Invertible Normalizing Flows

ECE57000: Artificial Intelligence, Fall 2019

David I. Inouye

Announcements

- Quiz moved to Friday
  - Same content (i.e., up to DCGANs, not today)

GAN Limitation:
Cannot compute density values

- Evaluation of GANs is challenging
  - (Explicit density models could use test log likelihood)
  - "I think this looks better than that"

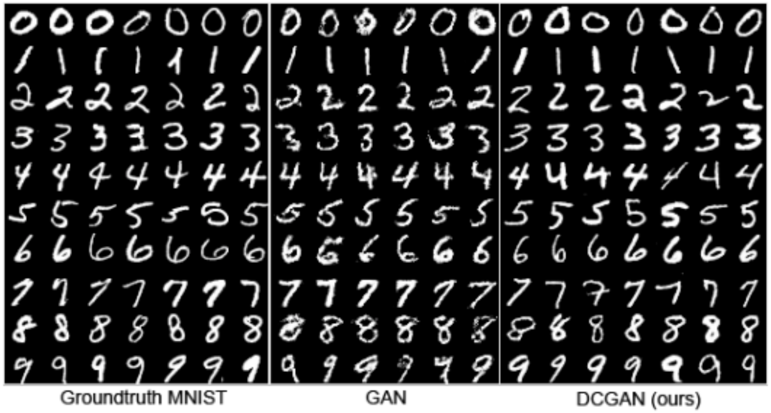

  - Inception scores
    - Train separate image classifier
    - See if passing fakes to classifier produces a high confidence prediction

- Cannot use for classification or outlier detection

GAN Limitation: Challenging to train because of careful balance between discriminator and generator

1. Assumptions on possible D and G
   1. Theory – All possible D and G
   2. Reality – Only functions defined by a neural network
2. Assumptions on optimality
   1. Theory – Both optimizations are solved perfectly
   2. Reality – The inner maximization is only solved approximately, and this interacts with outer minimization
3. Assumption on expectations
   1. Theory – Expectations over true distribution
   2. Reality – Empirical expectations over finite sample; for images, much of the high-dimensional space does not have samples

- **GANs can be very difficult/finicky to train**

GAN Limitation: Cannot go from observed to latent space, i.e. $x \to z$ not possible/easy

- Cannot manipulate an observed image in latent space
  - Cannot do the following, $x \to z,\ z' = z + 3,\ z' \to x'$
  - Rather, must start from fake image based on random z

All fake images->

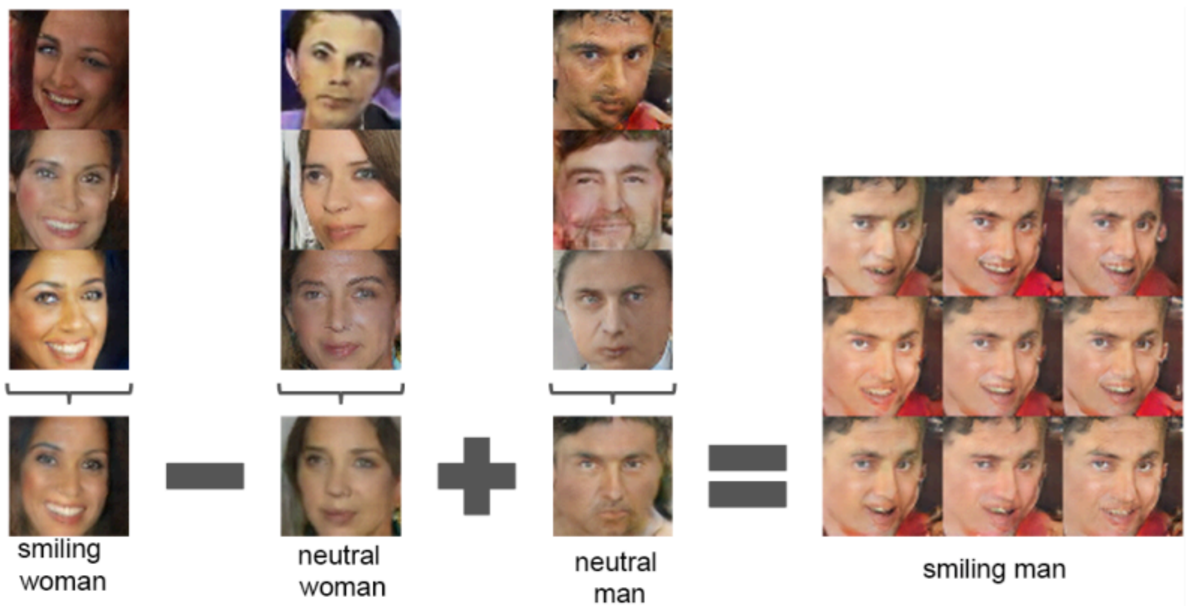

<u>Normalizing flows</u> use **invertible** deep models for the generator which allow more capabilities

- Transforming between observed/input and latent space is easy
  - $x = G\,(z)$
  - $z = G^{-1}(x)$

- Simple sampling like GANs
  - $z \sim \text{SimpleDistribution}$
  - $x = G(z) \sim \hat{p}_g(x)$, which is estimated distribution

- **Exact density** is computable via change of variables
  - Standard maximum likelihood estimation can be used for training

Highly realistic random samples from powerful flow model (GLOW)

Figure 1: Synthetic celebrities sampled from our model; see Section 3 for architecture and method, and Section 5 for more results.

https://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf

Interpolation between **real images** using GLOW

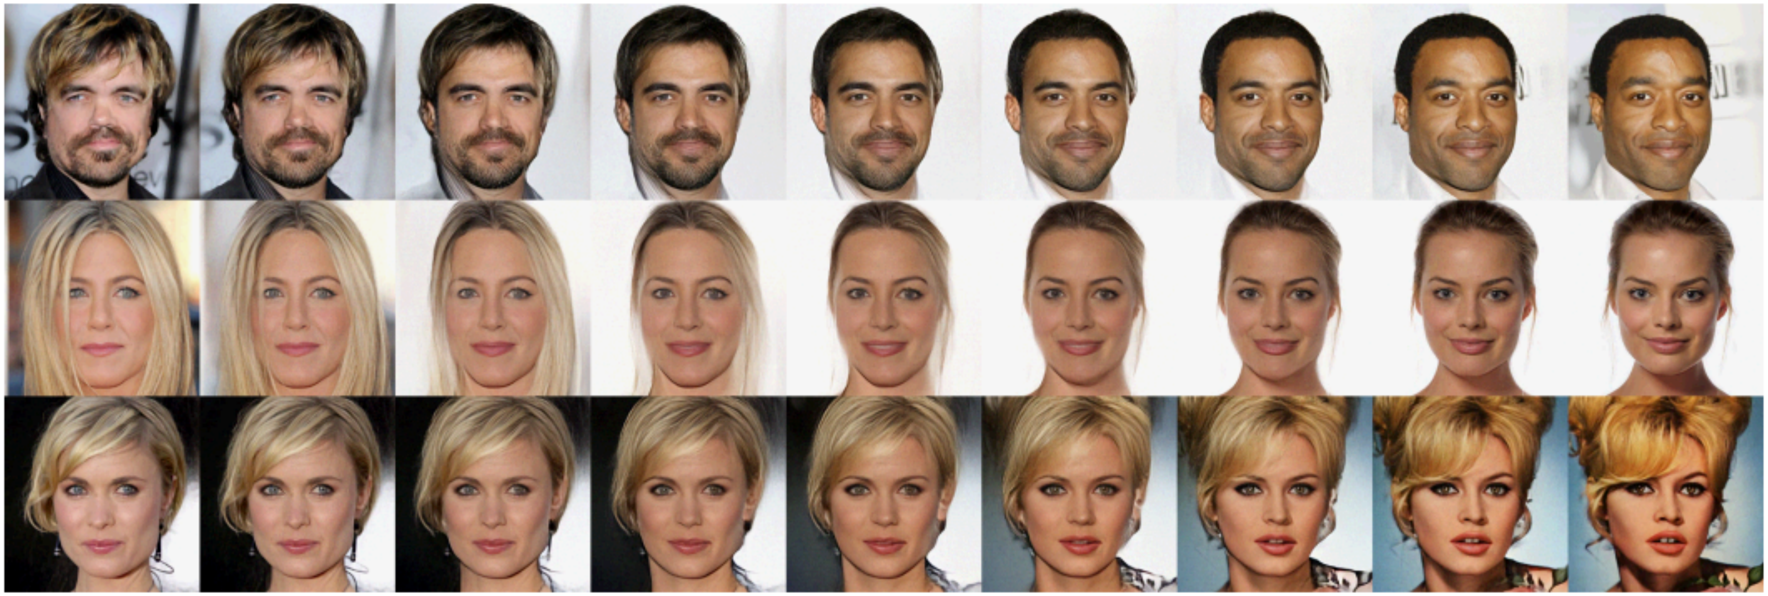

Figure 5: Linear interpolation in latent space between real images.

https://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf

Transformations of real image along various features

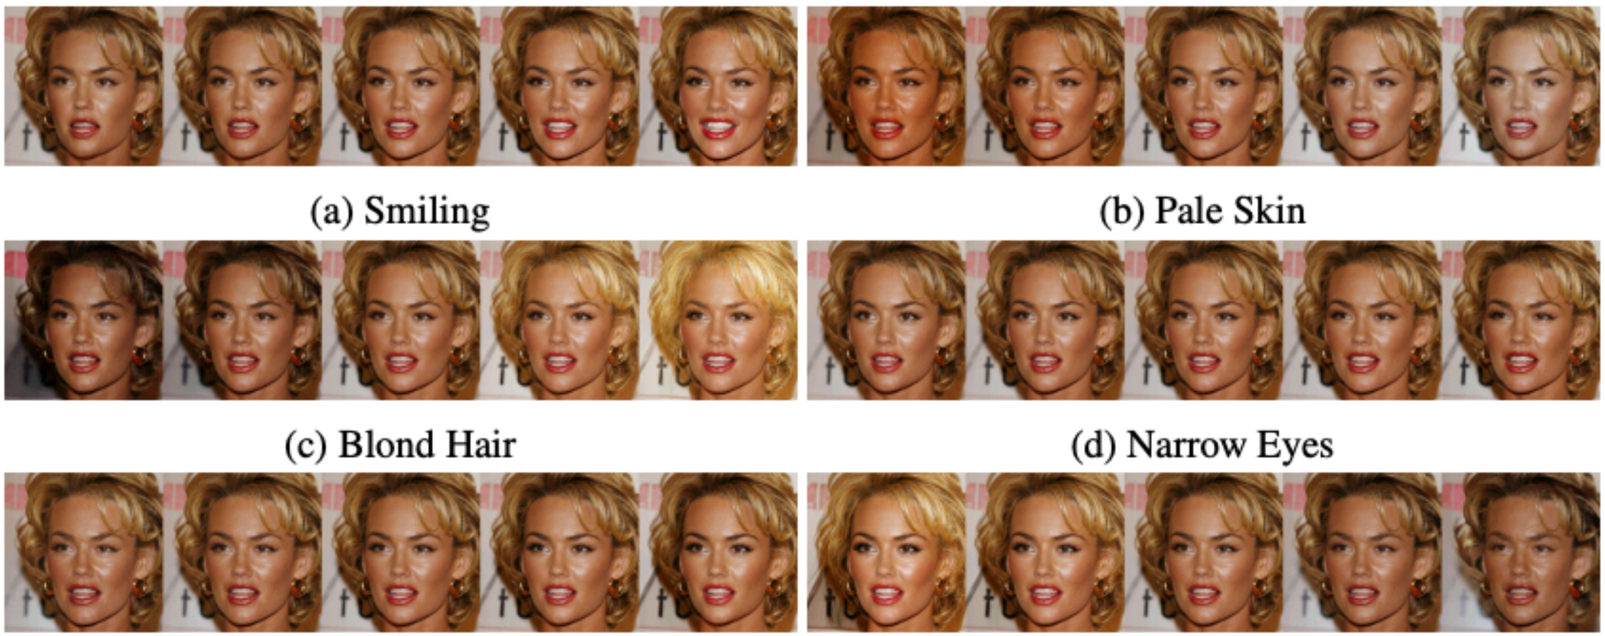

(a) Smiling
(b) Pale Skin
(c) Blond Hair
(d) Narrow Eyes
(e) Young
(f) Male

Figure 6: Manipulation of attributes of a face. Each row is made by interpolating the latent code of an image along a vector corresponding to the attribute, with the middle image being the original image

https://papers.nips.cc/paper/8224-glow-generative-flow-with-invertible-1x1-convolutions.pdf

Back to maximum likelihood estimation (MLE): How can we compute the likelihood for normalizing flows?

- Suppose
  - $z \sim \text{Uniform}([0,1]), \text{i.e.}, p_z(z) = 1$
  (latent space is uniform)
  - $G(\mathrm{z}) = 2\mathrm{z}$
  - Thus, $x = G(z) = 2z$.
- What is the density function of x, what is $p_x(x)$?