

Lecture 4

This document provides practice problems that are similar to those that will be asked during the final exam. Please note that the document reflects the style and not the number of the questions that will be on the exam.

Problem 1

The Hedge algorithm is an application of a more general concept known as multiplicative weight updates. Which other machine learning technique is also based on this update rule?

- (a) Bootstrap aggregation.
- (b) Adaptive boosting.
- (c) Gaussian processes.

Problem 2

Consider the setup of sequential decision making with expert advice, where at each stage $k > 0$ experts provide their advice, and we need to decide which expert to follow. We further assume that there is at least one perfect expert that always predicts the correct answer. We use the halving algorithm, where at each round, we follow the advice suggested by the majority of all experts who up to now have not made any mistakes.

What is the best worst-case upper bound on the number of mistakes that we will make?

- (a) $\lceil \log_2(k) \rceil$
- (b) $\lceil \log_2(k) \rceil - 1$
- (c) $\lfloor \log_2(k) \rfloor$
- (d) $\lfloor \log_2(k) \rfloor - 1$

Problem 3

Which of the following statements about the Exp3 algorithm are correct?

- (a) Exp3 is used in settings where we have full information of the losses of all actions.
- (b) The regret bound of Exp3 is concave and sublinear in time T .
- (c) The average regret bound of Exp3 is convex in time T and converges to zero as $T \rightarrow \infty$.
- (d) Exp3 can be obtained by applying the Hedge algorithm to a modified loss function.
- (e) Exp3 creates an unbiased estimator of the entire loss vector at every iteration.

Problem 4 (optional)

Suppose two players repeatedly play against each other in a zero-sum game with actions $\{(x_t, y_t) \in \Delta_n \times \Delta_m\}_{t=1}^T$, where Δ_n denotes the n -dimensional simplex and Δ_m , the m -dimensional simplex. At each round, the first player plays best response, that is $x_t = \arg \max_{x \in \Delta_n} x^\top A y_t$, whereas the second player plays according to the Hedge algorithm (i.e., MWA) and suffers the loss $x_t^\top A y_t$ at every iteration. The matrix $A \in \mathbb{R}^{n \times m}$ has elements in $[0, 1]$. Which of the following statements are true?

- (a) The sequence $\{(x_t, y_t)\}_{t=1}^T$ converges to a Nash equilibrium as $T \rightarrow \infty$.
- (b) The averages $\frac{1}{T} \sum_{t=1}^T x_t$ and $\frac{1}{T} \sum_{t=1}^T y_t$ converges to a Nash equilibrium for $T \rightarrow \infty$.
- (c) None of the above.