

Learning Navigation Behaviors End-to-End With AutoRL

Hao-Tien Lewis Chiang , Aleksandra Faust , Marek Fiser, and Anthony Francis

Abstract—We learn end-to-end point-to-point and path-following navigation behaviors that avoid moving obstacles. These policies receive noisy lidar observations and output robot linear and angular velocities. The policies are trained in small, static environments with AutoRL, an evolutionary automation layer around reinforcement learning (RL) that searches for a deep RL reward and neural network architecture with large-scale hyper-parameter optimization. AutoRL first finds a reward that maximizes task completion and then finds a neural network architecture that maximizes the cumulative of the found reward. Empirical evaluations, both in simulation and on-robot, show that AutoRL policies do not suffer from the catastrophic forgetfulness that plagues many other deep reinforcement learning algorithms, generalize to new environments and moving obstacles, are robust to sensor, actuator, and localization noise, and can serve as robust building blocks for larger navigation tasks. Our path-following and point-to-point policies are, respectively, 23% and 26% more successful than comparison methods across new environments.

Index Terms—Autonomous agents, collision avoidance, deep learning in robotics and automation, motion and path planning.

I. INTRODUCTION

ASSISTIVE robots, last-mile delivery, warehouse navigation, and robots in offices require robust robot navigation in dynamic environments (Fig. 1). While good methods exist for robot navigation sub-tasks such as localization and mapping, motion planning, and control, current local navigation methods typically must be tuned for each new robot and environment [7]. For example, vision-based navigation is robust to noisy sensors but typically relies on high-level motion primitives, such as “go straight” and “turn left” that abstract away robot dynamics [29]. Robot navigation requires behaviors that are useful to existing navigation stacks, easily transfer from simulation to physical robots or new large-scale environments with minimum tuning, and robustly avoid obstacles despite noisy sensors and actuators.

Manuscript received September 10, 2018; accepted January 18, 2019. Date of publication February 18, 2019; date of current version February 28, 2019. This letter was recommended for publication by Associate Editor C. Ekenna and Editor N. Amato upon evaluation of the reviewers’ comments. This work was supported by Google. (Hao-Tien Lewis Chiang and Aleksandra Faust contributed equally to this work.) (Corresponding author: Aleksandra Faust.)

The authors are with the Robotics at Google, Google AI, Mountain View, CA 94043 USA (e-mail: lewispro@google.com; faust@google.com; mfiser@google.com; centaur@google.com).

This letter has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. The enclosed video demonstrates AutoRL point to point and path following policies evaluations in simulation and on real robots. Contact lewispro@google.com and faust@google.com for further questions about this work.

Digital Object Identifier 10.1109/LRA.2019.2899918




Fig. 1. Navigation behaviors on a Fetch: (a) path following in a narrow corridor (b) point to point in a dynamic environment.

We define robot navigation behaviors as intelligent agents that depend on the robot’s sensors, actuators, dynamics and geometry, without relying on foreknowledge of its environment [3]. Many navigation tasks can be accomplished using two basic behaviors: local planning, i.e., point to point, which creates trajectories from the robot’s current configuration to a given target [12], and path following, which stays near a given guidance path [8]. Both behaviors need to produce dynamically feasible trajectories that are robust to many kinds of noise and inaccuracies, rely only on observations coming from primitive sensors, and avoid unexpected obstacles. In this letter, we present a reliable method to implement these navigation behaviors by learning end to end policies that directly map sensors to controls, and we show that these policies transfer from simulation to physical robots and new environments while robustly avoiding obstacles.

To do so, we rely on deep Reinforcement Learning (RL), which has shown great progress in visual navigation [37] and video games [28]. RL learns policies that map between observations and controls to maximize the cumulative return. Deep RL has enabled learning tasks that have been difficult or impossible to program manually, from matching human performance [28] to locomotion of humanoid figures [17]. Despite the promise of deep RL, training these policies is difficult and requires careful consideration when selecting the reward function and choosing the neural network architecture. Specifically, training even simple tasks can fail if rewards are sparse - that is, if the success conditions of the true objective are hard to discover [4]; this is true for navigation in large spaces. Reward shaping [21], [36] addresses this problem by introducing a proxy reward function that is less sparse than the true objective. However, poorly chosen shaped rewards can lead to pathologies where agents learn to exploit the reward function or to terminate the episode early



Fig. 2. (a) Training and (b, c, d) evaluation environments with 40 moving obstacles (white circles). The evaluation environments are 6.8 to 30.1 times bigger than the training, generated from real office building plans.

[31]. Designing good reward functions is well understood for some areas (e.g., video games [28]), but for most tasks, including navigation, it remains a black art. Similarly, selection of neural network architecture often consists of trial and error, and remains another source of training difficulties.

Hyperparameter tuning can improve learning behavior: gradient-free, evolutionary methods, such as Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [18], can transform an algorithm that fails to converge into one that produces near-optimal policies. To simplify deep RL training, we present AutoRL, which automates the search for both the shaped reward and the neural network architecture. AutoRL combines deep RL with gradient-free hyperparameter optimization. It takes a parameterized reward function along with the true objective, and then uses large-scale hyperparameter optimization to shape the reward. With this reward fixed, it then optimizes network layer sizes to identify the most successful policies for the given task. Applied to point-to-point and path-following tasks, AutoRL learns robust policies that cope with both static and dynamic obstacles, even though the tasks are trained in simulation in small environments with only static obstacles. The specific contributions of this work are the AutoRL method as well as point-to-point and path-following navigation policies.

We evaluate these policies in three large building environments (Fig. 2) with respect to various noise levels and a number of moving obstacles. Our simulation results show the path following and point to point policies have higher success rate in novel environments, on average 23% and 26% higher over the baseline, respectively, and are robust to noise. In experiments on a real robot (Fig. 1), we observe reliable obstacle avoidance in dynamic environments and collision-free navigation along 80+m paths.

II. RELATED WORK

Motion planning: Practical robot navigation typically uses a layered navigation stack: given an occupancy map generated by a mapping algorithm such as SLAM [6], a global planner such as PRM [22] or RRT [24] finds a near-optimal path that a local planner (e.g., [14]) executes. This “stacked” approach has demonstrated reliable operation over long distances [26]. Typical planners for dynamic environments include Artificial Potential Field (APF), velocity obstacle, and sampling-based methods [9]. Many of those methods, such as DRT [8] or PEARL [11], require perfect knowledge of obstacle position or dynamics, which is difficult to obtain when the robot has limited and

noisy sensors, like 1D lidar. Like these approaches, we assume reasonably accurate localization, but we couple sensing directly with controls to remove the need for knowing obstacle location or dynamics.

Deep learning for navigation: Machine learning for navigation has seen a recent boom, particularly in point-to-point navigation with vision [29], [37]. Some approaches discretize state and action spaces to enable simple RL algorithms [1], whereas others learn low-level controllers end to end using deep learning [7]. Our prior work on PRM-RL uses a point-to-point policy trained with DDPG as the local planner for PRM [12]. Like these approaches, we map from sensors to velocity control, but we train with AutoRL, which tunes the agent’s reward and network architectures to improve the policy performance and time needed for hand-engineering. Learning by optimizing multiple objective functions, rules, or perceptual features has proven effective for variants of the navigation problem [10]; other work uses hierarchical RL to break the problem apart into tractable components [5]. Unlike these methods, we automate the tuning by learning the proxy reward function and neural network architecture instead of imposing it. Some recent work features learning based sensor to controls navigation in the presence of moving obstacles. For example, Intention-Net [15] maps camera observations and intents to controls, while Pfeiffer *et al.* [32] navigate from lidar inputs. Both of those methods require demonstrations for training, while AutoRL learns by reinforcement.

Reinforcement learning with sparse rewards: Three main approaches handle sparse rewards in RL: curriculum learning, bootstrapping, and reward shaping. Curriculum methods make tasks progressively more difficult during training, e.g., by moving the goal further from the start in navigation [13], [20]. In contrast, AutoRL does not change task difficulty during training. Bootstrapping allows RL policies to be initialized from either hand-engineered policies or recorded episodes [35]. AutoRL starts with an uninitialized policy and trains it from scratch. Reward shaping learns from a parameterized dense reward [30]. AutoRL is a reward-shaping method. While reward shaping can be done with curriculum and bootstrapping [20], [35], AutoRL uses hyperparameter optimization rooted in evolutionary algorithms to accomplish this. Further, curriculum learning and bootstrapping can be applied in addition to the hyperparameter optimization.

Hyperparameter optimization in deep learning: AutoML methods learn the neural network architecture for a given problem. AutoML achieves this with gradient descent [2], rein-

forcement learning [38], and even hyper-parameter optimization [34]. AutoRL extends the AutoML approach to RL and subjects both the network architecture and the reward function to hyper-parameter optimization.

III. METHODS

To learn point-to-point and path-following navigation behaviors, we choose continuous reinforcement learning with partial state observations, which we model with Partially Observable Markov Decision Process (POMDP). We model the agent as an (O, A, D, R, γ) tuple representing a partially observable Markov decision process (POMDP) with continuous observations and actions. Observations, actions, and dynamics are determined by the robot and are common to both tasks; we describe them in Section III-A. Point to point and path following share observations, actions, and dynamics, but have different objectives, that are modeled through rewards in POMDP. Sections III-B and III-C describe the true task objectives, specific rewards, and additional observations specific to each behavior. Section III-D describes the hyper-parameters of a reinforcement learning algorithm, and Section III-E presents the AutoRL algorithm that learns both the hyperparameters and the policy.

A. POMDP Setup

The observations, $\mathbf{o} = (\mathbf{o}_l, \mathbf{o}_g)^{\theta_n} \in O$, are θ_n pairs of 1-D lidar vectors, \mathbf{o}_l , and goal set, \mathbf{o}_g , observed over the last θ_n steps. The agent is a unicycle robot and is controlled by a 2-dimensional continuous action vector, $\mathbf{a} = (v, \phi) \in A$, that encodes the robot's linear and angular velocity. The dynamics, D , is encoded in the simulator or implicit in the real world. The remaining factors encode the task: $\gamma \in (0, 1)$ is a scalar discount factor, whereas the structure of the reward R is more complicated and we discuss it next.

The agent's goal is to complete a true navigation objective. For point-to-point navigation this is arriving at a goal location, while for path following this is traversing the entire path by reaching all its waypoints. We can formalize this problem as learning a policy that maximizes the probability of reaching the true objective, G ,

$$\tilde{\pi} = \arg \max_{\pi} \mathbb{P}(G(s)|\pi), \quad (1)$$

where $\mathbb{P}(G(s)|\pi)$ means that true objective G is reachable from the state s under control of policy π . At the same time, RL learns a policy that maximizes the cumulative reward. We could use the true objective as a reward, but it is sparse, and there are other requirements, such as dynamical feasibility, smooth trajectories, avoiding obstacles, and sensory/motor limitations. We can formulate these requirements as parameterized atomic rewards which provide more timely feedback to aid learning. To that end we represent rewards with

$$R_{\theta_r}(s, \mathbf{a}) = \sum_{i=1}^{n_r} r_i(s, \mathbf{a}, \theta_{r_i}), \quad (2)$$

where $r_i(s, \mathbf{a}, \theta_{r_i})$ is a parameterized atomic reward, and $\theta_r = [\theta_{r_1} \cdots \theta_{r_{n_r}}]$ becomes a hyperparameter to be tuned.

B. Point-to-Point (P2P) Task

The goal of the P2P behavior is to navigate a robot to a goal position without collision. We assume the robot is well-localized using traditional methods. The P2P behavior can be used as a local planner for sampling-based planners [12] in order to navigate large environments. We require the agent to navigate to goals that are not within clear line of sight, but not far enough to require higher-level knowledge of the environment. For example, we expect it to navigate around a corner, but not to a room in a maze of hallways.

The true objective of P2P is to maximize the probability of reaching the goal during an episode,

$$G_{P2P}(s) = \mathbb{I}(\|s - s_g\| < d_{P2P}), \quad (3)$$

where \mathbb{I} is an indicator function, s_g is the goal pose, and d_{P2P} the goal size. The goal observation \mathbf{o}_g is the relative goal position in polar coordinates, which is readily available from localization. The reward for P2P is:

$$R_{\theta_{r_{P2P}}} = \theta_{r_{P2P}}^T [r_{\text{step}} r_{\text{goalDist}} r_{\text{collision}} r_{\text{turning}} r_{\text{clearance}} r_{\text{goal}}], \quad (4)$$

where r_{step} is a constant penalty step with value 1, r_{goalDist} is the negative Euclidean distance to the goal, $r_{\text{collision}}$ is 1 when the agent collides with obstacles and 0 otherwise, r_{turning} is the negative angular speed, $r_{\text{clearance}}$ is the distance to the closest obstacle, and r_{goal} is 1 when the agent reaches the goal and 0 otherwise.

C. Path-Following (PF) Task

The goal of the PF behavior is to follow a guidance path represented by a sequence of waypoints in the workspace. The collision-free guidance path does not need to be dynamically feasible, and can be generated by off-the-shelf path planners like PRMs or A*, or even manually. In navigation stack terms, PF is trajectory tracking.

Guidance paths for real-world navigation are often long (100+ m) and have many waypoints with varied separation. The varied input size and non-uniformity pose a challenge for RL methods using neural networks. To address this problem, we augment the original guidance path \mathcal{P}_o with intermediate waypoints that are linearly interpolated with a constant separation d_{ws} between consecutive waypoints. The result is a new guidance path \mathcal{P} consisting of both original and new waypoints. The separation, d_{ws} , between waypoints is a hyper-parameter that AutoRL searches for. The waypoint w_i is considered reached iff 1) the previous waypoint w_{i-1} is already reached and 2) the robot is within d_{wr} of w_i .

The true objective of PF is to reach as many waypoints per episode as possible:

$$G_{PF}(s) = \frac{\sum_{w \in \mathcal{P}} \mathbb{I}(\|s - w\| < d_{wr})}{\|\mathcal{P}\|}. \quad (5)$$

where \mathbb{I} is an indicator function that returns 1 if waypoint w is reached and 0 otherwise. The goal observation of PF, \mathbf{o}_g , is a partial path consisting of the N_{partial} un-reached waypoints. The reward for PF is:

$$R_{\theta_{r_{PF}}} = \theta_{r_{PF}}^T [r_{\text{step}} r_{\text{dist}} r_{\text{collision}} r_{\text{clearance}}]. \quad (6)$$

where r_{step} is a penalty step with value 1. r_{dist} is the Euclidean distance to the first un-reached waypoint. $r_{\text{collision}}$ is 1 when the agent collide with obstacles and 0 otherwise. $r_{\text{clearance}}$ is a penalty with value 1 when the the agent is within $d_{\text{clearance}}$ m away from obstacles and 0 otherwise.

D. Reinforcement Learning Parametrization

With observations, action space, true objectives, and parameterized rewards defined, training deep RL requires selecting neural network architecture. Network architecture affects the quality of the trained agent: the capacity of the network determines what the agent can learn. In this work, we use feed-forward fully-connected networks and fix network depth, leaving the size of each layer as our tunable hyperparameter. However, the hyperparameter tuning technique can be applied to any network architecture.

Let $FF(\theta)$, for $\theta \in \mathbb{R}^n$ be a feed-forward fully-connected neural network with RELU units and n layers, where the i^{th} layer contains θ_i neurons. Let us denote the learnable weights of the feed-forward network $FF(\theta)$ with W_θ . Let

$$RL(\text{Actor}(\theta_\pi), \text{Critic}(\theta_Q), R(\theta_r))$$

be a parameterized actor-critic reinforcement learning algorithm that learns policy $\pi(s|W_\pi) = FF(\theta_\pi)$ and critic $Q(s, \mathbf{a}|W_Q) = FF(\theta_Q)$. Actor and critics network architectures are parameterized with θ_π, θ_Q , and reward function, R given in Eq. (2), parameterized with vector θ_r . Let $Obj(\theta_\pi, \theta_Q, \theta_r|G) \in \mathbb{R}$ be the true objective the trained actor $\pi(s|W_\pi)$ achieved for the corresponding $RL(\text{Actor}(\theta_\pi), \text{Critic}(\theta_Q), R(\theta_r))$.

E. AutoRL

We automate RL hyperparameter selection with the AutoRL evolutionary search procedure, described in Algorithms 1 and 2. We split shaping in two phases, reward shaping (Line 1, Alg. 1) and network shaping (Line 2), because the search space grows exponentially in the number of tuning parameters. First, we find the best reward function w.r.t. task's true objective for a fixed actor and critic. Then, we find the best actor and critic w.r.t. to the previously selected reward (Algorithm 1).

For reward shaping, actor and critic network shapes are fixed sizes $\theta_\pi, \theta_Q \in I(\mathbf{n}_{\min}, \mathbf{n}_{\max})$ where $I(\mathbf{a}, \mathbf{b})$ is a closed interval in n -dimensional space bounded by points $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$. We run n_g trials, at most n_{mc} in parallel (Line 3, Alg. 2). At each trial i , we initialize the reward function $\theta_r(i)$ from $I(0, 1)^n$, based on all completed trials according to a black-box gradient-free optimization algorithm [18] (Line 5). The first n_{mc} trials select reward weights randomly. Next, we train asynchronous instances $RL(\text{Actor}(\theta_\pi), \text{Critic}(\theta_Q), R(\theta_r^i))$ (Line 13). After each agent is trained, its policy is evaluated according to the true task objective Eq. (1) (Line 14). For P2P that is Eq. (3) and for PF it is Eq. (5). Upon completion of all n_g trials, the best reward (Line 17)

$$\tilde{\theta}_r = \arg \max_i Obj(\theta_\pi, \theta_Q, \theta_r^i|G) \quad (7)$$

corresponds to the trial with the highest true objective.

Algorithm 1: AutoRL.

Output: $\pi(s|W_{\theta_\pi})$: Trained policy.

- 1: /* Select best reward. */
 - $-, \tilde{\theta}_r, -, - \leftarrow \text{HPSelector}(\theta_r = \text{None})$
 - 2: /* Select best NN architecture. */
 - $\tilde{\pi}, -, \theta_\pi^*, \theta_Q^* \leftarrow \text{HPSelector}(\tilde{\theta}_r)$
 - 3: **return** $\tilde{\pi}$
-

Algorithm 2: HPSelector: Hyper-parameter selector.

Property: n_g : Num. of generations (trials).

Property: n_{mc} : Num. of parallel trials.

Property: $\mathbf{n}_{\min}, \mathbf{n}_{\max}$: Min. and max. of neurons per layer.

Input: θ_r : Reward hyper-parameters.

Output: $\tilde{\pi}$: Policy trained with the selected hyper-parameters.

Output: $\tilde{\theta}_r$: Selected reward hyper-parameters.

Output: $\tilde{\theta}_\pi, \tilde{\theta}_Q$: Selected NN architecture hyper-parameters.

- 1: $T \leftarrow \emptyset$ /* Initialize the experience. */
 - 2: Initialize NN architecture: $\mathbf{n}_{\min} \leq \theta_\pi, \theta_Q \leq \mathbf{n}_{\max}$
 - 3: **for** $i = 1, \dots, n_g$ running n_{mc} trials in parallel. **do**
 - 4: **if** θ_r is None /* Tune rewards */ **then**
 - 5: $\theta_r^i \leftarrow \text{Select}(\min = 0, \max = 1, T)$ with [18]
 - 6: $\theta_\pi^i, \theta_Q^i \leftarrow \theta_\pi, \theta_Q$ /* NN hyper-parameters are fixed. */
 - 7: $obj_{fn} \leftarrow$ true objective (Eq. (3) or Eq. (5))
 - 8: **else**
 - 9: $\theta_r^i \leftarrow \theta_r$ /* Reward hyper-parameters are fixed. */
 - 10: $\theta_\pi^i, \theta_Q^i \leftarrow \text{Select}(\mathbf{n}_{\min}, \mathbf{n}_{\max}, T)$ with [18]
 - 11: $obj_{fn} \leftarrow$ cumulative reward.
 - 12: **end if**
 - 13: /* Train the agent with the selected hyper-parameters. */ $\pi^i = RL(\text{Actor}(\theta_\pi^i), \text{Critic}(\theta_Q^i), r(s, \theta_r^i))$
 - 14: $e^i \leftarrow$ Evaluate π^i w.r.t. obj_{fn} objective function.
 - 15: $T \leftarrow T \cup (e^i, \pi^i, \theta_\pi^i, \theta_Q^i)$ /* Save the trial. */
 - 16: **end for**
 - 17: /* Select best hyper-parameters according to Eq. (7) or (8). */ $\tilde{\pi}, \tilde{\theta}_r, \tilde{\theta}_\pi, \tilde{\theta}_Q \leftarrow \arg \max_{e \in T} T$
 - 18: **return** $\tilde{\pi}, \tilde{\theta}_r, \tilde{\theta}_\pi, \tilde{\theta}_Q$ /* Eq. (9) */
-

Next, we repeat a similar process to find the best actor and critic architecture w.r.t. to $\tilde{\theta}_r$ (Line 2, Alg. 1). In this case, the optimization objective is to maximize the cumulative reward (Line 11, Alg. 2). This time, at each trial we train asynchronously $RL(\text{Actor}(\theta_\pi^j), \text{Critic}(\theta_Q^j), R(\tilde{\theta}_r))$ and evaluate the objective w.r.t. Eq. (2). For P2P that is Eq. (4), and for PF it is Eq. (6). Lastly the best actor and critic architecture corresponds to the trial with the best return,

$$\tilde{\theta}_\pi, \tilde{\theta}_Q = \arg \max_j Obj(\theta_\pi^j, \theta_Q^j, \tilde{\theta}_r), \quad (8)$$

(Line 17). The final policy trained by AutoRL is

$$\tilde{\pi}(s|W_{\tilde{\theta}_\pi}) = RL(\text{Actor}(\tilde{\theta}_\pi), \text{Critic}(\tilde{\theta}_Q), R(\tilde{\theta}_r)) \quad (9)$$

Algorithm 2 scales linearly with the number of trials. It is important to note that the number of concurrent trials, n_{mc} , needs to be much smaller than the total number of trials, n_g , in order for the algorithm to have enough completed experience when selecting the next set of parameters in Lines 5 and 10. If there are no completed trials, the parameters are randomly selected. Overall, the Algorithm requires n_{mc} processors, and runs $\frac{n_g}{n_{mc}}$ times longer than vanilla RL.

IV. RESULTS

A. Setup

The training environment, generated from a real building floor plan, is 23 m by 18 m and contains no moving obstacles (Fig. 2a). The simulated differential drive robot a point mass, with a 64-beam 1D lidar with 220 degrees field of view with Gaussian distributed noise $\mathcal{N}(0, \sigma_{\text{Lidar}})$. The robot’s action space is the linear and angular velocities with bounds of $[-0.2, 1.0]$ m/s and $[-1.0, 1.0]$ radian/s, respectively. The details of our noise model are in the supplementary materials.

Actor and critic are three-layers deep for both tasks. We choose wide and shallow feed-forward fully-connected networks, because their fast inference time makes them feasible for on-board high frequency robot control. The critic consists of a two-layer observation embedding joined with the action network by two fully connected layers. We select DDPG for the learning algorithm, and Vizier [16] with CMA-ES [18] for hyperparameter tuning.

We compare P2P agent with vanilla hand-tuned DDPG [25], artificial potential fields (APF) [23], Dynamic Window Approach (DWA) [14], and behavior cloning (BC) [33]. Behavior cloning uses the same neural network architecture as learned with AutoRL, but relies on a supervised loss instead of a reward. The PF agent is compared with guidance paths generated with PRMs [22] with straight line local planner, and combined vanilla DDPG, APF, and DWA, to guide the robot, resulting in methods PRM-RL [12], PRM-APF [8], and PRM-guided DWA, denoted as PRM-DWA [14*] to differentiate from DWA without a guidance path. We test the PF and P2P policies in three previously unseen large environments (Figs. 2b, 2c, 2d). For the P2P policy, the start and goal are randomly selected to be between 5 and 10 meters apart. For the PF, the starts and goals are randomly chosen, requiring at least 35 meters Euclidean distance between start and goal.

B. Training

We train the P2P and PF agents for 1000 trials, running 100 agents in parallel, each for 5 million training steps. Each agent takes about 12 hours to train, with the complete AutoRL run completing in several days. For reproducibility, the supplementary materials contain the parameters found by AutoRL.

Fig. 3 shows the impact of the AutoRL on the PF task. Across both reward and network shaping there are certain parameter values where the agent does not perform well (Figs. 3a and 3c), although during the network shaping there are fewer bad trials. The learning curve of the non-shaped objective resembles the one depicted in Fig. 3b. The best agent after reward shaping

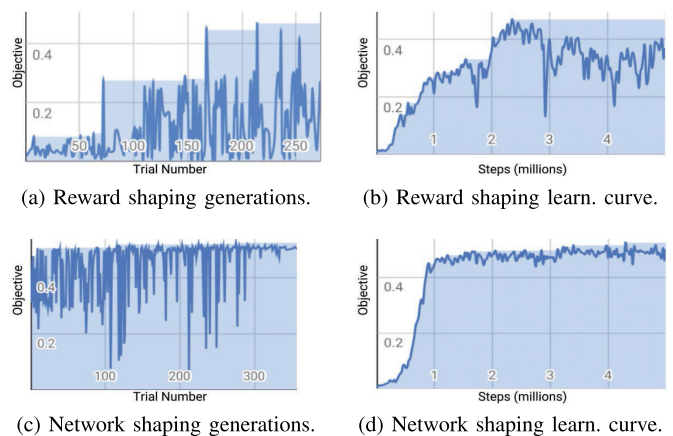




Fig. 3. size Results of AutoRL for the path following policy. Trail numbers over objective values (a) and (c). Training steps over objective (b) and (d).



(a) PF with AutoRL (Ours), RL [25], PRM-APF [8], PRM-DWA [14]*, PRM-RL [12]



(b) P2P with AutoRL (Ours), RL [25], APF [23], DWA [14], BC [33]

Fig. 4. Success rates for PF and P2P tasks in three environments.

(Fig. 3b) shows signs on catastrophic forgetting that plagues DDPG training. Notice, that the best agent after both shaping phases (Fig. 3d) does not exhibit such forgetfulness, and the training could have been terminated sooner. The PF AutoRL policy achieves a true objective of 0.56, while the non-shaped policy reaches an objective of 0.26. The P2P agent trained with AutoRL has training objective of 0.90, while the hand-tuned one is 0.54. This is a promising result for the utility of AutoRL.

C. Generalization to New Environments

1) *PF*: AutoRL has near perfect success rate, 98.7% on average, across all three environments (Fig. 4 a), and is the only method that can transfer to unseen, large, real building sized environments. PF with AutoRL is 11% more successful than hand tuned policy that has average success rate of 88.3%; and 23% more successful than best non-learned baseline (PRM-DWA with 80% success rate). PRM-APF is consistently under performing. The primary failure causes for both PRM-APF and

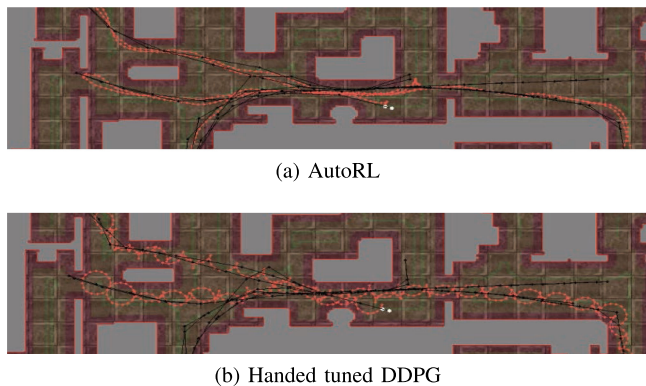


Fig. 5. PF policy sample trajectories (red) and guidance paths (black) found by (a) AutoRL and (b) hand-tuned RL.

PRM-DWA are wall collisions and getting stuck in a local minimum, especially when the guidance path generated by PRM is close to obstacles, which creates a local minimum [8]. PRM-RL performs worse than PRM-DWA and AutoRL, since it does not use hyperparameter tuning.

Curiously, the hand-tuned policy creates longer trajectories than the AutoRL one. Inspecting the trajectories (Fig. 5), it is clear that DDPG learns completely different behaviors under the two parameterizations. The AutoRL policy exhibits smooth, forward-moving behavior while the hand-tuned policy alternates between forward and reverse motion, resulting in a twirling behavior with longer paths. Since the 220 degree field of view lidar cannot detect obstacles in the back, the twirling behavior has a lower success rate.

2) *P2P*: Fig. 4 b shows that P2P’s generalization results are consistent with path following’s. The AutoRL policy’s success rate of 89% is highest (on average 71%, 26%, and 27% higher than APF, DWA, and BC respectively) across all environments. AutoRL exhibits smooth forward motion, while the failure cases are mostly due to complex static obstacles. Once again, despite promising success rate of 86.5%, the hand-tuned policy exhibits twirling, leading to high path lengths and subpar performance in noisy conditions.

D. Moving Obstacle Avoidance

We evaluate the shaped PF and P2P policies in a large environment (Fig. 2c) among up to 40 moving obstacles. The moving obstacles motion follows the social force model (SFM) [19] to mimic pedestrian motion.

1) *PF*: AutoRL policy outperforms hand-tuned DPPG and both PRM-APF and PRM-DWA in all numbers of obstacles (Fig. 6a). Note that although SFM avoids collision with the robot, simply executing the path using PRM-APF without reacting to obstacles (no repulsive potential) results in no success. This implies the robot must also partake in collision avoidance with moving obstacles. PRM-DWA can avoid moving obstacles and its performance is steady as the number of obstacles increases.

AutoRL PF performs uniformly across path guidance lengths up to 80 meters, and decreases slightly in the presence of moving obstacles, while PRM-APF degrades rapidly with guidance path

length (Figs. 6b and 6c). PRM-DWA’s success is also uniform across the guidance path lengths, but is lower than AutoRL’s.

2) *P2P*: P2P policy’s success degrades with the goal distance increase (Fig. 7) and its success rate is not affected with the number of obstacles (Fig. 7 b). Hand tuned RL, APF’s and DWA’s performance degrade rapidly with goal distance as well, which indicates that the main P2P failure cause, regardless of its implementation, is the inability to navigate among complex static obstacles rather than collision with moving obstacles. This is not surprising, since local planners are not designed to avoid complex static obstacles.

E. Robustness to Noise

Fig. 8 isolates one noise source at a time in order to investigate the impact on performance in an environment with 20 moving obstacles. Results show that PF and P2P policies are resilient to noise, even when the lidar noise is more than three times the radius of the robot ($\sigma_{\text{lidar}} = 1$ m). On the other hand, lidar noise heavily influences the success of APF and DWA (for both PRM and local planner variants). This is expected since obstacle clearance is used to compute the repulsive potential for APF and compute the objective function of DWA. In addition, the action of APF is computed by taking the gradient of the potential; such a greedy approach often guides the robot to collision or local minima.

F. Physical Robot Experiments

First, we investigate the difference between simulation and reality when PF and P2P policies were deployed on the Fetch robot [27]. For the PF task (Fig. 9 a), we manually specify a sequence of waypoints (black dots) as the guidance path (80.6 m in length). For the P2P task (Fig. 9 a) the start and goal are about 13.4 m apart. The robot reaches the goal without collision in all three runs, navigating roughly 240 m without collision. Figs. 9 a and 9 b (magenta) show one of robot trajectories. These trajectories are very close to the simulated ones (green for PF and red for P2P) but exhibit more turning than simulation, likely caused by the delay between sensing and action execution.

Next, we test the policy’s performance in a very narrow corridor (only 0.3 m wider than the robot, which is 0.3 m in diameter), and with moving obstacles (Fig. 1a), qualitatively over four trials. The enclosed video demonstrates the Fetch robot executing these policies. The robot reliably navigates in the corridor all four times. We also execute the P2P policy in a highly unstructured dynamic environment with a person playing with a dog (Fig. 1b). The robot stops and avoids obstacles reliably. Notice in the attached video how the agent adapts when blocked, moves away from the goal, and around obstacles. The only failure case we observed was caused by the obstacles undetectable by the 1D lidar such as human feet, which are below lidar’s field of view.

V. DISCUSSION

AutoRL is not sample efficient: it took 12 days to train 1000 agents. However, AutoRL learns high-quality navigation behaviors that transfer well across environments and are easy to




Fig. 6. PF success rate over (a) number of moving obstacles, (b) guidance path length with no obstacles, and (c) guidance path length with 20 moving obstacles for AutoRL (Ours), hand tuned RL [25], PRM-APF [8], and PRM-DWA [14]*.




Fig. 7. P2P policy: AutoRL (Ours), hand tuned RL [25], APF [23], and DWA [14] success rates (a) without and (b) with 20 moving obstacles over goal distance.

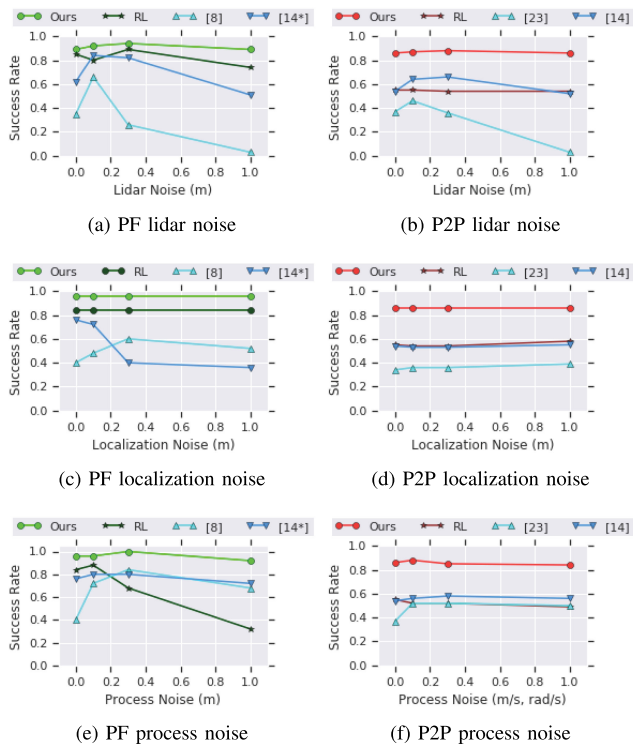


Fig. 8. Success rate over (top) lidar, (middle) localization noise, and (bottom) process noise with 20 moving obstacles. (left) Path Following compared to the hand tuned RL, PRM-APF [8] and PRM-DWA [14]*. (right) P2P compared to hand tuned RL [25], APF [23] and DWA [14] navigating from random starts to goals, 5-10 m apart.

deploy on-robot. For navigation agents, the extra training cost is justified by better quality policies.

End-to-end learning with AutoRL effectively creates tightly-coupled perception, planning and controls. The end-to-end POMDP setup enables robustness to noise. The comparison with the agents on the opposite side of the spectrum, traditional

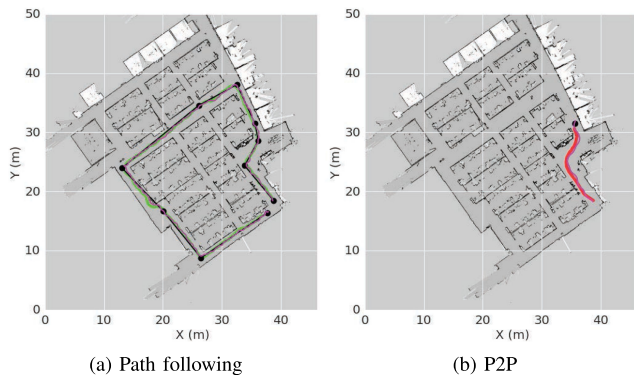


Fig. 9. (a) Path Following and (b) P2P on a real robot in an office building. The actual robot trajectories (magenta), the guidance path (black) and the trajectory in simulation (green for Path Following and red for P2P) overlaid over the 2D lidar map.

and learned, highlights the benefits of end-to-end learning. The hand-tuned DDPG agent, although it produces suboptimal trajectories, is still robust to noise (8). The traditional APF and DWA, not designed for tight integration between the controls and sensor, are brittle in the presence of noise.

Robustness to noise has an additional perk of overcoming guidance path imperfections. We have observed during hand-tuned training that path following is very sensitive to appropriate waypoint spacing and waypoint radius. It is only after including these parameters in reward tuning that the path-following agents learn high-quality behaviors. This is likely because the tuning finds the optimal distance from a waypoint w.r.t. to robot’s noise and abilities, essentially deciding to give a credit to reaching a waypoint for what is feasible on that particular robot.

The P2P policy found by AutoRL is more robust to local minima than APF and DWA, likely because the agent learned that following a wall usually leads to completing the goal. As demonstrated in the enclosed video, the agent is also willing to move away from the goal to avoid local minima. The failure point of P2P policy remains its inability to avoid large-scale local minima, such as moving from one room to another, which it was not designed to do.

VI. CONCLUSIONS

This letter presents AutoRL, an evolutionary automation layer around Deep Reinforcement Learning that, using large-scale hyperparameter optimization, selects reward and neural network architecture that are best for a given goal-oriented task. AutoRL is used to learn two navigation building blocks, P2P and PF

end-to-end behaviors. The resulting policies, although computationally expensive to train, exhibit more desirable behaviors compared to RL with hand-crafted hyperparameters and non-learned baselines. They generalize to new environments with moving obstacles, are robust to noise and are deployed to a physical robot without tuning. In the future, we plan to evaluate AutoRL on kinodynamic robots and high dimensional robots such as mobile manipulators.

ACKNOWLEDGMENT

The authors would like to thank J. C. Kew, O. Ramirez, L. Tapia, V. Vanhoucke, and C. Harris for helpful discussions.

REFERENCES

- [1] N. Altuntaş, E. Imal, N. Emanet, and C. N. Öztürk, "Reinforcement learning-based mobile robot navigation," *Turkish J. Elect. Eng. Comput. Sci.*, vol. 24, no. 3, pp. 1747–1767, 2016.
- [2] M. Andrychowicz *et al.*, "Learning to learn by gradient descent by gradient descent," in *Proc. Advances Neural Inf. Process. Syst.*, 2016, pp. 3981–3989.
- [3] E. A. Antonelo and B. Schrauwen, "On learning navigation behaviors for small mobile robots with reservoir computing architectures," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 4, pp. 763–780, Apr. 2015.
- [4] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," in *Proc. Neural Inf. Process. Syst.*, 2016, pp. 1471–1479.
- [5] B. Bischoff *et al.*, "Hierarchical reinforcement learning for robot navigation," in *Proc. Eur. Symp. Artif. Neural Netw.*, 2013, pp. 227–232.
- [6] C. Cadena *et al.*, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016.
- [7] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, "Deepdriving: Learning affordance for direct perception in autonomous driving," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 2722–2730.
- [8] H.-T. Chiang, N. Malone, K. Lesser, M. Oishi, and L. Tapia, "Path-guided artificial potential fields with stochastic reachable sets for motion planning in highly dynamic environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2015, pp. 2347–2354.
- [9] H.-T. L. Chiang, B. HomChaudhuri, L. Smith, and L. Tapia, "Safety, challenges, and performance of motion planners in dynamic environments," in *Proc. Int. Symp. Robot. Res.*, 2017, pp. 1–16.
- [10] F. Fathinezhad, V. Derhami, and M. Rezaeian, "Supervised fuzzy reinforcement learning for robot navigation," *Appl. Soft Comput.*, vol. 40, pp. 33–41, 2016.
- [11] A. Faust, H.-T. Chiang, N. Rackley, and L. Tapia, "Avoiding moving obstacles with stochastic hybrid dynamics using pearl: Preference appraisal reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2016, pp. 484–490.
- [12] A. Faust *et al.*, "PRM-RL: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, Brisbane, Qld, Australia, 2018, pp. 5113–5120.
- [13] C. Florensa, D. Held, M. Wulfmeier, M. Zhang, and P. Abbeel, "Reverse curriculum generation for reinforcement learning," in *Proc. 1st Annu. Conf. Robot Learn.*, S. Levine, V. Vanhoucke, and K. Goldberg, Eds. 2017, pp. 482–495.
- [14] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robot. Autom. Mag.*, vol. 4, no. 1, pp. 23–33, Mar. 1997.
- [15] W. Gao, D. Hsu, W. S. Lee, S. Shen, and K. Subramanian, "Intention-net: Integrating planning and deep learning for goal-directed autonomous navigation," in *Proc. 1st Annu. Conf. Robot Learn.*, S. Levine, V. Vanhoucke, and K. Goldberg, Eds. Nov. 2017, pp. 185–194.
- [16] D. Golovin, B. Solnik, S. Moitra, G. Kochanski, J. Karro, and D. Sculley, "Google vizier: A service for black-box optimization," in *Proc. ACM Int. Conf. Knowl. Discovery Data Mining*, 2017, pp. 1487–1495.
- [17] A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu, "Automated curriculum learning for neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 1311–1320.
- [18] N. Hansen, A. Ostermeier, and A. Gawelczyk, "On the adaptation of arbitrary normal mutation distributions in evolution strategies: The generating set adaptation," in *Proc. Int. Conf. Genetic Algorithms*, 1995, pp. 57–64.
- [19] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Phys. Rev. E*, vol. 51, no. 5, pp. 4282–4286, 1995.
- [20] B. Ivanovic, J. Harrison, A. Sharma, M. Chen, and M. Pavone, "Barc: Backward reachability curriculum for robotic reinforcement learning," *CoRR*, abs/1806.06161, 2018.
- [21] K. Judah, A. P. Fern, P. Tadepalli, and R. Goetschalckx, "Imitation learning with demonstrations and shaping rewards," in *Proc. 28th AAAI Conf. Artif. Intell.*, C. E. Brodley and P. Stone, Eds. 2014, pp. 1890–1896.
- [22] L. E. Kavraki, P. Švestka, J. C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Trans. Robot. Autom.*, vol. 12, no. 4, pp. 566–580, Aug. 1996.
- [23] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. Robot. Res.*, vol. 5, no. 1, pp. 90–98, 1986.
- [24] S. M. LaValle, J. James, and J. Kuffner, "Randomized kinodynamic planning," *Int. J. Robot. Res.*, vol. 20, no. 5, pp. 378–400, 2001.
- [25] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," 2015, arXiv:1509.02971.
- [26] E. Marder-Eppstein, E. Berger, T. Foote, B. Gerkey, and K. Konolige, "The office marathon: Robust navigation in an indoor office environment," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2010, pp. 300–307.
- [27] M. Wise, M. Ferguson, D. King, E. Diehr, and D. Dymesich, "Fetch & freight: Standard platforms for service robot applications," in *Proc. Workshop Auton. Mobile Service Robots Int. Joint Conf. Artif. Intell.*, 2016.
- [28] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [29] A. Mousavian, A. Toshev, M. Fiser, J. Kosecka, and J. Davidson, "Visual representations for semantic target driven navigation," *CoRR*, abs/1805.06066, 2018.
- [30] A. Y. Ng, D. Harada, and S. Russell, "Policy invariance under reward transformations: Theory and application to reward shaping," in *Proc. 16th Int. Conf. Mach. Learn.*, 1999, vol. 99, pp. 278–287.
- [31] M. Olsen, N. Siegelmann-Danieli, and H. Siegelmann, "Robust artificial life via artificial programmed death," *Artif. Intell.*, vol. 172, no. 6, pp. 884–898, 2008.
- [32] M. Pfeiffer, M. Schaeuble, J. I. Nieto, R. Siegwart, and C. Cadena, "From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2017, pp. 1527–1533.
- [33] D. A. Pomerleau, "Efficient training of artificial neural networks for autonomous navigation," *Neural Comput.*, vol. 3, no. 1, pp. 88–97, Mar. 1991.
- [34] E. Real *et al.*, "Large-scale evolution of image classifiers," in *Proc. 34th Int. Conf. Mach. Learn.*, D. Precup and Y. W. Teh, Eds. Sydney, NSW, Australia, Aug. 2017, pp. 2902–2911.
- [35] W. D. Smart and L. P. Kaelbling, "Effective reinforcement learning for mobile robots," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2002, vol. 4, pp. 3404–3410.
- [36] P. Su, D. Vandyke, M. Gašić, N. Mrkšić, T. Wen, and S. Young, "Reward shaping with recurrent neural networks for speeding up on-line policy learning in spoken dialogue systems," in *Proc. Special Interest Group Discourse Dialogue*, 2016, pp. 417–421.
- [37] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, "Value iteration networks," in *Proc. Neural Inf. Process. Syst.*, 2016, pp. 2154–2162.
- [38] J. X. Wang *et al.*, "Learning to reinforcement learn," 2016, arXiv:1611.05763.