Neural Networks and Evolutionary Strategies

Generative Adversarial Networks

Ing. Ivan Gruber

Department of Cybernetics
Faculty of Applied Sciences
University of West Bohemia

DEPARTMENT OF CYBERNETICS

Content

DEPARTMENT OF CYBERNETICS

Supervised vs Unsupervised learning

Supervised

- Data: (x, y)
- Goal: Learn a function to map x→y
- Examples: Classification, regression, object detection, semantic segmentation

Unsupervised

- Data: x
- Goal: Learn some underlying hidden structure of the data
- Examples: Clustering, dimensionality-reduction, feature learning, density estimation, etc

Discriminative vs Generative models

Discriminative

- Model differences between classes
- Decision boundaries between classes
- Learn conditional probability $p(x|y)$
- Examples: Logical Regression, SVM, kNN, traditional NN

Generative

- Model characteristics of each class
- Distribution of each class
- Learn joint probability $p(x, y)$
- Can generate unseen content!
- Examples: Naive Bayes, Markov random fields, GANs

Generative Adversarial Networks

- Introduced by Goodfellow et al.[1]
- Can be utilized in unsupervised learning tasks
- Two main parts: generator G, and discriminator D

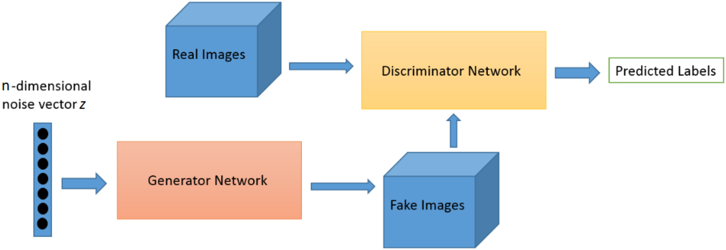

[1] Goodfellow, Ian, et al. "Generative adversarial nets."Advances in neural information processing systems. 2014.

Generator vs Discriminator

Generator

- Input: n-dimensional vector z
- Output: Fake image x_f
- Goal: To produce as realistic output as possible

Discriminator

- Input: Real image x_r or Fake image x_f
- Output: Predict label
- Goal: Distinguish between real and fake images

Goal: Nash equilibrium = The discriminator predicts "real"or "fake"with probability 0.5 for any sample.

Adversarial Loss

The expected value function of the discriminator:

$$V(G, D) = \frac{1}{2} E_{x \sim p_r}[\log D(x)] + \frac{1}{2} E_{z \sim p_z}[\log(1 - D(G(z)))], \quad (1)$$

$$\min_G (\max_D E(G, D)). \quad (2)$$

The best possible discriminator is the one which maximizes:

$$E_{x \sim p_r}[\log D(x)] + E_{z \sim p_z}[\log(1 - D(G(z)))]. \quad (3)$$

The best possible generator is the one which minimizes:

$$E_{z \sim p_z}[\log(1 - D(G(z)))]. \quad (4)$$

DEPARTMENT OF CYBERNETICS

Training Problems - Lack of Convergence

- Stems from an unbalance speed of the training
- The generator is trained faster:
  - The generator becomes superior to the discriminator
  - The generator produces perfect images (from discriminator's point of view)
  - The discriminator is unable to reveal fakes
- The discriminator is trained faster:
  - The discriminator becomes superior to the generator
  - The discriminator flawlessly reveals all fakes
  - The generator does not know what to improve
- Prevention: Heuristic strategies

DEPARTMENT OF CYBERNETICS

Training Problems - Mode Collapse

- No lever to force the generator to generate different outputs
- The generator generates only a few different outputs perfectly and omits the rest

- Prevention: Wasserstein distance, Conditional GAN

Wasserstein Distance

- Replaces standard Adversarial loss
- Minimum cost of transporting mass
- Also distance between two different distributions:

$$W(p_r, p_g) = \inf_{\gamma \in \Pi(p_r, p_g)} E_{(x,y)\sim\gamma}[||x - y||] \tag{5}$$

- Critic replaces discriminator - learn w to find optimal f_w
- Wasserstein loss:

$$L(p_r, p_g) = W(p_r, p_g) = \max_{w \in \mathbf{W}} E_{x\sim p_r}[f_w(x)] - E_{z\sim p_z}[f_w(G(z))] \tag{6}$$

DEPARTMENT OF CYBERNETICS

Applications and examples

- GAN ZOO
- GAN in Keras, Local Example
- StyleGAN - This person does not exist
- StyleGAN - This rental does not exist
- StyleGAN - These cats do not exist
- StyleGAN - This car does not exist
- StyleGAN - This waifu does not exist
- Which person is real?

Autoencoders

- Unsupervised learning - minimization of reconstruction loss
- Feed-forward network using "bottleneck"structure
- Two main parts: Encoder E, and Decoder D

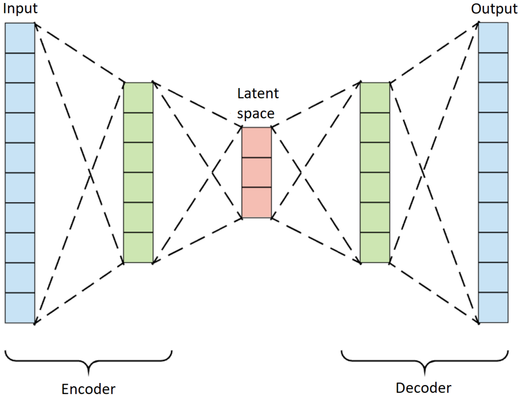

Encoder, Decoder, and Latent space

Encoder

- Input: Data x
- Output: latent code z
- Goal: Compress data into a feature vector representation while maintaining important information

Decoder

- Input:latent code z
- Output: Decoded data $\hat{x}$
- Goal: Generate an output map (with the same size as the original input) via upsampling procedure

Latent space

- No restriction applied to the latent space
- Over-training leads to dictionarization

DEPARTMENT OF CYBERNETICS

Variational Autoencoder

- Incorporates regularization by explicitly learning a joint distribution over data via forcing the latent space to follow a Gaussian distribution
- Regularization is added to the loss function
- Encourages the decoder to learn reconstruct data, while enforce the encoder to follow a Gaussian distribution
- Pros: Addition of probability allows unseen data generation
- Cons: Blurrier samples, harder to train

Latent-space arithmetic

- Good- trained encoder naturally holds big clustering ability across image attributes despite the lack of any additional information about them
- Occurs despite the unsupervised manner of the training and any additional constraints on the latent space
- Phenomenon occurs with VAE, GAN, VAEGAN, etc.
- Encoding and Decoding is highly non-linear process, however, some sort of linearity is preserved

DEPARTMENT OF CYBERNETICS

Latent-space arithmetic - Example 1

- Local example

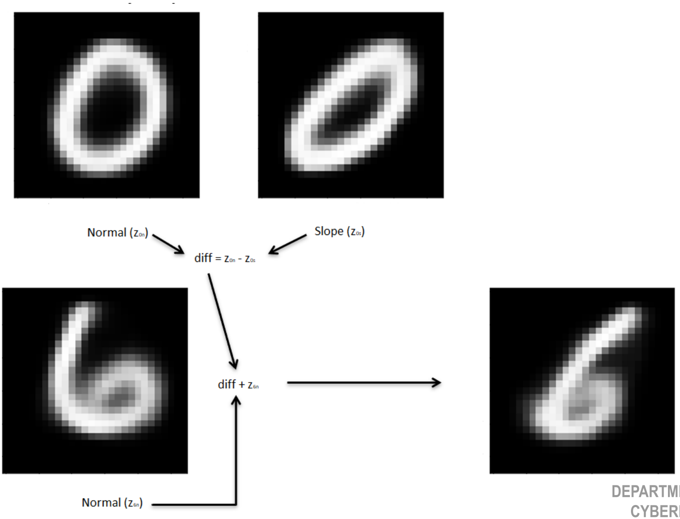

Latent-space arithmetic - Example 2

AE-GAN

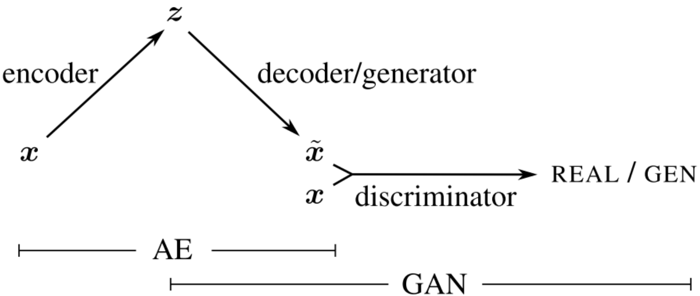
z
encoder
decoder/generator
x
x̃
x
discriminator
REAL / GEN
AE
GAN

Application and examples

- VAE - document background generation
- VAE - Logo detection
- U-Net - Semantic segmentation of historical documents
- TL-GAN - Latent-space arithmetic

Conditional GAN

- Additional condition on generated images
- Labels act as an extension to the latent space z

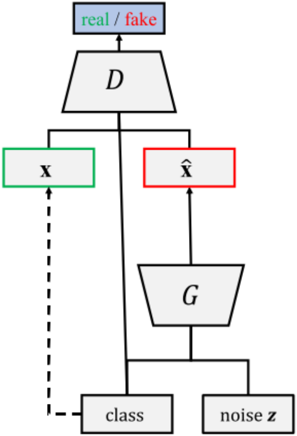

DEPARTMENT OF CYBERNETICS

Loss Function

Standard Adversarial loss:

$$V(G, D) = \frac{1}{2} E_{x \sim p_r}[\log D(x)] + \frac{1}{2} E_{z \sim p_z}[\log(1 - D(G(z)))]. \quad (7)$$

cGAN loss:

$$V(G, D) = \frac{1}{2} E_{x \sim p_r}[\log D(x|y)] + \frac{1}{2} E_{z \sim p_z}[\log(1 - D(G(z|y)))]. \quad (8)$$

DEPARTMENT OF CYBERNETICS

Application and examples

- Anime Girl generation
- Image-to-sketch translation
- New-environment generation
- Face aging
- New-pose generation
- Image inpainting

Thank you for your attention!

Questions?

DEPARTMENT OF
CYBERNETICS