

Kernel distances for distinguishing and sampling from probability distributions

Dougal Sutherland
University College London

3/12/2019 at 10:30 am
CoRE A 301

Abstract

Probability distributions are the core object of statistical machine learning, and one of the basic properties we can consider is distances between them. In this talk, we will consider using these distances for two important tasks, and show how to design distances which will be useful for each. First, we study the problem of two-sample testing, where we wish to determine whether (and how) two different datasets meaningfully differ. We then study this framework in the setting of training generative models, such as generative adversarial networks (GANs), which learn to sample from complex distributions such as those of natural images.

The distances used are defined in terms of kernels, but we parameterise these kernels as deep networks for flexibility. This combination gives both theoretical and practical benefits over staying purely in either framework, and we obtain state-of-the-art results for unsupervised image generation on CelebA and ImageNet with our novel Scaled MMD GAN.

Bio

Dougal Sutherland is a postdoctoral researcher at the Gatsby Computational Neuroscience Unit, University College London, working with Arthur Gretton. He received his PhD in 2016 from Carnegie Mellon University, advised by Jeff Schneider. His research focuses on problems of learning about distributions from samples, including training implicit generative models, density estimation, two-sample testing, and distribution regression. His work combines kernel frameworks with deep learning, and aims for theoretical grounding of practical results.