

Generative Adversarial Networks (GANs, Goodfellow et al, 2014) have had a profound impact on the generative modelling community. Both their potential to synthesize compellingly realistic images and the challenges in training these models have been thoroughly studied. Despite this, our understanding of the potential limitations and promise the adversarial training framework is just beginning. In this talk I will briefly introduce Generative Adversarial Networks and highlight some recent advances in our understanding of how they work. I will then review some of my recent work in pushing forward the frontier of applications of the adversarial training framework, including our efforts to incorporate inference into the GAN framework; to learn many-to-many transformations between unpaired domains; as well as our recent framing of mutual information estimation as an adversarial training problem.