Machine Learning (CSE 446): Generative Adversarial Networks (GANs)

Sham M Kakade

University of Washington
cse446-staff@cs.washington.edu

Announcements

- Weds is the final.
- One page of notes.
- List of topics posted tomorrow.

Class results with random Fourier features, HW3, Q7

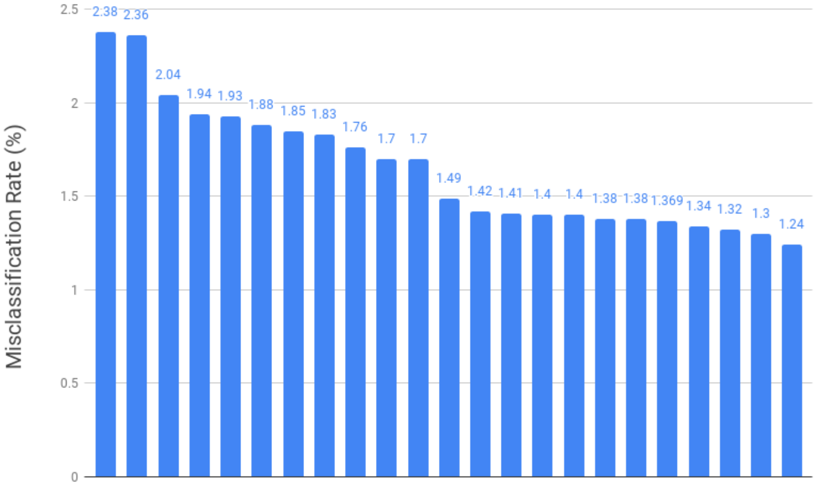

Image Generation

These are computer generated images from the “bigGAN”.

Figure 1: Class-conditional samples generated by our model.

Classification

- SPAM detection.
- x is an email.
- Suppose $Q(x)$ is the distribution over true emails.
- Suppose $G(x)$ is spammer's generative distribution over emails.
- usual supervised learning: make a dataset of (x, y) pairs, say with 50% labeled true and 50% spam.
- suppose $\Pr_\theta(Y = 1|X) = D_\theta(X)$ is our model's probably of that $Y = 1$, true email.

Building Our Classifier

- Maximum likelihood is:
$$\max_{\theta} \mathbb{E} \log \Pr_{\theta}(Y = 1|X)$$
- every θ corresponds to some model D.
- If our data is equally split, where all the $Y = 1$'s come from Q and and all the $Y = 0$'s come from G, then:
$$\max_{\theta} \mathbb{E} \log \Pr_{\theta}(Y = 1|X) =$$
$$\max_{D} \left(\mathbb{E}_{x \sim Q}[\frac{1}{2} \log D(x)] + \frac{1}{2} \mathbb{E}_{x \sim G}[\log(1 - D(x))] \right)$$
- Likelihood function:
$L(D, G) =$

The Spammer's Job

- If the spam detector D was fixed, then our spammer wants to:

$$\min_G L(G, D)$$

- The spammer wants to make the discriminator D's likelihood small.

We have a game:

- Let's think more abstractly: D is procedure to spot fake images, G is a procedure to generate images.
- The game can be viewed as:

$$\min_G \max_D L(G, D) =$$
$$\min_G \max_D \left(\mathbb{E}_{x\sim Q}[\frac{1}{2} \log D(x)] + \frac{1}{2}\mathbb{E}_{x\sim G}[\log(1 - D(x))] \right)$$

- Is this a powerful idea?
  Remember AD:

the Discriminator

- just *binary* image classification "fake or not"
- we know how to do supervised learning
- for a given G and samples from the truth Q, we can learn D.

the Generator

- what is a model for generating images?
  and how do we update G the model?
- let say $z \sim N(0, \mathbb{I})$ where I is a $d \times d$ matrix.
- a network:

the Generator network

Learning

- Fix G, update D.
  easy: just supervised learning.
- Now fix D. Recall:

$$\begin{aligned} L(G,D) &= \left(\mathbb{E}_{x\sim Q}[\frac{1}{2}\log D(x)] + \frac{1}{2}\mathbb{E}_{x\sim G}[\log(1-D(x))]\right) \\ &= \left(\mathbb{E}_{x\sim Q}[\frac{1}{2}\log D(x)] + \frac{1}{2}\mathbb{E}_{z\sim \mathrm{Normal}(0,\mathbb{I})}[\log(1-D(G(z)))]\right) \end{aligned}$$

- We can estimate this loss easily. Why?
- Update G: how?

Convergence/Comments

- Does it converge?
  Subtle: we are not "hill climbing" on an on an objective.
- at best, we can get to an equilibrium.
- Comparison to EM and likelihood based approaches:
  - Computing the the probability of x under G is difficult.
    difficulty for EM.
  - "mode collapse" with GANs
    sampling distribution not reflective of the truth.
- NLP generative methods do not use GANs!
  (better results with direct training approaches)

Whichfaceisreal.com

We can get creative

Style transfer. From the "cycleGAN"

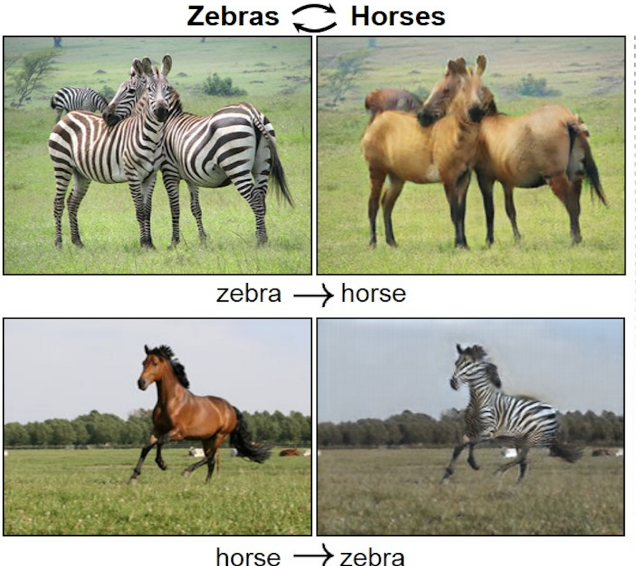

Thank you!!

Thank you for the hard work!

- Good luck on the final and have a great spring break.
- You have a good toolkit.
  Please participate in the larger ML community!