
Finding Options that Minimize Planning Time

Yuu Jinnai¹ David Abel¹ D Ellis Hershkowitz² Michael L. Littman¹ George Konidaris¹

Abstract

We formalize the problem of selecting the optimal set of options for planning as that of computing the smallest set of options so that planning converges in less than a given maximum of value-iteration passes. We first show that the problem is *NP*-hard, even if the task is constrained to be deterministic—the first such complexity result for option discovery. We then present the first polynomial-time boundedly suboptimal approximation algorithm for this setting, and empirically evaluate it against both the optimal options and a representative collection of heuristic approaches in simple grid-based domains.

1. Introduction

Markov Decision Processes or MDPs (Puterman, 1994) are a widely used expressive model of sequential decision-making. However, MDPs are computationally expensive to solve (Papadimitriou & Tsitsiklis, 1987; Littman, 1997; Goldsmith et al., 1997). One approach to solving such problems is to add high-level, temporally extended actions—often formalized as options (Sutton et al., 1999)—to the set of actions available to the agent. The right set of options allows planning to probe more deeply into the search space with a single computation. Thus, if options are chosen appropriately, planning algorithms can find good plans with less computation.

Indeed, previous work has offered substantial support that abstract actions can accelerate planning (Mann & Mannor, 2014; Silver & Ciosek, 2012). However, little is known about how to find the right set of options for planning. Prior work often seeks to codify an intuitive notion of what underlies an effective option, such as identifying relatively novel states (Şimşek & Barto, 2004), identifying bottleneck states or high-betweenness states (Şimşek et al., 2005; Şimşek &

Barto, 2009; Bacon, 2013; Moradi et al., 2012), finding repeated policy fragments (Pickett & Barto, 2002), or finding states that often occur on successful trajectories (McGovern & Barto, 2001; Bakker & Schmidhuber, 2004). While such intuitions often capture important aspects of the role of options in planning, the resulting algorithms are somewhat heuristic in that they are not based on optimizing any precise performance-related metric; consequently, their relative performance can only be evaluated empirically.

We aim to formalize what it means to find the set of options that is optimal for planning, and to use the resulting formalization to develop an algorithm with performance guarantees and a principled theoretical foundation. Specifically, we consider the problem of finding the smallest set of options so that planning converges in fewer than ℓ value iterations (VI). Our main result is that this problem is *NP*-hard. More precisely, the problem:

1. is $2^{\log^{1-\epsilon} n}$ -hard to approximate for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{poly} \log n})$,¹ where n is the input size;
2. is $\Omega(\log n)$ -hard to approximate even for deterministic MDPs unless $P = NP$;
3. has an $O(n)$ -approximation algorithm;
4. has an $O(\log n)$ -approximation algorithm for deterministic MDPs.

In Section 4, we introduce A-MOMI, a polynomial-time approximation algorithm that has $O(n)$ suboptimality in general and $O(\log n)$ suboptimality for deterministic MDPs. The expression $2^{\log^{1-\epsilon} n}$ is only slightly smaller than n : if $\epsilon = 0$ then $\Omega(2^{\log n}) = \Omega(n)$. Thus, the inapproximability results claim that A-MOMI is close to the best possible approximation factor.

In addition, we consider the complementary problem of finding a set of k options that minimize the number of VI iterations until convergence. We show that this problem is also *NP*-hard, even for a deterministic MDP and introduce A-MIMO, a polynomial time approximation algorithm.

¹This is a standard complexity assumption: See, for example, Diniz et al. (2012).

¹Brown University, Providence, RI, United States ²Carnegie Mellon University, Pittsburgh, PA, United States. Correspondence to: Yuu Jinnai <yuu.jinnai@brown.edu>.

Finally, we empirically evaluate the performance of two heuristic approaches for option discovery, betweenness options (Şimşek & Barto, 2009) and Eigenoptions (Machado et al., 2017), against the proposed approximation algorithms and the optimal options in standard grid domains.

2. Background

We first provide background on Markov Decision Processes (MDPs), planning, and options.

2.1. Markov Decision Processes and Planning

An MDP is a five tuple: $\langle \mathcal{S}, \mathcal{A}, R, T, \gamma \rangle$, where \mathcal{S} is a finite set of states; \mathcal{A} is a finite set of actions; $R : \mathcal{S} \times \mathcal{A} \rightarrow [0, \text{RMAX}]$ is a reward function; $T : \mathcal{S} \times \mathcal{A} \rightarrow \text{Pr}(\mathcal{S})$ is a transition function, denoting the probability of arriving in state $s' \in \mathcal{S}$ after executing action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$; and $\gamma \in [0, 1]$ is a discount factor, expressing the agent’s preference for immediate over delayed rewards.

An action-selection strategy is modeled by a *policy*, $\pi : \mathcal{S} \rightarrow \text{Pr}(\mathcal{A})$, mapping states to a distribution over actions. Typically, the goal of planning in an MDP is to *solve* the MDP—that is, to compute an optimal policy. A policy π is evaluated according to the Bellman equation, denoting the long term expected reward received by executing π :

$$V^\pi(s) = R(s, \pi(s)) + \gamma \sum_{s' \in \mathcal{S}} T(s, \pi(s), s') V^\pi(s'). \quad (1)$$

We denote $\pi^*(s) = \arg \max_\pi V^\pi(s)$ and $V^*(s) = \max_\pi V^\pi(s)$ as the optimal policy and value function, respectively.

The core problem we study is *planning*, namely, computing a near optimal policy for a given MDP. The main variant of the planning problem we study is the *value-planning problem*:

Definition 1 (Value-Planning Problem): **Given** an MDP $M = \langle \mathcal{S}, \mathcal{A}, R, T, \gamma \rangle$ and a non-negative real-value ϵ , **return** a value function, V such that $|V(s) - V^*(s)| < \epsilon$ for all $s \in \mathcal{S}$.

The value-planning problem can be solved in time polynomial in the size of the state space (Littman et al., 1995).

2.2. Options and Value Iteration

Temporally extended actions offer great potential for mitigating the difficulty of solving complex MDPs, either through planning or reinforcement learning (Sutton et al., 1999). However, it is possible that options that are useful for learning are not necessarily useful for planning, and vice versa. In fact, we don’t have an explicit metric for measuring the

quality of an option set for planning. Therefore, identifying techniques that produce good options in these scenarios is an important open problem in the literature.

We use the standard definition of options (Sutton et al., 1999):

Definition 2 (option): An option o is defined by a triple: $(\mathcal{I}, \pi, \beta)$ where:

- $\mathcal{I} \subseteq \mathcal{S}$ is a set of states where the option can initiate,
- $\pi : \mathcal{S} \rightarrow \text{Pr}(\mathcal{A})$ is a policy,
- $\beta : \mathcal{S} \rightarrow [0, 1]$, is a termination condition.

We let \mathcal{O}_{all} denote the set containing all options.

In planning, options have a well defined transition and reward model for each state named the multi-time model, introduced by Precup & Sutton (1998):

$$T_\gamma(s, o, s') = \sum_{t=0}^{\infty} \gamma^t \text{Pr}(s_t = s', \beta(s_t) | s, o). \quad (2)$$

$$R_\gamma(s, o) = \mathbb{E}_{o_\pi} \left[r_1 + \gamma r_2 + \dots + \gamma^{k-1} r_k \mid s, o \right]. \quad (3)$$

We use the multi-time model for value iteration. The algorithm computes a sequence of functions V_0, V_1, \dots, V_b using the Bellman optimality operator on the multi-time model:

$$V_{i+1}(s) = \max_{o \in \mathcal{A} \cup \mathcal{O}} \left(R_\gamma(s, o) + \sum_{s' \in \mathcal{S}} T_\gamma(s, o, s') V_i(s') \right). \quad (4)$$

The problem we consider is to find a set of options to add to the set of primitive actions that minimize the number of iterations required for VI to converge:²

Definition 3 ($L_{\epsilon, V_0}(\mathcal{O})$): The number of iterations $L_{\epsilon, V_0}(\mathcal{O})$ of VI using the joint action set $\mathcal{A} \cup \mathcal{O}$, with \mathcal{O} a non-empty set of options, is the smallest b at which $|V_{b'}(s) - V^*(s)| < \epsilon$ for all $s \in \mathcal{S}$, $b' \geq b$.

2.2.1. POINT OPTIONS

The options formalism is immensely general. Due to its generality, a single option can actually encode several completely unrelated sets of different behaviors. Consider the nine-state example MDP pictured in Figure 1; a single option can in fact initiate, make decisions in, and terminate along entirely independent trajectories. As we consider more complex MDPs (which, as discussed earlier, is often a motivation for introducing options), the number of inde-

²We can ensure $|V^*(s) - V_i(s)| < \epsilon$ by running VI until $|V_{i+1}(s) - V_i(s)| < \epsilon(1 - \gamma)/2\gamma$ for all $s \in \mathcal{S}$ (Williams & Baird, 1993).




Figure 1: A single option can encode multiple unrelated behaviors. The dark circles indicate where the option can be initiated (s_1 & s_6) and terminated (s_2 & s_9), whereas the lighter circles denote the states visited by the option policy when applied in the respective initiating state.

pendent behaviors that can be encoded by a single option increases further still.

As a result, it can be difficult to reason about the impact of adding a single option, in the traditional sense. As the MDP grows larger, a combinatorial number of different behaviors can emerge from “one” option. Consequently, it is difficult to address the question: which *single* option helps planning the most? As MDPs grow large, one option can encode a large number of possible, independent behaviors. Thus, we instead introduce and study “point options”, which only allow for a single continuous stream of behavior:

Definition 4 (Point option): A **point option** is any option whose initiation set and termination set are each true for exactly one state each:

$$|\{s \in \mathcal{S} : \mathcal{I}(s) = 1\}| = 1, \quad (5)$$

$$|\{s \in \mathcal{S} : \beta(s) > 0\}| = 1, \quad (6)$$

$$|\{s \in \mathcal{S} : \beta(s) = 1\}| = 1. \quad (7)$$

We let \mathcal{O}_p denote the set containing all point options.

For simplicity, we denote the initiation state as \mathcal{I}_o and the termination state as β_o for a point option o .

To plan with a point option from state s , the agent runs value iteration using a model (Eq. 2, 3) in addition to the backup operations by primitive actions where k is the duration of the option. We assume that the model of each option is given to the agent and ignore the computation cost for computing the model for the options.

Point options are a useful subclass to consider for several reasons. First, a point option is a simple model for a temporally extended action. Second, the policy of the point option can be calculated as a path-planning problem for deterministic MDPs. Third, any other options with a single termination state with termination probability 1 can be represented as

a collection of point options. Fourth, a point option has constant computational overhead per iteration.

3. Complexity Results

Our main results focus on two computational problems:

1. **MINOPTIONMAXITER (MOMI):** Which set of options lets value iteration converge in at most ℓ iterations?
2. **MINITERMAXOPTION (MIMO):** Which set of k or fewer options minimizes the number of iterations to convergence?

More formally, MOMI is defined as follows.

Definition 5 (MOMI): The **MINOPTIONMAXITER** problem:

Given an MDP M , a non-negative real-value ϵ , an initial value function V_0 , and an integer ℓ **return** \mathcal{O} that minimizes $|\mathcal{O}|$ subject to $\mathcal{O} \subseteq \mathcal{O}_p$ and $L_{\epsilon, V_0}(\mathcal{O}) \leq \ell$.

We then consider the complementary optimization problem **MINITERMAXOPTION (MIMO)**: compute a set of k options which minimizes the number of iterations:

Definition 6 (MIMO): The **MINITERMAXOPTION** problem:

Given an MDP M , a non-negative real-value ϵ , an initial value function V_0 , and an integer k **return** \mathcal{O} that minimizes $L_{\epsilon, V_0}(\mathcal{O})$, subject to $\mathcal{O} \subseteq \mathcal{O}_p$ and $|\mathcal{O}| \leq k$.

We now introduce our main result, which shows that both MOMI and MIMO are *NP-hard*.

Theorem 1. *MOMI and MIMO are NP-hard.*

Proof. We consider a problem OI-DEC which is a decision version of MOMI and MIMO. The problem asks if we can solve the MDP within ℓ iterations using at most k point options.

Definition 7 (OI-DEC):

Given an MDP M , a non-negative real-value ϵ , an initial value function V_0 , and integers k and ℓ , **return** ‘Yes’ if there exists an option set \mathcal{O} such that $\mathcal{O} \subseteq \mathcal{O}_p$, $|\mathcal{O}| \leq k$ and $L_{\epsilon, V_0}(\mathcal{O}) \leq \ell$. ‘No’ otherwise.

We prove the theorem by reduction from the decision version of the set-cover problem—known to be NP-complete—to OI-DEC. The set-cover problem is defined as follows.

Definition 8 (SetCover-DEC):

Given a set of elements \mathcal{U} , a set of subsets $\mathcal{X} = \{X \subseteq \mathcal{U}\}$, and an integer k , **return** ‘Yes’ if there exists a




Figure 2: Reduction from SetCover-DEC to OI-DEC. The example shows the reduction from an instance of SetCover-DEC which asks if we can pick two subsets from $\mathcal{X} = \{X_1, X_2\}$ where $X_1 = \{1, 2, 3\}$, $X_2 = \{3, 4, 5\}$ to cover all elements $\mathcal{U} = \{1, 2, 3, 4, 5\}$. The SetCover-DEC can be reduced to an instance of OI-DEC where the question is whether the MDP can be solved with 2 iterations of VI by adding at most two point options. The answer of OI-DEC is ‘Yes’ (adding point options from X_1 and X_2 to g will solve the problem), thus the answer of the SetCover-DEC is ‘Yes’. Here the set of initial states corresponds to the cover for the SetCover-DEC.

cover $\mathcal{C} \subseteq \mathcal{X}$ such that $\bigcup_{X \in \mathcal{C}} X = \mathcal{U}$ and $|\mathcal{C}| \leq k$.
‘No’ otherwise.

If there is some $u \in \mathcal{U}$ that is not included in at least one of the subsets X , then the answer is ‘No’. Assuming otherwise, we construct an instance of a shortest path problem (a special case of an MDP problem) as follows (Figure 2). There are four types of states in the MDP: (1) $u_i \in \mathcal{U}$ represents one of the elements in \mathcal{U} , (2) $X_i \in \mathcal{X}$ represents one of the subsets in \mathcal{X} , (3) $X'_i \in \mathcal{X}'$: we make a copy for every state $X_i \in \mathcal{X}$ and call them X'_i , (4) a goal state g . Thus, the state set is $\mathcal{U} \cup \mathcal{X} \cup \mathcal{X}' \cup \{g\}$. We build edges between states as follows: (1) $e(u, X) \in E$ iff $u \in X$: For $u \in \mathcal{U}$ and $X \in \mathcal{X}$, there is an edge between u and X . (2) $\forall X_i \in \mathcal{X}$, $e(X_i, X'_i) \in E$: For every $X_i \in \mathcal{X}$, we have an edge from X_i to X'_i . (3) $\forall e(X', g) \in E$: for every $X' \in \mathcal{X}'$ we have an edge from X' to the goal g . This construction can be done in polynomial time.

Let M be the MDP constructed in this way. We show that $\text{SetCover}(\mathcal{U}, \mathcal{X}, k) = \text{OI-DEC}(M, V_0 = 0, k, 2)$. Note that by construction every state X_i , X'_i , and g converges to its optimal value within 2 iterations as it reaches the goal state g within 2 steps. A state $u \in \mathcal{U}$ converges within 2 steps if and only if there exists a point option (a) from X to g where $u \in X$, (b) from u to X' where $u \in X$, or (c) from u to g . For options of type (b) and (c), we can find an option of type (a) that makes u converge within 2 steps by setting the initial state of the option to $\mathcal{I}_o = X$, where $u \in X$, and the termination state to $\beta_o = g$. Let \mathcal{O} be the solution of $\text{OI-DEC}(M, k, 2)$. If there exists an option of type (b) or

(c), we can swap them with an option of type (a) and still maintain a solution. Thus, it is sufficient to consider that every option is type (a). Let \mathcal{C} be a set of initial states of each option in \mathcal{O} ($\mathcal{C} = \{\mathcal{I}_o | o \in \mathcal{O}\}$). This construction exactly matches the solution of the SetCover-DEC. \square

3.1. Generalizations of MOMI and MIMO

A natural question is whether Theorem 1 extends to more general option-construction settings. We consider two possible extensions, which we believe offer significant coverage of finding optimal options for planning in general.

We first consider the case where the options are not necessarily point options. There is little sense in considering MOMI where one can choose any option since clearly the best option is the option whose policy is the optimal policy. Thus, using the space of all options \mathcal{O}_{all} we generalize MOMI as follows (MIMO_{gen} are defined analogously):

Definition 9 (MOMI_{gen}):

Given an MDP M , a non-negative real-value ϵ , an initial value function V_0 , $\mathcal{O}' \subseteq \mathcal{O}_{all}$, and an integer ℓ , **return** \mathcal{O} minimizing $|\mathcal{O}|$ subject to $L_{\epsilon, V_0}(\mathcal{O}) \leq \ell$ and $\mathcal{O} \subseteq \mathcal{O}'$.

Theorem 2. MOMI_{gen} and MIMO_{gen} are NP-hard.

The proof follows from the fact that MOMI_{gen} is a superset of MOMI and MIMO_{gen} is a superset of MIMO.

We next consider the multi-task generalization, where we aim to find a smallest number of options which the expected number of iterations to solve a problem M sampled from a distribution of MDPs, D , is bounded:

Definition 10 (MOMI_{multi}):

Given A distribution of MDPs D , $\mathcal{O}' \subseteq \mathcal{O}_{all}$, a non-negative real-value ϵ , an initial value function V_0 , and an integer ℓ , **return** \mathcal{O} that minimizes $|\mathcal{O}|$ such that $E_{M \sim D}[L_M(\mathcal{O})] \leq \ell$ and $\mathcal{O} \subseteq \mathcal{O}'$.

Theorem 3. MOMI_{multi} and MIMO_{multi} are NP-hard.

The proof follows from the fact that MOMI_{multi} is a superset of MOMI_{gen} and MIMO_{multi} is a superset of MIMO_{gen} .

In light of the computational difficulty of both problems, the appropriate approach is to find tractable approximation algorithms. However, even approximately solving MOMI is hard. More precisely:

Theorem 4.

1. MOMI is $\Omega(\log n)$ hard to approximate even for deterministic MDPs unless $P = NP$.
2. MOMI_{gen} is $2^{\log^{1-\epsilon} n}$ -hard to approximate for any $\epsilon > 0$ even for deterministic MDPs unless $NP \subseteq \text{DTIME}(n^{\text{poly} \log n})$.

3. MOMI is $2^{\log^{1-\epsilon} n}$ -hard to approximate for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{poly} \log n})$.

Proof. See appendix. \square

Note that an $O(n)$ -approximation is achievable by the trivial algorithm that returns a set of all candidate options. Thus, Theorem 4 roughly states that there is no polynomial time approximation algorithms other than the trivial algorithm for MOMI.

In the next section we show that an $O(\log n)$ -approximation is achievable for MOMI if the MDP is deterministic and the agent is given a set of all point options. Thus, together, these two results give a formal separation between the hardness of abstraction in MDPs with and without stochasticity.

In summary, *the problem of computing optimal behavioral abstractions for planning is computationally intractable.*

4. Approximation Algorithms

We now provide polynomial-time approximation algorithms, A-MIMO and A-MOMI, to solve MOMI and MIMO, respectively. Both algorithms have bounded suboptimality slightly worse than a constant factor for deterministic MDPs.

We assume that (1) there is exactly one absorbing state $g \in \mathcal{S}$ with $T(g, a, g) = 1$ and $R(g, a) = 0$, and every optimal policy eventually reaches g with probability 1, (2) there is no cycle with a positive reward involved in the optimal policy’s trajectory. That is, $V_+^\pi(s) := \mathbb{E}[\sum_{t=0}^{\infty} \max\{0, R(s_t, a_t)\}] < \infty$ for all policies π . Note that we can convert a problem with multiple goals to a problem with a single goal by adding a new absorbing state g to the MDP and adding a transition from each of the original goals to g .

Unfortunately, these algorithms are computationally harder than solving the MDP itself, and are thus not practical for planning. Instead, they are useful for analyzing and evaluating heuristically generated options. If the option set generated by the heuristic methods outperforms the option set found by the following algorithms, then one can claim that the option set found by the heuristic is close to the optimal option set (for that MDP). Our algorithms have a formal guarantee on bounded suboptimality if the MDP is deterministic, so any heuristic method that provably exceeds our algorithm’s performance will also guarantee bounded suboptimality. We also believe these algorithms may be a useful foundation for future option discovery methods.

4.1. A-MOMI

We now describe a polynomial-time approximation algorithm, A-MOMI, based on using set cover to solve MOMI.

The overview of the procedure is as follows.

1. Compute an asymmetric distance function $d_\epsilon(s, s') : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{N}$ representing the number of iterations for a state s to reach its ϵ -optimal value if we add a point option from a state s' to a goal state g .
2. For every state s_i , compute a set of states X_{s_i} within $\ell - 1$ distance of reaching s_i . The set X_{s_i} represents the states that converge within ℓ steps if we add a point option from s_i to g .
3. Let \mathcal{X} be a set of X_{s_i} for every $s_i \in \mathcal{S} \setminus X_g^+$, where X_g^+ is a set of states that converges within ℓ without any options (thus can be ignored).
4. Solve the set-cover optimization problem to find a set of subsets that covers the entire state set using the approximation algorithm by Chvatal (1979). This process corresponds to finding a minimum set of subsets $\{X_{s_i}\}$ that makes every state in \mathcal{S} converge within ℓ steps.
5. Generate a set of point options with initiation states set to one of the center states in the solution of the set-cover, and termination states set to the goal.

We compute a distance function $d_\epsilon : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{N}^3$, defined as follows:

Definition 11 (Distance $d_\epsilon(s_i, s_j)$): $d_\epsilon(s_i, s_j)$ is the smallest number b such that for all $b' \geq b$, $V_{b'}^\pi(s_i)$ is ϵ -optimal if we add a point option from s_j to g , minus one.

More formally, let $d'_\epsilon(s_i)$ denote the number of iterations needed for the value of state s_i to satisfy $|V(s_i) - V^*(s_i)| < \epsilon$, and let $d'_\epsilon(s_i, s_j)$ be an upper bound of the number of iterations needed for the value of s_i to satisfy $|V(s_i) - V^*(s_i)| < \epsilon$, if the value of s_j is initialized such that $|V(s_j) - V^*(s_j)| < \epsilon$. We define $d_\epsilon(s_i, s_j) := \min(d'_\epsilon(s_i) - 1, d'_\epsilon(s_i, s_j))$. For simplicity, we use d to denote the function d_ϵ . Consider the following example.

Example. Table 3b is a distance function for the MDP shown in Figure 3a. For a deterministic MDP, $d_0(s)$ corresponds to the number of edge traversals from state s to g , where we have edges only for those that corresponds to the state transition by the optimal actions. The quantity $d_0(s, s') - 1$ is the minimum of $d_0(s)$ and one plus the number of edge traversals from s to s' . \diamond

Note that we only need to solve the MDP once to compute d . $d(s, s')$ can be computed once you solved the MDP without any options and store all value functions V_i ($i = 1, \dots, b$) until convergence as a function of V_1 : $V_i(s) = f(V_1(s_0), V_1(s_1), \dots)$. If we add a point option

³Formally, d satisfies the triangle inequality, but does not satisfy the symmetry and the indiscernibles.




Figure 3: (a) Options discovered by A-MIMO with $k = 2$ are denoted by the dashed lines. (b) $d_0(s, s')$ for Figure 3a

from s' to g , then $V_1(s') = V^*(s')$. Thus, $d(s, s')$ is the smallest i where $V_i(s)$ reaches ϵ -optimal if we replace $V_1(s')$ with $V^*(s')$ when computing $V_i(s)$ as a function of V_1 .

Example. We use the MDP shown in Figure 3a as an example. Consider the problem of finding a set of options so that the MDP can be solved within 2 iterations. We generate an instance of a set-cover optimization problem. The set of elements for the set cover is the set of states of the MDP that do not reach their optimal value within ℓ steps without any options $\mathcal{S} \setminus X_g^+$. Here, we denote a set of nodes that can be solved within ℓ steps by X_g^+ . In the example, $\mathcal{U} = \mathcal{S} \setminus X_g^+ = \{s_1, s_2, s_3, s_4\}$. A state s is included in a subset $X_{s'}$ iff $d(s, s') \leq \ell - 1$. For example, $X_{s_1} = \{s_1\}$, $X_{s_2} = \{s_1, s_2\}$. Thus, the set of subsets are given as: $X_{s_1} = \{s_1\}$, $X_{s_2} = \{s_1, s_2\}$, $X_{s_3} = \{s_3\}$, $X_{s_4} = \{s_3, s_4\}$. In this case, the approximation algorithm finds the optimal solution $\mathcal{C} = \{X_{s_2}, X_{s_4}\}$ for the set-cover optimization problem $(\mathcal{U}, \mathcal{X})$. We generate a point option for each state in \mathcal{C} . Thus, the output of the algorithm is a set of two point options from s_2 and s_4 to g . \diamond

Theorem 5. A-MOMI has the following properties:

1. A-MOMI runs in polynomial time.
2. It guarantees that the MDP is solved within ℓ iterations using the option set acquired by A-MOMI \mathcal{O} .
3. If the MDP is deterministic, the option set is at most $O(\log n)$ times larger than the smallest option set pos-

sible to solve the MDP within ℓ iterations.

Proof. See the supplementary material. \square

Note that the approximation bound for a deterministic MDP will inherit any improvements to the approximation algorithm for set cover. Set cover is known to be *NP*-hard to approximate up to a factor of $(1 - o(1)) \log n$ (Dinur & Steurer, 2014), thus there may be an improvement on the approximation ratio for the set cover problem, which will also improve the approximation ratio of A-MOMI.

4.2. A-MIMO

The outline of the approximation algorithm for MIMO (A-MIMO) is as follows.

1. Compute $d_\epsilon(s, s') : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{N}$ for each pair of states.
2. Using this distance function, solve an asymmetric k -center problem, which finds a set of center states that minimizes the maximum number of iterations for every state to converge.
3. Generate point options with initiation states set to the center states in the solution of the asymmetric k -center, and termination states set to the goal.

As in A-MOMI, we first compute the function d . Then, we exploit this characteristic of d and solve the asymmetric k -center problem (Panigrahy & Vishwanathan, 1998) on (\mathcal{U}, d, k) to get a set of centers, which we use as initiation states for point options. The asymmetric k -center problem is a generalization of the metric k -center problem where the function d obeys the triangle inequality, but is not necessarily symmetric:

Definition 12 (AsymKCenter):

Given a set of elements \mathcal{U} , a function $d : \mathcal{U} \times \mathcal{U} \rightarrow \mathbb{N}$, and an integer k , **return** \mathcal{C} that minimizes $P(\mathcal{C}) = \max_{s \in \mathcal{U}} \min_{c \in \mathcal{C}} d(s, c)$ subject to $|\mathcal{C}| \leq k$.

We solve the problem using a polynomial-time approximation algorithm proposed by Archer (2001). The algorithm has a suboptimality bound of $O(\log^* k)^4$ where $k < |\mathcal{U}|$. It is known that the problem cannot be solved within a factor of $\log^* |\mathcal{U}| - \theta(1)$ unless $P = NP$ (Chuzhoy et al., 2005). As the procedure by Archer (2001) often finds a set of options smaller than k , we generate the rest of the options by greedily adding $\lceil \log k \rceil$ options at once. See the supplementary material for details. Finally, we generate a set of point options with initiation-states set to one of the centers and the termination state set to the goal state of the MDP. That is, for every c in \mathcal{C} , we generate a point option starting from c to the goal state g .

⁴ \log^* is the number of times the logarithm function must be iteratively applied before the result is less than or equal to 1.

Example. Consider an MDP shown in Figure 3a. The distance d_0 for the MDP is shown in Table 3b. Note that $d(s, s') \leq d(s, g)$ holds for every s, s' pair. Let us first consider finding one option ($k = 1$). This process corresponds to finding a column with the smallest maximum value in the Table 3b. The optimal point option is from s_5 to g as it has the smallest maximum value in the column. If $k = 2$, an optimal set of options is from s_2 and s_4 to g . Note that the optimal option for $k = 1$ is not in the optimal option set of size 2. This example shows that the strategy of greedily adding options does not find the optimal set. In fact, the improvement $L_{\epsilon, V_0}(\emptyset) - L_{\epsilon, V_0}(\mathcal{O})$ on by the greedy algorithm can be arbitrary small (i.e. 0) compared to the optimal option (see Proposition 1 in the supplementary material for a proof). \diamond

Theorem 6. *A-MIMO has the following properties:*

1. *A-MIMO runs in polynomial time.*
2. *If the MDP is deterministic, it has a bounded suboptimality of $O(\log^* k)$.*
3. *The number of iterations to solve the MDP using the option set acquired is upper bounded by $P(\mathcal{C})$.*

Proof. See the supplementary material. \square

5. Experiments

We evaluate the performance of the value-iteration algorithm using options generated by the approximation algorithms on several grid-based simple domains.

We ran the experiments on an 11×11 four-room domain and a 9×9 grid world with no walls. In both domains, the agent’s goal is to reach a specific square. The agent can move in the usual four directions but cannot cross walls.

Visualizations: First, we visualize a variety of option types, including the optimal point options, those found by our approximation algorithms, and several option types proposed in the literature. We computed the optimal set of point options by enumerating every possible set of point options and picking the best. As an optimal set of options is not unique, we picked one arbitrarily. We are only able to find optimal solutions up to four options within 10 minutes, while the approximation algorithm could find any number of options within a few minutes. For eigenoptions, we ignored the eigenvector corresponding to the smallest eigenvalue ($\lambda_0 = 0$) in the graph Laplacian because it has a constant value for every state. Both betweenness options and eigenoptions are polynomial time algorithm, thus run in a few minutes. Figure 4 shows the optimal and bounded suboptimal set of options computed by A-MIMO. See the supplementary material for visualizations for the 9×9 grid domain.

Figure 4e shows the four bottleneck states with highest shortest-path betweenness centrality in the state-transition graph (Şimşek & Barto, 2009). Interestingly, the optimal options are quite close to the bottleneck states in the four-room domain, suggesting that bottleneck states are also useful for planning as a heuristic to find important subgoals.

Figure 4f shows the set of subgoals discovered by graph Laplacian analysis following Machado et al. (2017). While they proposed to generate options to travel between subgoals for reinforcement learning, we generate a set of point options from each subgoal to the goal state as that is a better in the planning setting.

Quantitative Evaluation: Next, we run value iteration using the set of options generated by A-MIMO and A-MOMI. Figures 5a and 5b show the number of iterations on the four-room and the 9×9 grids using k options. The experimental results suggest that the suboptimal algorithm finds set of options similar to, but not quite as good as, the optimal ones. For betweenness options and eigenoptions, we evaluated every subset of options among the four and present results for the best subset found. Because betweenness options are placed close to the optimal options, the performance is close to optimal especially when the number of options is small.

In addition, we used A-MOMI to find a minimum option set to solve the MDP within the given number of iterations. Figures 5c and 5d show the number of options generated by A-MOMI compared to the minimum number of options.

6. Related Work

Many heuristic algorithms have proposed to discover options (Iba, 1989; McGovern & Barto, 2001; Menache et al., 2002; Stolle & Precup, 2002; Şimşek & Barto, 2004; Şimşek & Barto, 2009; Konidaris & Barto, 2009; Machado et al., 2017; Eysenbach et al., 2019). For example, some investigate the use of bottleneck states (Stolle & Precup, 2002; Şimşek & Barto, 2009; Menache et al., 2002; Lehnert et al., 2018). Stolle & Precup (2002) proposed to set states with high visitation counts as subgoal states, which identifies bottleneck states in the four-room domain. Şimşek & Barto (2009) generalized the concept of a bottleneck to (shortest-path) betweenness of the graph to capture how pivotal the state is. Menache et al. (2002) used a learned model of the environment to run a Max-Flow/Min-Cut algorithm to the state-space graph to identify bottleneck states. These methods generate options that leverage the idea that subgoals are states visited most frequently. On the other hand, Şimşek & Barto (2004) proposed to generate options to encourage exploration by generating options to relatively novel states. Eysenbach et al. (2019) instead proposed learning a policy for each option so that the diversity of the trajectories by the set of options are maximized. These methods

Finding Options that Minimize Planning Time




Figure 4: Comparison of the optimal point options with options generated by the approximation algorithm A-MIMO. The green square represents the termination state and the blue squares the initiation states. Observe that the approximation algorithm is similar to that of optimal options. Note that the optimal option set is not unique: there can be multiple optimal option sets, and we visualize the one returned by the solver.




Figure 5: MIMO and MOMI evaluations. Parts (a)–(b) show the number of iterations for VI using options generated by A-MIMO. Parts (c)–(d) show the number of options generated by A-MOMI to ensure the MDP is solved within a given number of iterations. OPT: optimal set of options. APPROX: a bounded suboptimal set of options generated by A-MIMO an A-MOMI. BET: betweenness options. EIG: eigenoptions.

generate options to explore infrequently visited states.

While empirical results show that these algorithms are useful in some scenarios, the conditions under which the methods are effective is often unclear because the relationship between the objective of the skill discovery algorithm and that of the agent is often not established. In fact, Jong et al. (2008) sought to investigate the utility of skills empirically and pointed out that introducing skills might worsen the learning performance. Harb et al. (2017) proposed to formulate good options to be options which minimize the deliberation costs in the bounded rationality framework (Simon, 1957). Brunskill & Li (2014) targeted the lifelong reinforcement learning setting and proposed an option generation method for lifelong reinforcement learning. They analyzed the sample complexity of RMAX using options and proposed an option discovery targeting to minimize the sample complexity. Solway et al. (2014) formalized an optimal behavioral hierarchy as a model which fits the behavior of the agent in tasks the best.

For planning, several works have shown empirically that adding a particular set of options or macro-operators can speed up planning algorithms (Francis & Ram, 1993; Sutton & Barto, 1998; Silver & Ciosek, 2012; Konidaris, 2016). Mann et al. (2015) analyzed the convergence rate of approx-

imate value iteration with and without options, and showed that options lead to faster convergence if their durations are longer and the value function is initialized pessimistically.

7. Conclusions

We considered two fundamental theoretical questions concerning the use of behavioral abstractions to solve MDPs: (1) minimize the size of option set given a maximum number of iterations (MOMI) and (2) minimize the number of iterations given a maximum size of option set (MIMO). We showed that both problems are computationally intractable, even for deterministic MDPs. For each problem, we produced a polynomial-time algorithm for MDPs with bounded reward and goal states, and with bounded optimality for deterministic MDPs. Although these algorithms are not practical for a single-task planning, we believe they may be a useful foundation for future option discovery methods. In the future, we are interested in using the insights established here to develop principled option-discovery algorithms for model-based reinforcement learning. Since we now know which options minimize planning time, we can better guide model-based agents toward learning them and potentially reduce sample complexity considerably.

Acknowledgments

We would like to thank the anonymous reviewer for their advice and suggestions to improve the inapproximability result for MOMI. We would like to thank Maehara Takanori and Kazuki Yoshizoe for their advice on formulating the problems. This work is supported in part by an AFOSR Young Investigator Grant to George Konidaris, under agreement number FA9550-17-1-0124. D Ellis Hershkowitz is supported in part by NSF grants CCF-1618280, CCF-1814603, CCF-1527110 and NSF CAREER award CCF-1750808.

References

- Archer, A. Two $O(\log^* k)$ -approximation algorithms for the asymmetric k -center problem. In *International Conference on Integer Programming and Combinatorial Optimization*, pp. 1–14, 2001.
- Bacon, P.-L. On the bottleneck concept for options discovery. *Masters thesis, McGill University*, 2013.
- Bakker, B. and Schmidhuber, J. Hierarchical reinforcement learning based on subgoal discovery and subpolicy specialization. In *Proceedings of the 8th Conference on Intelligent Autonomous Systems*, pp. 438–445, 2004.
- Brunskill, E. and Li, L. PAC-inspired option discovery in lifelong reinforcement learning. In *Proceedings of the 31st International Conference on International Conference on Machine Learning*, pp. 316–324, 2014.
- Chuzhoy, J., Guha, S., Halperin, E., Khanna, S., Kortsarz, G., Krauthgamer, R., and Naor, J. S. Asymmetric k -center is $\log^* n$ -hard to approximate. *Journal of the ACM*, 52(4):538–551, 2005.
- Chvatal, V. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.
- Dinitz, M., Kortsarz, G., and Raz, R. Label cover instances with large girth and the hardness of approximating basic k -spanner. In *International Colloquium on Automata, Languages, and Programming*, pp. 290–301. Springer, 2012.
- Dinur, I. and Steurer, D. Analytical approach to parallel repetition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 624–633. ACM, 2014.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. Diversity is all you need: Learning skills without a reward function. In *Proceedings of the Seventh International Conference on Learning Representations*, 2019.
- Francis, A. G. and Ram, A. The utility problem in case-based reasoning. In *Case-Based Reasoning: Papers from the 1993 Workshop*, pp. 160–161, 1993.
- Goldsmith, J., Littman, M. L., and Mundhenk, M. The complexity of plan existence and evaluation in probabilistic domains. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, pp. 182–189, San Francisco, CA, 1997. Morgan Kaufmann Publishers.
- Harb, J., Bacon, P.-L., Klissarov, M., and Precup, D. When waiting is not an option: Learning options with a deliberation cost. *arXiv preprint arXiv:1709.04571*, 2017.
- Iba, G. A. A heuristic approach to the discovery of macro-operators. *Machine Learning*, 3(4):285–317, 1989.
- Jong, N. K., Hester, T., and Stone, P. The utility of temporal abstraction in reinforcement learning. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pp. 299–306, 2008.
- Konidaris, G. Constructing abstraction hierarchies using a skill-symbol loop. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, volume 2016, pp. 1648, 2016.
- Konidaris, G. and Barto, A. Skill discovery in continuous reinforcement learning domains using skill chaining. In *Advances in Neural Information Processing Systems*, pp. 1015–1023, 2009.
- Lehnert, L., Laroche, R., and van Seijen, H. On value function representation of long horizon problems. In *AAAI*, 2018.
- Littman, M. L. Probabilistic propositional planning: Representations and complexity. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pp. 748–754. AAAI Press/The MIT Press, 1997.
- Littman, M. L., Dean, T. L., and Kaelbling, L. P. On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 394–402, 1995.
- Machado, M. C., Bellemare, M. G., and Bowling, M. A Laplacian framework for option discovery in reinforcement learning. In *Proceedings of the Thirty-fourth International Conference on Machine Learning*, 2017.
- Mann, T. and Mannor, S. Scaling up approximate value iteration with options: Better policies with fewer iterations. In *International Conference on Machine Learning*, pp. 127–135, 2014.
- Mann, T. A., Mannor, S., and Precup, D. Approximate value iteration with temporally extended actions. *Journal of Artificial Intelligence Research*, 53:375–438, 2015.

- McGovern, A. and Barto, A. G. Automatic discovery of sub-goals in reinforcement learning using diverse density. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pp. 361–368, 2001.
- Menache, I., Mannor, S., and Shimkin, N. Q-cut - dynamic discovery of sub-goals in reinforcement learning. In *European Conference on Machine Learning*, pp. 295–306, 2002.
- Moradi, P., Shiri, M. E., Rad, A. A., Khadivi, A., and Hasler, M. Automatic skill acquisition in reinforcement learning using graph centrality measures. *Intelligent Data Analysis*, 16(1):113–135, 2012.
- Panigrahy, R. and Vishwanathan, S. An $O(\log^* n)$ approximation algorithm for the asymmetric p-center problem. *Journal of Algorithms*, 27(2):259–268, 1998.
- Papadimitriou, C. H. and Tsitsiklis, J. N. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, August 1987.
- Pickett, M. and Barto, A. Policyblocks: An algorithm for creating useful macro-actions in reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pp. 506–513, 2002.
- Precup, D. and Sutton, R. S. Multi-time models for temporally abstract planning. In *Advances in neural information processing systems*, pp. 1050–1056, 1998.
- Puterman, M. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 1994.
- Silver, D. and Ciosek, K. Compositional planning using optimal option models. In *Proceedings of the 29th International Conference on Machine Learning*, pp. 1063–1070, 2012.
- Simon, H. A. *Models of man; social and rational*. Wiley, 1957.
- Şimşek, Ö. and Barto, A. Using relative novelty to identify useful temporal abstractions in reinforcement learning. In *Proceedings of the 21st International Conference on Machine Learning*, pp. 751–758, 2004.
- Şimşek, Ö. and Barto, A. G. Skill characterization based on betweenness. In *Advances in Neural Information Processing Systems*, pp. 1497–1504, 2009.
- Şimşek, Ö., Wolfe, A., and Barto, A. Identifying useful sub-goals in reinforcement learning by local graph partitioning. In *Proceedings of the Twenty Second International Conference on Machine Learning*, pp. 816–823, 2005.
- Solway, A., Diuk, C., Córdova, N., Yee, D., Barto, A. G., Niv, Y., and Botvinick, M. M. Optimal behavioral hierarchy. *PLoS computational biology*, 10(8):e1003779, 2014.
- Stolle, M. and Precup, D. Learning options in reinforcement learning. In *International Symposium on Abstraction, Reformulation, and Approximation*, pp. 212–223, 2002.
- Sutton, R., Precup, D., and Singh, S. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1): 181–211, 1999.
- Sutton, R. S. and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- Williams, R. J. and Baird, L. C. Tight performance bounds on greedy policies based on imperfect value functions. Technical report, College of Computer Science, Northeastern University, 1993.

Finding Options that Minimize Planning Time (Appendix)

Yuu Jinnai¹ David Abel¹ D Ellis Hershkowitz² Michael L. Littman¹ George Konidaris¹

A. Appendix: Inapproximability of MOMI

In this section we prove Theorem 4:

Theorem 4.

1. MOMI is $\Omega(\log n)$ hard to approximate even for deterministic MDPs unless $P = NP$.
2. $MOMI_{gen}$ is $2^{\log^{1-\epsilon} n}$ -hard to approximate for any $\epsilon > 0$ even for deterministic MDPs unless $NP \subseteq DTIME(n^{\text{poly} \log n})$.
3. MOMI is $2^{\log^{1-\epsilon} n}$ -hard to approximate for any $\epsilon > 0$ unless $NP \subseteq DTIME(n^{\text{poly} \log n})$.

For Theorems 4.2 and 4.3 we reduce our problem to the Min-Rep, problem, originally defined by (Kortsarz, 2001). Min-Rep is a variant of the better studied label cover problem (Dinur & Safra, 2004) and has been integral to recent hardness of approximation results in network design problems (Dinitz et al., 2012; Bhattacharyya et al., 2012). Roughly, Min-Rep asks how to assign as few labels as possible to nodes in a bipartite graph such that every edge is “satisfied.”

Definition 1 (Min-Rep):

Given a bipartite graph $G = (A \cup B, E)$ and alphabets Σ_A and Σ_B for the left and right sides of G respectively. Each $e \in E$ has associated with it a set of pairs $\pi_e \subseteq \Sigma_A \times \Sigma_B$ which satisfy it. **Return** a pair of assignments $\gamma_A : A \rightarrow \mathcal{P}(\Sigma_A)$ and $\gamma_B : B \rightarrow \mathcal{P}(\Sigma_B)$ such that for every $e = (A_i, B_j) \in E$ there exists an $(a, b) \in \pi_e$ such that $a \in \gamma_A(A_i)$ and $b \in \gamma_B(B_j)$. The objective is to minimize $\sum_{A_i \in A} |\gamma_A(A_i)| + \sum_{B_j \in B} |\gamma_B(B_j)|$.

We illustrate a feasible solution to an instance of Min-Rep in Figure 1.

The crucial property of Min-Rep we use is that no polynomial-time algorithm can approximate Min-Rep well. Let $\tilde{n} = |A| + |B|$.

Lemma 1 (Kortsarz 2001). *Unless $NP \subseteq DTIME(n^{\text{poly} \log n})$, Min-Rep admits no $2^{\log^{1-\epsilon} \tilde{n}}$ polynomial-time approximation algorithm for any $\epsilon > 0$.*

As a technical note, we emphasize that all relevant quantities in Min-Rep are polynomially-bounded. In Min-Rep we have $|\Sigma_A|, |\Sigma_B| \leq \tilde{n}^{c'}$ for constant c' . It immediately follows that $\sum_e |\pi_e| \leq n^c$ for constant c .

A.1. Hardness of Approximation of MOMI with Deterministic MDP

Theorem 4.1 Proof. The optimization version of the set-cover problem cannot be approximated within a factor of $c \cdot \ln n$ by a polynomial-time algorithm unless $P = NP$ (Raz & Safra, 1997). The set-cover optimization problem can be reduced to MOMI with a similar construction for a reduction from SetCover-DEC to OI-DEC. Here, the targeted minimization values of the two problems are equal: $P(\mathcal{C}) = |\mathcal{C}|$, and the number of states in OI-DEC is equal to the number of elements in the set cover on transformation. Assume there is a polynomial-time algorithm within a factor of $c \cdot \ln n$ approximation for MOMI where n is the number of states in the MDP. Let SetCover(\mathcal{U}, \mathcal{X}) be an instance of the set-cover problem. We can convert the instance into an instance of MOMI($M, 0, 2$). Using the approximation algorithm, we get a solution \mathcal{O} where $|\mathcal{O}| \leq c \ln n |\mathcal{O}^*|$, where \mathcal{O}^* is the optimal solution. We construct a solution for the set cover \mathcal{C} from the solution to the MOMI \mathcal{O} (see the construction in the proof of Theorem 1). Because $|\mathcal{C}| = |\mathcal{O}|$ and $|\mathcal{C}^*| = |\mathcal{O}^*|$, where \mathcal{C}^* is the optimal solution for the set cover, we get $|\mathcal{C}| = |\mathcal{O}| \leq c \ln n |\mathcal{O}^*| = c \ln n |\mathcal{C}^*|$. Thus, we acquire a $c \cdot \ln n$ approximation solution for the set-cover problem within polynomial time, something only possible if $P=NP$. Thus, there is no polynomial-time algorithm with a factor of $c \cdot \ln n$ approximation for MOMI, unless $P=NP$. \square

A.2. Hardness of Approximation of $MOMI_{gen}$

We now show our hardness of approximation of $2^{\log^{1-\epsilon} n}$ for $MOMI_{gen}$, Theorem 4.2.¹

We start by describing our reduction from an instance of Min-Rep to an instance of $MOMI_{gen}$. The intuition behind our reduction is that we can encode choosing a label for a

¹We assume that \mathcal{O}' is a “good” set of options in the sense that there exists some set $\mathcal{O}^* \subseteq \mathcal{O}'$ such that $L_{\epsilon, V_0}(\mathcal{O}^*) \leq \ell$. We also assume, without loss of generality, that $\epsilon < 1$ throughout this section; other values of ϵ can be handled by re-scaling rewards in our reduction.

vertex in Min-Rep as choosing an option in our MOMI_{gen} instance. In particular, we will have a state for each edge in our Min-Rep instance and reward will propagate quickly to that state when value iteration is run only if the options corresponding to a satisfying assignment for that edge are chosen.

More formally, our reduction is as follows. Consider an instance of Min-Rep, MR, given by $G = (A \cup B, E)$, Σ_A , Σ_B and $\{\pi_e\}$. Our instance of MOMI_{gen} is as follows where $\gamma = 1$ and $l = 3$.²

- **State space** We have a single goal state S_g along with states S'_g and S''_g . For each edge e we create a state S_e . Let $\text{Sat}_A(e)$ consist of all $a \in \Sigma_A$ such that a is in some assignment in π_e . Define $\text{Sat}_B(e)$ symmetrically. For each edge $e \in E$ we create a set of $2 \cdot |\text{Sat}_A(e)|$ states, namely S_{ea} and S'_{ea} for every $a \in \text{Sat}_A(e)$. We do the same for $b \in \text{Sat}_B(e)$.
- **Actions and Transitions** We have a single action from S'_g to S_g , a single action from S''_g to S'_g . For each edge e we have the following deterministic actions: Every S'_{ea} has a single outgoing action to S_{ea} for $a \in \text{Sat}_A(e)$; Every S_{eb} has a single outgoing action to $S_{eb'}$ for $b \in \text{Sat}_B(e)$; Every S_{ea} has an outgoing action to S_{eb} if $(a, b) \in \pi_e$ and every S'_{eb} has a single outgoing action to S_g ; Lastly, we have a single action from S'_{ea} to S''_g for every $a \in \text{Sat}_A(e)$.
- **Reward** The reward of arriving in S_g is 1. The reward of arriving in every other state is 0.
- **Option Set** Our option set \mathcal{O}' is as follows. For each vertex $A_i \in A$ and each $a \in \Sigma_A$ we have an option $O(A_i, a)$: The initiation set of this option is every S_e where e is incident to A_i ; The termination set of this

²It is easy to generalize these results to $l \geq 4$ by replacing certain edges with paths.




Figure 1: An instance of Min-Rep with $\Sigma_A = \{a_1, a_2, a_3\}$ and $\Sigma_B = \{b_1, b_2, b_3\}$. Edge e is labeled with pairs in π_e . Feasible solution (γ_A, γ_B) illustrated where $\gamma_A(A_i)$ and $\gamma_B(B_j)$ below A_i and B_j in blue. Constraints colored to coincide with stochastic action colors in Figure 3.




Figure 2: Our MOMI_{gen} reduction applied to the Min-Rep problem in Figure 1. $e_1 = (A_1, B_1)$, $e_2 = (A_1, B_2)$, $e_3 = (A_2, B_2)$. Actions given in solid lines and each option in \mathcal{O}' represented in its own color as a dashed line from initiation to termination states. Notice that a single option goes from S_{e3b1} and S_{e2b1} to S_g .

option is every S_{ea} where A_i is incident to e ; The policy of this option takes the action from S'_{ea} to S_{ea} when in S'_{ea} and the action from S_e to S'_{ea} when in S_e . Symmetrically, for every vertex $B_j \in B$ and each $b \in \Sigma_B$ we have an option $O(B_j, b)$: The initiation set of this option is every S_{eb} where e is incident to B_j ; The termination set of this option is S_g ; The policy of this option takes the action from S_{eb} to S'_{eb} when in S_{eb} and from S'_{eb} to S_g when in S'_{eb} .

One should think of choosing option $O(v, x)$ as corresponding to choosing label x for vertex v in the input Min-Rep instance. Let $\text{MOMI}_{gen}(\text{MR})$ be the MDP output given instance MR of Min-Rep and see Figure 3 for an illustration of our reduction.

Let $\text{OPT}_{\text{MOMI}_{gen}}$ be the value of the optimal solution to $\text{MOMI}_{gen}(\text{MR})$ and let OPT_{MR} be the value of the optimal Min-Rep solution to MR. The following lemmas demonstrates the correspondence between a MOMI_{gen} and Min-Rep solution.

Lemma 2. $\text{OPT}_{\text{MOMI}_{gen}} \leq \text{OPT}_{\text{MR}}$

Proof. Given a solution (γ_A, γ_B) to MR, define $\mathcal{O}_{\gamma_A, \gamma_B} := \{O(v, x) : v \in V(G) \wedge (\gamma_A(v) = x \vee \gamma_B(v) = x)\}$ as the corresponding set of options. Let γ_A^* and γ_B^* be the optimal solutions to MR which is of cost OPT_{MR} .

We now argue that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is a feasible solution to $\text{MOMI}_{gen}(\text{MR})$ of cost OPT_{MR} , demonstrating that the op-

timal solution to $\text{MOMI}_{gen}(\text{MR})$ has cost at most OPT_{MR} . To see this notice that by construction the MOMI_{gen} cost of $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is exactly the Min-Rep cost of (γ_A^*, γ_B^*) .

We need only argue, then, that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is feasible for $\text{MOMI}_{gen}(\text{MR})$ and do so now. The value of every state in $\text{MOMI}_{gen}(\text{MR})$ is 1. Thus, we must guarantee that after 3 iterations of value iteration, every state has value 1. However, without any options every state except each S_e has value 1 after 3 iterations of value iteration. Thus, it suffices to argue that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ guarantees that every S_e will have value 1 after 3 iterations of value iteration. Since (γ_A^*, γ_B^*) is a feasible solution to MR we know that for every $e = (A_i, B_j)$ there exists an $\bar{a} \in \gamma_A^*(A_i)$ and $\bar{b} \in \gamma_B^*(B_j)$ such that $(\bar{a}, \bar{b}) \in \pi_e$; correspondingly there are options $O(A_i, \bar{a}), O(B_j, \bar{b}) \in \mathcal{O}_{\gamma_A^*, \gamma_B^*}$. It follows that, given options $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ from, S_e one can take option $O(A_i, \bar{a})$ then the action from $S_{e\bar{a}}$ to $S_{e\bar{b}}$ and then option $O(B_j, \bar{b})$ to arrive in S_g ; thus, after 3 iterations of value iteration the value of S_e is 1. Thus, we conclude that after 3 iterations of value iteration every state has converged on its value. \square

We now show that a solution to $\text{MOMI}_{gen}(\text{MR})$ corresponds to a solution to MR. For the remainder of this section $\gamma_A^{\mathcal{O}}(A_i) := \{a : O(A_i, a) \in \mathcal{O}\}$ and $\gamma_B^{\mathcal{O}}(B_j) := \{b : O(B_j, b) \in \mathcal{O}\}$ is the Min-Rep solution corresponding to option set \mathcal{O} .

Lemma 3. *For a feasible solution to $\text{MOMI}_{gen}(\text{MR})$, \mathcal{O} , we have $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution to MR of cost $|\mathcal{O}|$.*

Proof. Notice that by construction the Min-Rep cost of $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is exactly $|\mathcal{O}|$. Thus, we need only prove that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution for MR.

We do so now. Consider an arbitrary edge $e = (A_i, B_j) \in E$; we wish to show that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ satisfies e . Since \mathcal{O} is a feasible solution to $\text{MOMI}_{gen}(\text{MR})$ we know that after 3 iterations of value iteration every state must converge on its value. Moreover, notice that the value of every state in $\text{MOMI}_{gen}(\text{MR})$ is 1. Thus, it must be the case that for every S_e there exists a path of length 3 from S_e to S_g using either options or actions. The only such paths are those that take an option $O(A_i, a)$, then an action from S_{ea} to S_{eb} then option $O(B_j, b)$ where $(a, b) \in \pi_e$. It follows that $a \in \gamma_A^{\mathcal{O}}(A_i)$ and $b \in \gamma_B^{\mathcal{O}}(B_j)$. But since $(a, b) \in \pi_e$, we then know that e is satisfied. Thus, every edge is satisfied and so $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution to MR. \square

Theorem 4.2 Proof. Assume $\text{NP} \not\subseteq \text{DTIME}(n^{\text{poly} \log n})$ and for the sake of contradiction that there exists an $\varepsilon > 0$ for which polynomial-time algorithm $\mathcal{A}_{\text{MOMI}_{gen}}$ can $2^{\log^{1-\varepsilon} n}$ -approximate MOMI_{gen} . We use $\mathcal{A}_{\text{MOMI}_{gen}}$ to $2^{\log^{1-\varepsilon'} \tilde{n}}$ approximate Min-Rep for a fixed constant $\varepsilon' > 0$ in polynomial-time, thereby contradicting Lemma 1. Again,

\tilde{n} is the number of vertices in the graph of the Min-Rep instance.

We begin by noting that the relevant quantities in $\text{MOMI}_{gen}(\text{MR})$ are polynomially-bounded. Notice that the number of states n in the MDP in $\text{MOMI}_{gen}(\text{MR})$ is at most $O(\tilde{n}^2 |\Sigma_A| |\Sigma_B|) = \tilde{n}^c$ for some fixed constant c by the aforementioned assumption that Σ_A and Σ_B are polynomially-bounded in \tilde{n} .³

Our polynomial-time approximation algorithm to approximate instance MR of Min-Rep is as follows: Run $\mathcal{A}_{\text{MOMI}_{gen}}$ on $\text{MOMI}_{gen}(\text{MR})$ to get back option set \mathcal{O} . Return $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ as defined above as our solution to MR.

We first argue that our algorithm is polynomial-time in \tilde{n} . However, notice that for each vertex, we create a polynomial number of states. Thus, the number of states in $\text{MOMI}_{gen}(\text{MR})$ is polynomially-bounded in \tilde{n} and so $\mathcal{A}_{\text{MOMI}_{gen}}$ runs in time polynomial in \tilde{n} . A polynomial runtime of our algorithm immediately follows.

We now argue that our algorithm is a $2^{\log^{1-\varepsilon'} \tilde{n}}$ -approximation for Min-Rep for some $\varepsilon' > 0$. Applying Lemma 3, the approximation of $\mathcal{A}_{\text{MOMI}_{gen}}$ and then Lemma 2, we have that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution for MR with cost

$$\begin{aligned} \text{cost}_{\text{Min-Rep}}(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}}) &= |\mathcal{O}| \\ &\leq 2^{\log^{1-\varepsilon} n} \text{OPT}_{\text{MOMI}_{gen}} \\ &\leq 2^{\log^{1-\varepsilon} n} \text{OPT}_{\text{MR}} \end{aligned}$$

Thus, $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a $2^{\log^{1-\varepsilon} n}$ approximation for the optimal Min-Rep solution where n is the number of states in the MDP of $\text{MOMI}_{gen}(\text{MR})$. Now recalling that $n \leq \tilde{n}^c$ for fixed constant c . We therefore have that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a $2^{\log^{1-\varepsilon} \tilde{n}^c} = 2^{c^{1-\varepsilon} \log^{1-\varepsilon} \tilde{n}} \leq c' \cdot 2^{\log^{1-\varepsilon} \tilde{n}}$ approximation for a constant c' . Choosing ε sufficiently small, we have that $c' \cdot 2^{\log^{1-\varepsilon} \tilde{n}} \leq 2^{\log^{1-\varepsilon'} \tilde{n}}$ for sufficiently large \tilde{n} .

Thus, our polynomial-time algorithm is a $2^{\log^{1-\varepsilon'} \tilde{n}}$ -approximation for Min-Rep for $\varepsilon' > 0$, thereby contradicting Lemma 1. We conclude that MOMI_{gen} cannot be $2^{\log^{1-\varepsilon} n}$ -approximated. \square

A.3. Hardness of Approximation of MOMI with Stochastic MDP

We now show our hardness of approximation of $2^{\log^{1-\varepsilon} n}$ for MOMI, Theorem 4.3. We will notably use the stochasticity

³It is also worth noticing that since we create at most $O(\tilde{n}|\Sigma_A| + \tilde{n}|\Sigma_B|)$ options, the total number of options in \mathcal{O}' is at most polynomial in \tilde{n} .

of the input MDP to show this result.⁴

We begin by describing our reduction from an instance of Min-Rep to an instance of MOMI. The intuition behind our reduction is as follows. As in our reduction for MOMI_{gen} we will have vertex for each edge in our Min-Rep instance and reward will propagate quickly to that vertex when value iteration is run only if the options corresponding to a satisfying assignment for that edge are chosen. The challenge, however, is that since our options are now only point options (whereas in MOMI_{gen} they were arbitrary options) it seems that we can no longer constrain a solution to choose options exactly corresponding to a feasible Min-Rep solution.

To solve this issue we critically use stochasticity. Whether or not a given edge in a Min-Rep is satisfied is an or of ands: A fixed edge is satisfied when *one* of its satisfying assignments is met (an or) and a given satisfying assignment is met when both endpoints have the right labels (an and). We will exploit the fact that the value of a state in an MDP is a max over actions to encode the “or” in Min-Rep and we will use the fact that in a stochastic MDP the value of a (state, action) pair is the sum over states to encode the “and” in Min-Rep.

More formally, our reduction is as follows. Consider instance MR of Min-Rep given by $G = (A \cup B, E)$, Σ_A, Σ_B and $\{\pi_e\}$. Our instance of MOMI is as follows where $\gamma = 1$ and $l = 2$.⁵

- **State space** We have a goal state S_i for each $A_i \in A$. Again, let $\text{Sat}_A(e)$ consist of all $a \in \Sigma_A$ such that a is in some assignment in π_e . For each $A_i \in A$ and $a \in \text{Sat}_A(e)$ we will we add to our MDP states S_{ia} and S'_{ia} . We symmetrically do the same for all states in Σ_B . For each $e \in E$ we will also add a state S_e .⁶
- **Actions and Transitions** Every S_{ia} state has a single action to S'_{ia} and every S'_{ia} state has a single action to S_i . The same symmetrically holds for states from a $B_j \in B$. Every S_e for $e = (A_i, B_j)$ has $|\pi_{(A_i, B_j)}|$ actions associated with it, namely $\{\alpha_{(a,b)}\}$ where $(a, b) \in \pi_{(A_i, B_j)}$. Action $\alpha_{(a,b)}$ has a probability .5 of transitioning to state S_{ia} and a probability .5 of transitioning to state S_{jb} .
- **Reward** The reward of arriving in any S_i or S_j for $A_i \in A$ or $B_j \in B$ is 1 and 0 for every other state.

⁴We may assume without loss of generality $\varepsilon < .5$ throughout this section; rewards in our reduction can be re-scaled to handle larger ε .

⁵It is easy to generalize these results to $l \geq 3$ by replacing edges with paths.

⁶It is not hard to see that this construction can be modified so that we have only a single goal state if need be; we need only set every S_i and S_j to be the same state. We assume multiple goal states for ease of exposition.




Figure 3: Our MOMI reduction applied to the Min-Rep problem in Figure 1. $e_1 = (A_1, B_1)$, $e_2 = (A_1, B_2)$, $e_3 = (A_2, B_2)$. Stochastic options colored according to the pair in π_e to which they correspond, branching into the two states in which they arrive with equal probability. Deterministic action given as solid black arcs. Possible point options given as dashed arcs.

Notice that no point options have S_e as an initialization state since any such option would have a .5 probability of never terminating (and we assume our options always terminate). See Figure 3 for an illustration of our reduction. One should think of choosing a point option from S_{ia} to S_i as corresponding to choosing label a for A_i in the input Min-Rep instance. The same holds for label b for B_j and choosing a point option from S_{jb} to S_j . Let $\text{MOMI}(\text{MR})$ be the MOMI instance output by our reduction given instance MR of Min-Rep.

We now demonstrate that our reduction allows us to show that MOMI cannot be $2^{\log^{1-\varepsilon} n}$ -approximated for any $\varepsilon > 0$. Let OPT_{MOMI} be the value of the optimal solution to $\text{MOMI}(\text{MR})$ and let OPT_{MR} be the value of the optimal Min-Rep solution to MR. The following lemmas demonstrates the correspondence between a MOMI and Min-Rep solution.

Lemma 4. $\text{OPT}_{\text{MOMI}} \leq \text{OPT}_{\text{MR}}$

Proof. Our proof translates between point options in our reduction and assignments in the input Min-Rep instance in the natural way. Given a solution (γ_A, γ_B) to MR, define $\mathcal{O}_{\gamma_A, \gamma_B}$ as consisting of all point options from S_{ia} to S_i if $a \in \gamma_A(A_i)$ and all points options from S_{jb} to S_j if $b \in \gamma_B(B_j)$. Let γ_A^* and γ_B^* be the optimal solutions to MR which is of cost OPT_{MR} .

We claim that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is a feasible solution to $\text{MOMI}(\text{MR})$ of cost OPT_{MR} , demonstrating that the optimal solution to $\text{MOMI}(\text{MR})$ has cost at most OPT_{MR} . To see this notice that by construction the MOMI cost of $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is exactly

the Min-Rep cost of γ_A^*, γ_B^* .

We need only argue, then, that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ is feasible for MOMI(MR) and do so now. Notice that the value of every state in MOMI is 1. Thus, we must guarantee that after 2 iterations of value iteration, every state has value 1. However, without any options every state except for S_e where $e \in E$ has value 1 after 2 iterations of value iteration. Thus, it suffices to argue that $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ guarantees that every S_e will have value 1 after 2 iterations of value iteration. Since (γ_A^*, γ_B^*) is a feasible solution to MR we know that for every $e = (A_i, B_j)$ there exists $\bar{a} \in \gamma_A^*(A_i)$ and $\bar{b} \in \gamma_B^*(B_j)$ such that $(\bar{a}, \bar{b}) \in \pi_e$; correspondingly there is some action from S_e with a .5 probability of resulting in state $S_{i\bar{a}}$ and a .5 probability of resulting in state $S_{j\bar{b}}$ where $\mathcal{O}_{\gamma_A^*, \gamma_B^*}$ has a point option from $S_{i\bar{a}}$ to S_i and a point options from $S_{j\bar{b}}$ to S_j . That is, $V_1(S_{i\bar{a}}) = 1$ and $V_1(S_{j\bar{b}}) = 1$. Thus, after one iteration of value iteration the values of $S_{i\bar{a}}$ and $S_{j\bar{b}}$ are both 1 and so after two iterations of value iteration the value of S_e is

$$\begin{aligned} V_2(S_e) &= \max_{\alpha_{(a,b)}} .5 \cdot (V_1(S_{ia})) + .5 \cdot (V_1(S_{jb})) \\ &\geq .5 \cdot (V_1(S_{i\bar{a}})) + .5 \cdot (V_1(S_{j\bar{b}})) \\ &= 1. \end{aligned}$$

Thus, $V_2(S_e) = 1$ for every S_e and so we conclude that after two iterations of value iteration every state has converged on its value. \square

We now show that a solution to MOMI(MR) corresponds to a solution to MR. For the remainder of this section let $\gamma_A^{\mathcal{O}}(A_i) := \{a : O(S_{ja}, S_j) \in \mathcal{O}\}$ and $\gamma_B^{\mathcal{O}}(B_j) := \{b : O(S_{jb}, S_j) \in \mathcal{O}\}$ where for the remainder of this section $O(S, S')$ stands for a point option with initiation state S and termination state S' .

Lemma 5. *For any feasible solution \mathcal{O} to MOMI(MR) we have $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution to MR of cost $|\mathcal{O}|$.*

Proof. Notice that by construction the Min-Rep cost of $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is exactly $|\mathcal{O}|$. Thus, we need only prove that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution for MR.

We do so now. Consider an arbitrary edge $e = (A_i, B_j) \in E$; we wish to show that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ satisfies e . Since \mathcal{O} is a feasible solution we know that after two iterations of value iteration every state must converge on its value (up to an ϵ factor which we can ignore by our above assumption that $\epsilon < .5$). Moreover, notice that the value of every state in MOMI(MR) is 1. Thus, it must be the case that for every S_e we have $V_2(S_e) = 1$ for $e = (A_i, B_j)$. It follows, then, that there is some action $\alpha_{(\bar{a}, \bar{b})}$ where $(\bar{a}, \bar{b}) \in \pi_{(A_i, B_j)}$ such that

$$1 = V_2(S_e) = .5 \cdot (V_1(S_{i\bar{a}})) + .5 \cdot (V_1(S_{j\bar{b}})).$$

Since the value of every state is at most 1, it follows that $V_1(S_{i\bar{a}}) = V_1(S_{j\bar{b}}) = 1$. However, since $V_1(S_{i\bar{a}})$ and $V_1(S_{j\bar{b}})$ are both two hops from the only goal reachable from them (S_i and S_j respectively) it must be the case that there is some point option from $S_{i\bar{a}}$ to S_i and $S_{j\bar{b}}$ to S_j . Thus, by definition of $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ we then have $\bar{a} \in \gamma_A^{\mathcal{O}}$ and $\bar{b} \in \gamma_B^{\mathcal{O}}$. Since $(\bar{a}, \bar{b}) \in \pi_{(A_i, B_j)}$ it follows that arbitrary edge $e = (A_i, B_j)$ is satisfied. Thus, every edge in E is satisfied and so $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a feasible solution for MR. \square

Finally, we conclude the hardness of approximation of MOMI.

Theorem 4.3 Proof. Assume $\text{NP} \not\subseteq \text{DTIME}(n^{\text{poly} \log n})$ and for the sake of contradiction that there exists an $\epsilon > 0$ for which a polynomial-time algorithm $\mathcal{A}_{\text{MOMI}}$ can $2^{\log^{1-\epsilon} n}$ -approximate MOMI. We use $\mathcal{A}_{\text{MOMI}}$ to $2^{\log^{1-\epsilon'} \tilde{n}}$ approximate Min-Rep for a fixed constant $\epsilon' > 0$ in polynomial-time in \tilde{n} , thereby contradicting Lemma 1. Again, \tilde{n} is the number of vertices in the graph of the Min-Rep instance.

We begin by noting that the relevant quantities in MOMI(MR) are polynomially-bounded. Let $\tilde{n} := |A| + |B|$ be the number of vertices in our MR instance. Notice that the number of states in the MDP, n , in our MOMI(MR) instance is at most $O(\tilde{n} + 2|A||\Sigma_A| + |B||\Sigma_B| + |E|) \leq \tilde{n}^c$ for some fixed constant c by the aforementioned assumption that Σ_A and Σ_B are polynomially-bounded in \tilde{n} .⁷

Our polynomial-time approximation algorithm to approximate instance MR of Min-Rep is as follows: Run $\mathcal{A}_{\text{MOMI}}$ on MOMI(MR) to get back option set \mathcal{O} . Return $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ as defined above as our solution to MR.

We first argue that our algorithm is polynomial time in \tilde{n} . For each vertex in MR, we create a polynomial number of states and actions. Thus, the number of states in MOMI(MR) is polynomially-bounded in \tilde{n} and so $\mathcal{A}_{\text{MOMI}}$ runs in time polynomial in \tilde{n} . A polynomial runtime of our algorithm immediately follows.

We now argue that our algorithm is a $2^{\log^{1-\epsilon'} \tilde{n}}$ -approximation for Min-Rep for some $\epsilon' > 0$. Applying Lemma 5, the approximation of $\mathcal{A}_{\text{MOMI}}$ and then Lemma 4, we have that the Min-Rep cost of $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is

$$\begin{aligned} \text{cost}_{\text{Min-Rep}}(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}}) &= |\mathcal{O}| \\ &\leq 2^{\log^{1-\epsilon} n} \text{OPT}_{\text{MOMI}} \\ &\leq 2^{\log^{1-\epsilon} n} \text{OPT}_{\text{MR}} \end{aligned}$$

Thus, $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a $2^{\log^{1-\epsilon} n}$ approximation for the opti-

⁷It is worth noting, also, that since we create at most $\sum_e |\pi_e|$ actions for any state, the number of total actions in our MDP is at most polynomial in \tilde{n} .

mal Min-Rep solution where n is the number of states in the MDP of MOMI(MR). Now recalling that $n \leq \tilde{n}^c$ for fixed constant c . We therefore have that $(\gamma_A^{\mathcal{O}}, \gamma_B^{\mathcal{O}})$ is a $2^{\log^{1-\varepsilon} \tilde{n}^c} = 2^{c^{1-\varepsilon} \log^{1-\varepsilon} \tilde{n}} \leq c' \cdot 2^{\log^{1-\varepsilon} \tilde{n}}$ approximation for a constant c' . Choosing ε sufficiently small, we have that $c' \cdot 2^{\log^{1-\varepsilon} \tilde{n}} \leq 2^{\log^{1-\varepsilon'} \tilde{n}}$ for sufficiently large \tilde{n} .

Thus, our polynomial-time algorithm is a $2^{\log^{1-\varepsilon'} \tilde{n}}$ -approximation for Min-Rep for $\varepsilon' > 0$, thereby contradicting Lemma 1. We conclude that MOMI cannot be $2^{\log^{1-\varepsilon} n}$ -approximated. \square

A.4. A-MIMO

In this subsection we show the following theorem (we show Theorem 5 later):

Theorem 6. *A-MIMO has following properties:*

1. *A-MIMO runs in polynomial time.*
2. *If the MDP is deterministic, it has a bounded suboptimality of $O(\log^* k)$.*
3. *The number of iterations to solve the MDP using the acquired options is upper bounded by $P(\mathcal{C})$.*

Theorem 6.1. *A-MIMO runs in polynomial time.*

Proof. Each step of the procedure runs in polynomial time.

(1) Solving an MDP takes polynomial time. To compute d we need to solve MDPs at most $|\mathcal{S}|$ times. Thus, it runs in polynomial time.

(2) The approximation algorithm we deploy for solving the asymmetric-k center which runs in polynomial time (Archer, 2001). Because the procedure by Archer (2001) terminates immediately after finding a set of options which guarantees the suboptimality bounds, it tends to find a set of options smaller than k . In order to use the rest of the options effectively within polynomial time, we use a procedure Expand to greedily add a few options at once until it finds all k options. We enumerate all possible set of options of size $r = \lceil \log k \rceil$ (if $|\mathcal{O}| + \log k > k$ then we set $r = k - |\mathcal{O}|$) and add a set of options which minimizes ℓ (breaking ties randomly) to the option set \mathcal{O} . We repeat this procedure until $|\mathcal{O}| = k$. This procedure runs in polynomial time. The number of possible option set of size r is ${}_r C_n = O(n^r) = O(k)$. We repeat this procedure at most $\lceil k / \log k \rceil$ times, thus the total computation time is bounded by $O(k^2 / \log k)$.

(3) Immediate.

Therefore, A-MIMO runs in polynomial time. \square

Before we show that it is sufficient to consider a set of options with its terminal state set to the goal state of the MDP.

Lemma 6. *There exists an optimal option set for MIMO and MOMI with all terminal state set to the goal state.*

Proof. Assume there exists an option with terminal state set to a state other than the goal state in the optimal option set \mathcal{O} . By triangle inequality, swapping the terminal state to the goal state will monotonically decrease $d(s, g)$ for every state. By swapping every such option we can construct an option set \mathcal{O}' with $L_{\varepsilon, V_0}(\mathcal{O}') \leq L_{\varepsilon, V_0}(\mathcal{O})$. \square

Lemma imply that discovering the best option set among option sets with their terminal state fixed to the goal state is sufficient to find the best option set in general. Therefore, our algorithms seek to discover options with termination state fixed to the goal state.

Using the option set acquired, the number of iterations to solve the MDP is bounded by $P(\mathcal{C})$. To prove this we first generalize the definition of the distance function to take a state and a set of states as arguments $d_\varepsilon : \mathcal{S} \times 2^{\mathcal{S}} \rightarrow \mathbb{N}$. Let $d_\varepsilon(s, \mathcal{C})$ the number of iterations for s to converge ε -optimal if every state $s' \in \mathcal{C}$ has converged to ε -optimal: $d_\varepsilon(s, \mathcal{C}) := \min(d'_\varepsilon(s), 1 + d'_\varepsilon(s, \mathcal{C})) - 1$. As adding an option will never make the number of iterations larger,

Lemma 7.

$$d(s, \mathcal{C}) \leq \min_{s' \in \mathcal{C}} d(s, s'). \quad (1)$$

Using this, we show the following proposition.

Theorem 6.2. *The number of iterations to solve the MDP using the acquired options is upper bounded by $P(\mathcal{C})$.*

Proof. $P(\mathcal{C}) = \max_{s \in \mathcal{S}} \min_{c \in \mathcal{C}} d(s, c) \geq \max_{s \in \mathcal{S}} d(s, \mathcal{C}) = L_{\varepsilon, V_0}(\mathcal{O})$ (using Equation 1). Thus $P(\mathcal{C})$ is an upper bound for $L_{\varepsilon, V_0}(\mathcal{O})$. \square

The reason why $P(\mathcal{C})$ does not always give us the exact number of iterations is because adding two options starting from s_1, s_2 may make the convergence of s_0 faster than $d(s_0, s_1)$ or $d(s_0, s_2)$. Example: Figure 4 is an example of such an MDP. From s_0 it may transit to s_1 and s_2 with probability 0.5 each. Without any options, the value function converges to exactly optimal value for every state with 3 steps. Adding an option either from s_1 or s_2 to g does not shorten the iteration for s_0 to converge. However, if we add two options from s_1 and s_2 to g , s_0 converges within 2 steps, thus the MDP is solved with 2 steps.

The equality of the statement 1 holds if the MDP is deterministic. That is, $d(s, \mathcal{C}) = \min_{s' \in \mathcal{C}} d(s, s')$ for deterministic MDP.

Theorem 6.3. *If the MDP is deterministic, it has a bounded suboptimality of $O(\log^* k)$.*




Figure 4: An example of an MDP where $d(s, \mathcal{C}) < \min_{s' \in \mathcal{C}} d(s, s')$. Here the transition induced by the optimal policy is stochastic, thus from s_0 one may go to s_1 and s_2 by probability 0.5 each. Either adding an option from s_1 or s_2 to g does not make the convergence faster, but adding both makes it faster.

Proof. First we show $P(\mathcal{C}^*) = L_{\epsilon, V_0}(\mathcal{O}^*)$ for deterministic MDP. From $d(s, \mathcal{C}) = \min_{s' \in \mathcal{C}} d(s, s')$, $P(\mathcal{C}^*) = \max_{s \in \mathcal{S}} \min_{c \in \mathcal{C}^*} d(s, c) = \max_{s \in \mathcal{S}} d(s, \mathcal{C}^*) = L_{\epsilon, V_0}(\mathcal{O}^*)$.

The asymmetric k -center solver guarantees that the output \mathcal{C} satisfies $P(\mathcal{C}) \leq c(\log^* k + O(1))P(\mathcal{C}^*)$ where n is the number of nodes (Archer, 2001). Let MIMO(M, ϵ, k) be an instance of MIMO. We convert this instance to an instance of asymmetric k -center AsymKCenter(\mathcal{U}, d, k), where $|\mathcal{U}| = |\mathcal{S}|$. By solving the asymmetric k -center with the approximation algorithm, we get a solution \mathcal{C} which satisfies $P(\mathcal{C}) \leq c(\log^* k + O(1))P(\mathcal{C}^*)$. Thus, the output of the algorithm \mathcal{O} satisfies $L_{\epsilon, V_0}(\mathcal{O}) = P(\mathcal{C}) \leq c(\log^* k + O(1))P(\mathcal{C}^*) = c(\log^* k + O(1))L_{\epsilon, V_0}(\mathcal{O}^*)$. Thus, $L_{\epsilon, V_0}(\mathcal{O}) \leq c(\log^* k + O(1))L_{\epsilon, V_0}(\mathcal{O}^*)$ is derived. \square

Proposition 1 (Greedy Strategy). *Let an option set \mathcal{O} be a set of point option constructed by greedily adding one point option which minimizes the number of iterations. An improvement $L_{\epsilon, V_0}(\emptyset) - L_{\epsilon, V_0}(\mathcal{O})$ by the greedy algorithm can be arbitrary small (i.e. 0) compared to the optimal option set.*

Proof. We show by the example in a shortest-path problem in Figure 5. The MDP can be solved within 4 iterations without options: $L_{\epsilon, V_0}(\emptyset) = 4$. With an optimal option set of size $k = 2$ the MDP can be solved within 2 iterations: $L_{\epsilon, V_0}(\mathcal{O}^*) = 2$ (an initiation state of each option in optimal option set is denoted by $*$ in the Figure). On the other hand, a greedy strategy may not improve L at all. No single point option does not improve L . Let's say we picked a point option from s_1 to g . Then, there is no single point option we can add to that option to improve L in the second iteration. Therefore, the greedy procedure returns \mathcal{O} which has $L_{\epsilon, V_0}(\emptyset) - L_{\epsilon, V_0}(\mathcal{O}) = 0$. Therefore,

$(L_{\epsilon, V_0}(\emptyset) - L_{\epsilon, V_0}(\mathcal{O})) / (L_{\epsilon, V_0}(\emptyset) - L_{\epsilon, V_0}(\mathcal{O}^*))$ can be arbitrary small non-negative value (i.e. 0).




Figure 5: Example of MIMO where the improvement of a greedy strategy can be arbitrary small compared to the optimal option set. \square

A.5. A-MOMI

In this subsection we show the following theorem:

Theorem 5. *A-MOMI has the following properties:*

1. A-MOMI runs in polynomial time.
2. It guarantees that the MDP is solved within ℓ iterations using the option set acquired by A-MOMI \mathcal{O} .
3. If the MDP is deterministic, the option set is at most $O(\log n)$ times larger than the smallest option set possible to solve the MDP within ℓ iterations.

Theorem 5.1. *A-MOMI runs in polynomial time.*

Proof. Each step of the procedure runs in polynomial time.

(1) Solving an MDP takes polynomial time (Littman et al., 1995). To compute d we need to solve MDPs at most $|\mathcal{S}|$ times. Thus, it runs in polynomial time.

(4) We solve the set cover using a polynomial time approximation algorithm (Chvatal, 1979) which runs in $O(n^3)$, thus run in polynomial time.

(2), (3), and (5) Immediate. \square

Theorem 5.2. *A-MOMI guarantees that the MDP is solved within ℓ iterations using the option set \mathcal{O} .*

Proof. A state $s \in X_g^+$ reaches optimal within ℓ steps by definition. For every state $s \in \mathcal{S} \setminus X_g^+$, the set cover guarantees that we have $X_{s'} \in \mathcal{C}$ such that $d(s, s') < \ell$. As we generate an option from s' to g , s' reaches to optimal value with 1 step. Thus, s reaches to ϵ -optimal value within $d(s, s') + 1 \leq \ell$. Therefore, every state reaches ϵ -optimal value within ℓ steps. \square

Theorem 5.3. *If the MDP is deterministic, the option set is at most $O(\log n)$ times larger than the smallest option set possible to solve the MDP within ℓ iterations.*

Proof. Using a suboptimal algorithm by Chvatal (1979) we get \mathcal{C} such that $|\mathcal{C}| \leq O(\log n)|\mathcal{C}^*|$. Thus, $|\mathcal{O}| = |\mathcal{C}| \leq O(\log n)|\mathcal{C}^*| = O(\log n)|\mathcal{O}^*|$. \square

Appendix: Experiments

We show the figures for experiments. Figure 6 shows the options found by solving MIMO optimally/suboptimally in four room domain. Figure 7 shows the options in 9x9 grid domain.

References

- Archer, A. Two $O(\log^* k)$ -approximation algorithms for the asymmetric k -center problem. In *International Conference on Integer Programming and Combinatorial Optimization*, pp. 1–14, 2001.
- Bhattacharyya, A., Grigorescu, E., Jung, K., Raskhodnikova, S., and Woodruff, D. P. Transitive-closure spanners. *SIAM Journal on Computing*, 41(6):1380–1425, 2012.
- Chvatal, V. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.
- Dinitz, M., Kortsarz, G., and Raz, R. Label cover instances with large girth and the hardness of approximating basic k -spanner. In *International Colloquium on Automata, Languages, and Programming*, pp. 290–301. Springer, 2012.
- Dinur, I. and Safra, S. On the hardness of approximating label-cover. *Information Processing Letters*, 89(5):247–254, 2004.
- Kortsarz, G. On the hardness of approximating spanners. *Algorithmica*, 30(3):432–450, 2001.
- Littman, M. L., Dean, T. L., and Kaelbling, L. P. On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 394–402, 1995.
- Raz, R. and Safra, S. A sub-constant error-probability low-degree test, and a sub-constant error-probability pcp characterization of np . In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pp. 475–484. ACM, 1997.




Figure 6: Comparison of the optimal point options vs. options generated by the approximation algorithm A-MIMO. We observed that the approximation algorithm is similar to that of optimal options. Note that optimal option set is not unique: there can be multiple optimal option set, and we are visualize one of them returned by the solver.




Figure 7: Comparison of the optimal point options for planning vs. bottleneck options proposed for reinforcement learning in the four room domain. Initiating conditions are shown in blue, the goal in green.