Optimization and Uncertainty

Summer term 2021

Prof. Dr. Martin Hoefer
Tim Koglin, Lisa Wilhelmi

GOETHE UNIVERSITÄT FRANKFURT AM MAIN
Institute of Computer Science
Algorithms and Complexity

Assignment 10

Issued: 06.07.2021
Due: 13.07.2021, **10:00h**

Exercise 10.1 *Unknown time horizon* (4 points)

All no-regret algorithms presented in the lecture assumed that the time horizon T is known. Indeed, they also work for unknown time horizons by applying a slight modification: Restart the original no-regret algorithm repeatedly, where each restart marks a new phase. Phase Π_k, $k \geq 0$, consists of steps $2^k, ..., 2^{k+1} - 1$, i.e., 2^k steps in total. At the beginning of each phase, restart the original no-regret algorithm with $T' = 2^k$. This process ends when the unknown time horizon T is reached. In case the last phase is not completed, the remaining steps of the phase yield zero costs.

Assume that the original algorithm (that knows T) has a regret of at most $\alpha\sqrt{T}$ on any sequence of length T. Show that the modified algorithm has regret of at most $\frac{\sqrt{2}}{\sqrt{2}-1}\alpha\sqrt{T}$.

Remark: For RWM for Experts, *we would set* $\alpha = 2\sqrt{\ln n}$, *for* Exp3 $\alpha = 3\sqrt{n \ln n}$. *In both cases, we lose only a constant factor in the regret.*

Exercise 10.2 *Follow the Leader* (3 points)

Construct an instance of Online Convex Optimization such that the regret of Follow-the-leader (FTL) is linear in T after $T > 0$ rounds. Proof correctness of your answer.

Exercise 10.3 *FTRL with Entropical regularization* (5 points)

Recall that in the Online Convex Optimization problem experts are vectors $\boldsymbol{v}$ in a convex and compact set $D \subseteq \mathbb{R}^d$. In what follows, consider $D = \{\boldsymbol{v} \in \mathbb{R}^d \mid v_i \geq 0 \text{ for all } i,\ \sum_{i=1}^d v_i = 1\}$. The Follow-the-regularized-leader (FTRL) algorithm picks in each round t a vector $\boldsymbol{w}^{(t)} \in D$ that minimizes the cost $R(\boldsymbol{w}^{(t)}) + \sum_{t'=1}^{t-1} c_{t'}(\boldsymbol{w}^{(t')})$. Here, consider linear cost functions, i.e., $c_{t'}(\boldsymbol{w}^{(t')}) = \sum_{i=1}^d \ell_i^{(t')} w_i^{(t')}$ where $\ell_i^{(t')}$ is the cost of expert i in round t'. For non-negative vectors $\boldsymbol{v} \in D$, Entropical regularization is defined by $R(\boldsymbol{v}) = \frac{1}{\eta} \sum_{i=1}^d v_i \ln v_i$, where $\eta > 0$ is a scaling factor.

Show that FTRL over D with Entropical regularization leads to choices of $\boldsymbol{w}^{(t)}$ exactly such that element $w_k^{(t)}$ is proportional to $\exp\left(-\eta \sum_{t'=1}^{t-1} \ell_k^{(t')}\right)$ for all $0 < k \leq d$.

Hint: Use the method of Lagrange multipliers which, in this case, works as follows: For $\boldsymbol{v} \in D$ *to be a local optimum of some objective function* $F : D \to \mathbb{R}$ *subject to some condition* $g : D \to \mathbb{R}$ *(here,* $g := \sum_{i=1}^d v_i = 1$*), there must exists a Lagrange multiplier* $\lambda \in \mathbb{R}$ *such that* $\boldsymbol{v}$ *is a stationary point of* $F(\boldsymbol{v}) + \lambda g(\boldsymbol{v})$ *with respect to* v_i*, for all* i*.*

Remark: By only changing the parametrization in the RWM algorithm for Experts*, one can show that it actually uses exactly the same update rule.*

Exercise 10.4 *FTL for web search* (3 points)

Consider an algorithm Γ that faces the following problem motivated by web search: Suppose there are n users that all search for the same keyword. There are k different results that they might be interested in. Whenever a user arrives, Γ displays these k results in a chosen order. Afterwards, it is revealed to Γ which of the k results the user was interested in and Γ incurs a cost of j if this result was the j^{th} result in the displayed order.

Model this problem as an Online Convex Optimization problem so that Γ can be realized by a Follow-the-leader (FTL) algorithm.

Exercise 10.5 *σ-strong convexity* (3 points)

Prove the following statement: If regularizer R is σ-strongly convex and cost functions $c_1, c_2, \ldots$ are convex, then $R + \sum_t c_t$ is σ-strongly convex.

The assignments and further information on the course are provided on our website:
http://algo.cs.uni-frankfurt.de/lehre/oau/sommer21/oau21.shtml

Contacts for submissions and questions: {`koglin,wilhelmi`}`@em.uni-frankfurt.de`.