of an agent includes the standard robot sensors (accelerometer, gyroscope, magnetometer, velocimeter, and robot-centric lidar observation of 10 directions exclusive and exhaustive for a full 360 view) and robot specific sensors (such as joint positions and velocities) as s_{robot} , and 10 waypoints $z \in \mathbb{R}^{2 \times 10}$ generated by the waypoints generator that guides the agent to the goal location as described in Sec. IV-A.

2) *Actions*: The actions of the agent is to control each robot to achieve a desired task. The robot we use for our experiments include three types: *point*, *car*, and *doggo* [1]. The states and actions for each robot type is as follows: *point* (42, 2), *car* (54, 2), and *doggo* (78, 12).

3) *Rewards*: As described in Sec. IV-B, the reward function of our path-conditioned RL consists of conventional reward function $f(s_t, a_t)$, and $h(s_t, a_t, z_t)$ which is calculated from the waypoints. Since the $h(s_t, a_t, z_t)$ is already described, here we define the reward term for original RL setting referring to [1] as:

$$f(s_t, a_t) = w_3 \mathbb{I}_{\text{collision}} + w_4 \mathbb{I}_{\text{goal}}, \quad (4)$$

where $\mathbb{I}_{\text{collision}}$ and \mathbb{I}_{goal} are indicator functions of whether collision between the agent and the obstacles occurs, and whether the agent reaches goal location.

TABLE I: Average number of collision with obstacles in an episode.

Number of training data	1K	10K	50K	100K
Number of collisions in an episode	3.05	1.57	1.15	0.96

4) *Terminal conditions*: An episode terminates with following three conditions: the location of the agent \mathbf{x}_t is sufficiently close to the goal location, the agent explores too far from the goal location, or the number of steps of an episode is over a specified threshold.

VI. EXPERIMENTAL RESULTS

This section presents results from experiments designed to answer the following questions:

- Can we learn to solve the desired task of reaching specified goal while avoiding collision with randomly-placed obstacles?
- Does the combination of waypoints generator and RL-based controller improve performance on RL?
- Does the waypoints generator helps low-level controller generalize to novel environments?

A. Performance of Waypoints Generator

We first demonstrate the effectiveness of the waypoints generator which generates waypoints for an RL agent. In order to evaluate the quality of the waypoints, we move the agent to the closest waypoint for each step, and evaluate the average number of times of collision over 1000 episodes.

Table I shows the results for average number of collisions. It suggests that the quality of waypoints improves when we add training data, and the number of collisions in an episode saturates around 1.00. The inaccuracy of the waypoints generator will be taken care of by the RL agent by minimizing the collision penalty defined in Eq. (4).

B. Sample Efficiency

We test the proposed method for sample efficiency as we believe the path-conditioned RL makes the agent explore efficiently. To verify this, we compare the proposed method to an RL agent without waypoints, which we define as baseline. Since the baseline does not use waypoints, it requires a reward signal instead of moving towards goal location along with waypoints. In order to do that, we add a goal distance penalty to encourages the agent to reach the goal, and also add the relative location of the goal $\mathbf{x}_t^{\text{rel}} = \mathbf{x}_t - \mathbf{x}_{\text{goal}}$ to an input to the RL agent. Therefore, the reward function of the baseline will be:

$$r_{\text{baseline}}(s_t, a_t) = f(s_t, a_t) + w_5 d_{\text{goal}}, \quad (5)$$

where $f(s_t, a_t)$ is the same reward as defined in Eq. (4), and the second term is distance penalty to the goal location computed by $d_{\text{goal}} = \|\mathbf{x}_t - \mathbf{x}_{\text{goal}}\|$. For a fair comparison, we evaluate both agents on this reward function. Both agents are trained with the SAC algorithm [28].

Figure 4 shows the training curves of resulting episodic returns. It suggests that the use of waypoints improves convergence performance with respect to the training without waypoints. Looking at the performance gap between *point* and *doggo*, we can see that decoupling the planning and control improves sample efficiency especially for high dimensional dynamical system.

C. Generalization to Novel Environments

Next, we evaluate the generalization of our method with respect to novel environments. As for environment, we prepared five different types as depicted in Fig. 5. We first see whether a trained agent can generalize to different size of a field and number of obstacles by preparing *pillar* ($W \times H \times N$) environment, where W and H is the width and height of the field, and N is number of obstacles. Note that the agent is trained with a *pillar* (2, 2, 10) environment. Then, we evaluate whether the agent can generalize to quickly moving obstacles in the *gremlin* environment, where boxes move circularly [1]. Finally, we evaluate the agent on simple maze environments: *two-room* and *four-room*, where the agent must take a detour to reach the goal location because the rooms are divided by walls.

We compared our method with baseline agent, which is the same with the one trained in the previous experiment in Sec. VI-B. The robot type we used is *doggo*, which has the most complex dynamical system provided by Safety Gym. For a fair comparison, we used a converged model for both ours and baseline trained with three million time steps. Since the distribution of input image to the waypoints generator is different from trained settings in Sec. VI-A, we also evaluate the agent trained by our method with fine-tuned waypoints generators. Note that only the waypoints model is fine-tuned, and we do not fine-tune the RL agents.

Table II shows the result of the experiments. Overall, our method generalizes to novel environments better than the baseline. Since our method decouples planning and control, the fine-tuned waypoints improves the performance in most novel environments by mitigating the inaccuracy of the waypoints. Thus, we verified our method improves generalization capability.

D. Online Replanning

In order to verify the performance of our online replanning strategy described in Sec. IV-C, we conducted a comparative study to evaluate the performance among different frequencies to replan waypoints. Figure 6 shows the results of training curves of average steps to reach goal over different frequencies to update waypoints. This results demonstrate that if the replanning frequency is too slow, then the agent does not learn well because of the inaccuracy of the waypoints from which the agent acquires rewards. Our proposed method provides more informative reward signals to an agent by replanning the waypoints while the update frequency of 8.5 is better than 4.

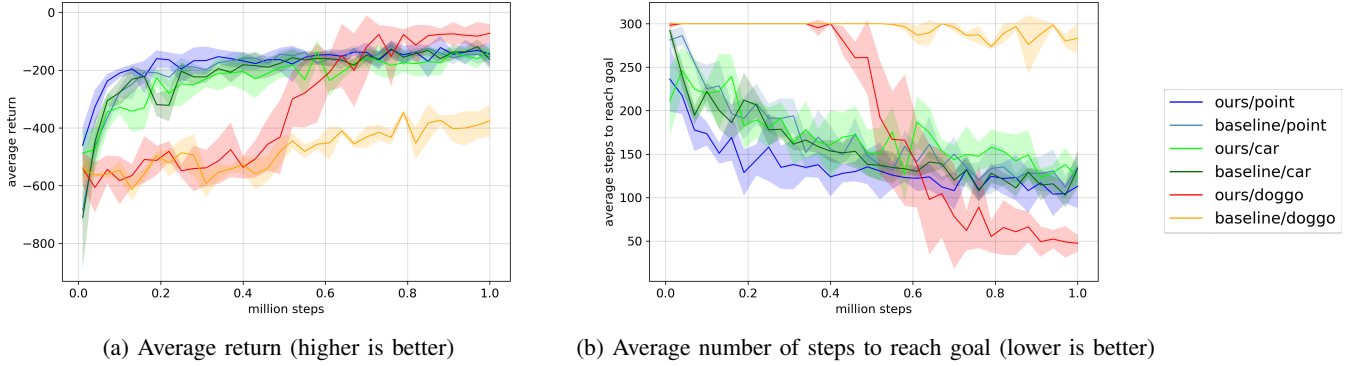


Fig. 4: Training curves on each robot type. The solid lines represent average returns over five instances with different random seeds. The shaded region represents the standard deviation of the five instances. Our approach outperforms baseline on both sample efficiency and final performance.

TABLE II: Performance on generalization task in simulation. The results are averaged over 200 episodes with random agent, goal, obstacles locations. The bold number indicates the best results among three different methods. In the “Ours Fine-tuned”, we used the fine-tuned waypoints generator to improve the quality of waypoints. Note we do not retrain the RL agent.

Environment	Goal reach rate			Steps to reach goal		
	Baseline	Ours	Ours Fine-tuned	Baseline	Ours	Ours Fine-tuned
<i>gremlin</i>	0.67	0.99	1.00	131.3	33.7	33.8
<i>pillar</i> (3, 3, 25)	0.78	0.96	0.97	152.1	66.3	59.6
<i>pillar</i> (4, 4, 40)	0.63	0.97	0.95	273.7	76.9	94.1
<i>two-room</i>	0.37	0.59	0.99	397.1	268.4	59.0
<i>four-room</i>	0.67	0.9	0.98	231.7	64.6	50.6

VII. CONCLUSION

In this paper, we present a sample efficient and generalizable method for navigation task for robots with complex dynamics in obstacle-cluttered environments. The method combines traditional path planning algorithm, supervised learning, and reinforcement learning. The SL learns to produce the output of the traditional path planning algorithm as a form of waypoints, and RL takes the waypoints to generate primitive actions to control robots. In the experiments performed in the Safety Gym suite, we demonstrated that the proposed method 1) learns more sample efficiently than RL without waypoints, 2) generalize well to unknown environments, 3) can adapt to moving obstacles in the environment while ensuring safety of the agent.

In our future work, we plan to extend this method to real applications for mobile robots and robotic arms.

REFERENCES

- [1] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking Safe Exploration in Deep Reinforcement Learning. 2019.
- [2] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- [3] Anthony Stentz. Optimal and Efficient Path Planning for Partially-Known Environments. Technical report, 1994.
- [4] F. Arambula Cosío and M.A. Padilla Castañeda. Autonomous robot navigation using adaptive potential fields. *Mathematical and Computer Modelling*, 40(9):1141 – 1156, 2004.
- [5] Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.
- [6] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016.
- [7] G. Huang, Z. Liu, L. v. d. Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, July 2017.
- [8] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [9] Michael Everett, Yu Fan Chen, and Jonathan P How. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3052–3059. IEEE, 2018.
- [10] Björn Lötjens, Michael Everett, and Jonathan P How. Safe reinforcement learning with model uncertainty estimates. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8662–8668. IEEE, 2019.
- [11] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [12] John Schulman, Sergey Levine, Pieter Abbeel, Michael I Jordan, and Philipp Moritz. Trust region policy optimization. In *ICML*, volume 37, pages 1889–1897, 2015.
- [13] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [14] Aleksandra Faust, Kenneth Oslund, Oscar Ramirez, Anthony Francis, Lydia Tapia, Marek Fiser, and James Davidson. Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and

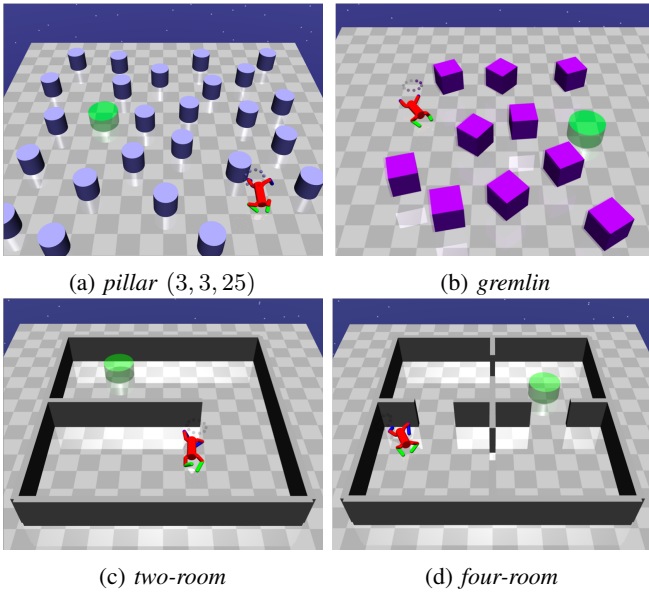


Fig. 5: Novel environments to verify generalization capability of proposed methods.

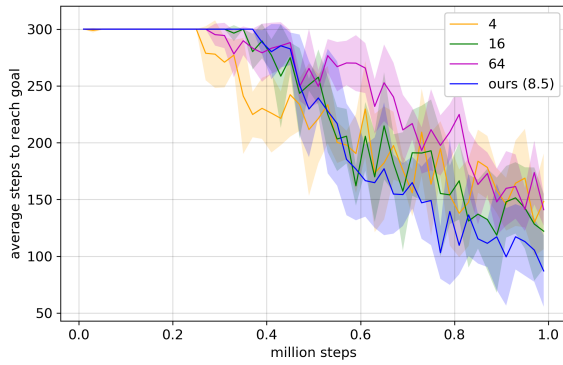


Fig. 6: Training curves comparison among different frequencies to update waypoints. The training is done by *doggo*.

- sampling-based planning. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [15] Garrett Thomas, Melissa Chien, Aviv Tamar, Juan Aparicio Ojea, and Pieter Abbeel. Learning robotic assembly from cad. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2018.
- [16] Hao-Tien Lewis Chiang, Aleksandra Faust, Marek Fiser, and Anthony Francis. Learning navigation behaviors end-to-end with autorl. *IEEE Robotics and Automation Letters*, 4(2):2007–2014, 2019.
- [17] Kei Ota, Devesh K. Jha, Tomoaki Oiki, Mamoru Miura, Takashi Nammoto, Daniel Nikovski, and Toshisada Mariyama. Trajectory optimization for unknown constrained systems using reinforcement learning. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019.
- [18] Lydia E Kavraru, Petr Svostka, Jean-Claude Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics and Automation*, 12(4), 1996.
- [19] Sergey Levine and Vladlen Koltun. Guided policy search. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 1–9, 2013.
- [20] Somil Bansal, Varun Tolani, Saurabh Gupta, Jitendra Malik, and Claire Tomlin. Combining optimal control and learning for visual navigation

- in novel environments. 2019.
- [21] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. *Proceedings of International Joint Conference on Artificial Intelligence*, Aug 2017.
- [22] A. Kanazaki, J. Nitta, and Y. Sasaki. Goselo: Goal-directed obstacle and self-location map for robot navigation using reactive neural networks. *IEEE Robotics and Automation Letters*, 3(2):696–703, April 2018.
- [23] Mayank Bansal, Alex Krizhevsky, and Abhijit Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. In *Proc. Robotics: Science and Systems (RSS)*, 2019.
- [24] Ahmed H Qureshi and Michael C Yip. Deeply informed neural sampling for robot motion planning. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6582–6588. IEEE, 2018.
- [25] Ahmed H Qureshi, Anthony Simeonov, Mayur J Bency, and Michael C Yip. Motion planning networks. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 2118–2124. IEEE, 2019.
- [26] Timothée Lesort, Natalia Díaz-Rodríguez, Jean-François Goudou, and David Filliat. State Representation Learning for Control: An Overview. *CoRR*, 2018.
- [27] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, July 1968.
- [28] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of International Conference on Machine Learning (ICML)*, pages 1861–1870, 2018.
- [29] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [30] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5026–5033. IEEE, 2012.

VIII. APPENDIX

A. Experiments Details

1) *Simulation Setup*: We evaluate our methods on recently proposed Safety Gym environments [1]. As for obstacles, we located N number of “pillars”, where $N = 10$ for Sec. VI-B and $N = \{25, 40\}$ for Sec. VI-C, which are immovable objects and the agent gets penalty when it collides with them. The location of an agent, a goal, and obstacles are randomly sampled when an episode starts, and the goal and obstacles are fixed during an episode. The field size ($W \times H$) is $W = H = 2$ [m] as default, and we change the size only when we conduct generalization experiments as described in Sec. VI-C.

2) *Terminal Conditions*: For the maximum steps for one of terminal conditions, we used $150 \times W$, where W is the width of the field. When the environment is *two-room* or *four-room*, we doubled the episode length because the agent needs to take a detour to reach a goal location. Also, the episode terminates when the distance between the goal and the agent is less than 0.3 [m], as implemented in Safety Gym. The last condition is whether the agent exceeds the limited region $W \times H$, which are defined above.

3) *Waypoints Generator*: As for the distance condition of updating the waypoints d_{update} , which is defined in Sec. IV-C, we use $d_{\text{update}} = 0.3$ for all experiments.

4) *Rewards*: Table. III summarizes the coefficients of reward terms defined in Eq. (3),(4),(5).

TABLE III: Coefficients of each reward term w_i used in our experiments.

Term	Value	Description
$w_1 d_{\text{path}}$	-0.1	Distance penalty to the closest reference path
$w_2 n_{\text{progress}}$	0.5	Distance reward going toward the goal
$w_3 \mathbb{I}_{\text{collision}}$	-1.0	Obstacle collision penalty
$w_4 \mathbb{I}_{\text{goal}}$	1.0	Goal reach reward
$w_5 d_{\text{goal}}$	-1.0	Distance penalty from the goal

TABLE IV: The architecture of the waypoints generator.

Shape	Layer size	Stride	Type
$64 \times 64 \times 6$	-	-	GOSELO-style input
$32 \times 32 \times 32$	$3 \times 3 \times 32$	2	Convolution
$16 \times 16 \times 64$	$3 \times 3 \times 64$	2	Convolution
$8 \times 8 \times 64$	$3 \times 3 \times 64$	2	Convolution
64	-	-	Global Average Pooling
256	-	256	Fully Connected
$2 \times n_{\text{way-points}}$	-	-	Output

B. Training Details

1) *Waypoints Generator*: Table IV summarizes the architecture of the CNN. This architecture is almost same with previous works for Atari environments [29].

The dataset of the waypoints generator is generated using A* search on each environment type by discretizing ($W \times H$) field by 0.1 [m], and collision detection is done by MuJoCo’s function [30]. To mitigate the jerky path caused by the planning on discretized space, we short-cut the produced vertices with randomly selected two vertices, and replace them with new vertices, that connects the two vertices with fixed distance, if those new vertices do not collide with obstacles.

2) *Reinforcement Learning*: We used SAC agent for an RL agent over all experiments. The network architecture and hyper-parameters of SAC agent is strictly same with original paper [28].