

DSC291: Machine Learning with Few Labels

Deep generative modeling
Generative adversarial learning

Zhiting Hu

Lecture 13, February 8, 2023

UC San Diego

HALICIOĞLU DATA SCIENCE INSTITUTE

Outline

- Generative adversarial networks (GANs)
- Normalizing Flow

Generative modeling

- In generative modeling, we'd like to train a network that models a distribution, such as a distribution over images.
- One way to judge the quality of the model is to sample from it.
- This field has seen rapid progress:



2009



2015



2018

Generative modeling

- In generative modeling, we'd like to train a network that models a distribution, such as a distribution over images.
- One way to judge the quality of the model is to sample from it.
- This field has seen rapid progress:




Generative modeling

- Modern approaches to generative modeling:
 - Variational Auto-encoder (Lecture #5)
 - Auto-regressive models (e.g., language model) (Lecture #6)
 - Generative adversarial networks (today)
 - Reversible architectures (today)
 - Diffusion models (later if time permits)

Implicit Generative Models


- **Implicit generative models** implicitly define a probability distribution
- Start by sampling the **code vector** \mathbf{z} from a fixed, simple distribution (e.g. spherical Gaussian)
- The **generator network** computes a differentiable function G mapping \mathbf{z} to an \mathbf{x} in data space



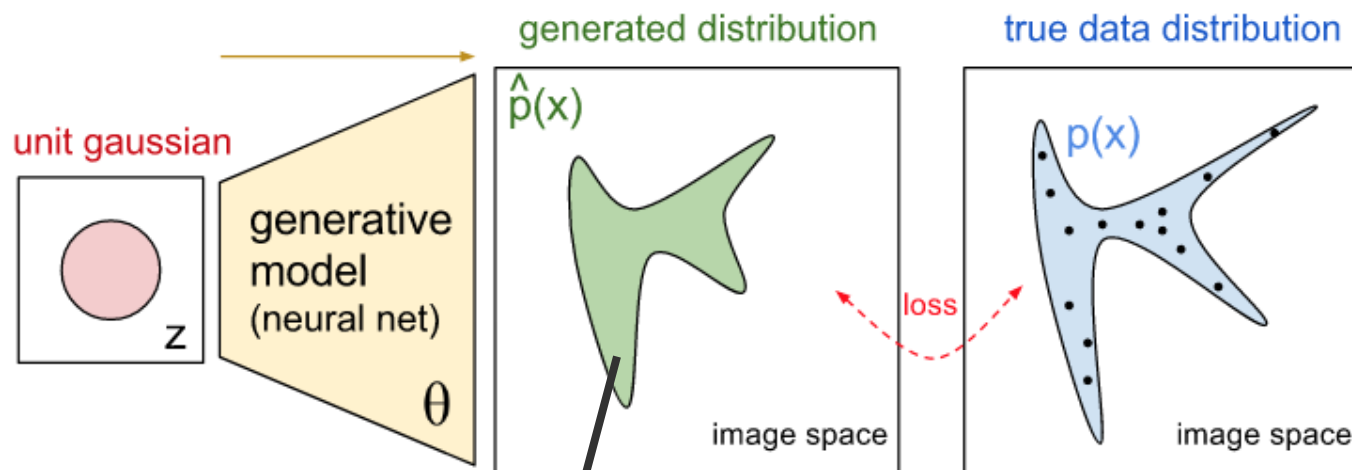
- a stochastic process to simulate data \mathbf{x}
- Intractable to evaluate likelihood

Implicit Generative Models

A 1-dimensional example:



Implicit Generative Models




<https://blog.openai.com/generative-models/>

Implicit Generative Models

- The advantage of implicit generative models: if you have some criterion for evaluating the quality of samples, then you can compute its gradient with respect to the network parameters, and update the network's parameters to make the sample a little better
- The idea behind **Generative Adversarial Networks (GANs)**: train two different networks
 - The generator network tries to produce realistic-looking samples
 - The discriminator network tries to figure out whether an image came from the training set or the generator network
- The generator network tries to fool the discriminator network

Generative Adversarial Nets (GANs)

- Generative model $\mathbf{x} = G_{\theta}(\mathbf{z})$, $\mathbf{z} \sim p(\mathbf{z})$
 - Maps noise variable \mathbf{z} to data space \mathbf{x}
 - Defines an implicit distribution over \mathbf{x} : $p_{g_{\theta}}(\mathbf{x})$
- Discriminator $D_{\phi}(\mathbf{x})$
 - Output the probability that \mathbf{x} came from the data rather than the generator




Generative Adversarial Nets (GANs)

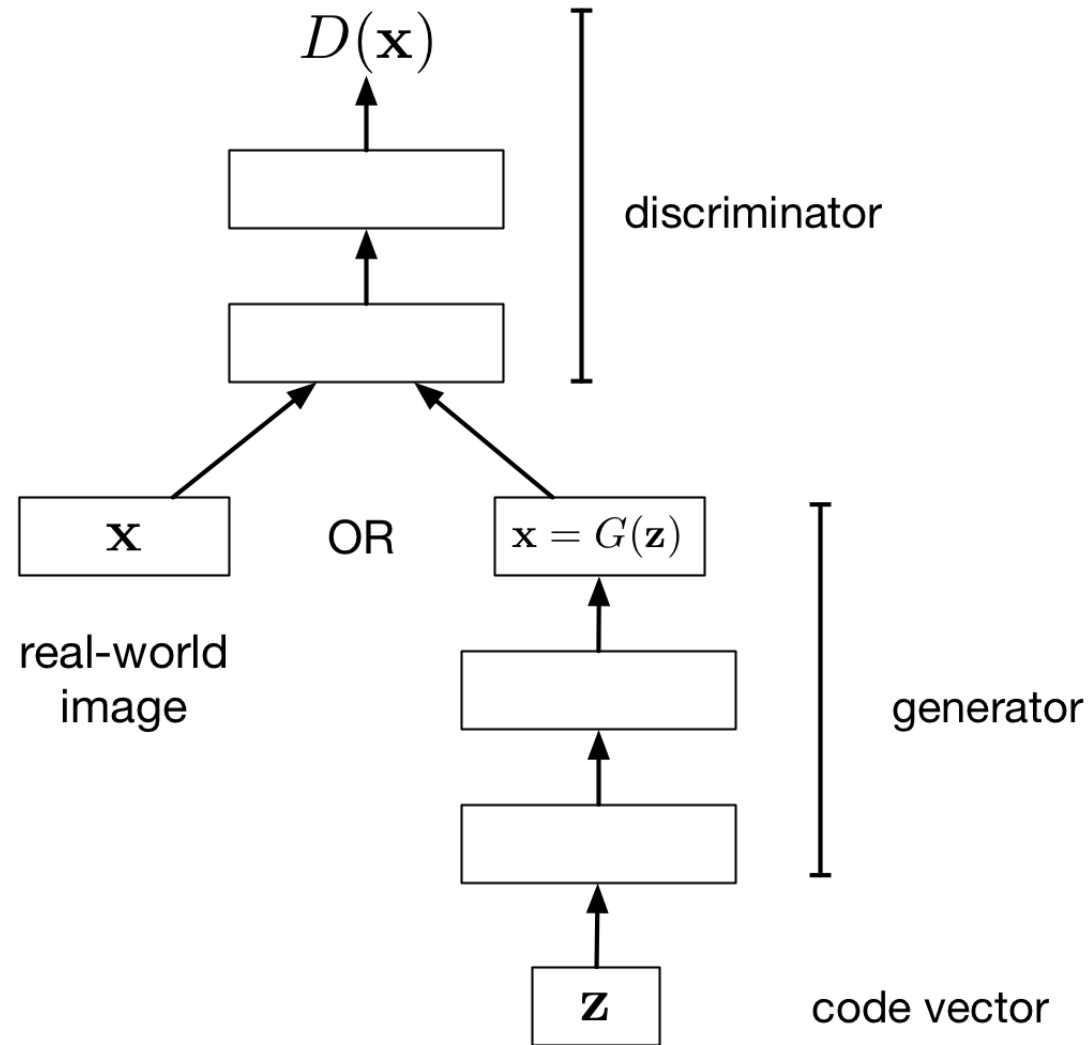
- Learning
 - A **minimax** game between the generator and the discriminator
 - Train D to maximize the probability of assigning the correct label to both training examples and generated samples
 - Train G to fool the discriminator

$$\max_D \mathcal{L}_D = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim G(\mathbf{z}), \mathbf{z} \sim p(\mathbf{z})} [\log(1 - D(\mathbf{x}))]$$

$$\min_G \mathcal{L}_G = \mathbb{E}_{\mathbf{x} \sim G(\mathbf{z}), \mathbf{z} \sim p(\mathbf{z})} [\log(1 - D(\mathbf{x}))].$$

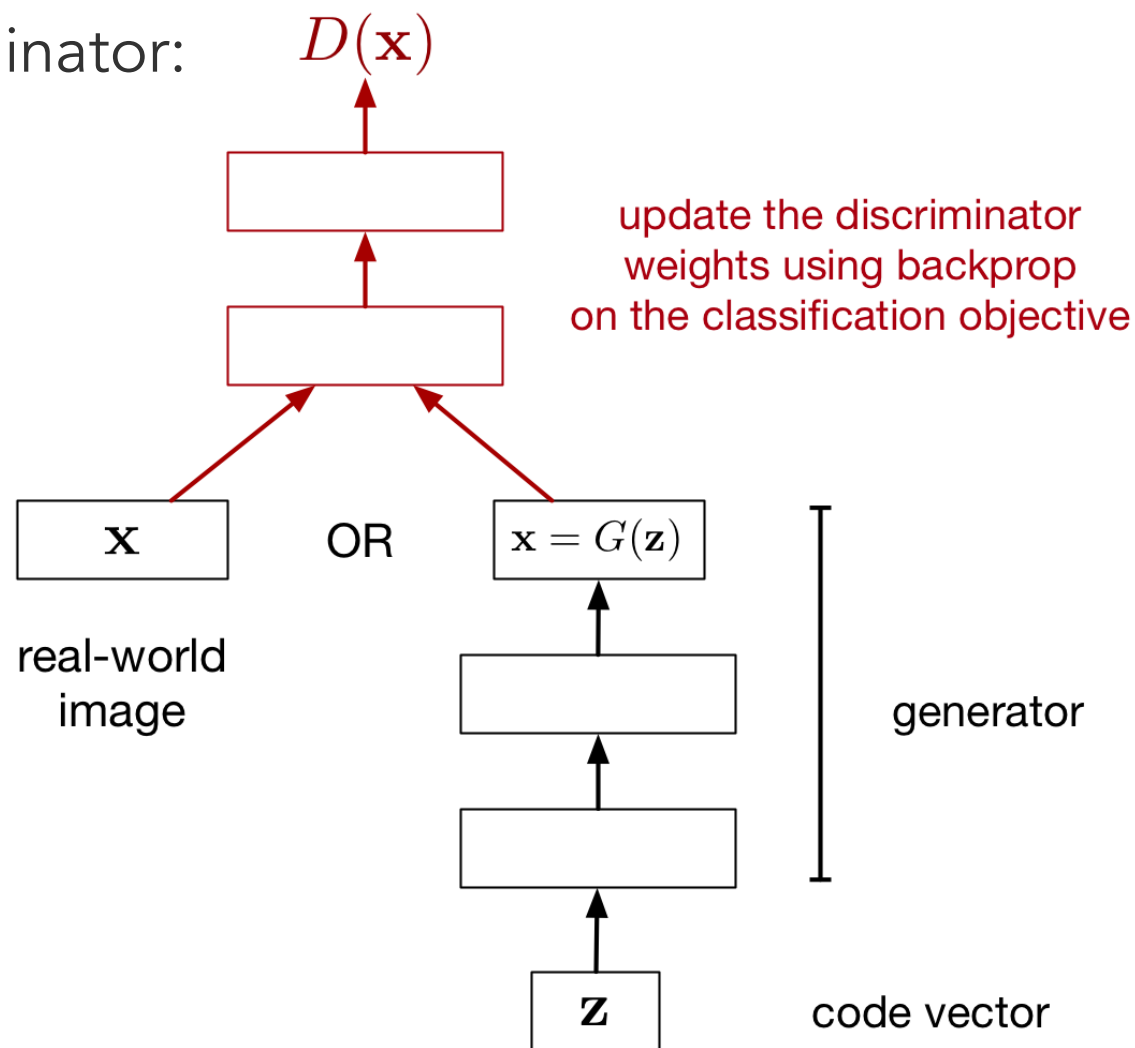


Generative Adversarial Nets (GANs)



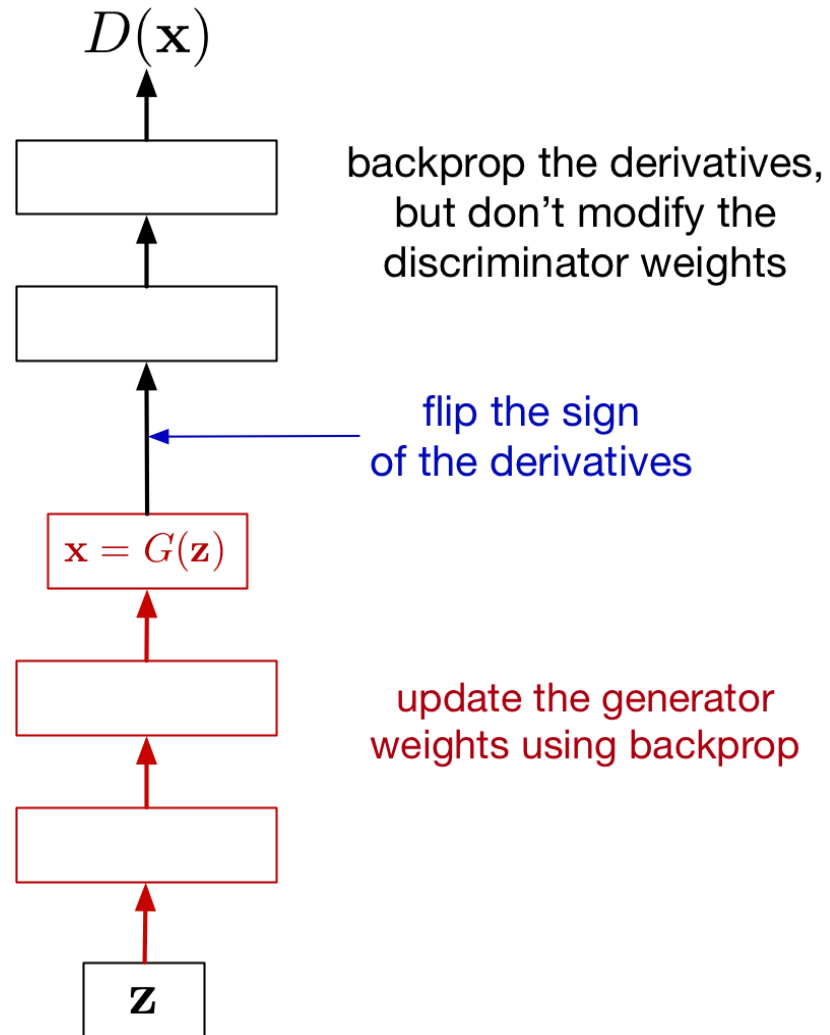
Generative Adversarial Nets (GANs)

Updating the discriminator:




Generative Adversarial Nets (GANs)

Updating the generator:



Generative Adversarial Nets (GANs)

Alternating training of the generator and discriminator:



Optimality of GANs

- Objectives:

$$\max_D \mathcal{L}_D = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim G(\mathbf{z}), \mathbf{z} \sim p(\mathbf{z})} [\log(1 - D(\mathbf{x}))]$$

$$\min_G \mathcal{L}_G = \mathbb{E}_{\mathbf{x} \sim G(\mathbf{z}), \mathbf{z} \sim p(\mathbf{z})} [\log(1 - D(\mathbf{x}))].$$

- Global optimality: $p_g = p_{data}$
- Proof:

Optimality of GANs

Proposition 1. *For G fixed, the optimal discriminator D is*

$$D_G^*(\mathbf{x}) = \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + p_g(\mathbf{x})} \quad (2)$$

Optimality of GANs

Proposition 1. *For G fixed, the optimal discriminator D is*

$$D_G^*(\mathbf{x}) = \frac{p_{data}(\mathbf{x})}{p_{data}(\mathbf{x}) + p_g(\mathbf{x})} \quad (2)$$

Proof. The training criterion for the discriminator D , given any generator G , is to maximize the quantity $V(G, D)$

$$\begin{aligned} V(G, D) &= \int_{\mathbf{x}} p_{data}(\mathbf{x}) \log(D(\mathbf{x})) dx + \int_{\mathbf{z}} p_{\mathbf{z}}(\mathbf{z}) \log(1 - D(g(\mathbf{z}))) dz \\ &= \int_{\mathbf{x}} p_{data}(\mathbf{x}) \log(D(\mathbf{x})) + p_g(\mathbf{x}) \log(1 - D(\mathbf{x})) dx \end{aligned} \quad (3)$$

For any $(a, b) \in \mathbb{R}^2 \setminus \{0, 0\}$, the function $y \rightarrow a \log(y) + b \log(1 - y)$ achieves its maximum in $[0, 1]$ at $\frac{a}{a+b}$.

Optimality of GANs

- The minimax game can now be reformulated as

$$\begin{aligned} C(G) &= \max_D V(G, D) \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - D_G^*(G(\mathbf{z})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [\log(1 - D_G^*(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[\log \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] \end{aligned}$$

Optimality of GANs

- The minimax game can now be reformulated as

$$\begin{aligned} C(G) &= \max_D V(G, D) \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - D_G^*(G(\mathbf{z})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [\log(1 - D_G^*(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[\log \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] \end{aligned}$$

Theorem 1. *The global minimum of the virtual training criterion $C(G)$ is achieved if and only if $p_g = p_{\text{data}}$. At that point, $C(G)$ achieves the value $-\log 4$.*

Optimality of GANs

- The minimax game can now be reformulated as

$$\begin{aligned} C(G) &= \max_D V(G, D) \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_z} [\log(1 - D_G^*(G(\mathbf{z})))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D_G^*(\mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_g} [\log(1 - D_G^*(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \left[\log \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] + \mathbb{E}_{\mathbf{x} \sim p_g} \left[\log \frac{p_g(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_g(\mathbf{x})} \right] \end{aligned}$$

Theorem 1. *The global minimum of the virtual training criterion $C(G)$ is achieved if and only if $p_g = p_{\text{data}}$. At that point, $C(G)$ achieves the value $-\log 4$.*

$$\begin{aligned} C(G) &= -\log(4) + KL \left(p_{\text{data}} \left\| \frac{p_{\text{data}} + p_g}{2} \right\| \right) + KL \left(p_g \left\| \frac{p_{\text{data}} + p_g}{2} \right\| \right) \\ &= -\log(4) + 2 \cdot JSD(p_{\text{data}} \| p_g) \quad \text{Jensen-Shannon Divergence} \end{aligned}$$

A better loss function

- We introduced the minimax cost function for the generator:

$$\mathcal{J}_G = \mathbb{E}_{\mathbf{z}}[\log(1 - D(G(\mathbf{z})))]$$

- One problem with this is **saturation**.
- Here, if the generated sample is really bad, the discriminator's prediction is close to 0, and the generator's cost is flat.

A better loss function: non-saturating GAN


- Original minimax cost:

$$\mathcal{J}_G = \mathbb{E}_{\mathbf{z}}[\log(1 - D(G(\mathbf{z})))]$$

- Modified generator cost:

$$\mathcal{J}_G = \mathbb{E}_{\mathbf{z}}[-\log D(G(\mathbf{z}))]$$

- This fixes the saturation problem.



Wasserstein GAN (WGAN)


- If our data are on a **low-dimensional** manifold of a high dimensional space, the model's manifold and the true data manifold can have a **negligible intersection in practice**

Wasserstein GAN (WGAN)

- If our data are on a **low-dimensional** manifold of a high dimensional space, the model's manifold and the true data manifold can have a **negligible intersection in practice**
- The loss function and gradients may not be continuous and well behaved

Wasserstein GAN (WGAN)

- If our data are on a **low-dimensional** manifold of a high dimensional space, the model's manifold and the true data manifold can have a **negligible intersection in practice**
- The loss function and gradients may not be continuous and well behaved
- The **Wasserstein Distance** is well defined
 - Earth Mover's Distance
 - Minimum transportation cost for making one pile of dirt in the shape of one probability distribution to the shape of the other distribution



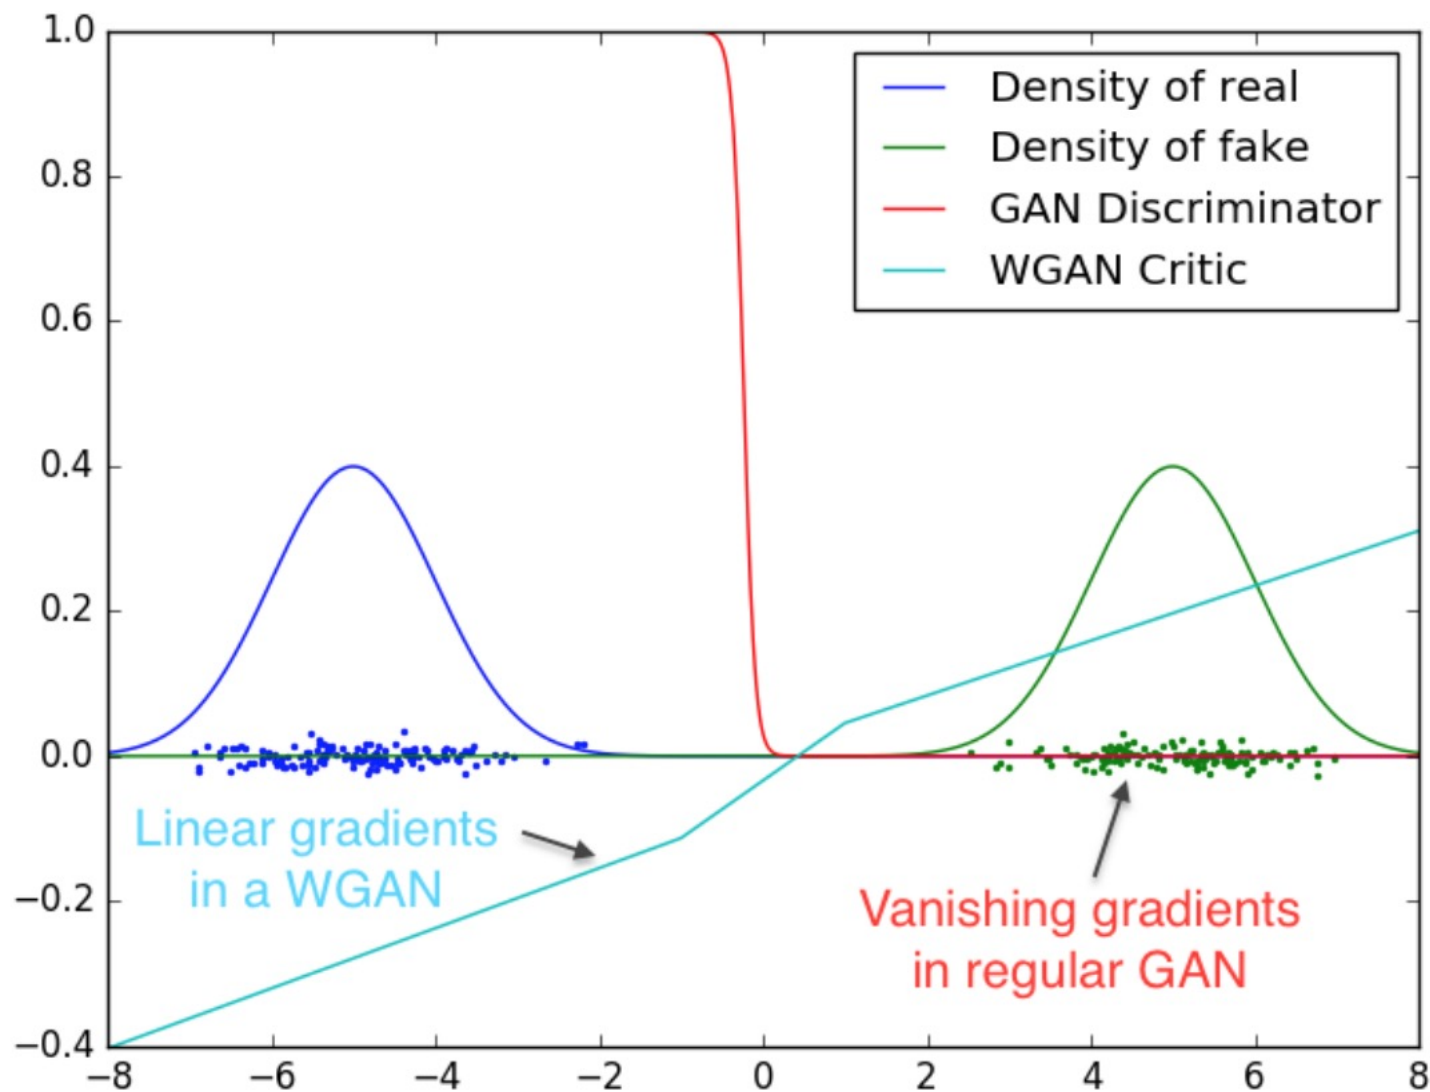
Wasserstein GAN (WGAN)

- Objective

$$W(p_{data}, p_g) = \frac{1}{K} \sup_{\|D\|_L \leq K} \mathbb{E}_{x \sim p_{data}} [D(x)] - \mathbb{E}_{x \sim p_g} [D(x)]$$

- $\|D\|_L \leq K$: K- Lipschitz continuous
- Use gradient-clipping to ensure D has the Lipschitz continuity

WGAN vs Vanilla GAN



Standard Equation and GANs

- Recall SE:

$$\min_{q, \theta} -\alpha \mathbb{H}(q) + \beta \mathbb{D} \left(q(\mathbf{x}), p_{\theta}(\mathbf{x}) \right) - \mathbb{E}_{q(\mathbf{x})} \left[f(\mathbf{x}) \right]$$

- In MLE, f is a fixed function

$$f := f_{data}(\mathbf{x}; \mathcal{D}) = \log \mathbb{E}_{\mathbf{x}^* \sim \mathcal{D}} [\mathbb{1}_{\mathbf{x}^*}(\mathbf{x})]$$

- Intuitively, see f as a similarity metric that measures similarity of sample \mathbf{x} against real data \mathcal{D}
- Instead of the above manually fixed metric, can we learn a metric f_{ϕ} ?

Standard Equation and GANs


- Augment the standard objective to account for ϕ :

$$\min_{\theta} \max_{\phi} \min_q -\alpha \mathbb{H}(q) + \beta \mathbb{D} \left(q(\mathbf{x}), p_{\theta}(\mathbf{x}) \right) - \mathbb{E}_{q(\mathbf{x})} \left[f_{\phi}(\mathbf{x}) \right] + \mathbb{E}_{p_d(\mathbf{x})} \left[f_{\phi}(\mathbf{x}) \right]$$

- Set $\alpha = 0, \beta = 1$. Under mild conditions, the objective recovers:
 - Vanilla GAN [Goodfellow et al., 2014], when \mathbb{D} is JS-divergence and f_{ϕ} is a binary classifier
 - f -GAN [Nowozin et al., 2016], when \mathbb{D} is f -divergence
 - W-GAN [Arjovsky et al., 2017], when \mathbb{D} is Wasserstein distance and f_{ϕ} is a 1-Lipschitz function

Progressive GAN



Low resolution images



Progressive GAN

Low resolution images

add in
additional
layers





Progressive GAN

Low resolution images

add in
additional
layers



High resolution images



BigGAN

BigGAN

- GANs benefit dramatically from **scaling**

BigGAN

- GANs benefit dramatically from **scaling**
- 2x – 4x more parameters
- 8x larger batch size
- Simple architecture changes that improve scalability



BigGAN

- GANs benefit dramatically from **scaling**
- 2x – 4x more parameters
- 8x larger batch size
- Simple architecture changes that improve scalability



BigGAN

- GANs benefit dramatically from **scaling**
- 2x → 4x more parameters
- 8x → 16x more parameters
- Similar quality to human



Outline


- Generative Adversarial Networks (GANs)
 - Vanilla GAN, Wasserstein GAN, Progressive GAN, BigGAN
- Normalizing Flow (NF)
 - Basic Concepts
 - GLOW

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**


Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**



Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**



Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z} \sim p(\mathbf{z})$$

$$\mathbf{x} = f(\mathbf{z})$$

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z} \sim p(\mathbf{z})$$

$$\mathbf{x} = f(\mathbf{z})$$

inference: $\mathbf{z} = f^{-1}(\mathbf{x})$

Transformation function f

-----> • Invertible

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z} \sim p(\mathbf{z})$$

$$\mathbf{x} = f(\mathbf{z})$$

Transformation function f

inference: $\mathbf{z} = f^{-1}(\mathbf{x})$

-----> • Invertible

density:
$$p(\mathbf{x}) = p(\mathbf{z}) \left| \det \frac{d\mathbf{z}}{d\mathbf{x}} \right|$$
$$= p(f^{-1}(\mathbf{x})) \left| \det \frac{df^{-1}}{d\mathbf{x}} \right|$$

$\det \frac{df^{-1}}{d\mathbf{x}}$ -- Jacobian determinant

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z} \sim p(\mathbf{z})$$

$$\mathbf{x} = f(\mathbf{z})$$

inference: $\mathbf{z} = f^{-1}(\mathbf{x})$

density:
$$p(\mathbf{x}) = p(\mathbf{z}) \left| \det \frac{d\mathbf{z}}{d\mathbf{x}} \right|$$
$$= p(f^{-1}(\mathbf{x})) \left| \det \frac{df^{-1}}{d\mathbf{x}} \right|$$

Transformation function f

-----> • Invertible

-----> • Jacobian determinant easy to compute
e.g., choose $df^{-1}/d\mathbf{x}$ to be a triangular matrix

$$\det \frac{df^{-1}}{d\mathbf{x}} \text{ -- Jacobian determinant}$$

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z}_0 \sim p(\mathbf{z}_0)$$

$$\mathbf{x} = \mathbf{z}_K = f_K \circ f_{K-1} \circ \cdots \circ f_1(\mathbf{z}_0)$$

Transformation function f_i

inference: $\mathbf{z}_i = f_i^{-1}(\mathbf{z}_{i-1})$

density: $p(\mathbf{z}_i) = p(\mathbf{z}_{i-1}) \left| \det \frac{d\mathbf{z}_{i-1}}{d\mathbf{z}_i} \right|$

-----> • Invertible

-----> • Jacobian determinant easy to compute
e.g., choose $df_i^{-1}/d\mathbf{z}_i$ to be a triangular matrix

Normalizing Flow (NF)

- Transforms a simple distribution into a complex one by applying a sequence of **transformation functions**

$$\mathbf{z}_0 \sim p(\mathbf{z}_0)$$

$$\mathbf{x} = \mathbf{z}_K = f_K \circ f_{K-1} \circ \dots \circ f_1(\mathbf{z}_0)$$

Transformation function f_i

inference: $\mathbf{z}_i = f_i^{-1}(\mathbf{z}_{i-1})$

----->

- Invertible

density: $p(\mathbf{z}_i) = p(\mathbf{z}_{i-1}) \left| \det \frac{d\mathbf{z}_{i-1}}{d\mathbf{z}_i} \right|$

----->


- Jacobian determinant easy to compute
e.g., choose $df_i^{-1}/d\mathbf{z}_i$ to be a triangular matrix

training: maximizes data log-likelihood

$$\log p(\mathbf{x}) = \log p(\mathbf{z}_0) + \sum_{i=1}^K \log \left| \det \frac{d\mathbf{z}_{i-1}}{d\mathbf{z}_i} \right|$$

GLOW


- [Kingma and Dhariwal., 2018]



One step of flow in the Glow model

GLOW

- [Kingma and Dhariwal., 2018]



One step of flow in the Glow model

Key Takeaways


- GANs:
 - Implicit generative model
 - Minimax formulation
 - non-saturating GANs
 - WGAN
- Normalizing Flow
 - Transforms a simple distribution into a complex one by applying a sequence of transformation functions

Questions?

Backups


Deep generative models

- Define probabilistic distributions over a set of variables
- "Deep" means multiple layers of hidden variables!



Early forms of deep generative models


- Hierarchical Bayesian models
 - Sigmoid belief nets [Neal 1992]



$$p\left(x_{kn} = 1 \mid \boldsymbol{\theta}_k, \mathbf{z}_n^{(1)}\right) = \sigma\left(\boldsymbol{\theta}_k^T \mathbf{z}_n^{(1)}\right)$$
$$p\left(z_{in}^{(1)} = 1 \mid \boldsymbol{\theta}_i, \mathbf{z}_n^{(2)}\right) = \sigma\left(\boldsymbol{\theta}_i^T \mathbf{z}_n^{(2)}\right)$$

Early forms of deep generative models

- Hierarchical Bayesian models
 - Sigmoid belief nets [Neal 1992]
- Neural network models
 - Helmholtz machines [Dayan et al., 1995]



[Dayan et al. 1995]

Early forms of deep generative models

- Hierarchical Bayesian models
 - Sigmoid belief nets [Neal 1992]
- Neural network models
 - Helmholtz machines [Dayan et al.,1995]
 - Predictability minimization [Schmidhuber 1995]




Figure courtesy: Schmidhuber 1996

Early forms of deep generative models

- Training of DGMs via an EM style framework

- Sampling / data augmentation

$$\mathbf{z} = \{\mathbf{z}_1, \mathbf{z}_2\}$$

$$\mathbf{z}_1^{new} \sim p(\mathbf{z}_1 | \mathbf{z}_2, \mathbf{x})$$

$$\mathbf{z}_2^{new} \sim p(\mathbf{z}_2 | \mathbf{z}_1^{new}, \mathbf{x})$$

- Variational inference

$$\log p(\mathbf{x}) \geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}, \mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}) || p(\mathbf{z})) := \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x})$$

$$\max_{\boldsymbol{\theta}, \boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{x})$$


- Wake sleep

$$\text{Wake: } \min_{\boldsymbol{\theta}} \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})]$$

$$\text{Sleep: } \min_{\boldsymbol{\phi}} \mathbb{E}_{p_\theta(\mathbf{x}|\mathbf{z})}[\log q_\phi(\mathbf{z}|\mathbf{x})]$$

Resurgence of deep generative models


- Restricted Boltzmann machines (RBMs) [Smolensky, 1986]
 - Building blocks of deep probabilistic models




Resurgence of deep generative models

- Restricted Boltzmann machines (RBMs) [Smolensky, 1986]
 - Building blocks of deep probabilistic models
- Deep belief networks (DBNs) [Hinton et al., 2006]
 - Hybrid graphical model
 - Inference in DBNs is problematic due to explaining away
- Deep Boltzmann Machines (DBMs) [Salakhutdinov & Hinton, 2009]
 - Undirected model

Deep Belief Network



Deep Boltzmann Machine



Resurgence of deep generative models

- Variational autoencoders (VAEs) [Kingma & Welling, 2014]
/ Neural Variational Inference and Learning (NVIL) [Mnih & Gregor, 2014]


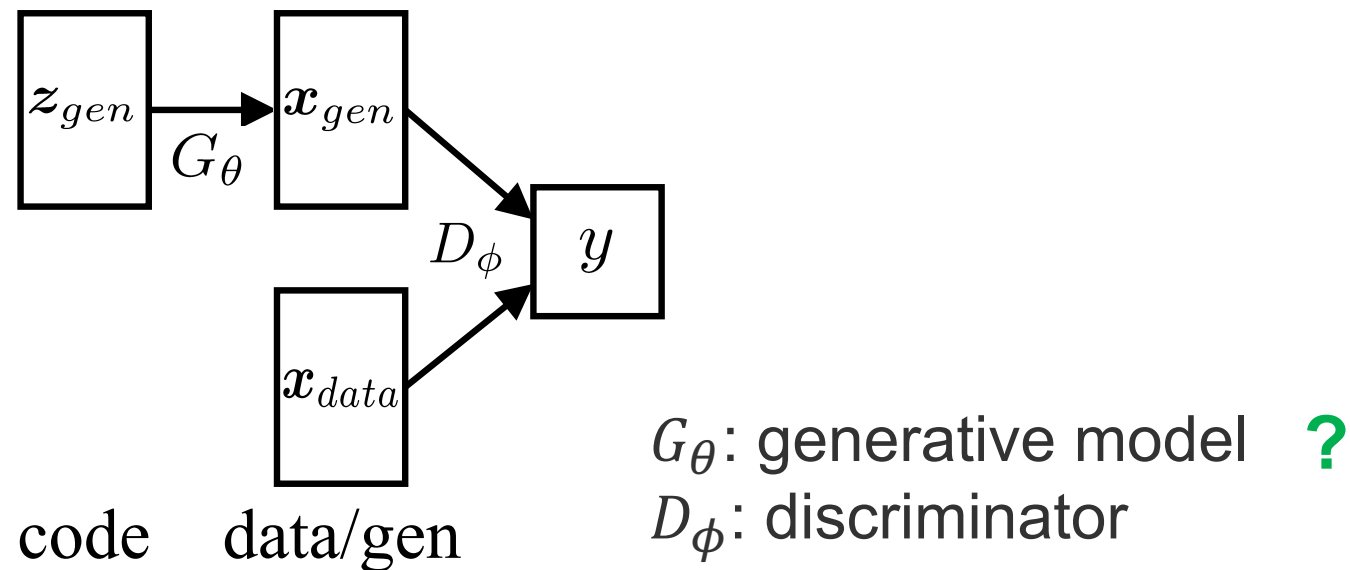


Figure courtesy: Kingma & Welling, 2014

Resurgence of deep generative models

- Variational autoencoders (VAEs) [Kingma & Welling, 2014]
/ Neural Variational Inference and Learning (NVIL) [Mnih & Gregor, 2014]
- Generative adversarial networks (GANs) [Goodfellow et al., 2014]



Resurgence of deep generative models

- Variational autoencoders (VAEs) [Kingma & Welling, 2014]
/ Neural Variational Inference and Learning (NVIL) [Mnih & Gregor, 2014]
- Generative adversarial networks (GANs) [Goodfellow et al., 2014]
- Generative moment matching networks (GMMNs) [Li et al., 2015; Dziugaite et al., 2015]

Resurgence of deep generative models

- Variational autoencoders (VAEs) [Kingma & Welling, 2014]
/ Neural Variational Inference and Learning (NVIL) [Mnih & Gregor, 2014]
- Generative adversarial networks (GANs) [Goodfellow et al., 2014]
- Generative moment matching networks (GMMNs) [Li et al., 2015; Dziugaite et al., 2015]
- Autoregressive neural networks

