

Harnessing Reinforcement Learning for Neural Motion Planning

Tom Jurgenson¹ and Aviv Tamar¹

¹Department of Electrical Engineering, Technion

Abstract—Motion planning is an essential component in most of today’s robotic applications. In this work, we consider the *learning* setting, where a set of solved motion planning problems is used to improve the efficiency of motion planning on different, yet similar problems. This setting is important in applications with rapidly changing environments such as in e-commerce, among others. We investigate a general deep learning based approach, where a neural network is trained to map an image of the domain, the current robot state, and a goal robot state to the next robot state in the plan. We focus on the learning algorithm, and compare supervised learning methods with reinforcement learning (RL) algorithms. We first establish that supervised learning approaches are inferior in their accuracy due to insufficient data on the boundary of the obstacles, an issue that RL methods mitigate by actively exploring the domain. We then propose a modification of the popular DDPG RL algorithm that is tailored to motion planning domains, by exploiting the known model in the problem and the set of solved plans in the data. We show that our algorithm, dubbed DDPG-MP, significantly improves the accuracy of the learned motion planning policy. Finally, we show that given enough training data, our method can plan significantly faster on novel domains than off-the-shelf sampling based motion planners. Results of our experiments are shown in <https://youtu.be/wHQ4Y4mBRb8>.

I. INTRODUCTION

Motion planning – the problem of finding a collision free trajectory for a robot – is a fundamental component in almost all robotic applications deployed today [28, 27]. Sampling based motion planners, such as probabilistic roadmaps [22] and rapidly exploring random trees [27], have been studied extensively, can be guaranteed to converge to an optimal solution [21], and are common practice in various robotic domains and off-the-shelf software implementations [46, 45]. However, for domains where the environment can change rapidly, such as in e-commerce applications, industry 4.0 [26], or home robotics, it is desired to plan *fast*, and the computational burden of sampling based planners can be limiting.

Consider, for example, planning pick-and-place motions for a robot in an e-commerce warehouse that handles various products. Building on the insight that the changes in the task will mostly be limited, e.g., just the product will change, while the robot and surrounding will stay the same, several recent studies have explored a *learning* setting for motion planning, where data from previously solved motion plans is used to speed-up planning in new domains [6, 20, 10, 12].

In recent years, deep learning has proven capable of learning complex patterns in data for various decision making domains such as computer vision, protein folding, and

games [24, 14, 43]. Motivated by these successes, we focus here on approaches that we collectively term *neural motion planners* [35, 5, 50, 38, 39], which use deep learning to approximate a motion planning computation. In a neural motion planner, a deep neural network is trained to map features of the domain (e.g., an image), the current robot state, and a goal robot state to the next robot state in the motion plan. By training on a set of motion planning domains, the network is hypothesized to learn the patterns which make for a successful motion plan, and, once trained, such a network can be used to quickly predict a motion plan in novel domains without running a heavy motion planning computation.

In this work, we investigate the algorithmic aspects of training a neural motion planner to solve nontrivial motion planning problems. We first consider a supervised learning approach, where previous plans are simply imitated [39, 16]. We observe that for high-dimensional domains that require high precision, the success of this approach is limited, which we attribute to insufficient data distribution on the boundary of the obstacles. We propose that reinforcement learning (RL; [47]) has the potential to overcome this problem, since the exploration process in RL will naturally drive the agent to investigate important areas in the domain, such as obstacle boundaries.

While RL algorithms are known to require extensive computation, the motion planning problem presents several features that can be exploited to dramatically improve RL performance. Our main contribution is an RL algorithm that exploits two features of motion planning problems: the fact that a model of the dynamics is known, and the fact that we can collect demonstration data offline through sampling based planners, to reduce the variance in training, and to perform efficient exploration.

We show that our method leads to significantly better accuracy, and much faster training times. In particular, we demonstrate predicting motion plans with almost perfect accuracy on a 4-dimensional robotic arm domain with very challenging narrow passages. With enough training time, our method learns to plan significantly faster than off-the-shelf sampling based planners. Our results suggest that with suitable training algorithms, deep learning can provide competitive results for learning in motion planning, opening the door to further investigations of network architectures and combinations of planning and learning.

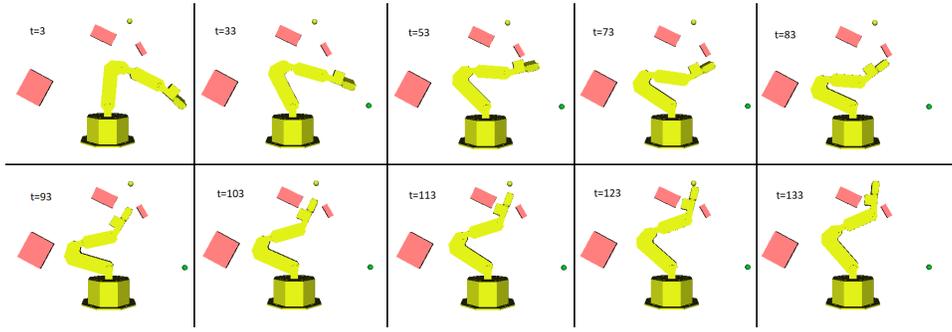


Fig. 1: DDPG-MP solves a narrow passage work-space. The figure shows several joint configurations along a successful plan executed by DDPG-MP for a narrow passage work-space (dubbed the *hard* scenario in Section VII). The episode starts at top-left image and ends at the bottom right. The green and yellow spheres mark the starting-pose and end-pose of the end-effector, t marks the time-step. Note the high precision required to navigate through the two rightmost obstacles without collision.

II. BACKGROUND

We begin with background on motion planning, imitation learning, and RL.

A. Motion Planning

We follow a standard motion planning formulation [28]. Let C denote the configuration space (joint space) of the robot, let S denote the task space (position and orientation) of the end effector, and let $FK : C \rightarrow S$ be the robot’s forward kinematics. Let $F \subset C$ denote the free space, and $\phi : C \rightarrow \{True, False\}$ a collision predicate that returns *True* if $c \in F$. In the motion planning problem, the robot starts from a joint position $c_0 \in C$, and is required to produce a trajectory plan $T : [0, 1] \rightarrow C$ in the free space that ends with the end-effector at a goal pose $g \in S$:

$$T(0) = c_0, \quad FK(T(1)) = g, \quad \forall a \in [0, 1] : \phi(T(a)) = True. \quad (1)$$

Popular motion planning algorithms rely on sampling the workspace, and effectively return a list of configurations c_0, \dots, c_T which make up the motion plan. A continuous plan is then derived by interpolation or smoothing [27]. We denote by $W = \{F, c_0, g\}$ a tuple that describes the *workspace* – the parameters of a motion planning problem for a given robotic domain.

B. Reinforcement Learning

In Reinforcement Learning (RL), an agent interacts with an environment E by observing the current state and executing an action in discrete time steps. Once an action is executed the state changes (the change can be either deterministic or stochastic), the agent receives a reward and the process repeats until termination. This can be modeled as a Markov Decision Process (MDP [47]) with a state space S , action space A , an initial state distribution $\rho(s_1)$, transition dynamics $p(s_{t+1}|s_t, a_t)$, and scalar reward function $r(s_t, a_t)$.

The goal of the agent is to interact with E by selecting actions according to some policy $\pi(a|s)$ (possibly stochastic) which maximizes the γ -discounted future rewards defined for time step t as $R_t = \mathbb{E} [\sum_{i=t}^{\infty} \gamma^{i-t} r(s_i, a_i)]$. Let $Q^\pi(s, a)$ be

the expected discounted future reward obtained by following policy π , after executing action a from state s : $Q^\pi(s, a) = \mathbb{E} [R_t | s_t = a, a_t = a, \pi]$. An important relation between current and future Q^π is the Bellman equation:

$$Q^\pi(s_t, a_t) = \mathbb{E}_{s_{t+1} \sim p} [r(s_t, a_t) + \gamma \mathbb{E}_{a_{t+1} \sim \pi(s_{t+1})} [Q^\pi(s_{t+1}, a_{t+1})]] \quad (2)$$

Let $Q^*(s, a)$ be the maximal expected discounted reward obtained by any strategy, after executing action a from state s : $Q^*(s, a) = \max_{\pi} \mathbb{E} [R_t | s_t = a, a_t = a, \pi]$. For the optimal policy π^* we have $Q^{\pi^*}(s, a) = Q^*(s, a)$ for all states and actions. The goal of an RL algorithm is to find such a policy.

C. Actor-Critic Algorithm: DDPG

Actor-critic is a popular RL approach [47]. The actor $\pi(s|a)$ models a policy, while the critic $\hat{Q}(s, a)$ approximates the state-action value of the actor. The learning process interleaves improving the critic (for example, by using temporal difference learning [47]), with updating the policy towards actions with higher state-action values.

Deep Deterministic Policy Gradient (DDPG [31]) is an actor-critic RL algorithm designed for continuous state and action spaces, where π and \hat{Q} are modeled using neural networks, and π is a deterministic policy with weights θ . DDPG is an off-policy algorithm: it learns by observing actions taken by a stochastic exploration policy β , and storing the observed transitions in a *replay buffer* [33]. DDPG alternates between learning \hat{Q} and learning π , fixing one while training the other. Specifically, the update of π ascends in the direction which maximizes \hat{Q} , motivated by the policy gradient theorem:

$$\nabla_{\theta} J \approx \mathbb{E}_{s \sim \beta} [\nabla_{\theta} \hat{Q}(s, a)], \quad (3)$$

where J denotes the expected discounted return [31].

D. Hindsight Experience Replay

In a motion planning setting, the policy needs to be *goal-conditioned*, i.e., taking in a goal state as an input parameter. The recent hindsight experience replay (HER) [4], is an

extension of off-policy RL methods such as DDPG to goal-conditioned policies. The idea is that, upon failing an episode (not reaching the goal), the agent adds to the replay-buffer an episode comprised of the same states and actions seen in the failed episode, but with the original goal replaced with the last encountered step – making the episode a success in hindsight. By making success more frequent, the agent is trained by an implicit curriculum, which eventually learns to cover the goal space. Indeed, HER has led to state-of-the-art results in goal-conditioned RL [4, 44, 30].

E. Imitation Learning

In imitation learning [41], a decision making policy is learned from observing expert behavior. Following the RL notation, given a data set of N expert demonstrations $\{s_i^{exp}, a_i^{exp}\}_{i=1}^N$, the problem is to find a policy $\pi = P(a|s)$ that imitates the expert behavior. A popular imitation learning method is behavioral cloning [37], where π is learned by maximizing the data log likelihood $\sum_{i=1}^N \log \pi(a_i^{exp}|s_i^{exp})$.

DAGger [40] is an imitation learning algorithm that corrects the *distribution mismatch problem*, where the distribution of states in the learning corpus may be different than the distribution of states encountered by the learned policy. In DAGger, after training a policy using behavioral cloning, the policy is executed, and additional expert data is first collected for all states visited, and then added to the agent’s data-set. Afterwards, a new policy is trained on the updated data-set and the process repeats to convergence.

III. PROBLEM FORMULATION

We consider the problem of learning for motion planning. In this setting, we assume that motion planning workspaces are generated from some (unknown) distribution $P(W)$. We are given a training data set of N workspaces $\{W_i\}_{i=1}^N$ drawn from P , and the goal is to use this data to plan faster in an unseen test workspace W_{test} , also drawn from P .

We focus on an approach termed neural motion planning (NMP) [39]. In NMP, a workspace W is described by a set of features, or context vector, w . The goal in NMP is to train a neural network policy π that maps w and the current configuration of the robot c_t to the next step in a motion plan c_{t+1} . Once trained, such a network can produce a motion plan by unrolling the predictions to produce a list of states c_0, \dots, c_T . We say that the motion plan is successful if it satisfies the motion planning problem conditions, i.e., $FK(c_T) = g$, and $\phi(c_t) = True \quad \forall 0 \leq t \leq T$.¹ In the remainder of this work, we investigate how to best train the neural network in NMP.

IV. IMITATION LEARNING VS. RL

A straightforward method for training an NMP is using imitation learning. In this approach, a standard motion planner (e.g., RRT*) is used to solve all the training work spaces. These motion plans are used to train a neural network using, e.g., behavioral cloning. This approach has been shown to

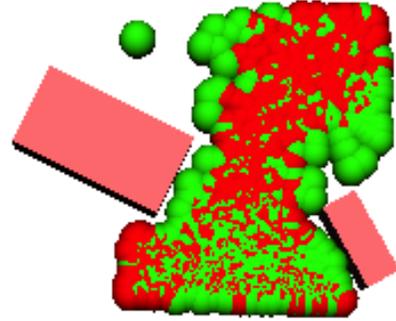


Fig. 2: Data Distribution on Edges of Obstacles. Spheres indicate visited end-effector poses seen during training. Green spheres were generated by an RL agent (DDPG) while red spheres are imitation learning targets (i.e., produced by the motion planner). Note that near obstacle boundaries, green spheres are more dominant, indicating the sparsity of data that imitation learning agents have on the edges of obstacles.

work well for discrete planning problems [16, 9], and recently has been investigated in the context of NMP [39, 5, 50].

In our investigation, we found that for domains that require high precision in the motion plan, such as a robot moving through tight passages, the performance of imitation learning is severely limited (see explicit results in Section VII). We attribute this finding to the distribution of the training data, as depicted in figure 2 for the end effector of a robot arm moving in a narrow passage. In NMP, it is fundamental that the neural network learns not to hit obstacles. However, since the expert produces perfect motion plans, there is no information about hitting an obstacle in the training data. Note in Figure 2 that the data distribution for imitation learning does not cover the obstacle boundaries, thus there is no reason to expect that a network trained with this data will learn not to hit obstacles. This is an instance of the distribution mismatch problem described by Ross et al. [40]. While in principle the DAGger algorithm [40] could mitigate this issue, it is too costly to compute in practice as it requires running motion planning on every sample in the data. One can instead acquire expert data only for states that resulted in a collision. However, as we report in the supplementary material (Section C), this approach did not lead to a significant improvement, which we believe is due to these samples not being well-balanced with the rest of the data.

We propose that RL is more suitable for training NMPs. In RL, the agent starts out with a policy that does not solve the task, and through trial and error it gradually collects data on how not to hit obstacles, thereby improving its policy. Indeed, as depicted in Figure 2, the data distribution in RL has much more presence near the obstacle boundaries – the critical decision points for motion planning. Therefore, in this work we pursue an RL approach to NMP, as we describe next.

¹In practice, we can only verify that $\|FK(c_T) - g\| < \epsilon$, for a small ϵ .

V. RL FOR MOTION PLANNING

In this section we first describe an RL formulation for NMP, and then propose a specialized RL algorithm for this domain.

We consider an episodic RL setting [47], where at each episode a workspace W is randomly selected from the training set, and the task is to plan a successful motion plan in it. At time step t of an episode, the robot configuration is c_t and the context vector for W is w . We define the state s_t to be the tuple $\{w, c_t\}$, and the action to be the difference in the robot configuration:²

$$c_{t+1} = c_t + a_t. \quad (4)$$

Note that this implies deterministic dynamics, since the context vector w is assumed to remain constant. The next state is:

$$s_{t+1} = \{w, c_t + a_t\} \doteq f(s_t, a_t). \quad (5)$$

Upon hitting an obstacle, or after reaching the goal, the episode terminates. We define a reward such that maximizing it produces a correct solution to the motion planning problem, by encouraging movement toward the goal and not hitting obstacles. Note that any transition $\{c_t, c_{t+1}\}$ can belong to one of three classes: free space movement, denoted as T_{free} ; collision, denoted as T_{col} ; and reaching the goal, denoted as T_{goal} . The transition reward is:

$$r_t^T = \begin{cases} -\epsilon, & (c_t, c_{t+1}) \in T_{free}, \\ 1, & (c_t, c_{t+1}) \in T_{goal}, \\ -1, & (c_t, c_{t+1}) \in T_{col}. \end{cases} \quad (6)$$

The constant $\epsilon > 0$ needs to be small enough such that the robot prefers moving in free space than colliding; we empirically set it to 0.01 in our experiments.

Some motion planning problems may include additional rewards. For example, in our experiments we wanted to encourage the robot not waste effort by trying to move a joint beyond its limits. To model this, we use a reward of r_t^D (domain specific reward), which in our case is the norm of the wasted movement (full details in section B). Finally, the full reward of the system is $r_t = r_t^T + r_t^D$.

A. The DDPG-MP Algorithm

We next present Deep Deterministic Policy Gradient for Motion Planning (DDPG-MP) – a deep RL algorithm tailored for training an NMP. To motivate the algorithm, we observe that in the context of RL, the motion planning setting admits the following unique features:

- 1) The dynamics (Eq. (4)) and reward (Eq. (6)) are known.
- 2) Similar to imitation learning, we can obtain expert demonstrations for the work spaces in our training data.

Since our problem is deterministic, we build on the DDPG algorithm [31], a strong and popular deep RL algorithm for continuous control. We add to DDPG two modifications that exploit the motion planning features described above: we

propose a model-based actor update that reduces variance, and use the expert demonstrations to perform efficient exploration. We next describe each idea in detail (for pseudo-code see Section A in the supplementary material).

B. Model-Based Actor Update

When training the actor network, the model is updated in the direction which maximizes Q^π . However, Q^π is not known, and we only have access to an approximation \hat{Q}^π learned by the critic, resulting in errors (due to bias and variance) in the policy gradient estimation of the actor. A key observation in our work is that we can reduce errors in the actor update by using the known dynamics model. Note that in a deterministic domain with a deterministic policy, by definition (cf. Eq. (2)) we have that for any $k \in 0, 1, \dots$:

$$Q^\pi(s_t, a_t) = r(s_t, a_t) + \dots + \gamma^{k-1}r(s_{t+k}, \pi(s_{t+k})) + \gamma^k Q^\pi(s_{t+k+1}, \pi(s_{t+k+1})).$$

Thus, if we know the reward function and transition function, we can estimate Q^π as $r(s_t, a_t) + \dots + \gamma^{k-1}r(s_{t+k}, \pi(s_{t+k})) + \gamma^k \hat{Q}^\pi(s_{t+k+1}, \pi(s_{t+k+1}))$, and we expect that as k grows, the error in this approximation will reduce [17] (for $k \rightarrow \infty$ the error is zero). In our experiments, we found that $k = 1$ is enough to significantly improve the actor update, and in the following we focus on this case. Extending our result to $k > 1$ is straightforward.

Following the above derivation, keeping in mind that transitioning to a goal or obstacle terminates the episode, and using the specific form of the dynamics in our problem (5), we obtain the following actor update:

$$\nabla_\theta J = \mathbb{E}_{s \sim \beta} \left[\nabla_\theta \left(r(s_t, \pi(s_t)) + \gamma \mathbb{I}_{(c_t, c_{t+1}) \in T_{free}} Q(f(s_t, \pi(s_t)), \pi(f(s_t, \pi(s_t)))) \right) \right] \quad (7)$$

Note, however, that both the indicator function in (7) and our reward function (6) are non differentiable, and their step-function structure means that the reward gradients in (7) will not be informative.

We therefore propose to use a smoothed reward function \tilde{r} and a smooth approximation \tilde{p} to $\mathbb{I}_{(c_t, c_{t+1}) \in T_{free}}$ instead. While one can derive explicit smoothing functions for these, in our work we opted for a general approach using supervised learning. Specifically, we represent $\tilde{r}(w, c_t, c_{t+1})$ and $\tilde{p}(w, c_t, c_{t+1})$ as neural networks, and train them using random samples of robot configurations. Technical implementation details regarding these models are given in Section D of the supplementary material.³

C. Targeted Exploration via Expert Demonstrations Scoring

Another disadvantage of DDPG is sample inefficiency that occurs because of uninformed exploration. DDPG’s exploration policy selects a random action in every step with low probability to discover novel moves. This exploration policy makes succeeding in narrow passages highly unlikely, since

²Our work can be extended to kinodynamic motion planning by including the robot’s dynamics. We defer this to future work.

³There are many possible choices for reward smoothing. We chose to model it as a neural network since it was easy to implement in our setup.

these often require a very specific action sequence with no guiding reward signal along the way. Since the probability to randomly generate such a sequence of actions is extremely low, the agent can never reach the other side of the passage to collect the reward. Worse, randomly exploring this narrow corridor leads in practice to many collisions, making the agent avoid entering the passage in the first place.

We propose to overcome the issue by providing the agent with motion planner demonstrations for failed work spaces. This is especially important in tight passages since instead of failing over-and-over, after the first failure the agent gets a useful signal that it can use to learn appropriate actions. Note that this strategy implicitly generates a curriculum: easier failure cases resolve earlier during training and the agent stops receiving demonstrations for those, while harder cases are still supplied with demonstrations. At the end of training we expect to succeed in most attempts and thus not use this exploration strategy, which is inline with other exploration strategies where the noise is being reduced over time, e.g using an ϵ -greedy approach with a decreasing ϵ over time [47].

Let $\tau_W = (c_0^\tau, a_0, c_1^\tau, a_1 \dots a_{N-1}, c_N^\tau)$ be a trajectory produced by a motion planner, which is a solution to a failed workspace query $W = \{F, c_0, g\}$, i.e $c_0^\tau = c_0$ and $FK(c_N^\tau) = g$. In order to incorporate the information of τ_W to the replay buffer, we only need to assign a reward to every transition, to obtain $(c_i^\tau, a_i, r_i, c_{i+1}^\tau)$. Because τ_W is known to be a valid trajectory, this process is straightforward – by applying Eq. (6), to obtain r_i . This strategy can be seen as a ‘smart’ form of exploration, where instead of trying random actions, we are guided by the motion planner to try actions that result in success. We have found that this allows us to reduce the magnitude of the random exploration required for DDPG to work, thereby obtaining a much more stable algorithm.

VI. RELATED WORK

The idea of learning from experience to plan faster in motion planning has been studied under various approaches, such as using a library of past trajectories [7, 20, 6, 10], learning the sampling procedure [32, 19, 25], learning a latent representation of obstacles [19], and learning to select goals [12].

The idea of using a neural network for motion planning (termed here as NMP) dates back to the previous decade [15, 49]. Following recent advances in deep learning [24], interest in these methods has rekindled [35, 5, 50, 38, 39]. The potential advantages of NMP are the possibility of generalizing to different domains, working directly on raw perceptual signals such as images, and a low memory footprint: once trained, the NMP is a compact prediction model that does not require storing a large library of trajectories. While most previous work on NMP concerned the neural network architecture [e.g., 29, 39, 50, 5], in this work we focus on a parallel investigation – the training algorithm for NMP. We believe that with better architectures, our results can be further improved.

RL for continuous control has been explored extensively [42, 31]. Our method builds on two main ideas for

improving RL: using expert demonstrations, and a model-based update for the policy gradient. While similar ideas have been explored in the literature [e.g., 34, 36, 17], our formulation is tailored for the NMP setting, and is, to the best of our knowledge, novel.

Using neural networks for planning in discrete tasks has been explored for navigation [48], task planning and combinatorial optimization [16, 9] and games [43]. Interestingly, the seminal AlphaGo program [43] exploited expert demonstrations and the known structure of the game to improve RL. Our work exploits similar properties of continuous motion planning domains, therefore requiring different methods.

VII. RESULTS

In this section we evaluate the DDPG-MP algorithm on various NMP domains, and show that it obtains state-of-the-art results in training an NMP model. In addition, we investigate the components that contribute to DDPG-MP’s success using an ablation study. Finally, we show that using DDPG-MP we can significantly improve upon the planning times of conventional motion planners. We investigate the following questions:

- 1) Is RL more suitable for NMP than imitation learning?
- 2) Can DDPG-MP obtain state-of-the-art results for NMP?
- 3) Can NMP models trained with DDPG-MP generalize to unseen work spaces with a high success rate?
- 4) Can we use NMP models to plan faster than sampling based motion planners?

We next describe the experimental settings we used.

NMP Architecture: As outlined above, in this work we wish to disentangle the questions of NMP architectures from the NMP training algorithm. Therefore, we chose popular neural network architectures that were the same for all methods we tested. Specifically, for domains without image inputs we used fully connected neural networks, which are popular for continuous control [13]. For domains with image inputs we used convolutional layers followed by fully connected layers [33, 31]. We provide full details of the network architecture in Section B of the supplementary material.

Simulation environment: In our choice of a domain, we opted for an environment that can present challenging planning problems, yet be simple enough visually to not require complex perception efforts. We chose a 4DoF WidowX robotic arm [1], with movement restricted the XZ plane, and box-shaped obstacles in various sizes, positions, and orientations. Restricting movement to the plane allows us to capture the domain using a simple 2-dimensional image of the obstacles, which is natural for processing with convolutional neural networks, as discussed above. By varying the positioning of the obstacles, we can generate planning problems with various difficulty levels. We note that this setup is significantly more challenging than the point robots and car-like robots explored in previous NMP work [e.g., 39].

We used Tensorflow [2] for training the neural networks. Our code is available at <https://github.com/tomjur/ModelBasedDDPG>.

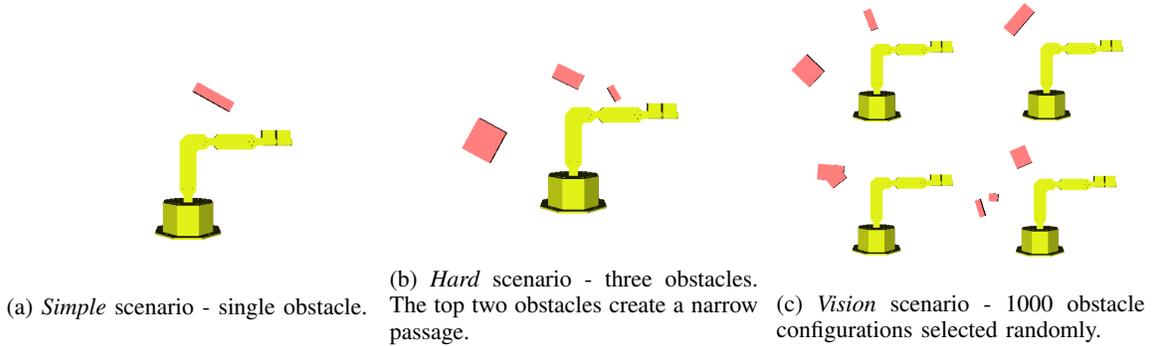


Fig. 3: Scenarios - the *simple* (left), *hard* (middle) and four random samples from the *vision* (right) scenarios used in our experiments. Note that the obstacles’ proximity to the arm make for a tight work space, especially in the *hard* scenario.

Evaluation metric: Evaluating the performance of agents trained with RL is a delicate matter [18], and in particular, how to choose the ‘best’ agent during training requires some validation metric. To address this, for each training run, the model iteration with highest test success-rate is taken and tested again on 1000 new validation work-spaces. This asserts that the reported scores are not just a lucky sample, but accurately capture the success rate of the model (the validation success-rates are mostly lower than the best test success-rates).

Scenarios: Within the simulation described above, we conducted experiments on three scenarios, which we term *simple*, *hard* and *vision*, as depicted in Figure 3. The *simple* scenario (Figure 3a) contains just a single obstacle. The *hard* scenario (Figure 3b) was manually designed with a challenging obstacle configuration: three obstacles, such that the two top obstacles create a narrow passage where a sequence of very specific movements is required in order to cross. For the NMP problem, in both *simple* and *hard* scenarios the obstacle configuration is fixed, and the only variation is in the starting and goal positions. Formally, $W = (F, c_0, g)$ contains a fixed F , and c_0 and g are sampled from a distribution of feasible trajectories. The context vector w is simply g .

The *vision* scenario, shown in Figure 3c, is used to investigate generalization to unseen obstacle configurations. We generated a data set of 1000 random obstacle configurations, partitioned to 80% train and 20% test. We report success rates on the test set. For this scenario all three parts of $W = (F, c_0, g)$ are randomly sampled, and the context vector in this case contains both an image I_F of the obstacle configuration, and the goal pose g .

We note that when sampling random work spaces (either obstacles or starting positions and goals), trivial problems, such as problems that do not require crossing obstacles, or where the goal is very close to the starting point, are frequent. To focus on interesting planning problems, we used rejection sampling to bias our testing to problems that require crossing from one side of an obstacle to the other. We note that in the *hard* and *vision* scenarios, the arm often needs to cross more than a single obstacle, making the task even more challenging. The difficulty of our domain is best appreciated

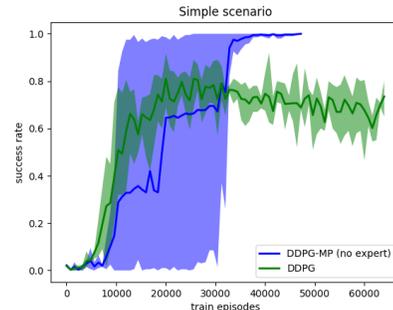


Fig. 4: *Simple* scenario success rate: visual comparison between DDPG in green, to DDPG-MP (no expert) in blue. Each curve shows the average test success rate of 3 runs, with the upper and lower bounds the maximal and minimal values respectively. Each data-point is the count of successful work-spaces out of 200 test work-spaces shown to the model during testing.

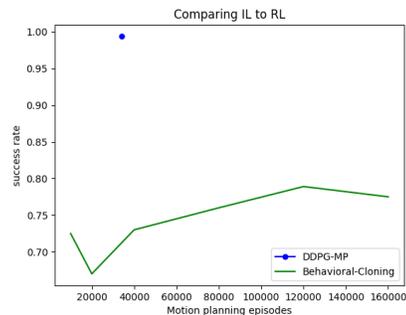


Fig. 5: Comparison between IL method to our method on the *simple* workspace.

in the accompanying video <https://youtu.be/wHQ4Y4mBRb8>. Full details of the sampling process are given in Section B. We next present our results.

A. Imitation Learning vs. RL

As mentioned above, we have found that imitation learning does not provide accurate enough training for NMP. We

Model	Test	Validation
DDPG	0.865	0.82533
DDPG-MP (no expert)	1.0	0.99366

TABLE I: Success rates *simple* scenario

demonstrate this here for the *simple* scenario, while similar results were obtained for other scenarios. We collected up to 160K expert demonstrations and trained imitation learning agents using behavioral cloning and DAgger. The results are presented in Figure 5, for various sizes of the training data set. The best behavioral-cloning and DAgger models reached a success rate of only 0.78 and 0.8 respectively. We also observed that the success-rate gains were diminishing, meaning that there is little contribution in adding more data. RL, on the other hand, was able to obtain a near-perfect NMP with a fraction of the expert demonstrations. This is explained in Figure 2 by observing the data distribution of RL and imitation learning, and noting that for RL the data is better distributed in the important areas – near obstacles.

B. RL for NMP

Next we evaluate DDPG-MP for NMP. We compare to two baselines: the original DDPG [31], and DDPG with HER [4], the current state-of-the-art in RL for learning goal-conditioned policies. Figures 4 and 6 plot the test success rate during training for the *simple* and *hard* scenarios. Tables II and III show the validation success rates for the *hard* and *vision* scenarios. Observe that DDPG-MP is the model with highest success rate in every experiment conducted, and on the *hard* scenario, alternative approaches only reach 0.31 success rate, while DDPG-MP reached a near-perfect success rate of 0.97. These results demonstrate the effectiveness of DDPG-MP in learning neural network policies for motion planning, especially when the problem is challenging in the sense that it requires navigating through narrow passages.

In the following we perform an ablation study to identify the components that make DDPG-MP so successful.

C. Investigating the Model Based Actor Update

In this section, we investigate the importance of the model-based actor update in DDPG-MP. To study this, we ablate the smart exploration component of DDPG-MP, and instead use the vanilla DDPG exploration. We call this algorithm DDPG-MP-no-expert. Thus, the only difference between DDPG-MP-no-expert and conventional DDPG is in changing the actor objective function from Eq. (3) to Eq. (7). As shown in figure 4 and table I, DDPG-MP-no-expert significantly improves the precision and speed-of-convergence in comparison to DDPG. Specifically, DDPG trained on the full 64K work-spaces reaches a validation success rate of 0.825. Meanwhile, DDPG-MP-no-expert needed on average 33082.66 episodes to reach a test success score of 1.0 and stop, achieving an average validation score of 0.9936.

Model	Test	Validation
DDPG	0.36833	0.31766
DDPG+HER	0.3633	0.286
DDPG-MP (no expert)	0.83	0.81933
DDPG-MP+HER (no expert)	0.72833	0.68866
DDPG-MP (full)	0.99	0.9733

TABLE II: Success rates *hard* scenario

D. Investigating Exploration via Expert Demonstrations

The *hard* scenario is a challenging motion planning environment to learn: the narrow corridor requires a very precise sequence of actions without much room for errors. Moreover, since most work-space queries (which are presented randomly to the agent) do not include this portion of the state space, it is rare for the agent to see states in the corridor and is therefore prone to the problem of catastrophic forgetting. Under regular RL exploration strategies this task is either impossible or requires a large amount of episodes which the exploration happens to give the correct sequence over and over.

The targeted exploration method via expert demonstrations of Section V-C is shown to be an effective solution to this problem. In this section, we both measure the contribution of this strategy, and compare it to HER, a popular state-of-the-art exploration strategy.

We start by comparing the baselines DDPG with and without HER to DDPG-MP (full) as shown in Figure 6a. We can see that DDPG-MP reaches a near-perfect validation success-rate of 0.9733, while the DDPG and DDPG+HER only reach 0.318 and 0.286 each, demonstrating that the complete DDPG-MP algorithm is better in this scenario than both baselines.

Next, we would like to understand the contribution of using demonstrations to explore. For this purpose, in Figure 6b we compare DDPG-MP with a vanilla exploration strategy (i.e., replacing the expert demonstrations exploration with the standard DDPG exploration) to DDPG-MP with a HER exploration strategy (i.e., replacing the expert demonstrations with HER) and the complete DDPG-MP algorithm.

Figure 6 shows the success rate over the number of training episodes, and Table II shows the average validation success-rate, both clearly show that DDPG-MP (full) is the best model. Comparing DDPG-MP+HER (no expert) to DDPG-MP (no expert) we see that there is no benefit for using HER, and it even degrades the performance. We conclude that using HER is either not beneficial in this case, or requires more fine-tuning in order to make it work for these types of challenging motion planning scenarios. In Section E of the supplementary material we provide an in-depth analysis of why HER is not well suited for such motion planning scenarios.

Finally, we analyze our proposed exploration strategy by comparing DDPG-MP (no expert) to DDPG-MP (full). DDPG-MP (full) reaches a near-perfect success rate of 0.9733 compared to 0.81933 of DDPG-MP (no expert), and also learns 4X faster. These results clearly show that DDPG-MP’s exploration strategy is beneficial.

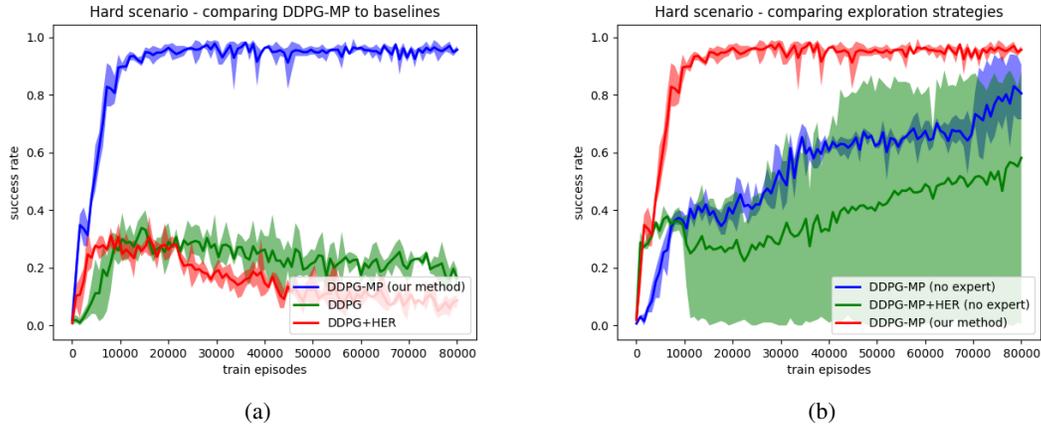


Fig. 6: *Hard* scenario results. 6a: Comparison with baselines. We compare DDPG and DDPG+HER in green and red, to DDPG-MP (our method) in blue. 6b: Comparing exploration strategies. We compare DDPG-MP with a vanilla DDPG exploration strategy (no expert) in blue, DDPG-MP with HER exploration strategy (DDPG-MP+HER) in green, and our exploration strategy utilizing expert demonstrations DDPG-MP (our method) in red. For both figures: each curve shows the average test success rate of 3 runs, with the upper and lower bounds the maximal and minimal values respectively. Each data-point is the count of successful work-spaces out of 200 test work-spaces shown to the model during testing.

Model	Test	Validation
DDPG	0.9233	0.89733
DDPG+HER	0.955	0.925
DDPG-MP (full)	0.96166	0.93566

TABLE III: Success rates for the *vision* scenario.

E. *Vision* scenario: Generalization to Unseen Obstacles

The previous two experiments showed that DDPG-MP can generalize to unseen work spaces, where the obstacles are fixed but the start and goal configurations are varied. In this experiment we also investigate the effect of changing the obstacle configurations – can we train NMP models that generalize to previously unseen obstacle configurations? To answer this question we use the *vision* scenario, where with each work space, a visual input of the obstacles configuration is given to the model. We hypothesize that if unseen work spaces share some similarities with the work spaces seen during training, DDPG-MP will be able to generalize and solve the unseen scenarios with high accuracy.

Table III shows that all models reach high success rates, with DDPG-MP getting the highest validation success rate of 0.93. We note that the individual scenarios that comprise the *vision* scenarios were selected from an easier distribution of obstacle configurations compared to the *hard* scenario, and this may be the reason that models such as DDPG and DDPG+HER are almost on-par with DDPG-MP.

F. *Vision* Scenario: Planning Times

Finally, we show that an NMP model trained with DDPG-MP has preferable running times to sample-based motion planners, for workspaces with an obstacle configuration not seen during training. To do this we compare the planning times it takes to produce trajectories for 100 random workspaces with unseen obstacle configurations (from the *vision* scenario

test set). DDPG-MP takes 8.55 seconds to compute, while it takes 50.93 seconds for OpenRAVE’s RRT motion planner – a 6X speedup in favor of DDPG-MP. The experiment was conducted on a Ubuntu 16 desktop machine, with a 12-core Intel i7-8700k 3.7GHz CPU, 32GB RAM and NVIDIA GeForceGTX 1080Ti. We note that OpenRAVE uses C++ code, while our code runs in non-optimized python. We expect that with dedicated optimization these results would improve.

These planning times should be contrasted with the accuracy rate of the NMP approach. For a 93% success rate with a 6X speedup, a naive strategy that first tries to sample a plan using NMP, and if it fails (by running it through a collision checker) falls back to a conventional motion planner, a significant speedup on average would still be guaranteed. That said, we believe that smarter ways of combining the NMP with a conventional planner could be devised (e.g., along the lines of [16], which used an NMP as a search heuristic for discrete planning domains).

VIII. CONCLUSION

We presented DDPG-MP: an RL based algorithm for training a neural motion planner. We showed that our method significantly improves the training time and final accuracy of the learned planner, compared to both imitation learning and state-of-the-art RL approaches. We further showed that our approach can be used to train neural networks that predict accurate motion plans given an image of the obstacles in the domain, and do so considerably faster than conventional motion planners.

In this work, we focused only on the learning algorithm for training a neural motion planner. Our results pave the way for further investigations of neural network architectures for motion planning, and combinations of learning based approaches with classical methods for motion planning.

REFERENCES

- [1] Trossen Robotics WidowX. <https://www.trossenrobotics.com/widowxrobotarm>. Accessed: 2019-01-30.
- [2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <http://tensorflow.org/>. Software available from tensorflow.org.
- [3] A. F. Agarap. Deep learning using rectified linear units (relu). *CoRR*, abs/1803.08375, 2018. URL <http://arxiv.org/abs/1803.08375>.
- [4] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. In *Advances in Neural Information Processing Systems*, pages 5048–5058, 2017.
- [5] M. J. Bency. *Towards Neural Network Embeddings of Optimal Motion Planners*. PhD thesis, UC San Diego, 2018.
- [6] D. Berenson, P. Abbeel, and K. Goldberg. A robot path planning framework that learns from experience. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3671–3678. IEEE, 2012.
- [7] M. S. Branicky, R. A. Knepper, and J. J. Kuffner. Path and trajectory diversity: Theory and algorithms. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 1359–1364. IEEE, 2008.
- [8] D. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *CoRR*, abs/1511.07289, 2015. URL <http://arxiv.org/abs/1511.07289>.
- [9] H. Dai, E. B. Khalil, Y. Zhang, B. Dilkina, and L. Song. Learning combinatorial optimization algorithms over graphs. *arXiv preprint arXiv:1704.01665*, 2017.
- [10] D. Dey, T. Y. Liu, B. Sofman, and J. A. Bagnell. Efficient optimization of control libraries. In *AAAI*, volume 1, page 4, 2012.
- [11] R. Diankov and J. Kuffner. Openrave: A planning architecture for autonomous robotics. *Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34*, 79, 2008.
- [12] A. D. Dragan, G. J. Gordon, and S. S. Srinivasa. Learning from experience in manipulation planning: Setting the right goals. In *Robotics Research*, pages 309–326. Springer, 2017.
- [13] Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International Conference on Machine Learning*, pages 1329–1338, 2016.
- [14] R. Evans, J. Jumper, J. Kirkpatrick, L. Sifre, et al. De novo structure prediction with deep-learning based scoring. In *In Thirteenth Critical Assessment of Techniques for Protein Structure Prediction (Abstracts)*, 1-4 December 2018.
- [15] R. Glasius, A. Komoda, and S. C. Gielen. Neural network dynamics for path planning and obstacle avoidance. *Neural Networks*, 8(1):125–133, 1995.
- [16] E. Groshev, A. Tamar, M. Goldstein, S. Srivastava, and P. Abbeel. Learning generalized reactive policies using deep neural networks. In *2018 AAAI Spring Symposium Series*, 2018.
- [17] N. Heess, G. Wayne, D. Silver, T. Lillicrap, T. Erez, and Y. Tassa. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, pages 2944–2952, 2015.
- [18] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger. Deep reinforcement learning that matters. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [19] B. Ichter and M. Pavone. Robot motion planning in learned latent spaces. *arXiv preprint arXiv:1807.10366*, 2018.
- [20] N. Jetchev and M. Toussaint. Trajectory prediction in cluttered voxel environments. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 2523–2528. IEEE, 2010.
- [21] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- [22] L. Kavraki, P. Svestka, and M. H. Overmars. *Probabilistic roadmaps for path planning in high-dimensional configuration spaces*, volume 1994. Unknown Publisher, 1994.
- [23] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <http://arxiv.org/abs/1412.6980>.
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [25] Y.-L. Kuo, A. Barbu, and B. Katz. Deep sequential models for sampling-based planning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6490–6497. IEEE, 2018.
- [26] H. Lasi, P. Fettke, H.-G. Kemper, T. Feld, and M. Hoffmann. Industry 4.0. *Business & Information Systems Engineering*, 6(4):239–242, 2014.
- [27] J.-C. Latombe. *Robot motion planning*, volume 124. Springer Science & Business Media, 2012.
- [28] S. M. LaValle. *Planning algorithms*. Cambridge university press, 2006.
- [29] D. V. Lebedev, J. J. Steil, and H. J. Ritter. The dynamic wave expansion neural network model for robot motion planning in time-varying environments. *Neural Networks*, 18(3):267–285, 2005.
- [30] A. Levy, R. Platt, and K. Saenko. Hierarchical reinforcement learning with hindsight. *arXiv preprint arXiv:1805.08180*, 2018.
- [31] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

- [32] S. R. Martin, S. E. Wright, and J. W. Sheppard. Offline and online evolutionary bi-directional rrt algorithms for efficient re-planning in dynamic environments. In *Automation Science and Engineering, 2007. CASE 2007. IEEE International Conference on*, pages 1131–1136. IEEE, 2007.
- [33] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- [34] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *arXiv preprint arXiv:1804.02717*, 2018.
- [35] M. Pfeiffer, M. Schaeuble, J. 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 *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1527–1533. IEEE, 2017.
- [36] M. Pfeiffer, S. Shukla, M. Turchetta, C. Cadena, A. Krause, R. Siegwart, and J. Nieto. Reinforced imitation: Sample efficient deep reinforcement learning for map-less navigation by leveraging prior demonstrations. *arXiv preprint arXiv:1805.07095*, 2018.
- [37] D. A. Pomerleau. ALVINN: An autonomous land vehicle in a neural network. In *NIPS*, pages 305–313, 1989.
- [38] A. H. Qureshi and M. C. Yip. Deeply informed neural sampling for robot motion planning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6582–6588. IEEE, 2018.
- [39] A. H. Qureshi, M. J. Bency, and M. C. Yip. Motion planning networks. *arXiv preprint arXiv:1806.05767*, 2018.
- [40] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.
- [41] S. Schaal, A. Ijspeert, and A. Billard. Computational approaches to motor learning by imitation. *Phil. Trans. of the Royal Society of London: Biological Sciences*, 358(1431):537–547, 2003.
- [42] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [43] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- [44] A. Srinivas, A. Jabri, P. Abbeel, S. Levine, and C. Finn. Universal planning networks. *arXiv preprint arXiv:1804.00645*, 2018.
- [45] I. A. Şucan and S. Chitta. MoveIt! URL <http://moveit.ros.org>.
- [46] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. doi: 10.1109/MRA.2012.2205651. <http://ompl.kavrakilab.org>.
- [47] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- [48] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel. Value iteration networks. In *Advances in Neural Information Processing Systems*, pages 2154–2162, 2016.
- [49] S. X. Yang and M. Meng. An efficient neural network approach to dynamic robot motion planning. *Neural Networks*, 13(2):143–148, 2000.
- [50] H. Zhang, E. Heiden, R. Julian, Z. He, J. J. Lim, and G. S. Sukhatme. Auto-conditioned recurrent mixture density networks for complex trajectory generation. *arXiv preprint arXiv:1810.00146*, 2018.