Learning with Continuous Experts using Drifting Games

Indraneel Mukherjee and Robert E. Schapire

Department of Computer Science, Princeton University, Princeton NJ 08540

Abstract

We consider the problem of learning to predict as well as the best in a group of experts making continuous predictions. We assume the learning algorithm has prior knowledge of the maximum number of mistakes of the best expert. We propose a new master strategy that achieves the best known performance for online learning with continuous experts in the mistake bounded model. Our ideas are based on drifting games, a generalization of boosting and online learning algorithms. We also prove new lower bounds based on the drifting games framework which, though not as tight as previous bounds, have simpler proofs and do not require an enormous number of experts.

1. Expert Learning Model

Our model is a repeated game with the following players: a *Master* against m *experts* and an *adversary*. In each round t, the following happen:

- The master chooses real weights $w_1^t, \ldots, w_m^t$ over the experts.
- Each expert i makes a prediction $x_i^t \in [-1, +1]$.
- The master predicts $\hat{y}^t \triangleq \text{sign}(\sum_i w_i^t x_i^t) \in \{-1, 0, +1\}$.
- The adversary then chooses a label $y^t \in \{-1, +1\}$, causing expert i to suffer loss $\frac{1}{2}|y^t - x_i^t|$, and the master to suffer loss $\mathbf{1}(y^t \neq \hat{y})$.

It is guaranteed that some expert will suffer less than k total loss. The goal of the master is to come up with a strategy to choose distributions $\mathbf{w}^t$ in each round, so as to minimize his loss against the worst possible adversary. Note that our model restricts the master to predict weighted linear combinations of the experts' advice. However, our lower bounds hold for *any* master algorithm, showing that the master in our model strategies are sufficiently powerful.

2. Previous and Related Work

With binary experts, outputting predictions in $\{-1, +1\}$, this problem was essentially solved entirely by Cesa-Bianchi et al. [2] who proposed the Binomial Weights (BW) algorithm. However, their work cannot be applied to our setting since here the experts are continuous, with predictions in $[-1, +1]$. In such a setting, other methods, notably exponential-weight algorithms [1, 4, 5], can be used instead. However, such algorithms do not enjoy the same level of tight optimality of the BW algorithm, and it has been an open problem since the introduction of BW as to whether this method can be generalized to continuous experts.

3. Drifting Games and the OS algorithm

Our algorithm is based on the *drifting games* framework introduced by Schapire [6]. This framework generalizes a number of online and boosting learning algorithms, including boost-by-majority [3], AdaBoost [4], the weighted majority algorithm [5] and Binomial Weights [2]. Schapire [6] suggests a general master strategy, OS, for playing any drifting game. In our model, given the number of rounds T, the OS algorithm is based on a set of potential functions $\phi_t : \mathbb{R} \to \mathbb{R}$ defined recursively as follows:

- $\phi_T(x) = \mathbf{1}(x \leq 2k - T)$.
- $\phi_{t-1}(x) = \min_{w \in \mathbb{R}} \max_{z \in [-1,+1]} (\phi_t(x+z) + wz)$

Denoting by s_i^t the loss suffered by expert i in the first $T - t$ rounds, the OS algorithm chooses w_i^t as follows

$$w_i^t \in \arg\min_{w \in \mathbb{R}} \max_{z \in [-1,+1]} (\phi_{t+1}(2s_i^t + t - T + z) + w).$$

Schapire [6] argues that the (general)OS algorithm is nearly optimal when the number of players is very large for general Drifting Games. For our model, the following result follows immediately from his work.

Theorem 1 (Drifting Games [6])**.** *The loss suffered by the OS algorithm is upper bounded by $\phi_0(0)$ where ϕ is defined as above. Further, whenever $\epsilon > T/\sqrt{m}$, any master algorithm playing for T rounds can be forced to make $\phi_0(0) - \epsilon$ mistakes.*

Intuition The potentials capture the minimax value of a game against one expert that is randomized. At any stage, the average of the potentials yields the minimax value of a game against a bunch of such experts that are distributed, i.e. predicting independently of each other. Such experts turn out to be slightly more powerful, as the upper-bound in Theorem 1 implies. Near-optimality follows from the fact that a large number of experts can, on average, simulate the individual randomness. The advantage of this approach is, by making the experts distributed, it becomes easier to analyse.

4. Our contributions

Our main contribution is explicitly computing the potentials.

Theorem 2 (Drifting Games for $[-1, +1]$ Experts)**.** *The values of the potential can be computed as follows:*

$$\phi_{T-t}(s + 2k - T) = \begin{cases} 1 & \text{if } s \leq 0 \\ 1 - 2^{-t} \sum_{i=0}^{s-1} \binom{t}{\lceil \frac{t+i}{2} \rceil} & \text{else.} \end{cases} \quad (1)$$

Further, the OS strategy for this game can be computed efficiently.

The recurrence defining ϕ_t is non-trivial to compute since there are infinitely many choices for any expert response z. We manage to show Theorem 2 nonetheless by using the next technical result, whose proof holds more generally and maybe of independent interest.

Theorem 3 (Piecewise Convexity)**.** *For every round t, ϕ_t is piecewise convex with pieces breaking at integers, i.e., for every integer n, ϕ_t is convex in $[n, n+1]$.*

4.1 Improved Lower Bounds

The dependence of ϵ on T, m in Theorem 1 leads to poor optimality results for our problem. This is because Schapire's [6] proof, based on sampling, holds for more general drifting games and has to rely on approximate concentration bounds. By carefully tailoring his proof to our model, we are able to derandomize his approach, leading to ϵ as small as T/m. Further, using knowledge about the potential functions, the dependence reduces to $\sqrt{k}/m$ in Theorem 1. Consequently, we are able to show that our strategy never makes $O(\log k)$ more mistakes than the optimal strategy when $m = \Omega(2^k)$.

5. Comparison with abstaining experts

Our analysis yields the same upper bound for continuous and abstaining experts(outputting predictions in $\{-1, 0, +1\}$). Even our lower bounds use only abstaining experts, and we believe that they can be made to match the upper bounds. This would imply the surprising result that continuous experts are no more powerful than abstaining experts for sufficiently many experts.

References

[1] Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *Journal of the Association for Computing Machinery*, 44(3):427–485, May 1997.

[2] Nicolò Cesa-Bianchi, Yoav Freund, David P. Helmbold, and Manfred K. Warmuth. On-line prediction and conversion strategies. *Machine Learning*, 25:71–110, 1996.

[3] Yoav Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121(2):256–285, 1995.

[4] Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, August 1997.

[5] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.

[6] Robert E. Schapire. Drifting games. *Machine Learning*, 43(3):265–291, June 2001.