
Continuous Online Learning and New Insights to Online Imitation Learning

Jonathan Lee*

University of California, Berkeley
jonathan_lee@berkeley.edu

Ching-An Cheng*

Georgia Institute of Technology
cacheng@gatech.edu

Ken Goldberg

University of California, Berkeley
goldberg@berkeley.edu

Byron Boots

Georgia Institute of Technology
bboots@cc.gatech.edu

Abstract

Online learning is a powerful tool for analyzing iterative algorithms. However, the classic adversarial setup sometimes fails to capture certain regularity in online problems in practice. Motivated by this, we establish a new setup, called Continuous Online Learning (COL), where the gradient of online loss function changes continuously across rounds with respect to the learner’s decisions. We show that COL covers and more appropriately describes many interesting applications, from general equilibrium problems (EPs) to optimization in episodic MDPs. Using this new setup, we revisit the difficulty of achieving sublinear dynamic regret. We prove that there is a fundamental equivalence between achieving sublinear dynamic regret in COL and solving certain EPs, and we present a reduction from dynamic regret to both static regret and convergence rate of the associated EP. At the end, we specialize these new insights into online imitation learning and show improved understanding of its learning stability.

1 Introduction

Online learning [1, 2] studies the interactions between a learner (i.e. an algorithm) and an opponent through regret minimization. It has proven to be a powerful framework for analyzing and designing iterative algorithms. However, while classic online learning setups focus on bounding the worst case, many applications are not naturally adversarial. This reality gap exists especially for iterative algorithms that are designed to solve optimization problems concerning Markov decision processes (MDPs). Because the objective is often stated in expectation for these problems, continuity properties often arise naturally from the smoothing effect of taking expectations over randomness. When such properties are ignored, theoretical analyses can be overly conservative.

To this end, we propose a new setup for online learning, called Continuous Online Learning (COL). In contrast to the standard adversarial setup that treats losses as adversarial, COL concerns online learning problems where the per-round losses change continuously with respect to the learner’s decisions, and it models adversity, such as stochasticity and bias, as corruption in the feedback signals of these continuous loss sequences. This modified setup natively captures regularity in online losses, while still being able to handle adversity that appears in common problems. As a result, certain concepts that are difficult to analyze in the classic adversarial setup (e.g. sublinear dynamic regret with an adaptive opponent) become attainable in COL.

*Equal contribution

The goal of this paper is to establish COL and to study, particularly, conditions and efficient algorithms for achieving sublinear dynamic regret. Our first result shows that achieving sublinear dynamic regret in COL, interestingly, is equivalent to solving certain equilibrium problems (EPs), which are known to be PPAD-complete² [3]. In other words, achieving sublinear dynamic regret that is polynomial in the dimension of the decision set can be extremely difficult in general. Nevertheless, based on the solution concept of EP, we present a *reduction* from sublinear dynamic regret to static regret and convergence to the solution of the associated EP. This reduction allows us to quickly derive non-asymptotic dynamic regret bounds of popular online learning algorithms based on their known static regret rates.

Using these insights of COL, we revisit online imitation learning (IL) [4] and show it can be framed as a COL problem. We demonstrate that, by using standard analyses of COL, we are able to recover and improve existing understanding of online IL algorithms [4, 5, 6]. In particular, we characterize existence and uniqueness of solutions, and present convergence and dynamic regret bounds for a common class of IL algorithms in deterministic and stochastic settings. A more detailed version of this paper with additional theoretical results and proofs omitted here can be found in the full technical report [7].

2 Continuous Online Learning

We recall, generally, an online learning problem repeats the following steps: in round n , the learner plays a decision x_n from a convex and compact decision set \mathcal{X} , the opponent chooses a loss function $l_n : \mathcal{X} \rightarrow \mathbb{R}$ based on the decisions of the learner, and then information about l_n (e.g. $\nabla l_n(x_n)$) is revealed to the learner to inform the next decision. Classically, this abstract setup studies the *adversarial* setting where l_n can be almost arbitrarily chosen except for minor restrictions like convexity [8, 9]. Often the performance is measured relatively through *static regret*,

$$\text{Regret}_N^s := \sum_{n=1}^N l_n(x_n) - \min_{x \in \mathcal{X}} \sum_{n=1}^N l_n(x). \quad (1)$$

Recently, interest has emerged in algorithms that can make nearly optimal decisions at each round. The regret is therefore measured on-the-fly and suitably named *dynamic regret*,

$$\text{Regret}_N^d := \sum_{n=1}^N l_n(x_n) - \sum_{n=1}^N l_n(x_n^*), \quad (2)$$

where $x_n^* \in \arg \min_{x \in \mathcal{X}} l_n(x)$. As dynamic regret by definition upper bounds static regret, minimizing the dynamic regret is a more difficult problem.

At a high level, one can view online learning as a protocol to describe iterative algorithms, i.e., an algorithm receives some feedback, updates its decision, tries it out and receives a performance measure, and then repeats. Indeed, this idea has made online learning a ubiquitous tool to analyze a wide range of problems. But often in these problems, the loss sequence has certain correlations; if the algorithm outputs the same decision, regardless of which iteration it is in, its performance will be measured similarly. This structure of regularity, however, is missing the classic adversarial setup. While it is possible to introduce ad-hoc constraints to limit the amount of adversity in the classic setup, as in [2, 10, 11, 12, 13, 14, 15], such a scheme often leads to case-by-case analyses and can hardly model problems where the adversity depends also on the learner's decision, like online IL of interest here (see Section 5). This mismatch between practice and theory makes studying certain convergence concepts difficult, such as sublinear dynamic regret which is useful to understand the performance of the last iterate produced by the algorithm.

COL differs from the classic setup mainly in the way the loss and the feedback are defined, so that it can inherently model regularity that shows up in the loss sequence of problems in practice. In COL, we suppose that the opponent possesses a bifunction $f : (x, x') \mapsto f_x(x') \in \mathbb{R}$, for $x, x' \in \mathcal{X}$, that is *unknown* to the learner. This bifunction is used by the opponent to determine the per-round losses: in round n , if the learner chooses x_n , then the opponent responds with

$$l_n(x) := f_{x_n}(x). \quad (3)$$

Finally, the learner suffers $l_n(x_n)$ and receives feedback about l_n . For $f_x(x')$, we treat x as the *query argument* that proposes a question (i.e. an optimization objective $f_x(\cdot)$), and treat x' as the *decision*

²In short, they are NP problems whose solutions are known to exist, but it is open as to if they belong to P.

argument whose performance is evaluated. This bifunction f generally can be defined online as queried, with only one limitation that the same loss function $f_x(\cdot)$ must be selected by the opponent whenever the learner plays the same decision x . Thus, the opponent can be adaptive, but in response to only the learner's current decision. We assume, for all $x \in \mathcal{X}$, $\|\nabla f_x(x)\|_* < G$ for some $G < \infty$. In Section 5, we will discuss how the bifunction provides a natural interpretation for certain difficult objectives such as in online IL.

In addition to the restriction in (3), we impose regularity into f to relate l_n across rounds (so that seeking sublinear dynamic regret becomes well defined.³)

Definition 1. We say an online learning problem is *continuous* if l_n is set as in (3) by a bifunction f satisfying, $\forall x' \in \mathcal{X}$, $\nabla f_x(x')$ is a continuous map in x ⁴.

The continuity may appear to restrict COL to purely deterministic settings, but adversity such as stochasticity can be incorporated via an important *nuance* in the relationship between loss and feedback. In the classical online learning setting, the adversity is incorporated in the loss: the losses l_n and decisions x_n may themselves be generated adversarially or stochastically and then they directly determine the feedback, e.g., given as full information (receiving l_n or $\nabla l_n(x_n)$) or bandit (just $l_n(x_n)$). The (expected) regret is then measured with respect to these intrinsically adversarial losses l_n . By contrast, in COL, we always measure regret with respect to the true underlying bifunction $l_n = f_{x_n}$. Instead we give the opponent the freedom to add an additional stochastic or adversarial component into the feedback; e.g., in first-order feedback, the learner could receive $g_n = \nabla l_n(x_n) + \xi_n$, where ξ_n is a probabilistically bounded and potentially adversarial vector, which can be used to model noise or bias in feedback. In other words, the COL setting models a true underlying loss with regularity, but allows adversity to be modeled within the feedback, analogous to stochastic feedback oracles in convex optimization. This additional structure is especially important for studying dynamic regret, as it allows us to always consider regret with respect to the true f_{x_n} while still incorporating the possibility of stochasticity and adversity.

3 Equivalence and Hardness of Continuous Online Learning

We first ask what extra information the COL formulation entails. We present this result as an equivalence between achieving sublinear dynamic in COL and solving several mathematical programming problems. Particularly, suppose $\mathcal{X} \subset \mathbb{R}^d$; we are interested in whether sublinear dynamic regret with polynomial dependency on d is even possible. It turns out, *in general*, this is difficult, at least as hard as a set of difficult problems known to be PPAD-complete [3], even when $f_x(\cdot)$ is convex and continuous.

Theorem 1. *Let f be given in Definition 1 for a convex and compact decision set $\mathcal{X} \subset \mathbb{R}^d$. Suppose $f_x(\cdot)$ is convex and continuous. For any f satisfying the above assumption, if there is an algorithm that achieves sublinear dynamic regret that is in $\text{poly}(d)$ in the associated COL, then it solves all PPAD problems in polynomial time. In particular, achieving sublinear dynamic regret is equivalent to solving the equilibrium problem $\text{EP}(\mathcal{X}, \Phi)$ with $\Phi(x, x') = f_x(x') - f_x(x)$ and the variational inequality $\text{VI}(\mathcal{X}, F)$ with $F(x) = \nabla f_x(x)$.*

Theorem 1 is an excerpt of [7, Theorem 1] in the technical report. We recall an EP problem, $\text{EP}(\mathcal{X}, \Phi)$, is defined by a variable set $\mathcal{X} \subseteq \mathbb{R}^d$ and a bifunction $\Phi : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that⁵ $\Phi(x, x) \geq 0$, $\Phi(\cdot, x)$ is continuous, and $\Phi(x, \cdot)$ is convex. Its goal is to find a point $x^* \in \mathcal{X}$ such that

$$\Phi(x^*, x) \geq 0, \quad \forall x \in \mathcal{X}.$$

Similarly, the goal of a VI problem, $\text{VI}(\mathcal{X}, F)$ with $F : \mathcal{X} \rightarrow \mathbb{R}^d$, is to find a point $x^* \in \mathcal{X}$ such that

$$\langle F(x^*), x - x^* \rangle \geq 0, \quad \forall x \in \mathcal{X}$$

By definition one can see that the VI problem $\text{VI}(\mathcal{X}, F)$ is also an EP problem $\text{EP}(\mathcal{X}, \Phi_F)$ with $\Phi_F(x, x') := \langle F(x), x' - x \rangle$.

In other words, Theorem 1 states that, based on the identification $\Phi(x, x') = f_x(x') - f_x(x)$ and $F(x) = \nabla f_x(x)$, achieving sublinear dynamic regret is essentially equivalent to finding an

³Otherwise the opponent can define $f_x(\cdot)$ pointwise for each x to make $l_n(x_n) - l_n(x_n^*)$ constant.

⁴We define $\nabla f_x(x')$ as the derivative with respect to x' .

⁵The convexity and continuity can be relaxed further, e.g., as hemi-continuity.

equilibrium $x^* \in X^*$, in which X^* denotes the set of solutions of the EP and VI (one can show these two solution sets coincide [7]). Therefore, a *necessary* condition for sublinear dynamic regret is that X^* is non-empty, which is true when $\nabla f_x(x)$ is continuous in x and \mathcal{X} is compact [16].

Theorem 1 also implies that extra structure on COL is necessary for designing *efficient* algorithms that achieve sublinear dynamic regret and find these solutions. Specifically, we are interested in algorithms whose dynamic regret is sublinear *and* polynomial in d . The requirement of polynomial dependency is important to properly define the problem. Without it, sublinear dynamic regret can be achieved already (at least asymptotically), e.g. by simply performing a grid search that discretizes \mathcal{X} (as \mathcal{X} is compact and ∇f is continuous) albeit with an exponentially large constant.

Based on this equivalence, we can strengthen the structural properties of COL so that they are conducive to designing such efficient algorithms.

Definition 2. We say a COL problem with f is (α, β) -regular if for some $\alpha, \beta \in [0, \infty)$, $\forall x \in \mathcal{X}$,

1. $f_x(\cdot)$ is a α -strongly convex function.
2. $\nabla f_x(x)$ is a β -Lipschitz continuous map.

Leveraging these, we can identify similar structural properties in the equivalent problems.

Proposition 1. *If the COL problem with f is (α, β) -regular, then the map $\nabla f_x(x)$ is $(\alpha - \beta)$ -strongly monotone. That is, for all $x, y \in \mathcal{X}$,*

$$\langle \nabla f_x(x) - \nabla f_y(y), x - y \rangle \geq (\alpha - \beta) \|x - y\|^2$$

It is well known that strong monotonicity implies that $\text{VI}(\mathcal{X}, \nabla f)$ has a unique solution. It also implies that fast linear convergence is possible for deterministic feedback in VI problems [16].

4 Reduction by Regularity

We present a reduction from minimizing dynamic regret to minimizing static regret and convergence to X^* . Intuitively, this is possible because Theorem 1 suggests achieving sublinear dynamic regret should not be harder than finding $x^* \in X^*$. Define $\text{Regret}_N^s(x^*) := \sum_{n=1}^N l_n(x_n) - l_n(x^*)$.

Theorem 2. *Define $D_{\mathcal{X}} := \max_{x, x' \in \mathcal{X}} \|x - x'\|$. Let $x^* \in X^*$ and $\Delta_n := \|x_n - x^*\|$. If f is (α, β) -regular for $\alpha, \beta \in [0, \infty)$, then for all N ,*

$$\text{Regret}_N^d \leq \min\{G \sum_{n=1}^N \Delta_n, \text{Regret}_N^s(x^*)\} + \sum_{n=1}^N \min\{\beta D_{\mathcal{X}} \Delta_n, \frac{\beta^2}{2\alpha} \Delta_n^2\}$$

Theorem 2 roughly shows that when an equilibrium x^* exists (e.g. given by the sufficient conditions in the previous section), it provides a stabilizing effect to the problem, so the dynamic regret behaves almost like the static regret when the decisions are around x^* .

This relationship can be used as a powerful tool for understanding the dynamic regret of existing algorithms designed for EPs and VIs. These include, e.g., mirror descent [17], mirror-prox [18, 19], conditional gradient descent [20], Mann iteration [21], etc. Interestingly, many of those are also standard tools in online learning with static regret bounds that are well known [9].

We can apply Theorem 2 in different ways, depending on the known convergence of an algorithm. For algorithms whose convergence rate of Δ_n to zero is known, Theorem 2 essentially shows that their dynamic regret is at most $O(\sum_{n=1}^N \Delta_n)$. For the algorithms with only known static regret bounds, we can use a corollary.

Corollary 1. *If f is (α, β) -regular and $\alpha > \beta$, it holds $\text{Regret}_N^d \leq \text{Regret}_N^s(x^*) + \frac{\beta^2 \widetilde{\text{Regret}}_N^s(x^*)}{2\alpha(\alpha - \beta)}$, where $\widetilde{\text{Regret}}_N^s(x^*)$ is the static regret of the linear online learning problem with $l_n(x) = \langle \nabla f_n(x_n), x \rangle$.*

The purpose of Corollary 1 is not to give a tight bound, but to show that for nicer problems with $\alpha > \beta$, achieving sublinear dynamic regret is not harder than achieving sublinear static regret under linear losses. For tighter bounds, we still refer to Theorem 2 to leverage the equilibrium convergence.

Finally, we remark Theorem 2 is directly applicable to expected dynamic regret (the right-hand side of the inequality will be replaced by its expectation) when the learner only has access to stochastic feedback, because the COL setup in non-anticipating. Similarly, high-probability bounds can be obtained based on martingale convergence theorems (see [22] for a COL example). In these cases, we note that the regret is defined with respect to l_n in COL, *not* the sampled losses.

5 Application to Online Imitation Learning

In this section, we investigate an application of the COL framework in the sequential decision problem of online IL [4]. We consider an episodic MDP with state space \mathcal{S} , action space \mathcal{A} , and finite horizon T . For any $s, s' \in \mathcal{S}$ and $a \in \mathcal{A}$, the transition dynamics \mathcal{P} gives the conditional density, denoted by $\mathcal{P}(s'|s, a)$, of transitioning to s' from state s and action a . The reward of state s and action a is denoted as $r(s, a)$. A policy π is a mapping from \mathcal{S} to a density over \mathcal{A} . We suppose the MDP starts from some fixed initial state distribution. We denote the probability of being in state s at time t under policy π as $d_t^\pi(s)$, and we define the average state distribution under π as $d^\pi(s) = \frac{1}{T} \sum_{t=1}^T d_t^\pi(s)$.

In IL, we assume that \mathcal{P} and r are unknown to the learner, but, during training time, the learner is given access to an expert policy π^* and full knowledge of a supervised learning loss function $c(s, \pi; \pi^*)$, defined for each state $s \in \mathcal{S}$. The objective of IL is to solve

$$\min_{\pi \in \Pi} \mathbb{E}_{s \sim d^\pi} [c(s, \pi; \pi^*)], \quad (4)$$

where Π is the set of allowable parametric policies, which will be assumed convex; note that it is often the case that $\pi^* \notin \Pi$.

As d^π is not known analytically, optimizing (4) directly leads to a reinforcement learning problem and therefore can be sample inefficient. *Online IL*, such as the popular DAGGER algorithm [4], bypasses this difficulty by reducing (4) into a sequence of supervised learning problems. Below we describe a general construction of online IL: at the n th iteration (1) execute the learner's current policy π_n in the MDP to collect state, action samples; (2) update π_{n+1} with information of the stochastic approximation of $l_n(\pi) = \mathbb{E}_{d^{\pi_n}} [c(s, \pi; \pi^*)]$ based the samples collected in the first step. Importantly, we remark that in these empirical risks, the states are sampled according to d^{π_n} of the learner's policy.

The use of online learning to analyze online IL is well established [4]. As studied in [5, 6], these online losses can be formulated through a bifunction formulation, $l_n(\pi) = f_{\pi_n}(\pi) = \mathbb{E}_{s \sim d^{\pi_n}} [c(s, \pi; \pi^*)]$, and the policy class Π can be viewed as the decision set \mathcal{X} . Naturally, this online learning formulation results in many online IL algorithms resembling standard online learning algorithms, such as follow-the-leader (FTL), which uses full information feedback $l_n(\cdot) = \mathbb{E}_{s \sim d^{\pi_n}} [c(s, \cdot; \pi^*)]$ at each round, [4] and mirror descent [23], which uses the first-order feedback $\nabla l_n(\pi_n) = \mathbb{E}_{d^{\pi_n}} [\nabla_{\pi_n} c(s, \pi_n; \pi^*)]$. This feedback can also be approximated by unbiased samples. The original work by Ross et al. [4] analyzed FTL in the static regret case by immediate reductions to known static regret bounds of FTL. However, a crucial objective is understanding when these algorithms converge to useful solutions in terms of policy performance, which more recent work has attempted to address [5, 6, 24]. According to these refined analyses, dynamic regret is a more appropriate solution concept to online IL when $\pi^* \notin \Pi$, which is the common case in practice.

Below we frame online IL in the proposed COL framework and study its properties based on the properties of COL that we obtained in the previous sections. We have already shown that the per-round loss $l_n(\cdot)$ can be written as the evaluation of a bifunction $f_{\pi_n}(\cdot)$. This COL problem is actually an (α, β) -regular COL problem when the expected supervised learning loss $\mathbb{E}_{s \sim d^{\pi_n}} [c(s, \pi; \pi^*)]$ is strongly convex in π and the state distribution d^π is Lipschitz continuous (see [4, 5, 6]). We can then leverage our results in the COL framework to immediately answer an interesting question in the online IL problem.

Proposition 2. *When $\alpha > \beta$, there exists a unique policy $\hat{\pi}$ that is optimal on its own distribution:*

$$\mathbb{E}_{s \sim d_{\hat{\pi}_n}} [c(s, \hat{\pi}; \pi^*)] = \min_{\pi \in \Pi} \mathbb{E}_{s \sim d_{\hat{\pi}_n}} [c(s, \pi; \pi^*)].$$

This result is immediate from the fact that $\alpha > \beta$ implies that $\nabla f_{\pi}(\pi)$ is a μ -strongly monotone VI with $\mu = \beta - \alpha$ by Proposition 1, which is guaranteed to have a unique solution [16].

Furthermore, we can improve upon the known sufficient conditions required to find this policy through online gradient descent and give a non-asymptotic convergence guarantee through a reduction to strongly monotone VIs. We will additionally assume that f is L -smooth in π , satisfying $\|\nabla f_{\pi'}(\pi_1) - \nabla f_{\pi'}(\pi_2)\| \leq L\|\pi_1 - \pi_2\|$ for any fixed query argument π' .

We then apply the projection algorithm [16], which is equivalent to online gradient descent studied in [23, 6]. Let P_Π denote the Euclidean projection onto Π . The online gradient descent algorithm can

be described as computing the following at each round: $\pi_{n+1} = P_{\Pi}(\pi_n - \eta_n g_n)$ or equivalently

$$\pi_{n+1} = \arg \min_{\pi \in \Pi} \eta_n \langle g_n, \pi \rangle + \frac{1}{2} \|\pi - \pi_n\|^2.$$

Proposition 3 ([16]). *If $\alpha > \beta$ and the stepsize is chosen such that $\eta < \frac{2\mu}{(L+\beta)^2}$, then, under the online gradient descent algorithm with deterministic feedback $g_n = \nabla l_n(\pi_n)$, it holds that*

$$\|\pi_n - \hat{\pi}\|^2 \leq (1 + (L + \beta)^2 \eta^2 - 2\mu\eta)^{n-1} \|\pi_1 - \hat{\pi}\|^2$$

By Theorem 2, Regret_N^d will therefore be sublinear (in fact $\text{Regret}_N^d = O(1)$) and the policy converges linearly to the policy that is optimal on its own distribution, $\hat{\pi}$. The only condition required on the problem itself is $\alpha > \beta$ while the state-of-the-art sufficient condition of [6] additionally requires $\frac{\alpha}{L} > \frac{2\beta}{\alpha}$. The result also gives a new non-asymptotic convergence rate to $\hat{\pi}$.

The above result only considers the case when the feedback is deterministic; i.e., there is no sampling error due to executing the policy on the MDP, and the risk $\mathbb{E}_{d^{\pi_n}}[c(s, \pi; \pi^*)]$ is known exactly at each round. While this is a standard starting point in analysis of online IL algorithms [4], we are also interested in the more realistic stochastic case, which has so far not been analyzed for the online gradient descent algorithm in online IL. It turns out that the COL framework can be easily leveraged here too to provide a sublinear dynamic regret bound.

At round n , we consider observing the empirical risk $\tilde{l}_n(\pi) = \frac{1}{T} \sum_{t=1}^T c(s_t, \pi; \pi^*)$ where $s_t \sim d_t^{\pi_n}$. Note that $\mathbb{E}[\tilde{l}_n(\pi) | \pi_n] = l_n(\pi)$ and it is easy to show that the first-order feedback $\nabla \tilde{l}_n(\pi_n)$ can be modeled as the expected gradient with an additive zero-mean noise: $g_n = \nabla l_n(\pi_n) + \epsilon_n$. For simplicity, we assume $\mathbb{E}[\|\epsilon_n\|^2] < \infty$.

Proposition 4. *If $\alpha > \beta$ and the stepsize is chosen as $\eta_n = \frac{1}{\sqrt{n}}$, then, under online gradient descent with stochastic feedback, it holds that $\mathbb{E}[\text{Regret}_N^d] = O(\sqrt{N})$.*

The proof leverages the reduction to static regret in Corollary 1. It is immediate from the fact that the online IL problem is (α, β) -regular (see Proposition 9 in the full technical report [7] for details). The dynamic regret is worse than that of the deterministic case, but it is still sublinear. This is the price paid for stochastically sampling from the MDP.

6 Conclusion

We present COL, a new class of online learning problems where the gradient varies continuously across rounds with respect to the learner's decisions. We show that this setting can be equated with certain equilibrium problems (EPs) and variational inequalities (VIs). Leveraging this insight, we present conditions for achieving sublinear dynamic regret. Furthermore, we show a reduction from dynamic regret to static regret and the convergence to equilibrium point. This insight suggests that, when these conditions are met, we may employ standard algorithms from the EP literature to achieve interpretable, sublinear dynamic regret rates. Lastly, we apply our theoretical results to the online imitation learning problem, showing that interesting novel results.

There are several directions for future research on this topic. Our current analyses focus on classical algorithms in online learning. We suspect that the use of adaptive or optimistic methods [25] can accelerate convergence to equilibria if some coarse model of the bifunction can be estimated. This is especially relevant in applications on episodic MDPs where the expected losses are exactly determined by an underlying reward function and transition dynamics. In addition to online IL, there are also several iterative optimization problems with MDPs that are interesting to consider in the COL setting. First, the problems of online system identification and structured prediction have also been posed as adversarial online learning and analyzed under static regret [26, 4]. We also note that the classic fitted Q-iteration [27, 28] for reinforcement learning also uses a similar setup. In round n , the loss can be written as $l_n(Q) = \mathbb{E}_{s,a \sim \mu_{\pi(Q_n)}} \mathbb{E}_{s' \sim \mathcal{P}(s,a)} [(Q(s,a) - r(s,a) - \gamma \max_{a'} Q_n(s', a'))^2]$, where $\mu_{\pi(Q_n)}$ is the state-action distribution induced by running a policy $\pi(Q_n)$ based on the Q-function estimate Q_n of the learner. These problem settings can all be posed as COL problems and it would be interesting to see how their algorithms and analyses can be reconciled with those of EP problems via this reduction.

References

- [1] Geoffrey J Gordon. Regret bounds for prediction problems. In *Conference on Learning Theory*, volume 99, pages 29–40, 1999.
- [2] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning*, pages 928–936, 2003.
- [3] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [4] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *International conference on artificial intelligence and statistics*, pages 627–635, 2011.
- [5] Ching-An Cheng and Byron Boots. Convergence of value aggregation for imitation learning. In *International Conference on Artificial Intelligence and Statistics*, pages 1801–1809, 2018.
- [6] Jonathan Lee, Michael Laskey, Ajay Kumar Tanwani, Anil Aswani, and Ken Goldberg. A dynamic regret analysis and adaptive regularization algorithm for on-policy robot imitation learning. In *Workshop on the Algorithmic Foundations of Robotics*, 2018.
- [7] Ching-An Cheng, Jonathan Lee, Ken Goldberg, and Byron Boots. Online learning with continuous variations: Dynamic regret and reductions. *arXiv preprint arXiv:1902.07286*, 2019.
- [8] Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [9] Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [10] Aryan Mokhtari, Shahin Shahrampour, Ali Jadbabaie, and Alejandro Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 7195–7201. IEEE, 2016.
- [11] Tianbao Yang, Lijun Zhang, Rong Jin, and Jinfeng Yi. Tracking slowly moving clairvoyant: optimal dynamic regret of online learning with true and noisy gradient. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 449–457. JMLR. org, 2016.
- [12] Rishabh Dixit, Amrit Singh Bedi, Ruchi Tripathi, and Ketan Rajawat. Online learning with inexact proximal online gradient descent algorithms. *IEEE Transactions on Signal Processing*, 67(5):1338–1352, 2019.
- [13] Omar Besbes, Yonatan Gur, and Assaf Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- [14] Ali Jadbabaie, Alexander Rakhlin, Shahin Shahrampour, and Karthik Sridharan. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, pages 398–406, 2015.
- [15] Lijun Zhang, Tianbao Yang, Jinfeng Yi, Jing Rong, and Zhi-Hua Zhou. Improved dynamic regret for non-degenerate functions. In *Advances in Neural Information Processing Systems*, pages 732–741, 2017.
- [16] Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.
- [17] Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- [18] Arkadi Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

- [19] Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- [20] Martin Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*, pages 427–435, 2013.
- [21] W Robert Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4(3):506–510, 1953.
- [22] Ching-An Cheng, Remi Tachet des Combes, Byron Boots, and Geoff Gordon. A reduction from reinforcement learning to no-regret online learning. *arXiv preprint arXiv:1911.05873*, 2019.
- [23] Wen Sun, Arun Venkatraman, Geoffrey J Gordon, Byron Boots, and J Andrew Bagnell. Deeply aggravated: Differentiable imitation learning for sequential prediction. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3309–3318. JMLR. org, 2017.
- [24] Ching-An Cheng, Xinyan Yan, Evangelos A Theodorou, and Byron Boots. Accelerating imitation learning with predictive models. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- [25] Ching-An Cheng, Xinyan Yan, Nathan Ratliff, and Byron Boots. Predictor-corrector policy optimization. In *International Conference on Machine Learning*, pages 1151–1161, 2019.
- [26] Arun Venkatraman, Martial Hebert, and J Andrew Bagnell. Improving multi-step prediction of learned time series models. In *Conference on Artificial Intelligence*, 2015.
- [27] Geoffrey J Gordon. Stable function approximation in dynamic programming. In *Machine Learning Proceedings 1995*, pages 261–268. Elsevier, 1995.
- [28] Martin Riedmiller. Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pages 317–328. Springer, 2005.