

Theoretically-Grounded Policy Advice from Multiple Teachers in Reinforcement Learning Settings with Applications to Negative Transfer

Yusen Zhan¹, Haitham Bou Ammar², and Matthew E. Taylor¹

¹Washington State University, Pullman, Washington

²Princeton University, Princeton, New Jersey

yusen.zhan@wsu.edu, hammar@princeton.edu, taylor@m@eecs.wsu.edu

Abstract

Policy advice is a transfer learning method where a student agent is able to learn faster via advice from a teacher. However, both this and other reinforcement learning transfer methods have little theoretical analysis. This paper formally defines a setting where multiple teacher agents can provide advice to a student and introduces an algorithm to leverage both autonomous exploration and teacher’s advice. Our regret bounds justify the intuition that good teachers help while bad teachers hurt. Using our formalization, we are also able to quantify, for the first time, when negative transfer can occur within such a reinforcement learning setting.

1 Introduction

Reinforcement Learning (RL) has become a popular framework for autonomous behavior generation from limited feedback [Sutton and Barto, 1998]. Typical RL methods learn in isolation increasing their learning times and sample complexities. Transfer learning aims to significantly improve learning by providing informative knowledge from an external source. The source of such knowledge varies from source agents to humans providing advice [Erez and Smart, 2008; Taylor and Stone, 2009]. In this paper, we focus on a framework referred to as action advice or the *advice model* [Taylor *et al.*, 2014]. Here, the agent (i.e., student), learning in a task, has access to a teacher (another agent or human) which can provide action suggestions to facilitate learning. Given “good-enough” teachers, such advice models have shown multiple benefits over standard RL techniques. For example, others [Taylor *et al.*, 2014; Zimmer *et al.*, 2014] show reduced learning times and sample complexities for successful behavior.

These methods, however, suffer from two main drawbacks. First, validation results are empirical in nature and not formally-grounded. We do not have fundamental understanding of these methods. Consequently, it is difficult to formally comprehend why these methods work. Second, most of these techniques require the availability of a “good-enough” (optimal) teacher to benefit the student. Unfortunately, access to such teachers is difficult in a variety of complex domains,

reducing the applicability of policy advice in real-world settings.

In this paper, we remedy the aforementioned drawbacks by proposing a new framework for policy advice. Our method formally generalizes current single-teacher advice models to the multi-teacher setting. Our algorithm also remedies the need for optimal teachers by exploiting both the student’s and the teacher’s knowledge. Even if the teacher is not optimal, a student, using our algorithm, is still capable of acquiring optimal behavior in a task; a property not supported by some state-of-the-art methods, e.g., learning from demonstration. We theoretically and empirically analyze the performance of the proposed method and derive, for the first time, regret bounds quantifying the successfulness of action advice. We also provide theoretical justification for current methods (i.e., single-teacher models) as special case of our formulation¹. Our contributions can be summarized as:

- defining (formally) multi-teacher advice models,
- introducing novel algorithms leveraging teacher and student knowledge,
- deriving the regret analysis showing reduced sample complexities,
- deriving theoretical guarantees for single teacher advice models, and
- quantifying negative transfer under such advice model.

Interestingly, these theoretical results justify a well-known intuition inherent to advice models: “good teachers help while bad teachers hurt.” The results show that students can still achieve optimal behavior when being advised by bad teachers. They, however, pay an extra cost in terms of their learning times or sample complexities, relative to an optimal teacher. This should inspire researchers to adopt high quality teacher policies or avoid “bad teachers” if possible.

Given our formalization, we also derive a relation to negative transfer. We quantify, for the first time, the occurrence of negative transfer in action advice models, shedding the light on failure modes of these methods. Consequently, these results yield two claims about transfer learning. First, high quality transfer knowledge may still cause negative transfer when the target algorithm is able to outperform the source

¹The full version: <https://arxiv.org/pdf/1604.03986v1>

knowledge. Second, expert knowledge is important for the researchers to determine whether or not to transfer because evaluation of the transfer knowledge is usually expensive (it is equivalent to evaluating the teacher policy in the target MDP).

2 Preliminaries

2.1 Online Reinforcement Learning & Regret Model

In RL, an agent must sequentially select actions to maximize its total expected return. Such problems are formalized as a Markov decision Process (MDP), defined as $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R} \rangle$, where \mathcal{S} and \mathcal{A} denote the finite state and actions spaces with a total size of $|\mathcal{S}|$ and $|\mathcal{A}|$ respectively, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ represents the probability transition kernel describing the task dynamics, and $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function quantifying the performance of the agent. The total expected return of an agent following an algorithm, \mathcal{G} , to compute the optimal action-selection rule from a starting state $s \in \mathcal{S}$ after T time steps is defined as:

$$\mathcal{R}^{\mathcal{G}}(s, T) = \mathbb{E} \left[\sum_{t=0}^T \mathcal{R}(s_t, a_t) \right], \quad (1)$$

with $s_t \in \mathcal{S}$ and $a_t \in \mathcal{A}$. The goal is to determine an optimal policy, $\pi^* : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the total expected return.

Regret Model: Similar to standard online learning, we quantify the performance of the algorithm, \mathcal{G} , by measuring its regret with respect to optimality. We define the regret of a state s after T time steps in terms of the expected reward as:

$$\Delta^{\mathcal{G}}(s, T) = \lambda^* T - \mathcal{R}^{\mathcal{G}}(s, T), \quad (2)$$

where λ^* is the optimal reward acquired by following an optimal algorithm \mathcal{G}^* at each time step. In the general case when no reachability assumptions are imposed, it is easy to construct MDPs in which algorithms suffer high regret. Following Puterman [2005], we remedy this problem by considering weakly-communicating MDPs² defined as follows.

Definition 1. An MDP is called *weakly communicating* in such a case where the state set \mathcal{S} can be decomposed into two subsets, \mathcal{S}_1 and \mathcal{S}_2 . In \mathcal{S}_1 any state is reachable from every other state under a deterministic policy, π , while states in \mathcal{S}_2 are transient under all policies.

The optimal gain, λ^* in Equation 2, is state independent. That is, any $s \in \mathcal{S}$, shares the same optimal expected reward [Puterman, 2005], which can be solved for using:

$$\mathbf{h}^* + \lambda^* \mathbf{e} = \max_{\mathbf{a} \in \mathcal{A}} \{ \mathcal{R}(s, \mathbf{a}) + \mathcal{P}_{s, \mathbf{a}}^T \mathbf{h}^* \},$$

where \mathbf{h}^* is an $|\mathcal{S}|$ dimensional vector typically referred to as the bias vector, $\mathcal{P}_{s, \mathbf{a}}$ denotes the probability to transition from s applying action \mathbf{a} , and $\mathbf{e} \in \mathbb{R}^{|\mathcal{S}|}$ is a unit dimensional vector. When needed, we explicitly write the dependency of λ^* and \mathbf{h}^* as $\mathbf{h}^*(s; \mathcal{M})$ and $\lambda^*(\mathcal{M})$. We also define the span of \mathbf{h} as: $\text{sp}(\mathbf{h}) = \max_{s \in \mathcal{S}} \mathbf{h}(s) - \min_{s \in \mathcal{S}} \mathbf{h}(s)$.

²Please note that weakly-communicating MDPs are considered the most general among subclasses of MDPs, see Puterman [2005].

Algorithm 1 REGAL.C: Constrained Optimization

Input: parameter H , dataset \mathcal{D}_i and current time T

Output: $\hat{\pi}_{i+1}$

- 1: $t_i =$ current time T
- 2: Use \mathcal{D}_i to update the state transition probabilities by

$$\hat{P}_{s, \mathbf{a}}^t(s') = \frac{N(s, \mathbf{a}, s'; t)}{\max\{N(s, \mathbf{a}; t), 1\}} \quad (3)$$

- 3: With $t = t_i$, \mathcal{M}^i is the set of MDPs s.t.

$$\|P_{s, \mathbf{a}} - \hat{P}_{s, \mathbf{a}}^t\|_1 \leq \sqrt{\frac{12|\mathcal{S}| \log(2|\mathcal{A}|t/\delta)}{\max\{N(s, \mathbf{a}; t), 1\}}}$$

- 4: Select $M^i \in \mathcal{M}^i$ by following optimization equation over $\forall M \in \mathcal{M}^i$,

$$\max \lambda^*(M) \text{ s.t. } \text{sp}(\mathbf{h}^*(M)) \leq H$$

- 5: $\hat{\pi}_{i+1} =$ average reward optimal policy for M^i (value iteration)

- 6: **return** $\hat{\pi}_{i+1}$
-

Finally, we follow Bartlett and Tewari [2009] to define reachability in weakly communicating MDPs using the *one-way diameter*: $\text{diam}_{\text{one-way}}(\mathcal{M}) = \max_{s \in \mathcal{S}} \min_{\pi} T_{s_1 \rightarrow \bar{s}}^{\pi}$, with $T_{s_1 \rightarrow \bar{s}}^{\pi}$ being the expected number of steps needed for reaching $\bar{s} = \arg \max_{s \in \mathcal{S}} \mathbf{h}^*(s; \mathcal{M})$ from $s_1 \in \mathcal{S}$.

2.2 Algorithms for Weakly-Communicating MDPs

REGAL.C is an on-line algorithm for weakly communicating MDPs developed by Bartlett and Tewari [2009]. The basic idea is that the REGAL.C can estimate the true MDP with high probability in order to learn an ϵ -optimal policy with high probability. Let $N(s, \mathbf{a}, s'; t)$ be the number of state-action-state triples (s, \mathbf{a}, s') that have been visited at time t . Further, let t_i to denote the initial time of the iteration i . For brevity, we use $N_i(s, \mathbf{a}, s')$ and $N_i(s, \mathbf{a})$ to denote $N(s, \mathbf{a}, s'; t_i)$ and $N(s, \mathbf{a}; t_i)$ at iteration i . We also use $v_i(s, \mathbf{a}) = N_{i+1}(s, \mathbf{a}) - N_i(s, \mathbf{a})$ to denote the number of times a state-action pair (s, \mathbf{a}) is visited during iteration i . For each iteration i , REGAL acquires a dataset \mathcal{D}_i as input and updates the transition probability (see Equation 3). It then constructs a set of MDPs \mathcal{M}^i to select from using $\max \lambda^*(M)$ s.t. $\text{sp}(\mathbf{h}^*(M)) \leq H$, where H is the upper bound on the span $\text{sp}(\mathbf{h}^*(\mathcal{M}))$. Given the MDP, REGAL.C uses value iteration for acquiring the optimal policy. These steps are summarized in Algorithm 1.

2.3 Single Teacher Advice Model

The single teacher advice model is a framework in which a student learning in an environment benefits from a teacher's advice to speed-up learning. We define such a framework as the tuple of $\langle \pi^{\mathcal{S}}, \mathbf{b}, \mathcal{S}, \mathbf{f}_d \rangle$. Here, $\pi^{\mathcal{S}}$ denotes the teacher's policy, \mathbf{b} represents the budget constraining the teacher's advice, \mathcal{S} is the student, and \mathbf{f}_d is a function controlling the advice from the teacher to the student. Apart from considering single teacher models, previous work assumed optimal teachers where students always execute recommended actions. It is

easy to construct complex settings in which access to optimal teachers is difficult. Consequently, we extend these works to the more realistic settings of sub-optimal teachers, as we detail later.

3 Multiple Teacher Advice Model

In this section we start by extending the single teacher model of Taylor *et al.* [2014] to the multiple *non-optimal* teacher setting. Our advice model for m teachers is defined as the tuple $\langle \Pi, \mathcal{B}, \mathcal{G}, \mathbf{f}_d \rangle$, where $\Pi = \{\pi_1^{\mathcal{S}}, \pi_2^{\mathcal{S}}, \dots, \pi_m^{\mathcal{S}}\}$ is the set of $m \in \mathbb{N}$ teacher policies, and $\mathcal{B} = \{\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_m\}$ denotes the set of budgets. It is easy to see that in case $\Pi = \{\pi^{\mathcal{S}}\}$ and $\mathcal{B} = \{\mathbf{b}\}$, we can easily recover the special case single teacher model. We also generalize the work of Taylor *et al.* [2014] by making no restrictive assumptions on the optimality of any of the teachers. We measure the performance of the teacher with respect to a base policy $\pi^{\mathcal{B}}$ in terms of regret:

Definition 2. *Given a teacher’s policy, $\pi^{\mathcal{S}} \in \Pi$, and a base policy $\pi^{\mathcal{B}}$, then the regret of following $\pi^{\mathcal{S}}$ is related to that acquired by following $\pi^{\mathcal{B}}$ using:*

$$\Delta^{\mathcal{S}}(s, T) = \rho \Delta^{\mathcal{B}}(s, T),$$

where $\rho \geq 0$ denotes the regret ratio, $\Delta^{\mathcal{S}}(s, T) = \lambda^* T - \mathcal{R}^{\mathcal{S}}(s, T)$ and $\Delta^{\mathcal{B}}(s, T) = \lambda^* T - \mathcal{R}^{\mathcal{B}}(s, T)$.

The above definition captures the three interesting cases quantifying the performance of an advice-based algorithm. If the teacher is optimal, i.e., when $\Delta^{\mathcal{S}}(s, T) = 0$, ρ is also 0. In case $0 < \rho \leq 1$, then $\Delta^{\mathcal{S}}(s, T) \leq \Delta^{\mathcal{B}}(s, T)$ indicating the the teacher’s policy is at least as good as the base policy $\pi^{\mathcal{B}}$. Finally, when $\rho > 1$, $\Delta^{\mathcal{S}}(s, T) > \Delta^{\mathcal{B}}(s, T)$ implying the underperformance of the teacher. Consequently, with the correct choice of the teacher by ρ one can still achieve successful advice even in such a generalized setting.

4 Efficient Multi-teacher Advice

In this section, we propose a new algorithm which combines the advice policy and the MDPs information collected so far. This allows for an accurate framework outperforming state-of-the-art techniques for policy advice. On a high level, our algorithm consists of three main steps. First, a combined policy is constructed based on multiple teachers. Second, data depending on both teacher’s advice as well as MDP information is collected. Third, a new policy is computed online.

Next, we outline each of the three steps and describe our novel algorithm. Having achieved an accurate advice model, we then rigorously analyze the theoretical aspects of our method and show a decrease in sample complexities compared to current techniques.

4.1 The Grand-Teacher

Our method of policy advice constructs a grand teacher combining all teacher policies in a meta-policy. To construct the grand-teacher, we use an ensemble method and design two meta-policy variations: online and offline-constructions. Next, we detail each of the two variations.

Online Grand-Teacher: In the *online* construction, whenever the student observes an unvisited state, $s \in \mathcal{S}$, each

Algorithm 2 Offline Construction of the Meta-Teacher

Input: The set of states in the MDP, \mathcal{S} .

- 1: **while** $\exists s \in \mathcal{S}$ is not visited **do**
 - 2: Follow a policy in the MDP
 - 3: **if** Current state s is not visited **then**
 - 4: Query all teachers for advice and select action a using Majority Vote.
 - 5: **return** $\pi^{\text{grand-teacher}}$
-

teacher provides its policy advice of the form $\pi_i^{\mathcal{S}}$, for all $i \in \{1, \dots, m\}$ with m being the total number of teachers. The student then selects and stores the majority action from all teachers for that state s . As far as budget is concerned, it is easy to see that we only require to know advice for each state in \mathcal{S} , thus $\mathbf{b}_1 = \dots = \mathbf{b}_m = |\mathcal{S}|$. Though easy to implement and test, the online construction suffers from the potentially unrealistic need for the continuous availability of online teachers.

Offline Grand-Teacher: To eliminates the need for an online teacher at each visit of a new state, the offline procedure traverses the states in the MDP for constructing the meta-advice policy. The main steps of this construction is summarized in Algorithm 2.

Note that Algorithm 2 is capable of constructing an offline meta-teacher but requires extra exploration in the MDP. We next show that $\mathcal{O}\left(|\mathcal{S}| \log \frac{|\mathcal{S}|}{\delta}\right)$ steps are enough to explore each state in the MDP with high probability:

Theorem 1 (Sample Complexity). *If Algorithm 2 independently and uniformly explores each state $s \in \mathcal{S}$, then with probability of at least $1 - \delta$, $\mathcal{O}\left(|\mathcal{S}| \log \frac{|\mathcal{S}|}{\delta}\right)$ steps are sufficient to visit each state at least one time.*

4.2 Multi-Teacher Advice Algorithm

To improve current methods and arrive at a more realistic advice framework, we now introduce our algorithm combining the grand-teacher’s policy and information attained by the student from the MDP.

Our algorithm is based on the following intuition. At the beginning of the learning process, a student requires guidance as it typically has little to no information of the task to be solved. As time progress and the student explores, the MDP can be effectively exploited for successful learning. Unfortunately, such a process is not well modeled using current methods. Here, we remedy this problem by introducing an algorithm which follows the teacher’s advice at the very beginning and then switches to a policy computed by an algorithm operating within the MDP. That is, the teacher guides the student at the beginning of the learning process and as the student gathers more experience, the teacher’s influence diminishes over time by switching into a policy computed by REGAL.C. The overall procedure is summarized in Algorithm 3. Note that our algorithm is inspired by DAGGER [Ross *et al.*, 2010] in the sense that policies are updated by collecting data using a mixture of action selection rules (i.e., student and teacher policies). Contrary to DAGGER, however, our method col-

Algorithm 3

Input: $\pi^{\mathfrak{S}}$ = the grand-teacher policy, $\hat{\pi}_1$ =any policy

Output: π , the ϵ -optimal policy

- 1: $T = 0$
 - 2: **for** $i = 1$ to m **do**
 - 3: Let $\pi_{i+1} = \beta_i \pi^{\mathfrak{S}} + (1 - \beta_i) \hat{\pi}_i$
 - 4: Follow π_i until T_i -steps
 - 5: Get dataset D_i
 - 6: $T = T + T_i$
 - 7: $\hat{\pi}_{i+1} = \text{REGAL.C}(D_i, T)$ {See Algorithm 1}
 - 8: **return** π_{m+1}
-

lects all trajectories opposed to only collecting inconsistent actions, allowing for more accurate and efficient updates.

To leverage both the teacher’s and learned policies, we set a mixed policy of the form $\pi_{i+1} = \beta_i \pi^{\mathfrak{S}} + (1 - \beta_i) \hat{\pi}_i$, for $0 \leq \beta_i \leq 1$ to guide the student’s dataset collection while allowing the teacher to fractionally control exploration needed to collect data at the next iteration. β should typically be set so as to decay exponentially over time. This decreases the student’s reliance on the teacher and allows it to exploit the knowledge gathered from the MDP to learn better behaving policies than that of the teacher. It is for this reason that our algorithm, contrary to other methods, does not impose any optimality restrictions on the teacher. Having collected the dataset, Algorithm 3 uses REGAL.C (Algorithm 1) to update $\hat{\pi}_i$.

4.3 Theoretical Guarantees

In this section we formally derive the regret exhibited by our algorithm. At a high level, we provide two theoretical results. In the first, we consider the general teacher case, while in the second we derive a corollary of the regret for optimal teachers. We show, for the first time, *better than constant* improvements compared to standard learners.

Theorem 2. *Assume Algorithm 3 is running for total T steps in a weakly communicating MPD \mathcal{M} starting from an initial state $\mathbf{s} \in \mathcal{S}$. Let H be a parameter such that $H \geq sp(\mathbf{h}^*(\mathcal{M}))$. Then, with a probability of at least $1 - \delta$, the total regret is given by: $\Delta(\mathbf{s}, T) =$*

$$\mathcal{O} \left((1 - \beta + \rho\beta) H |\mathcal{S}| \sqrt{|\mathcal{A}| T \log \frac{|\mathcal{A}| T}{\delta}} \right), \text{ where } \beta \in [0, 1]$$

such that $1 - \beta = \max_{1 \leq i \leq m} \{1 - \beta_i\}$, and $\rho \geq 0$ is the ratio between the teacher’s regret $\Delta^{\mathfrak{S}}$ and the regret exhibited by REGAL.C $\Delta^{\text{REGAL.C}}$ such that $\Delta^{\mathfrak{S}} \leq \rho \Delta^{\text{REGAL.C}}$.

Proof. Due to the space limits, we provide a proof sketch. The proof is based on the regret bound of REGAL.C. We introduce the regret ratio to reduce the grand-teacher’s regret to the REGAL.C’s regret. Then, we apply the Hoeffding’s inequality to arrive at the statement of the theorem. \square

Theorem 2 implies that the teacher improves learning as long as it is “good.” Namely, if $0 < \rho \leq 1$, $1 - \beta + \rho\beta \leq 1$, $\beta \in [0, 1]$ which implies the student can enjoy a fraction of REGAL.C’s regret. However, if $\rho > 1$, $1 - \beta + \rho\beta > 1$, the student suffers more regret than the original REGAL.C

algorithm. This justifies our intuition that good teachers assist learning while poor ones hamper learning. Moreover, if there exists prior knowledge that a teacher has poor performance, it would be better off for the student to neglect its advice as it will suffer extra regret.

If the teacher’s $\rho = 0$, we have the following Corollary:

Corollary 1. *If the teacher is optimal, then with at least a probability of $1 - \delta$ the total regret is given by: $\Delta(\mathbf{s}, T) =$*

$$\mathcal{O} \left((1 - \beta) H |\mathcal{S}| \sqrt{|\mathcal{A}| T \log \frac{|\mathcal{A}| T}{\delta}} \right).$$

Remark 1. *Please note that the above theoretical results are more than a constant improvement to the regret. Notice that β depends on the number of iterations which can be bounded by $|\mathcal{S}|$ and $|\mathcal{A}|$ of the input MDP \mathcal{M} [Auer et al., 2009]. Further, ρ depends on the input teacher’s policy which is also an input to Algorithm 3. Consequently, it can be shown that these regret improvements exceed simple constant bounds.*

5 Negative Transfer

To formalize the relation to negative transfer, we recognize that the regret ratio can be written as:

$$\rho = \frac{\Delta^{\mathfrak{S}}(\mathbf{s}, T)}{\Delta^{\mathfrak{B}}(\mathbf{s}, T)} = \frac{\lambda^* T - \mathcal{R}^{\mathfrak{S}}(\mathbf{s}, T)}{\lambda^* T - \mathcal{R}^{\mathfrak{B}}(\mathbf{s}, T)} \quad (4)$$

This suggests that we can estimate the ratio by calculating λ^* and $\mathcal{R}^{\pi}(\mathbf{s}, T)$, given a policy π . So, we use

$$\rho(\pi_1, \pi_2, T) = \frac{\lambda^* T - \mathcal{R}^{\pi_1}(\mathbf{s}, T)}{\lambda^* T - \mathcal{R}^{\pi_2}(\mathbf{s}, T)}$$

to denote the regret ratio between policy π_1 and π_2 until step T . At this stage, we define:

- Negative transfer from policy π_1 to π_2 until T steps: $\rho(\pi_1, \pi_2, T) > 1$.
- Positive transfer from policy π_1 to π_2 until T steps: $\rho(\pi_1, \pi_2, T) \leq 1$

To formalize negative transfer, our goal at this stage is to relate $\rho(\cdot)$ to a metric between source and target tasks. For that sake, we define: $d_t^s(\pi_s) = \hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T) - \hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)$, with $\hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T)$ and $\hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)$ being the agent’s estimates of the rewards in the source and the target after T steps. Consequently, an estimate $\hat{\rho}$ to ρ can be derived as:

$$\begin{aligned} \hat{\rho}(\pi_s, \pi_t, T) &= \frac{\lambda^* T - \hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T)}{\lambda^* T - \hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)} \\ &= \frac{\lambda^* T + \left(\hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T) - \hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T) \right) - \hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T)}{\lambda^* T - \hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)} \\ &= \frac{\lambda^* T + d_t^s(\pi_s) - \hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T)}{\lambda^* T - \hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)}. \end{aligned}$$

$\hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T)$ and $\hat{\mathcal{R}}_t^{\pi_t}(\mathbf{s}, T)$ can be bounded by the Empirical Bernstein bound [Audibert et al., 2007]. With a probability $1 - \delta$, we have

$$\left| \hat{\mathcal{R}}_t^{\pi_s}(\mathbf{s}, T) - \mathbb{E}_{\pi_s} \left[\sum_{t=0}^T \mathcal{R}_s(\mathbf{s}_t, \mathbf{a}_t) \right] \right| \leq \epsilon_1,$$

with $\epsilon_1 = \bar{\sigma} \sqrt{\frac{2 \log(3/\delta)}{n_s}} + \frac{6R_{max} \log(3/\delta)}{n_s}$, $\bar{\sigma} = \sqrt{1/n_s \sum_{i=1}^{n_s} (R_i - \bar{R})^2}$ is the standard deviation of the sample, we derive

$$\frac{\lambda^* T + d_t^s(\pi_s) - C_2}{\lambda^* T - C_4} \leq \hat{\rho}(\pi_s, \pi_t, T) \leq \frac{\lambda^* T + d_t^s(\pi_s) - C_1}{\lambda^* T - C_3} \quad (5)$$

with C_1, C_2, C_3 , and C_4 are constants. Consequently, for negative transfer:

$$\hat{\rho}(\pi_s, \pi_t, T) \geq \frac{\lambda^* T + d_t^s(\pi_s) - C_2}{\lambda^* T - C_4} > 1.$$

Then, assuming enough samples, negative transfer occurs if:

$$d_t^s(\pi_s) > \left\{ \mathbb{E}_{\pi_s} \left[\sum_{t=0}^T \mathcal{R}_s(s_t, \mathbf{a}_t) \right] - \mathbb{E}_{\pi_t} \left[\sum_{t=0}^T \mathcal{R}_t(s_t, \mathbf{a}_t) \right] \right\} \quad (6)$$

The condition sheds light on the negative transfer in the sense of metric notation and provides a formal way to determine negative transfer. First, if the condition in Eq. 6 holds after evaluation, researchers should avoid the source policy π_s to the target tasks since it may cause negative transfer. Second, if the researchers have enough expert knowledge about their working domain and transfer information, usually they can avoid this evaluation phase in practice. In short, Eq. 6 provides a formal way to understand negative transfer and justify the intuition (adopt high quality source knowledge and avoid bad teachers) in the transfer practice.

6 Experimental Results

Given the above theoretical successes, this section provides empirical validation on three domains:

Combination Lock: We use the domain described in Figure 2 which is a variation from [Whitehead, 1991]. The experimental setting follows the caption description.

Grid World is an RL benchmark in which an agent has to navigate an $m \times m$ grid world with the goal of reaching a goal state. We employ an 11×11 grid world with a four room layout as introduced in Sutton and Barto [1998]. The agent begins in the lower left corner of the map and navigates to the goal state being the upper right corner. To navigate, the agent has access (in each cell) to four actions transitioning it to the: north, south, west and east. Applying an action, it then transitions in that direction with a probability of 0.8 and in the other three with a probability of 0.2. In case the direction is blocked, the agent stays in the same state. Finally, the agent receives a reward of 0 once reaching the goal state and a reward of -1 in all others.

Block Dude is a game where an agent again navigates a maze to reach a goal state. Reaching the goal directly is impossible due to the presence of blocks restricting its movement. The agent, however, can move to the left, right, and upwards. To reach the goal state, it needs to pick-up blocks and relocate them in correct positions. We use the default level 1 BURLAP [MacGlashan, 2014] in which there are two blocks and 3×25 maze. The agent receives a reward of $+1$ in the goal state and a reward of -1 in all other states.

6.1 Experimental Setup & Results

To construct the grand teacher, we set the total number of teachers $k = 10$. For each teacher, the budget, b_i , is set to the total number of states. In Algorithm 3, the maximum number of iterations and the size of each dataset, \mathcal{D}_i , were set to 10 and 200, respectively. Values of $p^i = 0.5^i$ for $i = 1, \dots, 10$ were used to determine $\beta_i = p^i$. For Algorithm 1, the confidence δ was set to 0.8 and H to 1000. The optimal gain λ^* and the optimal bias vector \mathbf{h} can be approximated using the value function \mathcal{V} [Puterman, 2005]. Let \mathcal{V}^l be the value function at iteration l , $l = 0, 1, \dots$. The optimal gain $\lambda^* \approx \text{sp}(\mathcal{V}^{l+1} - \mathcal{V}^l)$, where $\text{sp}(\mathcal{V}^l) = \max_{s \in \mathcal{S}} \mathcal{V}^l(s) - \min_{s \in \mathcal{S}} \mathcal{V}^l(s)$ and the optimal bias vector $\mathbf{h}^* \approx \mathcal{V}^l - l\lambda^l$, when l is large enough. To smooth the natural variance in the student’s performance, each learning curve is averaged over 10 independent trials of student learning. To better evaluate our method, we adopt six experimental settings by considering different teachers and learning algorithms. For teachers, we consider three forms. The first, referred to as “optimal teacher” provides optimal actions and is used by the grand teachers. The second, referred to as “worst teacher” advises the student to take actions with the lowest Q-values, while the last randomly selects action suggestions from the set of allowed moves. We also compare our method to REGAL.C (no advice), optimal policy (without learning), and Azar’s method [Azar *et al.*, 2013]. Please note that Azar’s method can not converge to the optimal policy and suffers loss as its performance is restricted by the teacher.

Performance, measured by the average reward, is reported in Figure 1. First, it is clear that given optimal teachers, our method exactly traces the optimal policy achieving a regret of 0. It is also important to note that in all three domains, even if the teacher was not optimal, and contrary to current techniques, our method is capable of acquiring optimal behavior. This is achievable as our method allows for learning within the multiple teacher framework.

7 Related Work on Transfer Learning

Few theoretical results on transfer and policy advice have been achieved. Closest to this work is that in Taylor *et al.* [2014], where the authors only provide empirical validations to their approach without drawing on any theoretical analysis. Given the theoretical derivations in this paper, we in fact note that the method [Taylor *et al.*, 2014] is a special case of ours considering only one-teacher advice models.

Another method considering advice under multiple teachers is that in Azar *et al.* [2013]. Azar *et al.* propose a method capable of selecting the best policy from a set of teacher policies and derive sub-linear regret of the form $\mathcal{O}(\sqrt{T})$ with T being the total number of rounds. One drawback of their method, however, is the assumption of a “good-enough” teacher which can guide the student to optimality. Such a method may suffer huge regret if the overall quality of teacher policies is poor. It also can not obtain better policies than those of the teacher. Our algorithm remedies these problems by allowing agents to further improve, which gives them the opportunity to surpass the teacher’s performance.

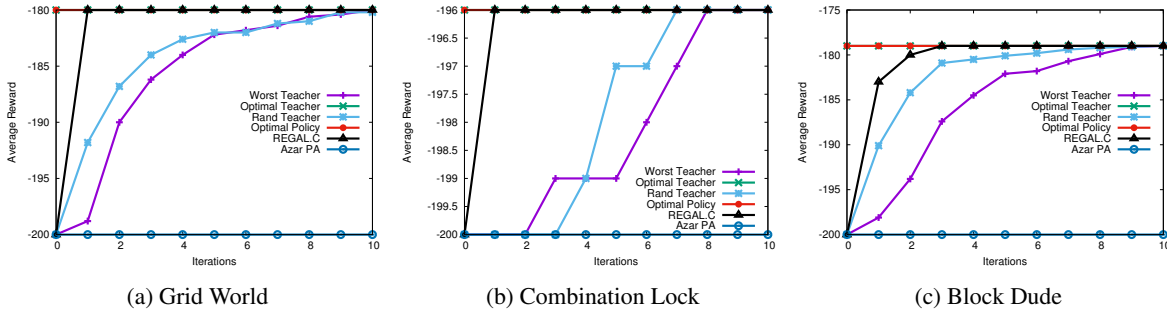


Figure 1: Our method with optimal teacher has similar performance as the optimal policy. And the REGAL algorithm (no advice) outperform random teacher and worst teacher group which justifies that the poorer teachers do harm the learning. Azar’s method depends on the quality of the teachers — when the teachers are very poor, the algorithm shows no learning.

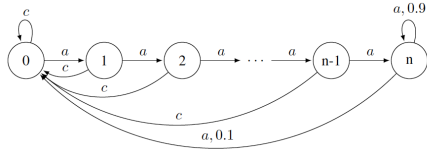


Figure 2: There are $n + 1$ states in the MDP. The last state has only one action and the rest have two. The agent receives reward -1 for all actions, except when taking action a in state n , $\mathcal{R}(n, a) = 1$. The agent stays in state n with probability 0.9 and goes to state 0 with probability 0.1 . The optimal policy is to take action a in each state. Since there are $n + 1$ states, the budget \mathfrak{B} is at least $n + 1$ to achieve zero regret.

Human advice is also a good source of policy advice. Usually, this method adopts the human advice as the teacher’s policy to improve the learning performance. However, these works focus on empirical validations [Cakmak and Lopes, 2012; Griffith *et al.*, 2013].

Probabilistic policy reuse is similar to our method in which the algorithm follows its own knowledge with probability $1 - \epsilon$ and teacher’s policy with probability ϵ [Fernández and Veloso, 2006]. However, ϵ is not decaying over time, making the algorithm divergent if teacher policies are not optimal. Cederborg *et al.* introduce a policy shaping algorithm using human teachers, but focus on providing rewards rather than action advice [Cederborg *et al.*, 2015]. Both of these works rely solely on empirical results.

Work on transfer for RL is also related to this paper, where we can consider policy advice as an instance of transferring from teachers to students [Lazaric, 2012]. Here, Ferrante *et al.*, for instance, propose a method to transfer high quality samples from source to target tasks using bi-simulation measures [Ferrante *et al.*, 2008]. Their method only transfers samples once, while our approach gradually provides advice to the student. Due to space constraints, we refer the reader to Taylor and Stone [2009] for a comprehensive survey.

Lifelong reinforcement learning has drawn significant attention to the transfer community recently. Brunskill and Li studied online discovery problems in a lifelong learning setting [Brunskill and Li, 2015]. Bou-Ammar *et al.* also studied

such a problem and introduced constraints on the policy to compute “safe” policies [Bou-Ammar *et al.*, 2015]. Contrary to these works, in this paper, we focus on the single agent setting operating within one task.

Finally, Learning from Demonstration [Argall *et al.*, 2009] (LfD) is also related to our work, but LfD usually assumes that the expert is optimal and the student only tries to mimic the expert.

8 Conclusion and Future Work

In this paper, we formally defined the multi-teacher advice model and introduced a new algorithm which leverages teacher and student’s own knowledge in the weakly communicating MDPs. We theoretically analyzed our algorithm and showed, for the first time, that the agent can achieve optimality even when starting from non-optimal teachers. Our results provide a theoretical justification for the intuition that “bad” teachers can hurt the learning process of the student. Also, we formally established the condition of negative transfer, shedding light on future transfer learning research, where for example, researchers can choose “good teachers” based on the Eq 6 and avoid negative transfer with prior expert knowledge.

In future, we plan on adopting other online reinforcement learning algorithms (e.g., REGAL.D [Bartlett and Tewari, 2009], R-max [Brafman and Tennenholtz, 2003], or E^3 [Kearns and Singh, 2002]) to replace REGAL.C. We will provide better methods to construct the “grand-teacher” without exploring the whole MDP. Also, extensions to large-scale MDPs may be an interesting direction for future research as well.

9 Acknowledgements

This research has taken place in part at the Intelligent Robot Learning (IRL) Lab, Washington State University. IRL research is supported in part by grants AFRL FA8750-14-1-0069, AFRL FA8750-14-1-0070, NSF IIS-1149917, NSF IIS- 1319412, USDA 2014-67021-22174, and a Google Research Award.

References

[Argall *et al.*, 2009] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot

- learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009.
- [Audibert *et al.*, 2007] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Tuning bandit algorithms in stochastic environments. In *Algorithmic Learning Theory*, pages 150–165. Springer, 2007.
- [Auer *et al.*, 2009] Peter Auer, Thomas Jaksch, and Ronald Ortner. Near-optimal regret bounds for reinforcement learning. In *Advances in neural information processing systems*, pages 89–96, 2009.
- [Azar *et al.*, 2013] Mohammad Gheshlaghi Azar, Alessandro Lazaric, Brunskill Emma, et al. Regret bounds for reinforcement learning with policy advice. In *ECML/PKDD-European Conference on Machine Learning and Principles and practice of knowledge Discovery in Database*, 2013.
- [Bartlett and Tewari, 2009] Peter L Bartlett and Ambuj Tewari. Regal: A regularization based algorithm for reinforcement learning in weakly communicating MDPs. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 35–42. AUAI Press, 2009.
- [Bou-Ammar *et al.*, 2015] Haitham Bou-Ammar, Rasul Tutunov, and Eric Eaton. Safe policy search for life-long reinforcement learning with sublinear regret. *CoRR*, abs/1505.05798, 2015.
- [Brafman and Tennenholtz, 2003] Ronen I Brafman and Moshe Tennenholtz. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *The Journal of Machine Learning Research*, 3:213–231, 2003.
- [Brunskill and Li, 2015] Emma Brunskill and Lihong Li. The online discovery problem and its application to life-long reinforcement learning. *CoRR*, abs/1506.03379, 2015.
- [Cakmak and Lopes, 2012] Maya Cakmak and Manuel Lopes. Algorithmic and human teaching of sequential decision tasks. In *AAAI Conference on Artificial Intelligence (AAAI-12)*, 2012.
- [Cederborg *et al.*, 2015] Thomas Cederborg, Ishaan Grover, Charles L. Isbell, and Andrea L Thomaz. Policy Shaping With Human Teachers. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [Erez and Smart, 2008] Tom Erez and William D Smart. What does shaping mean for computational reinforcement learning? In *Development and Learning, 2008. ICDL 2008. 7th IEEE International Conference on*, pages 215–219. IEEE, 2008.
- [Fernández and Veloso, 2006] Fernando Fernández and Manuela Veloso. Probabilistic policy reuse in a reinforcement learning agent. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 720–727. ACM, 2006.
- [Ferrante *et al.*, 2008] Eliseo Ferrante, Alessandro Lazaric, and Marcello Restelli. Transfer of task representation in reinforcement learning using policy-based proto-value functions. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*, pages 1329–1332. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [Griffith *et al.*, 2013] Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles Isbell, and Andrea L Thomaz. Policy shaping: Integrating human feedback with reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 2625–2633, 2013.
- [Kearns and Singh, 2002] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002.
- [Lazaric, 2012] Alessandro Lazaric. Transfer in reinforcement learning: a framework and a survey. In *Reinforcement Learning*, pages 143–173. Springer, 2012.
- [MacGlashan, 2014] James MacGlashan. The Brown-UMBC Reinforcement Learning and Planning (BURLAP) <http://burlap.cs.brown.edu/index.html>, 2014.
- [Puterman, 2005] Martin L Puterman. Markov decision processes: Discrete stochastic dynamic programming (wiley series in probability and statistics). 2005.
- [Ross *et al.*, 2010] Stéphane Ross, Geoffrey J Gordon, and J Andrew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. *arXiv preprint arXiv:1011.0686*, 2010.
- [Sutton and Barto, 1998] Richard S Sutton and Andrew G Barto. *Introduction to reinforcement learning*. MIT Press, 1998.
- [Taylor and Stone, 2009] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *The Journal of Machine Learning Research*, 10:1633–1685, 2009.
- [Taylor *et al.*, 2014] Matthew E. Taylor, Nicholas Carboni, Anestis Fachantidis, Ioannis Vlahavas, and Lisa Torrey. Reinforcement learning agents providing advice in complex video games. *Connection Science*, 26(1):45–63, 2014.
- [Whitehead, 1991] Steven D Whitehead. Complexity and cooperation in q-learning. In *Proceedings of the Eighth International Workshop on Machine Learning*, pages 363–367, 1991.
- [Zimmer *et al.*, 2014] Matthieu Zimmer, Paolo Viappiani, and Paul Weng. Teacher-Student Framework: a Reinforcement Learning Approach. In *AAMAS Workshop Autonomous Robots and Multirobot Systems*, 2014.