

Reinforcement Learning in POMDPs with Memoryless Options and Option-Observation Initiation Sets

Denis Steckelmacher,¹ Diederik M. Roijers,¹ Anna Harutyunyan,¹
Peter Vrancx,² H el ene Plisnier,¹ Ann Now e¹

¹ Artificial Intelligence Lab, Vrije Universiteit Brussel (VUB), Pleinlaan 2, 1050 Brussels, Belgium

² PROWLER.io, 66-68 Hills Road, Cambridge CB2 1LA, United Kingdom; work conducted while affiliated to the VUB

Abstract

Many real-world reinforcement learning problems have a hierarchical nature, and often exhibit some degree of partial observability. While hierarchy and partial observability are usually tackled separately (for instance by combining recurrent neural networks and options), we show that addressing both problems simultaneously is simpler and more efficient in many cases. More specifically, we make the initiation set of options conditional on the previously-executed option, and show that options with such Option-Observation Initiation Sets (OOIs) are at least as expressive as Finite State Controllers (FSCs), a state-of-the-art approach for learning in POMDPs. OOIs are easy to design based on an intuitive description of the task, lead to explainable policies and keep the top-level and option policies memoryless. Our experiments show that OOIs allow agents to learn optimal policies in challenging POMDPs, while being much more sample-efficient than a recurrent neural network over options.

1 Introduction

Real-world applications of reinforcement learning (RL) face two main challenges: complex long-running tasks and partial observability. Options, the particular instance of Hierarchical RL we focus on, addresses the first challenge by factoring a complex task into simpler sub-tasks (Barto et al. 2003; Roy et al. 2006; Tessler et al. 2016). Instead of learning what action to perform depending on an observation, the agent learns a top-level policy that repeatedly selects options, that in turn execute sequences of actions before returning (Sutton et al. 1999). The second challenge, partial observability, is addressed by maintaining a belief of what the agent thinks the full state is (Kaelbling et al. 1998; Cassandra et al. 1994), reasoning about possible future observations (Littman et al. 2001; Boots et al. 2011), storing information in an external memory for later reuse (Peshkin et al. 1999; Zaremba et al. 2015; Graves et al. 2016), or using recurrent neural networks (RNNs) to allow information to flow between time-steps (Bakker 2001; Mnih et al. 2016).

In *planning*, options are already used to increase the planning horizon under uncertainty (He et al. 2011; Lim et al. 2011) and to construct policies robust to partial observability (Omidshafiei et al. 2017), but *learning* algo-

rithms for hierarchical partially observable tasks are not yet ideal. HQ-Learning decomposes a task into a sequence of fully-observable subtasks (Wiering et al. 1997), which precludes cyclic tasks from being solved. Using recurrent neural networks in options and for the top-level policy (Sridharan et al. 2010) addresses both challenges, but brings in the design complexity of RNNs (J ozefowicz et al. 2015; Angeline et al. 1994; Mikolov et al. 2014). RNNs also have limitations regarding long time horizons, as their memory decays over time (Hochreiter et al. 1997).

In her PhD thesis, Precup (2000, page 126) suggests that options may already be close to addressing partial observability, thus removing the need for more complicated solutions. In this paper, we prove this intuition correct by:

1. Showing that standard options do not suffice in POMDPs;
2. Introducing Option-Observation Initiation Sets (OOIs), that make the initiation sets of options conditional on the previously-executed option;
3. Proving that OOIs make options at least as expressive as Finite State Controllers (Section 3.2), thus able to tackle challenging POMDPs.

In contrast to existing HRL algorithms for POMDPs (Wiering et al. 1997; Theodorou 2002; Sridharan et al. 2010), OOIs handle repetitive tasks, do not restrict the action set available to sub-tasks, and keep the top-level and option policies memoryless. A wide range of robotic and simulated experiments in Section 4 confirm that OOIs allow partially observable tasks to be solved optimally, demonstrate that OOIs are much more sample-efficient than a recurrent neural network over options, and illustrate the flexibility of OOIs regarding the amount of domain knowledge available at design time. In Section 4.5, we demonstrate the robustness of OOIs to sub-optimal option sets. While it is generally accepted that the designer provides the options and their initiation sets, we show in Section 4.4 that random initiation sets, combined with learned option policies and termination functions, allow OOIs to be used without any domain knowledge.

1.1 Motivating Example

OOIs are designed to solve complex partially-observable tasks that can be decomposed into a set of fully-observable

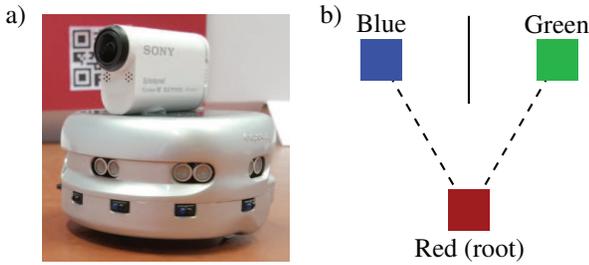


Figure 1: Robotic object gathering task. a) Khepera III, the two-wheeled robot used in the experiments. b) The robot has to gather objects from two terminals separated by a wall, and to bring them to the root.

sub-tasks. For instance, a robot with first-person sensors may be able to avoid obstacles, open doors or manipulate objects even if its precise location in the building is not observed. We now introduce such an environment, on which our robotic experiments of Section 4.3 are based.

A Khepera III robot¹ has to gather objects from two terminals separated by a wall, and to bring them to the root (see Figure 1). Objects have to be gathered one by one from a terminal until it becomes empty, which requires many journeys between the root and a terminal. When a terminal is emptied, the other one is automatically refilled. The robot therefore has to alternatively gather objects from both terminals, and the episode finishes after the terminals have been emptied some random number of times. The root is colored in red and marked by a paper QR-code encoding 1. Each terminal has a screen displaying its color and a dynamic QR-code (1 when full, 2 when empty). Because the robot cannot read QR-codes from far away, the state of a terminal cannot be observed from the root, where the agent has to decide to which terminal it will go. This makes the environment partially observable, and requires the robot to remember which terminal was last visited, and whether it was full or empty.

The robot is able to control the speed of its two wheels. A wireless camera mounted on top of the robot detects bright color blobs in its field of view, and can read nearby QR-codes. Such low-level actions and observations, combined with a complicated task, motivate the use of hierarchical reinforcement learning. Fixed options allow the robot to move towards the largest red, green or blue blob in its field of view. The options terminate as soon as a QR-code is in front of the camera and close enough to be read. The robot has to learn a policy over options that solves the task.

The robot may have to gather a large number of objects, alternating between terminals several times. The repetitive nature of this task is incompatible with HQ-Learning (Wiering et al. 1997). Options with standard initiation sets are not able to solve this task, as the top-level policy is memoryless (Sutton et al. 1999) and cannot remember from which terminal the robot arrives at the root, and whether that terminal was full or empty. Because the terminals are a dozen feet

¹<http://www.k-team.com/mobile-robotics-products/old-products/khepera-iii>

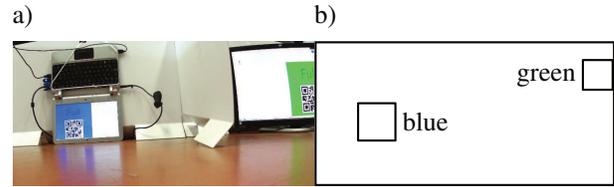


Figure 2: Observations of the Khepera robot. a) Color image from the camera. b) Color blobs detected by the vision system, as observed by the robot. QR-codes can only be decoded when the robot is a couple of inches away from them.

away from the root, almost a hundred primitive actions have to be executed to complete any root/terminal journey. Without options, this represents a time horizon much larger than usually handled by recurrent neural networks (Bakker 2001) or finite history windows (Lin et al. 1993).

OOIs allow each option to be selected conditionally on the previously executed one (see Section 3.1), which is much simpler than combining options and recurrent neural networks (Sridharan et al. 2010). The ability of OOIs to solve complex POMDPs builds on the time abstraction capabilities and expressiveness of options. Section 4.3 shows that OOIs allow a policy for our robotic task to be learned to expert level. Additional experiments demonstrate that both the top-level and option policies can be learned by the agent (see Section 4.4), and that OOIs lead to substantial gains over standard initiation sets even if the option set is reduced or unsuited to the task (see Section 4.5).

2 Background

This section formally introduces Markov Decision Processes (MDPs), Options, Partially Observable MDPs (POMDPs) and Finite State Controllers, before presenting our main contribution in Section 3.

2.1 Markov Decision Processes

A discrete-time Markov Decision Process (MDP) $\langle S, A, R, T, \gamma \rangle$ with discrete actions is defined by a possibly-infinite set S of states, a finite set A of actions, a reward function $R(s_t, a_t, s_{t+1}) \in \mathbb{R}$, that provides a scalar reward r_t for each state transition, a transition function $T(s_t, a_t, s_{t+1}) \in [0, 1]$, that outputs a probability distribution over new states s_{t+1} given a (s_t, a_t) state-action pair, and $0 \leq \gamma < 1$ the discount factor, that defines how sensitive the agent should be to future rewards.

A stochastic memoryless policy $\pi(s_t, a_t) \in [0, 1]$ maps a state to a probability distribution over actions. The goal of the agent is to find a policy π^* that maximizes the expected cumulative discounted reward $E_{\pi^*} [\sum_t \gamma^t r_t]$ obtainable by following that policy.

2.2 Options

The options framework, defined in the context of MDPs (Sutton et al. 1999), consists of a set of options O where each option $\omega \in O$ is a tuple $\langle \pi_\omega, I_\omega, \beta_\omega \rangle$, with $\pi_\omega(s_t, a_t) \in$

$[0, 1]$ the memoryless option policy, $\beta_\omega(s_t) \in [0, 1]$ the termination function that gives the probability for the option ω to terminate in state s_t , and $I_\omega \subseteq S$ the initiation set that defines in which states ω can be started (Sutton et al. 1999).

The memoryless top-level policy $\mu(s_t, \omega_t) \in [0, 1]$ maps states to a distribution over options and allows to choose which option to start in a given state. When an option ω is started, it executes until termination (due to β_ω), at which point μ selects a new option based on the now current state.

2.3 Partially Observable MDPs

Most real-world problems are not completely captured by MDPs, and exhibit at least some degree of partial observability. A Partially Observable MDP (POMDP) $\langle \Omega, S, A, R, T, W, \gamma \rangle$ is an MDP extended with two components: the possibly-infinite set Ω of observations, and the $W : S \rightarrow \Omega$ function that produces observations x based on the unobservable state s of the process. Two different states, requiring two different optimal actions, may produce the same observation. This makes POMDPs remarkably challenging for reinforcement learning algorithms, as memoryless policies, that select actions or options based only on the current observation, typically no longer suffice.

2.4 Finite State Controllers

Finite State Controllers (FSCs) are commonly used in POMDPs. An FSC $\langle \mathcal{N}, \psi, \eta, \eta^0 \rangle$ is defined by a finite set \mathcal{N} of nodes, an action function $\psi(n_t, a_t) \in [0, 1]$ that maps nodes to a probability distribution over actions, a successor function $\eta(n_{t-1}, x_t, n_t) \in [0, 1]$ that maps nodes and observations to a probability distribution over next nodes, and an initial function $\eta^0(x_1, n_1) \in [0, 1]$ that maps initial observations to nodes (Meuleau et al. 1999).

At the first time-step, the agent observes x_1 and activates a node n_1 by sampling from $\eta^0(x_1, \cdot)$. An action is performed by sampling from $\psi(n_1, \cdot)$. At each time-step t , a node n_t is sampled from $\eta(n_{t-1}, x_t, \cdot)$, then an action a_t is sampled from $\psi(n_t, \cdot)$. FSCs allow the agent to select actions according to the entire history of past observations (Meuleau et al. 1999), which has been shown to be one of the best approaches for POMDPs (Lin et al. 1992). OOIs, our main contribution, make options at least as expressive and as relevant to POMDPs as FSCs, while being able to leverage the hierarchical structure of the problem.

3 Option-Observation Initiation Sets

Our main contribution, Option-Observation Initiation Sets (OOIs), make the initiation sets of options conditional on the option that has just terminated. We prove that OOIs make options at least as expressive as FSCs (thus suited to POMDPs, see Section 3.2), even if the top-level and option policies are memoryless, while options without OOIs are strictly less expressive than FSCs (see Section 3.3). In Section 4, we show on one robotic and two simulated tasks that OOIs allow challenging POMDPs to be solved optimally.

3.1 Conditioning on Previous Option

Descriptions of partially observable tasks in natural language often contain allusions at sub-tasks that must be se-

quenced or cycled through, possibly with branches. This is easily mapped to a policy over options (learned by the agent) and sets of options that may or may not follow each other.

A good memory-based policy for our motivating example, where the agent has to bring objects from two terminals to the root (see Section 1.1), can be described as “go to the green terminal, then go to the root, then go back to the green terminal if it was full, to the blue terminal otherwise”, and symmetrically so for the blue terminal. This sequence of sub-tasks, that contains a condition, is easily translated to a set of options. Two options, ω_{GF} and ω_{GE} , sharing a single policy, go from the green terminal to the root (using low-level motor actions). ω_{GF} is executed when the terminal is full, ω_{GE} when it is empty. At the root, the option that goes back to the green terminal can only follow ω_{GF} , not ω_{GE} . When the green terminal is empty, going back to it is therefore forbidden, which forces the agent to switch to the blue terminal when the green one is empty.

We now formally define our main contribution, Option-Observation Initiation Sets (OOIs), that allow to describe which options may follow which ones. We define the initiation set I_ω of option ω so that the set \mathcal{O}_t of options available at time t depends on the observation x_t and previously-executed option ω_{t-1} :

$$I_\omega \subseteq \Omega \times (O \cup \{\emptyset\})$$

$$\mathcal{O}_t \equiv \{\omega \in O : (x_t, \omega_{t-1}) \in I_\omega\}$$

with $\omega_0 = \emptyset$, Ω the set of observations and O the set of options. \mathcal{O}_t allows the agent to condition the option selected at time t on the one that has just terminated, even if the top-level policy does not observe ω_{t-1} . The top-level and option policies remain memoryless. Not having to observe ω_{t-1} keeps the observation space of the top-level policy small, instead of extending it to $\Omega \times O$, without impairing the representational power of OOIs, as shown in the next sub-section.

3.2 OOIs Make Options as Expressive as FSCs

Finite State Controllers are state-of-the-art in policies applicable to POMDPs (Meuleau et al. 1999). By proving that options with OOIs are as expressive as FSCs, we provide a lower bound on the expressiveness of OOIs and ensure that they are applicable to a wide range of POMDPs.

Theorem 1. *OOIs allow options to represent any policy that can be expressed using a Finite State Controller.*

Proof. The reduction from any FSC to options requires one option $\langle n'_{t-1}, n_t \rangle$ per ordered pair of nodes in the FSC, and one option $\langle \emptyset, n_1 \rangle$ per node in the FSC. Assuming that $n_0 = \emptyset$ and $\eta(\emptyset, x_1, \cdot) = \eta^0(x_1, \cdot)$, the options are defined by:

$$\beta_{\langle n'_{t-1}, n_t \rangle}(x_t) = 1 \tag{1}$$

$$\pi_{\langle n'_{t-1}, n_t \rangle}(x_t, a_t) = \psi(n_t, a_t) \tag{2}$$

$$\mu(x_t, \langle n'_{t-1}, n_t \rangle) = \eta(n'_{t-1}, x_t, n_t) \tag{3}$$

$$I_{\langle \emptyset, n_1 \rangle} = \Omega \times \{\emptyset\}$$

$$I_{\langle n'_{t-1}, n_t \rangle} = \Omega \times \{\langle n'_{t-2}, n_{t-1} \rangle : n'_{t-1} = n_{t-1}\}$$

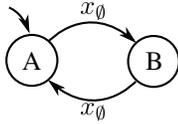


Figure 3: Two-nodes Finite State Controller that emits an infinite sequence ABAB... based on an uninformative observation x_\emptyset . This FSC cannot be expressed using options without OOIs.

Each option corresponds to an edge of the FSC. Equation 1 ensures that every option stops after having emitted a single action, as the FSC takes one transition every time-step. Equation 2 maps the current option to the action emitted by the destination node of its corresponding FSC edge. We show that μ and $I_{\langle n'_{t-1}, n_t \rangle}$ implement $\eta(n_{t-1}, x_t, n_t)$, with $\omega_{t-1} = \langle n'_{t-2}, n_{t-1} \rangle$, by:

$$\mu(x_t, \langle n'_{t-1}, n_t \rangle) = \begin{cases} \eta(n_{t-1}, x_t, n_t) & \langle n'_{t-2}, n_{t-1} \rangle \in I_{\langle n'_{t-1}, n_t \rangle} \\ & \Leftrightarrow n'_{t-1} = n_{t-1} \\ 0 & \langle n'_{t-2}, n_{t-1} \rangle \notin I_{\langle n'_{t-1}, n_t \rangle} \\ & \Leftrightarrow n'_{t-1} \neq n_{t-1} \end{cases}$$

Because η maps nodes to nodes and μ selects options representing pairs of nodes, μ is extremely sparse and returns a value different from zero, $\eta(n_{t-1}, x_t, n_t)$, only when $\langle n'_{t-2}, n_{t-1} \rangle$ and $\langle n'_{t-1}, n_t \rangle$ agree on n_{t-1} . \square

Our reduction uses options with trivial policies, that execute for a single time-step, which leads to a large amount of options to compensate. In practice, we expect to be able to express policies for real-world POMDPs with much less options than the number of states an FSC would require, as shown in our simulated (Section 4.4, 2 options) and robotic experiments (Section 4.3, 12 options). In addition to being sufficient, the next sub-section proves that OOIs are necessary for options to be as expressive as FSCs.

3.3 Original Options are not as Expressive as FSCs

While options with regular initiation sets are able to express some memory-based policies (Sutton et al. 1999, page 7), the tiny but valid Finite State Controller presented in Figure 3 cannot be mapped to a set of options and a policy over options (without OOIs). This proves that options without OOIs are strictly less expressive than FSCs.

Theorem 2. *Options without OOIs are not as expressive as Finite State Controllers.*

Proof. Figure 3 shows a Finite State Controller that emits a sequence of alternating A's and B's, based on a constant uninformative observation x_\emptyset . This task requires memory because the observation does not provide any information about what was the last letter to be emitted, or which one must now be emitted. Options having memoryless policies,

options executing for multiple time-steps are unable to represent the FSC exactly. A combination of options that execute for a single time-step cannot represent the FSC either, as the options framework is unable to represent memory-based policies with single-time-step options (Sutton et al. 1999). \square

4 Experiments

The experiments in this section illustrate how OOIs allow agents to perform optimally in environments where options without OOIs fail. Section 4.3 shows that OOIs allow the agent to learn an expert-level policy for our motivating example (Section 1.1). Section 4.4 shows that the top-level and option policies required by a repetitive task can be learned, and that learning option policies allow the agent to leverage random OOIs, thereby removing the need for designing them. In Section 4.5, we progressively reduce the amount of options available to the agent, and demonstrate how OOIs still allow good memory-based policies to emerge when a sub-optimal amount of options are used.

All our results are averaged over 20 runs, with standard deviation represented by the light regions in the figures. The source code, raw experimental data, run scripts, and plotting scripts of our experiments, along with a detailed description of our robotic setup, are available on GitHub.² A video detailing our robotic experiment is available on YouTube.³

4.1 Learning Algorithm

All our agents learn their top-level and option policies (if not provided) using a simplified version of the Option-Critic architecture (Bacon et al. 2017), where the critic is replaced by Monte-Carlo returns, and a single neural network learns the policies and termination functions of all the options. Our neural network π takes three inputs and produces one output. The inputs are problem-specific observation features \mathbf{x} , the one-hot encoded current option ω ($\omega = \mathbf{0}$ when executing the top-level policy), and a mask, \mathbf{mask} . The output \mathbf{y} is the joint probability distribution over selecting actions or options (so that the same network can be used for the top-level and option policies), while terminating or continuing the current option:

$$\begin{aligned} \mathbf{h}_1 &= \tanh(\mathbf{W}_1[\mathbf{x}^T \omega^T]^T + \mathbf{b}_1), \\ \hat{\mathbf{y}} &= \sigma(\mathbf{W}_2 \mathbf{h}_1 + \mathbf{b}_2) \circ \mathbf{mask}, \\ \mathbf{y} &= \frac{\hat{\mathbf{y}}}{\mathbf{1}^T \hat{\mathbf{y}}}, \end{aligned}$$

with W_i and b_i the trainable weights and biases of layer i , σ the sigmoid function, and \circ the element-wise product of two vectors. The fraction ensures that a valid probability distribution is produced by the network. The initiation sets of options are implemented using the \mathbf{mask} input of the neural network, a vector of $2 \times (|A| + |O|)$ integers, the same dimension as the \mathbf{y} output. When executing the top-level policy

²github.com/vub-ai-lab/options-with-oois

³youtu.be/VprJZEOD5NE

($\omega = \mathbf{0}$), the mask forces the probability of primitive actions to zero, preserves option ω_i according to I_{ω_i} , and prevents the top-level policy from terminating. When executing an option policy ($\omega \neq \mathbf{0}$), the mask only allows primitive actions to be executed. For instance, if there are two options and three actions, $\text{mask} = \begin{matrix} \text{end} \\ \text{cont} \end{matrix} \begin{pmatrix} 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix}$ when executing any of the options. When executing the top-level policy, $\text{mask} = \begin{matrix} \text{end} \\ \text{cont} \end{matrix} \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ a & b & 0 & 0 & 0 \end{pmatrix}$, with $a = 1$ if and only if the option that has just finished is in the initiation set of the first option, and $b = 1$ according to the same rule but for the second option. The neural network π is trained using Policy Gradient, with the following loss:

$$\mathcal{L}(\pi) = - \sum_{t=0}^T (\mathcal{R}_t - V(x_t, \omega_t)) \log(\pi(x_t, \omega_t, a_t))$$

with $a_t \sim \pi(x_t, \omega_t, \cdot)$ the action executed at time t . The return $\mathcal{R}_t = \sum_{\tau=t}^T \gamma^\tau r_\tau$, with $r_\tau = R(s_\tau, a_\tau, s_{\tau+1})$, is a simple discounted sum of future rewards, and ignores changes of current option. This gives the agent information about the complete outcome of an action or option, by directly evaluating its flattened policy. A baseline $V(x_t, \omega_t)$ is used to reduce the variance of the \mathcal{L} estimate (Sutton et al. 2000). $V(x_t, \omega_t)$ predicts the expected cumulative reward obtainable from x_t in option ω_t using a separate neural network, trained on the monte-carlo return obtained from x_t in ω_t .

4.2 Comparison with LSTM over Options

In order to provide a complete evaluation of OOIs, a variant of the π and V networks of Section 4.1, where the hidden layer is replaced with a layer of 20 LSTM units (Hochreiter et al. 1997; Sridharan et al. 2010), is also evaluated on every task. We use 20 units as this leads to the best results in our experiments, which ensures a fair comparison of LSTM against OOIs. In all experiments, the LSTM agents are provided the same set of options as the agent with OOIs. Not providing any option, or less options, leads to worse results. Options allow the LSTM network to focus on important observations, and reduces the time horizon to be considered. Shorter time horizons have been shown to be beneficial to LSTM (Bakker 2001).

Despite our efforts, LSTM over options only manages to learn good policies in our robotic experiment (see Section 4.3), and requires more than twice the amount of episodes as OOIs to do so. In our repetitive task, dozens of repetitions seem to confuse the network, that quickly diverges from any good policy it may learn (see Section 4.4). On TreeMaze, a much more complex version of the T-maze task, originally used to benchmark LSTM agents (Bakker 2001), the LSTM agent learns the optimal policy after more than 100K episodes (not shown on the figures). These results illustrate how learning with recurrent neural networks is sometimes difficult, and how OOIs allow to reliably obtain good results, with minimal engineering effort.

4.3 Object Gathering

The first experiment illustrates how OOIs allow an expert-level policy to be learned for a complex robotic partially-

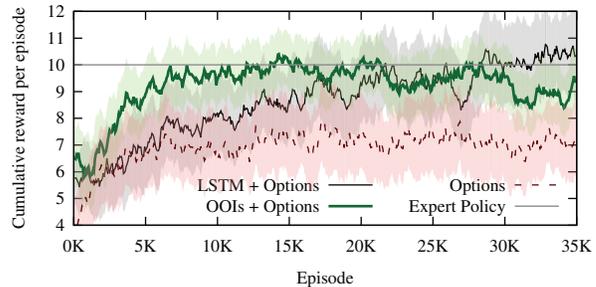


Figure 4: Cumulative reward per episode obtained on our object gathering task, with OOIs, without OOIs, and using an LSTM over options. OOIs learns an expert-level policy much quicker than an LSTM over options. The LSTM curve flattens-out (with high variance) after about 30K episodes.

observable repetitive task. The experiment takes place in the environment described in Section 1.1. A robot has to gather objects one by one from two terminals, green and blue, and bring them back to the root location. Because our actual robot has no effector, it navigates between the root and the terminals, but only pretends to move objects. The agent receives a reward of +2 when it reaches a full terminal, -2 when the terminal is empty. At the beginning of the episode, each terminal contains 2 to 4 objects, this amount being selected randomly for each terminal. When the agent goes to an empty terminal, the other one is re-filled with 2 to 4 objects. The episode ends after 2 or 3 emptyings (combined across both terminals). Whether a terminal is full or empty is observed by the agent only when it is at the terminal. The agent therefore has to remember information acquired at terminals in order to properly choose, at the root, to which terminal it will go.

The agent has access to 12 memoryless options that go to red ($\omega_{R1..R4}$), green ($\omega_{G1..G4}$) or blue objects ($\omega_{B1..B4}$), and terminate when the agent is close enough to them to read a QR-code displayed on them. The initiation set of $\omega_{R1,R2}$ is $\omega_{G1..G4}$, of $\omega_{R3,R4}$ is $\omega_{B1..B4}$, and of ω_{G_i,B_i} is $\omega_{R_i} \forall i = 1..4$. This description of the options and their OOIs is purposefully uninformative, and illustrates how little information the agent has about the task. The option set used in this experiment is also richer than the simple example of Section 3.1, so that the solution of the problem, not going back to an empty terminal, is not encoded in OOIs but must be learned by the agent.

Agents with and without OOIs learn top-level policies over these options. We compare them to a *fixed* agent, using an expert top-level policy that interprets the options as follows: $\omega_{R1..R4}$ go to the root from a full/empty green/blue terminal (and are selected accordingly at the terminals depending on the QR-code displayed on them), while $\omega_{G1..G4,B1..B4}$ go to the green/blue terminal from the root when the previous terminal was full/empty and green/blue. At the root, OOIs ensure that only one option amongst *go to green after a full green*, *go to green after an empty blue*, *go to blue after a full blue* and *go to blue after an empty green*

is selected by the top-level policy: the one that corresponds to what color the last terminal was and whether it was full or empty. The agent goes to a terminal until it is empty, then switches to the other terminal, leading to an average cumulative reward of 10^4 .

When the top-level policy is learned, OOs allow the task to be solved, as shown in Figure 4, while standard initiation sets do not allow the task to be learned. Because experiments on a robot are slow, we developed a small simulator for this task, and used it to produce Figure 4 after having successfully asserted its accuracy using two 1000-episodes runs on the actual robot. The agent learns to properly select options at the terminals, depending on the QR-code, and to output a proper distribution over options at the root, thereby matching our expert policy. The LSTM agent learns the policy too, but requires more than twice the amount of episodes to do so. The high variance displayed in Figure 4 comes from the varying amounts of objects in the terminals, and the random selection of how many times they have to be emptied.

Because fixed option policies are not always available, we now show that OOs allow them to be learned at the same time as the top-level policy.

4.4 Modified DuplicatedInput

In some cases, a hierarchical reinforcement learning agent may not have been provided policies for several or any of its options. In this case, OOs allow the agent to learn its top-level policy, the option policies and their termination functions. In this experiment, the agent has to learn its top-level and option policies to copy characters from an input tape to an output tape, removing duplicate B's and D's (mapping ABCCEDD to ABCCED for instance; B's and D's always appear in pairs). The agent only observes a single input character at a time, and can write at most one character to the output tape per time-step.

The input tape is a sequence of N symbols $x \in \Omega$, with $\Omega = \{A, B, C, D, E\}$ and N a random number between 20 and 30. The agent observes a single symbol $x_t \in \Omega$, read from the i -th position in the input sequence, and does not observe i . When $t = 1$, $i = 0$. There are 20 actions ($5 \times 2 \times 2$), each of them representing a symbol (5), whether it must be pushed onto the output tape (2), and whether i should be incremented or decremented (2). A reward of 1 is given for each correct symbol written to the output tape. The episode finishes with a reward of -0.5 when an incorrect symbol is written.

The agent has access to two options, ω_1 and ω_2 . OOs are designed so that ω_2 cannot follow itself, with no such restriction on ω_1 . No reward shaping or hint about what each option should do is provided. The agent automatically discovers that ω_1 must copy the current character to the output, and that ω_2 must skip the character without copying it. It also learns the top-level policy, that selects ω_2 (skip) when observing B or D and ω_1 is allowed, ω_1 otherwise (copy).

⁴ $\frac{2+3}{2} \times (-2 + \frac{2+4}{2} \times 2)$, 2 or 3 emptyings of terminals that contain 2 to 4 objects. Average confirmed experimentally from 1000 episodes using the policy, $p > 0.30$.

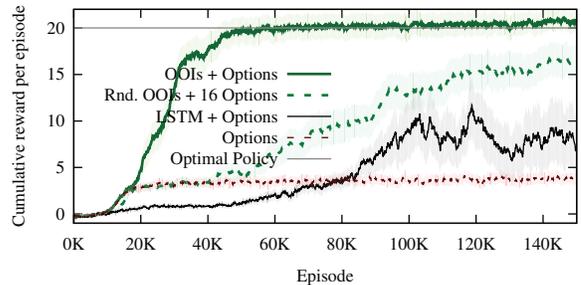


Figure 5: Cumulative reward per episode obtained on modified DuplicatedInput, with random or designed OOs, without OOs and using an LSTM over options. Despite our efforts, an LSTM over options repeatedly learns then forgets optimal policies, as shown by the high variance of its line.

Figure 5 shows that an agent with two options and OOs learns the optimal policy for this task, while an agent with two options and only standard initiation sets ($I_\omega = \Omega \forall \omega$) fails to do so. The agent without OOs only learns to copy characters and never skips any (having two options does not help it). This shows that OOs are necessary for learning this task, and allow to learn top-level and option policies suited to our repetitive partially observable task.

When the option policies are learned, the agent becomes able to adapt itself to random OOs, thereby removing the need for designing them. For an agent with N options, each option has $\frac{N}{2}$ randomly-selected options in its initiation set, with the initiation sets re-sampled for each run. The agents learn how to leverage their option set, and achieve good results on average (16 options used in Figure 5, more options lead to better results). When looking at individual runs, random OOs allow optimal policies to be learned, but several runs require more time than others to do so. This explains the high variance and noticeable steps shown in Figure 5.

The next section shows that an improperly-defined set of human-provided options, as may happen in design phase, still allows the agent to perform reasonably well. Combined with our results with random OOs, this shows that OOs can be tailored to the exact amount of domain knowledge available for a particular task.

4.5 TreeMaze

The optimal set of options and OOs may be difficult to design. When the agent learns the option policies, the previous section demonstrates that random OOs suffice. This experiment focuses on human-provided option policies, and shows that a sub-optimal set of options, arising from a misspecification of the environment or normal trial-and-error in design phase, does not prevent agents with OOs from learning reasonably good policies.

TreeMaze generalizes the T-maze environment (Bakker 2001) to arbitrary heights. The agent starts at the root of the tree-like maze depicted in Figure 6, and has to reach the extremity of one of the 8 leaves. The leaf to be reached (the goal) is chosen uniformly randomly before each episode,

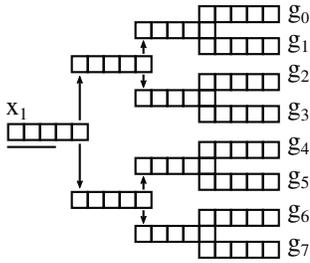


Figure 6: TreeMaze environment. The agent starts at x_1 and must go to one of the leaves. The leaf to be reached is indicated by 3 bits observed at time-steps 1, 2 and 3.

and is indicated to the agent using 3 bits, observed one at a time during the first 3 time-steps. The agent receives no bit afterwards, and has to remember them in order to navigate to the goal. The agent observes its position in the current corridor (0 to 4) and the number of T junctions it has already crossed (0 to 3). A reward of -0.1 is given each time-step, $+10$ when reaching the goal. The episode finishes when the agent reaches any of the leaves. The optimal reward is 8.2.

We consider 14 options with predefined memoryless policies, several of them sharing the same policy, but encoding distinct states (among 14) of a 3-bit memory where some bits may be unknown. 6 partial-knowledge options $\omega_{0--}, \omega_{1--}, \omega_{00-}, \dots, \omega_{11-}$ go right then terminate. 8 full-knowledge options $\omega_{000}, \omega_{001}, \dots, \omega_{111}$ go to their corresponding leaf. OOI are defined so that any option may only be followed by itself, or one that represents a memory state where a single 0 or - has been flipped to 1. Five agents have to learn their top-level policy, which requires them to learn how to use the available options to remember to which leaf to go. The agents do not know the name or meaning of the options. Three agents have access to all 14 options (with, without OOIs, and LSTM). The agent with OOIs (8) only has access to full-knowledge options, and therefore cannot disambiguate unknown and 0 bits. The agent with OOIs (4) is restricted to options $\omega_{000}, \omega_{010}, \omega_{100}$ and ω_{110} and therefore cannot reach odd-numbered goals. The options of the (8) and (4) agents terminate in the first two cells of the first corridor, to allow the top-level policy to observe the second and third bits.

Figure 7 shows that the agent with OOIs (14) consistently learns the optimal policy for this task. When the number of options is reduced, the quality of the resulting policies decreases, while still remaining above the agent without OOIs. Even the agent with 4 options, that cannot reach half the goals, performs better than the agent without OOIs but 14 options. This experiment demonstrates that OOIs provide measurable benefits over standard initiation sets, even if the option set is largely reduced.

Combined, our three experiments demonstrate that OOIs lead to optimal policies in challenging POMDPs, consistently outperform LSTM over options, allow the option policies to be learned, and can still be used when reduced or no domain knowledge is available.

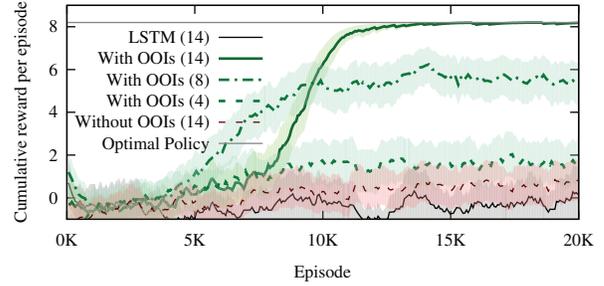


Figure 7: Cumulative reward per episode obtained on TreeMaze, using 14, 8 or 4 options. Even with an insufficient amount of options (8 or 4), OOIs lead to better performance than no OOIs but 14 options. LSTM over options learns the task after more than 100K episodes (not shown).

5 Conclusion and Future Work

This paper proposes OOIs, an extension of the initiation sets of options so that they restrict which options are allowed to be executed after one terminates. This makes options as expressive as Finite State Controllers. Experimental results confirm that challenging partially observable tasks, simulated or on physical robots, one of them requiring exact information storage for hundreds of time-steps, can now be solved using options. Our experiments also illustrate how OOIs lead to reasonably good policies when the option set is improperly defined. Furthermore, we show that learning the top-level and option policies in parallel allow random OOIs to be used, thereby providing a solution to partial observable problems without requiring engineering work.

Options with OOIs also perform surprisingly well compared to an LSTM network over options. While LSTM over options does not require the design of OOIs, their ability to learn without any a-priori knowledge comes at the cost of sample efficiency and explainability. Furthermore, random OOIs are as easy to use as an LSTM and lead to superior results (see Section 4.4). OOIs therefore provide a compelling alternative to recurrent neural networks over options, applicable to a wide range of problems, able to use domain knowledge when available, but not requiring it.

Finally, the compatibility between OOIs and a large variety of reinforcement learning algorithms leads to many future research opportunities. For instance, we have obtained very encouraging results with CACLA (Van Hasselt et al. 2007), in continuous-action hierarchical POMDPs, and we are working on extending OOIs to multi-objective POMDPs (Roijers et al. 2015).

Acknowledgments

The first author is “Aspirant” with the Research Foundation - Flanders (FWO, Belgium), grant number 1129317N. The second author is “Postdoctoral Fellow” with the FWO, grant number 12J0617N. Thanks to Joris Scharpf for his very helpful input on this paper. Thanks to Finn Lattimore, who gave a computer to the first author, when his died while attending UAI in Australia.

References

- Angeline, P. J.; Saunders, G. M.; and Pollack, J. B. 1994. An evolutionary algorithm that constructs recurrent neural networks. *IEEE Transactions on Neural Networks* 5(1):54–65.
- Bacon, P.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 1726–1734.
- Bakker, B. 2001. Reinforcement learning with long short-term memory. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14.
- Barto, A. G., and Mahadevan, S. 2003. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems: Theory and Applications* 13(12):341–379.
- Boots, B.; Siddiqi, S. M.; and Gordon, G. J. 2011. Closing the learning-planning loop with predictive state representations. *The International Journal of Robotics Research* 30(7):954–966.
- Cassandra, A.; Kaelbling, L.; and Littman, M. 1994. Acting optimally in partially observable stochastic domains. In *Proceedings of the 12th AAAI National Conference on Artificial Intelligence*, volume 2, 1023–1028.
- Graves, A.; Wayne, G.; Reynolds, M.; Harley, T.; Danihelka, I.; Grabska-Barwińska, A.; Gómez Colmenarejo, S.; Grefenstette, E.; Ramalho, T.; Agapiou, J.; Badia, A. P.; Moritz Hermann, K.; Zwols, Y.; Ostrovski, G.; Cain, A.; King, H.; Summerfield, C.; Blunsom, P.; Kavukcuoglu, K.; and Hassabis, D. 2016. Hybrid computing using a neural network with dynamic external memory. *Nature* 538(7626):471–476.
- He, R.; Brunskill, E.; and Roy, N. 2011. Efficient planning under uncertainty with macro-actions. *Journal of Artificial Intelligence Research (JAIR)* 40:523–570.
- Hochreiter, S., and Schmidhuber, J. J. 1997. Long short-term memory. *Neural Computation* 9(8):1–32.
- Józefowicz, R.; Zaremba, W.; and Sutskever, I. 2015. An empirical exploration of recurrent network architectures. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, 2342–2350.
- Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101:99–134.
- Lim, Z. W.; Hsu, D.; and Lee, W. S. 2011. Monte carlo value iteration with macro-actions. In *Advances in Neural Information Processing Systems (NIPS)*, 1287–1295.
- Lin, L.-J., and Mitchell, T. M. 1992. *Memory approaches to reinforcement learning in non-Markovian domains*. Carnegie-Mellon University. Department of Computer Science.
- Lin, L.-J., and Mitchell, T. M. 1993. Reinforcement learning with hidden states. *From Animals to Animats* 2:271–280.
- Littman, M. L.; Sutton, R. S.; and Singh, S. 2001. Predictive Representations of State. In *Advances in Neural Information Processing Systems (NIPS)*, volume 14, 1555–1561.
- Meuleau, N.; Peshkin, L.; Kim, K.-e.; and Kaelbling, L. P. 1999. Learning Finite-State Controllers for partially observable environments. In *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI)*, 427–436.
- Mikolov, T.; Joulin, A.; Chopra, S.; Mathieu, M.; and Ranzato, M. 2014. Learning longer memory in recurrent neural networks. *CoRR* abs/1412.7753.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T. P.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous Methods for Deep Reinforcement Learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, volume 48, 1928–1937.
- Omidshafiei, S.; Agha-Mohammadi, A.; Amato, C.; Liu, S.; How, J. P.; and Vian, J. 2017. Decentralized control of multi-robot partially observable markov decision processes using belief space macro-actions. *International Journal of Robotics Research* 36(2):231–258.
- Peshkin, L.; Meuleau, N.; and Kaelbling, L. 1999. Learning policies with external memory. In *Proceedings of the 16th International Conference on Machine Learning (ICML)*, 307–314.
- Precup, D. 2000. *Temporal Abstraction in Reinforcement Learning*. Ph.D. Dissertation, University of Massachusetts.
- Roijsers, D. M.; Whiteson, S.; and Oliehoek, F. A. 2015. Point-based planning for multi-objective POMDPs. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence IJCAI 2015*, 1666–1672.
- Roy, N.; Gordon, G.; and Thrun, S. 2006. Planning under uncertainty for reliable health care robotics. *Springer Tracts in Advanced Robotics* 24:417–426.
- Sridharan, M.; Wyatt, J.; and Dearden, R. 2010. Planning to see: A hierarchical approach to planning visual actions on a robot using POMDPs. *Artificial Intelligence* 174:704–725.
- Sutton, R.; Precup, D.; and Singh, S. 1999. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112:181–211.
- Sutton, R.; McAllester, D.; Singh, S.; and Mansour, Y. 2000. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13.
- Tessler, C.; Givony, S.; Zahavy, T.; Mankowitz, D. J.; and Mannor, S. 2016. A Deep Hierarchical Approach to Long-term Learning in Minecraft. In *13th European Workshop on Reinforcement Learning (EWRL)*.
- Theocharous, G. 2002. *Hierarchical learning and planning in partially observable Markov decision processes*. Ph.D. Dissertation, Michigan State University.
- Van Hasselt, H., and Wiering, M. A. 2007. Reinforcement learning in continuous action spaces. *Proceedings of the 2007 IEEE Symposium on Approximate Dynamic Programming and Reinforcement Learning* 272–279.
- Wiering, M., and Schmidhuber, J. 1997. HQ-Learning. *Adaptive Behavior* 6(2):219–246.
- Zaremba, W., and Sutskever, I. 2015. Reinforcement Learning Neural Turing Machines. *CoRR* abs/1505.00521.