

Training Generative Adversarial Networks with Kernel Methods

ARTHUR GRETTON

Gatsby Computational Neuroscience Unit, University College London, UK

Email: `arthur.gretton@gmail.com`

Generative adversarial networks (GANs) use neural networks as generative models, creating realistic samples that mimic real-life reference samples (for instance, images of faces, bedrooms, and more). These networks require an adaptive critic function while training, to teach the networks how to move improve their samples to better match the reference data. I will describe a kernel divergence measure, the maximum mean discrepancy, which represents one such critic function. With gradient regularisation, the MMD is used to obtain strong performance on challenging image generation tasks, including 160×160 CelebA and 64×64 ImageNet. In addition to adversarial network training, I'll discuss the challenge of benchmarking GAN performance.