


Module 9: Unsupervised learning

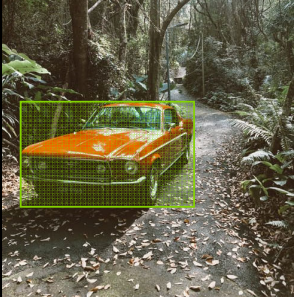
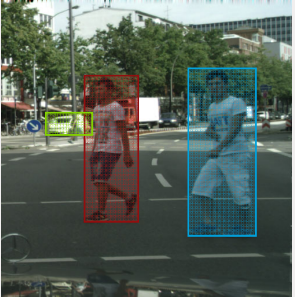
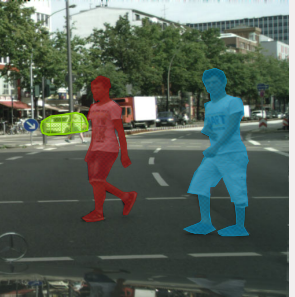
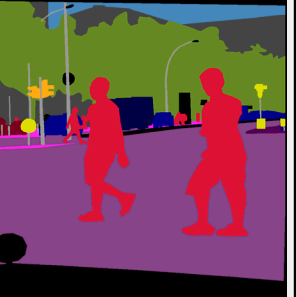
5LSM0: Convolutional neural networks for computer vision


Fons van der Sommen

Electrical Engineering / VCA research group




Last time

Classification+localization	Object detection	Instance segmentation	Semantic segmentation
 <p style="text-align: center; color: green; font-weight: bold;">Car</p>	 <p style="display: flex; justify-content: center; gap: 10px;"> Pedestrian Car </p> <p style="text-align: center; color: blue; font-weight: bold;">Pedestrian</p>	 <p style="display: flex; justify-content: center; gap: 10px;"> Pedestrian Car </p> <p style="text-align: center; color: blue; font-weight: bold;">Pedestrian</p>	 <p style="display: flex; justify-content: center; gap: 5px;"> Truck Road Vegetation </p> <p style="display: flex; justify-content: center; gap: 5px;"> Pedestrian Sidewalk Car </p>
Single object	Multiple objects		No objects, just pixels



2 5LSM0 Module 9: Unsupervised learning



Last time

Unpooling

Input: 4×4

Classification + localization

Output vector: 1000

0.07	Boat	Class scores	→ Softmax loss <i>Given: correct label</i>
0.81	Car		
0.02	Tree		
...	...	Box coordinates	→ L2 loss <i>Given: correct position</i>
(x, y, w, h)			

Transpose convolution

Input: 4×4

FCN for semantic segmentation

- Downsampling: Pooling, Strided convolution
- Upsampling: Unpooling, Strided transpose convolution

3 5LSM0 Module 9: Unsupervised learning

Last time

objects?

Fish (x,y,w,h)
Fish (x,y,w,h)
Fish (x,y,w,h)
...
Fish (x,y,w,h)

R-CNN

Fast R-CNN

Faster R-CNN

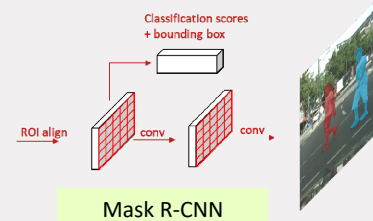
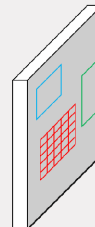
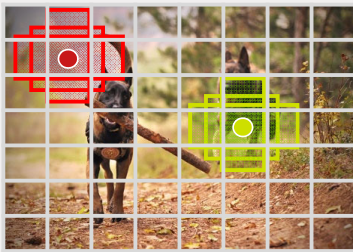
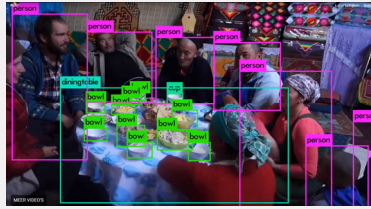
Object detection

Region proposals

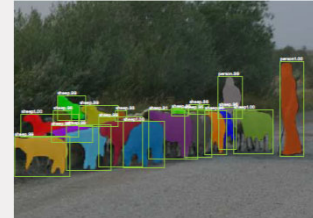
4 5LSM0 Module 9: Unsupervised learning

Last time

YOLO / SsSD



Instance segmentation



5

5LSMO Module 9: Unsupervised learning

This time

Unsupervised learning

- ... with convnets!

Generative models

- PixelRNN and PixelCNN
- Variational Autoencoders (VAEs)
- Generative Adversarial Networks (GANs)



6

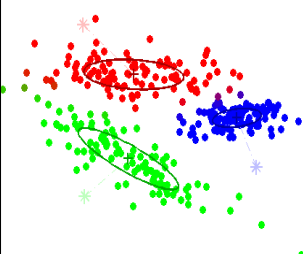
5LSMO Module 9: Unsupervised learning

Supervised vs Unsupervised

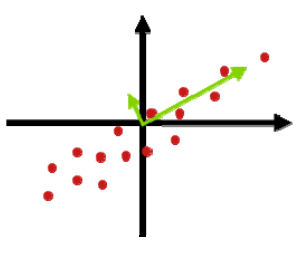
Supervised learning	Unsupervised learning
Labeled data Data $X \in \mathbb{R}^{N \times D}$ and targets $t \in \mathbb{R}^{N \times M}$	Just data, no labels Data $X \in \mathbb{R}^{N \times D}$
Goal: learn a function to map $X \rightarrow t$	Goal: learn some underlying hidden structure of the data
Examples: Classification ($X \in \mathbb{R}, t \in \mathcal{S}$) Regression ($X \in \mathbb{R}, t \in \mathbb{R}$) Segmentation ($X \in \mathbb{R}, t \in \mathcal{S}$)	Examples: Clustering Dimensionality reduction Density estimation



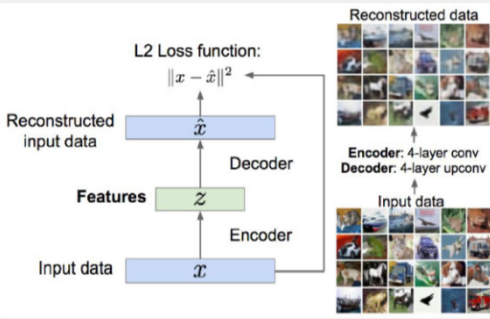
Unsupervised learning



Clustering



Dimensionality reduction



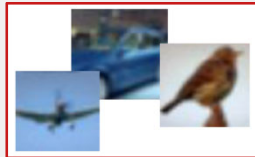
Feature learning



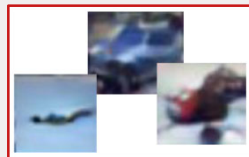
Generative models

Learn the underlying distribution of the (training data)

- So we can generate new samples



Training data $p_{data}(x)$



Generated data $p_{model}(x)$

Goal: learn model $p_{model}(x)$ which is close to the underlying distribution $p_{data}(x)$ of the training data.

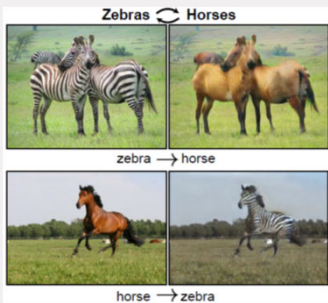
Form of density estimation

- Explicit density estimation: explicitly define and solve for p_{model}
- Implicit density estimation: learn a model that can sample from p_{model}



Generative models

Why are generative models so cool?

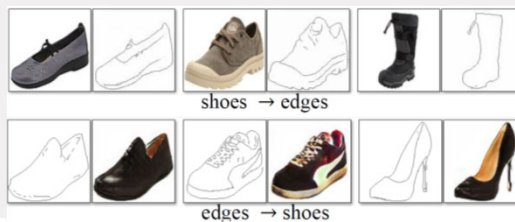


Zhu et al. "Unpaired image-to-image translation using cycle-consistent adversarial networks" CVPR 2017.

Yeh et al. "Semantic image inpainting with deep generative models." CVPR 2017.



These are not actual people



Karras, Tero, Samuli Laine, and Timo Aila "A style-based generator architecture for generative adversarial networks." arXiv preprint arXiv:1812.04948 (2018).



PixelRNN

Type of fully visible belief network

- Explicit density model
- Use chain rule to decompose likelihood of an image x into product of 1-d distributions

$$p(x) = \prod_{i=1}^n p(x_i | x_1, \dots, x_{i-1})$$

Likelihood of image x

Probability of the i^{th} pixel value given all previous pixel values

- Then maximize likelihood of training data

Complex distribution! → Use neural network to model this.

What does this even mean?



11 5LSM0 Module 9: Unsupervised learning

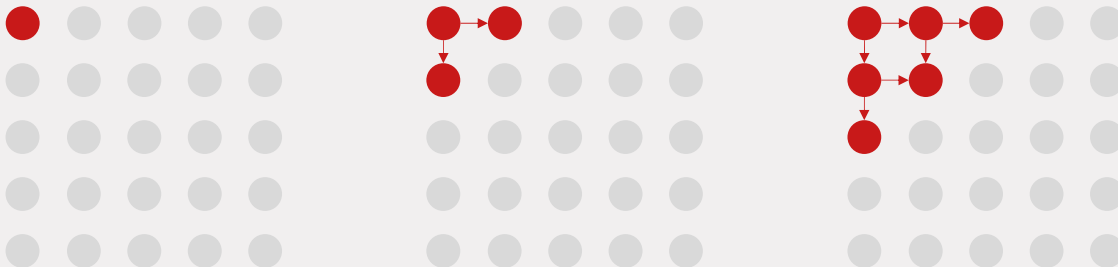
TU/e

Van der Oord et al., "Pixel Recurrent Neural Networks", ICML 2016F

PixelRNN

Generate image pixels from corner

- Dependency on previous pixels modeled using a Recurrent Neural Network (RNN) [module 11]



Drawback: sequential generation is super slow!



12 5LSM0 Module 9: Unsupervised learning

TU/e

Van der Oord et al., "Conditional Image Generation with PixelCNN Decoders", NIPS 2016

PixelCNN

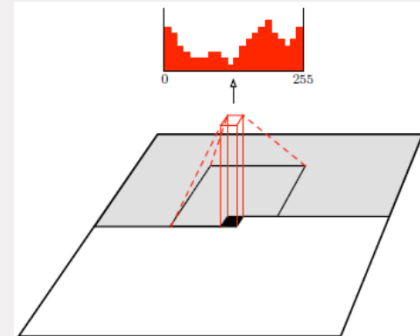
Generate image pixels from corner

- New: dependency on previous pixels now modeled using a CNN over a context region
- Training: maximize likelihood of training images

$$p(x) = \prod_{i=1}^n p(x_i | x_1, \dots, x_{i-1})$$

Training is faster than PixelRNN!

Generation is still slow due to sequential procedure

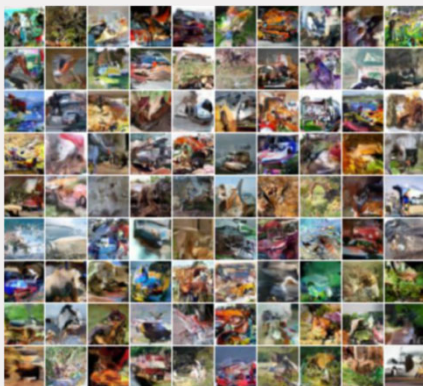


13 5LSM0 Module 9: Unsupervised learning

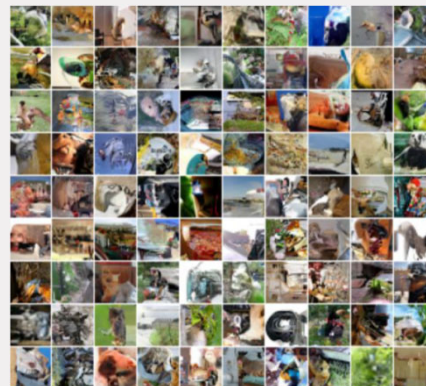
TU/e

Van der Oord et al., "Conditional Image Generation with PixelCNN Decoders", NIPS 2016

PixelCNN



Trained on 32×32 pixel CIFAR-10



Trained on 32×32 pixel ImageNet



14 5LSM0 Module 9: Unsupervised learning

TU/e

PixelRNN & pixelCNN

Pros:

- Can explicitly compute likelihood (quality) of your generated samples $p(x)$
- Explicit likelihood of training data gives good evaluation metric
- Good samples

Cons:

- Sequential generation \rightarrow slow..

Can't we generate an image in a single pass?

- Variational Autoencoders!



15 5LSM0 Module 9: Unsupervised learning

TU/e

Variational autoencoders

So far we have defined a tractable density function

$$p(x) = \prod_{i=1}^n p(x_i | x_1, \dots, x_{i-1})$$

Variational AutoEncoders (VAEs) define intractable density function

$$p_{\theta}(x) = \int p(z) p_{\theta}(x|z) dz$$

We cannot optimize this directly



16 5LSM0 Module 9: Unsupervised learning

TU/e

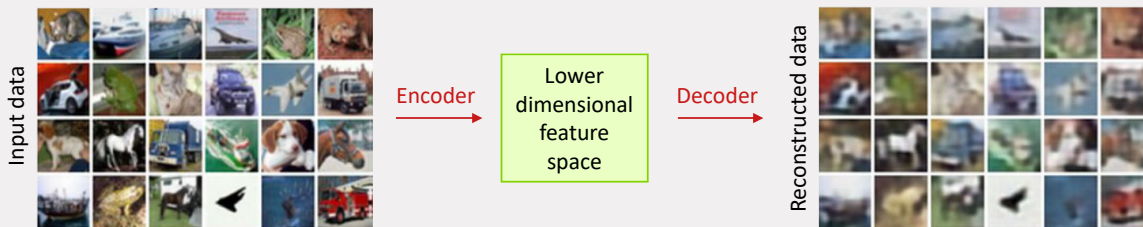
Variational autoencoders

Background: Autoencoders

- Learn a lower-dimensional feature representation

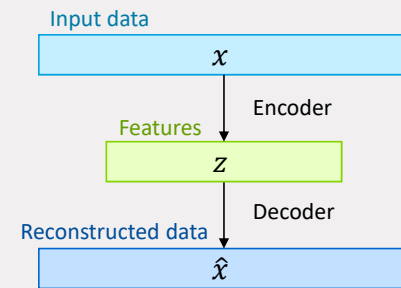
How to learn this transformation?

- Add a decoder to reconstruct the original data



17 5LSM0 Module 9: Unsupervised learning

TU/e



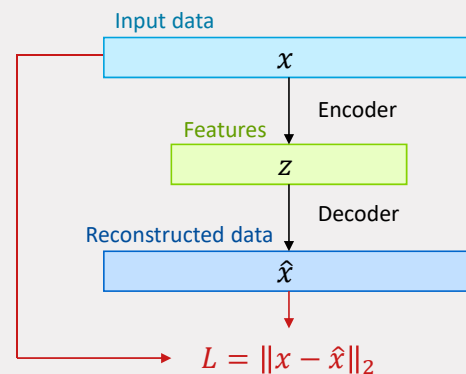
Variational autoencoders

Background: Autoencoders

- Learn a lower-dimensional feature representation

How to learn this transformation?

- Add a decoder to reconstruct the original data
- The mismatch between x and \hat{x} defines the loss



18 5LSM0 Module 9: Unsupervised learning

TU/e

Variational autoencoders

Background: Autoencoders

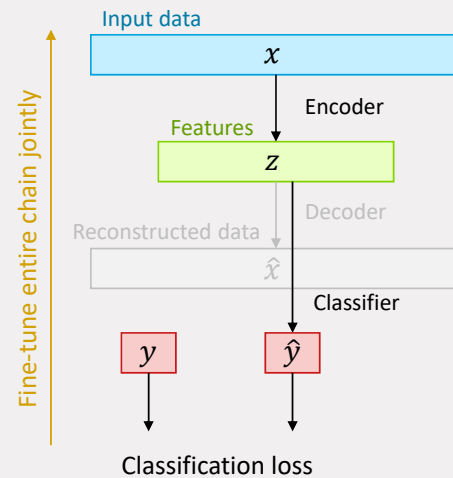
- Learn a lower-dimensional feature representation

How to learn this transformation?

- Add a decoder to reconstruct the original data
- The mismatch between x and \hat{x} defines the loss

Use the learned feature representation

- After training, discard the decoder
- Fine-tune for supervised learning task



We can use a large set of unlabeled data to learn a good initialization for a supervised model!



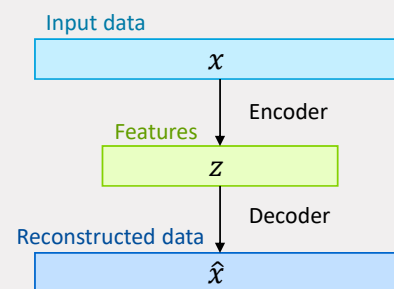
Variational autoencoders

Autoencoder summary

- Autoencoders can reconstruct data and can learn features to initialize a supervised model
- Features capture factors of variation in training data.

Can we generate new images from an autoencoder?

- Variational autoencoders!



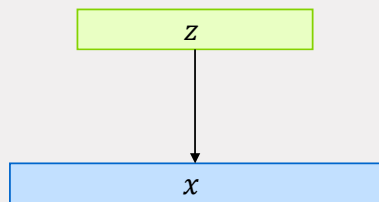
Variational autoencoders

Probabilistic spin on autoencoders

- Let us sample from the model to generate data!
- Assume training data $\{x^{(i)}\}_{i=1}^N$ is generated from underlying latent representation z

Sample from true conditional $p_{\theta^*}(z)$

Sample from true conditional $p_{\theta^*}(x|z^{(i)})$



Intuitive interpretation

- **z** contains information on how much each of a set of latent attributes is present in x .
- Example: if our training data contained images of faces, elements of z could represent the size of the ears, the position & color of the eyes, etc..



Variational autoencoders

→ Estimate the true parameters θ^* of this generative model

How should we represent this model?

- Choose prior $p(z)$ simple: Gaussian
- Conditional $p(x|z)$ is complex → Neural network
 - *This generates the image from the latent representation*

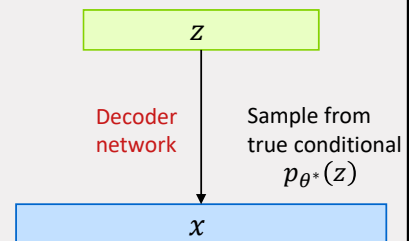
How to train this model?

- Earlier we saw: maximize likelihood of training data:

$$p_{\theta}(x) = \int p(z)p_{\theta}(x|z)dz$$

Q: Problem?

Sample from true conditional $p_{\theta^*}(x|z^{(i)})$



Variational autoencoders

Let's take a closer look

- Data likelihood:

$$p_{\theta}(x) = \int \underbrace{p(z)}_{\text{Simple Gaussian prior}} \underbrace{p_{\theta}(x|z)}_{\text{Output of decoder network}} dz$$

Intractable to compute for every z

- Posterior density:

$$p_{\theta}(z|x) = \frac{p_{\theta}(x|z)p_{\theta}(z)}{p_{\theta}(x)}$$

Intractable, so cannot compute the posterior → cannot optimize model

Solution:

- Define an additional encoder network $q_{\phi}(z|x)$ to approximate $p_{\theta}(z|x)$

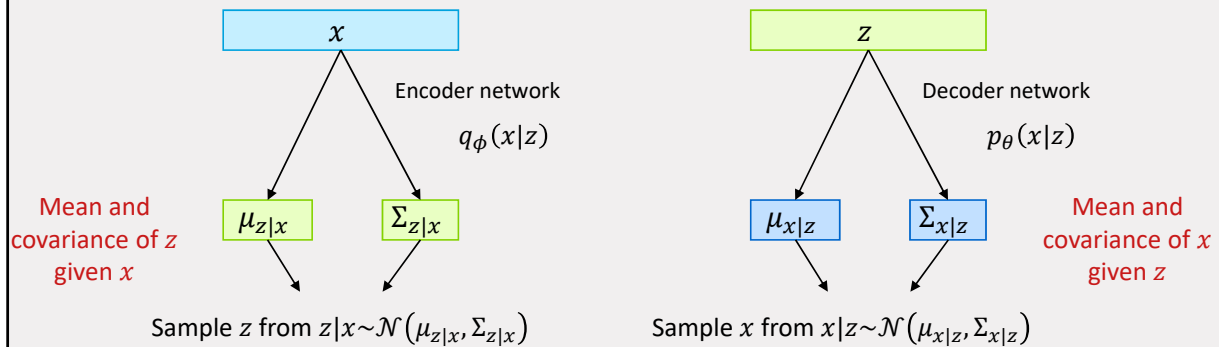


Variational autoencoders

Probabilistic generation of data

- Encoder and decoder networks are probabilistic

Kingma and Welling, "Auto-Encoding Variational Bayes", ICLR 2014



Variational autoencoders

$$\begin{aligned}
 \log p_{\theta}(x^{(i)}) &= \mathbf{E}_z[\log p_{\theta}(x^{(i)})] &&= \mathbf{E}_z \left[\log \frac{p_{\theta}(x^{(i)}|z)p_{\theta}(z)}{p_{\theta}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z \left[\log \frac{p_{\theta}(x^{(i)}|z)p_{\theta}(z)}{p_{\theta}(z|x^{(i)})} \cdot \frac{q_{\phi}(z|x^{(i)})}{q_{\phi}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z [\log p_{\theta}(x^{(i)}|z)] - \mathbf{E}_z \left[\log \frac{q_{\phi}(z|x^{(i)})}{p_{\theta}(z)} \right] + \mathbf{E}_z \left[\log \frac{q_{\phi}(z|x^{(i)})}{p_{\theta}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z [\log p_{\theta}(x^{(i)}|z)] - D_{KL}(q_{\phi}(z|x^{(i)}) || p_{\theta}(z)) + D_{KL}(q_{\phi}(z|x^{(i)}) || p_{\theta}(z|x^{(i)}))
 \end{aligned}$$

$p_{\theta}(x^{(i)})$ independent of z

Bayes' rule

Multiply by constant

Rewrite using log identities

Decoder network gives $p_{\theta}(x|z)$, can compute estimate of this term through sampling

This KL term (between Gaussians for encoder and z prior) has nice closed form solution

$p_{\theta}(z|x^{(i)})$ intractable, can't compute this KL term. But we know divergence always ≥ 0

25 SLSM0 Module 9: Unsupervised learning

Variational autoencoders

$$\begin{aligned}
 \log p_{\theta}(x^{(i)}) &= \mathbf{E}_z[\log p_{\theta}(x^{(i)})] &&= \mathbf{E}_z \left[\log \frac{p_{\theta}(x^{(i)}|z)p_{\theta}(z)}{p_{\theta}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z \left[\log \frac{p_{\theta}(x^{(i)}|z)p_{\theta}(z)}{p_{\theta}(z|x^{(i)})} \cdot \frac{q_{\phi}(z|x^{(i)})}{q_{\phi}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z [\log p_{\theta}(x^{(i)}|z)] - \mathbf{E}_z \left[\log \frac{q_{\phi}(z|x^{(i)})}{p_{\theta}(z)} \right] + \mathbf{E}_z \left[\log \frac{q_{\phi}(z|x^{(i)})}{p_{\theta}(z|x^{(i)})} \right] \\
 &= \mathbf{E}_z [\log p_{\theta}(x^{(i)}|z)] - D_{KL}(q_{\phi}(z|x^{(i)}) || p_{\theta}(z)) + D_{KL}(q_{\phi}(z|x^{(i)}) || p_{\theta}(z|x^{(i)}))
 \end{aligned}$$

Variational lower bound
 $\log p_{\theta}(x^{(i)}) \geq \mathcal{L}(x^{(i)}, \theta, \phi)$

Training: minimize lower bound
 $\theta^*, \phi^* = \arg \max_{\theta, \phi} \mathcal{L}(x^{(i)}, \theta, \phi)$

Force features to reconstruct the input data

Force features to follow our prior

Tractable lower bound $\mathcal{L}(x^{(i)}, \theta, \phi)$ which we can take the gradient of and optimize

always ≥ 0

26 SLSM0 Module 9: Unsupervised learning

Variational autoencoders

How to maximize this lower bound?

- Let's look at computing this bound for a given mini-batch of data

$$\mathcal{L}(x^{(i)}, \theta, \phi) = \mathbf{E}_z[\log p_\theta(x^{(i)}|z)] - D_{KL}(q_\phi(z|x^{(i)})||p_\theta(z))$$

For every minibatch of input data: compute this forward pass and then backprop

Use to compute KL term

Use to compute reconstruction term

Input data x

Encoder network $q_\phi(z|x)$

$\mu_{z|x}$ $\Sigma_{z|x}$

Sample z from $z|x \sim \mathcal{N}(\mu_{z|x}, \Sigma_{z|x})$

Decoder network $p_\theta(x|z)$

$\mu_{x|z}$ $\Sigma_{x|z}$

Sample x from $x|z \sim \mathcal{N}(\mu_{x|z}, \Sigma_{x|z})$

\hat{x}

27 5LSM0 Module 9: Unsupervised learning **TU/e**

Variational autoencoders

Generating data

- Sample from prior in feature space and decode!

Sample z from $z|x \sim \mathcal{N}(0, I)$

Decoder network $p_\theta(x|z)$

$\mu_{x|z}$ $\Sigma_{x|z}$

Sample x from $x|z \sim \mathcal{N}(\mu_{x|z}, \Sigma_{x|z})$

\hat{x}

Example: data manifold for $z \in \mathbb{R}^2$

Vary z_1

Vary z_2

28 5LSM0 Module 9: Unsupervised learning **TU/e**

Variational autoencoders

Generating data

- Sample from prior in feature space and decode!

Sample z from $z|x \sim \mathcal{N}(0, I)$

Sample x from $x|z \sim \mathcal{N}(\mu_{x|z}, \Sigma_{x|z})$

\hat{x}

Q1: Benefit of enforcing diagonal covariance?

Q2: What if we would have sampled from the feature space of the vanilla AE?

Example: Faces

29 5LSM0 Module 9: Unsupervised learning

Variational autoencoders

Intuitive explanation VAEs

- What if we would like to interpolate between two points in feature space when using a vanilla autoencoder
- We have computed a bunch of feature representations, so we know where in feature space certain objects are
- Is this interpolation straightforward when using normal autoencoders?

30 5LSM0 Module 9: Unsupervised learning

Variational autoencoders

Intuitive explanation VAEs

- What if we would like to interpolate between two points in feature space when using a vanilla autoencoder
- We have computed a bunch of feature representations, so we know where in feature space certain objects are
- Is this interpolation straightforward when using normal autoencoders?

We don't know how different categories get mapped to feature space!

31 5LSM0 Module 9: Unsupervised learning

Variational autoencoders

VAEs add two crucial components to fix this

1. A data sample x is not mapped to a single point in feature space, rather it is smeared out over a small region by means of a probability distribution
 - *This means that when decoding, there should be overlap between objects that look similar and, hence, forces them to "stick" together*

$$\mathcal{L}(x^{(i)}, \theta, \phi) = \mathbf{E}_z[\log p_\theta(x^{(i)}|z)] - D_{KL}(q_\phi(z|x^{(i)})||p_\theta(z))$$

2. A cost is introduced for distributions that deviate from a zero-mean, unity std dev Gaussian.
 - *This keeps the feature space compact*

The region responsible for our favorite boatcar connects to both car and boat samples in feature space!

32 5LSM0 Module 9: Unsupervised learning

So far...

Pixel RNN/CNNs

- PixelCNNs define tractable density function, optimize likelihood of training data:

Variational Autoencoders

- VAEs define intractable density function with latent z :

$$p_{\theta}(x) = \int p(z)p_{\theta}(x|z)dz$$

- We cannot optimize this directly, instead we derive a lower bound on the likelihood to maximize

$$p(x) = \prod_{i=1}^n p(x_i|x_1, \dots, x_{i-1})$$

What if we give up on explicitly modeling density altogether?

- We just want the ability to sample



33 5LSM0 Module 9: Unsupervised learning

TU/e

Generative Adversarial Networks

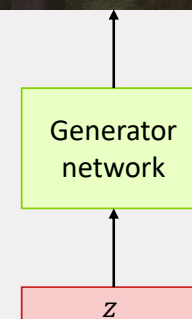
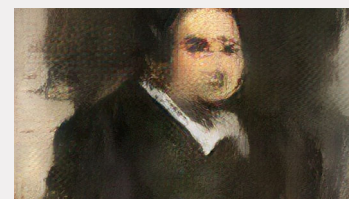
Generative Adversarial Networks (GANs)

- Don't work with an explicit density function!
- Instead, learn to generate from training distribution through a 2-player game...

Problem: Want to sample from complex, high-dimensional training distribution

Solution: Sample from a simple distribution, e.g. white noise, and learn to transformation from simple distribution to training distribution

Model this transformation with a CNN!



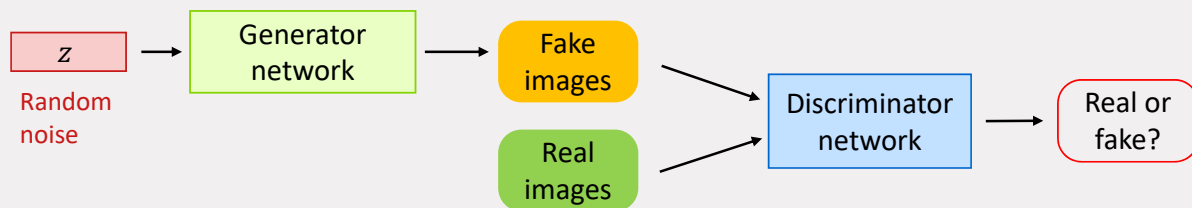
34 5LSM0 Module 9: Unsupervised learning

TU/e

Generative Adversarial Networks

It's a two player game

- Generator network: try to fool the discriminator by generating real-looking images
- Discriminator network: try to distinguish between real and fake images



35 5LSM0 Module 9: Unsupervised learning

TU/e

Generative Adversarial Networks

It's a two player game

- Generator network: try to fool the discriminator by generating real-looking images
- Discriminator network: try to distinguish between real and fake images

Train jointly in minimax game

- Objective function:

Discriminator outputs likelihood of data sample being real [0, 1]

$$\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

Generator wants to minimize

Discriminator wants to maximize

Expected discriminator output for real data

Expected discriminator output for fake data



36 5LSM0 Module 9: Unsupervised learning

TU/e

Generative Adversarial Networks

Q: What does this $D(G(z)) \rightarrow 0$ imply?

Minimax objective function

$$\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

Alternate between

- Gradient ascent on discriminator

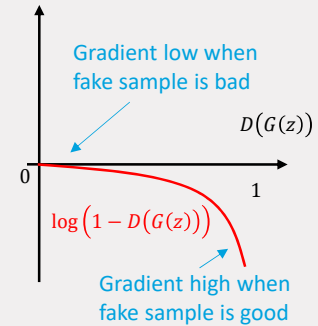
$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

- Gradient descent on generator

$$\min_{\theta_g} \left[\mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

In practice, this doesn't work very well...

Generator learns slow when it's doing a bad job...



Generative Adversarial Networks

Note: training GANs is very challenging. Better objective functions for this is an active area of research.

Minimax objective function

$$\min_{\theta_g} \max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

Alternate between

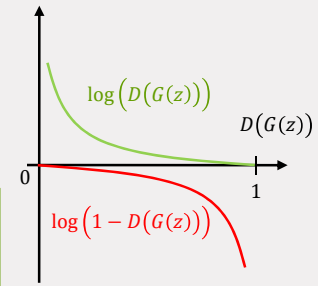
- Gradient ascent on discriminator

$$\max_{\theta_d} \left[\mathbb{E}_{x \sim p_{data}} \log D_{\theta_d}(x) + \mathbb{E}_{z \sim p(z)} \log \left(1 - D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

- Instead: Gradient ascent on generator

$$\max_{\theta_g} \left[\mathbb{E}_{z \sim p(z)} \log \left(D_{\theta_d}(G_{\theta_g}(z)) \right) \right]$$

Much better! Gradient large when generated sample is likely fake → better learning



Generative Adversarial Networks

Training GANs: putting it all together

Train discriminator
Train generator

```

for number of training iterations do
  for k steps do
    → Sample minibatch of m noise samples {z(1), ..., z(m)} from noise prior pg(z)
    → Sample minibatch of m examples {x(1), ..., x(m)} of real data
    → Update the discriminator by ascending its stochastic gradient:
      ∇θd  $\frac{1}{m} \sum_{i=1}^m \log D_{\theta_d}(x) + \log (1 - D_{\theta_d}(G_{\theta_g}(z)))$ 
    end for
    → Sample minibatch of m noise samples {z(1), ..., z(m)} from noise prior pg(z)
    → Update the generator by ascending its stochastic gradient:
      ∇θg  $\frac{1}{m} \sum_{i=1}^m \log (D_{\theta_d}(G_{\theta_g}(z)))$ 
    end for
  end for
end for
    
```

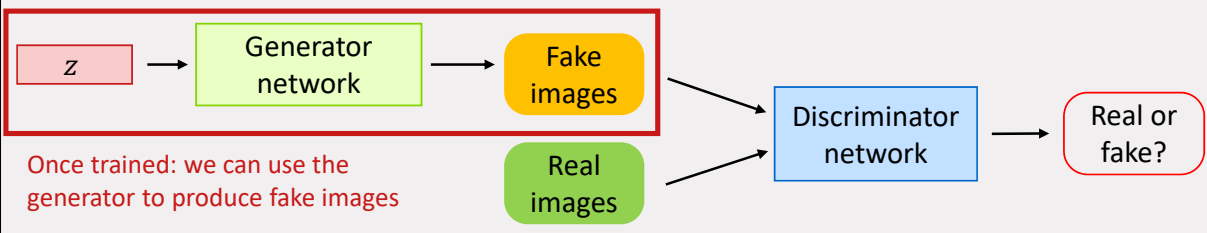
Discriminate between fake and real data!
Improve generated fakes to fool discriminator!



Generative Adversarial Networks

It's a two player game

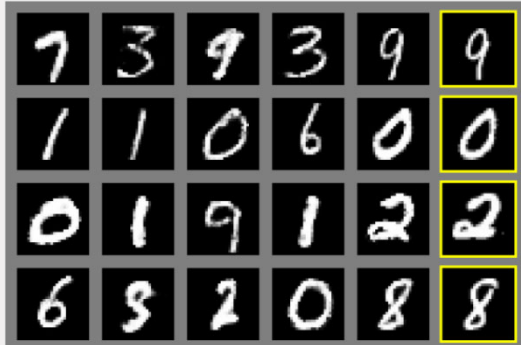
- Generator network: try to fool the discriminator by generating real-looking images
- Discriminator network: try to distinguish between real and fake images



Generative Adversarial Networks

GANs: first results (2014)

Ian Goodfellow et al., "Generative Adversarial Nets", NIPS 2014



Generated digits

Nearest neighbors



Generated faces

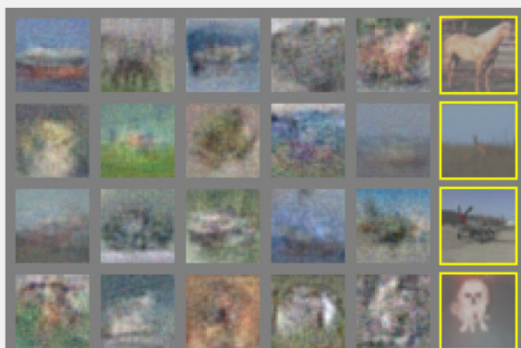
Nearest neighbors



Generative Adversarial Networks

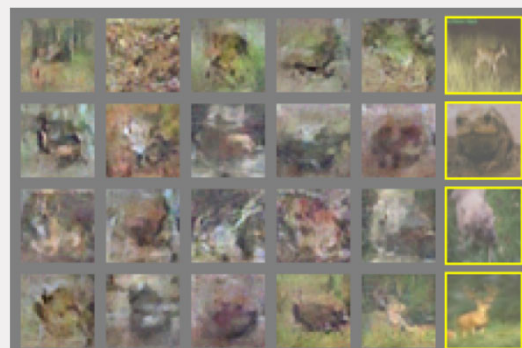
GANs: first results (2014)

Ian Goodfellow et al., "Generative Adversarial Nets", NIPS 2014



Results on CIFAR-10

Nearest neighbors



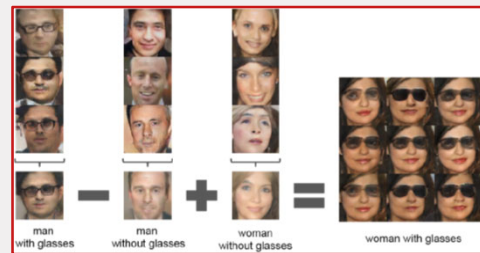
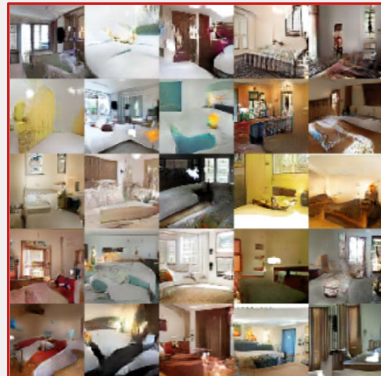
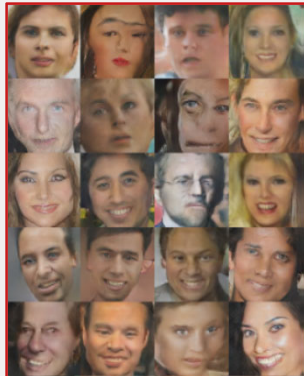
Nearest neighbors



Generative Adversarial Networks

GANs: major improvement (2016)

Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016



43 5LSM0 Module 9: Unsupervised learning



Generative Adversarial Networks

GANs: major improvement (2016)

Radford et al, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks", ICLR 2016

Interpolation



Q: Why does this work?

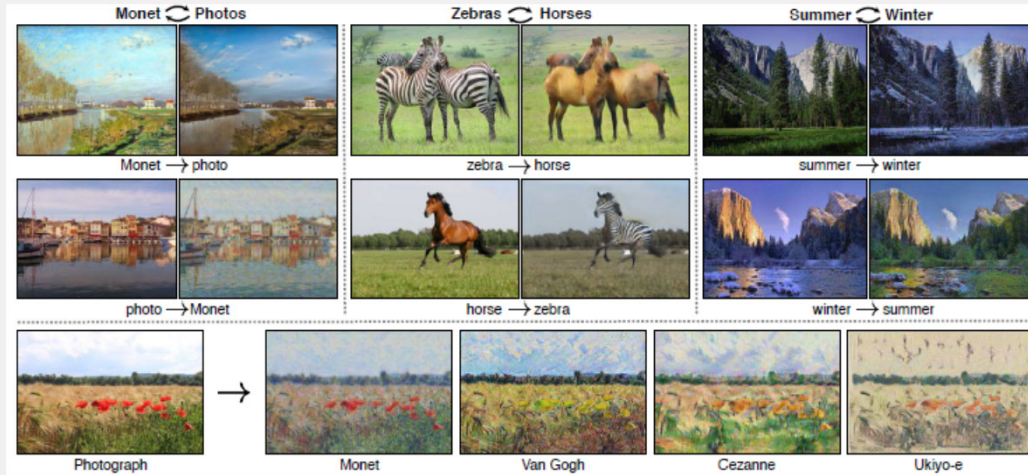


44 5LSM0 Module 9: Unsupervised learning



Generative Adversarial Networks

Zhu et al., "Unpaired image-to-image translation using cycle-consistent adversarial networks", CVPR 2017

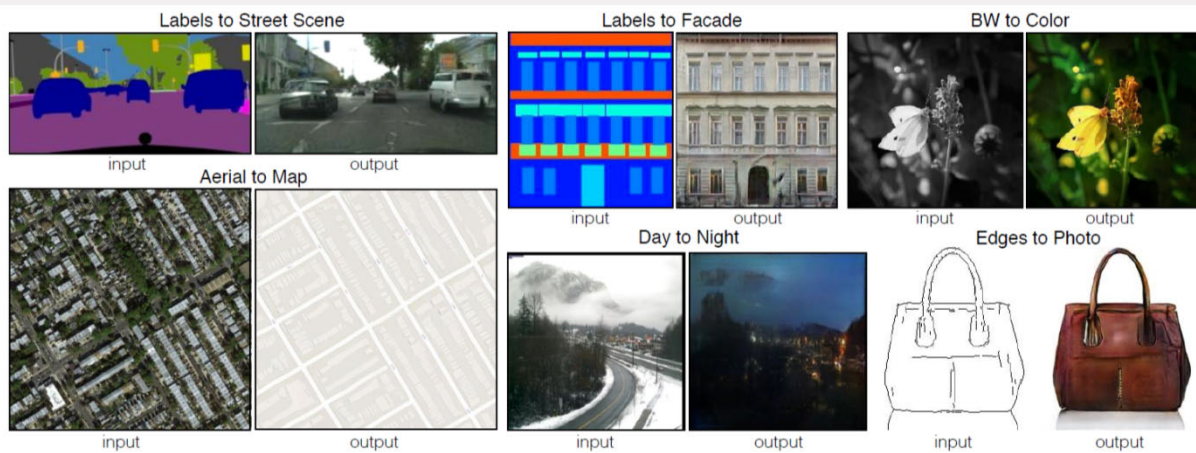


45 5LSM0 Module 9: Unsupervised learning



Generative Adversarial Networks

Isola et al., "Image-to-Image Translation with Conditional Adversarial Nets", CVPR 2017



46 5LSM0 Module 9: Unsupervised learning

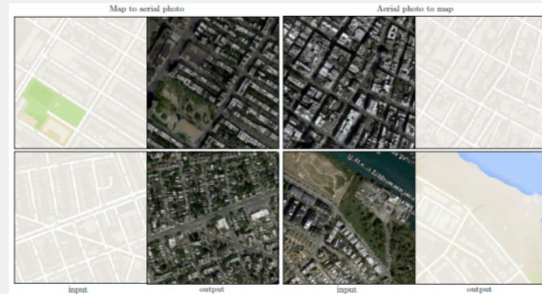


Generative Adversarial Networks



Isola et al., "Image-to-Image Translation with Conditional Adversarial Nets", CVPR 2017

<https://arxiv.org/abs/1611.07004>



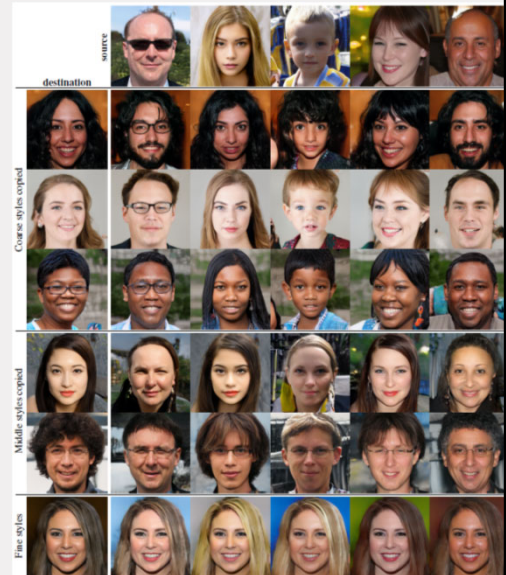
47 5LSM0 Module 9: Unsupervised learning



Generative Adversarial Networks



Karras et al., "A style-based generator architecture for generative adversarial networks." arXiv preprint arXiv:1812.04948 (2018).



48 5LSM0 Module 9: Unsupervised learning



Summary

Unsupervised learning with convnets

**UP NEXT:
SEMI-SUPERVISED LEARNING**

Generative models

- PixelRNN and PixelCNN
 - *Explicit density model, optimizes exact likelihood, good samples.*
 - *Inefficient sequential generation*
- Variational Autoencoders (VAEs)
 - *Optimize variational lower bound on likelihood. Useful latent representation, inference queries.*
 - *Current sample quality not the best.*
- Generative Adversarial Networks (GANs)
 - *Game-theoretic approach, best samples!*
 - *Can be tricky and unstable to train, no inference queries (e.g. $p(x)$, $p(x|z)$)*

