Proceedings of Machine Learning Research vol :1–30, 2026

On Randomized Algorithms in Online Strategic Classification

Chase Hutton CHUTTON6@UMD.EDU
University of Maryland

Adam Melrod AMELROD@MATH.HARVARD.EDU
Harvard University

Han Shao HANSHAO@UMD.EDU
University of Maryland

Abstract

Online strategic classification studies settings in which agents strategically modify their features to obtain favorable predictions (Ahmadi et al., 2023). For example, given a classifier that determines loan approval based on credit scores, applicants may open or close credit cards and bank accounts to obtain a positive prediction. The learning goal is to achieve low mistake or regret bounds despite such strategic behavior.

While randomized algorithms have the potential to offer advantages to the learner in strategic settings, they have been largely underexplored. In the realizable setting, no lower bound is known for randomized algorithms, and existing lower bound constructions for deterministic learners can be circumvented by randomization. In the agnostic setting, the best known regret upper bound is $O(T^{3/4}\log^{1/4}(T|\mathcal{H}|))$ (Ahmadi et al., 2023), which is far from the standard online learning rate of $O(\sqrt{T\log|\mathcal{H}|})$.

In this work, we provide refined upper and lower bounds for online strategic classification in both the realizable and agnostic settings. In the realizable setting, using a new construction, we prove a lower bound that, for $T > \mathrm{Ldim}(\mathcal{H})\Delta^2$, extends the existing deterministic lower bound of $\Omega(\mathrm{Ldim}(\mathcal{H})\Delta)$ to all algorithms. This is the first lower bound that applies to randomized algorithms, resolving an open question of Ahmadi et al. (2023). We also give the first randomized algorithm that improves on the known deterministic upper bound of $O(\mathrm{Ldim}(\mathcal{H})\cdot\Delta\log\Delta)$, achieving $O(\sqrt{T\cdot\mathrm{Ldim}(\mathcal{H})\log\Delta})$ expected mistakes. This beats the deterministic bound for $T < \mathrm{Ldim}(\mathcal{H})\,\Delta^2\log\Delta$.

In the agnostic setting, we give an improper randomized algorithm with expected regret $O(\sqrt{T\log|\mathcal{H}|})$ against adaptive adversaries, improving upon the previous $O(T^{3/4}\log^{1/4}(T|\mathcal{H}|))$ bound and matching the standard online learning rate. We also prove that this optimal rate requires improper learning.

Keywords: strategic classification, online learning, Littlestone dimension.

© 2026 C. Hutton, A. Melrod & H. Shao.