

# Deep Generative Models


## Lecture 9: Generative Adversarial Networks

---


Aditya Grover

UCLA

# Recap



$$\begin{aligned} \mathbf{x}^{(j)} &\sim P_{\text{data}} \\ j &= 1, 2, \dots, m \end{aligned}$$



- Model families

- Autoregressive Models:  $p_{\theta}(\mathbf{x}) = \prod_{i=1}^n p_{\theta}(x_i | \mathbf{x}_{<i})$
- Variational Autoencoders:  $p_{\theta}(\mathbf{x}) = \int p_{\theta}(\mathbf{x}, \mathbf{z}) d\mathbf{z}$
- Normalizing Flow Models:

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

- All the above families are based on maximizing likelihoods (or approximations)

# Why maximum likelihood?

$$\hat{\theta} = \arg \max_{\theta} \sum_{i=1}^n \log p_{\theta}(\mathbf{x}_i), \quad \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \sim p_{\text{data}}(\mathbf{x})$$

- **Optimal statistical efficiency.**
  - Assume sufficient model capacity, such that there exists a unique  $\theta^* \in \mathcal{M}$  that satisfies  $p_{\theta^*} = p_{\text{data}}$ .
  - The convergence of  $\hat{\theta}$  to  $\theta^*$  when  $n \rightarrow \infty$  is the “fastest” among all statistical methods when using maximum likelihood training.
- **Higher likelihood = better lossless compression.**
- Is the likelihood a good indicator of the quality of samples generated by the model?

**Case 1:** Optimal generative model will give best **sample quality** and highest test **log-likelihood**

For imperfect models, achieving high log-likelihoods might not always imply good sample quality, and vice-versa (Theis et al., 2016)

## Towards likelihood-free learning

**Case 2:** Great test log-likelihoods, poor samples. E.g., For a discrete noise mixture model  $p_{\theta}(\mathbf{x}) = 0.01p_{\text{data}}(\mathbf{x}) + 0.99p_{\text{noise}}(\mathbf{x})$

- 99% of the samples are just noise
- Taking logs, we get a lower bound

$$\begin{aligned}\log p_{\theta}(\mathbf{x}) &= \log[0.01p_{\text{data}}(\mathbf{x}) + 0.99p_{\text{noise}}(\mathbf{x})] \\ &\geq \log 0.01p_{\text{data}}(\mathbf{x}) = \log p_{\text{data}}(\mathbf{x}) - \log 100\end{aligned}$$

- For expected likelihoods, we know that
  - Lower bound

$$E_{p_{\text{data}}}[\log p_{\theta}(\mathbf{x})] \geq E_{p_{\text{data}}}[\log p_{\text{data}}(\mathbf{x})] - \log 100$$

- Upper bound (via non-negativity of KL)

$$E_{p_{\text{data}}}[\log p_{\text{data}}(\mathbf{x})] \geq E_{p_{\text{data}}}[\log p_{\theta}(\mathbf{x})]$$

- As we increase the dimension of  $\mathbf{x}$ , absolute value of  $\log p_{\text{data}}(\mathbf{x})$  increases proportionally but  $\log 100$  remains constant. Hence,  $E_{p_{\text{data}}}[\log p_{\theta}(\mathbf{x})] \approx E_{p_{\text{data}}}[\log p_{\text{data}}(\mathbf{x})]$  in very high dimensions


**Case 3:** Great samples, poor test log-likelihoods. E.g., Memorizing training set

- Samples look exactly like the training set (cannot do better!)
- Test set will have zero probability assigned (cannot do worse!)

### Takeaways:

- The above cases suggest that it might be useful to disentangle likelihoods and samples
- **Likelihood-free learning** consider objectives that do not depend directly on a likelihood function

## Comparing distributions via samples



$$S_1 = \{\mathbf{x} \sim P\}$$

vs.



$$S_2 = \{\mathbf{x} \sim Q\}$$


Given a finite set of samples from two distributions  $S_1 = \{\mathbf{x} \sim P\}$  and  $S_2 = \{\mathbf{x} \sim Q\}$ , how can we tell if these samples are from the same distribution? (i.e.,  $P = Q$ ?)

## Two-sample tests


- Given  $S_1 = \{\mathbf{x} \sim P\}$  and  $S_2 = \{\mathbf{x} \sim Q\}$ , a **two-sample test** considers the following hypotheses
  - Null hypothesis  $H_0: P = Q$
  - Alternative hypothesis  $H_1: P \neq Q$
- Test statistic  $T$  compares  $S_1$  and  $S_2$  e.g., difference in means, variances of the two sets of samples
- If  $T$  is larger than a threshold  $\alpha$ , then reject  $H_0$  otherwise we say  $H_0$  is consistent with observation.
- **Key observation:** Test statistic is **likelihood-free** since it does not involve the densities  $P$  or  $Q$  (only samples)



# Generative modeling and two-sample tests




$$x^{(j)} \sim P_{\text{data}} \\ j = 1, 2, \dots, m$$



- A priori we assume direct access to  $S_1 = \mathcal{D} = \{\mathbf{x} \sim p_{\text{data}}\}$
- In addition, we have a model distribution  $p_{\theta}$
- Assume that the model distribution permits efficient sampling (e.g., directed models). Let  $S_2 = \{\mathbf{x} \sim p_{\theta}\}$
- **Alternative notion of distance between distributions:**  
Train the generative model to minimize a two-sample test objective between  $S_1$  and  $S_2$


# Two-Sample Test via a Discriminator

- Finding a two-sample test objective in high dimensions is hard



- In the generative model setup, we know that  $S_1$  and  $S_2$  come from different distributions  $p_{\text{data}}$  and  $p_{\theta}$  respectively
- Key idea:** Learn a statistic that **maximizes** a suitable notion of distance between the two sets of samples  $S_1$  and  $S_2$

## Two-Sample Test via a Discriminator



- **Two-Sample Test via a Discriminator**

- Any function (e.g., neural network) which tries to distinguish “real” samples from the dataset and “fake” samples generated from the model
- Maximizes the two-sample test objective (in support of the alternative hypothesis  $p_{\text{data}} \neq p_\theta$ )

## Two-Sample Test via a Discriminator

- **Training objective for discriminator:**


$$\max_D V(G, D) = E_{\mathbf{x} \sim p_{\text{data}}} [\log D(\mathbf{x})] + E_{\mathbf{x} \sim p_G} [\log(1 - D(\mathbf{x}))]$$

- For a fixed generator  $G$ , the discriminator is performing binary classification with the cross entropy objective
  - Assign probability 1 to true data points  $\mathbf{x} \sim p_{\text{data}}$
  - Assign probability 0 to fake samples  $\mathbf{x} \sim p_G$
- Optimal discriminator

$$D_G^*(\mathbf{x}) = \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_G(\mathbf{x})}$$

# Generative Adversarial Networks

- A two player minimax game between a **generator** and a **discriminator**



- **Generator**
  - Directed, latent variable model with a deterministic mapping between  $z$  and  $x$  given by  $G_\theta$
  - Minimizes a two-sample test objective (in support of the null hypothesis  $p_{\text{data}} = p_\theta$ )

## Example of GAN objective

- **Training objective for generator:**

$$\min_G \max_D V(G, D) = E_{\mathbf{x} \sim p_{\text{data}}} [\log D(\mathbf{x})] + E_{\mathbf{x} \sim p_G} [\log(1 - D(\mathbf{x}))]$$

- For the optimal discriminator  $D_G^*(\cdot)$ , we have

$$\begin{aligned} & V(G, D_G^*(\mathbf{x})) \\ &= E_{\mathbf{x} \sim p_{\text{data}}} \left[ \log \frac{p_{\text{data}}(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_G(\mathbf{x})} \right] + E_{\mathbf{x} \sim p_G} \left[ \log \frac{p_G(\mathbf{x})}{p_{\text{data}}(\mathbf{x}) + p_G(\mathbf{x})} \right] \\ &= E_{\mathbf{x} \sim p_{\text{data}}} \left[ \log \frac{p_{\text{data}}(\mathbf{x})}{\frac{p_{\text{data}}(\mathbf{x}) + p_G(\mathbf{x})}{2}} \right] + E_{\mathbf{x} \sim p_G} \left[ \log \frac{p_G(\mathbf{x})}{\frac{p_{\text{data}}(\mathbf{x}) + p_G(\mathbf{x})}{2}} \right] - \log 4 \\ &= \underbrace{D_{KL} \left[ p_{\text{data}}, \frac{p_{\text{data}} + p_G}{2} \right] + D_{KL} \left[ p_G, \frac{p_{\text{data}} + p_G}{2} \right]}_{2 \times \text{Jensen-Shannon Divergence (JSD)}} - \log 4 \\ &= 2D_{JSD}[p_{\text{data}}, p_G] - \log 4 \end{aligned}$$

# Jenson-Shannon Divergence

- Also called as the symmetric KL divergence

$$D_{JSD}[p, q] = \frac{1}{2} \left( D_{KL} \left[ p, \frac{p+q}{2} \right] + D_{KL} \left[ q, \frac{p+q}{2} \right] \right)$$

- Properties

- $D_{JSD}[p, q] \geq 0$
- $D_{JSD}[p, q] = 0$  iff  $p = q$
- $D_{JSD}[p, q] = D_{JSD}[q, p]$
- $\sqrt{D_{JSD}[p, q]}$  satisfies triangle inequality  $\rightarrow$  Jenson-Shannon Distance

- Optimal generator for the JSD/Negative Cross Entropy GAN

$$p_G = p_{\text{data}}$$

- For the optimal discriminator  $D_{G^*}^*(\cdot)$  and generator  $G^*(\cdot)$ , we have

# The GAN training algorithm

- Sample minibatch of  $m$  training points  $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(m)}$  from  $\mathcal{D}$
- Sample minibatch of  $m$  noise vectors  $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \dots, \mathbf{z}^{(m)}$  from  $p_z$
- Update the discriminator parameters  $\phi$  by stochastic gradient **ascent**

$$\nabla_{\phi} V(G_{\theta}, D_{\phi}) = \frac{1}{m} \nabla_{\phi} \sum_{i=1}^m [\log D_{\phi}(\mathbf{x}^{(i)}) + \log(1 - D_{\phi}(G_{\theta}(\mathbf{z}^{(i)})))]$$

- Update the generator parameters  $\theta$  by stochastic gradient **descent**


$$\nabla_{\theta} V(G_{\theta}, D_{\phi}) = \frac{1}{m} \nabla_{\theta} \sum_{i=1}^m \log(1 - D_{\phi}(G_{\theta}(\mathbf{z}^{(i)})))$$

- Repeat for fixed number of epochs



# Alternating optimization in GANs

$$\min_{\theta} \max_{\phi} V(G_{\theta}, D_{\phi}) = E_{\mathbf{x} \sim p_{\text{data}}} [\log D_{\phi}(\mathbf{x})] + E_{\mathbf{z} \sim p(\mathbf{z})} [\log(1 - D_{\phi}(G_{\theta}(\mathbf{z})))]$$




## Frontiers in GAN research



- GANs have been successfully applied to several domains and tasks
- However, working with GANs can be very challenging in practice
  - Unstable optimization
  - Mode collapse
  - Evaluation
- Many bag of tricks applied to train GANs successfully

# Optimization challenges

- **Theorem (informal):** If the generator updates are made in function space and discriminator is optimal at every step, then the generator is guaranteed to converge to the data distribution
- **Unrealistic assumptions!**
- In practice, the generator and discriminator loss keeps oscillating during GAN training




Source: Mirantha Jayatilaka

- No robust stopping criteria in practice (unlike MLE)


# Mode Collapse

- GANs are notorious for suffering from **mode collapse**
- Intuitively, this refers to the phenomena where the generator of a GAN collapses to one or few samples (dubbed as “modes”)




Arjovsky et al., 2017

# Mode Collapse




- True distribution is a mixture of Gaussians



- The generator distribution keeps oscillating between different modes

# Mode Collapse



Source: Metz et al., 2017

- Fixes to mode collapse are mostly empirically driven: alternative architectures, alternative GAN loss, adding regularization terms, etc.
- <https://github.com/soumith/ganhacks>  
How to Train a GAN? Tips and tricks to make GANs work by Soumith Chintala

## Beauty lies in the eyes of the discriminator



Source: Robbie Barrat, Obvious

GAN generated art auctioned at Christie's.

**Expected Price:** \$7,000 – \$10,000

**True Price:** \$432,500