

THE UNIVERSITY OF BRITISH COLUMBIA

Topics in AI (CPSC 532L): **Multimodal Learning with Vision, Language and Sound**

Lecture 12: Generative Models Cont. (GANs)



Course Logistics

- Assignment 4 will be out tomorrow (Friday) and is due in a week

- Reminder: **Project presentations** on Thursday
 - Logistics: form is up ____
 - Send me slides to minimize laptop switching on the day

Last week ...

PixelCNNs define tractable density function, optimize likelihood of training data: $p(x) = \prod$ i=1

VAEs define intractable density function with latent variables z (that we need to marginalize):

$$p_{\theta}(x) = \int$$

cannot optimize directly, derive and optimize lower bound of likelihood instead

$$\left[p(x_i | x_1, ..., x_{i-1}) \right]$$

$\int p_{\theta}(z) p_{\theta}(x|z) dz$





PixelCNNs define tractable density function, optimize likelihood of training data: $p(x) = \prod$ i=1

VAEs define intractable density function with latent variables z (that we need to marginalize):

$$p_{\theta}(x) = \int p_{\theta}(z) p_{\theta}(x|z) dz$$

cannot optimize directly, derive and optimize lower bound of likelihood instead

What if we give up on explicitly modeling density, and just want to sample?

$$\left[p(x_i | x_1, ..., x_{i-1}) \right]$$







PixelCNNs define tractable density function, optimize likelihood of training data: $p(x) = \prod$ i=1

VAEs define intractable density function with latent variables z (that we need to marginalize):

$$p_{\theta}(x) = \int p_{\theta}(z) p_{\theta}(x|z) dz$$

cannot optimize directly, derive and optimize lower bound of likelihood instead

What if we give up on explicitly modeling density, and just want to sample?

GANs: don't work with any explicit density function

$$\left[p(x_i | x_1, ..., x_{i-1}) \right]$$





Generative Adversarial Networks (GANs)

Problem: Want to sample from complex, high-dimensional training distribution. There is no direct way to do this!

[Goodfellow et al., 2014]



Problem: Want to sample from complex, high-dimensional training distribution. There is no direct way to do this!

Solution: Sample from a simple distributions, e.g., random noise. Learn transformation to the training distribution

Goodfellow et al., 2014



Problem: Want to sample from complex, high-dimensional training distribution. There is no direct way to do this!

Solution: Sample from a simple distributions, e.g., random noise. Learn transformation to the training distribution

Question: What can we use to represent complex transformation function?

Goodfellow et al., 2014]



Problem: Want to sample from complex, high-dimensional training distribution. There is no direct way to do this!

Solution: Sample from a simple distributions, e.g., random noise. Learn transformation to the training distribution

Question: What can we use to represent complex transformation function?

Goodfellow et al., 2014]

Output: Sample from training distribution

Input: Random noise





Generator network: try to fool the discriminator by generating real-looking images **Discriminator** network: try to distinguish between real and fake images

[Goodfellow et al., 2014]





Generator network: try to fool the discriminator by generating real-looking images **Discriminator** network: try to distinguish between real and fake images



Goodfellow et al., 2014]

Real Images (from training set)





Generator network: try to fool the discriminator by generating real-looking images **Discriminator** network: try to distinguish between real and fake images

Train jointly in minimax game Minimax objective function:

$$\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log \underline{D_{\theta_d}(x)} + \mathbb{E}_{z \sim p(z)} \log(1 - \underline{D_{\theta_d}(G_{\theta_g}(z))}) \right]$$

$$\text{Discriminator output} \qquad \text{Discriminator output} \qquad \text{Discriminator output for real data x} \qquad \text{dependent of the data G}$$

- **Discriminator** (θ_d) wants to maximize objective such that D(x) is close to 1 (real) and D(G(z)) is close to 0 (fake)
- into thinking generated G(z) is real)

Goodfellow et al., 2014

Discriminator outputs likelihood in (0,1) of real image

or generated fake data G(z)

- **Generator** (θ_{α}) wants to minimize objective such that D(G(z)) is close to 1 (discriminator is fooled





Minimax objective function: $\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right]$

Alternate between:

1. Gradient **ascent** on discriminator

$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_z \right]$$

2. Gradient **descent** on generator

$$\min_{\theta_g} \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$$

[Goodfellow et al., 2014]

$\sum_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$

)))



Minimax objective function: $\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right]$

Alternate between:

1. Gradient **ascent** on discriminator

$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_z \right]$$

2. Gradient **descent** on generator

$$\min_{\theta_g} \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$$

In practice, optimizing this generator objective does not work well!

[Goodfellow et al., 2014]

$\sum_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$

)))



Minimax objective function: $\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right]$

Alternate between:

1. Gradient **ascent** on discriminator

$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z))) \right]$$

2. Gradient **descent** on generator

$$\min_{\theta_g} \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$$

In practice, optimizing this generator objective does not work well!

Goodfellow et al., 2014]

Gradient signal

where sample is

)))

already good $\log(1 - D(G(z)))$ When sample is likely fake, want to learn from it to improve generator. But gradient in this region 0.2 0.4 0.6 0.8 0.0 1.0 is relatively flat! D(G(z))

* slide from Fei-Dei Li, Justin Johnson, Serena Yeung, cs231n Stanford



dominated by region

Minimax objective function: $\min_{\theta_{g}} \max_{\theta_{d}} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_{d}}(x) + \mathbb{E}_{z \sim p(z)} \log(1 - D_{\theta_{d}}(G_{\theta_{g}}(z))) \right]$

Alternate between:

1. Gradient **ascent** on discriminator

$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z} \right]$$

2. Instead, gradient **ascent** on generator, different objective

$$\max_{\theta_g} \mathbb{E}_{z \sim p(z)} \log(D_{\theta_d}(G_{\theta_g}(z)))$$

Instead of minimizing likelihood of discriminator being correct, now maximize likelihood of discriminator being wrong.

Same objective of fooling discriminator, but now higher gradient signal for bad samples => works much better! Standard in practice.

Goodfellow et al., 2014

 $\sim_{p(z)} \log(1 - D_{\theta_d}(G_{\theta_g}(z)))$





Sampling GANs





Generated Samples



GANs with Convolutional Architectures



[Radford et al., 2016]



GANs with Convolutional Architectures

Interpolating between points in latent space



[Radford et al., 2016]



Smiling woman

Samples from the model



Neutral womai Neutral man



Radford et al., 2016]





Smiling woman

Samples from the model -

Average z vectors, do arithmetic







Radford et al., 2016]

Neutral man Neutral womai





Smiling woman

Samples from the model -

Average z vectors, do arithmetic







Radford et al., 2016]

Neutral man Neutral womai



Smiling man









Samples from the model





[Radford et al., 2016]

Glasses Man No Glasses Man No Glasses Woman





Samples from the model



Average z vectors, do arithmetic









[Radford et al., 2016]

Glasses Man No Glasses Man No Glasses Woman





Glasses Man

Samples from the model



Average z vectors, do arithmetic









Radford et al., 2016]

No Glasses Man No Glasses Woman

Radford et al, **ICLR 2016**













Year of the **GAN**

Better training and generation



(a) Church outdoor.



(b) Dining room.



(d) Conference room. (c) Kitchen. LSGAN. Mao et al. 2017.



BEGAN. Bertholet et al. 2017.

Source->Target domain transfer

Input



horse \rightarrow zebra



 $zebra \rightarrow horse$



apple \rightarrow orange



CycleGAN. Zhu et al. 2017.









→ summer Yosemite



→ winter Yosemite

Text -> Image Synthesis

this small bird has a pink breast and crown, and black almost all black with a red primaries and secondaries.



this magnificent fellow is crest, and white cheek patch.





Reed et al. 2017.

Many GAN applications









Pix2pix. Isola 2017. Many examples at https://phillipi.github.io/pix2pix/





Year of the GAN

- GAN Generative Adversarial Networks
- 3D-GAN Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Model
- acGAN Face Aging With Conditional Generative Adversarial Networks
- AC-GAN Conditional Image Synthesis With Auxiliary Classifier GANs
- AdaGAN AdaGAN: Boosting Generative Models
- AEGAN Learning Inverse Mapping by Autoencoder based Generative Adversarial Nets
- AffGAN Amortised MAP Inference for Image Super-resolution
- AL-CGAN Learning to Generate Images of Outdoor Scenes from Attributes and Semantic Layouts
- ALI Adversarially Learned Inference
- AM-GAN Generative Adversarial Nets with Labeled Data by Activation Maximization
- AnoGAN Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker
- ArtGAN ArtGAN: Artwork Synthesis with Conditional Categorial GANs
- b-GAN b-GAN: Unified Framework of Generative Adversarial Networks
- Bayesian GAN Deep and Hierarchical Implicit Models
- BEGAN BEGAN: Boundary Equilibrium Generative Adversarial Networks
- BiGAN Adversarial Feature Learning
- BS-GAN Boundary-Seeking Generative Adversarial Networks
- CGAN Conditional Generative Adversarial Nets
- CaloGAN CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calo with Generative Adversarial Networks
- CCGAN Semi-Supervised Learning with Context-Conditional Generative Adversarial Networks
- CatGAN Unsupervised and Semi-supervised Learning with Categorical Generative Adversarial Networks
- CoGAN Coupled Generative Adversarial Networks

odeling	•	Context-RNN-GAN - Contextual RNN-GANs for Abstract Reasoning Diagram Generation
	•	C-RNN-GAN - C-RNN-GAN: Continuous recurrent neural networks with adversarial training
	•	CS-GAN - Improving Neural Machine Translation with Conditional Sequence Generative Adversarial Nets
	•	CVAE-GAN - CVAE-GAN: Fine-Grained Image Generation through Asymmetric Training
	•	CycleGAN - Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks
	•	DTN - Unsupervised Cross-Domain Image Generation
	•	DCGAN - Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks
	•	DiscoGAN - Learning to Discover Cross-Domain Relations with Generative Adversarial Networks
	•	DR-GAN - Disentangled Representation Learning GAN for Pose-Invariant Face Recognition
	•	DualGAN - DualGAN: Unsupervised Dual Learning for Image-to-Image Translation
	•	EBGAN - Energy-based Generative Adversarial Network
	•	f-GAN - f-GAN: Training Generative Neural Samplers using Variational Divergence Minimization
Discovery	•	FF-GAN - Towards Large-Pose Face Frontalization in the Wild
	•	GAWWN - Learning What and Where to Draw
	•	GeneGAN - GeneGAN: Learning Object Transfiguration and Attribute Subspace from Unpaired Data
	•	Geometric GAN - Geometric GAN
	•	GoGAN - Gang of GANs: Generative Adversarial Networks with Maximum Margin Ranking
	•	GP-GAN - GP-GAN: Towards Realistic High-Resolution Image Blending
	•	IAN - Neural Photo Editing with Introspective Adversarial Networks
alorimeters	•	iGAN - Generative Visual Manipulation on the Natural Image Manifold
	•	IcGAN - Invertible Conditional GANs for image editing
	•	ID-CGAN - Image De-raining Using a Conditional Generative Adversarial Network
	•	Improved GAN - Improved Techniques for Training GANs
	•	InfoGAN - InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversaria
	•	LAGAN - Learning Particle Physics by Example: Location-Aware Generative Adversarial Networks for Physics

• LAPGAN - Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks

Synthesis

* slide from Fei-Dei Li, Justin Johnson, Serena Yeung, cs231n Stanford



al Nets

GANS

Don't work with an explicit density function Take game-theoretic approach: learn to generate from training distribution through 2-player game

Pros:

— Beautiful, state-of-the-art samples!

Cons:

- Trickier / more unstable to train
- Can't solve inference queries such as p(x), p(z|x)

Active area of research:

- Better loss functions, more stable training (Wasserstein GAN, LSGAN, many others)
- Conditional GANs, GANs for all kinds of applications