

Generative Models: Part I

Part 9: Generative Models

- Variational Autoencoders
- Generative Adversarial Networks

Jonas Glombitza, Martin Erdmann

RWTH Aachen

2nd Terascale Machine Learning School |



Outline

I. Variational Autoencoders

II. Introduction to Generative Adversarial Networks (GANs)

> Adversarial frameworks

III. Tutorial: Implementation of GANs

BREAK

IV.Latest developments & advanced techniques

- > Wasserstein GANs
- V.Application in physics research
 - Simulation acceleration
 - Style transfer (domain adaption)

VI.Tutorial: Implementation of Wasserstein GANs

Tutorial on Generative Models Glombitza | RWTH Aachen | 03/05/21 | 2nd Terascale Machine Learning School





Feel free to ask questions during the seminar! Just "raise" your hand...



Supervised and Unsupervised Learning

- Generative Models

3

- Variational Autoencoders
- Generative Adversarial Networks



Supervised Learning

- Situation
 - Large labeled data set (pair of input $\, x \,$ and output y)
- Typical Task:
 - Learn function to map input to specific output
 - Train model to predict the associated label
 - Achieve best generalization performance
 - Infer *conditional* probability density $p(\mathbf{x}|\mathbf{y})$



Tutorial on Generative Models



Unsupervised Learning



- Typical Situation: non labeled data set
- Tasks:
 - Learn (low dimensional) data encodings \rightarrow *autoencoders*
 - Estimate underlying probability density \rightarrow generative models
 - Clustering, anomaly detection find (non-) similar samples
- Infer *a priori* probability density p(x)
- Models typical trained without label information
 - Contrast: semi-supervised learning



Recap: Autoencoders





- Reconstruction of input data (approximation of identity function)
- Learning interesting representation (constraints to hidden layer)
- Objective function:

$$\mathcal{L}(\mathbf{x}, \mathbf{\hat{x}}) = \frac{1}{N} \sum_{i=1}^{N} (\mathbf{\hat{x}}_i - \mathbf{x}_i)^2$$

• Deep autoencoders often show underfitting \rightarrow use shortcuts!

Tutorial on Generative Models

6

Generative Models



Approximate data distribution P_r with another distribution P_{θ} θ = distribution parameters



Variational Autoencoder





- Learned representation is not an arbitrary function
 - > Impose prior distribution on the hidden (low dimensional) representation
- Trained decoder part can be used as *generator* (sample from prior distribution)
- Objective function = reconstruction error + divergence of hidden representation from a prior

$$\mathcal{L}(\mathbf{x}, \mathbf{\hat{x}}, \mathbf{z}) = \mathcal{L}_{recon.}(\mathbf{x}, \mathbf{\hat{x}}) + \mathcal{D}_{KL}[q(\mathbf{z}|\mathbf{x})|p(\mathbf{z})]$$
$$= \frac{1}{N} \sum_{\mathbf{x}} \left[||\mathbf{x} - \mathbf{\hat{x}}||^2 + \sum_{\mathbf{z}} q(\mathbf{z}|\mathbf{x}) \log \frac{p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})} \right]$$

Kullback-Leibler Divergence: \mathcal{D}_{KL} Measure of information loss if $p(\mathbf{z})$ is used instead of $p(\mathbf{z}|\mathbf{x})$

Tutorial on Generative Models

8

Variational Autoencoder





Gaussian prior: $z \sim \mathcal{N}(0, 1)$

- Encoder $q(\mathbf{z}|\mathbf{x})$ learns latent parameters $\mu_{\mathbf{z}}(\mathbf{x}), \sigma_{\mathbf{z}}(\mathbf{x})$ of Gaussian distribution
- Re-parametrization $z=\mu_{\mathbf{z}}(\mathbf{x})+\sigma_{\mathbf{z}}(\mathbf{x})\epsilon$ with $\epsilon\in\mathcal{N}(0,1)$
- Example: 2D Gaussian
 - Only two independently normal distributed parameters in hidden layer for each input

•
$$\mathcal{D}_{KL}[\mathcal{N}(\mu_{z}(\mathbf{x}), \sigma_{z}(\mathbf{x})) | \mathcal{N}(0, 1)]$$

$$= \frac{1}{N} \sum_{\mathbf{x}} \int_{\mathbf{z}} dz \, \mathcal{N}(\mu_{\mathbf{z}}, \sigma_{\mathbf{z}}) \log \frac{\mathcal{N}(0, 1)}{\mathcal{N}(\mu_{\mathbf{z}}(\mathbf{x}), \sigma_{\mathbf{z}}(\mathbf{x}))}$$

$$= \frac{1}{N} \sum_{\mathbf{x}} \frac{1}{2} (1 + \log \sigma_{\mathbf{z}}^{2}(\mathbf{x}) - \mu_{\mathbf{z}}^{2}(\mathbf{x}) - \sigma_{\mathbf{z}}^{2}(\mathbf{x}))$$

Tutorial on Generative Models

9

Variational Autoencoder



Objective: $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}, \mathbf{z}) = MSE(\mathbf{x}, \hat{\mathbf{x}}) + \mathcal{D}_{KL}[q(\mathbf{z}|\mathbf{x})|p(\mathbf{z})]$

- Mean-squared-error: \rightarrow How accurate input can be reconstructed
- KL-divergence: \rightarrow How close the latent variables match (unit Gaussian)
- Allows walk in latent space



VAE trained on MNIST using 2D Gaussian prior

Tutorial on Generative Models

10



Latent space of VAE trained on MNIST

- Samples of Variational Autoencoders often look noisy
 - Gaussian prior not always best choice / can use arbitrary prior distribution
 - Gaussian distributions can not capture all modes of the data
 - Mean-squared-error loss very inflexible

11

> Try adversarial approach \rightarrow Only train a generator



Introduction to GANs

Generative models

Basics of GANs



Generative Adversarial Networks - GANs



Künstliche Intelligenz

13

Auktionshaus versteigert erstmals KI-Gemälde

Kein Maler sondern ein Computer-Algorithmus hat das Porträt "Edmond de Belamy" erschaffen. Beim Auktionshaus Christie's zahlte ein Interessent dafür knapp 400.000 Euro.

26. Oktober 2018, 9:00 Uhr / Quelle: ZEIT ONLINE, AFP, fo / 75 Kommentare



Der verschwommene Druck "Edmond de Belamy" zeigt einen Mann in dunkler Kutte mit weißem Kragen, der an einen französischen Geistlichen erinnert. © Christie's/dpa



https://ailab.criteo.com/iclr-2019-stats-trends-and-best-papers/

How to train a Generator



I. Objective: learn to generate new samples following P_{θ}

II.Learn a function that transform a distribution $p(\mathbf{z})$ into P_{θ} using a generator G_{θ} $\mathbf{z} \in Z \rightarrow$ latent space

III.Generator G_{θ} is implemented as neural network with weights θ



Generative Adversarial Networks

I. Hard to formulate a supervised training loss

II.Use unsupervised training to train the generator

- > Objective: $P_{\theta} \approx P_r$
- Measure: given by second neural network
- → Generated samples of generator should be similar to real samples after training
- without reproducing training data

\rightarrow Adversarial approach:

Train 2 networks adversarial (against each other)





Generative Adversarial Networks

I. Generator

- Try to generate realistic samples
- II. Discriminator
 - Try to discriminate between fakes and realistic images
 - > Evaluate if $P_{\theta} \approx P_r$
- III.Discriminator returns probability if generated sample is real



"GANs is the most interesting idea in the last ten years in machine learning." - Y. LeCun



Train Discriminator

Train Generator



Tutorial on Generative Models

17



Train Discriminator

Train Generator



Tutorial on Generative Models

18

Train the Discriminator





Tutorial on Generative Models

19

Train the Generator

I. Optimal discriminator is freezed

$$Loss = \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [log(1 - D(G_\theta(\mathbf{z})))]$$

- II. **Minimize** loss: \rightarrow maximize binary cross entropy
 - \succ Tuning the generator weights θ
 - Discriminator should fail to discriminate

III. Best case: coin flipping

$$D(G(\mathbf{z})) = \frac{1}{2}$$
 $D(\mathbf{x}) =$

2





Tutorial on Generative Models

GAN Training



$$\min_{G} \max_{D} L(D,G) = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [log(D(\mathbf{x}))] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [log(1 - D(G(\mathbf{z})))]$$

Training 2 networks at the same time is challenging Losses of discriminator and generator are highly dependent

- I. Train generator and discriminator alternating
 - Min/Max game
 - Sum of both players is zero



- II.Finding Nash equilibrium is hard
 - Discriminator and generator need to have same quality
 - Minimize Jensen-Shannon divergence (assume optimal discriminator)



Optimal Evolution of GAN Training



 \rightarrow G generates samples which are more likely identified as data



Network Design



- Discriminator / classical DCNN for classification
 - Use sigmoid (1 output node) / softmax (2 output nodes) after last layer
- DCGANs (Deep Convolutional GANs) show improved stability
- Use Deep Convolutional generator and discriminator:

 Use batch normalization
 Remove fully connected hidden layers
 Use ReLU in the generator
 Use LeakyReLU in the discriminator
 Use transposed convolutions



Tutorial on Generative Models

23

Deep Convolutional GAN (DC-GAN)



- Topology of the generator: Ι.
 - Decrease feature space ≻
 - Increase spatial extent ≻
- II. Supports a simple structured latent space
- III. Use transposed convolutions + striding
- \rightarrow Shows improved training stability





A. Radford, L. Metz, S. Chintala - https://arxiv.org/abs/1511.06434

Implementation: Adversarial Training



Generate fake samples \tilde{x} similar to real samples x

I. Train discriminator

25

- > Sample noise z, feed it into the generator to generate fake samples $G(z) = \tilde{x}$
- \blacktriangleright Train discriminator to classify fake \tilde{x} and real samples x

II. Train generator using the discriminator feedback

- > Freeze parameters of discriminator
- \blacktriangleright Generate fake samples $G(z)=\tilde{x}$ and pass it to the discriminator

III.Unfreeze parameters of the discriminator

freeze / unfreeze layers in a model
for layer in model.layers:
 layer.trainable = True # False
 # update parameters for a single batch
 loss = model.train_on_batch(x, y)
Tutorial on Generative Models
Glombitza | RWTH Aachen | 03/05/21 | 2nd Terascale Machine Learning School

Tutorial on Generative Models Glombitza | RWTH Aachen | 03/05/21 | 2nd Terascale Machine Learning School

Generative Adversarial Networks

• Wrong global structure





Wrong body parts







ImageNet



Training Data

Samples

Manifold Hypothesis



Idea: Manifolds of meaningful pictures are highly concentrated with very little volume and embedded in a very high dimensional space

- I. Generation of images is a very challenging task
- II.Correlations / probability dimension are high dimensional

Example: Try to generate images randomly:



"To deal with a 14-dimensional space, visualize a 3-D space and say 'fourteen' to yourself very loudly. Everyone does it." - G. Hinton

Tutorial on Generative Models

27

Evolution of GANs - 2016



goldfinch

daisy

Monarch butterfly Odena, Olah, Shlens - arXiv:1610.09585

Evolution of GANs

Zhang, Goodfellow, Metaxas, Odena - arXiv:1805.08318

GANs not perfect...

Brock, Donahue, Simonyan - https://arxiv.org/abs/1809.11096

Karras, Alla, Laine, Lehtinen - arXiv:1710.10196

 \rightarrow Much better images in the next lecture...

Conditioning of GANs – Semi Supervised

- Constrain generator to learn conditional probability distribution
 - Reduce complexity of latent space, allow for interpretations
- > Feed generator and discriminator additional information (e.g. class labels: dog)
 - Force generated samples show specific characteristics (label dependencies)

Tutorial on Generative Models

Conditioning of GANs

InfoGAN

TH **MMM** (a) Rotation E E (b) Width

Chen et al. 2016

Tutorial on Generative Models

32 Glombitza | RWTH Aachen | 03/05/21 | 2nd Terascale Machine Learning School

Conditional image synthesis

 \succ Field of generative model is growing very fast

VAE

GAN

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

Next Lecture: Advanced Techniques

• Which picture is generated, which picture is part of CELEB A data set?

Summary

Generative Models

Generation of new samples using approximation of underlying data distribution

Variational Autoencoder

- Hidden representation follows low dimensional arbitrary prior distribution
- Trained decoder part can be used as generator to produce new samples

Generative Adversarial Networks

- Hand-coded loss is replaced by discriminator (tries to discriminate between fake samples and real samples)
- Adversarial training: generator and discriminator trained against each other
- Generator tries to fool discriminator
- Use conditioning to create prior on latent space

Tutorial on Generative Models

References & Further Reading

- Dr. David Walz, DLiPR, Summerterm 2017
- Kingma, Welling: Variational Autoencoder https://arxiv.org/abs/1312.6114
- Makzhani et al.: Adversarial Autoencoders https://arxiv.org/abs/1511.05644
- Goodfellow et al.: Generative Adversarial Networks https://arxiv.org/abs/1406.2661
- Odena et al.: AC-GAN https://arxiv.org/abs/1610.09585
- Radford et al.: DCGAN https://arxiv.org/abs/1511.06434
- Zhang et al.: SAGAN https://arxiv.org/pdf/1805.08318.pdf

Tutorial on Generative Models

Tutorial

Open jupyter notebooks in google colab: You can find the repository at: https://github.com/Napoleongurke/tutorial_generative_models

Open PART I: Vanilla_GAN.ipynb

Arithmetic in Latent Space

Idea: Discover structure of the latent space

• Can we do arithmetic with latent vectors of generated samples which represent different characteristics?

 $\mathbf{z}_{king} - \mathbf{z}_{man} + \mathbf{z}_{women} = \mathbf{z}_{queen}?$

• Average over several samples to get representation vector

Tutorial on Generative Models