

# THINKING WHILE MOVING: DEEP REINFORCEMENT LEARNING WITH CONCURRENT CONTROL

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

We study reinforcement learning in settings where sampling an action from the policy must be done concurrently with the time evolution of the controlled system, such as when a robot must decide on the next action while still performing the previous action. Much like a person or an animal, the robot must think and move at the same time, deciding on its next action before the previous one has completed. In order to develop an algorithmic framework for such concurrent control problems, we start with a continuous-time formulation of the Bellman equations, and then discretize them in a way that is aware of system delays. We instantiate this new class of approximate dynamic programming methods via a simple architectural extension to existing value-based deep reinforcement learning algorithms. We evaluate our methods on simulated benchmark tasks and a large-scale robotic grasping task where the robot must “think while moving.”

## 1 INTRODUCTION

In recent years, Deep Reinforcement Learning (DRL) methods have achieved tremendous success on a variety of diverse environments, including video games (Mnih et al., 2015), zero-sum games (Silver et al., 2016), robotic grasping (Kalashnikov et al., 2018), and in-hand manipulation tasks (OpenAI et al., 2018). While impressive, all of these examples use a *blocking* observe-think-act paradigm: the agent assumes that the environment will remain static while it thinks, so that its actions will be executed on the same states from which they were computed. This assumption breaks in the *concurrent* real world, where the environment state evolves substantially as the agent processes observations and plans its next actions. As an example consider a dynamic task such as catching a ball; it is not possible to pause the ball mid-air to wait for the agent to decide on the next control to command. In addition to solving dynamic tasks where blocking models would fail, thinking and acting concurrently can provide practical qualitative benefits such as smoother, more human-like motions and the ability to seamlessly plan for next actions while executing the current one.

Despite these potential benefits, most DRL approaches are mainly evaluated in blocking simulation environments. Blocking environments make the assumption that the environment state will not change between when the environment state is observed and when the action is executed. This assumption holds true in most simulated environments, which encompass popular domains such as Atari (Mnih et al., 2013) and Gym control benchmarks (Brockman et al., 2016). The system is treated in a sequential manner: the agent observes a state, freezes time while computing an action, and finally applies the action and unfreezes time. However, in dynamic real-time environments such as real-world robotics, the synchronous environment assumption is no longer valid. After observing the state of the environment and computing an action, the agent often finds that when it executes an action, the environment state has evolved from what it had initially observed; we consider this environment a *concurrent environment*.

In this paper, we introduce an algorithmic framework that can handle concurrent environments in the context of DRL. In particular, we derive a modified Bellman operator for concurrent MDPs and present the minimal set of information that we must augment state observations with in order to recover blocking performance with Q-learning. We introduce experiments on different simulated environments that incorporate concurrent actions, ranging from common simple control domains to vision-based robotic grasping tasks. Finally, we show an agent that acts concurrently in a real-world

robotic grasping task is able to achieve comparable task success to a blocking baseline while acting 49% faster.

## 2 RELATED WORK

**Minimizing Concurrent Effects** Although real-world robotics systems are inherently concurrent, it is sometimes possible to engineer them into approximately blocking systems. Different examples include choosing low-latency hardware (Abbeel et al., 2006), optimizing software infrastructure, or designing lightweight model architectures (Cruz et al., 2017), all of which minimize the time spent during state capture and policy inference. Another option is to design actions to be executed to completion via closed-loop feedback controllers and the system velocity is decelerated to zero before a state is recorded (Kalashnikov et al., 2018). In contrast to these works, we tackle the concurrent action execution directly in the learning algorithm. Our approach allows us to reduce the engineering effort required to design a blocking environment and can be applied to tasks where it is not possible to wait for the system to come to rest between deciding new actions.

**Algorithmic Approaches** Other works utilize algorithmic modifications to directly overcome the challenges of concurrent control. Previous work in this area can be grouped into four approaches: (1) learning policies that are robust to variable latencies (Tan et al., 2018), (2) including past history such as frame-stacking (Haarnoja et al., 2018), (3) learning dynamics models to predict the future state at which the action will be executed (Firoiu et al., 2018; Amiranashvili et al., 2018), and (4) using a time-delayed MDP framework (Walsh et al., 2007; Firoiu et al., 2018; Schuitema et al., 2010). In contrast to these works, our approach is able to (1) optimize for a specific latency regime as opposed to being robust to all of them, (2) directly consider the properties of the source of latency as opposed to force the network to learn them from high-dimensional inputs, (3) avoid learning explicit forward dynamics models in high-dimensional spaces, which can be costly and challenging, and (4) consider environments where actions are interrupted as opposed to discrete-time time-delayed environments where multiple actions are queued and each action is executed until completion.

**Continuous-time Reinforcement Learning** While previously mentioned related works largely operate in discrete-time environments, framing concurrent environments as continuous-time systems is a natural framework to apply. In the realm of continuous-time optimal control, path integral solutions (Kappen, 2005; Theodorou et al., 2010) are linked to different noise levels in system dynamics, which could potentially include latency that results in concurrent properties. Finite differences can approximate the Bellman update in continuous-time stochastic control problems (Munos & Bourgine, 1998) and continuous-time temporal difference learning methods (Doya, 2000) can utilize neural networks as function approximators (Coulom, 2002). The effect of time-discretization (converting continuous-time environments to discrete-time environments) is studied in Tallec et al. (2019), where the advantage update is scaled by the time discretization parameter. We build on top of many of the theoretical formulations and findings in these works, and show their applications to deep reinforcement learning methods on more complex, vision-based robotics tasks.

## 3 VALUE-BASED REINFORCEMENT LEARNING IN CONCURRENT ENVIRONMENTS

In this section, we first introduce the preliminaries necessary for the discrete- and continuous-time RL formulations. We then describe the MDP modifications necessary to represent concurrent actions and finally, present value-based RL algorithms that can cope with concurrent environments.

While concurrent environments affect DRL methods beyond value-based RL, we focus our scope on value-based methods due to their attractive sample-efficiency and off-policy properties for real-world robotic tasks.

### 3.1 DISCRETE-TIME REINFORCEMENT LEARNING PRELIMINARIES

We use standard reinforcement learning formulations in both discrete-time and continuous-time settings (Sutton & Barto, 1998). In the discrete-time case, at each time step  $i$ , the agent receives




Figure 1: In both diagrams, shaded nodes represent observable random variables and unshaded nodes represent unobserved random variables. **(a)**: in “blocking” MDPs, the environment state does not change while the agent records the current state and selects an action. This can be implemented by stopping time (in simulation), or ensuring that the entire system completely comes to rest before the state capture and action selection occur. **(b)**: in “concurrent” MDPs, we formulate state and action dynamics as continuous-time stochastic processes  $s(t)$  and  $a_i(t)$ . The specific intervals  $t_{AS''}, t_{AS}, t_{AS'}$  are unobservable random variables. At time  $t$ , the agent observes the state of the world  $s(t)$ , but by the time it selects an action  $a_i(t + t_{AS})$ , the last chosen action process  $a_{i-1}(t - H + t_{AS''})$  has “rolled over” to an unobserved state  $s(t + t_{AS})$ . An agent that concurrently selects actions from old states while in motion may need to interrupt a previous action before it has finished executing its current trajectory.

state  $s_i$  from a set of possible states  $\mathcal{S}$  and selects an action  $\mathbf{a}_i$  from some set of possible actions  $\mathcal{A}$  according to its policy  $\pi$ , where  $\pi$  is a mapping from  $\mathcal{S}$  to  $\mathcal{A}$ . The environment returns the next state  $s_{i+1}$  sampled from a transition distribution  $p(s_{i+1}|s_i, \mathbf{a}_i)$  and a reward  $r(s_i, \mathbf{a}_i)$ . The return for a given trajectory of states and actions is the total discounted return from time step  $i$  with discount factor  $\gamma \in (0, 1]$ :  $R_i = \sum_{k=0}^{\infty} \gamma^k r(s_{i+k}, \mathbf{a}_{i+k})$ . The goal of the agent is to maximize the expected return from each state  $s_i$ . The  $Q$ -function for a given policy  $\pi$  gives the expected return when selecting action  $\mathbf{a}$  at state  $s$ :  $Q(s, \mathbf{a}) = \mathbb{E}[R_i | s_i = s, \mathbf{a}_i = \mathbf{a}]$ . Similarly, the value function gives expected return from state  $s$ :  $V(s) = \mathbb{E}[R_i | s_i = s]$ .

The default blocking environment formulation is detailed in Figure 1a.

### 3.2 VALUE FUNCTIONS AND POLICIES IN CONTINUOUS TIME

For the continuous-time case, we start by formalizing a continuous-time MDP with the differential equation:

$$ds(t) = F(s(t), \mathbf{a}(t))dt + G(s(t), \mathbf{a}(t))d\beta \quad (1)$$

where  $\mathcal{S} = \mathbb{R}^d$  is a set of states,  $\mathcal{A}$  is a set of actions and  $F : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  describes the stochastic dynamics of the environment, and  $\beta$  is a Wiener process (Ross et al., 1996). Continuous-time functions  $s(t)$  and  $a_i(t)$  specify the state and  $i$ -th action taken by the agent. We refer to the values taken by these functions at time  $t$  as  $\mathbf{s}(t)$  and  $\mathbf{a}(t)$  respectively. Note that we use **bolded** font to indicate state/action function values ( $\mathbf{s}_i(t), \mathbf{a}(t)$ ) and *unbolded* font to indicate the function itself ( $s(t), a_i(t)$ ). The agent interacts with the environment through a state-dependent, deterministic policy function  $\pi$  and the return  $R$  of a trajectory  $\tau$  is given by (Doya, 2000):

$$R(\tau) = \int_{t=0}^{\infty} \gamma^t r(\mathbf{s}(t), \mathbf{a}(t))dt, \quad (2)$$

which leads to a continuous-time value function (Tallec et al., 2019):

$$\begin{aligned} V(\mathbf{s}(t)) &= \mathbb{E}_{\tau \sim \pi} [R(\tau) | \mathbf{s}(t)] \\ &= \mathbb{E}_{\tau \sim \pi} \left[ \int_{t=0}^{\infty} \gamma^t r(\mathbf{s}(t), \mathbf{a}(t))dt \right], \end{aligned} \quad (3)$$

and similarly, a continuous  $Q$ -function:

$$Q(\mathbf{s}(t), a, t, H) = \mathbb{E}_s \left[ \int_{t'=t}^{t'+H} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}(t'))dt' + \gamma^H V(\mathbf{s}(t+H)) \right], \quad (4)$$

where  $H$  is the constant sampling period between state captures (i.e. the duration of an action trajectory). The expectations are computed with respect to stochastic process defined in Eq. 1.

### 3.3 CONCURRENT ACTIONS IN CONTINUOUS AND DISCRETE-TIME

We consider Markov Decision Processes (MDPs) with concurrent actions, where actions are not executed to full completion. More specifically, by concurrent actions we mean that the state is captured while the previous action is still being executed. After state capture, the policy selects an action that is executed in the environment regardless of whether the previous action has completed.

In the continuous-time MDP case, concurrent actions can be considered as horizontally translating the action along the time dimension (Walsh et al., 2007), and the effect of concurrent actions is illustrated in Figure 1b.

### 3.4 VALUE-BASED CONCURRENT REINFORCEMENT LEARNING ALGORITHMS IN CONTINUOUS AND DISCRETE-TIME

We start our derivation from this continuous-time reinforcement learning standpoint, as it allows us to easily characterize the concurrent nature of the system. We then demonstrate that the conclusions drawn for the continuous case also apply to the more commonly-used discrete setting.

**Continuous Formulation** In order to further analyze the concurrent setting, we introduce the following notation. As shown in Figure 1b, an agent selects  $N$  action trajectories during an episode,  $a_1, \dots, a_N$ , where each  $a_i(t)$  is a continuous stochastic process generating controls as a function of time  $t$ . Let  $t_{AS}$  be the time duration of state capture, policy inference and any additional communication latencies. At time  $t$ , an agent begins computing the  $i$ -th trajectory  $a_i(t)$  from state  $\mathbf{s}(t)$ , while concurrently executing the previous selected trajectory  $a_{i-1}(t)$  over the time interval  $(t - H + t_{AS}, t + t_{AS})$ . At time  $t + t_{AS}$ , where  $t \leq t + t_{AS} \leq t + H$ , the agent switches to executing actions from  $a_i(t)$ . The continuous-time  $Q$ -function for the concurrent case from Eq. 4 can be expressed as following:

$$\begin{aligned}
 Q(\mathbf{s}(t), a, t, H) = & \underbrace{\mathbb{E}_s \left[ \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' \right]}_{\text{Executing action trajectory } a_{i-1}(t) \text{ until } t + t_{AS}} \\
 & + \underbrace{\mathbb{E}_s \left[ \int_{t'=t+t_{AS}}^{t'=t+H} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_i(t')) dt' \right]}_{\text{Executing action trajectory } a_i(t) \text{ until } t + H} + \underbrace{\mathbb{E}_s [\gamma^H V(\mathbf{s}(t + H))]}_{\text{Value function at } t + H} \quad (5)
 \end{aligned}$$

The first two terms correspond to expected discounted returns for executing the action trajectory  $a_{i-1}(t)$  from time  $(t, t + t_{AS})$  and the trajectory  $a_i(t)$  from time  $(t + t_{AS}, t + t_{AS} + H)$ . We can obtain a single-sample Monte Carlo estimator  $\hat{Q}$  by sampling random functions values  $s$ , which simply correspond to policy rollouts:

$$\begin{aligned}
 \hat{Q}(\mathbf{s}(t), a, t, H) = & \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \\
 & \gamma^{t_{AS}} \left[ \int_{t'=t+t_{AS}}^{t'=t+H} \gamma^{t'-t-t_{AS}} r(\mathbf{s}(t'), \mathbf{a}_i(t')) dt' + \gamma^{H-t_{AS}} V(\mathbf{s}(t + H)) \right] \quad (6)
 \end{aligned}$$

Next, for the continuous-time case, let us define a new concurrent Bellman backup operator:

$$\begin{aligned}
 \mathcal{T}_{nb}^\pi \hat{Q}(\mathbf{s}(t), a_{i-1}, t, t_{AS}) = & \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \\
 & \gamma^{t_{AS}} \mathbb{E}_s \hat{Q}(\mathbf{s}(t + t_{AS}), a_i, t + t_{AS}, H - t_{AS}). \quad (7)
 \end{aligned}$$

In addition to expanding the Bellman operator to take into account concurrent actions, we demonstrate that this modified operator maintain its contraction properties that are crucial for Q-learning convergence.

**Lemma 3.1.** *The concurrent continuous-time Bellman operator is a contraction.*

*Proof.* See Appendix A.2. □

**Discrete Formulation** In order to simplify the notation for the discrete-time case where the distinction between the action function  $a_i(t)$  and the value of that function at time step  $t$ ,  $\mathbf{a}_i(t)$ , is not necessary, we refer to the current state, current action, and previous action as  $\mathbf{s}_t$ ,  $\mathbf{a}_t$ ,  $\mathbf{a}_{t-1}$  respectively, replacing subindex  $i$  with  $t$ . Following this notation, we define the concurrent  $Q$ -function for the discrete-time case:

$$Q(\mathbf{s}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) = r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma \frac{t_{AS}}{H} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}} | \mathbf{s}_t, \mathbf{a}_{t-1})} Q(\mathbf{s}_{t+t_{AS}}, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}) \quad (8)$$

Where  $t_{AS}'$  is the ‘‘spillover duration’’ for action  $\mathbf{a}_t$  beginning execution at time  $t + t_{AS}$  (see Figure 1b). The concurrent Bellman operator is as follows:

$$\mathcal{T}_{nb}^\pi Q(\mathbf{s}_t, \mathbf{a}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) = r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma \frac{t_{AS}}{H} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}} | \mathbf{s}_t, \mathbf{a}_{t-1})} Q(\mathbf{s}_{t+t_{AS}}, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}). \quad (9)$$

Similarly to the continuous-time case, we demonstrate that this Bellman operator is a contraction.

**Lemma 3.2.** *The concurrent discrete-time Bellman operator is a contraction.*

*Proof.* See Appendix A.2. □

We refer the reader to Appendix A.1 for more detailed derivations of the  $Q$ -functions and Bellman operators. Crucially, Equation 9 implies that we can extend a conventional discrete-time  $Q$ -learning framework to handle MDPs with concurrent actions by providing the  $Q$  function with values of  $t_{AS}$  and  $\mathbf{a}_{t-1}$ , in addition to the standard inputs  $\mathbf{s}_t$ ,  $\mathbf{a}_t$ ,  $t$ .

### 3.5 DEEP $Q$ -LEARNING WITH CONCURRENT KNOWLEDGE

While we have shown that knowledge of the concurrent system properties ( $t_{AS}$  and  $\mathbf{a}_{t-1}$ , as defined previously for the discrete-time case) is theoretically sufficient, it is often hard to accurately predict  $t_{AS}$  during inference on a complex robotics system. In order to allow practical implementation of our algorithm on a wide range of RL agents, we consider three additional features encapsulating *concurrent knowledge* used to condition the  $Q$ -function: (1) Previous action ( $\mathbf{a}_{t-1}$ ), (2) Action selection time ( $t_{AS}$ ), and (3) Vector-to-go ( $VTG$ ), which we define as the remaining action to be executed at the instant the state is measured.

We note that  $\mathbf{a}_{t-1}$  is available across the vast majority of environments and it is easy to obtain. Using  $t_{AS}$ , which encompasses state capture, communication latency, and policy inference, relies on having some knowledge of the concurrent properties of the system. Calculating  $VTG$  requires having access to some measure of action completion at the exact moment when state is observed. When utilizing a first-order control action space, such as joint angle or desired pose,  $VTG$  is easily computable if proprioceptive state is measured and synchronized with state observation. In these cases,  $VTG$  is an alternate representation of the same information encapsulated by  $\mathbf{a}_{t-1}$  and the current state. We limit our analysis to environments where  $\mathbf{a}_{t-1}$ ,  $t_{AS}$ , and  $VTG$  are all obtainable and  $H$  is held constant.

## 4 EXPERIMENTS

In our experimental evaluation we aim to study the following questions: (1) Is concurrent knowledge defined in Section 3.5, both necessary and sufficient for a  $Q$ -function to recover the performance of a blocking unconditioned  $Q$ -function, when acting in a concurrent environment? (2) Which

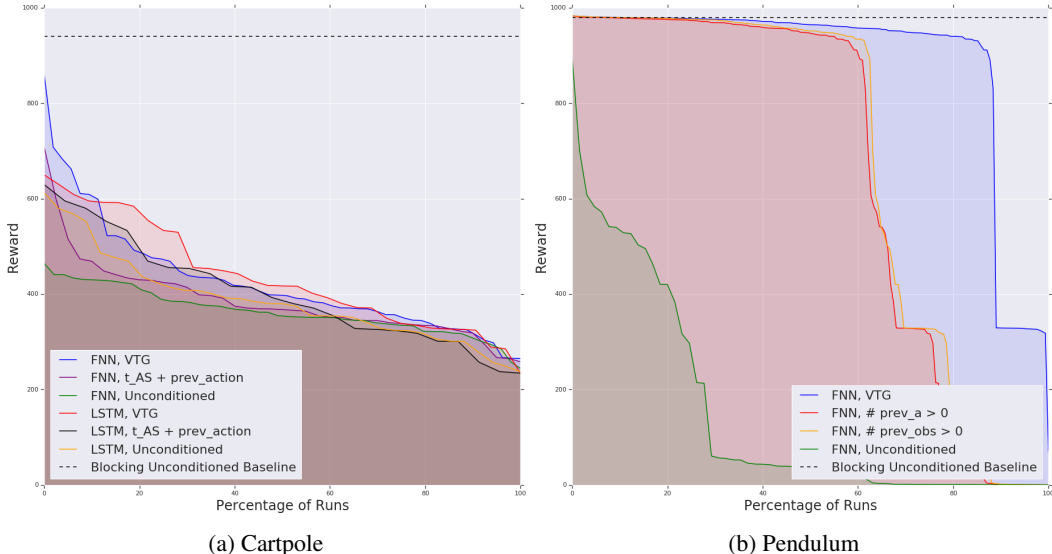


Figure 2: In concurrent versions of both Cartpole and Pendulum, we observe that providing the critic with VTG leads to more robust performance across all hyperparameters. **(a)** Environment rewards achieved by DQN with different network architectures [either a feedforward network (FNN) or a Long Short-Term Memory (LSTM) network] and different concurrent knowledge features [Unconditioned, Vector-to-go (VTG), or previous action and  $t_{AS}$ ] on the concurrent Cartpole task for every hyperparameter in a sweep, sorted in decreasing order. **(b)** Environment rewards achieved by DQN with a FNN and different frame-stacking and concurrent knowledge parameters on the concurrent Pendulum task for every hyperparameter in a sweep, sorted in decreasing order. Larger versions of these figures are provided in Appendix A.4.

representations of concurrent knowledge are most useful for a  $Q$ -function to act in a concurrent environment? (3) Can concurrent models improve smoothness and execution speed of a real-robot policy in a realistic, vision-based manipulation task?

#### 4.1 TOY FIRST-ORDER CONTROL PROBLEMS

First, we illustrate the effects of a concurrent control paradigm on value-based DRL methods through an ablation study on concurrent versions of the standard Cartpole and Pendulum environments. We use 3D MuJoCo based implementations in DeepMind Control Suite (Tassa et al., 2018) for both tasks. For the baseline learning algorithm implementations, we use the TF-Agents (Guadarrama et al., 2018) implementations of a Deep  $Q$ -Network agent, which utilizes a Feed-forward Neural Network (FNN), and a Deep  $Q$ -Recurrent Neutral Network agent, which utilizes a Long Short-Term Memory (LSTM) network. To approximate different difficulty levels of latency in concurrent environments, we utilize different parameter combinations for action execution steps and action selection steps ( $t_{AS}$ ). The number of action execution steps is selected from  $\{0\text{ms}, 5\text{ms}, 25\text{ms}, \text{or } 50\text{ms}\}$  once at environment initialization.  $t_{AS}$  is selected from  $\{0\text{ms}, 5\text{ms}, 10\text{ms}, 25\text{ms}, \text{or } 50\text{ms}\}$  either once at environment initialization or repeatedly at every episode reset. In addition to environment parameters, we allow trials to vary across model parameters: number of previous actions to store, number of previous states to store, whether to use VTG, whether to use  $t_{AS}$ ,  $Q$ -network architecture, and number of discretized actions. Further details are described in Appendix A.3.1.

To estimate the relative importance of different concurrent knowledge representations, we conduct an analysis of the sensitivity of each type of concurrent knowledge representations to combinations of the other hyperparameter values, shown in Figure 2a. While all combinations of concurrent knowledge representations increase learning performance over baselines that do not leverage this information, the clearest difference stems from including VTG. In Figure 2b we conduct a similar analysis but on a Pendulum environment where  $t_{AS}$  is fixed every environment; thus, we do not focus on  $t_{AS}$  for this analysis but instead compare the importance of VTG with frame-stacking previous




Figure 3: An overview of the robotic grasping task. A static manipulator arm attempts to grasp objects placed in bins front of it. In simulation, the objects are procedurally generated.

actions and observations. While frame-stacking helps nominally, the majority of the performance increase results from utilizing information from VTG.

#### 4.2 CONCURRENT QT-OPT ON LARGE-SCALE ROBOTIC GRASPING

Next, we evaluate scalability of our approach to a practical robotic grasping task. We simulate a 7 DoF arm with an over-the-shoulder camera, where a bin in front of the robot is filled with procedurally generated objects to be picked up by the robot. A binary reward is assigned if an object is lifted off a bin at the end of an episode. We train a policy with QT-Opt (Kalashnikov et al., 2018), a deep  $Q$ -Learning method that utilizes the cross-entropy method (CEM) to support continuous actions. In the blocking mode, a displacement action is executed until completion: the robot uses a closed-loop controller to fully execute an action, decelerating and coming to rest before observing the next state. In the concurrent mode, an action is triggered and executed without waiting, which means that the next state is observed while the robot remains in motion. Further details of the setup are shown in Figure 3 and explained in Appendix A.3.2.

Table 1 summarizes the performance for blocking and concurrent modes comparing unconditioned models against the concurrent knowledge models described in Section 3.5. Our results indicate that the VTG model acting in concurrent mode is able to recover baseline task performance of the blocking execution unconditioned baseline, while the unconditioned baseline acting in concurrent model suffers some performance loss. In addition to the success rate of the grasping policy, we also evaluate the speed and smoothness of the learned policy behavior. Concurrent knowledge models are able to learn faster trajectories: *episode duration*, which measures the total amount of wall-time used for an episode, is reduced by 31.3% when comparing concurrent knowledge models with blocking unconditioned models. When switching from blocking execution mode to concurrent execution mode, we see a significantly lower *action completion*, measured as the ratio from executed gripper displacement to commanded displacement, which expectedly indicates a switch to a concurrent environment. The concurrent knowledge models have higher action completions than the unconditioned model in the concurrent environment, which suggests that the concurrent knowledge models are able to utilize more efficient motions, resulting in smoother trajectories. The qualitative benefits of faster, smoother trajectories are drastically apparent when viewing video playback of learned policies<sup>1</sup>.

**Real robot results** In addition, we evaluate qualitative policy behaviors of concurrent models compared to blocking models on a real-world robot grasping task, which is shown in Figure 3b. As

<sup>1</sup><https://youtu.be/Gr2sZVwrX5w>

Table 1: Large-Scale Simulated Robotic Grasping Results

Blocking Actions	VTG	Previous Action	Grasp Success	Episode Duration	Action Completion
Yes	No	No	91.53% $\pm$ 1.04%	120.81s $\pm$ 9.13s	<b>89.53%</b> $\pm$ 2.267%
No	No	No	83.77% $\pm$ 9.27%	97.16s $\pm$ 6.28s	34.69% $\pm$ 16.80%
No	Yes	No	92.55% $\pm$ 4.39%	<b>82.98s</b> $\pm$ <b>5.74s</b>	47.28% $\pm$ 14.25%
No	No	Yes	92.70% $\pm$ 1.42%	87.15s $\pm$ 4.80s	50.09% $\pm$ 14.25%
No	Yes	Yes	<b>93.49%</b> $\pm$ <b>1.04%</b>	90.75s $\pm$ 4.15s	49.19% $\pm$ 14.98%

Table 2: Real-World Robotic Grasping Results.

Blocking Actions	VTG	Grasp Success	Policy Duration
Yes	No	<b>81.43%</b>	22.60s $\pm$ 12.99s
No	Yes	68.60%	<b>11.52s</b> $\pm$ <b>7.272s</b>

seen in Table 2, the models achieve comparable grasp success, but the concurrent model is 49% faster than the blocking model in terms of *policy duration*, which measures the total execution time of the policy (this excludes the infrastructure setup and teardown times accounted for in episode duration, which can not be optimized with concurrent actions). In addition, the concurrent VTG model is able to execute smoother and faster trajectories than the blocking unconditioned baseline, which is clear in video playback<sup>1</sup>.

## 5 DISCUSSION AND FUTURE WORK

We presented a theoretical framework to analyze concurrent systems where the robot must “think while moving”. Viewing this formulation through the lens of continuous-time value-based reinforcement learning, we showed that by considering concurrent knowledge about the time delay  $t_{AS}$  and the previous action, the concurrent continuous-time and discrete-time Bellman operators remained contractions and thus maintained  $Q$ -Learning convergence guarantees. While more information than  $t_{AS}$  and previous action may be helpful, we showed that  $t_{AS}$  and previous action (and different representations of this information) are the sole theoretical requirements for good learning performance. In addition, we introduced Vector-to-go (VTG), which incorporates the remaining previous action to be executed, as an alternative representation for information about the concurrent system that previous action and  $t_{AS}$  contain.

Our theoretical findings were supported by experimental results on  $Q$ -learning models acting in simulated control tasks that were engineered to support concurrent action execution. We conducted large-scale ablation studies on toy task concurrent 3D Cartpole and Pendulum environments, across model parameters as well as concurrent environment parameters. Our results indicated that VTG is the least hyperparameter-sensitive representation, and was able to recover blocking learning performance in concurrent settings. We extended these results to a complex concurrent large-scale simulated robotic grasping task, where we showed that the concurrent models were able to recover blocking execution baseline model success while acting 31.3% faster. We analyzed the qualitative benefits of concurrent models through a real-world robotic grasping task, where we showed that a concurrent model with comparable grasp success as a blocking baseline was able to learn smoother trajectories that were 49% faster.

An interesting topic to explore in future work is the possibility of increased data efficiency when training on off-policy data from various latency regimes. Another natural extension of this work is to evaluate DRL methods beyond value-based algorithms, such as on-policy learning and policy gradient approaches. Finally, concurrent methods may allow robotic control in dynamic environments where it is not possible for the robot to stop the environment before computing the action. In these scenarios, robots must truly think and act at the same time.



## REFERENCES

- Pieter Abbeel, Adam Coates, Morgan Quigley, and Andrew Y. Ng. An application of reinforcement learning to aerobatic helicopter flight. In Bernhard Schlkopf, John C. Platt, and Thomas Hofmann (eds.), *NIPS*, pp. 1–8. MIT Press, 2006. ISBN 0-262-19568-2. URL <http://dblp.uni-trier.de/db/conf/nips/nips2006.html#AbbeelCQN06>.
- Artemij Amiranashvili, Alexey Dosovitskiy, Vladlen Koltun, and Thomas Brox. Motion perception in reinforcement learning with dynamic objects. In *CoRL*, volume 87 of *Proceedings of Machine Learning Research*, pp. 156–168. PMLR, 2018. URL <http://dblp.uni-trier.de/db/conf/corl/corl2018.html#AmiranashviliDK18>.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016. URL <http://arxiv.org/abs/1606.01540>. cite arxiv:1606.01540.
- Rémi Coulom. *Reinforcement learning using neural networks, with applications to motor control*. PhD thesis, Institut National Polytechnique de Grenoble-INPG, 2002.
- Nicols Cruz, Kenzo Lobos-Tsunekawa, and Javier Ruiz del Solar. Using convolutional neural networks in robots with limited computational resources: Detecting nao robots while playing soccer. *CoRR*, abs/1706.06702, 2017. URL <http://dblp.uni-trier.de/db/journals/corr/corr1706.html#CruzLR17>.
- Kenji Doya. Reinforcement learning in continuous time and space. *Neural Computation*, 12(1): 219–245, 2000. URL <http://dblp.uni-trier.de/db/journals/neco/neco12.html#Doya00>.
- Vlad Firoiu, Tina Ju, and Joshua Tenenbaum. At Human Speed: Deep Reinforcement Learning with Action Delay. *arXiv e-prints*, October 2018.
- Sergio Guadarrama, Anoop Korattikara, Oscar Ramirez, Pablo Castro, Ethan Holly, Sam Fishman, Ke Wang, Ekaterina Gonina, Chris Harris, Vincent Vanhoucke, et al. Tf-agents: A library for reinforcement learning in tensorflow, 2018.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018. URL <http://dblp.uni-trier.de/db/journals/corr/corr1812.html#abs-1812-05905>.
- Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *CoRR*, abs/1806.10293, 2018. URL <http://dblp.uni-trier.de/db/journals/corr/corr1806.html#abs-1806-10293>.
- H J Kappen. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11):P11011–P11011, nov 2005. doi: 10.1088/1742-5468/2005/11/p11011. URL <https://doi.org/10.1088%2F1742-5468%2F2005%2F11%2Fp11011>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. 2013. URL <http://arxiv.org/abs/1312.5602>. cite arxiv:1312.5602Comment: NIPS Deep Learning Workshop 2013.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, February 2015. ISSN 00280836. URL <http://dx.doi.org/10.1038/nature14236>.

- Rémi Munos and Paul Bourgin. Reinforcement learning for continuous stochastic control problems. In M. I. Jordan, M. J. Kearns, and S. A. Solla (eds.), *Advances in Neural Information Processing Systems 10*, pp. 1029–1035. MIT Press, 1998. URL <http://papers.nips.cc/paper/1404-reinforcement-learning-for-continuous-stochastic-control-problems.pdf>.
- OpenAI, Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jzefowicz, Bob McGrew, Jakub W. Pachocki, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *CoRR*, abs/1808.00177, 2018. URL <http://dblp.uni-trier.de/db/journals/corr/corr1808.html#abs-1808-00177>.
- Sheldon M Ross, John J Kelly, Roger J Sullivan, William James Perry, Donald Mercer, Ruth M Davis, Thomas Dell Washburn, Earl V Sager, Joseph B Boyce, and Vincent L Bristow. *Stochastic processes*, volume 2. Wiley New York, 1996.
- Erik Schuitema, Lucian Busoniu, Robert Babuka, and Pieter P. Jonker. Control delay in reinforcement learning for real-time dynamic systems: A memoryless approach. *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3226–3231, 2010.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–, January 2016. URL <http://dx.doi.org/10.1038/nature16961>.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, March 1998. ISBN 0262193981. URL <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0262193981>.
- Correntin Tallec, Leonard Blier, and Yann Ollivier. Making Deep Q-learning Methods Robust to Time Discretization. *arXiv e-prints*, January 2019.
- Jie Tan, Tingnan Zhang, Erwin Coumans, Atıl İscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-Real: Learning Agile Locomotion For Quadruped Robots. *arXiv e-prints*, April 2018.
- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A. Riedmiller. Deepmind control suite. *CoRR*, abs/1801.00690, 2018. URL <http://dblp.uni-trier.de/db/journals/corr/corr1801.html#abs-1801-00690>.
- Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. Reinforcement learning of motor skills in high dimensions: A path integral approach. pp. 2397 – 2403, 06 2010. doi: 10.1109/ROBOT.2010.5509336.
- Thomas J. Walsh, Ali Nouri, Lihong Li, and Michael L. Littman. Planning and learning in environments with delayed feedback. In Joost N. Kok, Jacek Koronacki, Ramn Lpez de Mntaras, Stan Matwin, Dunja Mladenic, and Andrzej Skowron (eds.), *ECML*, volume 4701 of *Lecture Notes in Computer Science*, pp. 442–453. Springer, 2007. ISBN 978-3-540-74957-8. URL <http://dblp.uni-trier.de/db/conf/ecml/ecml2007.html#WalshNLL07>.

## A APPENDIX

### A.1 DEFINING BLOCKING BELLMAN OPERATORS

As introduced in Section 3.4, we define a continuous-time  $Q$ -function estimator with concurrent actions.

$$\hat{Q}(\mathbf{s}(t), a, t, H) = \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \quad (10)$$

$$\int_{t''=t+t_{AS}}^{t''=t+H} \gamma^{t''-t} r(\mathbf{s}(t''), \mathbf{a}_i(t'')) dt'' + \gamma^H V(\mathbf{s}(t+H)) \quad (11)$$

$$= \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \quad (12)$$

$$\gamma^{t_{AS}} \int_{t''=t+t_{AS}}^{t''=t+H} \gamma^{t''-t-t_{AS}} r(\mathbf{s}(t''), \mathbf{a}_i(t'')) dt'' + \gamma^H V(\mathbf{s}(t+H)) \quad (13)$$

$$= \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \quad (14)$$

$$\gamma^{t_{AS}} \left[ \int_{t''=t+t_{AS}}^{t''=t+H} \gamma^{t''-t-t_{AS}} r(\mathbf{s}(t''), \mathbf{a}_i(t'')) dt'' + \gamma^{H-t_{AS}} V(\mathbf{s}(t+H)) \right] \quad (15)$$

We observe that the second part of this equation (after  $\gamma^{t_{AS}}$ ) is itself a  $Q$ -function at time  $t+t_{AS}$ . Since the future state, action, and reward values at  $t+t_{AS}$  are not known at time  $t$ , we take the following expectation:

$$Q(\mathbf{s}(t), a, t, H) = \int_{t'=t}^{t'=t+t_{AS}} \gamma^{t'-t} r(\mathbf{s}(t'), \mathbf{a}_{i-1}(t')) dt' + \quad (16)$$

$$\gamma^{t_{AS}} \mathbb{E}_s \hat{Q}(\mathbf{s}(t), a, t+t_{AS}, H-t_{AS}) \quad (17)$$

which indicates that the  $Q$ -function in this setting is not just the expected sum of discounted future rewards, but it corresponds to an expected future  $Q$ -function.

In order to show the discrete-time version of the problem, we parameterize the discrete-time concurrent  $Q$ -function as:

$$\hat{Q}(\mathbf{s}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) = r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma^{\frac{t_{AS}}{H}} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}} | \mathbf{s}_t, \mathbf{a}_{t-1})} r(\mathbf{s}_{t+t_{AS}}, \mathbf{a}_t) + \quad (18)$$

$$\gamma^{\frac{H}{H}} \mathbb{E}_{p(\mathbf{s}_{t+H} | \mathbf{s}_{t+t_{AS}}, \mathbf{a}_t)} V(\mathbf{s}_{t+H}) \quad (19)$$

which with  $t_{AS} = 0$ , corresponds to a synchronous environment.

Using this parameterization, we can rewrite the discrete-time  $Q$ -function with concurrent actions as:

$$\hat{Q}(\mathbf{s}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) = r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma^{\frac{t_{AS}}{H}} [\mathbb{E}_{p(\mathbf{s}_{t+t_{AS}} | \mathbf{s}_t, \mathbf{a}_{t-1})} r(\mathbf{s}_{t+t_{AS}}, \mathbf{a}_t) + \quad (20)$$

$$\gamma^{\frac{H-t_{AS}}{H}} \mathbb{E}_{p(\mathbf{s}_{t+H} | \mathbf{s}_{t+t_{AS}}, \mathbf{a}_t)} V(\mathbf{s}_{t+H})] \quad (21)$$

$$= r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma^{\frac{t_{AS}}{H}} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}} | \mathbf{s}_t, \mathbf{a}_{t-1})} \hat{Q}(\mathbf{s}_t, \mathbf{a}_t, t+t_{AS}, t_{AS}, H-t_{AS}) \quad (22)$$

### A.2 CONTRACTION PROOFS FOR THE BLOCKING BELLMAN OPERATORS

#### Proof of the Discrete-time Blocking Bellman Update

**Lemma A.1.** *The traditional Bellman operator is a contraction, i.e.:*

$$\|T^\pi Q_\infty(\mathbf{s}, \mathbf{a}) - T^\pi Q_\infty(\mathbf{s}, \mathbf{a}')\| \leq c \|Q_1(\mathbf{s}, \mathbf{a}) - Q_2(\mathbf{s}, \mathbf{a}')\|, \quad (23)$$

where  $T^\pi Q(\mathbf{s}, \mathbf{a}) = r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_\pi Q(\mathbf{s}', \mathbf{a}')$  and  $0 \leq c \leq 1$ .

*Proof.* In the original formulation, we can show that this is the case as following:

$$\mathcal{T}^\pi Q_1(\mathbf{s}, \mathbf{a}) - \mathcal{T}^\pi Q_2(\mathbf{s}, \mathbf{a}) \quad (24)$$

$$= r(\mathbf{s}, \mathbf{a}) + \gamma \mathbb{E}_\pi [Q_1(\mathbf{s}', \mathbf{a}') - r(\mathbf{s}, \mathbf{a}) - \gamma \mathbb{E}_\pi Q_2(\mathbf{s}', \mathbf{a}')] \quad (25)$$

$$= \gamma \mathbb{E}_\pi [Q_1(\mathbf{s}', \mathbf{a}') - Q_2(\mathbf{s}', \mathbf{a}')] \quad (26)$$

$$\leq \gamma \sup_{\mathbf{s}', \mathbf{a}'} [Q_1(\mathbf{s}', \mathbf{a}') - Q_2(\mathbf{s}', \mathbf{a}')], \quad (27)$$

with  $0 \leq \gamma \leq 1$  and  $\|f\|_\infty = \sup_x [f(x)]$ .  $\square$

Similarly, we can show that the updated Bellman operators introduced in Section 3.4 are contractions as well.

### Proof of Lemma 3.2

*Proof.*

$$\mathcal{T}_{nb}^\pi Q_1(\mathbf{s}_t, \mathbf{a}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) - \mathcal{T}_{nb}^\pi Q_2(\mathbf{s}_t, \mathbf{a}_t, \mathbf{a}_{t-1}, t, t_{AS}, H) \quad (28)$$

$$= r(\mathbf{s}_t, \mathbf{a}_{t-1}) + \gamma \frac{t_{AS}}{H} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}}|\mathbf{s}_t, \mathbf{a}_{t-1})} Q_1(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}) \quad (29)$$

$$- r(\mathbf{s}_t, \mathbf{a}_{t-1}) - \gamma \frac{t_{AS}}{H} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}}|\mathbf{s}_t, \mathbf{a}_{t-1})} Q_2(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}) \quad (30)$$

$$= \gamma \frac{t_{AS}}{H} \mathbb{E}_{p(\mathbf{s}_{t+t_{AS}}|\mathbf{s}_t, \mathbf{a}_{t-1})} [Q_1(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}) - Q_2(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS})] \quad (31)$$

$$\leq \gamma \frac{t_{AS}}{H} \sup_{\mathbf{s}_t, \mathbf{a}_t, t+t_{AS}, t_{AS}', H-t_{AS}} [Q_1(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS}) - Q_2(\mathbf{s}_t, \mathbf{a}_t, t + t_{AS}, t_{AS}', H - t_{AS})] \quad (32)$$

$\square$

### Proof of Lemma 3.1

*Proof.* To prove that this the continuous-time Bellman operator is a contraction, we can follow the discrete-time proof, from which it follows:

$$\mathcal{T}_{nb}^\pi Q_1(\mathbf{s}(t), a_{i-1}, t, t_{AS}) - \mathcal{T}_{nb}^\pi Q_2(\mathbf{s}(t), a_{i-1}, t, t_{AS}) \quad (33)$$

$$= \gamma^{t_{AS}} \mathbb{E}_s [Q_1(\mathbf{s}(t), a_i, t + t_{AS}, H - t_{AS}) - Q_2(\mathbf{s}(t), a_i, t + t_{AS}, H - t_{AS})] \quad (34)$$

$$\leq \gamma^{t_{AS}} \sup_{\mathbf{s}(t), a_i, t+t_{AS}, H-t_{AS}} [Q_1(\mathbf{s}(t), a_i, t + t_{AS}, H - t_{AS}) - Q_2(\mathbf{s}(t), a_i, t + t_{AS}, H - t_{AS})] \quad (35)$$

$\square$

## A.3 EXPERIMENT IMPLEMENTATION DETAILS

### A.3.1 CARTPOLE AND PENDULUM ABLATION STUDIES

Here, we describe the implementation details of the toy task Cartpole and Pendulum experiments in Section 4.1.

For the environments, we use the 3D MuJoCo implementations of the Cartpole-Swingup and Pendulum-Swingup tasks in DeepMind Control Suite (Tassa et al., 2018). We use discretized action spaces for first-order control of joint position actuators. For the observation space of both tasks, we use the default state space of ground truth positions and velocities.

For the baseline learning algorithms, we use the TensorFlow Agents (Guadarrama et al., 2018) implementations of a Deep  $Q$ -Network agent, which utilizes a Feed-forward Neural Network (FNN), and a Deep  $Q$ -Recurrent Neural Network agent, which utilizes a Long Short-Term Memory (LSTM) network. Learning parameters such as `learning_rate`, `lstm_size`, and `fc_layer_size` were selected through hyperparameter sweeps.

To approximate different difficulty levels of latency in concurrent environments, we utilize different parameter combinations for action execution steps and action selection steps ( $t_{AS}$ ). The number of action execution steps is selected from  $\{0\text{ms}, 5\text{ms}, 25\text{ms}, \text{or } 50\text{ms}\}$  once at environment initialization.  $t_{AS}$  is selected from  $\{0\text{ms}, 5\text{ms}, 10\text{ms}, 25\text{ms}, \text{or } 50\text{ms}\}$  either once at environment initialization or repeatedly at every episode reset. The selected  $t_{AS}$  is implemented in the environment as additional physics steps that update the system during simulated action selection.

Frame-stacking parameters affect the observation space by saving previous observations and actions. The number of previous actions to store as well as the number of previous observations to store are independently selected from the range  $[0, 4]$ . Concurrent knowledge parameters, as described in Section 4, include whether to use VTG and whether to use  $t_{AS}$ . Including the previous action is already a feature implemented in the frame-stacking feature of including previous actions. Finally, the number of actions to discretize the continuous space to is selected from the range  $[3, 8]$ .

### A.3.2 LARGE SCALE SIMULATED ROBOTIC GRASPING

We simulate a 7 DoF arm with an over-the-shoulder camera (see Figure 3). A bin in front of the robot is filled with procedurally generated objects to be picked up by the robot. A binary reward is assigned if an object is lifted off a bin at the end of an episode. We train a policy with QT-Opt (Kalashnikov et al., 2018), a Deep  $Q$ -Learning method that utilizes the Cross-Entropy Method (CEM) to support continuous actions. States are represented in form of RGB images and actions are continuous Cartesian displacements of the gripper 3D positions and yaw. In addition, the policy commands discrete gripper open and close actions and may terminate an episode.

In blocking mode, a displacement action is executed until completion: the robot uses a closed loop controller to fully execute an action, decelerating and coming to rest before observing the next state. In concurrent mode, an action is triggered and executed without waiting, which means that the next state is observed while the robot remains in motion. It should be noted that in blocking mode, action completion is close to 100% unless the gripper moves are blocked by contact with the environment or objects; this causes average blocking mode action completion to be lower than 100%, as seen in Table 1.

### A.4 FIGURES

See Figure 4 and Figure 5 on the next page.

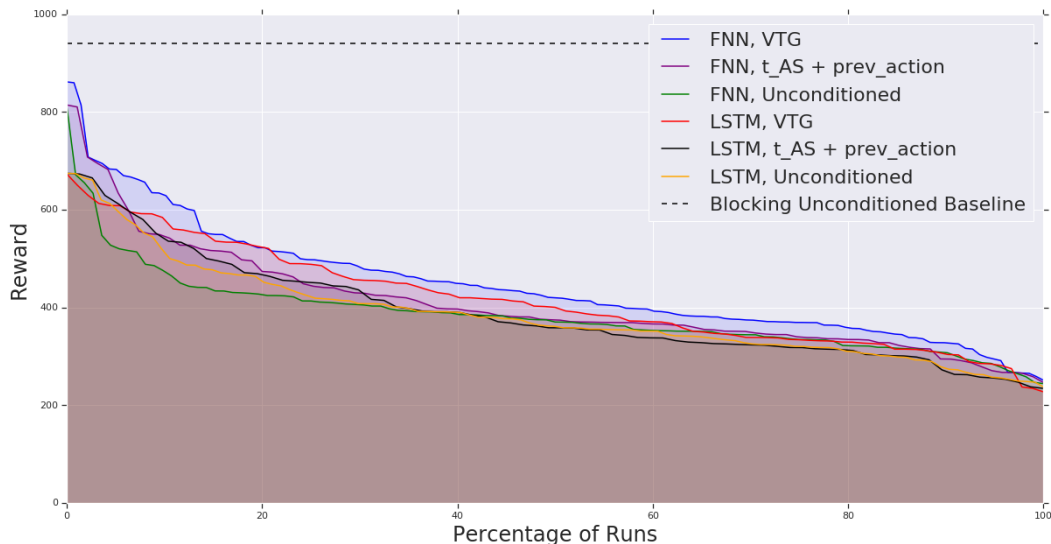


Figure 4: Environment rewards achieved by DQN with different network architectures [either a feedforward network (FNN) or a Long Short-Term Memory (LSTM) network] and different concurrent knowledge features [Unconditioned, vector-to-go (VTG), or previous action and  $t_{AS}$ ] on the concurrent Cartpole task for every hyperparameter in a sweep, sorted in decreasing order. Providing the critic with VTG information leads to more robust performance across all hyperparameters. This figure is a larger version of 2a.




Figure 5: Environment rewards achieved by DQN with a FNN and different frame-stacking and concurrent knowledge parameters on the concurrent Pendulum task for every hyperparameter in a sweep, sorted in decreasing order.