
CMS Blair Spearman Doctoral Prize
Prix de doctorat Blair-Spearman de la SMC
(Org: **Christopher Liaw** (University of Toronto))

CHRISTOPHER LIAW, Google

Optimal anytime regret with two experts

A fundamental problem in online learning is the problem of prediction with expert advice. The problem can be cast as a game with an algorithm, an adversary, and n experts, where at each time step t , the algorithm chooses a probability vector p_t over the n experts, the adversary chooses a cost vector c_t over the n experts, and the algorithm incurs a cost equal to the inner product between c_t and p_t . The goal is to minimize the regret: the algorithm's total cost minus the cost of the single best expert in hindsight. If the time horizon for the game is known in advance, optimal algorithms are known when there are exactly two, three, or four experts and asymptotically if the number of experts is large. In the anytime setting (where the game is played indefinitely), no optimal algorithm was previously known.

In this talk, I will talk about the first minimax optimal algorithm for minimizing regret in the anytime setting. This talk will consider the case of two experts, and prove that the optimal regret is $\gamma\sqrt{t}/2$ at all time steps t , where $\gamma \approx 1.30693$ is a natural constant that arose 35 years ago in studying fundamental properties of Brownian motion. The algorithm is designed by considering a continuous analogue of the regret problem, which is solved using ideas from stochastic calculus.

This talk is based on joint work with Nick Harvey, Ed Perkins, and Sikander Randhawa.