
Reinforcement and Imitation Learning via Interactive No-Regret Learning

Stéphane Ross J. Andrew Bagnell

stephaneross@cmu.edu dbagnell@ri.cmu.edu

The Robotics Institute
Carnegie Mellon University,
Pittsburgh, PA, USA

Abstract

Recent work has demonstrated that problems— particularly imitation learning and structured prediction— where a learner’s predictions influence the input-distribution it is tested on can be naturally addressed by an interactive approach and analyzed using no-regret online learning. **These approaches to imitation learning, however, neither require nor benefit from information about the cost of actions.** We extend existing results in two directions: first, we develop an interactive imitation learning approach that leverages cost information; second, we extend the technique to address reinforcement learning. The results provide theoretical support to the commonly observed successes of online approximate policy iteration. Our approach suggests a broad new family of algorithms and provides a unifying view of existing techniques for imitation and reinforcement learning.

1 Introduction

Imitation learning has become increasingly important in fields— notably robotics and game AI— where it is easier for an expert to demonstrate a behavior than to translate that behavior to code. [1] Perhaps surprisingly, it has also become central in developing predictors for complex output spaces, e.g. sets and lists [22], parse trees [10], image parsing [20, 25] and natural language understanding[12]. In these domains, a policy is trained to imitate an oracle on ground-truthed data. Iterative training procedures (*e.g.* DAGGER, SEARN, SMILE[24, 10, 21]) that interleave policy execution and learning have demonstrated impressive practical performance and strong theoretical guarantees that were not possible with batch supervised learning. **Most of these approaches to imitation learning, however, neither require nor benefit from information about the cost of actions;** rather they leverage only information provided about “correct” actions by the demonstrator.

While **iterative training corrects the compounding of error effect one sees in control and decision making applications, it does not address all issues that arise.** Consider, for instance, a problem of learning to drive near the edge of a cliff: methods like DAGGER consider all errors from agreeing with the expert driver equally. If driving immediately off the cliff makes the expert easy to imitate— because the expert simply chooses the go straight from then on— these approaches may learn that very poor policy. More generally, a method that only reasons about agreement with a demonstrator instead of the long term costs of various errors may poorly trade-off inevitable mistakes. Even a crude estimate of the cost-to-go (*e.g.* it’s very expensive to drive off the cliff)— may improve a learned policy’s performance at the user’s intended task.

SEARN, by contrast, **does** reason about cost-to-go, but uses rollouts from the current policy which can be impractical for imitation learning. SEARN additionally requires the use of stochastic policies.

We present a simple, general approach we term AGGREGATE (Aggregate Values to Imitate) that **leverages cost-to-go information in addition to correct demonstration**, and establish that previous

methods can be understood as **special cases of a more general no-regret strategy**. The approach provides much stronger guarantees than existing methods by providing a statistical *regret* rather than statistical *error* reduction. [8]

This general strategy of leveraging cost-sensitive no-regret learners can be extended to Approximate Policy Iteration (API) variants for reinforcement learning. We show that any no-regret learning algorithm can be used to develop stable API algorithms with guarantees as strong as any available in the literature. We denote the resulting algorithm NRPI. The results provide theoretical support to the commonly observed success of online policy iteration[29] despite a paucity of formal results: such online algorithms often enjoy no-regret guarantees or share similar stability properties. Our approach suggests a broad new family of algorithms and provides a unifying view of existing techniques for both imitation and reinforcement learning.

2 Imitation Learning with Cost-To-Go

2.1 Preliminaries

We consider in this work finite horizon¹ control problems in the form of a Markov Decision Process with states s and actions a . We assume the existence of a cost-function $C(s, a)$, bounded between 0 and 1, that we are attempting to optimize over a horizon of T decisions. We denote a class of policies Π mapping states² to actions.

We use $Q_t^\pi(s, a)$ to denote the expected future cost-to-go of executing action a in state s , followed by executing policy π for $t - 1$ steps. We denote by $d_\pi = \frac{1}{T} \sum_{t=1}^T d_\pi^t$ the time-averaged distribution over states induced by executing policy π in the MDP (d_π^t is the distribution of states at time t induced by executing policy π). The overall performance metric of total cost of executing π for T -steps is denoted $J(\pi) = \sum_{t=1}^T \mathbb{E}_{s \sim d_\pi^t} [C(s, \pi(s))]$. We assume system dynamics are either unknown or complex enough that we typically have only sample access to them. The resulting setting for learning policies by demonstration– or learning policies by approximate policy iteration– are **not** typical *i.i.d.* supervised learning problems as the learned policy strongly influences its own test distribution rendering optimization difficult.

2.2 Algorithm: AGGREGATE

We describe here a simple extension of the DAGGER technique of [24] that learns to choose actions to minimize the cost-to-go of the expert rather than the zero-one classification loss of mimicking its actions. In simplest form, on the first iteration AGGREGATE collects data by simply observing the expert perform the task, and in each trajectory, at a uniformly random time t , explores an action a in state s , and observes the cost-to-go Q of the expert after performing this action. (See Algorithm 1 below.)³

Each of these steps generates a **cost-weighted training example** (s, t, a, Q) [19] and AGGREGATE trains a policy $\hat{\pi}_2$ to minimize the expected cost-to-go on this dataset. At each following iteration n , AGGREGATE collects data through interaction with the learner as follows: for each trajectory, begin by using the current learner’s policy $\hat{\pi}_n$ to perform the task, interrupt at a uniformly random time t , explore an action a in the current state s , after which control is provided back to the expert to continue up to time-horizon T . This results in new examples of the cost-to-go of the expert (s, t, a, Q) , under the distribution of states visited by the current policy $\hat{\pi}_n$. This new data is aggregated with all previous data to train the next policy $\hat{\pi}_{n+1}$; more generally, this data can be used by a no-regret online learner to update the policy and obtain $\hat{\pi}_{n+1}$. This is iterated for some number of iterations N and the best policy found is returned. We optionally allow the algorithm to continue executing the expert’s actions with small probability β_n , instead of always executing $\hat{\pi}_n$, up to the random time t where an action is explored and control is shifted to the expert. The general AGGREGATE is detailed in Algorithm 1.

Main difference

Otherwise identical to Dagger

¹All our results can be easily extended to the infinite discounted horizon setting.

²More generally features of the state (and potentially time)– our derivations do not require full observability and hence carry over to featurized state of POMDPs.

³This cost-to-go may be estimated by rollout, or provided by the expert.

Observing the expert’s cost-to-go indicates how much cost we might expect to incur in the future if we take an action now and then can behave as well (or nearly so) as the expert henceforth. Under the assumption that the expert is a good policy, and that the policy class Π contains similarly good policies, this provides a rough estimate of what good policies in Π will be able to achieve at future steps. **By minimizing this cost-to-go at each step, we will choose policies that lead to situations where incurring low future cost-to-go is possible.** **For instance, we will be able to observe that if some actions put the expert in situations where falling off a cliff or crash is inevitable then these actions must be avoided at all costs in favor of those where the expert is still able to recover.**

Algorithm 1 AGGREGATE: Imitation Learning with Cost-To-Go

Initialize $\mathcal{D} \leftarrow \emptyset, \hat{\pi}_1$ to any policy in Π .
for $i = 1$ **to** N **do**
 Let $\pi_i = \beta_i \pi^* + (1 - \beta_i) \hat{\pi}_i$ #Optionally mix in expert’s own behavior.
 Collect m data points as follows:
 for $j = 1$ **to** m **do**
 Sample uniformly $t \in \{1, 2, \dots, T\}$.
 Start new trajectory in some initial state drawn from initial state distribution
 Execute current policy π_i up to time $t - 1$.
 Execute some exploration action a_t in current state s_t at time t
 Execute expert from time $t + 1$ to T , and **observe estimate of cost-to-go \hat{Q}** starting at time t
 end for
 Get dataset $\mathcal{D}_i = \{(s, t, a, \hat{Q})\}$ of states, times, actions, with expert’s cost-to-go.
 Aggregate datasets: $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_i$.
 Train cost-sensitive classifier $\hat{\pi}_{i+1}$ on \mathcal{D}
 (Alternately: use any online learner on the data-sets \mathcal{D}_i in sequence to get $\hat{\pi}_{i+1}$)
end for
Return best $\hat{\pi}_i$ on validation.

In AGGREGATE the problem of choosing the sequence of policies $\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_N$ over iterations is viewed as an online cost-sensitive classification problem. Our analysis below demonstrates that any no-regret algorithm on such problems can be used to update the sequence of policies and provide strong guarantees. To achieve this, **when the policy class Π is finite, randomized online learning algorithms like *weighted majority* [9] may be used.** When dealing with infinite policy classes (e.g. all linear classifiers), no-regret online cost-sensitive classification is not always computationally tractable. Instead, typically reductions of cost-sensitive classification to regression or ranking problems as well as convex upper bounds [8] on the classification loss lead to efficient no-regret online learning algorithms (e.g. gradient descent). The algorithm description suggests as the “default” learning strategy a **(Regularized)-Follow-The-Leader** online learner: it attempts to learn a good classifier for **all** previous data. This strategy for certain loss function (notably strongly convex surrogates to the cost-sensitive classification loss) [9] and any sufficiently stable batch learner [23, 27] ensures the no-regret property. It also highlights why the approach is likely to be particularly stable across rounds of interaction.

2.3 Training the Policy to Minimize Cost-to-Go

In standard “full-information” cost-sensitive classification, a cost vector is provided for each data-point in the training data that indicates the cost of predicting each class or label for this input. In our setting, **that implies for each sampled state we receive a cost-to-go estimate/rollout for all actions.** Training the policy at each iteration then simply corresponds to solving a cost-sensitive classification problem. That is, if we collect a dataset of m samples, $\{(s_i, t_i, \hat{Q}_i)\}_{i=1}^m$, **where \hat{Q}_i is a cost vector** of cost-to-go estimates for each action in state s_i at time t_i , then we solve the cost-sensitive classification problem: $\arg \min_{\pi \in \Pi} \sum_{i=1}^m \hat{Q}_i(\pi(s_i, t_i))$. Reductions of cost-sensitive classification to convex optimization problems can be used like weighted multi-class Support Vector Machines or ranking [8], to obtain problems that can be optimized efficiently while still guaranteeing good performance at this cost-sensitive classification task.

For instance, a simple approach is to transform this into an *argmax regression problem*: i.e., $\pi_n(s, t) = \arg \min_{a \in A} Q_n(s, t, a)$, for Q_n the learned regressor at itera-

tion n that minimizes the squared loss at predicting the cost-to-go estimates: $Q_n = \arg \min_{Q \in \mathcal{Q}} \sum_{(s_i, t_i, a_i, \hat{Q}_i) \in D} (Q(s_i, t_i, a_i) - \hat{Q}_i)^2$, where D is the dataset of all collected cost-to-go estimates so far, and \mathcal{Q} the class of regressors considered (e.g. linear regressors). This approach also naturally handles the more common situation in imitation learning where we only have *partial information* for a particular action chosen at a state. Alternate approaches include importance weighting techniques to transform the problem into a standard cost-sensitive classification problem [14, 11] and other online learning approaches meant to handle “bandit” feedback.

Local Exploration in Partial Information Setting In the *partial information* setting we must also select which action to explore for an estimate of cost-to-go. The uniform strategy is simple and effective but inefficient. The problem may be cast as a *contextual bandit problem* [2, 6] where features of the current state define the context of exploration. These algorithms, by choosing more carefully than at random, may be significantly more sample efficient. In our setting, in contrast with traditional bandit settings, we care only about the final learned policy and not the cost of explored actions along the way. Recent work [3] may be more appropriate for improving performance in this case. Many contextual bandit algorithms require a finite set of policies Π [2] or full realizability [18], and this is an open and very active area of research that could have many applications here.

2.4 Analysis

We analyze AGGREGATE, showing that the no-regret property of online learning procedures can be leveraged in this interactive learning procedure to obtain strong performance guarantees. Our analysis seeks to answer the following question: how well does the learned policy perform if we can repeatedly identify good policies that incur cost-sensitive classification loss competitive with the expert demonstrator on the aggregate dataset we collect during training?

Our analysis of AGGREGATE relies on connecting the iterative learning procedure with the (adversarial) online learning problem [9] and using the no-regret property of the underlying cost-sensitive classification algorithm choosing policies $\hat{\pi}_{1:N}$. Here, the *online learning problem* is defined as follows: at each iteration i , the learner chooses a policy $\hat{\pi}_i \in \Pi$ that incurs loss ℓ_i chosen by the adversary, and defined as $\ell_i(\pi) = \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \pi)]$ for $U(1:T)$ the uniform distribution on the set $\{1, 2, \dots, T\}$, $\pi_i = \beta_i \pi^* + (1 - \beta_i) \hat{\pi}_i$ and Q^* the cost-to-go of the expert. We can see that AGGREGATE at iteration i is exactly collecting a dataset \mathcal{D}_i , that provides an empirical estimate of this loss ℓ_i .

requires expert restarts

Let $\epsilon_{\text{class}} = \min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, a) - \min_a Q_{T-t+1}^*(s, a)]$ denote the minimum expected cost-sensitive classification regret achieved by policies in the class Π on all the data over the N iterations of training. Denote the online learning average regret of the sequence of policies chosen by AGGREGATE, $\epsilon_{\text{regret}} = \frac{1}{N} [\sum_{i=1}^N \ell_i(\hat{\pi}_i) - \min_{\pi \in \Pi} \sum_{i=1}^N \ell_i(\pi)]$.

We provide guarantees for the “uniform mixture” policy $\bar{\pi}$, that at the beginning of any trajectory samples a policy π uniformly randomly among the policies $\{\hat{\pi}_i\}_{i=1}^N$ and executes this policy π for the entire trajectory. This immediately implies good performance for the best policy $\hat{\pi}$ in the sequence $\hat{\pi}_{1:N}$, i.e. $J(\hat{\pi}) = \min_{i \in 1:N} J(\hat{\pi}_i) \leq J(\bar{\pi})$, and the last policy $\hat{\pi}_N$ when the distribution of visited states converges over the iterations of learning.

Assume the cost-to-go of the expert Q^* is non-negative and bounded by Q_{\max}^* , and $\beta_i \leq (1 - \alpha)^{i-1}$ for all i for some constant α ⁴. Then the following holds in the infinite sample case (i.e. if at each iteration of AGGREGATE we collected an arbitrarily large amount of data by running the current policy):

Theorem 2.1. *After N iterations of AGGREGATE:*

similar to Dagger

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi^*) + T[\epsilon_{\text{class}} + \epsilon_{\text{regret}}] + O\left(\frac{Q_{\max} T \log T}{\alpha N}\right).$$

Thus if a no-regret online algorithm is used to pick the sequence of policies $\hat{\pi}_{1:N}$, then as the number of iterations $N \rightarrow \infty$:

$$\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi^*) + T\epsilon_{\text{class}}$$

⁴The default parameter-free version of AGGREGATE corresponds to $\alpha = 1$, using $0^0 = 1$.

The proof of this result is presented in the Appendix. This theorem indicates that after sufficient iterations, AGGREGATE will find policies that perform the task nearly as well as the demonstrator if there are policies in Π that have small cost-sensitive classification regret on the aggregate dataset (i.e. policies with cost-sensitive classification loss not much larger than that of the bayes-optimal one on this dataset). Note that non-interactive supervised learning methods are unable to achieve a similar bound which degrades only linearly with T and the cost-sensitive classification regret. [24].

behavioural cloning

The analysis above abstracts away the issue of action exploration and learning from finite data. These issues come into play in a sample complexity analysis. Such analyses depend on many factors such as the particular reduction and exploration method. When reductions of cost-sensitive classification to simpler regression/ranking/classification [5] problems are used, our results can directly relate the task performance of the learned policy to the performance on the simpler problem. To illustrate how such results may be derived, we provide a result for the special case where actions are explored uniformly at random and the reduction of cost-sensitive classification to regression is used.

In particular, if $\hat{\epsilon}_{\text{regret}}$ denotes the empirical average online learning regret on the training regression examples collected over the iterations, and $\hat{\epsilon}_{\text{class}}$ denotes the empirical regression regret of the best regressor in the class on the aggregate dataset of regression examples when compared to the bayes-optimal regressor, we have that:

Theorem 2.2. *N iterations of AGGREGATE, collecting m regression examples (s, a, t, Q) per iteration, guarantees that with probability at least $1-\delta$:*

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi^*) + 2\sqrt{|A|T\sqrt{\hat{\epsilon}_{\text{class}} + \hat{\epsilon}_{\text{regret}} + O(\sqrt{\log(1/\delta)/Nm})}} + O\left(\frac{Q_{\max}T \log T}{\alpha N}\right).$$

Thus if a no-regret online algorithm is used to pick the sequence of regressors $\hat{Q}_{1:N}$, then as the number of iterations $N \rightarrow \infty$, with probability 1:

$$\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi^*) + 2\sqrt{|A|T\sqrt{\hat{\epsilon}_{\text{class}}}}$$

The detailed proof is presented in the Appendix. This result demonstrates how the task performance of the learned policies may be related all the way down to the regret on the regression loss at predicting the observed cost-to-go during training. In particular, it relates task performance to the square root of the online learning regret, on this regression loss, and the regression regret of the best regressor in the class to the bayes-optimal regressor on this training data.⁵

2.5 Discussion

AGGREGATE as a reduction: AGGREGATE can be interpreted as a *regret* reduction of imitation learning to no-regret online learning.⁶ We present a statistical regret reduction, as here, performance is related directly to the online, cost-sensitive classification regret on the aggregate dataset. By minimizing cost-to-go, we obtain regret reduction, rather than a weaker error reduction as in [24] when simply minimizing immediate classification loss.

Limitations: As just mentioned, in cases where the expert is much better than any policy in Π , the expert's cost-to-go may be a very optimistic estimate of the true future cost after taking a certain action. The approach may fail to learn policies that perform well, even if policies that can perform the task (albeit not as well as the expert) exist in the policy class. Consider again the driving scenario, where one may choose one of two roads to reach a goal: a shorter route that involves driving on a very narrow road next to cliffs on either side, and a longer route which is safer and risks no cliff. If in this example, the expert takes the short route faster and no policy in the class Π can drive without falling on the narrow road, but there exists policies that can take the longer road and safely reach the

⁵The appearance of the square root is particular to the use of this reduction to squared-loss regression and implies relative slow convergence to good performance. Other cost-sensitive classification reductions and regression losses (e.g. [17, 7]) do not introduce this square root and still allow efficient learning.

⁶Unfortunately *regret* here has two different meanings common in the literature: the first is in the statistical sense of doing nearly as well as the Bayes-optimal predictor. [8] The second use is in the online, adversarial, no-regret sense of competing against any hypothesis on a particular sequence without statistical assumptions. [9]

goal, this algorithm would fail to find these policies. The reason for this is that, as we minimize cost-to-go of the expert, we would always favor policies that heads toward the shorter narrow road. But once we are on that road, inevitably at some point we will encounter a scenario where no policies in the class can predict the same low cost-to-go actions as the expert (i.e. making ϵ large in the previous guarantee). The end result is that we may learn a policy that takes the short narrow road and eventually falls off the cliff, in these pathological scenarios.

Comparison to SEARN: AGGREGATE shares deep commonalities with SEARN but by providing a reduction to online learning allows much more general schemes to update the policy at each iteration that may be more convenient or efficient rather than the particular stochastic mixing update of SEARN. These include deterministic ones that provide upper convex bounds on performance. In fact, SEARN may be thought as a particular case of AGGREGATE, where the policy class is the set of distributions over policies, and the online coordinate descent algorithm (Frank-Wolfe) of [13] is used to update the distribution over policies at each iteration. Both collect data in a similar fashion at each iteration by executing the current policy up to a random time and then collecting cost-to-go estimates for explored actions in the current state. A distinction is that SEARN collects cost-to-go of the **current** policy after execution of the random action, instead of the cost-to-go of the expert. Interestingly, SEARN is usually used in practice with the approximation of collecting cost-to-go of the expert [10], rather than the current policy. Our approach can be seen as providing a theoretical justification for what was previously a heuristic.

finite set?

3 Reinforcement Learning via No-Regret Policy Iteration

A relatively simple modification of the above approach enables us to develop a family of sample-based approximate policy iteration algorithms. Conceptually, we make a swap: from executing the current policy and then switching to the expert to observe a cost-to-go; to, executing the expert policy while collecting cost-to-go of the learner's current policy. We denote this family of algorithms *No-Regret Policy Iteration* NRPI and detail and analyze it below.

This alternate has similar guarantees to the previous version, but may be preferred when no policy in the class is as good as the expert or when only a distribution of "important states" is available. In addition it can be seen to address a general model-free reinforcement learning setting where we simply have a state exploration distribution we can sample from and from which we collect examples of the current policy's cost-to-go. This is similar in spirit to how Policy Search by Dynamic Programming (PSDP) [4, 28] proceeds, and in some sense, the algorithm we present here provides a generalization of PSDP. However, by learning a stationary policy instead of a non-stationary policy, NRPI can generalize across time-steps and potentially lead to more efficient learning and practical implementation in problems where T is large or infinite.

Following [4, 15] we assume access to a state exploration distribution ν_t for all times t in $1, 2, \dots, T$. As will be justified by our theoretical analysis, these state exploration distributions should ideally be (close to) that of a (near-)optimal policy in the class Π . In the context where an expert is present, then this may simply be the distribution of states induced by the expert policy, i.e. $\nu_t = d_{\pi^*}^t$. In general, this may be the state distributions induced by some base policy we want to improve upon, or be determined from prior knowledge of the task.

Given the exploration distributions $\nu_{1:T}$, NRPI proceeds as follows. At each iteration n , it collects cost-to-go examples by sampling uniformly a time $t \in \{1, 2, \dots, T\}$, sampling a state s_t from ν_t , and then executes an exploration action a in s_t followed by execution of the current learner's policy π_n for time $t + 1$ to T , to obtain a cost-to-go estimate (s, a, t, Q) of executing a followed by π_n in state s at time t .⁷ Multiple cost-to-go estimates are collected this way and added in dataset D_n . After enough data has been collected, we update the learner's policy, to obtain π_{n+1} , using any no-regret online learning procedure, on the loss defined by the cost-sensitive classification examples in the new data D_n . This is iterated for a large number of iterations N . Initially, we may start with π_1 to be any guess of a good policy from the class Π , or use the expert's cost-to-go at the first iteration, to avoid having to specify an initial policy. This algorithm is detailed in Algorithm 2.

⁷In the particular case where $\nu_t = d_{\pi}^t$ of an exploration policy π , then to sample s_t , we would simply execute π from time 1 to $t - 1$, starting from the initial state distribution.

Algorithm 2 NRPI Algorithm

Initialize $\mathcal{D} \leftarrow \emptyset$, $\hat{\pi}_1$ to any policy in Π .
for $i = 1$ **to** N **do**
 Collect m data points as follows:
 for $j = 1$ **to** m **do**
 Sample uniformly $t \in \{1, 2, \dots, T\}$.
 Sample state s_t from exploration distribution ν_t .
 Execute some exploration action a_t in current state s_t at time t
 Execute $\hat{\pi}_i$ from time $t + 1$ to T , and observe estimate of cost-to-go \hat{Q} starting at time t
 end for
 Get dataset $\mathcal{D}_i = \{(s, a, t, \hat{Q})\}$ of states, actions, time, with current policy's cost-to-go.
 Aggregate datasets: $\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_i$.
 Train cost-sensitive classifier $\hat{\pi}_{i+1}$ on \mathcal{D} (*Alternately: use any online learner on the data-sets \mathcal{D}_i in sequence to get $\hat{\pi}_{i+1}$*)
end for
Return best $\hat{\pi}_i$ on validation.

3.1 Analysis

Consider the loss function L_n given to the online learning algorithm within NRPI at iteration n . Assuming infinite data, it assigns the following loss to each policy $\pi \in \Pi$:

$$L_n(\pi) = \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^{\hat{\pi}_n}(s, \pi)].$$

This loss represents the expected cost-to-go of executing π immediately for one step followed by current policy $\hat{\pi}_n$, under the exploration distributions $\nu_{1:T}$.

This sequence of losses over the iterations of training corresponds to an online cost-sensitive classification problem, as in the previous AGGREGATE algorithm. Let ϵ_{regret} be the average regret of the online learner on this online cost-sensitive classification problem after N iterations of NRPI:

$$\epsilon_{\text{regret}} = \frac{1}{N} \sum_{i=1}^N L_i(\pi_i) - \min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N L_i(\pi).$$

For any policy $\pi \in \Pi$, denote the average L_1 or variational distance between ν_t and d_π^t over time steps t as $D(\nu, \pi) = \frac{1}{T} \sum_{t=1}^T \|\nu_t - d_\pi^t\|_1$. Note that if $\nu_t = d_\pi^t$ for all t , then $D(\nu, \pi) = 0$.

Denote by Q_{\max} a bound on cost-to-go (which is always $\leq TC_{\max}$). Denote $\hat{\pi}$ the best policy found by NRPI over iterations, and $\bar{\pi}$ the uniform mixture policy over $\pi_{1:N}$ defined as before. Then NRPI achieves the following guarantee:

Theorem 3.1. *For any policy $\pi' \in \Pi$:*

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi') + T\epsilon_{\text{regret}} + TQ_{\max}D(\nu, \pi')$$

If a no-regret online cost-sensitive classification algorithm is used: $\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi') + TQ_{\max}D(\nu, \pi')$

NRPI thus finds policies that are as good as any other policy $\pi' \in \Pi$ whose state distribution $d_{\pi'}^t$ is close to ν_t on average over time t . Importantly, if $\nu_{1:T}$ corresponds to the state distribution of an optimal policy in class Π , then this theorem guarantees that NRPI will find an optimal policy (within the class Π) in the limit.

This theorem provides a similar performance guarantee to the results for PSDP presented in [4]. NRPI has the advantage of learning a single policy for test execution instead one at each time allowing for improved generalization and more efficient learning. NRPI imposes stronger requirements: it uses a no-regret online cost-sensitive classification procedure instead of simply a cost-sensitive supervised learner. For finite policy classes Π , or using reductions of cost-sensitive classification as mentioned previously, we may still obtain convex online learning problems for which efficient no-regret strategies exist or use the simple aggregation of data-sets with any sufficiently stable batch learner. [23, 27]

The result presented here can be interpreted as a reduction of model-free reinforcement learning to no-regret online learning. It is a regret reduction, as performance is related directly to the online regret at the cost-sensitive classification task. However performance is strongly limited by the quality of the exploration distribution.⁸

4 Discussion and Future Work

Contribution. The work here provides theoretical support for two seemingly unrelated empirical observations. First, and perhaps most crucially, much anecdotal evidence suggests that approximate policy iteration– and especially online variants [29]– is more effective and stable than theory and counter-examples to convergence might suggest. This cries out for some explanation; we contend that it can be understood as such online algorithms often enjoy no-regret guarantees or share similar stability properties than can ensure relative performance guarantees.

Similarly, practical implementation of imitation learning-for-structured-prediction methods like SEARN rely on what was previously considered a heuristic of using the expert demonstrator as an estimate of the future cost-to-go. The resulting good performance can be understood as a consequence of this heuristic being a special case of AGGREGATE where the online Frank-Wolfe algorithm [13] is used to choose policies. Moreover, stochastic mixing is but one of several approaches to achieving good online performance and deterministic variants have proven more effective in practice. [25]

The resulting algorithms make suggestions for batch approaches as well: they suggest, for instance, that approximate policy iteration procedures (as well as imitation learning ones) are likely to be more stable and effective if they train not only on the cost-to-go of the most recent policy but also on previous policies. At first this may seem counter-intuitive, however, it prevents the oscillations and divergences that at times plague batch approximate dynamic programming algorithms by ensuring that each learned policy is good across many states.

From a broad point of view, this work forms a piece of a growing picture that online algorithms and no-regret analyses– in contrast with the traditional *i.i.d.* or batch analysis– are important for understanding learning with application to control and decision making [26, 22, 25, 24]. At first glance, online learning seems concerned with a very different adversarial setting. By understanding these methods as attempting to ensure both good performance and robust, stable learning across iterations [23, 27], they become a natural tool for understanding the dynamics of interleaving learning and execution when our primary concern is generalization performance.

Limitations. It is important to note that any method relying on cost-to-go estimates can be impractical as collecting each estimate for a single state-action pair may involve executing an entire trajectory. In many settings, minimizing imitation loss with DAGGER [24], is more practical as we can observe the action chosen by the expert in every visited state along a trajectory and thus collect T data points per trajectory instead of single one. This is less crucial in structured prediction settings where the cost-to-go of the expert may often be quickly computed which has lead to the success of the heuristic analyzed here. A potential combination of the two approaches, where first simple imitation loss minimization provides a reasonable policy, and then this is refined using AGGREGATE (e.g. through additional gradient descent steps) thus using fewer (expensive) iterations.

In the reinforcement learning setting, the bound provided is as strong as that provided by [4, 16] for an arbitrary policy class. However, as TQ_{\max} is generally $O(T^2)$, this only provides meaningful guarantees when d_{π}^t is very close to ν_t (on average over time t). Previous methods like [4, 15, 28] provide a much stronger, *multiplicative* error guarantee when we consider competing against the bayes optimal policy in a fully observed MDP. It is not obvious how the current algorithm and analysis can extend to that variant of the bound.

Future Work. Much work remains to be done: there are a wide variety of no-regret learners and their practical trade-offs are almost completely open. Future work must explore this set to identify which methods are most effective in practice.

⁸One would naturally consider adapting the exploration distributions $\nu_{1:T}$ over the iterations of training. It can be shown that if $\nu_{1:T}^i$ are the exploration distributions at iteration i , and we have a mechanism for making $\nu_{1:T}^i$ converge to the state distributions of an optimal policy in Π as $i \rightarrow \infty$, then we would always be guaranteed to find an optimal policy in Π . Unfortunately, no known method can guarantee this.

References

- [1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 2009.
- [2] P. Auer, N. Cesa-Bianchi, Y. Freund, and R.E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 2002.
- [3] O. Avner, S. Mannor, and O. Shamir. Decoupling exploration and exploitation in multi-armed bandits. In *ICML*, 2012.
- [4] J. A. Bagnell, A. Y. Ng, S. Kakade, and J. Schneider. Policy search by dynamic programming. In *Advances in Neural Information Processing Systems*, 2003.
- [5] A. Beygelzimer, V. Dani, T. Hayes, J. Langford, and B. Zadrozny. Error limiting reductions between classification tasks. In *ICML*, 2005.
- [6] A. Beygelzimer, J. Langford, L. Li, L. Reyzin, and R. E. Schapire. Contextual bandit algorithms with supervised learning guarantees. In *AISTATS*, 2011.
- [7] A. Beygelzimer, J. Langford, and P. Ravikumar. Error-correcting tournaments. In *ALT*, 2009.
- [8] A. Beygelzimer, J. Langford, and B. Zadrozny. Machine learning techniques reductions between prediction quality metrics. In *Performance Modeling and Engineering*, pages 3–28. Springer US, 2008.
- [9] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [10] H. Daumé III, J. Langford, and D. Marcu. Search-based structured prediction. *Machine Learning*, 2009.
- [11] M. Dudik, J. Langford, and L. Li. Doubly robust policy evaluation and learning. In *ICML*, 2011.
- [12] F. Duvallet, T. Kollar, and A. Stentz. Imitation learning for natural language direction following through unknown environments. In *ICRA*, 2013.
- [13] E. Hazan and S. Kale. Projection-free online learning. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 521–528, 2012.
- [14] D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of American Statistics Association*, 1952.
- [15] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the 19th International Conference on Machine Learning (ICML)*, 2002.
- [16] S. M. Kakade. *On the sample complexity of reinforcement learning*. PhD thesis, University College London, 2003.
- [17] J. Langford and A. Beygelzimer. Sensitive error correcting output codes. In *COLT*, 2005.
- [18] L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *WWW*, 2010.
- [19] P. Mineiro. Error and regret bounds for cost-sensitive multiclass classification reduction to regression, 2010.
- [20] D. Munoz, J. A. Bagnell, and M. Hebert. Stacked hierarchical labeling. In *ECCV*, 2010.
- [21] S. Ross and J. A. Bagnell. Efficient reductions for imitation learning. In *AISTATS*, 2010.
- [22] S. Ross and J. A. Bagnell. Agnostic system identification for model-based reinforcement learning. In *ICML*, 2012.
- [23] S. Ross and J. A. Bagnell. Stability conditions for online learnability. In *Under Review. 15th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2012.
- [24] S. Ross, G. J. Gordon, and J. A. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011.
- [25] S. Ross, D. Munoz, M. Hebert, and J. A. Bagnell. Learning message-passing inference machines for structured prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
- [26] S. Ross, J. Zhou, Y. Yue, D. Dey, and J. A. Bagnell. Learning policies for contextual submodular prediction. In *ICML*, 2013.
- [27] Ankan Saha, Prateek Jain, and Ambuj Tewari. The interplay between stability and regret in online learning. *arXiv preprint arXiv:1211.6158*, 2012.
- [28] Bruno Scherrer. Approximate policy iteration schemes: A comparison. In *Proceedings of the International Conference on Machine Learning*, 2014.
- [29] R. Sutton. The right way to do reinforcement learning with function approximation. In *Approximate the Policy or the Value Function: Reinforcement Learning Workshop at Neural Information Processing Systems (NIPS)*, 2000.

Appendix: Proofs and Detailed Bounds

In this appendix, we provide the proofs and detailed analysis of the algorithms for imitation learning and reinforcement learning provided in the main document.

Lemmas

We begin with a classical and useful general lemma that is needed for bounding the expected loss under different distributions. This will be used several times throughout. Here this will be useful for bounding the expected loss under the state distribution of $\hat{\pi}$ (which optional queries the expert a fraction of the time during its execution) in terms of the expected loss under the state distribution of π_i :

Lemma 4.1. *Let P and Q be any distribution over elements $x \in \mathcal{X}$, and $f : \mathcal{X} \rightarrow \mathbb{R}$, any bounded function such that $f(x) \in [a, b]$ for all $x \in \mathcal{X}$. Let the range $r = b - a$. Then $|\mathbb{E}_{x \sim P}[f(x)] - \mathbb{E}_{x \sim Q}[f(x)]| \leq \frac{r}{2} \|P - Q\|_1$*

Proof. We provide the proof for \mathcal{X} discrete, a similar argument can be carried for \mathcal{X} continuous, using integrals instead of sums.

$$\begin{aligned} & |\mathbb{E}_{x \sim P}[f(x)] - \mathbb{E}_{x \sim Q}[f(x)]| \\ &= \left| \sum_x P(x)f(x) - Q(x)f(x) \right| \\ &= \left| \sum_x f(x)(P(x) - Q(x)) \right| \end{aligned}$$

Additionally, since for any real $c \in \mathbb{R}$, $\sum_x P(x)c = \sum_x Q(x)c$, then we have for any c :

$$\begin{aligned} & \left| \sum_x f(x)(P(x) - Q(x)) \right| \\ &= \left| \sum_x (f(x) - c)(P(x) - Q(x)) \right| \\ &\leq \sum_x |f(x) - c| |P(x) - Q(x)| \\ &\leq \max_x |f(x) - c| \sum_x |P(x) - Q(x)| \\ &= \max_x |f(x) - c| \|P - Q\|_1 \end{aligned}$$

This holds for all $c \in \mathbb{R}$. This upper bound is minimized for $c = a + \frac{r}{2}$, making $\max_x |f(x) - c| \leq \frac{r}{2}$. This proves the lemma. \square

The L_1 distance between the distribution of states encountered by $\hat{\pi}_i$, the policy chosen by the online learner, and π_i , the policy used to collect data that continues to execute the expert's actions with probability β_i is bounded as follows:

Lemma 4.2. $\|d_{\pi_i} - d_{\hat{\pi}_i}\|_1 \leq 2 \min(1, T\beta_i)$.

Proof. Let d the distribution of states over T steps conditioned on π_i picking the expert π^* at least once over T steps. Since π_i always executes $\hat{\pi}_i$ (never executes the expert action) over T steps with probability $(1 - \beta_i)^T$ we have $d_{\pi_i} = (1 - \beta_i)^T d_{\hat{\pi}_i} + (1 - (1 - \beta_i)^T) d$. Thus

$$\begin{aligned} & \|d_{\pi_i} - d_{\hat{\pi}_i}\|_1 \\ &= (1 - (1 - \beta_i)^T) \|d - d_{\hat{\pi}_i}\|_1 \\ &\leq 2(1 - (1 - \beta_i)^T) \\ &\leq 2T\beta_i \end{aligned}$$

The last inequality follows from the fact that $(1 - \beta)^T \geq 1 - \beta T$ for any $\beta \in [0, 1]$. Finally, since for any 2 distributions p, q , we always have $\|p - q\|_1 \leq 2$, then $\|d_{\pi_i} - d_{\hat{\pi}_i}\|_1 \leq 2 \min(1, T\beta_i)$. \square

Below we use the *performance difference lemma* [4, 15, 16] that is useful to bound the change in total cost-to-go. This general result bounds the difference in performance of any two policies. We present this results and its proof here for completeness.

Lemma 4.3. *Let π and π' be any two policy and denote V'_t and Q'_t the t -step value function and Q -value function of policy π' respectively, then:*

$$J(\pi) - J(\pi') = T \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi}^t} [Q'_{T-t+1}(s, \pi) - V'_{T-t+1}(s)]$$

for $U(1 : T)$ the uniform distribution on the set $\{1, 2, \dots, T\}$.

Proof. Let π_t denote the non-stationary policy that executes π in the first t time steps, and then switches to execute π' at time $t + 1$ to T . Then we have $J(\pi) = J(\pi_T)$ and $J(\pi') = J(\pi_0)$. Thus:

$$\begin{aligned} & J(\pi) - J(\pi') \\ &= \sum_{t=1}^T [J(\pi_t) - J(\pi_{t-1})] \\ &= \sum_{t=1}^T [\mathbb{E}_{s \sim d_{\pi_t}^t} [Q_{T-t+1}'(s, \pi) - V_{T-t+1}'(s)]] \\ &= T \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_t}^t} [Q_{T-t+1}'(s, \pi) - V_{T-t+1}'(s)] \end{aligned}$$

□

AGGREGATE Reduction Analysis

Let $\epsilon_{\text{class}} = \min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \pi) - \min_a Q_{T-t+1}^*(s, a)]$ denote the minimum expected cost-sensitive classification regret achieved by policies in the class Π on all the data over the N iterations of training. Denote the online learning average regret on the cost-to-go examples of the sequence of policies chosen by AGGREGATE, $\epsilon_{\text{regret}} = \frac{1}{N} [\sum_{i=1}^N \ell_i(\hat{\pi}_i) - \min_{\pi \in \Pi} \sum_{i=1}^N \ell_i(\pi)]$, where $\ell_i(\pi) = \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \pi)]$. Assume the cost-to-go of the expert Q^* is non-negative and bounded by Q_{\max}^* , and that β_i are chosen such that $\beta_i \leq (1 - \alpha)^{i-1}$ for some α . Then we have the following:

Theorem 4.4. *After N iterations of AGGREGATE:*

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi^*) + T[\epsilon_{\text{class}} + \epsilon_{\text{regret}}] + O\left(\frac{Q_{\max}^* T \log T}{\alpha N}\right).$$

Thus if a no-regret online algorithm is used to pick the sequence of policies $\hat{\pi}_{1:N}$, then as the number of iterations $N \rightarrow \infty$:

$$\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi^*) + T\epsilon_{\text{class}}$$

Proof. For every policy $\hat{\pi}_i$, we have:

$$\begin{aligned} & J(\hat{\pi}_i) - J(\pi^*) \\ &= T \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] \\ &= \sum_{t=1}^T \mathbb{E}_{s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] \\ &\leq \sum_{t=1}^T \mathbb{E}_{s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] + Q_{\max}^* \sum_{t=1}^T \|d_{\hat{\pi}_i}^t - d_{\pi_i}^t\|_1 \\ &\leq \sum_{t=1}^T \mathbb{E}_{s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] + 2Q_{\max}^* \sum_{t=1}^T \min(1, t\beta_i) \\ &\leq \sum_{t=1}^T \mathbb{E}_{s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] + 2TQ_{\max}^* \min(1, T\beta_i) \\ &= T \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] + 2TQ_{\max}^* \min(1, T\beta_i) \end{aligned}$$

where we use lemma 4.3 in the first equality, lemma 4.1 in the first inequality, and a similar argument to lemma 4.2 for the second inequality.

Since β_i are non-increasing, define n_β the largest $n \leq N$ such that $\beta_n > 1/T$. Then:

$$\begin{aligned} & J(\bar{\pi}) - J(\pi^*) \\ &= \frac{1}{N} \sum_{i=1}^N [J(\hat{\pi}_i) - J(\pi^*)] \\ &\leq \frac{1}{N} \sum_{i=1}^N [T \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - V_{T-t+1}^*(s)] + 2TQ_{\max}^* \min(1, T\beta_i)] \\ &= T[\min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \pi) - V_{T-t+1}^*(s)]] + T\epsilon_{\text{regret}} \\ &\quad + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \\ &\leq T[\min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \pi) - \min_a Q_{T-t+1}^*(s, a)]] + T\epsilon_{\text{regret}} \\ &\quad + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \\ &= T\epsilon_{\text{class}} + T\epsilon_{\text{regret}} + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \end{aligned}$$

Again, $J(\hat{\pi}) \leq J(\bar{\pi})$ since the minimum is always better than the average, i.e. $\min_i J(\hat{\pi}_i) \leq \frac{1}{N} \sum_{i=1}^N J(\hat{\pi}_i)$. Finally, we have that when $\beta_i = (1 - \alpha)^{i-1}$, $[n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \leq \frac{\log(T)+2}{\alpha}$. This proves the first part of the theorem.

The second part follows immediately from the fact that $\epsilon_{\text{regret}} \rightarrow 0$ as $N \rightarrow \infty$, and similarly for the extra term $O\left(\frac{Q_{\max}^* T \log T}{\alpha N}\right)$. □

Finite Sample AGGREGATE with Q-function approximation

We here consider the finite sample case where actions are explored uniformly randomly and the reduction of cost-sensitive classification to squared loss regression is used. We consider learning an estimate Q-value function \hat{Q} of the expert's cost-to-go, and we consider a general case where the cost-to-go predictions may depend on features $f(s, a, t)$ of the state s , action a and time t , e.g. \hat{Q} could be a linear regressor s.t. $\hat{Q}_{T-t+1}(s, a) = w^\top f(s, a, t)$ is the estimate of the cost-to-go $Q_{T-t+1}^*(s, a)$, and w are the parameters of the linear regressor we learn. Given such estimates \hat{Q} , we consider executing the policy $\hat{\pi}$, such that in state s at time t , $\hat{\pi}(s, t) = \min_{a \in A} \hat{Q}_{T-t+1}(s, a)$.

Theorem 4.5. *After N iterations of AGGREGATE, collecting m regression examples (s, a, t, Q) per iteration, guarantees that with probability at least $1 - \delta$:*

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi^*) + 2\sqrt{|A|T} \sqrt{\hat{\epsilon}_{\text{class}} + \hat{\epsilon}_{\text{regret}} + O(\sqrt{\log(1/\delta)/Nm})} + O\left(\frac{Q_{\max} T \log T}{\alpha N}\right).$$

Thus if a no-regret online algorithm is used to pick the sequence of regressors $\hat{Q}_{1:N}$, then as the number of iterations $N \rightarrow \infty$, with probability 1:

$$\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi^*) + 2\sqrt{|A|T} \sqrt{\hat{\epsilon}_{\text{class}}}$$

Proof. Consider $\bar{\pi}$, the bayes-optimal non-stationary policy that minimizes loss on the cost-to-go examples. That is, $\bar{\pi}(s, t) = \min_{a \in A} Q_{T-t+1}^*(s, a)$, i.e. it picks the action with minimum expected expert cost-to-go conditioned on being in state s and time t . Additionally, given the observed noisy Q-values from each trajectory, the bayes-optimal regressor is simply the Q-value function Q^* of the expert that predicts the expected cost-to-go.

At each iteration i , we execute a policy $\hat{\pi}_i$, such that $\hat{\pi}_i(s, t) = \arg \min_{a \in A} \hat{Q}_{T-t+1}^i(s, a)$, where \hat{Q}^i is the current regressor at iteration i from the base online learner. The cost-sensitive regret of policy $\hat{\pi}_i$, compared to $\bar{\pi}$, can be related to the regression regret of \hat{Q}_i as follows:

Consider any state s and time t . Let $\hat{a}_i = \hat{\pi}_i(s, t)$ and consider the action a' of any other policy. We have that:

$$\begin{aligned} & Q_{T-t+1}^*(s, \hat{a}_i) - Q_{T-t+1}^*(s, a) \\ & \leq \hat{Q}_{T-t+1}^i(s, \hat{a}_i) - \hat{Q}_{T-t+1}^i(s, a') + Q_{T-t+1}^*(s, \hat{a}_i) - \hat{Q}_{T-t+1}^i(s, \hat{a}_i) + \hat{Q}_{T-t+1}^i(s, a') - Q_{T-t+1}^*(s, a') \\ & \leq Q_{T-t+1}^*(s, \hat{a}_i) - \hat{Q}_{T-t+1}^i(s, \hat{a}_i) + \hat{Q}_{T-t+1}^i(s, a') - Q_{T-t+1}^*(s, a') \\ & \leq 2 \max_{a \in A} |Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)| \end{aligned}$$

Additionally, for any joint distribution D over (s, t) , and $U(A)$ the uniform distribution over actions, we have that:

$$\begin{aligned} & \mathbb{E}_{(s,t) \sim D} [\max_{a \in A} |Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|]^2 \\ & \leq \mathbb{E}_{(s,t) \sim D} [\max_{a \in A} |Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2] \\ & \leq \mathbb{E}_{(s,t) \sim D} [\sum_{a \in A} |Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2] \\ & = |A| \mathbb{E}_{(s,t) \sim D, a \sim U(A)} [|Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2] \end{aligned}$$

Thus we obtain that for every $\hat{\pi}_i$:

$$\begin{aligned} & \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - Q_{T-t+1}^*(s, \bar{\pi})] \\ & \leq 2 \mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t} [\max_{a \in A} |Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|] \\ & \leq 2\sqrt{|A|} \sqrt{\mathbb{E}_{t \sim U(1:T), s \sim d_{\hat{\pi}_i}^t, a \sim U(A)} [|Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2]} \end{aligned}$$

Thus

$$\begin{aligned}
& J(\bar{\pi}) - J(\pi^*) \\
&= \frac{T}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - Q_{T-t+1}^*(s, \pi^*)] \\
&\leq \frac{T}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - Q_{T-t+1}^*(s, \pi^*)] + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \\
&\leq \frac{T}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t} [Q_{T-t+1}^*(s, \hat{\pi}_i) - Q_{T-t+1}^*(s, \tilde{\pi})] + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \\
&\leq \frac{2\sqrt{|A|}T}{N} \sum_{i=1}^N \sqrt{\mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t, a \sim U(A)} [|Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2]} \\
&\quad + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i] \\
&\leq 2\sqrt{|A|}T \sqrt{\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t, a \sim U(A)} [|Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2]} \\
&\quad + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i]
\end{aligned}$$

Now in state s at time t , when performing a and then following the expert, consider the distribution $d_{s,a,t}$ over observed cost-to-go Q , such that $\mathbb{E}_{Q \sim d_{s,a,t}} [Q] = Q_{T-t+1}^*(s, a)$.

For any regressor \hat{Q} , define the expected squared loss in predictions of the observed cost-to-go at iteration i as $\ell_i(\hat{Q}) = \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t, a \sim U(A), Q \sim d_{s,t,a}} [|Q - \hat{Q}_{T-t+1}(s, a)|^2]$. Then since for any random variable X with mean μ , if we have an estimate $\hat{\mu}$ of the mean, $|\hat{\mu} - \mu|^2 = \mathbb{E}_x [(x - \hat{\mu})^2 - (x - \mu)^2]$, we have that:

$$\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi_i}^t, a \sim U(A)} [|Q_{T-t+1}^*(s, a) - \hat{Q}_{T-t+1}^i(s, a)|^2] = \frac{1}{N} \sum_{i=1}^N \ell_i(\hat{Q}^i) - \ell_i(Q^*)$$

Now, in the finite sample case, consider collecting m samples at each iteration i : $\{(s_{ij}, a_{ij}, t_{ij}, Q_{ij})\}_{j=1}^m$. The expected squared loss ℓ_i is estimated as $\hat{\ell}_i(\hat{Q}) = \frac{1}{m} \sum_{j=1}^m (\hat{Q}_{T-t_{ij}+1}(s_{ij}, a_{ij}) - Q_{ij})^2$, and the no-regret algorithm is run on the estimated loss $\hat{\ell}_i$.

Define $Y_{i,j} = \ell_i(\hat{Q}^i) - (\hat{Q}_{T-t_{ij}+1}^i(s_{ij}, a_{ij}) - Q_{ij})^2 - \ell_i(Q^*) + (Q_{T-t_{ij}+1}^*(s_{ij}, a_{ij}) - Q_{ij})^2$, the difference between the expected squared loss and the empirical square loss at each sample for both \hat{Q}^i and Q^* . Conditioned on previous trajectories, each $Y_{i,j}$ has expectation 0. Then the sequence of random variables $X_{km+l} = \sum_{i=1}^k \sum_{j=1}^m Y_{i,j} + \sum_{j=1}^l Y_{(k+1),j}$, for $k \in \{0, 1, 2, \dots, N-1\}$ and $l \in \{1, 2, \dots, m\}$, forms a martingale, and if the squared loss at any sample is bounded by ℓ_{\max} , we obtain that $|X_i - X_{i+1}| \leq 2\ell_{\max}$. By Azuma-Hoeffding's inequality, this implies that with probability at least $1 - \delta$, $\frac{1}{Nm} X_{Nm} \leq 2\ell_{\max} \sqrt{\frac{2 \log(1/\delta)}{Nm}}$.

Denote the empirical average online regret on the training squared loss $\hat{\epsilon}_{\text{regret}} = \frac{1}{N} \sum_{i=1}^N \hat{\ell}_i(\hat{Q}^i) - \min_{\hat{Q} \in \mathcal{Q}} \frac{1}{N} \sum_{i=1}^N \hat{\ell}_i(\hat{Q})$. Let \tilde{Q}^* be the bayes-optimal regressor on the finite training data, and define the empirical regression regret of the best regressor in the class as $\hat{\epsilon}_{\text{class}} = \min_{\hat{Q} \in \mathcal{Q}} \frac{1}{N} \sum_{i=1}^N [\hat{\ell}_i(\hat{Q}) - \hat{\ell}_i(\tilde{Q}^*)]$.

Then we obtain that with probability at least $1 - \delta$:

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \ell_i(\hat{Q}^i) - \ell_i(Q^*) \\
&= \frac{1}{N} \sum_{i=1}^N \hat{\ell}_i(\hat{Q}^i) - \hat{\ell}_i(Q^*) + \frac{1}{Nm} X_{Nm} \\
&\leq \frac{1}{N} \sum_{i=1}^N \hat{\ell}_i(\hat{Q}^i) - \hat{\ell}_i(Q^*) + 2\ell_{\max} \sqrt{\frac{2 \log(1/\delta)}{Nm}} \\
&\leq \min_{\hat{Q} \in \mathcal{Q}} \frac{1}{N} \sum_{i=1}^N [\hat{\ell}_i(\hat{Q}) - \hat{\ell}_i(Q^*)] + \hat{\epsilon}_{\text{regret}} + 2\ell_{\max} \sqrt{\frac{2 \log(1/\delta)}{Nm}} \\
&\leq \hat{\epsilon}_{\text{class}} + \hat{\epsilon}_{\text{regret}} + 2\ell_{\max} \sqrt{\frac{2 \log(1/\delta)}{Nm}}
\end{aligned}$$

where the last inequality follows from the fact that $\sum_{i=1}^N \hat{\ell}_i(\tilde{Q}^*) \leq \sum_{i=1}^N \hat{\ell}_i(Q^*)$.

Combining with the above, we obtain that with probability at least $1 - \delta$:

$$J(\bar{\pi}) - J(\pi^*) \leq 2\sqrt{|A|}T \sqrt{\hat{\epsilon}_{\text{class}} + \hat{\epsilon}_{\text{regret}} + 2\ell_{\max} \sqrt{\frac{2 \log(1/\delta)}{Nm}}} + \frac{2TQ_{\max}^*}{N} [n_\beta + T \sum_{i=n_\beta+1}^N \beta_i]$$

□

NRPI Reduction Analysis

We here provide the proof of the result for NRPI, sampled from state exploration distributions $\nu_{1:T}$.

To analyze this version, we begin with an alternate version of the performance difference lemma (lemma 4.3) presented before:

Lemma 4.6. *Let π and π' be any two policy and denote V_t and Q_t the t -step value function and Q -value function of policy π respectively, then:*

$$J(\pi) - J(\pi') = T \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi'}^t} [V_{T-t+1}(s) - Q_{T-t+1}(s, \pi')]$$

for $U(1:T)$ the uniform distribution on the set $\{1, 2, \dots, T\}$.

Proof. By applying lemma 4.3 to $J(\pi') - J(\pi)$, we obtain:

$$J(\pi') - J(\pi) = T \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi}^t} [Q_{T-t+1}(s, \pi') - V_{T-t+1}(s)]$$

This proves the lemma. □

Now denote the loss L_n used by the online learner at iteration n , s.t.:

$$L_n(\pi) = \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^{\hat{\pi}_n}(s, \pi)].$$

and ϵ_{regret} the average regret after the N iterations of NRPI:

$$\epsilon_{\text{regret}} = \frac{1}{N} \sum_{i=1}^N L_i(\pi_i) - \min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N L_i(\pi).$$

For any policy $\pi \in \Pi$, denote the average L_1 distance between ν_t and d_{π}^t over time steps t as:

$$D(\nu, \pi) = \frac{1}{T} \sum_{t=1}^T \|\nu_t - d_{\pi}^t\|_1.$$

Assume the cost-to-go of the learned policies $\pi_1, \pi_2, \dots, \pi_N$ are non-negative and bounded by Q_{\max} , for any state s , action a and time t (in the worst case this is TC_{\max}). Denote $\hat{\pi}$ the best policy found by NRPI over the iterations, and $\bar{\pi}$ the uniform mixture policy over $\pi_{1:N}$ defined as before. Then we have to following guarantee with this version of NRPI with learner's cost-to-go:

Theorem 4.7. *For any $\pi' \in \Pi$:*

$$J(\hat{\pi}) \leq J(\bar{\pi}) \leq J(\pi') + T\epsilon_{\text{regret}} + TQ_{\max}D(\nu, \pi')$$

Thus, if a no-regret online cost-sensitive classification algorithm is used, then:

$$\lim_{N \rightarrow \infty} J(\bar{\pi}) \leq J(\pi') + TQ_{\max}D(\nu, \pi')$$

Proof. Let Q_t^i denote the t -step Q -value function of policy $\hat{\pi}_i$. Then for every $\hat{\pi}_i$ we have:

$$\begin{aligned} & J(\hat{\pi}_i) - J(\pi') \\ &= T \mathbb{E}_{t \sim U(1:T), s \sim d_{\pi'}^t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] \\ &= \sum_{t=1}^T \mathbb{E}_{s \sim d_{\pi'}^t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] \\ &\leq \sum_{t=1}^T \mathbb{E}_{s \sim \nu_t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] + Q_{\max} \sum_{t=1}^T \|\nu_t - d_{\pi'}^t\|_1 \\ &= T \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] + TQ_{\max}D(\nu, \pi') \end{aligned}$$

where we use lemma 4.6 in the first equality, and lemma 4.1 in the first inequality.

Thus:

$$\begin{aligned} & J(\bar{\pi}) - J(\pi') \\ &= \frac{1}{N} \sum_{i=1}^N [J(\hat{\pi}_i) - J(\pi')] \\ &\leq \frac{1}{N} \sum_{i=1}^N [T \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] + TQ_{\max}D(\nu, \pi')] \\ &\leq T \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^i(s, \hat{\pi}_i) - Q_{T-t+1}^i(s, \pi')] - T \min_{\pi \in \Pi} \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{t \sim U(1:T), s \sim \nu_t} [Q_{T-t+1}^i(s, \pi)] \\ &\quad + TQ_{\max}D(\nu, \pi') \\ &= T\epsilon_{\text{regret}} + TQ_{\max}D(\nu, \pi') \end{aligned}$$

Again, $J(\hat{\pi}) \leq J(\bar{\pi})$ since the minimum is always better than the average, i.e. $\min_i J(\hat{\pi}_i) \leq \frac{1}{N} \sum_{i=1}^N J(\hat{\pi}_i)$. This proves the first part of the theorem.

The second part follows immediately from the fact that $\epsilon_{\text{regret}} \rightarrow 0$ as $N \rightarrow \infty$. □