

The Intentional Unintentional Agent: Learning to Solve Many Continuous Control Tasks Simultaneously

Serkan Cabi Sergio Gómez Colmenarejo Matthew W. Hoffman

Misha Denil Ziyu Wang Nando de Freitas

DeepMind

{cabi,sergomez,mwhoffman,mdenil,ziyu,nandodefritis}@google.com

Abstract: This paper introduces the Intentional Unintentional (IU) agent. This agent endows the deep deterministic policy gradients (DDPG) agent for continuous control with the ability to solve several tasks simultaneously. Learning to solve many tasks simultaneously has been a long-standing, core goal of artificial intelligence, inspired by infant development and motivated by the desire to build flexible robot manipulators capable of many diverse behaviours. We show that the IU agent not only learns to solve many tasks simultaneously but it also learns faster than agents that target a single task at-a-time. In some cases, where the single task DDPG method completely fails, the IU agent successfully solves the task. To demonstrate this, we build a playroom environment using the MuJoCo physics engine, and introduce a grounded formal language to automatically generate tasks.

Keywords: Deep deterministic policy gradients, control, multi-task, physics

1 Introduction

Imagine a toddler in a playroom trying to bring two blocks together. While purposely focusing on this task, the infant is accomplishing many other goals incidentally both simpler and more complex: gazing, extending the arms, sitting, bending sideways, reaching, grasping, navigating around obstacles, dodging a looming object thrown by a sibling, sensing texture, sensing temperature, and so on. Over the first year, infants also display a wide range of spontaneous movements: kicks, stomps, sways, flaps, flails, rocks, rubs, nods, shakes, bounces, bangs, waves, wiggles and so on [1, 2].

We hypothesize that a single stream of experience offers agents the opportunity to learn and perfect many policies both on purpose and incidentally, thus accelerating the acquisition of grounded knowledge.

To investigate this hypothesis, we propose a deep actor-critic architecture, trained with deterministic policy gradients [3, 4], for learning several policies concurrently. The architecture enables the agent to attend to one task on-policy, while unintentionally learning to solve many other tasks off-policy. Importantly, the policies learned unintentionally can be harnessed for intentional use even if those policies were never followed before.

More precisely, this intentional-unintentional architecture, shown in Figure 1, consists of two neural networks. The actor neural network has multiple-heads representing different policies with shared lower-level representations. The critic network represents several state-action value functions, sharing a common representation for the observations.

The architecture alone does not suffice for investigating our hypothesis. We also need a flexible way of generating many diverse tasks and a suitable environment. To this end, we introduce an automatic procedure to generate semantic goals for the agent. We also introduce a physical environment, with gravity, a ground, rigid objects, and a simple embodied agent. While some characteristics of the physical world change with experiences, the laws of physics and the body remain fixed to allow for transfer and continual learning.

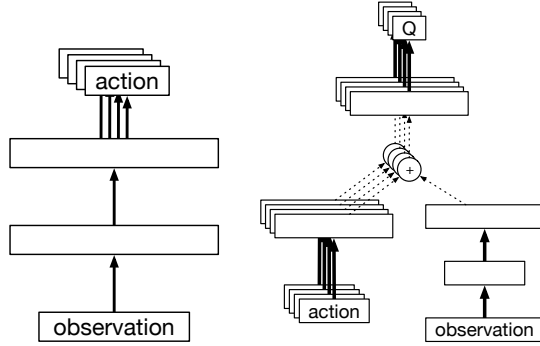


Figure 1: The IU architecture. The actor network on the left consists of two shared MLP-tanh layers followed by non-shared MLP-tanh layers to produce the multivariate actions for each policy (4 policies in this diagram). The right hand side shows the critic network. The action vectors provided by the policies are fed into a non-shared MLP-tanh layer, which is then point-wise added to the output of a two layer MLP-tanh network applied to the observation. The resulting activations are processed by non-shared linear layers to produce the Q values.

While being specific, this study aims to address generality in artificial intelligence by designing agents capable of doing many things, to overcome the problem of sparse rewards associated with conventional reinforcement learning by generating and controlling a stream of reward functions. We design embodied agents situated in a physical environment resembling a playroom, not only as a consequence of being inspired by infant development, but also as a result of being interested in eventual transfer to the world of flexible manufacturing with robots.

1.1 Related work

Thinking about an agent as immersed in a stream of multivariate rewards provides us with a powerful alternative to the conventional univariate reward reinforcement learning framework.

Recently, Jaderberg et al. [5] consider Asynchronous Advantage Actor Critic (A3C) agents [6] immersed in a sensorimotor stream. The agents are designed to achieve an extrinsic scalar reward, but are endowed with *auxiliary control tasks* and *auxiliary predictive tasks*. The auxiliary control tasks (pixel changes and simple network features) are shown to enable the A3C agent to learn to achieve the scalar reward faster in domains where the action-space is discrete. This paper will show more dramatic gains for continuous action spaces.

Sutton et al. [7] introduced the horde architecture to learn grounded knowledge from an unsupervised sensorimotor stream. While not focusing on the issue of representation, with the value functions being trained separately with different weights, the Horde provides much inspiration for this paper.

The concept of general value functions introduced in the horde architecture was further explored by Schaul et al. [8], and has connections with research on options, successor representations and hierarchical RL [9, 10, 11]. At the time of writing this paper, van Seijen et al. [12] introduced a linear decomposition of reward functions, related to [13], whereby several action-value functions are learned separately with Deep-Q-Networks (DQN) [14].

Multivariate reward feedback appears in sequential multi-objective decision making [15], in predictive decision making for 3D games [16], and in agents that use auxiliary predictive tasks to reduce sample complexity [17, 18]. Learning multiple control problems using various policy gradient methods has also been studied elsewhere [19, 20].

Our work is related to learning neural networks with a static or adaptive curricula [21, 22, 23], and to learning curricula for training neural networks using bandit techniques and Bayesian optimization [24, 25].

In developmental psychology, there are many studies on incidental activity and its consequences on motor development, including fetal and neonatal movement [26], twitching during sleep [27], stereotypies and flails in infants, and gross motor play [2].

Findings in developmental psychology have inspired the design of robotic systems in the field of developmental robotics [28]. It has been also been argued that robots provide a platform for examining many of these findings. We believe that rich physical simulators are a viable alternative, provided that we keep advancing environments and task generation mechanisms as done in this paper. The importance of bodies and physical environments in the study of artificial intelligence has been championed by many, notably by Brooks [29].

The automatic construction of grounded reward functions with formal languages has become a topic of great interest in recent months [30, 31, 32, 33].

2 The Intentional Unintentional Agent

Policy gradient algorithms form a very popular, if not the most popular, class of continuous action reinforcement learning algorithms. The fundamental basis of many of these algorithms is the *policy gradient theorem* [34]. However, this approach necessitates the use of stochastic policies which can complicate the process of learning off-policy. More recently, a deterministic policy gradient theorem has been formulated by [3] which removes this need. This approach was later extended in [4] to modern deep neural network actor-critic architectures, with *scalar rewards*.

In our setting, the agent perceives a *stream of rewards* r_t^i , indexed by i at time t . To learn the actor neural network parameterized by θ , we are interested in simultaneously maximizing the expected value of all tasks, that is

$$J(\theta) = \mathbb{E}_{\rho^\beta} \left[\sum_i Q_\mu^i(\mathbf{s}, \mu_\theta^i(\mathbf{s})) \right],$$

where μ_θ^i is the actor’s policy associated with the i -th task, mapping a state vector \mathbf{s} to a continuous action vector \mathbf{a} , and Q^i is the action-value critic associated with this task. The expectation above is taken with respect to ρ^β , the stationary distribution of some behavior policy $\beta(\mathbf{a}|\mathbf{s})$. We refer to the task whose behavior the agent follows during training as the *intentional* task, and to the remaining tasks as *unintentional*. Philosophically, in our work, the unintentional tasks are not thought of as playing a mere auxiliary role, but they can themselves become the intentional task. Our tasks are of a semantic nature, for example “*move the red block east of the blue block*”, and hence it is sensible to learn a library of such tasks for potential future re-use. Note that due to the fact that multiple policies are being learned at once we must necessarily be learning off-policy and will do so for all tasks including the intentional tasks. The corresponding gradient for the actor is

$$\nabla_\theta J(\theta) \approx \mathbb{E}_{\rho^\beta} \left[\sum_i \nabla_\theta \mu_\theta^i(\mathbf{s}) \nabla_{\mathbf{a}^i} Q_\mu^i(\mathbf{s}, \mathbf{a}^i) \Big|_{\mathbf{a}^i = \mu_\theta^i(\mathbf{s})} \right]. \quad (1)$$

The behaviour policy is effectively given by the intentional policies as we will detail shortly. Trajectories of the intentional policy during all stages of its training are stored in a replay buffer. Both intentional and unintentional policies are trained from the transitions in this replay buffer using off-policy learning. Given an observation \mathbf{s}_t , the behavior policy produces the action vector \mathbf{a}_t . In response, the environment returns a *reward vector* \mathbf{r}_t , with one scalar component for each task, and the next state observation \mathbf{s}_{t+1} . The tuple $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{r}_t, \mathbf{s}_{t+1})$ is stored in a replay buffer. Note that reward observations do not enter into the gradient estimate as they have instead been captured by the action-value function Q_μ . However, since this quantity is never directly observed, we will instead replace this function with a parameterized critic Q_w which must be trained by an appropriate policy evaluation mechanism. Here we update the critic in order to simultaneously minimize the temporal difference error of all tasks.

To update the critic and actor for each policy, we sample a mini-batch of tuples uniformly at random from the replay buffer and perform stochastic gradient descent with respect to both the actor and critic losses. Combining these updates we have,

$$\delta_j^i = r_j^i + \gamma Q_w^i(\mathbf{s}_{j+1}, \mu_{\theta'}^i(\mathbf{s}_{j+1})) - Q_w^i(\mathbf{s}_j, \mathbf{a}_j), \quad (2)$$

$$w \leftarrow w + \alpha_{\text{critic}} \sum_j \sum_i \delta_j^i \nabla_w Q_w^i(\mathbf{s}_j, \mathbf{a}_j), \quad (3)$$

$$\theta \leftarrow \theta + \alpha_{\text{actor}} \sum_j \sum_i \nabla_\theta \mu_\theta^i(\mathbf{s}_j) \nabla_{\mathbf{a}^i} Q_w^i(\mathbf{s}_j, \mathbf{a}^i) \Big|_{\mathbf{a}^i = \mu_\theta^i(\mathbf{s}_j)}. \quad (4)$$

Here j represents the sampled set of indices for the mini-batch. Note that in equation (2), the same action sampled from the replay buffer, \mathbf{a}_j , is used to update every critic Q^i . We have also adopted target parameters θ' and w' which stabilize learning of the critic, as in [4]. The target parameters are periodically updated to reflect the current values of the optimization parameters. In practice, rather than using fixed learning rates for the actor and critic we instead use dynamic learning rates, e.g. via the Adam optimizer [35].

We can now discuss the behavior policy used during training. In this work we use

$$\beta(\mathbf{a}|\mathbf{s}) = \mu_{\theta}^i(\mathbf{s}) + Z,$$

where Z is a random variable introduced for exploration; following [4] we utilize temporally correlated noise from an Ornstein-Uhlenbeck process. As before, the index i denotes the intentional task that is being followed. We considered different variations on this process, for example following a single task, selecting a random task to follow at the beginning of each episode, and switching between tasks whenever the current task is successful. However, for clarity of presentation, we focus on having a single intentional policy solving one of the hardest tasks in the physical playroom domain. We will return to this point in Section 5.

3 Experimental Setup

Understanding and manipulating relations among physical objects has been a long standing challenge in machine learning [36]. In this work we have embraced the complexity of realistic physical interactions, while maintaining a simple but diverse set of relations. Our experiments are set in a virtual physical world in which the agent can interact with a variety of objects and obtain rewards by satisfying many procedurally-generated semantic relations between the objects. The interactions among the physical objects involve complex contact forces, which can pose significant challenges for control algorithms due to resulting trajectories involving collisions, rotations, object-to-object sliding, multi-object dragging, getting stuck due to walls, etc.

Crucially, the agent is physically embodied in this domain and its actions have consistent dynamics throughout the space. Embodiment, object commonality and consistent physics enable the agent to learn features that effectively generalize between different tasks. In this section we will first describe the physical environment in which the agent is situated and then detail the method by which we automatically generate tasks and grounded rewards.

3.1 The physical playroom domain

We implemented a domain using the MuJoCo physics engine [37] which consists of a tabletop on which N objects are placed. Each object may have different properties associated with it, for example color, size, friction, density, shape and so on, which allow the objects to be either partially or fully specified.

The agent is embodied as an actuated “fist” whose action space consists of 2-dimensional velocities. In other words, the agent can move its fist through the playroom domain by setting its immediate velocity. The agent is able to interact with other objects and affect their positions only indirectly by way of contact with these objects. Although the fist is only able to move in 2-dimensions the other objects can exhibit more complex behavior based on the speed and angle at which contacts occur.

The playroom consists of a square arena with 80 cm edges, populated with N cubic objects with 12 cm edges. Figure 3.1 shows an example setting of this domain with $N = 3$ cubes of different colors. In the figure we can also see the agent’s actuated fist, displayed as a white sphere. The fist can be thought of as a directly controllable object, and as a result we will be able to directly express relations between this object and other objects in the scene.

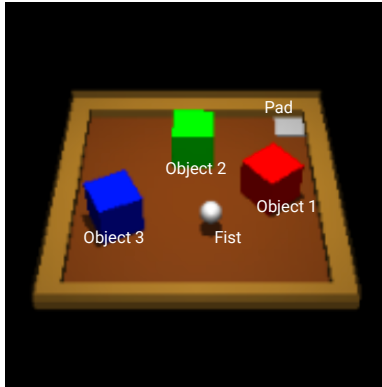


Figure 2: Example of the playroom environment showing three objects and the velocity controlled actuator (fist). Also shown is a goal pad location in the upper right corner.

We have also added a border to the environment in order to reduce the occurrence of irrecoverable actions wherein objects are pushed into corners and consequently cannot be retrieved. The border allows the fist to move into the corners but not the blocks. Finally, we will also consider relations between objects and a given goal position. These goals are implemented as immovable goal objects for the purposes of allowing relational expressions; in particular this object is implemented as a *pad* sitting at ground level and hence does not physically interact with any object. An example of one such object is displayed in Figure 3.1 as a white pad in the upper right corner.

3.2 Automatic reward generation with formal language

Given the above domain we are now interested in generating rewards based on the properties of objects as well as relations between these objects. We will define a number of property functions $p : \mathcal{O} \times \mathcal{S} \rightarrow \{0, 1\}$ where \mathcal{S} is the set of possible world states. Each property acts as an indicator over objects $o \in \mathcal{O}$, taking value 1 when the object satisfies a certain property. It is instructive to think of these properties as defining sets of objects in the scene. Often these properties will be independent of the state $s \in \mathcal{S}$ in which case we can write them simply as $p(o)$. For example we might define $p_{\text{blue}}(o)$ to be a property of blue objects, which is independent of their location, velocity, etc.

Next we can introduce binary relations between objects which can be thought of as functions of the form $b : \mathcal{O} \times \mathcal{O} \times \mathcal{S} \rightarrow \{0, 1\}$. For example, we can introduce a “nearness” relation which holds when two objects are close to each other, that is when their center points are within some specified distance. By combining properties and relations we can write rewards of the form

$$r_{\text{red_near_blue}}(\mathbf{s}) = \sum_{o_1, o_2} p_{\text{red}}(o_1) p_{\text{blue}}(o_2) b_{\text{near}}(o_1, o_2, \mathbf{s}).$$

where o_1 and o_2 represent object identifiers and we are summing over all pairs of objects such that the properties select a particular pair. The above reward represents the task “bring a red object near a blue object”. If the objects identified in the scene are uniquely identified by their properties this may also be more succinctly written as

$$r_{\text{red_near_blue}}(\mathbf{s}) = b_{\text{near}}(\text{red}, \text{blue}, \mathbf{s}).$$

where `red` and `blue` are the uniquely identified objects. While properties are frequently independent of the state, relations such as `near` will depend on this state, specifically the positions of its two input objects o_1 and o_2 .

We can generate many different rewards by logically combining these operations. In the experiments that follow we will make use of the following atoms:

1. color-based properties: $p_{\text{red}}, p_{\text{blue}}, p_{\text{green}}$;
2. properties identifying the fist p_{fist} and the goal state p_{goal} ;
3. near and far relations: $b_{\text{near}}, b_{\text{far}}$, parameterized by a distance parameter ϵ such that whenever the distance between two blocks is less than ϵ they are considered to be near;
4. directional relations: $b_{\text{north}}, b_{\text{east}}, b_{\text{south}}, b_{\text{west}}$ such that $b_{\text{north}}(a, b)$ is 1 whenever object a is north of object b .

Finally, more complicated rewards can be expressed by and-ing these relations. A “gather to pad” task which collects blocks of the three different colors to a goal pad location can be written as

$$r_{\text{gather_to_pad}}(\mathbf{s}) = b_{\text{near}}(\text{red}, \text{pad}, \mathbf{s}) b_{\text{near}}(\text{blue}, \text{pad}, \mathbf{s}) b_{\text{near}}(\text{green}, \text{pad}, \mathbf{s}).$$

One can also simply gather blocks together or move them far apart from each other, without having to specify a pad or coordinates.

4 Results

The agent observes the relative position of the fist actuator to each object in the scene. The environment is partially observable due to unobservable rotations of the cubic objects. Actions are taken by moving the fist in 2-dimensions, which involves setting the velocities of the actuator. We followed the training protocol described in Appendix A of the DDPG paper of Lillicrap et al. [4]. The actor and




Figure 3: **Left:** Test performance for the task of gathering two blocks together in the playroom, using a varying number of additional tasks. **Right:** What happens when an extra block is added to the environment, causing significant physical interference. In both cases, the more tasks the IU agent solves simultaneously, the faster it learns the intentional task.

critic multi-layer perceptrons, shown in Figure 1, have 200 and 400 units in each layer respectively and standard hyperbolic tangent activations.

We first consider tasks involving two uncontrolled objects (identified by the red and blue properties), as well as the fist actuator. Our tasks consist of 6 relations: north, south, east, west, near, and far, between each pair of objects including the fist. There are 3 distinct pairs of objects, which results in 18 total reward functions. We are interested in how the use of extra reward functions affects the ability of our agent to learn a single intentional policy.

Figure 4 (left) illustrates the ability of our agent to learn to maximize the single intentional reward $b_{\text{near}}(\text{red}, \text{blue}, \text{s})$, when given access to additional signals. We consider three scenarios: no additional unintentional tasks (i.e. standard DDPG), all near and far tasks (6), and finally all 18 tasks. We see that by considering all 18 tasks simultaneously, the IU agent learns to control the system under a single intentional reward much more rapidly. In these plots the test-time reward is shown, averaged over 50 runs, where the error bounds show the min and max rewards for a given evaluation.

Figure 4 shows the test performance of the 18 policies learned simultaneously. Note that some tasks are easier than others, explaining the different average rewards to which each of the tasks converge. Those tasks which are more difficult require more time to obtain non-zero values and hence converge to an average return less than 1.0. All tasks, however, are able to improve over their initial baseline. Note also that all tasks are capable of being learned, even those that are unintentional, which we can see from the fact that the reward increases for each of these 18 policies.

Next we examine the ability of the agents to learn when an additional (green) cube is added to the playroom. The additional object results in 6 distinct pairs of objects and as a result increases the number of tasks to 36 (6×6). The results are shown in the right-most plot of Figure 4. We see that although learning is slower, the system is still able to learn to achieve the gathering task despite the distracting block, which can of course cause significant interference. Once again, the more tasks the agent solves the faster it learns.

In Figure 4, we display the performance of the agent on tasks with three blocks. Here the intentional task involves placing all three blocks in the upper right corner pad. We display the average reward for this task when solving 1, 7, and 43 tasks. DDPG when following just a single task is incapable of succeeding at this task. However, the intentional unintentional agent succeeds when following 7 and 43 additional rewards. On the right panel of Figure 4, we display the performance of the agent when trained on a larger task space, that is a ground with sides that are 50% larger in length.

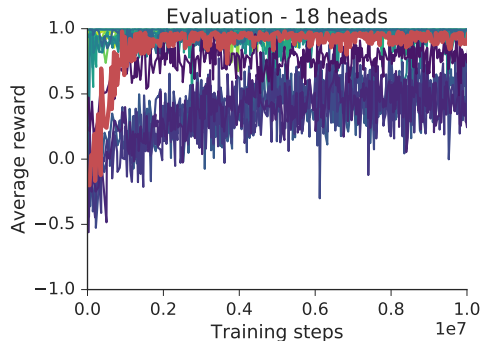


Figure 4: Test performance of the 18 policies 17 of which are unintentional and hence trained off-policy, with the intentional one shown in red, learned simultaneously in the experiment corresponding to Figure 3 (left).




Figure 5: **Left:** Test performance for the task of moving three blocks to the corner pad of the playroom. **Right:** Same task after increasing the length of the playroom sides by 50%. When the length is doubled the task can not be learned at all. For these very hard exploration task, the IU agent performs reasonably, while DDPG completely fails.

We observe that eventually, as the complexity of exploration increases, even the IU agents struggle. However by following 43 tasks the IU agent is still capable of gaining some reward, whereas DDPG completely fails.

Finally, in Figure 4 we display frames from example trajectories for various policies learned in the 3-block domain. Shown below these trajectories are illustrations of the object positions over time. The IU agent is able to learn reasonable action sequences which move the objects to the goal pad.

5 Discussion

Our experiments demonstrate that when acting according to the policy associated with one of the hardest tasks, we are able to learn all other tasks off-policy. The results for the playroom domain also showed that by increasing the number of tasks, all actors and critics learn faster. In fact, in some settings, learning with many goals was essential to solve hard many-body control tasks.

It is worth entertaining what happens when instead of following the policy associated with one of the hardest tasks, we follow a different behaviour strategy. For instance, what happens if we act according to a uniformly random mixture of actors? Is it advantageous to switch policies each time a policy succeeds in encountering a reward? Is it sound to choose behavior actors adaptively and in accordance with the critics' values?

The above questions are related to the problem of learning curricula. We conducted experiments to explore each of these questions, but found that the naive strategy of choosing the one of the hardest tasks for the behaviour policy works best. There are reasons for this. The behaviour policy determines what information is written to replay memory. These memories in turn are used to update the remaining actor-critics off-policy by gradient descent. When acting according to policies solving simple tasks, the replay memory ends up consisting mostly of experiences associated with the simple tasks, and consequently, the IU agent fails to explore. Populating the memory with rich experiences is essential for learning to solve tasks that involve more exploration.

Our experimental setup was restricted in the sense that the tasks had a nested structure. That is, approaching a block is needed to bring two blocks together, which in turn is needed to bring three blocks together, and so on. Our conclusions regarding curricula might be different in the absence of this nesting of tasks. This should be studied in future work.

It is not difficult to make the hardest task sufficiently difficult, from an exploration perspective, that even the IU agent fails. For example, if we increase the size of the playground and ask the actuator to bring ten blocks together, all existing control agents are likely to fail. The solution to this hard exploration problem appears to be one of task decomposition, either through hierarchical reinforcement learning or an understanding of objects and relations.

In this paper we focused on tasks that are sufficiently hard that not even popular continuous control algorithms such as DDPG can solve them. For this reason we obviated other experimental factors, such as perception from pixels, articulated bodies with more degrees of freedom, and diverse sets of objects. Each of these constitute important challenges that should be addressed in future research.

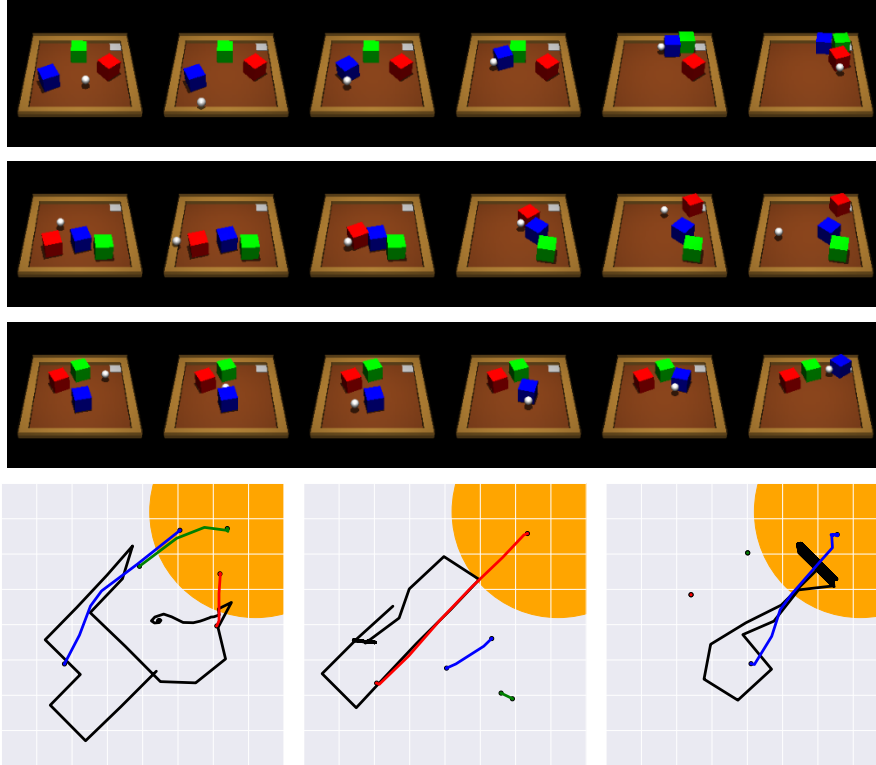


Figure 6: IU agent executing three different policies. The top film-strip shows the agent solving the hardest task: gathering all three blocks into the upper right corner. The two strips below this show the agent moving a single block into the corner; red and blue respectively. Below these temporal illustrations is a figure illustrating a coordinate-centric view of each of these three tasks. Here the black line denotes the position of the fist as a function of time, and the colored lines represent the positions of each of the three colored blocks. We can see from the first plot that the agent moves each block to reach the goal, shown as an orange semicircle. For the second and third tasks, only the red and blue blocks are acted on respectively, although we can see in the second plot that the agent unintentionally moves the other two blocks slightly. Videos can be seen at goo.gl/eNnP5W.

This work also did not touch upon the topic of policy re-use. That is, once the various policies are learned how do we construct other controllers that can harness these policies to solve new tasks. A simple way to achieve this is to train agents that combine all the policies, either via weighted combinations or more sophisticated deep networks, to construct new policies. Related to this, our work did not address the problem of pruning irrelevant policies.

In regard to hierarchical RL agents, it would be worthwhile investigating architectures where intrinsic motivation provides the ultimate reward, and is used to guide the automatic online addition or pruning of actors in the IU agent architecture.

6 Conclusions

We present in this work a novel architecture for learning several tasks at once. We also propose a flexible way of generating tasks in simulated physical environments. While these environments pose significant challenges because of complex contact forces among the objects, continuous action spaces, and difficult exploration, the body, objects and physical laws are shared among the tasks. In these domains, the more tasks, the faster the IU agents learn. In fact, they can learn to solve complex tasks where popular DDPG agents completely fail.

References

- [1] E. Thelen. Rhythmical stereotypies in normal human infants. *Animal Behaviour*, 27:699–715, 1979.
- [2] K. E. Adolph and S. R. Robinson. Motor development. In R. M. Lerner, editor, *Handbook of Child Psychology and Developmental Science*, volume 2, pages 114–157. John Wiley and Sons, 2015.
- [3] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller. Deterministic policy gradient algorithms. In *International Conference on Machine Learning*, 2014.
- [4] 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. In *International Conference on Learning Representations*, 2016.
- [5] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In *International Conference on Learning Representations*, 2017.
- [6] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.
- [7] R. S. Sutton, J. Modayil, M. Delp, T. Degris, P. M. Pilarski, A. White, and D. Precup. Horde: A scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 761–768, 2011.
- [8] T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *International Conference on Machine Learning*, pages 1312–1320, 2015.
- [9] T. G. Dietterich. The MAXQ method for hierarchical reinforcement learning. In *International Conference on Machine Learning*, 1998.
- [10] R. S. Sutton, D. Precup, and S. P. Singh. Intra-option learning about temporally abstract actions. In *International Conference on Machine Learning*, pages 556–564, 1998.
- [11] T. D. Kulkarni, K. Narasimhan, A. Saeedi, and J. Tenenbaum. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pages 3675–3683, 2016.
- [12] H. van Seijen, M. Fatemi, J. Romoff, R. Larocche, T. Barnes, and J. Tsang. Hybrid reward architecture for reinforcement learning. Technical report, Microsoft Maluuba, 2017.
- [13] S. Russell and A. Zimdars. Q-decomposition for reinforcement learning agents. In *International Conference on Machine Learning*, 2003.
- [14] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [15] D. M. Roijers, P. Vamplew, S. Whiteson, and R. Dazeley. A survey of multi-objective sequential decision-making. *J. Artif. Int. Res.*, 48(1):67–113, 2013.
- [16] A. Dosovitskiy and V. Koltun. Learning to act by predicting the future. In *International Conference on Learning Representations*, 2017.
- [17] X. Li, L. Li, J. Gao, X. He, J. Chen, L. Deng, and J. He. Recurrent reinforcement learning: a hybrid approach. *Preprint arXiv:1509.03044*, 2015.
- [18] G. Lample and D. S. Chaplot. Playing FPS games with deep reinforcement learning. *Preprint arXiv:1609.05521*, 2016.

- [19] H. Bou-Ammar, E. Eaton, P. Ruvolo, and M. E. Taylor. Online multi-task learning for policy gradient methods. In *ICML*, 2014.
- [20] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. *CoRR*, abs/1707.01495, 2017.
- [21] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *International Conference on Machine Learning*, pages 41–48, 2009.
- [22] W. Zaremba and I. Sutskever. Learning to execute. *Preprint arXiv:1410.4615*, 2014.
- [23] S. Reed and N. de Freitas. Neural programmer-interpreters. In *International Conference on Learning Representations*, 2016.
- [24] Y. Tsvetkov, M. Faruqui, W. Ling, B. MacWhinney, and C. Dyer. Learning the curriculum with Bayesian optimization for task-specific word representation learning. In *Association for Computational Linguistics*, 2016.
- [25] A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu. Automated curriculum learning for neural networks. *Preprint arXiv:1704.03003*, 2017.
- [26] J. W. Sparling, J. Van Tol, and N. C. Chescheir. Fetal and neonatal hand movement. *Physical Therapy*, 79(1):24, 1999.
- [27] M. S. Blumberg, H. G. Marques, and F. Iida. Twitching in sensorimotor development from sleeping rats to robots. *Current Biology*, 23(12):R532–R537, 2013.
- [28] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science*, 15(4):151–190, 2003.
- [29] R. A. Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.
- [30] M. L. Littman, U. Topcu, J. Fu, C. Isbell, M. Wen, and J. MacGlashan. Environment-independent task specifications via GLTL. Technical report, Brown University, 2017.
- [31] H. Yu, H. Zhang, and W. Xu. A deep compositional framework for human-like language acquisition in virtual environment. Technical report, Baidu Research, 2017.
- [32] K. M. Hermann, F. Hill, S. Green, F. Wang, R. Faulkner, H. Soyer, D. Szepesvari, W. Czarnecki, M. Jaderberg, D. Teplyashin, M. Wainwright, C. Apps, D. Hassabis, and P. Blunsom. Grounded language learning in a simulated 3D world. Technical report, DeepMind, 2017.
- [33] M. Denil, S. G. Colmenarejo, S. Cabi, D. Saxton, and N. de Freitas. Programmable agents. Technical report, DeepMind, 2017.
- [34] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, 1999.
- [35] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [36] T. Winograd. Understanding natural language. *Cognitive Psychology*, 3(1):1 – 191, 1972.
- [37] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *International Conference on Intelligent Robots and Systems*, 2012.