Les mathématiques de l'apprentissage machine
« Atelier sur les défis modernes en théorie d'apprentissage »
23 au 26 avril 2018

Mathematics of Machine Learning
"Workshop on Modern Challenges of Learning Theory"
April 23-26, 2018

A distributional perspective on reinforcement learning

Marc G. Bellemare *

marc.g.bellemare@gmail.com

This talk will present an overview of our recent research on distributional reinforcement learning. Our starting point is our recent ICML paper, in which we argued for the fundamental importance of the value distribution: the distribution of random returns received by a reinforcement learning agent. This is in contrast to the common approach, which models the expectation of this return, or value. Back then, we were able to design a new algorithm that learns the value distribution through a TD-like bootstrap process and achieved state-of-the-art performance on games from the Arcade Learning Environment (ALE). However, this left open the question as to why the distributional approach should perform better at all. We've since delved deeper into what makes distributional RL work: first by improving the original using quantile regression, which directly minimizes the Wasserstein metric; and second by unearthing surprising connections between the original C51 algorithm and the distant cousin of the Wasserstein metric, the Cramer distance.

*Google Brain Montréal, 1253, Avenue McGill College, Montréal, QC H3B 2Y5, CANADA