

Lectures 1&2

Generative convolutional networks: a short survey

Convolutional networks (ConvNets) are most commonly used for image recognition. In recent years, however, there is a growing interest in generative ConvNets, and in particular in ConvNets that are trained to generate or to transform images. As it turns out, ConvNets are as good at generating images, as they are at recognizing them. The lecture will overview the mechanisms behind such generative ConvNets as well as some important details specific to them. In particular, the lecture will cover the so-called perceptual loss functions and the adversarial learning approach, which both enhance the realism of ConvNet outputs greatly. I will also discuss some innate priors inside generative ConvNet architectures that make them particularly well suited for image generation.