CRUNCH Seminars at Brown, Division of Applied Mathematics

Friday – June 8, 2018

GANGs: Generative adversarial network games
(Oliehoek et al., arXiv:1712.00679)
Yang Liu

Generative Adversarial Networks (GAN) have become one of the most successful frameworks for unsupervised generative modeling. As GANs are difficult to train much research has focused on this. However, very little of this research has directly exploited game- theoretic techniques. We introduce Generative Adversarial Network Games (GANGs), which explicitly model a finite zero-sum game between a generator (G) and classifier (C) that use mixed strategies. The size of these games precludes exact solution methods, therefore we define resource-bounded best responses (RBBRs), and a resource- bounded Nash Equilibrium (RB-NE) as a pair of mixed strategies such that neither G or C can find a better RBBR. The RB-NE solution concept is richer than the notion of 'local Nash equilibria' in that it captures not only failures of escaping local optima of gradient descent, but applies to any approximate best response computations, including methods with random restarts. To validate our approach, we solve GANGs with the Parallel Nash Memory algorithm, which provably monotonically converges to an RB-NE. We compare our results to standard GAN setups, and demonstrate that our method deals well with typical GAN problems such as mode collapse, partial mode coverage and forgetting.