Generative Models



Desiderata for generative models

Probability evaluation: given a sample, it is computationally efficient to evaluate the probability of this sample.
 χ () (χ)

• Flexible model family: it is easy to incorporate any neural network models.

$$\chi = L_{\theta}(2)$$

$$Z - N(U, J)$$

• **Easy sampling:** it is computationally efficient to sample a data from the probabilistic model.

$$P_{\Theta}(\mathbf{x})$$

Desiderata for generative models



Slide credit to Yang Song

Taxonomy of generative models



Image credits to Andrej Risteski

Key challenge for building generative models



Slide credit to Yang Song

Slide credit to Yang Song

Key challenge for building generative models

Approximating the normalizing constant

- Variational auto-encoders [Kingma & Welling 2014, Rezende et al. 2014]
- Energy-based models [Ackley et al. 1985, LeCun et al. 2006]

Using restricted neural network models

- Autoregressive models [Bengio & Bengio 2000, van den Oord et al. 2016]
- Normalizing flow models [Dinh et al. 2014, Rezende & Mohamed 2015]

Generative adversarial networks (GANs)

• Model the generation process, not the probability distribution [Goodfellow et al. 2014]











Variational Autoencoder



Architecture



VAE Training

evidence some bound

- Training via optimizing ELBO
 - $L(\phi, \theta; x) = \mathbb{E}_{z \sim q(z|x;\phi)}[\log p(z|x;\theta)] KL(q(z|x;\phi)||p(z))$
 - Likelihood term + KL penalty
- KL penalty for Gaussians has closed form.
- Likelihood term (reconstruction loss):
 - Monte-Carlo estimation
 - Draw samples from $q(z | x; \phi) \begin{pmatrix} z_1, \dots, z_k \end{pmatrix}$
 - Compute gradient of θ :
 - No gradient! • $x \sim N(f(z; \theta); I)$ • $p(x) = \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2} ||x - f(z;\theta)||_2^2)$



VAE Training

• Likelihood term (reconstruction loss):

- Gradient for ϕ . Loss: $L(\phi) = \mathbb{E}_{z \sim q(z;\phi)} \left[\log p(x \mid z) \right]$
- Reparameterization trick:
- $z \sim N(\mu, \Sigma) \Leftrightarrow z = \mu + \epsilon, \epsilon \sim N(0, \Sigma)$ • $L(\phi) \propto \mathbb{E}_{z \sim q(z|\phi)} \left[\|f(z; \theta) - x\|_2^2 \right]$
 - $L(\phi) \propto \mathbb{E}_{z \sim q(z|\phi)} \left[\|f(z;\theta) x\|_2^2 \right]$ $\propto \mathbb{E}_{\epsilon \sim N(0,I)} \left[\|f(\mu(x;\phi) + \sigma(x;\phi) \cdot \epsilon;\theta) x\|_2^2 \right]$
- Monte-Carlo estimate for $\nabla L(\phi)$
 - practice. rse (Sample St. -- EK) is ensush
- End-to-end training

JN



VAE vs. AE

- AE: classical unsupervised representation learning method.
- VAR: a probabilistic model of AE
 - AE + Gaussian noise on z
 - KL penalty: L_2 constraint on the latent vector z



 $Z \sim N(0, I)$ $P_0(x(Z))$

Conditioned VAE

• Semi-supervised learning: some labels are also available



conditioned generation

Comments on VAE

- Pros:
 - Flexible architecture
 - Stable training
- Cons:
 - Inaccurate probability evaluation (approximate inference)

Energy-Based Models



Energy-based Models

- Goal of generative models:
 - a probability distribution of data: P(x)
- Requirements
 - $P(x) \ge 0$ (non-negative)
 - $\int_{x} P(x)dx = 1$

chouse came NN

- Energy-based model:
 - Energy function: $E(x; \theta)$, parameterized by θ

•
$$P(x) = \frac{1}{z} \exp(-E(x;\theta))$$
 (why exp?)
• $z = \int_{z} \exp(-E(x;\theta)) dx$

Boltzmann Machine

y JS binary

Generative model



• When y_i is binary, patterns are affecting each other through W



Boltzmann Machine: Training

- Objective: maximum likelihood learning (assume T = 1):
 - Probability of one sample:

$$P(y) = \frac{\exp(\frac{1}{2}y \psi)}{\sum_{y'} \exp(y' \forall Wy')}$$

• Maximum log-likelihood:

 $L(W) = \frac{1}{N} \sum_{Y \in Y} \frac{1}{2} \frac{y^7 W y}{2 \log 2} \log \sum_{y'} \frac{1}{2} \frac{y^7 W y}{y'} \log \sum_{y'} \frac{1}{2} \frac{y^7 W y}{y'}$ **Boltzmann Machine: Training** $V - \frac{1}{2} d x d$ $V - \frac{1}{2} d x d$ V - $E_{y'} [\underline{y_i'}]$ $=) u_{P} Monte - Caulo$ $Sample S = \{\underline{y_i'}, ..., \underline{y_k}\}$ $=) V_{W_i} L \approx \frac{1}{N_{xe0}} \overline{\Sigma_{e0}} y_i \cdot y_j - \frac{1}{N_{xe0}} \overline{\Sigma_{e0}} y_i \cdot y_j$

Boltzmann Machine: Training

Sampling:
$$\mathcal{M}(\mathcal{M}($$

, initialize $\mathcal{Y}(\mathbf{D}) \in \mathcal{P}^d$
tor $t=1,...,T$
i iterate over $i \in \{1,...,d\}, \mathcal{Y}_i(t) \sim \mathbb{P}(\cdot | \mathcal{Y}_{j\neq i}(t))$
 $= \sum_{i=1}^{n} \langle \mathcal{Y}(\mathbf{O}), -.., \mathcal{Y}(T) \rangle$



Restricted Bolzmann Machine

- A structured Boltzmann Machine
 - Hidden neurons are only connected to visible neurons
 - No intra-layer connections,
 - Invented by Paul Smolensky in '89
 - Became more practical after Hinton invested fast learning algorithms in mid 2000

latent Unvicible



Restricted Bolzmann Machine





Restricted Bolzmann Machine

- Sampling:
 - Randomly initialize visible neurons v_0
 - Iterative sampling between hidden neurons and visible neurons
 - Get final sample (v_∞,h_∞)
- Training:
 - MLE
 - Sampling to approximate gradient



Deep Bolzmann Machine

- Can we have a **deep** version of RBM?
 - Deep Belief Net ('06)
 - Deep Boltzmann Machine ('09)
- Sampling?
 - Forward pass: bottom-up
 - Backward pass: top-down
- Deep Bolzmann Machine
 - The very first deep generative model
 - Salakhudinov & Hinton



Deep Bolzmann Machine

Deep Boltzmann Machine



Gaussian visible units (raw pixel data)

Training Samples						
- Car	1	B	JA.		营	
\$	Y	X	1	A	8	
1		(m)		*	P.	
Ŕ	P	Ŷ	-	100	R	
A		X	1	II.	(I)	
(Lin)	245	Ť	A.		14 a	

Generated Samples \checkmark \checkmark

Summary

- Pros: powerful and flexible
 - An arbitrarily complex density function $p(x) = \frac{1}{z} \exp(-E(x))$
- Cons: hard to sample / train
 - Hard to sample:
 - MCMC sampling
 - Partition function
 - No closed-form calculation for likelihood
 - Cannot optimize MLE loss exactly
 - MCMC sampling

Normalizing Flows



Intuition about easy to sample

- Goal: design p(x) such that
 - Easy to sample
 - Tractable likelihood (density function)
- Easy to sample
 - Assume a continuous variable *z*
 - e.g., Gaussian $z \sim N(0,1)$, or uniform $z \sim \text{Unif}[0,1]$
 - x = f(z), x is also easy to sample

Intuition about tractable density

- Goal: design $f(z; \theta)$ such that
 - Assume *z* is from an "easy" distribution
 - $p(x) = p(f(z; \theta))$ has tractable likelihood
- Uniform: $z \sim \text{Unif}[0,1]$
 - Density p(z) = 1
 - x = 2z + 1, then p(x) = ? (). 5



Intuition about tractable density



Change of variable

• Suppose x = f(z) for some general non-linear $f(\cdot)$

• The linearized change in volume is determined by the Jacobian of $f(\ \cdot\)$:



Normalizing Flow Jun. Fra bijertay

- Idea
 - Sample z_0 from an "easy" distribution, e.g., standard Gaussian Apply \widetilde{K} bijections $z_i = f_i(z_{i-1})$, z_0 , -. $\widehat{Z}_k = k$

 - The final sample $x = f_K(z_K)$ has tractable desnity
- Normalizing Flow

•
$$z_0 \sim N(0,I), z_i = f_i(z_{i-1}), x = Z_K$$
 where $x, z_i \in \mathbb{R}^d$ and f_i is invertible

Every revertible function produces a normalized density function

$$\underbrace{p(z_i) = p(z_{i-1})}_{\bullet} \left| \det \left(\frac{\partial f_i}{\partial z_{i-1}} \right) \right|^{-1} / \underbrace{\mathcal{Q}(2_{\mathcal{Q}})}_{\bullet} - \underbrace{\mathcal{Q}(\mathcal{Z}_{\mathcal{Q}})}_{\bullet}$$



Normalizing Flow

- Generation is trivial
 - Sample z_0 then apply the transformations
- Log-likelihood

•
$$\log p(x) = \log p(Z_{k-1}) - \log \left| \det \left(\frac{\partial f_K}{\partial z_{K-1}} \right) \right|$$

• $\log p(x) = \log p(z_0) - \sum_i \log \left| \det \left(\frac{\partial f_i}{\partial z_{i-1}} \right) \right|$ $O(d^3)!!!!$



Normalizing Flow

- Naive flow model requires extremely expensive computation
 - Computing determinant of $d \times d$ matrices
- Idea:
 - Design a good bijection $f_i(z)$ such that the determinant is easy to compute

u, v EDd, AEPard.

<u>9</u>×=<u>1</u>

O(d)

- Technical tool: Matrix Determinant Lemma:
 - $\det(A + uv^{\top}) = (1 + v^{\top}A^{-1}u) \det A$
- Model:

•
$$\frac{\det(A + uv^{\mathsf{T}}) = (1 + v^{\mathsf{T}}A^{-1}u) \det A}{\mathsf{Model}}$$

•
$$\frac{f_{\theta}(z) = z + u \odot h(w^{\mathsf{T}}z + b)}{h(\cdot) \text{ chosen to be tanh}(\cdot)(0 < h'(\cdot) < 1)}$$

•
$$\frac{\theta}{h(\cdot)} = [u, w, b], \det\left(\frac{\partial f}{\partial z}\right) = \det(I + h'(w^{\mathsf{T}}z + b)uw^{\mathsf{T}}) = 1 + h'(w^{\mathsf{T}}z + b)u^{\mathsf{T}}w^{\mathsf{T}}$$

- Computation in O(d) time
- Remarks:

Plannar Flow

- $u^{\dagger}w > -1$ to ensure invertibility
- Require normalization on u and w

1[UI] / [[WI]

Planar Flow (Rezende & Mohamed, '16)

- $f_{\theta}(z) = z + uh\left(w^{\mathsf{T}}z + b\right)$
- 10 planar transformations can transform simple distributions into a more complex one



Extensions

- Other flow models uses triangular Jacobian (NICE, Dinh et al. '14)
- Invertible 1x1 convolutions (Kingma et al. '18)
- Auto-regressive flow:
 - WaveNet (Deepmind '16)
 - PixelCNN (Deepmind '16)

Summary



- Pros:
 - Easy to sample by transforming from a simple distribution
 - Easy to evaluate the probability
 - Easy training (MLE)
- Con
 - Most restricted neural network structure
 - Trade expressiveness for tractability