Learning Deep Generative Models

Russ Salakhutdinov

Machine Learning Department Carnegie Mellon University Canadian Institute for Advanced Research

Carnegie Mellon University



Impact of Deep Learning

- Speech Recognition
- ► Computer Vision
- Recommender Systems
- Language Understanding
- Drug Discovery and Medical Image Analysis

Statistical Generative Models



Grover and Ermon, DGM Tutorial

Statistical Generative Models



Training Data(CelebA)

Model Samples (Karras et.al., 2018)

4 years of progression on Faces



Brundage et al., 2017

Conditional Generation

Conditional generative model P(zebra images| horse images)



► Style Transfer



Input Image

Monet

Van Gogh Zhou el al., Cycle GAN 2017

Conditional Generation

Conditional generative model P(zebra images| horse images)



► Failure Case



Zhou el al., Cycle GAN 2017

Face2Face

Source actor



Real-time reenactment

Target actor

Face2Face, Thies et al, CVPR 2016



Explicit Density p(x)

8

Talk Roadmap

- Fully Observed Models
- Undirected Deep Generative Models
 - Restricted Boltzmann Machines (RBMs),
 - Deep Boltzmann Machines (DBMs)
- Directed Deep Generative Models
 - Variational Autoencoders (VAEs)
 - Normalizing Flows
- Generative Adversarial Networks (GANs)

Fully Observed Models

Density Estimation by Autoregression

$$p(x_1, \dots, x_d) = \prod_{i=1}^d p(x_i | x_{i-1}, \dots, x_1) \approx \prod_{i=1}^d p(x_i | g(x_{i-1}, \dots, x_1))$$

$$\emptyset \rightarrow h_1 = g(\emptyset, h_0) \rightarrow \square p(x_1)$$
Each conditional can be a deep neural network
$$(x_1 \rightarrow h_2 = g(x_1, h_1) \rightarrow \square p(x_2 | x_1)$$

$$(x_{d-1} \rightarrow h_d = g(x_{d-1}, h_{d-1}) \rightarrow \square p(x_d | x_d, \dots, x_1)$$

Ordering of variables is crucial

NADE (Uria 2013), MADE (Germain 2017), MAF (Papamakarios 2017), PixelCNN (van den Oord, et al, 2016)

Fully Observed Models

Density Estimation by Autoregression



PixelCNN (van den Oord, et al, 2016)

NADE (Uria 2013), MADE (Germain 2017), MAF (Papamakarios 2017), PixelCNN (van den Oord, et al, 2016)

WaveNet

Generative Model of Speech Signals











Talk Roadmap

- Fully Observed Models
- Undirected Deep Generative Models
 - Restricted Boltzmann Machines (RBMs),
 - Deep Boltzmann Machines (DBMs)
- Directed Deep Generative Models
 - Variational Autoencoders (VAEs)
 - Normalizing Flows
- Generative Adversarial Networks (GANs)

Restricted Boltzmann Machines



- RBM is a Markov Random Field with
 - Stochastic binary visible variables $\mathbf{v} \in \{0, 1\}^D$.
 - Stochastic binary hidden variables $\mathbf{h} \in \{0,1\}^F$.
 - Bipartite connections

Markov random fields, Boltzmann machines, log-linear models.

Maximum Likelihood Learning



$$P_{\theta}(\mathbf{v}) = \frac{P^{*}(\mathbf{v})}{\mathcal{Z}(\theta)} = \frac{1}{\mathcal{Z}(\theta)} \sum_{\mathbf{h}} \exp\left[\mathbf{v}^{\top} W \mathbf{h} + \mathbf{a}^{\top} \mathbf{h} + \mathbf{b}^{\top} \mathbf{v}\right]$$

Maximize log-likelihood objective:

$$L(\theta) = \frac{1}{N} \sum_{n=1}^{N} \log P_{\theta}(\mathbf{v}^{(n)})$$

Derivative of the log-likelihood:

$$\begin{split} \frac{\partial L(\theta)}{\partial W_{ij}} &= \frac{1}{N} \sum_{n=1}^{N} \frac{\partial}{\partial W_{ij}} \log \left(\sum_{\mathbf{h}} \exp \left[\mathbf{v}^{(n)\top} W \mathbf{h} + \mathbf{a}^{\top} \mathbf{h} + \mathbf{b}^{\top} \mathbf{v}^{(n)} \right] \right) - \frac{\partial}{\partial W_{ij}} \log \mathcal{Z}(\theta) \\ &= \mathbf{E}_{P_{data}} \left[v_i h_j \right] - \mathbf{E}_{P_{\theta}} \left[v_i h_j \right] \\ \mathbf{h}; \theta) &= P(\mathbf{h} | \mathbf{v}; \theta) P_{data}(\mathbf{v}) \\ &= \frac{1}{N} \sum_{n} \delta(\mathbf{v} - \mathbf{v}^{(n)}) \end{split}$$
 Difficult to compute: exponentially many configurations

 $P_{data}(\mathbf{v},$

 $P_{data}(\mathbf{v})$

Learning Features



RBMs for Real-valued & Count Data

4 million **unlabelled** images



Learned features (out of 10,000)





REUTERS

Associated Press

Reuters dataset: 804,414 unlabeled newswire stories

Bag-of-Words



ru m ye SO

Learned features: ``topics''

russian	clinton	computer	trade
russia	house	system	country
moscow	president	product	import
yeltsin	bill	software	world
soviet	congress	develop	economy

stock wall street point dow

Deep Boltzmann Machines



Built from **unlabeled** inputs.

(Salakhutdinov 2008, Salakhutdinov & Hinton 2009)

Deep Boltzmann Machines

Learn simpler representations, then compose more complex ones



Built from **unlabeled** inputs.

(Salakhutdinov 2008, Salakhutdinov & Hinton 2009)

Model Formation

 $P_{\theta}(\mathbf{v}, \mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \mathbf{h}^{(3)}) = \frac{1}{\mathcal{Z}(\theta)} \exp\left[\mathbf{v}^{\top} W^{(1)} \mathbf{h}^{(1)} + \mathbf{h}^{(1)^{\top}} W^{(2)} \mathbf{h}^{(2)} + \mathbf{h}^{(2)^{\top}} W^{(3)} \mathbf{h}^{(3)}\right]$



Same as RBMs $\theta = \{W^1, W^2, W^3\}$ model parameters

- Dependencies between hidden variables
- All connections are undirected
- Maximum Likelihood Learning:

 $\frac{\partial \log P_{\theta}(\mathbf{v})}{\partial W^{1}} = \mathbb{E}_{P_{data}}[\mathbf{vh^{1}}^{\top}] - \mathbb{E}_{P_{\theta}}[\mathbf{vh^{1}}^{\top}]$

Both expectations are intractable

Approximate Learning

$$P_{\theta}(\mathbf{v}, \mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \mathbf{h}^{(3)}) = \frac{1}{\mathcal{Z}(\theta)} \exp\left[\mathbf{v}^{\top} W^{(1)} \mathbf{h}^{(1)} + \mathbf{h}^{(1)}^{\top} W^{(2)} \mathbf{h}^{(2)} + \mathbf{h}^{(2)}^{\top} W^{(3)} \mathbf{h}^{(3)}\right]$$



Maximum Likelihood Learning:

$$\frac{\partial \log P_{\theta}(\mathbf{v})}{\partial W^{1}} = \mathbb{E}_{P_{data}}[\mathbf{vh^{1}}^{\top}] - \mathbb{E}_{P_{\theta}}[\mathbf{vh^{1}}^{\top}]$$



Approximate Learning

$$P_{\theta}(\mathbf{v}, \mathbf{h}^{(1)}, \mathbf{h}^{(2)}, \mathbf{h}^{(3)}) = \frac{1}{\mathcal{Z}(\theta)} \exp\left[\mathbf{v}^{\top} W^{(1)} \mathbf{h}^{(1)} + \mathbf{h}^{(1)}^{\top} W^{(2)} \mathbf{h}^{(2)} + \mathbf{h}^{(2)}^{\top} W^{(3)} \mathbf{h}^{(3)}\right]$$



Maximum Likelihood Learning:



Good Generative Model?

► CIFAR Dataset





Samples

Training

Learning Part-Based Representations

Deep Belief Network





Groups of parts

Object Parts

Trained on face images.

Lee, Grosse, Ranganath, Ng, ICML 2009

Learning Part-Based Representations

Faces Cars Elephants Chairs

Lee, Grosse, Ranganath, Ng, ICML 2009

Talk Roadmap

- Fully Observed Models
- Undirected Deep Generative Models
 - Restricted Boltzmann Machines (RBMs),
 - Deep Boltzmann Machines (DBMs)
- Directed Deep Generative Models
 - Variational Autoencoders (VAEs)
 - Normalizing Flows
- Generative Adversarial Networks (GANs)

Helmholtz Machines

▶ Hinton, G. E., Dayan, P., Frey, B. J. and Neal, R., Science 1995



Input data

- ▶ Kingma & Welling, 2014
- ▶ Rezende, Mohamed, Daan, 2014
- Mnih & Gregor, 2014
- ▶ Bornschein & Bengio, 2015
- ► Tang & Salakhutdinov, 2013

Helmholtz Machines

Helmholtz Machine



Deep Boltzmann Machine



Input data

Deep Directed Generative Models



Latent Variable Models



 Conditional distributions are parameterized by deep neural networks

Directed Deep Generative Models

Directed Latent Variable Models with Inference Network



- Maximum log-likelihood objective $\max_{\theta} \sum_{\mathbf{x} \in \mathcal{D}} \log p_{\theta}(\mathbf{x})$
- Marginal log-likelihood is intractable:

$$\log p_{\theta}(\mathbf{x}) = \log \int p_{\theta}(\mathbf{x}, \mathbf{z}) \mathrm{d}\mathbf{z}$$

Key idea: Approximate true posterior p(z|x) with a simple, tractable distribution q(z|x) (inference/recognition network).

Grover and Ermon, DGM Tutorial

Variational Autoencoders (VAEs)

The VAE defines a generative process in terms of ancestral sampling through a cascade of hidden stochastic layers:

$$p(\mathbf{x}|\boldsymbol{\theta}) = \sum_{\mathbf{h}^{1},...,\mathbf{h}^{L}} p(\mathbf{h}^{L}|\boldsymbol{\theta})p(\mathbf{h}^{L-1}|\mathbf{h}^{L},\boldsymbol{\theta})\cdots p(\mathbf{x}|\mathbf{h}^{1},\boldsymbol{\theta})$$
Each conditional term denotes a nonlinear relationship
$$\mathbf{W}^{3} \neq P(\mathbf{h}^{2}|\mathbf{h}^{3})$$

$$\mathbf{h}^{2} \neq P(\mathbf{h}^{1}|\mathbf{h}^{2})$$

$$\mathbf{h}^{2} \neq P(\mathbf{h}^{1}|\mathbf{h}^{2})$$

$$\mathbf{h}^{2} \neq P(\mathbf{h}^{1}|\mathbf{h}^{2})$$

$$\mathbf{h}^{2} \neq P(\mathbf{h}^{1}|\mathbf{h}^{2})$$

$$\mathbf{h}^{3} \neq P(\mathbf{h}^{2}|\mathbf{h}^{1})$$

Input data

Variational Autoencoders (VAEs)

► Single stochastic (Gaussian) layer, followed by many deterministic layers





$$p(\mathbf{z}) = \mathcal{N}(0, I)$$
$$p_{\theta}(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\mu(\mathbf{z}, \theta), \Sigma(\mathbf{z}, \theta))$$

Deep neural network parameterized by θ . (Can use different noise models)

$$q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mu(\mathbf{x},\phi), \Sigma(\mathbf{x},\phi))$$

Deep neural network parameterized by ϕ .

Variational Bound

► VAE is trained to maximize the variational lower bound:

$$\begin{split} \log p_{\theta}(\mathbf{x}) &= \log \int p_{\theta}(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \log \int q_{\phi}(\mathbf{z} | \mathbf{x}) \frac{p_{\theta}(\mathbf{x}, \mathbf{z})}{q_{\phi}(\mathbf{z} | \mathbf{x})} d\mathbf{z} \\ &= \log \mathbb{E}_{q_{\phi}(\mathbf{z} | \mathbf{x})} \left[\frac{p_{\theta}(\mathbf{x}, \mathbf{z})}{q_{\phi}(\mathbf{z} | \mathbf{x})} \right] \\ &\geq \mathbb{E}_{q_{\phi}(\mathbf{z} | \mathbf{x})} \log \left[\frac{p_{\theta}(\mathbf{x}, \mathbf{z})}{q_{\phi}(\mathbf{z} | \mathbf{x})} \right] \\ &= \log p_{\theta}(\mathbf{x}) - \mathrm{KL} \left(q_{\phi}(\mathbf{z} | \mathbf{x}) | | p_{\theta}(\mathbf{z} | \mathbf{x}) \right) = \mathcal{L}(\mathbf{x}) \end{split}$$

- Trading off the data log-likelihood and the KL divergence from the true posterior
- Hard to optimize the variational bound with respect to the q recognition network (high variance)
- ▶ Key idea of Kingma and Welling is to use reparameterization trick

Reparameterization

• Assume that the recognition distribution is Gaussian:

$$q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mu(\mathbf{x},\phi), \Sigma(\mathbf{x},\phi))$$

q_φ(z|x) ► Alternatively, we can express this in term of auxiliary variable:

$$\mathbf{z}(\epsilon, \mathbf{x}, \phi) = \Sigma(\mathbf{x}, \phi)^{1/2} \epsilon + \mu(\mathbf{x}, \phi), \quad \epsilon \sim \mathcal{N}(0, I)$$

- ► The recognition distribution can be expressed as a deterministic mapping
 - \blacktriangleright Distribution of ϵ does not depend on φ

Deterministic Encoder

 $\mathbf{z}(\epsilon, \mathbf{x}, \phi)$

 $p_{\theta}(\mathbf{x}|\mathbf{z})$

Computing Gradients

The gradients of the variational bound w.r.t the recognition (similar w.r.t the generative) parameters:

$$\begin{aligned} \nabla_{\phi} \mathcal{L}(\mathbf{x}) &= \nabla_{\phi} \mathbb{E}_{\mathbf{z} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} \left[\log \frac{p_{\theta}(\mathbf{x}, \mathbf{z})}{q_{\phi}(\mathbf{z}|\mathbf{x})} \right] & \text{Autoencoder} \\ &= \nabla_{\phi} \mathbb{E}_{\epsilon \sim \mathcal{N}(0,I)} \left[\log \frac{p_{\theta}(\mathbf{x}, \mathbf{z}(\epsilon, \mathbf{x}, \phi))}{q_{\phi}(\mathbf{z}(\epsilon, \mathbf{x}, \phi)|\mathbf{x})} \right] \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(0,I)} \left[\nabla_{\phi} \log \frac{p_{\theta}(\mathbf{x}, \mathbf{z}(\epsilon, \mathbf{x}, \phi))}{q_{\phi}(\mathbf{z}(\epsilon, \mathbf{x}, \phi)|\mathbf{x})} \right] \\ & \text{Gradients can be} \\ & \text{computed by backprop} \end{aligned}$$
 The mapping \mathbf{z} is a deterministic neural net for fixed ε

VAE Assumptions

► Remember the variational bound:

$$\mathcal{L}(\mathbf{x}) = \log p(\mathbf{x}) - D_{\mathrm{KL}} \left(q(\mathbf{h}|\mathbf{x}) \right) || p(\mathbf{h}|\mathbf{x}) \rangle$$

- ► The variational assumptions must be approximately satisfied.
- The posterior distribution must be approximately factorial (common practice) and predictable with a feed-forward net.
- We can relax these assumptions using a tighter lower bound on marginal loglikelihood.

Importance Weighted Autoencoder

Improve VAE by using the following k-sample importance weighting of the log-likelihood:

$$\mathcal{L}_{k}(\mathbf{x}) = \mathbb{E}_{\mathbf{z}_{1},\mathbf{z}_{2},...,\mathbf{z}_{k} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} \left[\log \frac{1}{k} \sum_{i=1}^{k} \frac{p_{\theta}(\mathbf{x},\mathbf{z}_{i})}{q_{\phi}(\mathbf{z}_{i}|\mathbf{x})} \right]$$
$$= \mathbb{E}_{\mathbf{z}_{1},\mathbf{z}_{2},...,\mathbf{z}_{k} \sim q_{\phi}(\mathbf{z}|\mathbf{x})} \left[\log \frac{1}{k} \sum_{i=1}^{k} w_{i} \right]$$
$$\bullet \text{ where multiple z are sampled from the unnormalized importance weights}}$$

► Can improve the tightness of the bound.

Burda, Grosse, Salakhutdinov, ICLR 2016, Mnih & Rezende, ICML 2016

 $p_{\theta}(\mathbf{x}|\mathbf{z})$

Tighter Lower Bound

- ► Using more samples can only improve the tightness of the bound.
- ► For all k, the lower bounds satisfy:

 $\log p(\mathbf{x}) \ge \mathcal{L}_{k+1}(\mathbf{x}) \ge \mathcal{L}_k(\mathbf{x})$

• Moreover if $p(\mathbf{h}, \mathbf{x})/q(\mathbf{h}|\mathbf{x})$ is bounded, then:

 $\mathcal{L}_k(\mathbf{x}) \to \log p(\mathbf{x}), \text{ as } k \to \infty$

Burda, Grosse, Salakhutdinov, ICLR 2016, Mnih & Rezende, ICML 2016

Talk Roadmap

- Fully Observed Models
- Undirected Deep Generative Models
 - Restricted Boltzmann Machines (RBMs),
 - Deep Boltzmann Machines (DBMs)
- Directed Deep Generative Models
 - Variational Autoencoders (VAEs)
 - Normalizing Flows
- Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GAN)

- Implicit generative model for an unknown target density p(x)
- Converts sample from a known noise density p_Z(z) to the target p(x)



Unknown target density p(x) of data over domain \mathcal{X} , e.g. $\mathbb{R}^{32 \times 32}$



Noise density $p_Z(z)$ over space \mathcal{Z}



Distribution of generated samples should follow target density p(x)

[Slide Credit: Manzil Zaheer]

GAN Formulation

► GAN consists of two components



<u>Goal</u>: Produce samples indistinguishable from true data



 $D:\mathcal{X}\to\mathbb{R}$



Goal: Distinguish true and generated data apart Goodfellow et al, 2014

Wasserstein GAN

► WGAN optimization

$$\min_{G} \max_{D} W(G, D) = \mathbb{E}_{x \sim p} \left[D(x) \right] - \mathbb{E}_{z \sim p_{Z}} \left[D(G(z)) \right]$$

- ► Difference in expected output on real vs. generated images
 - ► Generator attempts to drive objective ≈ 0
- More stable optimization

Compare to training DBMs
$$\frac{\partial \log P_{\theta}(\mathbf{v})}{\partial W^{1}} = \mathbb{E}_{P_{data}}[\mathbf{vh^{1}}^{\top}] - \mathbb{E}_{P_{\theta}}[\mathbf{vh^{1}}^{\top}]$$

D outputs:

D(x) = 1 real D(x) = 0 generated

Arjovsky et al., 2017

Modelling Point Cloud Data



43

Zaheer et al. Point Cloud GAN 2018

Interpolation in Latent Space







Zaheer et al. Point Cloud GAN 2018

Normalizing Flows

Directed Latent Variable Invertible models



The mapping between x and z is deterministic and invertible:

$$egin{array}{rcl} \mathbf{x} &=& \mathbf{f}_{ heta}(\mathbf{z}) \ \mathbf{z} &=& \mathbf{f}_{ heta}^{-1}(\mathbf{x}) \end{array}$$

► Use change-of-variables to relate densities between z and x

$$p_X(\mathbf{x};\theta) = p_Z(\mathbf{z}) \left| \det \frac{\partial \mathbf{f}_{\theta}^{-1}(\mathbf{x})}{\partial X} \right|_{X=\mathbf{x}}$$

Grover and Ermon DGM Tutorial, NICE (Dinh et al. 2014), Real NVP (Dinh et al. 2016)

Normalizing Flows

► Invertible transformations can be composed:

$$\mathbf{z}^{M} \triangleq \mathbf{f}_{\theta}^{M} = \mathbf{f}_{\theta}^{M} \cdot \otimes \cdot \mathbf{f}_{\theta}^{1} \cdot (\mathbf{z}^{0}) \mathcal{P}_{X}(\mathbf{x}; \boldsymbol{\beta}_{X}(\mathbf{x}; \boldsymbol{\beta}; \boldsymbol{$$

► Planar Flows

$$f(\mathbf{z}) = \mathbf{z} + \mathbf{u}g(\mathbf{w}^{\top}\mathbf{z} + b)$$

$$m = 0 \qquad m = 1 \qquad m = 2 \qquad m = 10$$

Rezendre and Mohamed, 2016, Grover and Ermon DGM Tutorial

Normalizing Flows

Maximum log-likelihood objective

$$\max_{\theta} \max_{\theta} \log p_{X}(\mathcal{D}; \theta) = \sum_{\mathbf{x} \in \mathcal{D}} \left(\log p_{Z}(\mathbf{z}) \log \left| \log \left| \log \frac{\partial (\mathbf{f}_{\theta}) + 1}{\partial X \partial X} \right|_{X} \right|_{X \to \mathbf{x}} \right)$$

- Exact log-likelihood evaluation via inverse transformations
- Sampling from the model

$$\mathbf{z} \ \mathbf{z} \sim p_Z(\mathbf{z}), \quad \mathbf{x} = \mathbf{f}_{\theta}(\mathbf{z})$$

► Inference over the latent representations: $\mathbf{z} \neq \mathbf{f}_{\theta} \mathbf$

Rezendre and Mohamed, 2016, Grover and Ermon DGM Tutorial

Example: GLOW

 Generative Flow with Invertible 1x1 Convolutions https://blog.openai.com/glow/



Example: GLOW



Remove Beard

Decrease Age

Thank you

Important Breakthroughs

► Deep Belief Networks, 2006 (Unsupervised)



- ► Theoretical Breakthrough:
 - Adding additional layers improves variational lower-bound.
- Efficient Learning and Inference with multiple layers

Hinton, G. E., Osindero, S. and Teh, Y., A fast learning algorithm for deep belief nets, Neural Computation, 2006.

Learning Part-Based Representations

Deep Belief Network





Groups of parts

Object Parts

Trained on face images.

Lee, Grosse, Ranganath, Ng, ICML 2009

Learning Part-Based Representations

Faces Cars Elephants Chairs

Lee, Grosse, Ranganath, Ng, ICML 2009

Important Breakthroughs

- Deep Convolutional Nets for Vision (Supervised)
 - Krizhevsky, A., Sutskever, I. and Hinton, G. E., ImageNet Classification with Deep Convolutional Neural Networks, NIPS, 2012.





1.2 million training images 1000 classes



- Deep Nets for Speech (Supervised)
 - Hinton et. al. Deep Neural Networks for Acoustic Modeling in Speech Recognition: The Shared Views of Four Research Groups, IEEE Signal Processing Magazine. 2012.

GAN Formulation: Discriminator

Discriminator's objective: Tell real and generated data apart like a classifier

$$\max_{D} \mathbb{E}_{x \sim p} \left[\log D(x) \right] + \mathbb{E}_{z \sim p_{Z}} \left[\log \left(1 - D(G(z)) \right) \right]$$



GAN Formulation: Generator

Generator's objective: Fool the best discriminator

$$\min_{G} \max_{D} \mathbb{E}_{x \sim p} \left[\log D(x) \right] + \mathbb{E}_{z \sim p_{Z}} \left[\log \left(1 - D(G(z)) \right) \right]$$



GAN Formulation: Optimization

Overall GAN optimization

$$\min_{G} \max_{D} V(G, D) = \mathbb{E}_{x \sim p} \left[\log D(x) \right] + \mathbb{E}_{z \sim p_{Z}} \left[\log \left(1 - D(G(z)) \right) \right]$$

The generator-discriminator are iteratively updated using SGD to find "equilibrium" of a "min-max objective" like a game

 $G \leftarrow G - \eta_G \nabla_G V(G, D)$

 $D \leftarrow D - \eta_D \nabla_D V(G, D)$



[Slide Credit: Manzil Zaheer]