
Boosting for Control of Dynamical Systems

Naman Agarwal¹ Nataly Brukhim^{2,1} Elad Hazan^{2,1} Zhou Lu²

Abstract

We study the question of how to aggregate controllers for dynamical systems in order to improve their performance. To this end, we propose a framework of boosting for online control. Our main result is an efficient boosting algorithm that combines weak controllers into a provably more accurate one. Empirical evaluation on a host of control settings supports our theoretical findings.

1. Introduction

In many learning scenarios it is significantly easier to come up with a mildly accurate rule of thumb than state of the art performance. This motivation led to the development of ensemble methods and boosting (Schapire & Freund, 2012), a theoretically sound methodology to combine rules of thumb (often referred to as weak learners) into a substantially more accurate learner.

The application of boosting has transformed machine learning across a variety of applications, including supervised learning: classification (Freund & Schapire, 1997), regression (Mason et al., 2000), online learning (Beygelzimer et al., 2015b;a), agnostic learning (Kanade & Kalai, 2009), recommendation systems (Freund et al., 2003) and many more.

While the same motivation for boosting exists for dynamical systems, i.e. it is often easy to come up with a reasonable predictor or a controller for a dynamical system, the theory and practice of boosting faces significant challenges in these settings due to the existence of a state. Formally, a dynamical system is specified by a rule $x_{t+1} = f(x_t, u_t) + w_t$. The problem of optimal control in dynamical systems requires the design of a sequence of controls $\{u_t\}$ so as to minimize a certain objective (for example making the states follow a certain trajectory). As can be seen readily, the decisions made by a controller affects the future trajectory

of the system, and hence it is often not a-priori clear how to obtain a meaningful guarantee when switching between or aggregating different controllers.

In this paper we propose a framework for formalizing boosting in the context of optimal control of dynamical systems. The first crucial insight comes from the newly emerging literature on non-stochastic control, which allows a non-counterfactual description of the dynamics. Then leveraging techniques from online learning with memory (Anava et al., 2015) and online gradient boosting (Beygelzimer et al., 2015b), we provide a boosting algorithm for controlling systems with state and prove appropriate theoretical guarantees on its performance.

The notion of boosting in online learning. Notice that unlike supervised learning, in the case of dynamical systems it is not immediately clear what metric should be used to measure the improvement provided by boosting algorithms. The theory of online gradient boosting (Beygelzimer et al., 2015b) suggests to boost online learning by improving the *expressivity* of the predictors. Given a set of weak online learners, online boosting guarantees prediction (as measured by mistake bounds or regret) that is as good as a predictor in a larger class.

Furthermore, a robust way to study the optimal control problem is to study it in the online non-stochastic setting (Agarwal et al., 2019), where both the objective to be minimized and the perturbations to the system get revealed in an online fashion¹.

Motivated by this, we take an online boosting approach to optimal control. Our boosting algorithm when given access to (weak) controllers from a certain class, provides a boosted controller than can provably perform (in terms of regret) as well as a controller from the larger class of a committee (or convex combination) of the weak controllers. Furthermore, we provide an alternate boosting algorithm, which is more efficient in terms of the number of weak controllers required, and allows for the utilization of weak controllers designed for handling only quadratic losses into strong controllers that can compete against the more general class of smooth and strongly-convex losses.

¹Google AI Princeton ²Department of Computer Science, Princeton University. Correspondence to: <namanagarwal@google.com, {nbrukhim,ehazan,zhou}@princeton.edu>.

¹We formally define the notion of non-stochastic control in Section 2.1.

We provide the relevant background and setup in Section 2. We formally describe our algorithms and provide formal guarantees on their performance in Section 3. We conclude with experimental evaluation of our methods on a variety of control and learning tasks in dynamical systems.

Dynamics with bounded memory. For our boosting techniques to induce bounded regret we require that the dynamical systems considered have negligible dependence on history beyond a certain amount of time steps in the past. We quantify this property as H -bounded memory and formally define it in Definition 2.1. This assumption is analogous to a standard assumption in Reinforcement Learning, called *mixability* of the underlying Markov Decision Process. The so called *episodic* setting in RL is also often used to circumvent long-term dependencies.

In control theory, the bounded memory assumption manifests in two forms. The stronger notion called *stability*, posits that the system remains bounded over any sequence of actions. This is often considered to be a very strong assumption. A weaker assumption commonly used is *stabilizability* or *controllability*, which posits the existence of a controller or a policy which generates stable actions. Upon action with such a controller, the system de-facto exhibits a bounded memory. In this paper, we assume that all policy classes we work with induce a bounded memory on the dynamical system (or analogously are *fast-mixing* or *stabilizing*).

Boosting vs. overparametrization of deep controllers.

As opposed to supervised learning, in control there are many situations in which there is limited availability of training data, if at all. While deep neural networks have proven extremely successful for large data regimes in supervised learning, we experimentally find that for control, boosting small network controllers results in superior performance to training of a large network controller.

2. Background and Setup

Our treatment of boosting applies only to certain dynamical systems and control methods. The main requirement we place is a memory bound on the effect of past actions. This section formally describes the dynamics setting and underlying assumptions for our algorithms to apply.

2.1. Non-stochastic Control

A general framework of robust optimal control has emerged recently which rests on analyzing optimal control in an online non-stochastic setting (Agarwal et al., 2019; Hazan et al., 2019; Simchowitz et al., 2020). In this framework, at each round $t \in [T]$ the controller observes the state of the system $x_t \in \mathbb{R}^k$, and outputs an action $u_t \in \mathcal{U} \subset \mathbb{R}^d$, where

\mathcal{U} is a convex bounded set. The adversary then reveals a convex cost function and the loss $c_t(x_t, u_t)$ is incurred by the controller. The system then transitions to a new state x_{t+1} according to the following law with f representing the dynamics of the system,

$$x_{t+1} = f(x_t, u_t) + w_t, \quad (1)$$

where $w_t \in \mathbb{R}^k$ is an adversarially chosen perturbation to the dynamics that the system suffers at each time step. The costs and the perturbations are not known to the controller in advance and are assumed to be revealed to the controller after it has committed to the action u_t . The task of the controller is to minimize regret defined in the following way

$$\text{Regret} = \sum_t c_t(x_t, u_t) - \min_{\pi \in \Pi} \sum_t c_t(x_t(w_{1:t}, \pi), u_t^\pi)$$

Here Π represents a class of policies that we wish to compare to. Furthermore $x_t(w_{1:t}, \pi)$ is the state that the system would have reached when executing π on the perturbed dynamics with the *same* perturbations $\{w_1 \dots w_t\}$. Observe that this is a counterfactual notion of regret, since the actions performed by the comparator affect the future cost suffered by the comparator.

Note that the assumption of observable perturbation is without loss of generality when the underlying system f is known to the controller and the state is fully observable. We do not make these assumptions in the paper but rather work with the setting of the complete observation of w as in (Agarwal et al., 2019). Furthermore, we make no distributional assumptions on w_t and only assume $\|w_t\|_2 \leq W$, for some $W > 0$.

Perturbation Based Policies The reference class of policies Π we consider in this paper is comprised of policies π that map a sequence of perturbations $w_1 \dots w_{t-1}$ to an action u_t . Note that this class of policies is only more general than the standard notion of policies which map the current x_t to an action u_t . In particular, it captures linear policies for linear dynamical systems (Section 5). A crucial property of such policies is that the decisions depend directly on the underlying dynamics and do not depend directly on the control feedback, but only implicitly via the perturbations.

Another important limitation we place on policies and dynamics in this paper is memory boundedness, as we now define.

Definition 2.1 ((H, ε) -Bounded Memory). *Given a dynamical system as given in Equation 1, for a sequence of actions u_1, \dots, u_T and any time t , let \hat{x}_t be the state reached by the system if we artificially set $x_{t-H+1} = 0$ and simulate the system with the actions u_{t-H+1}, \dots, u_t .*

The sequence of actions $\{u_1 \dots u_T\}$ is considered to be of (H, ε) -bounded memory if for all t ,

$$|c_t(\hat{x}_t, u_t) - c_t(x_t, u_t)| \leq \varepsilon.$$

Assumption 2.2 (Bounded Memory of Convex Combinations). For a given dynamical system the class of (H, ε) -memory bounded sequences is closed under convex combination.

Note that the notion of bounded memory is a slightly stronger notion than that of *controllability* in the sense that it is applicable to changing policies as well. This notion is also referred to as sequential strong stability in the work of (Cohen et al., 2018). The key point is that the effect of the distant past (beyond H most recent actions and disturbances) on the current state is negligible. Multiple previous works exhibit policy classes which produce bounded memory actions (Cohen et al., 2018; Agarwal et al., 2019; Cohen et al., 2019). Concretely, in Section 5 we describe the GPC controller from (Agarwal et al., 2019) which is shown to be memory bounded as well as satisfy Assumption 2.2.

Reduction to Bounded Memory Functions For a sequence of actions, we now define a proxy function which only penalizes the last H actions ²:

$$\ell_t(\mathbf{0}, u_{t-H+1}, \dots, u_t) := c_t(\hat{x}_t, u_t) \quad (2)$$

(H, ε) -bounded memory of actions now ensures that minimizing regret over the proxy costs ℓ_t , which have finite memory, is sufficient to minimize overall regret. Having reduced the control of dynamical systems to minimizing regret over functions with memory, we are now ready to discuss the technique of Online Boosting, which we will apply on the proxy cost function.

2.2. Online Boosting

The presence of state in non-stochastic control makes online boosting more challenging. We first give a brief background on online boosting for the regression setting (Beygelzimer et al., 2015a), and in Section 3 discuss a reduction that enables the use of a similar technique for non-stochastic control.

Informally, online boosting refers to a meta-learning algorithm which is given black-box oracle access to an online (weak) learning algorithm \mathcal{A} for a function class Π and linear losses, with regret R , and is given a bound N on the total number of calls made in each iteration to copies of \mathcal{A} . The algorithm then obtains an online learning algorithm \mathcal{A}' for a richer function class $\Pi' = \text{conv}(\Pi)$ (i.e. the convex

²Each ℓ_t naturally depends on the sequence of $\{w_t : w_{t-H}\}$ chosen by the adversary. We suppress this dependence for notational convenience.

hull of Π), and any convex losses, with a (possibly larger) regret R' .

The online booster maintains N instances of the weak learning algorithm, denoted $\mathcal{A}_1, \dots, \mathcal{A}_N$. In each round $t \in [T]$, an adversary selects an example x_t from a compact feature space \mathcal{X} , and a loss function $\ell_t : \mathcal{X} \rightarrow \mathbb{R}^k$, and presents x_t to the learner. The policy regret $R(T)$ of each weak learner \mathcal{A}_i is assumed to be bounded as

$$\sum_{t=1}^T \ell_t(\mathcal{A}_i(x_t)) - \min_{\pi \in \Pi} \sum_{t=1}^T \ell_t(\pi(x_t)) \leq R(T),$$

with the regret $R(T)$ being a non-decreasing sub-linear function of T . Note the slight abuse of notation here; $\mathcal{A}_i(\cdot)$ is not a function but rather the output of the online learning algorithm \mathcal{A}_i computed on the given example using its internal state. In each round t , the online booster takes some convex combination of all the predictions made by the learners, and outputs the boosted prediction. To update $\mathcal{A}_1, \dots, \mathcal{A}_N$ at every iteration $t \in [T]$, the booster passes a carefully chosen loss function to each of the weak learners. Specifically, each learner \mathcal{A}_i is fed with a residual loss function $\ell_t^i(y) = \nabla(y_t^{i-1}) \cdot y$, where y_t^{i-1} is a convex combination of previous weak learner predictions, $\mathcal{A}_1(x_t), \dots, \mathcal{A}_{i-1}(x_t)$.

The work of (Beygelzimer et al., 2015a) proves that this technique results in a regret bound of

$$\sum_{t=1}^T \ell_t(y_t) - \min_{\pi \in \text{conv}(\Pi)} \sum_{t=1}^T \ell_t(\pi(x_t)) \leq R(T) + O\left(\frac{T}{N}\right).$$

where y_1, \dots, y_T are the predictions outputted by the booster. Note that although the regret of the boosting algorithm is larger by $O(T/N)$ than the regret of the weak learners, it is now achieved against the best predictor available in a richer class. This is especially meaningful when the class of predictors Π is e.g., neural networks, a highly non-convex class. Thus, by boosting such predictors, the resulting algorithm is guaranteed to have low regret with respect to the convex hull of Π . Our method is based on the online boosting technique, as detailed next.

3. Algorithms and Main Results

This section describes our algorithms for boosting in dynamical systems. The main idea of our methods is to leverage the memory boundedness and reduce online control of dynamical systems to online learning with finite memory (Anava et al., 2015). We achieve this by constructing a proxy cost function which only takes into account the H most recent rounds of the system (see Equation 2). We then extend the online boosting methodology (discussed in Subsection 2.2) to apply to these losses with memory. Bounded memory ensures that minimizing regret over our constructed proxy costs is sufficient to minimize overall regret.

Algorithm	Class	Loss	Regret
DBoost 1	conv(Π)	linear	$R + T/N$
DBoost 2	Π	quadratic	$R + T(1 - \frac{\alpha}{\beta})^N$

Table 1. Main results summary. Boosting uses N weak controllers which have low regret $R = o(T)$, against a reference class of predictors Π . The *DynaBoost1* algorithm allows to compete with the best committee (convex combination) of weak controllers conv(Π). *DynaBoost2*, which is more efficient (requires smaller N), suited for losses that are α -strongly convex and β -smooth. *DynaBoost1* and *DynaBoost2* assume weak controller guarantees hold w.r.t. linear and quadratic losses, respectively.

We propose two algorithms (1, 2) for boosting online control, given access to *weak* controllers (see definitions below) which obtain low regret against a policy class Π and class of losses \mathcal{L} . For the first algorithm we assume \mathcal{L} to be the class of *linear* losses as detailed in Subsection 3.1. For the second algorithm we assume \mathcal{L} to be the class of *quadratic* losses as detailed in Subsection 3.1.

Although the second method requires stronger assumptions, its advantage is that it is more efficient in terms of the number of copies N of weak controllers required to achieve low regret.

3.1. DynaBoost1: Boosting Online Control

Consider the non-stochastic control setting described in Subsection 2.1, for a dynamical system as defined in Equation 1. DynaBoost1 is presented as Algorithm 1 and assumes an oracle access to a weak controller, which is defined as follows:

Definition 3.1. Let \mathcal{A}_i be an online learning algorithm for a dynamical system as defined in Equation 1 and a reference policy class Π . The learner \mathcal{A}_i is a **weak controller** with respect to a class of loss functions \mathcal{L} if

1. The sequence of actions produced by \mathcal{A}_i is of (H, ε) -bounded memory (see Definition 2.1).
2. When run with losses ℓ_t^i chosen from the class of loss functions \mathcal{L} , it produces a sequence of actions u_1, \dots, u_T s.t.,

$$\sum_{t=1}^T \ell_t^i(u_1, \dots, u_t) - \min_{\pi \in \Pi} \sum_{t=1}^T \ell_t^i(u_1^\pi, \dots, u_t^\pi) \leq R(T).$$

where action u_t^π is obtained by applying $\pi \in \Pi$ the best policy in hindsight, and the regret $R(T)$ is a non-decreasing sub-linear function of the horizon T .

We can now construct the proxy *linear* cost functions $\ell_t^i(u_{t-H+1}, \dots, u_t)$ which only consider the H most recent rounds (see line 11 of Algorithm 1), thus obtaining the

Algorithm 1 DynaBoost 1

- 1: Maintain N weak learners $\mathcal{A}_1, \dots, \mathcal{A}_N$.
- 2: Set step length $\eta_i = \frac{2}{i+1}$ for $i \in [N]$.
- 3: **for** $t = 1, \dots, T$ **do**
- 4: Receive the state x_t .
- 5: Define $u_t^0 = \mathbf{0}$.
- 6: **for** $i = 1$ to N **do**
- 7: Define $u_t^i = (1 - \eta_i)u_t^{i-1} + \eta_i \mathcal{A}_i(x_t)$.
- 8: **end for**
- 9: Output action $u_t = u_t^N$.
- 10: Receive loss ℓ_t , suffer $\ell_t(u_1, \dots, u_t)$.
- 11: Define linear loss function:

$$\ell_t^i(\mathbf{u}_1 \dots \mathbf{u}_H) \triangleq \sum_{j=1}^H \nabla_{j,i}^\top \mathbf{u}_j$$

where, $\nabla_{j,i} := \nabla_{t-H+j} \ell_t(\mathbf{0}, u_{t-H+1}^{i-1}, \dots, u_t^{i-1})$.

- 12: Pass loss function $\ell_t^i(\cdot)$ to weak controller \mathcal{A}_i .
- 13: **end for**

following regret guarantee,

$$\sum_{t=1}^T \ell_t^i(u_{t-H+1}, \dots, u_t) - \min_{\pi \in \Pi} \sum_{t=1}^T \ell_t^i(u_{t-H+1}^\pi, \dots, u_t^\pi) \leq R(T) + 2T\varepsilon. \quad (3)$$

Before stating our main theorem, we need the following definition. We say that a loss function ℓ is β -smooth if for all u_1, \dots, u_H and $\tilde{u}_1, \dots, \tilde{u}_H$, it holds that,

$$\begin{aligned} & \ell(u_1, \dots, u_H) - \ell(\tilde{u}_1, \dots, \tilde{u}_H) \\ & \leq \sum_{j=1}^H \nabla_{j,i} \ell(\tilde{u}_1, \dots, \tilde{u}_H)^\top (u_j - \tilde{u}_j) + \frac{\beta}{2} \sum_j \|u_j - \tilde{u}_j\|_2^2 \end{aligned} \quad (4)$$

Under these assumptions we can now give our main theorem, providing a regret bound for Algorithm 1.

Theorem 3.2. Let \mathcal{L}' be the class of β -smooth loss functions. Assume oracle access to N copies of a weak controller \mathcal{A} (see Definition 3.1) satisfying Equation 3. Let $D_{\mathcal{U}}$ be the diameter of the action set \mathcal{U} . Then, there exists a boosting algorithm (Algorithm 1) which produces a sequence of actions u_t for which the following regret bound holds with respect to the reference class conv(Π),

$$\begin{aligned} & \sum_{t=1}^T \ell_t(u_1, \dots, u_t) - \min_{\pi \in \text{conv}(\Pi)} \sum_{t=1}^T \ell_t(u_1^\pi, \dots, u_t^\pi) \\ & \leq \frac{2\beta D_{\mathcal{U}}^2 HT}{N} + R(T) + 4T\varepsilon. \end{aligned}$$

The proof is given in Section 4.

3.2. DynaBoost2: Fast-Boosting Online Control

We now present our results for the case when the loss functions we compete with are strongly convex. In this case we prove that the excess regret of boosting goes down exponentially in the number of weak learners. The weak learners required for this result are stronger in the sense that they are able to have low regret against quadratic functions as opposed to linear functions in the previous part. Due to this, the boosted algorithm does not compete with an expanded class of predictors but rather just with the original class of predictors Π .

In addition to assumptions in the previous subsection we will need the following additional assumptions. We say a loss function ℓ is α -strongly convex when for all u_1, \dots, u_H and $\tilde{u}_1, \dots, \tilde{u}_H$,

$$\begin{aligned} & \ell(u_1, \dots, u_H) - \ell(\tilde{u}_1, \dots, \tilde{u}_H) \\ & \geq \sum_{j=1}^H \nabla_{j,i} \ell(\tilde{u}_1, \dots, \tilde{u}_H)^\top (u_j - \tilde{u}_j) + \frac{\alpha}{2} \sum_j \|u_j - \tilde{u}_j\|_2^2 \end{aligned} \quad (5)$$

Furthermore we say ℓ is G -bounded if for all $u_1, \dots, u_H \in \mathcal{U}$ we have that $|\ell(u_1, \dots, u_H)| \leq G$.

Under these assumptions we can now give our main theorem, providing a regret bound for Algorithm 2.

Theorem 3.3. *Let \mathcal{L}' be the class of α strongly convex and G -bounded loss functions. Assume oracle access to N copies of a weak controller \mathcal{A} (see Definition 3.1) satisfying Equation 3 with respect to the class \mathcal{L} of α -strongly convex quadratic functions. Then, there exists a boosting algorithm (Algorithm 2) which produces a sequence of actions u_t for which the following regret bound holds with respect to the reference class Π ,*

$$\begin{aligned} & \sum_{t=1}^T \ell_t(u_1, \dots, u_t) - \min_{\pi \in \Pi} \sum_{t=1}^T \ell_t(u_1^\pi, \dots, u_t^\pi) \\ & \leq (1 - \frac{\alpha}{\beta})^N 2GT + R(T) + 4T\varepsilon. \end{aligned}$$

We provide the proof of Theorem 3.2 next. The proof of Theorem 3.3 which follows a similar argument is deferred to the Appendix.

4. Proof of Theorem 3.2

Proof. First, note that for any $i = 1, 2 \dots N$, since $\ell_t^i \in \mathcal{L}$, the loss function encountered by the weak controller

Algorithm 2 DynaBoost 2

- 1: Maintain N weak learners $\mathcal{A}_1, \dots, \mathcal{A}_N$.
- 2: Set step length $\eta_i = \frac{\alpha}{\beta}$ for $i \in [N]$.
- 3: **for** $t = 1, \dots, T$ **do**
- 4: Receive the state x_t .
- 5: Define $u_t^0 = \mathbf{0}$.
- 6: **for** $i = 1$ to N **do**
- 7: Define $u_t^i = (1 - \eta_i)u_t^{i-1} + \eta_i \mathcal{A}_i(x_t)$.
- 8: **end for**
- 9: Output action $u_t = u_t^N$.
- 10: Receive loss ℓ_t , suffer $\ell_t(u_1, \dots, u_t)$.
- 11: Define quadratic loss function

$$\ell_t^i(\mathbf{u}_1 \dots \mathbf{u}_H) \triangleq \sum_{j=1}^H \frac{\eta_i \beta}{2} \|\mathbf{u}_j - u_{t-H+j}^{i-1}\|^2 +$$

$$\sum_{j=1}^H \left(\nabla_{j,i}^\top (\mathbf{u}_j - u_{t-H+j}^{i-1}) \right).$$

- 12: where, $\nabla_{j,i} := \nabla_{t-H+j} \ell_t(\mathbf{0}, u_{t-H+1}^{i-1}, \dots, u_t^{i-1})$.
- 13: Pass loss function $\ell_t^i(\cdot)$ to weak controller \mathcal{A}_i .
- 14: **end for**

(defined in Line 11 of 1), is a linear function, we have that:

$$\begin{aligned} & \min_{\pi \in \text{conv}(\Pi)} \sum_{t=1}^T \ell_t^i(u_{t-H+1}^\pi, \dots, u_t^\pi) \\ & = \min_{\pi \in \Pi} \sum_{t=1}^T \ell_t^i(u_{t-H+1}^\pi, \dots, u_t^\pi) \end{aligned}$$

Now let π be any function in $\text{conv}(\Pi)$. Observe that by the equality above and the regret bound of the weak controller (Equation 3), we get,

$$\begin{aligned} & \sum_{t=1}^T \left(\ell_t^i(u_{t-H+1}, \dots, u_t) - \ell_t^i(u_{t-H+1}^\pi, \dots, u_t^\pi) \right) \\ & \leq R(T) + 2T\varepsilon. \end{aligned} \quad (6)$$

Denote $j^- = t - j + 1$ for brevity. Define for any $i \in [N]$, $t \in [T]$, and any $\ell_t \in \mathcal{L}'$ loss function encountered by the booster,

$$\Delta_{t,i} \triangleq \ell_t(\mathbf{0}, u_{H-}^i, \dots, u_t^i) - \ell_t(\mathbf{0}, u_{H-}^\pi, \dots, u_t^\pi).$$

Consider the following calculations for $\Delta_{t,i}$:

$$\begin{aligned}
 \Delta_{t,i} &= \ell_t \left(\mathbf{0}, u_{H^-}^{i-1} + \eta_i (\mathcal{A}_i(x_{H^-}) - u_{H^-}^{i-1}), \dots, \right. \\
 &\quad \left. u_t^{i-1} + \eta_i (\mathcal{A}_i(x_t) - u_t^{i-1}) \right) - \ell_t(\mathbf{0}, u_{H^-}^\pi, \dots, u_t^\pi) \\
 &\text{(by substituting } u_t^i \text{ as in line 7 of Algorithm 1)} \\
 &\leq \ell_t(\mathbf{0}, u_{H^-}^{i-1}, \dots, u_t^{i-1}) - \ell_t(\mathbf{0}, u_{H^-}^\pi, \dots, u_t^\pi) + \\
 &\quad \sum_{j=1}^H \left(\eta_i \nabla_{j,i}^\top (\mathcal{A}_i(x_{t-H+j}) - u_{t-H+j}^{i-1}) + \right. \\
 &\quad \left. \frac{\eta_i^2 \beta}{2} \|\mathcal{A}_i(x_{t-H+j}) - u_{t-H+j}^{i-1}\|^2 \right) \\
 &\text{(by convexity and } \beta\text{-smoothness of } \ell_t, \\
 &\quad \text{and definition of } \nabla_{j,i} \text{ (line 11, Algorithm 1))}
 \end{aligned}$$

Denote $\Delta_i = \sum_t \Delta_{t,i}$. Then, by summing over $t \in [T]$, and applying the weak-controller regret bound (Equation 6), we have,

$$\begin{aligned}
 \Delta_i &\leq \sum_{t=1}^T \left((\ell_t(\mathbf{0}, u_{H^-}^{i-1}, \dots, u_t^{i-1}) - \ell_t(\mathbf{0}, u_{H^-}^\pi, \dots, u_t^\pi)) + \right. \\
 &\quad \left. \eta_i \sum_{j=1}^H \nabla_{j,i}^\top (u_{t-H+j}^\pi - u_{t-H+j}^{i-1}) \right) + \\
 &\quad \eta_i (R(T) + 2T\varepsilon) + \frac{\eta_i^2 \beta D_U^2 HT}{2} \\
 &\leq (1 - \eta_i) \Delta_{i-1} + \eta_i (R(T) + 2T\varepsilon) + \frac{\eta_i^2 \beta D_U^2 HT}{2}
 \end{aligned}$$

where we used the bound $\|\mathcal{A}_i(x_{t-H+j}) - u_{t-H+j}^{i-1}\|^2 \leq 2D_U$. For $i = 1$, since $\eta_1 = 1$, the above bound implies that $\Delta_1 \leq \frac{\beta D_U^2 HT}{2} + (R(T) + 2T\varepsilon)$. Starting from this base case, by induction on $i \geq 1$ it follows that $\Delta_i \leq \frac{2\beta D_U^2 HT}{i} + (R(T) + 2T\varepsilon)$. Applying the above bound for $i = N$ yields the desired result for truncated memory losses. Lastly, using Assumption 2.2 completes the proof. \square

5. Case Studies

For the sake of clarity, we precisely spell out the application of our boosting algorithms with two choices of weak learning methods to illustrate the general technique of applying our boosting algorithm.

5.1. Boosting Deep Controllers

Consider a controller based on a Recursive Neural Network (RNN) for the non-stochastic control problem with dynamics (1). As motivated earlier we explicitly enforce the H -memory bounded property via the choice of the sequence

length of the RNN. Formally, the weak learners in this setting are deep neural networks RNN_θ that map a sequence of H past perturbations $w_{t:t-H} = w_t, \dots, w_{t-H}$ to control:

$$u_{t+1} = \text{RNN}_\theta(w_{t-H:t}).$$

Here by θ we denote the internal weights of the network.

When used inside Algorithm 1, each weak learner $\text{RNN}_\theta = \mathcal{A}_i$ is an instance of neural net that is initialized arbitrarily. Iteratively, the network RNN_θ receives w_t and predicts u_{t+1}^i using a sequential feed forward computation over $w_{t-H:t}$. It then receives the residual loss function $\ell_t^i(\cdot)$. It then applies the back-propagation algorithm to update its internal weights.

5.2. Boosting for Linear Dynamical Systems

A linear dynamical system is governed by the dynamics equation

$$x_{t+1} = Ax_t + Bu_t + w_t, \quad (7)$$

The system is assumed to be known and strongly stable (See Definition 3.3 in (Agarwal et al., 2019)). We use the controller presented in (Agarwal et al., 2019) (referred to as Gradient Perturbation Controller (GPC)) as the weak learners. The GPC controller parameterizes the control actions u_t via the following equation:

$$u_t = -Kx_t + \sum_{i=1}^H M^i w_{t-i} \quad (8)$$

where K is a fixed pre-computed matrix (depending only on A, B) and $M = (M^1, \dots, M^H)$ are parameters governing the controller over which the controller learns. As shown in (Agarwal et al., 2019), K can be selected such that the strong stability property of the system implies that the actions u_t are $(O(\log(T/\varepsilon)), \varepsilon)$ -bounded memory (see Theorem 5.3 in (Agarwal et al., 2019)). Furthermore it can be easily checked that the actions also satisfy Assumption 2.2.

Having setup the weak controller thus we feed it inside Algorithm 1. Similar to the setting with the deep networks, iteratively, the controller receives w_t and predicts u_{t+1}^i using the GPC prediction. Furthermore, it then receives the residual loss function $\ell_t^i(\cdot)$ and the internal parameters are updated according to the GPC update.

6. Experiments

We have tested our framework of online boosting given in Algorithm 1 in various control settings, as detailed below.

The first weak-controller we have tested is the Gradient Perturbation Controller (GPC) discussed above (see Subsection 5.2), presented in Figure 1. In addition, we also give results for a RNN-based controller (see Subsection 5.1), presented

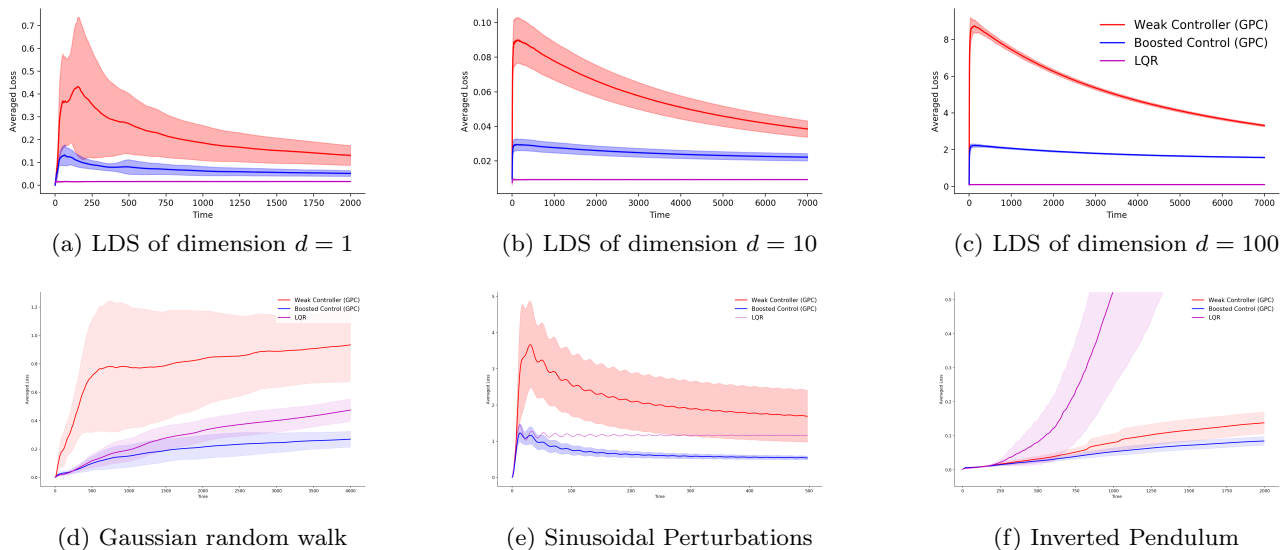


Figure 1. Online Boosting Control with GPC weak-controllers

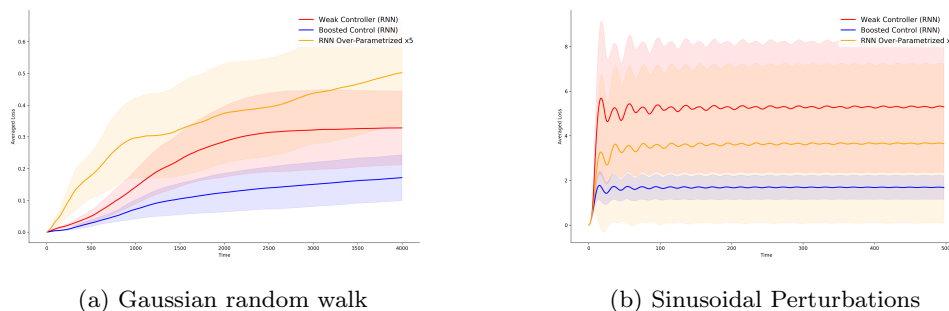


Figure 2. Online Boosting Control with RNN weak-controllers

in Figure 2. The weak-controller baselines, and the weak controllers fed to the boosting method, are the exact same controllers, with identical configuration per setting. Note that in all settings, weak-controllers performance (plotted in red) can be improved by applying boosting (plotted in blue).

We begin with experiments on a Linear Dynamical System (as in Equation 7) where the matrices A, B are generated randomly. We then present experiments for a non-linear dynamics as well (Inverted Pendulum setting). The cost function used in all settings is $c(x, u) = \|x\|_2^2 + \|u\|_2^2$. The GPC weak-controller is designed as in Equation 8, following (Agarwal et al., 2019), with the pre-fixed matrix K set to 0. The RNN weak-controller, using an LSTM architecture, with 5 hidden units. In all figures, we plot the averaged results for a fixed system, which differs per setting, over 20 experiment runs with different stochasticity. Confidence intervals of .95 are plotted in each setting as well.

Sanity check experiments. To demonstrate the effectiveness of the system in terms of both (i) its ability to reach close to a known optimal controller, and (ii) its performance in different dimensions, we present the results of this setting, shown in the first row of Figure 1. For the system used in each dimension $d \in \{1, 10, 100\}$ (with $d = k$ in all settings), each noise term w_t is normally i.i.d. distributed with zero mean, and 0.1^2 variance. We set the memory length to $H = 5$, and use $N = 5$ weak-learners in all the experiments. The Linear Quadratic Regulator (LQR) is known to be optimal in this setting and therefore this experiment only serves as a sanity check.

Correlated disturbances experiments. We now consider more challenging LDS settings in which the disturbances w_t are correlated across time. In the "Gaussian random walk" setting, each noise term is distributed normally, with the previous noise term as its mean (specifically,

$w_{t+1} \sim \mathcal{N}(w_t, 0.3^2)$), and is clipped to the range $[-1, 1]$. In the "Sinusoidal Perturbations" setting, the sine function is applied to the time index, such that, $w_t = \sin(t)/2\pi$.

Note that in these settings the LQR method is no longer optimal due to perturbations being correlated across time. The RNN-based controllers perform better than GPC-based controllers in the "Gaussian random walk" setting, whereas in the "Sinusoidal Perturbations" setting, GPC outperforms RNNs. However, in both cases, Boosting improves upon its corresponding weak-controller.

Boosting vs. Over-Parameterization In Figure 2, the "Over-parametrized RNN" baseline refers to a baseline controller of the same architecture and hyper-parameters as the RNN-weak controller, but with a larger hidden layer. We demonstrate that by using a larger network with overall same number of parameters as the boosted RNN controller, boosting achieves superior performance. Notice that enlarging the size of the network might result in a controller that is outperformed even by the smaller RNN controller, as in Figure 2(a). Overall, this experiment implies that the strength of our method does not stem from using more parameters, but rather from the way in which the weak-controllers are maintained by the boosting framework.

Inverted Pendulum experiment. The inverted pendulum, a highly nonlinear unstable system, is a commonly used benchmark for control methods. The objective of the control system is to balance the inverted pendulum by applying torque that will stabilize it in a vertically upright position. Here we follow the dynamics that was implemented in (Brockman et al., 2016). The LQR baseline solution is obtained from the linear approximation of the system dynamics, whereas our baseline and boosted controllers are not restricted to that approximation. We add correlated disturbances obtained from a Gaussian random walk, as above, such that $w_t \sim \mathcal{N}(w_{t-1}, 5e-3)$, where the noise values are then clipped to the range $[-0.5, 0.5]$.

7. Conclusions

We have described a framework for boosting of algorithms that have state information, and two efficient algorithms that provably enhance weak learnability in different ways. These can be applied to a host of control problems in dynamical systems. Preliminary experiments in simulated control look promising, of boosting for both linear and deep controllers.

References

Agarwal, N., Bullins, B., Hazan, E., Kakade, S. M., and Singh, K. Online control with adversarial disturbances. *arXiv preprint arXiv:1902.08721*, 2019.

Anava, O., Hazan, E., and Mannor, S. Online learning for adversaries with memory: price of past mistakes. In *Advances in Neural Information Processing Systems*, pp. 784–792, 2015.

Beygelzimer, A., Hazan, E., Kale, S., and Luo, H. Online gradient boosting. In *Advances in neural information processing systems*, pp. 2458–2466, 2015a.

Beygelzimer, A., Kale, S., and Luo, H. Optimal and adaptive algorithms for online boosting. In *International Conference on Machine Learning*, pp. 2323–2331, 2015b.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.

Cohen, A., Hasidim, A., Koren, T., Lazic, N., Mansour, Y., and Talwar, K. Online linear quadratic control. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1029–1038. PMLR, 2018.

Cohen, A., Koren, T., and Mansour, Y. Learning linear-quadratic regulators efficiently with only \sqrt{T} regret. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1300–1309, Long Beach, California, USA, 09–15 Jun 2019. PMLR.

Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, August 1997. ISSN 0022-0000. doi: 10.1006/jcss.1997.1504. URL <http://dx.doi.org/10.1006/jcss.1997.1504>.

Freund, Y., Iyer, R., Schapire, R. E., and Singer, Y. An efficient boosting algorithm for combining preferences. *J. Mach. Learn. Res.*, 4:933–969, December 2003. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=945365.964285>.

Hazan, E., Kakade, S. M., and Singh, K. The nonstochastic control problem. *arXiv preprint arXiv:1911.12178*, 2019.

Kanade, V. and Kalai, A. Potential-based agnostic boosting. In *Advances in neural information processing systems*, pp. 880–888, 2009.

Mason, L., Baxter, J., Bartlett, P. L., and Frean, M. R. Boosting algorithms as gradient descent. In *Advances in neural information processing systems*, pp. 512–518, 2000.

Schapire, R. E. and Freund, Y. *Boosting: Foundations and algorithms*. MIT press, 2012.

Simchowitz, M., Singh, K., and Hazan, E. Improper learning for non-stochastic control, 2020.