
Online Optimization in Games via Control Theory: Connecting Regret, Passivity and Poincaré Recurrence

Yun Kuen Cheung¹ Georgios Piliouras²

Abstract

We present a novel control-theoretic understanding of online optimization and learning in games, via the notion of *passivity*. Passivity is a fundamental concept in control theory, which abstracts energy conservation and dissipation in physical systems. It has become a standard tool in analysis of general feedback systems, to which game dynamics belong. Our starting point is to show that all continuous-time Follow-the-Regularized-Leader (FTRL) dynamics, which include the well-known Replicator Dynamic, are *lossless*, i.e. it is passive with no energy dissipation. Interestingly, we prove that passivity implies bounded regret, connecting two fundamental primitives of control

theory and online optimization.

The observation of energy conservation in FTRL inspires us to present a family of lossless learning dynamics, each of which has an underlying energy function with a simple gradient structure. This family is closed under *convex combination*; as an immediate corollary, any convex combination of FTRL dynamics is lossless and thus has bounded regret. This allows us to extend the framework of Fox & Shamma (2013) to prove not just global asymptotic stability results for game dynamics, but Poincaré recurrence results as well. Intuitively, when a lossless game (e.g. graphical constant-sum game) is coupled with lossless learning dynamic, their interconnection is also lossless, which results in a pendulum-like energy-preserving recurrent behavior, generalizing (Piliouras & Shamma, 2014; Mertikopoulos et al., 2018).

¹Royal Holloway University of London, Egham, United Kingdom ²Singapore University of Technology and Design, Singapore. Correspondence to: Yun Kuen Cheung <yunkuen.cheung@rhul.ac.uk>, Georgios Piliouras <georgios@sutd.edu.sg>.

Proceedings of the 38th International Conference on Machine Learning, PMLR 139, 2021. Copyright 2021 by the author(s).

| | Lossless Learning Dynamic | Physical System |
|-------------------------------|--|--|
| examples | Replicator Dynamic, Follow-the-Regularized-Leader (FTRL) dynamics | gravity |
| state | $\mathbf{q} = \int \mathbf{p}(t) dt =$ cumulative payoffs | $h =$ vertical height |
| energy | $E(\mathbf{q}) =$ storage function | $V(h) =$ (negative) potential energy |
| gradient of energy | $\mathbf{x} = \nabla E(\mathbf{q}) =$ mixed strategy of the agent | $F = \nabla V(h) =$ gravitational force |
| invariant property | convex combination (CC): any CC of storage functions produces a lossless learning dynamic that yields the same CC of mixed strategies | linear combination (LC): any LC of potential energy function is a potential energy function that yields the same LC of gravitational forces |
| another analogue | $\mathbf{p} =$ instantaneous payoffs | $v =$ velocity |
| change of energy value | $\int \langle \mathbf{x}, \mathbf{p} \rangle dt$ | $\int \langle F, v \rangle dt$ |

Figure 1. Analogues between lossless online learning dynamic and physical system. The evolution of a system of learning dynamics can be thought of as capturing the movements of particles, thus tools from control theory can find direct application in the study of learning dynamics in games.

1. Introduction

Online optimization aims at designing algorithms that can maximize performance in unpredictable and even adversarially evolving environments. The standard benchmark for success in these environments is minimizing regret, which is defined as the difference between the accumulated performance of the algorithm and that of the best action in hindsight. One of the most important achievements of the field has been to establish that such regret minimizing algorithms exist (Cesa-Bianchi & Lugosi, 2006; Shalev-Shwartz, 2012). Amongst the most well-known such algorithms is the class of Follow-the-Regularized-Leader (FTRL) algorithms, which include as special cases ubiquitous meta-algorithms such as Multiplicative Weights Update (Freund & Schapire, 1999; Arora et al., 2005) and Gradient Descent (Shalev-Shwartz, 2012). It is well known that such algorithms can achieve $\mathcal{O}(\sqrt{T})$ regret by employing slowly decreasing step-sizes, and that this bound is effectively optimal given arbitrary payoff sequences. When applying such algorithms in games, as the sequence of payoffs becomes more predictable, stronger regret guarantees are possible (Rakhlin & Sridharan, 2013; Foster et al., 2016; Syrgkanis et al., 2015; Bailey & Piliouras, 2019). In continuous-time model, FTRL dynamics are once again optimal, achieving bounded regret in general settings (Kwon & Mertikopoulos, 2017; Mertikopoulos et al., 2018; Bailey et al., 2020). Hence both from the perspective of optimization as well as game theory, FTRL constitute effectively an optimal choice.

Control theory, on the other hand, is motivated by a seemingly unrelated set of questions. It aims to develop methodologies for stabilizing complex processes and machines. Due to its intrinsic connections to real-world systems, control theory revolves around concepts with a strong grounding in physical systems. A fundamental property of numerous physical systems is *passivity*, which is typically defined in terms of energy dissipation, conservation and transformation (Willems, 1972a;b; Ortega et al., 2013). Passivity is an “input-output” property of a system, and expresses that a system which is supplied with bounded energy can only output bounded energy. Passive systems come equipped with a *storage function* that accumulates the supplied energy but perhaps with some loss (cf. energy loss due to friction in mechanical systems). Overall, passivity encodes a useful notion of stability, since such system cannot explode into unpredictable out-of-control motion as it would correspond to unbounded energy output.

Although the fields of online optimization and control theory are both well developed with long and distinct histories, their interconnection is still rather nascent. Online algorithms can be abstractly thought as input-output operators where the input is a stream of payoffs, and the output is a stream of behavioral outcomes. Both notions of regret and

passivity are similar properties of such input-output algorithms/operators, and encode a notion of predictability and stability around a reference frame. In regret, the reference frame is given by the cumulative payoffs of past actions, in passivity by energy level sets. This raises our first questions:

Are there formal connections between regret and passivity? Moreso, can we interpret the optimal regret of FTRL dynamics from a passivity perspective? Are there similarly optimal learning dynamics / input-output operators?

Any formal connection across the two fields is clearly valuable, as it allows for a fusion of ideas and methodologies between two well-developed fields, and expedite the progress on areas of interest that are common to both, such as game theory (Fudenberg & Levine, 1998; Cesa-Bianchi & Lugosi, 2006; Marden & Shamma, 2015; 2018). For related issues on the intersection of learning, control theory and games, see (Shamma, 2020) and the references therein.

Notably, Fox & Shamma (2013) proposed a control-theoretic framework for analyzing learning in games. One of their key contributions is to identify a *modular approach*, where an analysis can be performed by studying a learning operator (which converts payoff input to strategy output) and a game operator (which converts strategy input to payoff output) *independently*, while the whole game dynamic is a feedback interconnection system of the two operators (see Figure 2). By focusing on coupling learning heuristics that are strictly passive with passive games (e.g. zero-sum games), the resulting strictly passive systems were shown to converge to equilibria, generalizing and unifying numerous prior results, e.g. (Hofbauer & Sandholm, 2009).

The modular approach has allowed numerous works which study learning or game operators separately (Mabrok & Shamma, 2016; Park et al., 2018; Mabrok, 2018; Gadjev & Pavel, 2019). Despite this progress, settings of critical importance for AI such as understanding the perfectly recurrent non-equilibrating behaviors of Gradient Descent (and other FTRL dynamics) in zero-sum games has so far remained outside the reach of these techniques (Piliouras & Shamma, 2014; Mertikopoulos et al., 2018; Balduzzi et al., 2018; Vlatakis-Gkaragkounis et al., 2019; Perolat et al., 2020). This raises our second question:

Can the pendulum-like cyclic behaviors of FTRL dynamics in zero-sum games be understood and generalized via passivity?

Our Contributions. We provide affirmative answers to both questions raised above. We show that any finitely passive¹ learning dynamic guarantees constant regret (Theorem 8). By using the notion of *convex conjugate* from

¹A finitely passive operator is a passive operator with a storage function which is bounded from below.




Figure 2. A feedback interconnection system that captures a game dynamic, by interconnecting a learning operator and a game operator. \hat{r} is (random) perturbation to payoffs; we set it zero here.

convex analysis, we show that any continuous-time FTRL dynamic is finitely passive and lossless (Theorem 7); the same holds for certain *escort replicator dynamics* (Harper, 2011) (see Appendix D). These generalize (Mabrok, 2018) which showed that Replicator Dynamic is finitely lossless. Combining the two theorems above immediately recovers the result in (Mertikopoulos et al., 2018) that any FTRL dynamic guarantees constant regret. We note that in the analysis of Mabrok (2018), the state space (i.e. the domain of the storage function) is the space of mixed strategies, while we turn to a new state space of *cumulative payoffs* for FTRL. This choice is crucial for the generalization to FTRL, and it permits a cleaner proof via tools in convex analysis.

A key observation that enables the above results is that FTRL dynamic admits a storage function with a simple gradient structure, which will be described formally in Section 5. This motivates us to study a new family of learning dynamics, which is in one-one correspondence with the family of storage functions with that gradient structure. By observing that such storage functions are closed under convex combination, we discover that any convex combination of FTRL dynamics is finitely passive and lossless, and thus guarantees constant regret (Theorem 14). “Convex combination of FTRL dynamics” means: Suppose there are k FTRL dynamics, indexed from 1 to k . When we use the j -th one, it converts the cumulative payoffs at time t to a mixed strategy \mathbf{x}_j^t . A convex combination of these FTRL dynamics converts the cumulative payoffs at time t to $\sum_{j=1}^k \alpha_j \cdot \mathbf{x}_j^t$, where α_j ’s are positive constants satisfying $\sum_{j=1}^k \alpha_j = 1$.

Convex combinations of lossless dynamics are directly analogous to linear combinations of conservative vector fields in analyzing physical dynamics (see Figure 1). This technique is also of practical relevance, since we might want to mix-and-match different dynamics to elicit their advantages. For instance, different learning dynamics may lean toward either *exploitation* or *exploration*. By combining them via convex combination with our own choice of α_j ’s, we can control our desired balance between exploitation and exploration (see Example 12).

We also show that for every graphical constant-sum game (e.g. two-person zero-sum game) that admits a fully-

mixed Nash equilibrium, it corresponds to a finitely lossless game operator (Proposition 15). Thus, the game dynamic of any convex combinations of FTRL dynamics in such a game corresponds to a finitely lossless operator. This implies that the game dynamic is almost perfectly recurrent, via the notion of *Poincaré recurrence* (Theorem 17). This distinguishes our work from (Fox & Shamma, 2013) and its subsequent works about learning in games, as they mostly concern stability/convergence, while we study recurrence.

Roadmap. In Sections 2 and 3, we present the necessary background, including the definitions of different operators, the notions of (lossless) passivity and storage function, and some basic results about passivity. In Section 4, we show the desirable properties (finitely losslessness, constant regret) of FTRL dynamics. In Section 5, we present a characterization of lossless learning dynamics via the above-mentioned gradient structure of storage functions, and we discuss some properties of convex combinations of such learning dynamics. Poincaré recurrences of learning in graphical constant-sum games are presented in Section 6. All missing proofs can be found in the supplementary material.

2. Preliminary

In this section, we define the operators depicted in Figure 2. We first define learning dynamic and its corresponding learning operator. When there are multiple agents and each agent is using one learning dynamic, their learning operators can be concatenated naturally to form a *merged learning operator*. Then we define game operator, whose interconnection with a merged learning operator in the manner of Figure 2 is called a *dynamical game system*.

We use a bold lower case to denote a vector variable. Let $\mathbb{R}^+ := [0, +\infty)$ denote the set of non-negative real numbers. In this paper, every function from \mathbb{R}^+ to \mathbb{R}^d is assumed to be square integrable, and we call it a *function-of-time*. An (input-output) *operator* is a mapping whose input and output are both functions-of-time. Let Δ^n denote the probability simplex over n actions, i.e. $\Delta^n := \{(x_1, \dots, x_n) \mid \sum_{j=1}^n x_j = 1; \text{ for } 1 \leq j \leq n, x_j \geq 0\}$. $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the inner product of the vectors \mathbf{a}, \mathbf{b} of same dimension d , i.e. $\langle \mathbf{a}, \mathbf{b} \rangle = \sum_{j=1}^d a_j b_j$.

Learning Dynamic and Learning Operator. We focus on the following type of continuous-time learning dynamics. An agent has an action set A ; let $n := |A|$. The process starts at time 0. For any time $t \geq 0$ and for each action $j \in A$, the agent computes the *cumulative payoff* if she chooses action j in the time interval $[0, t]$, denoted by $q_j(t)$. Formally, let $p_j(\tau)$ denote the (*instantaneous*) *payoff* to action j at time τ . Then $q_j(t) := q_j^0 + \int_0^t p_j(\tau) d\tau$, where q_j^0 is a constant chosen at the beginning of the process. Let $\mathbf{q}(t) := (q_1(t), \dots, q_n(t))$. For any $t \geq 0$, the agent uses a *conversion function* f which takes $\mathbf{q}(t)$ as input, and

outputs a mixed strategy $\mathbf{x}(t) \in \Delta^n$ over the n actions. The process can be expressed compactly as an ordinary differential equations (ODE) system:

$$\begin{aligned} \mathbf{q}(0) &= \mathbf{q}^0 && \text{(Initial condition)} \\ \dot{\mathbf{q}}(t) &= \mathbf{p}(t) && \text{(Cumulative payoff/state update)} \\ \mathbf{x}(t) &= f(\mathbf{q}(t)). && \text{(Behavioral/strategy output)} \end{aligned} \quad (1)$$

When $\mathbf{q}^0 = (q_1^0, \dots, q_n^0)$ is already given, the conversion function f specifies the learning dynamic. The learning dynamic can be viewed as an operator, which takes the function-of-time $\mathbf{p} : \mathbb{R}^+ \rightarrow \mathbb{R}^n$ as input, and the output is another function-of-time $\mathbf{x} : \mathbb{R}^+ \rightarrow \Delta^n$.

Regret. For any $\mathbf{p} : \mathbb{R}^+ \rightarrow \mathbb{R}^n$, the *regret* of a learning dynamic at time $T > 0$ is

$$\left(\max_{j \in A} \int_0^T p_j(\tau) d\tau \right) - \int_0^T \langle \mathbf{p}(\tau), \mathbf{x}(\tau) \rangle d\tau.$$

We say a learning dynamic *guarantees constant regret* if for any $\mathbf{p} : \mathbb{R}^+ \rightarrow \mathbb{R}^n$ and any $T > 0$, the regret at time T is bounded from above by a constant that depends on \mathbf{q}^0 only.

Replicator Dynamic and FTRL Learning Dynamics.

When the conversion function f is the logit choice map

$$f_{\text{RD}}(\mathbf{q}) = \left(\frac{\exp(q_1)}{\mathcal{N}}, \frac{\exp(q_2)}{\mathcal{N}}, \dots, \frac{\exp(q_n)}{\mathcal{N}} \right), \quad (2)$$

where $\mathcal{N} = \sum_{j=1}^n \exp(q_j)$, the learning dynamic (1) is equivalent to the well-known Replicator Dynamic (Hofbauer & Sigmund, 1998; Sandholm, 2010), which is the continuous-time analogue of Multiplicative Weights Update.

A FTRL learning dynamic is specified by a strictly convex *regularizer function* $h : \Delta^n \rightarrow \mathbb{R}$, which determines the conversion function as below:

$$f(\mathbf{q}) = \arg \max_{\mathbf{x} \in \Delta^n} \{ \langle \mathbf{q}, \mathbf{x} \rangle - h(\mathbf{x}) \}. \quad (3)$$

It is known that Replicator Dynamic is a special case of FTRL, by setting $h(\mathbf{x}) = \sum_{j=1}^n x_j \log x_j$. Online Gradient Descent (OGD) is another commonly studied learning dynamic that is also a special case of FTRL with L^2 regularization, i.e. $h(\mathbf{x}) = \frac{1}{2} \sum_{j=1}^n (x_j)^2$. (Hazan et al., 2016) When $n = 2$, the conversion function of OGD is

$$f_{\text{OGD}}(\mathbf{q}) = \begin{cases} (1, 0) & \text{if } q_1 - q_2 \geq 1; \\ (0, 1) & \text{if } q_1 - q_2 \leq -1; \\ \left(\frac{q_1 - q_2 + 1}{2}, \frac{q_2 - q_1 + 1}{2} \right) & \text{otherwise.} \end{cases} \quad (4)$$

Merged Learning Operator. When a system has $m \geq 2$ agents, and each agent uses a learning dynamic, we concatenate the corresponding learning operators together to form a *merged learning operator* (MLO). Precisely, the

input to the MLO is $\hat{\mathbf{p}} = (\mathbf{p}^1, \dots, \mathbf{p}^m)$, and its output is $\hat{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^m)$, where \mathbf{x}^i is the output of the learning operator of agent i when its input is \mathbf{p}^i . In this paper, we use hat notations (e.g. $\hat{\mathbf{p}}, \hat{\mathbf{x}}$) to denote variables formed by such concatenations of variables of individual agents.

Game Operator and Dynamical Game System. Here, we provide a general definition of game operators, and leave the discussion about *graphical constant-sum games*, which appear in our Poincaré recurrence results, to Section 6.

A game has m agents. Each agent i has n_i actions. After each agent chooses a mixed strategy over her own actions, the game determines a payoff vector $\hat{\mathbf{p}} \in \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_m}$, where $\hat{p}_{k\ell}$ is the payoff to action ℓ of agent k . Let $\Delta := \Delta^{n_1} \times \dots \times \Delta^{n_m}$, and $\mathcal{P} := \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_m}$. We can think of the game as a function $G : \Delta \rightarrow \mathcal{P}$. Its game operator takes a function-of-time $\hat{\mathbf{x}} : \mathbb{R}^+ \rightarrow \Delta$ as input, and it outputs a function-of-time $\hat{\mathbf{p}} : \mathbb{R}^+ \rightarrow \mathcal{P}$, where $\hat{\mathbf{p}}(t) \equiv G(\hat{\mathbf{x}}(t))$.

A *dynamical game system* (DGS) comprises of $m \geq 2$ agents. Each agent uses a learning dynamic of the form (1). The agents' learning operators are concatenated to form a MLO, which has input $\hat{\mathbf{p}} = (\mathbf{p}^1, \dots, \mathbf{p}^m)$ and output $\hat{\mathbf{x}} = (\mathbf{x}^1, \dots, \mathbf{x}^m)$ when $\hat{\mathbf{r}} \equiv \mathbf{0}$. The MLO is interconnected with a game operator in the manner of Figure 2. The game operator has input $\hat{\mathbf{x}}$ and output $\hat{\mathbf{p}}$.

3. Passivity

To motivate the notions of passivity and energy, consider an electrical network connected to a power source, where the voltage and current across the network at time τ are respectively $v(\tau)$ and $i(\tau)$. Let $E(t)$ denote the energy stored in the network at time t . We have $E(t) \leq E(0) + \int_0^t v(\tau) \cdot i(\tau) d\tau$. The reason for the inequality (but not an exact equality) is that energy might dissipate from the network. In this setting, the function-of-time v is the input, while the function-of-time i is the output, so the network is indeed an operator.

Passivity of State Space System. To generalize the above idea to passivity of an ODE system, we need several mathematical notations. Let \mathbb{L}_2 denote the Hilbert space of square integrable functions mapping \mathbb{R}^+ to \mathbb{R}^n with inner product: $\langle f, g \rangle_T := \int_0^T \langle f(t), g(t) \rangle dt$. Let $\mathbb{L}_{2,e} := \{f : \mathbb{R}^+ \rightarrow \mathbb{R}^n \mid \langle f, f \rangle_T < \infty \text{ for all } T \in \mathbb{R}^+\}$. An (input-output) operator is simply a mapping $S : \mathbb{U} \rightarrow \mathbb{Y}$, where $\mathbb{U}, \mathbb{Y} \subset \mathbb{L}_{2,e}$.

We consider the following type of operators, which can be represented by an ODE system called *state space system* (SSS) of the following general form:

$$\begin{aligned} \mathbf{z}(0) &= \mathbf{z}^0; \\ \dot{\mathbf{z}}(t) &= g_1(\mathbf{z}(t), \mathbf{u}(t)); \\ \mathbf{y}(t) &= g_2(\mathbf{z}(t), \mathbf{u}(t)), \end{aligned} \quad (5)$$

where $\mathbf{z}^0 \in \mathcal{Z} \subset \mathbb{R}^{d_1}$, $\mathbf{z} : \mathbb{R}^+ \rightarrow \mathcal{Z}$, and $\mathbf{u}, \mathbf{y} : \mathbb{R}^+ \rightarrow \mathbb{R}^{d_2}$. The set \mathcal{Z} is called the set of *states*. As the notations suggest, \mathbf{u}, \mathbf{y} are the input and output of this operator respectively. When \mathbf{u} is fed into the operator, the first two equalities define a well-posed ODE system, so a unique solution of \mathbf{z} exists under mild conditions. Then \mathbf{y} is the output determined by a function g_2 of the unique solution \mathbf{z} and the input \mathbf{u} . The learning dynamic (1) is such an operator, by viewing $\mathbf{q}, \mathbf{p}, \mathbf{x}, f(\mathbf{q}(t))$ in (1) as $\mathbf{z}, \mathbf{u}, \mathbf{y}, g_2(\mathbf{z}(t), \mathbf{u}(t))$ in (5) respectively. We are ready to present the definition of passivity for such operators; note its reminiscence with the motivating example of electrical network.

Definition 1. A SSS is passive if there exists a storage function $L : \mathcal{Z} \rightarrow \mathbb{R}$ such that for all $\mathbf{z}^0 \in \mathcal{Z}$, $t \in \mathbb{R}^+$ and all input-output pairs $\mathbf{u} \in \mathbb{U}$, $\mathbf{y} \in \mathbb{Y}$, we have

$$L(\mathbf{z}(t)) \leq L(\mathbf{z}^0) + \int_0^t \langle \mathbf{u}(\tau), \mathbf{y}(\tau) \rangle d\tau. \quad (6)$$

If the equality always holds, then we say the SSS is lossless passive or simply lossless. If a SSS is passive (resp. lossless) via a storage function L that has a finite lower bound, we say it is finitely passive (resp. finitely lossless).

When a SSS is finitely passive/lossless, we may assume, without loss of generality, that the finite lower bound of its storage function is zero. We make this assumption on all finitely passive/lossless SSS in the rest of this paper.

Feedback Interconnection System. In Figure 2, we presented a feedback interconnection (FIC) system. While it is intuitive for readers to have a first understanding of the concept in control theory, for our analysis it is more convenient to use Figure 3. The FIC in Figure 3 consists of two SSS $S_1 : \mathbb{U} \rightarrow \mathbb{Y}$ and $S_2 : \mathbb{Y} \rightarrow \mathbb{U}$, with an external input source $\mathbf{r} \in \mathbb{U}$, while its output is $\mathbf{y}_1 \in \mathbb{Y}$; note that the FIC is an operator by definition. The variables are related via: $\mathbf{u}_1 = \mathbf{r} - \mathbf{y}_2$, $\mathbf{y}_1 = S_1(\mathbf{u}_1)$, $\mathbf{u}_2 = \mathbf{y}_1$, and $\mathbf{y}_2 = S_2(\mathbf{u}_2)$.

An important property of passive operators is that passivity is composable, i.e. the composition of two passive operators results in a passive system. Intuitively, if no operator in the system is able to produce energy, then the system as a whole cannot produce energy either. The theorem below formally captures this intuition for FIC systems.

Theorem 2 ((Fox & Shamma, 2013) Theorem 3.2). *Consider the FIC system in Figure 3. Suppose that for $i = 1, 2$, S_i is passive via storage function L_i . Then the FIC system is a passive operator via storage function $L_1 + L_2$. Precisely, for any $\mathbf{z}_1^0 \in \mathcal{Z}_1$, $\mathbf{z}_2^0 \in \mathcal{Z}_2$ and $t \in \mathbb{R}^+$,*

$$\begin{aligned} & L_1(\mathbf{z}_1(t)) + L_2(\mathbf{z}_2(t)) \\ & \leq L_1(\mathbf{z}_1^0) + L_2(\mathbf{z}_2^0) + \int_0^t \langle \mathbf{r}(\tau), \mathbf{y}_1(\tau) \rangle d\tau. \end{aligned}$$

If S_1, S_2 are lossless, then the FIC system is lossless via storage function $L_1 + L_2$, i.e. the inequality above becomes an equality.

DGS as a FIC system. In the context of DGS, Figure 3 is obtained after several modifications from Figure 2:

- In Figure 2, the game operator's output is $\hat{\mathbf{p}}$, which is added to $\hat{\mathbf{r}}$ to form the MLO's input. In Figure 3, the game operator's output is $-\hat{\mathbf{p}}$ instead, and it is subtracted from \mathbf{r} to form the MLO's input.
- The MLO's output is $\hat{\mathbf{x}}$ in Figure 2, but the MLO's output in Figure 3 has a constant shift $\hat{\mathbf{x}}^* = (\mathbf{x}^{*,1}, \dots, \mathbf{x}^{*,m}) \in \Delta$. Precisely, the output is $\hat{\mathbf{x}} - \hat{\mathbf{x}}^*$. We call this operator the "MLO with shift $\hat{\mathbf{x}}^*$ ".
- In Figure 3, the game operator's input is $\hat{\mathbf{x}} - \hat{\mathbf{x}}^*$ instead of $\hat{\mathbf{x}}$, while its output is $-\hat{\mathbf{p}}$ instead of $\hat{\mathbf{p}}$.

Basic Results about Passivity of Learning Operators.

By viewing a DGS as a FIC system, we are interested in MLO and game operators which possess good properties like passivity. The key advantage of this approach is it permits a DGS to be decoupled into two distinct operators, which can be analysed separately.

S_1 is a MLO. First, we show that it is passive if all learning operators possessed by the agents are passive. This allows us to turn our focus to analyzing whether each individual learning operator is passive/lossless or not.

Proposition 3. *Suppose for each agent i , her learning operator with shift $\mathbf{x}^{*,i}$ is passive (resp. lossless) via a storage function L^i . Then the MLO with shift $\hat{\mathbf{x}}^* = (\mathbf{x}^{*,1}, \dots, \mathbf{x}^{*,m})$ is passive (resp. lossless) via the storage function $\sum_{i=1}^m L^i$.*

By the next proposition, the choice of shift does *not* affect passivity of the learning operator. Thus, we say a learning dynamic is passive when its learning operator with any shift is passive.

Proposition 4. *Let S^a, S^b be two learning operators of the same learning dynamic, with shifts $\mathbf{x}^{*,a}, \mathbf{x}^{*,b}$ respectively. S^a is passive (resp. lossless) via storage function $L^a(\mathbf{q})$ if and only if S^b is passive (resp. lossless) via storage function $L^a(\mathbf{q}) - \langle (\mathbf{x}^{*,b} - \mathbf{x}^{*,a}), \mathbf{q} \rangle + c$, where c is any constant.*

The above proposition works even when the shifts are not mixed strategies, e.g. when $\mathbf{x}^{*,a}$ is the zero vector. We use E to denote a storage function of a learning operator with zero shift, and L to denote a storage function of a learning operator with a shift of a mixed strategy.

While a shift does not affect passivity, it *does* affect whether the learning operator is *finitely* passive or not. In order to prove that certain learning dynamics guarantee constant regret, we need their learning operators with some specific shifts to be finitely passive. This motivates the following definition of *finitely passive learning dynamics*. Let \mathbf{e}_j denote the vector with the j -th entry be one, and all other entries be zero.




Figure 3. A FIC system which includes two operators S_1 and S_2 . We assume $\mathbf{r} \equiv 0$ in the discussion below. S_1 is a merged learning operator which converts payoffs to mixed strategies of agents with a shift $\hat{\mathbf{x}}^*$, while S_2 is a game operator which converts the mixed strategies of agents (with shift $\hat{\mathbf{x}}^*$) to the negation of payoffs.

Definition 5. A learning dynamic is *finitely passive* (resp. *finitely lossless*) if for every action j , its learning operator with shift \mathbf{e}_j is finitely passive (resp. finitely lossless).

Proposition 6. If a learning dynamic is finitely passive (resp. finitely lossless), then for any mixed strategy \mathbf{x}^* , the learning operator with shift \mathbf{x}^* is finitely passive (resp. finitely lossless).

4. Passivity of Learning Operators

4.1. FTRL Dynamics are Finitely Lossless

We start the analysis by establishing a strong connection between FTRL dynamics and passivity. Specifically, FTRL dynamics are finitely lossless.

Theorem 7. Given any FTRL dynamic over n actions and with regularizer function h , let the convex conjugate of h be

$$h^*(\mathbf{q}) := \max_{\mathbf{x} \in \Delta^n} \{\langle \mathbf{q}, \mathbf{x} \rangle - h(\mathbf{x})\}. \quad (7)$$

Then for any $\mathbf{x}^* \in \Delta^n$, the learning operator with shift \mathbf{x}^* is finitely lossless via the storage function $L(\mathbf{q})$ given below:

$$L(\mathbf{q}) = h^*(\mathbf{q}) - \langle \mathbf{q}, \mathbf{x}^* \rangle + h(\mathbf{x}^*). \quad (8)$$

In particular, for any action j , the learning operator with shift \mathbf{e}_j is finitely lossless, and hence any FTRL dynamic is finitely lossless by Definition 5.

Proof: By the ‘‘maximizing argument’’ identity in p. 149 of (Shalev-Shwartz, 2012), we have $\nabla h^*(\mathbf{q}(t)) = \mathbf{x}(t)$. Hence, $\nabla L(\mathbf{q}(t)) = \nabla h^*(\mathbf{q}(t)) - \mathbf{x}^* = \mathbf{x}(t) - \mathbf{x}^*$. By the chain rule, $\frac{dL(\mathbf{q}(t))}{dt} = \langle \mathbf{x}(t) - \mathbf{x}^*, \mathbf{p}(t) \rangle$, and hence $L(\mathbf{q}(t)) = L(\mathbf{q}(0)) + \int_0^t \langle \mathbf{x}(\tau) - \mathbf{x}^*, \mathbf{p}(\tau) \rangle d\tau$, certifying losslessness. Moreover, L is bounded below by zero, since by the definition of h^* , for any \mathbf{q} , $L(\mathbf{q}) = h^*(\mathbf{q}) - \langle \mathbf{q}, \mathbf{x}^* \rangle + h(\mathbf{x}^*) \geq (\langle \mathbf{q}, \mathbf{x}^* \rangle - h(\mathbf{x}^*)) - \langle \mathbf{q}, \mathbf{x}^* \rangle + h(\mathbf{x}^*) = 0$. \square

We summarize the properties of the operators of FTRL dynamics with various shifts in Figure 4.

| | FTRL with zero shift | FTRL with shift of a mixed strategy \mathbf{x}^* |
|-------------------------|--|--|
| input | \mathbf{p} (payoff) | \mathbf{p} |
| output | \mathbf{x} (mixed strategy) | $\mathbf{x} - \mathbf{x}^*$ |
| state | \mathbf{q} (cumulative payoff) | \mathbf{q} |
| storage function | $E(\mathbf{q}) = h^*(\mathbf{q})$ (see (7)) | $L(\mathbf{q}) = h^*(\mathbf{q}) - \langle \mathbf{q}, \mathbf{x}^* \rangle + h(\mathbf{x}^*)$ |
| infimum | $-\infty$ | 0 |
| property | lossless | finitely lossless |

Figure 4. FTRL learning operators with various shifts.

4.2. Relationship to Regret

Theorem 8. Any finitely passive learning dynamic guarantees constant regret.

Proof: Let L^j denote the storage function of the learning operator with shift \mathbf{e}_j . Since the learning dynamic is finitely passive, we can assume the infimum of L^j is zero. By the definition of passivity,

$$L^j(\mathbf{q}(t)) \leq L^j(\mathbf{q}^0) + \int_0^t \langle \mathbf{x}(\tau) - \mathbf{e}_j, \mathbf{p}(\tau) \rangle d\tau.$$

Hence, the regret w.r.t. action j at time t satisfies:

$$\begin{aligned} & \int_0^t \langle \mathbf{e}_j, \mathbf{p}(\tau) \rangle d\tau - \int_0^t \langle \mathbf{x}(\tau), \mathbf{p}(\tau) \rangle d\tau \\ & \leq L^j(\mathbf{q}^0) - L^j(\mathbf{q}(t)) \leq L^j(\mathbf{q}^0). \end{aligned}$$

Thus, the regret up to time t is bounded by $\max_j \{L^j(\mathbf{q}^0)\}$, which is a constant that depends only on the initial state \mathbf{q}^0 . \square

An immediate corollary of Theorems 7 and 8, all FTRL dynamics guarantee constant regret.

Corollary 9 ((Mertikopoulos et al., 2018)). *Every FTRL dynamic guarantees constant regret.*

5. A Characterization of Lossless Learning Dynamics, and Their Convex Combinations

In Section 4.1, we showed that every FTRL dynamic is lossless. A FTRL dynamic is specified by a convex regularizer function h , while the storage function L has an implicit form in term of h . Here, we ask how we can specify the storage function directly to *generate* a lossless learning dynamic. Not all storage functions work, so we first seek some necessary conditions on them. These conditions are stated using the storage function E of the learning operator with zero shift. By Proposition 4, the lossless storage function of the learning operator with shift \mathbf{x}^* is $E(\mathbf{q}) - \langle \mathbf{q}, \mathbf{x}^* \rangle + c$.

Suppose there is a lossless learning dynamic in the form of (1); recall that when \mathbf{q}^0 is already given, the learning dynamic is specified by its conversion function f . Let E be the storage function of its learning operator with zero shift. We present two necessary conditions on E and f .

Necessary Condition 1 (NC1). (i) For any $j = 1, \dots, n$, $\nabla_j E(\mathbf{q}) \geq 0$. (ii) For any \mathbf{q} , $\nabla E(\mathbf{q}) = f(\mathbf{q})$.

Reason: Since the learning dynamic is lossless, $E(\mathbf{q}(t)) = E(\mathbf{q}(0)) + \int_0^t \langle \mathbf{x}(\tau), \mathbf{p}(\tau) \rangle d\tau$. Taking time-derivative on both sides yields $\frac{dE(\mathbf{q}(t))}{dt} = \langle \mathbf{x}(t), \mathbf{p}(t) \rangle = \langle f(\mathbf{q}(t)), \mathbf{p}(t) \rangle$. On the other hand, by the chain rule, we have $\frac{dE(\mathbf{q}(t))}{dt} = \langle \nabla E(\mathbf{q}(t)), \dot{\mathbf{q}} \rangle = \langle \nabla E(\mathbf{q}(t)), \mathbf{p}(t) \rangle$. Since these two equalities hold for any $\mathbf{p}(t)$ and $\mathbf{q}(t)$, they readily imply $\nabla E(\mathbf{q}) = f(\mathbf{q})$ for all \mathbf{q} , i.e. condition (ii) holds. Since $f(\mathbf{q}) \in \Delta^n$, condition (i) holds.

Necessary Condition 2 (NC2). For any real number r and any \mathbf{q} , $E(\mathbf{q} + r \cdot \mathbf{1}) = E(\mathbf{q}) + r$.

Reason: By NC1, $\nabla E(\mathbf{q}) = f(\mathbf{q})$, which is in Δ^n . Thus, the directional derivative of E along the direction $\mathbf{1}$ is $\langle \nabla E(\mathbf{q}), \mathbf{1} \rangle = \langle f(\mathbf{q}), \mathbf{1} \rangle = 1$.

Indeed, it is easy to verify that the above two necessary conditions are also sufficient.

Proposition 10. Any smooth function $E(\mathbf{q})$ which satisfies NC1(i) and NC2 is a lossless storage function of a learning operator with no shift. The conversion function f satisfies $f(\mathbf{q}) = \nabla E(\mathbf{q})$.

The family of smooth functions E satisfying NC1(i) and NC2 can be represented compactly as

$$\mathcal{E} = \{E : \mathbb{R}^n \rightarrow \mathbb{R} \mid \forall \mathbf{q}, \nabla E(\mathbf{q}) \geq \mathbf{0} \ \& \ \langle \nabla E(\mathbf{q}), \mathbf{1} \rangle = 1\}.$$

An interesting observation is that this family is closed under *convex combination*, i.e. if $E_1, \dots, E_k \in \mathcal{E}$, then for any real numbers $\alpha_1, \dots, \alpha_k \geq 0$ such that $\sum_{\ell=1}^k \alpha_\ell = 1$, $(\sum_{\ell=1}^k \alpha_\ell \cdot E_\ell) \in \mathcal{E}$. By Proposition 10, E_1, \dots, E_k are lossless storage functions of some learning operators with no shift. Suppose the conversion functions are f_1, \dots, f_k respectively. Then by Proposition 10 again, $(\sum_{\ell=1}^k \alpha_\ell \cdot E_\ell)$

is a lossless storage function of a learning operator with no shift, with conversion function $(\sum_{\ell=1}^k \alpha_\ell \cdot f_\ell)$. This motivates the following definition.

Definition 11. Given k learning dynamics, each over n actions, let f_ℓ denote the conversion function of the ℓ -th learning dynamic. Given any non-negative constants $\alpha_1, \dots, \alpha_k$ where $\sum_{\ell=1}^k \alpha_\ell = 1$, the convex combination of the k learning dynamics with parameters $\alpha_1, \dots, \alpha_k$ is a learning dynamic with conversion function $(\sum_{\ell=1}^k \alpha_\ell \cdot f_\ell)$.

Example 12. Suppose we are using the half-half convex combination of Replicator Dynamic (RD) and Online Gradient Descent (OGD), and $n = 2$. Recall the conversion functions of RD and OGD in (2) and (4). Suppose that at some time t , $\mathbf{q}(t) = (0.6, 0)$. Then the mixed strategy at time t with the half-half convex combination is

$$\begin{aligned} \mathbf{x}(t) &= \frac{1}{2} \cdot f_{\text{RD}}((0.6, 0)) + \frac{1}{2} \cdot f_{\text{OGD}}((0.6, 0)) \\ &\approx \frac{1}{2} [(0.6457, 0.3543) + (0.8, 0.2)] \approx (0.7228, 0.2772). \end{aligned}$$

By (4), $f_{\text{OGD}}(\mathbf{q})$ outputs a strategy with zero probability of choosing action 2 whenever $q_1 - q_2 \geq 1$. In contrast, $f_{\text{RD}}(\mathbf{q})$ maintains a tiny but positive probability of choosing action 2 even when q_1 is much larger than q_2 . We may say OGD leans toward exploitation while RD leans toward exploration. By combining the two learning dynamics via convex combination with our own choice of α 's, we obtain a new lossless learning dynamic with our desired balance between exploitation and exploration.

Convex combination not only preserves losslessness, but also finitely losslessness. Suppose there are several finitely lossless learning dynamics. By Definition 5, for every action j , the storage functions of their learning operators with shift \mathbf{e}_j have finite lower bounds. It is easy to verify that for a convex combination of these lossless learning dynamics, its learning operator with shift \mathbf{e}_j is lossless via the same convex combination of the storage functions mentioned above. Since the storage functions have finite lower bounds, their convex combination has a finite lower bound too.

Theorem 13. Given any k lossless (resp. finitely lossless) learning dynamics, any convex combination of them is a lossless (resp. finitely lossless) learning dynamic.

Theorems 7, 13 and 8 lead to the following interesting theorem.

Theorem 14. Any convex combination of any finitely lossless learning dynamics is a learning dynamic that guarantees constant regret. In particular, any convex combination of any FTRL learning dynamics is a learning dynamic that guarantees constant regret.

Remark. Note that the family of storage functions that generate *finitely* lossless learning dynamics

is $\{E \in \mathcal{E} \mid \forall j, E(\mathbf{q}) - \langle \mathbf{q}, \mathbf{e}_j \rangle \text{ is bounded from below}\}$. Again, this family is closed under convex combination. By Theorem 7, the family of FTRL learning dynamics is a subset of the family of finitely lossless learning dynamics. It is not clear if the two families are equal; we believe not. For instance, for the half-half convex combination of Replicator Dynamic and Online Gradient Descent, we cannot find any regularizer function that validates it is a FTRL dynamic.

6. Lossless DGS and Poincaré Recurrence

We have presented a potentially broad family of (finitely) lossless learning dynamics. In this section, our focus is on DGS in which agents use such learning dynamics. We first prove that for certain *graphical constant-sum games*, their game operators are finitely lossless. Thus, for any DGS comprising of such a game and finitely lossless learning dynamics, it admits a constant-of-motion by Theorem 2 (when $\mathbf{r} \equiv 0$), i.e. the sum of the two storage function values is a constant for any $t \geq 0$. Then we use this constant-of-motion and follow a principled approach proposed by Mertikopoulos et al. (2018) to show our main result here: the DGS is Poincaré recurrent. In the rest of this section, we first define graphical constant-sum game and Poincaré recurrence, then we apply the principled approach to prove our main result.

Graphical Constant-sum Game. There are m agents. Each agent i has n_i actions. In a *graphical game* (Kearns et al., 2001), there is a normal-form (matrix) game between every pair of agents, which we call an *edge-game*. The edge game between agents i, k is specified by two matrices, $\mathbf{A}^{ik} \in \mathbb{R}^{n_i \times n_k}$ and $\mathbf{A}^{ki} \in \mathbb{R}^{n_k \times n_i}$. When each agent $k \neq i$ chooses a mixed strategy $\mathbf{x}_k \in \Delta^{n_k}$, the payoff vector of agent i , which contains the payoff to each action of agent i , is the sum of the payoff vectors in all her edge-games, i.e.

$$\mathbf{p}_i = \sum_{1 \leq k \leq m, k \neq i} \mathbf{A}^{ik} \cdot \mathbf{x}_k. \quad (9)$$

A *graphical constant-sum game* (Daskalakis & Papadimitriou, 2009) is a graphical game such that for every pair of agents $\{i, k\}$, there exists a constant $c^{\{i,k\}}$ satisfying $\mathbf{A}_{j\ell}^{ik} + \mathbf{A}_{\ell j}^{ki} = c^{\{i,k\}}$ for any action j of agent i and action ℓ of agent k .

Game Operator. Recall that a game operator has an input of mixed strategies of different agents with shift $\hat{\mathbf{x}}^*$, while the output is the negative of the cumulative payoff vector to different actions. We point out one useful fact, which follows quite readily from (Piliouras & Shamma, 2014).

Proposition 15. *If $\hat{\mathbf{x}}^*$ is a Nash equilibrium of a graphical constant-sum game, then the game operator with shift $\hat{\mathbf{x}}^*$ is passive; the storage function is the zero function. Moreover, if $\hat{\mathbf{x}}^*$ is fully-mixed (i.e. every entry in $\hat{\mathbf{x}}^*$ is positive), then the game operator is lossless via the zero storage function.*

When the game is lossless, by Theorem 2, if it is coupled with passive (resp. lossless) learning dynamics, then the DGS is passive (resp. lossless).

Poincaré Recurrence. In a DGS, $\hat{\mathbf{p}}$ is a function of $\hat{\mathbf{x}}$ via the game operator, while $\hat{\mathbf{x}}$ is a function of $\hat{\mathbf{q}}$ via the conversion function. Thus, $\hat{\mathbf{p}}$ is a function of $\hat{\mathbf{q}}$. We say the ODE system (1) is *divergence-free* if $\sum_{i=1}^m \sum_j \frac{\partial p_{ij}}{\partial q_{ij}}$ is zero everywhere. When $\hat{\mathbf{p}}$ is derived using a graphical game via (9), p_{ij} does not depend on q_{ij} since for every $k \neq i$, \mathbf{x}_k is a function of \mathbf{q}_k only. Thus, the game dynamic is divergence-free.

Intuitively, an ODE system is Poincaré recurrent if almost all trajectories return arbitrarily close to their initial position infinitely often. In order to work formally with the notion of Poincaré recurrence, we need to define a measure on \mathbb{R}^n . We use the standard Lebesgue measure on \mathbb{R}^n . Liouville's formula states that divergence-free ODE system preserves volume (Weibull, 1995) (see Appendix A for more discussion of volume and Liouville's formula). Thus, the following Poincaré Recurrence Theorem is applicable if its bounded-orbit requirement is satisfied.

Theorem 16 (Poincaré, 1890; Barreira, 2006). *If a transformation preserves volume and has bounded orbits, then it is Poincaré recurrent, i.e. for each open set there exist orbits that intersect this set infinitely often.*

Given any $\epsilon > 0$, we can cover \mathbb{R}^n by countably many balls of radius ϵ , and apply the theorem to each ball. We conclude that almost every point returns to within an ϵ neighbourhood of itself. Since $\epsilon > 0$ is arbitrary, we conclude that almost every initial point is almost recurrent.

Poincaré Recurrence in DGS. Recall that in each learning operator of a DGS, the state is represented by a vector $\mathbf{q} \in \mathbb{R}^n$, which represents the cumulative payoffs. Clearly, it can be unbounded as time goes even in a graphical constant sum game², thus prohibiting us to use Theorem 16. Instead, as in (Mertikopoulos et al., 2018), we consider a transformation that maps $\mathbf{q} = (q_1, \dots, q_n)$ to $(q'_1, q'_2, \dots, q'_{n-1}) := (q_1 - q_n, \dots, q_{n-1} - q_n) \in \mathbb{R}^{n-1}$. It is well-known that for any FTRL dynamic with starting point $\mathbf{q}^0 = (q_1^0, \dots, q_n^0)$ and conversion function $f : \mathbb{R}^n \rightarrow \Delta^n$, it is equivalent to the following dynamic with state variables $\mathbf{q}' = (q'_1, \dots, q'_{n-1}) \in \mathbb{R}^{n-1}$, whereas $1 \leq j \leq n-1$:

$$\begin{aligned} q'_j(0) &= q_j^0 - q_n^0 \\ \dot{q}'_j(t) &= p_j(t) - p_n(t) \\ \mathbf{x}(t) &= f(q'_1(t), q'_2(t), \dots, q'_{n-1}(t), 0). \end{aligned} \quad (10)$$

Given a mixed strategy \mathbf{x}^* , if we cast the output of the above dynamic to be $(x_1(t) - x_1^*, \dots, x_{n-1}(t) - x_{n-1}^*)$,

²For instance, this happens for a two-person zero-sum game with every payoff entry to agent 1 is strictly positive.

then it is easy to verify that the learning operator is finitely lossless via the storage function $\bar{L}(q'_1, \dots, q'_{n-1}) = L(q'_1, \dots, q'_{n-1}, 0)$, where L is the storage function of the original learning operator with shift \mathbf{x}^* .

Theorem 17. *Poincaré recurrence occurs in the strategy space Δ for any dynamical game system where (1) each agent employs a learning dynamic which is a convex combination of FTRL; and (2) the underlying game is a graphical constant-sum game with a fully-mixed Nash equilibrium.*

See Figure 5 for an example of Poincaré recurrence under conditions (1) and (2). To prove the theorem, we first show that Poincaré recurrence occurs in the space that contains $\hat{\mathbf{q}}'$ ($\hat{\mathbf{q}}'$ is the concatenation of the \mathbf{q}' of all agents). This comprises of two major steps:

- The dynamic preserves *volume*, since the dynamic is divergence-free.
- For any starting point $\hat{\mathbf{q}}'(0)$, the dynamic remains bounded. To show this, we use the fact that \bar{L} is a constant-of-motion of the game dynamic, so for any $t \geq 0$, $\hat{\mathbf{q}}'(t)$ must stay within a level set of \bar{L} , which is bounded (see Appendix C).

Poincaré recurrence in the space that contains $\hat{\mathbf{q}}'$ implies Poincaré recurrence in the strategy space Δ , since the conversion function in (10) is continuous.

7. Conclusion

We present a control-theoretic perspective to understanding popular learning dynamics like FTRL and escort replicator dynamics. At the heart of it is the use of storage (energy) functions to govern how the dynamic turns history of pay-offs to strategic choices. This mirrors the study of physical dynamics, e.g. electrical networks (Khalil, 2015). Analysis via storage functions permits us to prove optimal regret bounds and inspires interesting generalizations of FTRL via convex combinations.

An important benefit of these control-theoretic tools is they allow decoupling of game dynamics into learning operators and game operators. This provides a framework to understand learning-in-games via a modular approach, by analyzing these operators separately. This technique can liberate us from analyzing each individual learning-in-game system in ad-hoc manner.

In our work, we initiate the study of connections between online optimization and control theory with continuous-time learning dynamics. An interesting problem is how such a connection can be generalized to discrete-time learning *algorithms*, e.g. Multiplicative Weights Update and its optimistic variant (Bailey & Piliouras, 2018; Cheung, 2018; Daskalakis & Panageas, 2019; Cheung & Piliouras, 2020). There does exist theory that generalizes passivity




Figure 5. The trajectories of learning in a graphical constant-sum game called *Cyclic Matching Pennies*. In this game, there are three

agents, each has two actions. $\mathbf{A}^{12} = \mathbf{A}^{23} = \mathbf{A}^{31} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$,

while $\mathbf{A}^{21} = \mathbf{A}^{32} = \mathbf{A}^{13} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$. The mixed strategies of the three agents are denoted by $(x_1, 1 - x_1)$, $(x_2, 1 - x_2)$ and $(x_3, 1 - x_3)$ respectively. All trajectories start with $x_1 = 0.9$, $x_2 = 0.88$ and $x_3 = 0.4$ (the black dot in the figure). The trajectories are simulated when all agents use Replicator Dynamics (blue), Online Gradient Descent (Black), and half-half convex combination of the former two dynamics (red) respectively. All trajectories are recurrent.

to discrete-time settings (Desoer & Vidyasagar, 1975). We hope our work inspires further studies in this direction.

Lastly, we believe that this control-theoretic perspective is also useful for understanding learning dynamics/algorithms which are not always passive. The perspective can help us spot under which situations the learning dynamics/algorithms create or dissipate energy. By avoiding situations where energy is created, it is possible that we can achieve stable outcomes in learning processes.

Acknowledgements

This research/project is supported in part by NRF2019-NRF-ANR095 ALIAS grant, grant PIE-SGP-AI-2018-01, NRF 2018 Fellowship NRF-NRFF2018-07, AME Programmatic Fund (Grant No. A20H6b0151) from the Agency for Science, Technology and Research (A*STAR) and the National Research Foundation, Singapore under its AI Singapore Program (AISG Award No: AISG2-RP-2020-016).

References

- Arora, S., Hazan, E., and Kale, S. The multiplicative weights update method: a meta algorithm and applications. Technical report, 2005.
- Bailey, J. and Piliouras, G. Fast and furious learning in zero-sum games: vanishing regret with non-vanishing step sizes. In *Advances in Neural Information Processing Systems*, pp. 12977–12987, 2019.
- Bailey, J. P. and Piliouras, G. Multiplicative weights update in zero-sum games. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pp. 321–338. ACM, 2018. doi: 10.1145/3219166.3219235.
- Bailey, J. P., Gidel, G., and Piliouras, G. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pp. 391–407. PMLR, 2020.
- Balduzzi, D., Racaniere, S., Martens, J., Foerster, J., Tuyls, K., and Graepel, T. The mechanics of n-player differentiable games. In *ICML*, 2018.
- Barreira, L. Poincare recurrence: old and new. In *XIVth International Congress on Mathematical Physics. World Scientific.*, pp. 415–422, 2006.
- Cesa-Bianchi, N. and Lugosi, G. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- Cheung, Y. K. Multiplicative weights updates with constant step-size in graphical constant-sum games. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 3532–3542, 2018.
- Cheung, Y. K. and Piliouras, G. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pp. 807–834. PMLR, 2019.
- Cheung, Y. K. and Piliouras, G. Chaos, extremism and optimism: Volume analysis of learning in games. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Daskalakis, C. and Panageas, I. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pp. 27:1–27:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi: 10.4230/LIPICs.ITCS.2019.27.
- Daskalakis, C. and Papadimitriou, C. H. On a network generalization of the minmax theorem. In *ICALP 2009: Proceedings of the 2009 International Colloquium on Automata, Languages, and Programming*, 2009.
- Desoer, C. and Vidyasagar, M. *Feedback Systems: Input-output Properties*. Electrical science series. Academic Press, 1975. ISBN 9780122120503.
- Foster, D. J., Lykouris, T., Sridharan, K., and Tardos, E. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems*, pp. 4727–4735, 2016.
- Fox, M. J. and Shamma, J. S. Population games, stable games, and passivity. *Games*, 4(4):561–583, 2013.
- Freund, Y. and Schapire, R. E. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- Fudenberg, D. and Levine, D. K. *The Theory of Learning in Games*. MIT Press Books. The MIT Press, 1998.
- Gadjov, D. and Pavel, L. A passivity-based approach to nash equilibrium seeking over networks. *IEEE Transactions on Automatic Control*, 64(3):1077–1092, 2019.
- Harper, M. Escort evolutionary game theory. *Physica D: Nonlinear Phenomena*, 240(18):1411 – 1415, 2011. ISSN 0167-2789. doi: 10.1016/j.physd.2011.04.008.
- Hazan, E. et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Hofbauer, J. and Sandholm, W. H. Stable games and their dynamics. *Journal of Economic theory*, 144(4):1665–1693, 2009.
- Hofbauer, J. and Sigmund, K. *Evolutionary Games and Population Dynamics*. Cambridge University Press, Cambridge, UK, 1998.
- Kearns, M. J., Littman, M. L., and Singh, S. P. Graphical models for game theory. In *UAI*, 2001.
- Khalil, H. K. *Nonlinear Control*. Pearson Education, 2015.
- Kwon, J. and Mertikopoulos, P. A continuous-time approach to online optimization. *Journal of Dynamics and Games*, 4(2):125–148, April 2017.

- Mabrok, M. A. Passivity Analysis of Replicator Dynamics and its Variations. *arXiv e-prints*, art. arXiv:1812.07164, Dec 2018.
- Mabrok, M. A. and Shamma, J. S. Passivity analysis of higher order evolutionary dynamics and population games. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 6129–6134, 2016.
- Marden, J. R. and Shamma, J. S. Game theory and distributed control. In *Handbook of game theory with economic applications*, volume 4, pp. 861–899. Elsevier, 2015.
- Marden, J. R. and Shamma, J. S. Game theory and control. *Annual Review of Control, Robotics, and Autonomous Systems*, 1:105–134, 2018.
- Mertikopoulos, P., Papadimitriou, C., and Piliouras, G. Cycles in adversarial regularized learning. In *ACM-SIAM Symposium on Discrete Algorithms*, 2018.
- Ortega, R., Perez, J. A. L., Nicklasson, P. J., and Siramirez, H. J. *Passivity-based control of Euler-Lagrange systems: mechanical, electrical and electromechanical applications*. Springer Science & Business Media, 2013.
- Park, S., Shamma, J. S., and Martins, N. C. Passivity and evolutionary game dynamics. In *2018 IEEE Conference on Decision and Control (CDC)*, pp. 3553–3560, 2018.
- Perolat, J., Munos, R., Lespiau, J.-B., Omidshafiei, S., Rowland, M., Ortega, P., Burch, N., Anthony, T., Balduzzi, D., De Vylder, B., et al. From poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. *arXiv preprint arXiv:2002.08456*, 2020.
- Piliouras, G. and Shamma, J. S. Optimization despite chaos: Convex relaxations to complex limit sets via poincaré recurrence. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pp. 861–873. SIAM, 2014.
- Poincaré, H. Sur le problème des trois corps et les équations de la dynamique. *Acta mathematica*, 13(1):A3–A270, 1890.
- Rakhlin, A. and Sridharan, K. Online learning with predictable sequences. In *COLT*, 2013.
- Sandholm, W. H. *Population Games and Evolutionary Dynamics*. MIT Press, 2010.
- Shalev-Shwartz, S. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012. ISSN 1935-8237. doi: 10.1561/22000000018.
- Shamma, J. S. Feedback control perspectives on learning, 2020. URL https://neurips.cc/virtual/2020/protected/invited_16168.html.
- Syrgkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. Fast convergence of regularized learning in games. In *Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS’15*, pp. 2989–2997, Cambridge, MA, USA, 2015. MIT Press.
- Vlatakis-Gkaragkounis, E.-V., Flokas, L., and Piliouras, G. Poincaré recurrence, cycles and spurious equilibria in gradient-descent-ascent for non-convex non-concave zero-sum games. In *Advances in Neural Information Processing Systems*, pp. 10450–10461, 2019.
- Weibull, J. W. *Evolutionary Game Theory*. MIT Press; Cambridge, MA: Cambridge University Press., 1995.
- Willems, J. C. Dissipative dynamical systems part i: General theory. In *Archive for Rational Mechanics and Analysis*, volume 45(5), pp. 321–351. 1972a.
- Willems, J. C. Dissipative dynamical systems part ii: Linear systems with quadratic supply rates. In *Archive for Rational Mechanics and Analysis*, volume 45(5), pp. 352–393. 1972b.