



THE WEIZMANN INSTITUTE OF SCIENCE
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Vision and Robotics Seminar

Room 1 ,Ziskind Building
on Thursday, Jun 13, 2019
at 12:15

Eitan Richardson
The Hebrew University of Jerusalem

On GANs and GMMs

Abstract:

A longstanding problem in machine learning is to find unsupervised methods that can learn the statistical structure of high dimensional signals. In recent years, GANs have gained much attention as a possible solution to the problem, and in particular have shown the ability to generate remarkably realistic high resolution sampled images. At the same time, many authors have pointed out that GANs may fail to model the full distribution ("mode collapse") and that using the learned models for anything other than generating samples may be very difficult. In this paper, we examine the utility of GANs in learning statistical models of images by comparing them to perhaps the simplest statistical model, the Gaussian Mixture Model. First, we present a simple method to evaluate generative models based on relative proportions of samples that fall into predetermined bins. Unlike previous automatic methods for evaluating models, our method does not rely on an additional neural network nor does it require approximating intractable computations. Second, we compare the performance of GANs to GMMs trained on the same datasets. While GMMs have previously been shown to be successful in modeling small patches of images, we show how to train them on full sized images despite the high dimensionality. Our results show that GMMs can generate realistic samples (although less sharp than those of GANs) but also capture the full distribution, which GANs fail to do. Furthermore, GMMs allow efficient inference and explicit representation of the underlying statistical structure. Finally, we discuss how GMMs can be used to generate sharp images.