

1163-60-1688

Philippe Casgrain* and **Anastasis Kratsios**. *Optimizing Optimizers: Generating regret-optimal gradient descent algorithms.*

The need for fast and robust optimization algorithms are of critical importance in all areas of machine learning. This paper treats the task of designing optimization algorithms as a stochastic control problem. Using the expected future regret as a metric for an algorithm's performance, we derive the necessary and sufficient dynamics that regret-optimal algorithms must satisfy as a discrete-time FBSDE. We study the existence and uniqueness of regret-optimal algorithms and derive bounds on rates of convergence to solutions of convex optimization problems. Though closed-form optimal dynamics cannot be obtained in general, we present fast numerical methods for approximating them, generating optimization algorithms which directly optimize their long-term regret. Lastly, these are benchmarked against various other commonly used optimization algorithms to demonstrate their effectiveness. (Received September 18, 2020)