

Variational Auto Encoders

Machine Learning WS2021/22

Prof. Gerhard Neumann

KIT, Institut für Anthropomatik und Robotik

Learning Outcome for today...

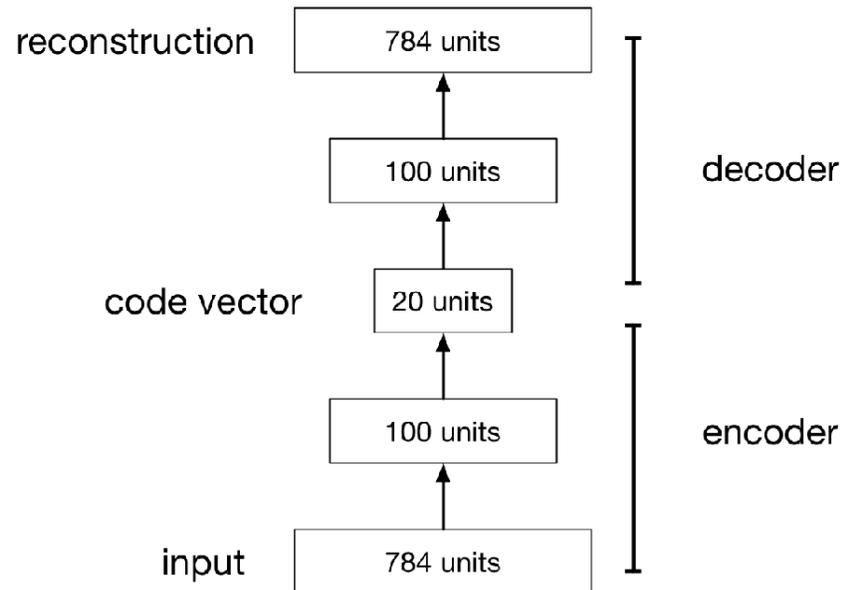
In this lecture, we will..

- Understand the **auto-encoders and what you can do with it**
- How to use them as generative model and its connection to latent variable models
- What **variational auto-encoders** are...
- ... and how to train them using **Variational Bayes**

Auto-encoders

Auto-encoders

- An **autoencoder** is a feed-forward neural net whose job it is to take an input x and predict x
- To make this non-trivial, we need to add a **bottleneck layer** whose dimension is much smaller than the input



Minimize reconstruction loss:

$$L(\theta) = \sum_i ||\text{dec}_\theta(\text{enc}_\theta(\mathbf{x}_i)) - \mathbf{x}_i||^2$$

- **Note:** the simplest auto-encoder only has 1 linear layer for the encoder and decoder -> PCA

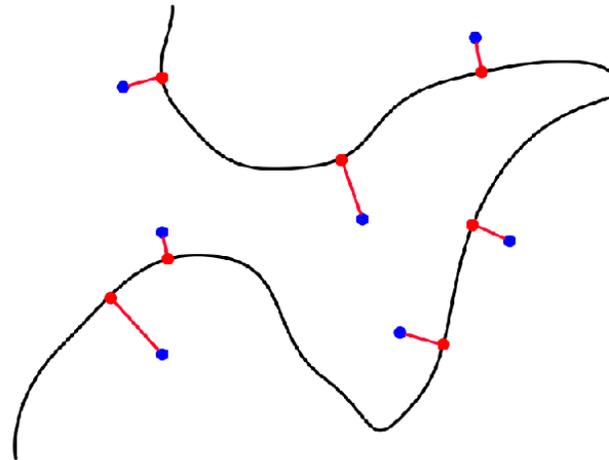
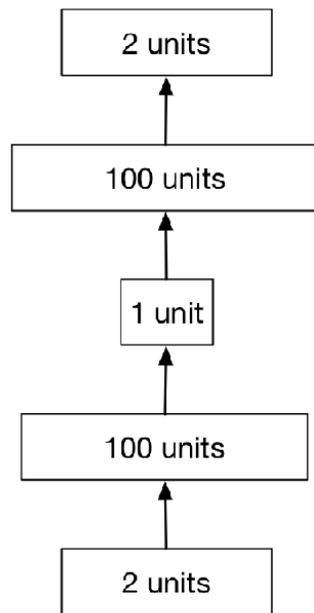
Auto-encoders

Why autoencoders?

- Map high-dimensional data to two dimensions for visualization
- Compression (i.e. reducing the file size)
 - Note: this requires a VAE, not just an ordinary autoencoder.
- Learn abstract features in an unsupervised way so you can apply them to a supervised task
 - Unlabeled data can be much more plentiful than labeled data
- Learn a semantically meaningful representation where you can, e.g., interpolate between different images.

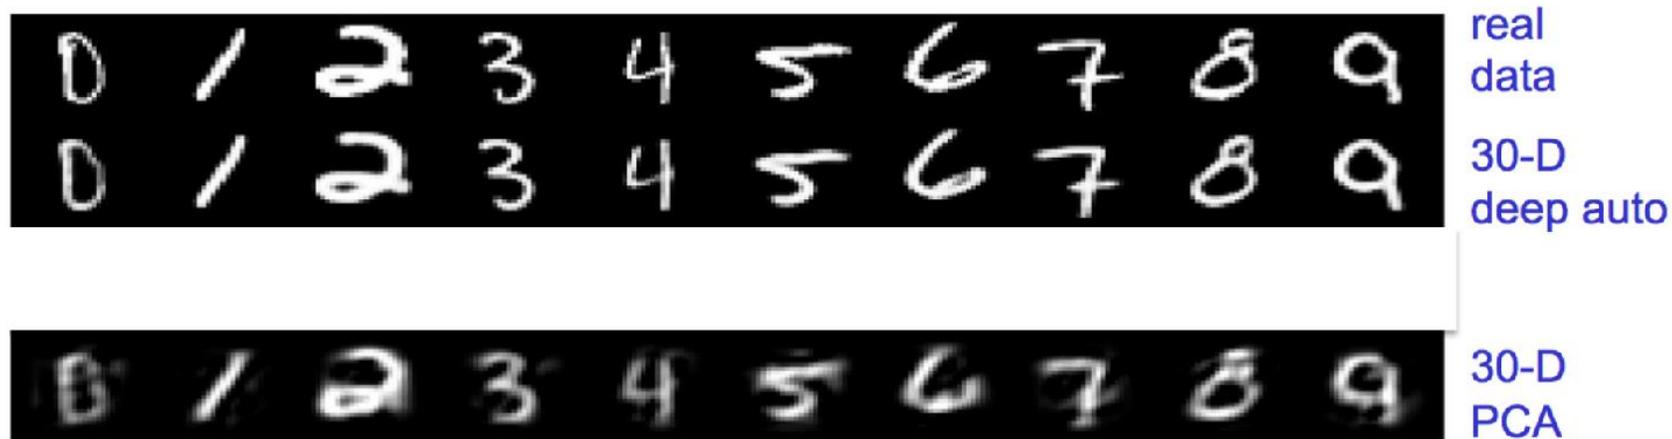
Deep auto-encoders

- Deep nonlinear autoencoders learn to project the data, not onto a linear subspace, but onto a **nonlinear manifold**
- This manifold is the image of the decoder.
- This is a kind of **nonlinear dimensionality reduction**.



Deep auto-encoders

- Nonlinear autoencoders can learn more powerful codes for a given dimensionality, compared with linear autoencoders (PCA)



Some limitations of autoencoders

- They're not generative models, so they don't define a distribution
- How to choose the latent dimension?

Generative Model

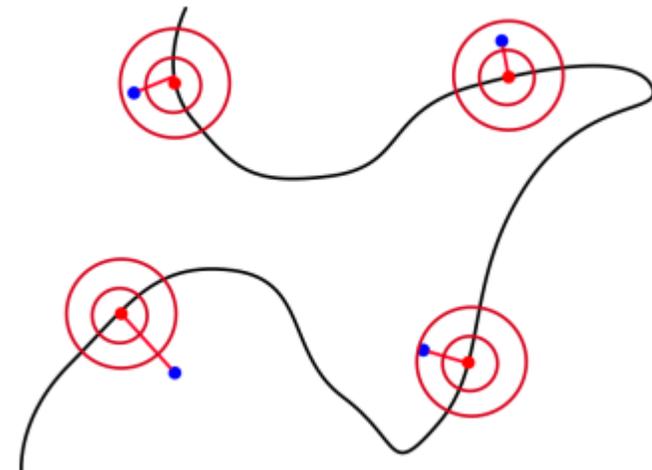
- Consider training a generator network with maximum likelihood.

$$p(\mathbf{x}) = \int p(\mathbf{z})p(\mathbf{x}|\mathbf{z})d\mathbf{z}$$

- One problem: if \mathbf{z} is low-dimensional and the decoder is deterministic, then $p(\mathbf{x}) = 0$ almost everywhere!
 - The model only generates samples over a low-dimensional sub-manifold of X .
- Solution: define a noisy observation model, e.g.

$$p(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_{\theta}(\mathbf{z}), \sigma^2 \mathbf{I})$$

where $\boldsymbol{\mu}_{\theta}(\mathbf{z})$ is the function computed by the decoder with parameters θ .



Latent Variable Models and Variational Bayes

Latent Variable Models

Recap: **latent variable** models

- Examples: mixture models, missing data, latent factors, **variational auto-encoders**
- Observed variables: \mathbf{x} , Latent variables: \mathbf{z}

- **Parametric model:** $p_{\theta}(\mathbf{x}, \mathbf{z}) = p(\mathbf{z})p_{\theta}(\mathbf{x}|\mathbf{z})$

- **Marginal distribution:** $p_{\theta}(\mathbf{x}) = \int_{\mathbf{z}} p(\mathbf{z})p_{\theta}(\mathbf{x}|\mathbf{z})d\mathbf{z}$

At least the integral $p_{\theta}(\mathbf{x}) = \int_{\mathbf{z}} p(\mathbf{z})p_{\theta}(\mathbf{x}|\mathbf{z})d\mathbf{z}$ is well-defined, but how can we compute it?

- The decoder function is very complicated, so there's no hope of finding a closed form.

Marginal Likelihood for Latent Variable Models

(Marginal) Log-Likelihood:

$$\text{LogLik}(\mathcal{D}) = \sum_{i=1}^N \log p(\mathbf{x}_i) = \sum_{i=1}^N \log \left(\int p(\mathbf{x}_i | \mathbf{z}) p(\mathbf{z}) d\mathbf{z} \right) \approx \sum_{i=1}^N \log \left(\frac{1}{M} \sum_{\mathbf{z}_j \sim p(\mathbf{z})} p(\mathbf{x}_i | \mathbf{z}_j) \right)$$

... which is **computationally infeasible** in most cases

- Requires a **lot of samples** \mathbf{z}_j for each \mathbf{x}_i due to **uninformed sampling of** \mathbf{z}_j (high variance in $p(\mathbf{x} | \mathbf{z})$)

Variational Bayes

Variational Bayes (VB) uses a lower bound of the marginal log-likelihood for the optimization

- For simplicity, let's consider only a single data-point first

$$\underbrace{\log p(\mathbf{x})}_{\text{marginal log-like}} = \underbrace{\int q(\mathbf{z}) \log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z})} d\mathbf{z}}_{\text{Lower Bound } \mathcal{L}(q)} + \underbrace{\int q(\mathbf{z}) \log \frac{q(\mathbf{z})}{p(\mathbf{z}|\mathbf{x})} d\mathbf{z}}_{\text{KL Divergence: } \text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x}))}$$

- Where $q(\mathbf{z})$ is called **the variational / auxiliary distribution**
 - This decomposition holds for any $q(\mathbf{z})$
 - By introducing $q(\mathbf{z})$, the optimization will become much simpler

- **Why is that the same?**

- We can use Bayes rule for $p(\mathbf{z}|\mathbf{x}) = \frac{p(\mathbf{x}, \mathbf{z})}{p(\mathbf{x})}$ and all terms except $p(\mathbf{x})$ cancel

Note

Expectation maximization is a special case of this decomposition

Simplification in EM:

- Posterior $p(\mathbf{z}|\mathbf{x})$ can be computed in closed form
- **Examples:** Gaussian Mixture Models, Probabilistic PCA

Optimization in Variational Bayes

VB optimizes the lower bound instead of the log-likelihood

$$\mathcal{L}(q, p) = \int q(\mathbf{z}) \log \left(p(\mathbf{x}|\mathbf{z}) \frac{p(\mathbf{z})}{q(\mathbf{z})} \right) d\mathbf{z} = \int q(\mathbf{z}) \log p(\mathbf{x}|\mathbf{z}) d\mathbf{z} - \text{KL}(q(\mathbf{z})||p(\mathbf{z}))$$

- Why is it a lower bound? Since $\text{KL}(q||p) \geq 0$ it follows that $\mathcal{L}(q, p) \leq \log p(\mathbf{x})$
- Its also called **Evidence lower bound (ELBO)**, as the marginal likelihood is often called evidence

Joint optimization of the lower bound $\mathcal{L}(q, p)$ w.r.t p and q using stochastic gradient descent

$$q^*, p^* = \arg \max_{q, p} \mathcal{L}(q, p)$$

- In practice, $q_\phi(\mathbf{z}), p_\varphi(\mathbf{x}|\mathbf{z})$ and $p_\varphi(\mathbf{z})$ will be parametrized distributions and we optimize over ϕ and φ
- We always improve the lower-bound, but there is no guarantee to improve the marginal likelihood
- Standard for most **continuous latent variable** models (e.g. Variational Auto-Encoder)
- Lower bound is only tight if $\text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x}))$ can be set to 0. Thats only true for EM.

Objective for the variational distribution

What does q learn?

$$\begin{aligned}\text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x})) &= \int q(\mathbf{z}) \log \frac{q(\mathbf{z})}{p(\mathbf{z}|\mathbf{x})} d\mathbf{z} = \int q(\mathbf{z}) \log \frac{p(\mathbf{x})q(\mathbf{z})}{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})} d\mathbf{z} = \int q(\mathbf{z}) \log \frac{q(\mathbf{z})}{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})} d\mathbf{z} + \text{const} \\ &= -\mathcal{L}(q, p) + \text{const}\end{aligned}$$

$$\Rightarrow \operatorname{argmin}_q \text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x})) = \operatorname{argmax}_q \mathcal{L}(q, p)$$

- By maximizing the variational lower bound w.r.t $q(\mathbf{z})$, the **variational distribution will approximate the posterior, i.e.,**

$$q(\mathbf{z}) \approx p(\mathbf{z}|\mathbf{x})$$

Full-Dataset Lower Bound

$$\mathcal{L}(\{q_i\}, p) = \frac{1}{N} \sum_i \int q_i(\mathbf{z}) \log p(\mathbf{x}_i | \mathbf{z}) d\mathbf{z} - \text{KL}(q_i(\mathbf{z}) || p(\mathbf{z}))$$

- Introduced individual variational distribution $q_i(\mathbf{z})$ for each data point
- ✓ **More directed sampling:**
 - Instead of sampling from the uninformed prior $p(\mathbf{z})$...
 - ... we can now sample from the variational distributions $q_i(\mathbf{z}) \approx p(\mathbf{z} | \mathbf{x}_i)$
 - Each $q_i(\mathbf{z})$ will produce **samples with high $p(\mathbf{x} | \mathbf{z})$ once optimized!**
- ✓ **Integral is outside the log:**
 - **only one sample from $q_i(\mathbf{z})$** needed to obtain unbiased estimate of the lower bound
 - i.e. suitable for **stochastic gradient descent** while the marginal loglikelihood is not

Special Case 1: Expectation Maximization

$$\underbrace{\log p(\mathbf{x})}_{\text{marginal log-like}} = \underbrace{\int q(\mathbf{z}) \log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z})} d\mathbf{z}}_{\text{Lower Bound } \mathcal{L}(q)} + \underbrace{\int q(\mathbf{z}) \log \frac{q(\mathbf{z})}{p(\mathbf{z}|\mathbf{x})} d\mathbf{z}}_{\text{KL Divergence: } \text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x}))}$$

Expectation-Maximization uses the same decomposition, but two separate optimization steps

– **Maximization Step:**

- Keep $q(\mathbf{z})$ fixed, maximize Lower bound $\mathcal{L}(q, p)$ w.r.t. model distribution $p_\varphi(\mathbf{x}|\mathbf{z})$ and $p_\varphi(\mathbf{z})$

$$\varphi^* = \arg \max_{\varphi} \mathcal{L}(q, p_\varphi)$$

– **Expectation Step:**

- Keep model $p_\varphi(\mathbf{x}|\mathbf{z})$ and $p_\varphi(\mathbf{z})$ fixed, minimize KL w.r.t q

$$q^* = \arg \min_q \text{KL}(q(\mathbf{z})||p(\mathbf{z}|\mathbf{x}))$$

$$\rightarrow q^*(\mathbf{z}) = p(\mathbf{z}|\mathbf{x})$$

Sidenote

- In EM the **KL can be set to zero** (i.e. we can compute posterior $p_\varphi(\mathbf{z}|\mathbf{x})$)
- Only works in special cases
 - e.g. discrete \mathbf{z} , GMMs
- In this case the **lower bound is tight**
- increasing **lower bound** always **increases marginal log-like**

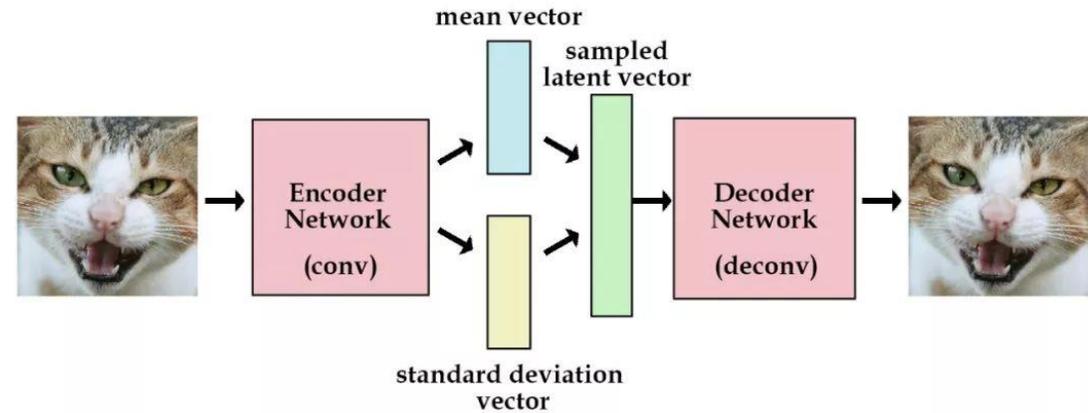
Special Case 2: Amortized Variational Inference

Instead of using an individual auxiliary distribution $q_i(z)$ per data-point x_i , we can use an “amortized” distribution $q_\phi(z|x_i)$ that is given by a DNN

$$\mathcal{L}(q, p) = \frac{1}{N} \sum_i \int q_\phi(z|x_i) \log p_\phi(x_i|z) dz - \text{KL}(q_\phi(z|x_i) || p_\phi(z))$$

This is the standard objective used for variational auto-encoders (VAE)

- Encoder $q_\phi(z|x)$
- Decoder $p_\phi(x|z)$
- Latent Prior $p_\phi(z)$



Optimization over the variational distribution

$$\mathcal{L}(q, p) = \frac{1}{N} \sum_i \int q_\phi(\mathbf{z}|\mathbf{x}_i) \log p_\phi(\mathbf{x}_i|\mathbf{z}) d\mathbf{z} - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i) || p_\phi(\mathbf{z}))$$

But: How can we **optimize over the sampling distribution** $q_\phi(\mathbf{z}|\mathbf{x})$?

- Different to standard max-likelihood: Here, **samples are not fixed but generated!**
- Standard gradients can be used (related to policy gradients), but very inefficient as it does not use gradient information of $\partial \log p_\phi(\mathbf{x}|\mathbf{z})/\partial \mathbf{z}$
- We need something more efficient: **Reparameterization trick!**

Basics: Reparameterization Trick

We want to optimize distribution $p_{\theta}(\mathbf{x})$ using the following expected objective

$$\theta^* = \operatorname{argmin}_{\theta} \mathbb{E}_{p_{\theta}} [f(\mathbf{x})] = \int p_{\theta}(\mathbf{x}) f(\mathbf{x}) d\mathbf{x}$$

... and we are given $\frac{df}{d\mathbf{x}}(\mathbf{x})$. How can we exploit this information?

We can reparametrize the expectation:

- Introduce random variable $\xi \sim q(\xi)$ where q is a simple, parameter-free distribution (e.g., $q(\xi) = \mathcal{N}(\mathbf{0}, \mathbf{I})$)
- If we can find a mapping $\mathbf{x} = \mathbf{h}_{\theta}(\xi)$ such that \mathbf{x} is distributed as $\mathbf{x} \sim p_{\theta}(\mathbf{x})$ then:

$$\int p_{\theta}(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} = \int q(\xi) f(\mathbf{h}_{\theta}(\xi)) d\xi$$

We moved the parameters from distribution $p_{\theta}(\mathbf{x})$ into a function $\mathbf{h}_{\theta}(\xi)$

Basics: Reparameterization Trick

Example: Lets assume $p_{\theta}(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $q(\boldsymbol{\xi}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. If we set

$$\mathbf{h}_{\theta}(\boldsymbol{\xi}) = \boldsymbol{\mu} + \mathbf{A}^T \boldsymbol{\xi}, \quad \text{with } \boldsymbol{\theta} = \{\boldsymbol{\mu}, \boldsymbol{\Sigma}\} \text{ and } \mathbf{A}^T \mathbf{A} = \boldsymbol{\Sigma}$$

Then $\mathbf{x}' = \mathbf{h}_{\theta}(\boldsymbol{\xi})$ is distributed with $p(\mathbf{x}') = \mathcal{N}(\boldsymbol{\mu}, \mathbf{A}^T \mathbf{I} \mathbf{A}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$

Reparameterization Trick:

$$\int p_{\theta}(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} = \int q(\boldsymbol{\xi}) f(\mathbf{h}_{\theta}(\boldsymbol{\xi})) d\boldsymbol{\xi}$$

Reparameterized Gradient:

$$\nabla_{\boldsymbol{\theta}} \mathbb{E}_{p_{\theta}} [f(\mathbf{x})] = \nabla_{\boldsymbol{\theta}} \int q(\boldsymbol{\xi}) f(\mathbf{h}_{\theta}(\boldsymbol{\xi})) d\boldsymbol{\xi} = \int q(\boldsymbol{\xi}) \frac{\partial \mathbf{h}_{\theta}}{\partial \boldsymbol{\theta}}(\boldsymbol{\xi}) \frac{\partial f}{\partial \mathbf{x}}(\mathbf{h}_{\theta}(\boldsymbol{\xi})) d\boldsymbol{\xi}$$

- We can now **use the gradient** $\frac{\partial f}{\partial \mathbf{x}}$ **to compute** $\nabla_{\boldsymbol{\theta}} \mathbb{E}_{p_{\theta}} [f(\mathbf{x})]$!

Back to optimization over the variational distribution

- Lower bound:
$$\mathcal{L}(q, p) = \frac{1}{N} \sum_i \int q_\phi(\mathbf{z}|\mathbf{x}_i) \log p_\varphi(\mathbf{x}_i|\mathbf{z}) d\mathbf{z} - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x}_i)||p_\varphi(\mathbf{z}))$$

- Distribution:
$$q_\phi(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}|\boldsymbol{\mu}_\phi(\mathbf{x}), \sigma_\phi^2(\mathbf{x})\mathbf{I})$$

- Reparametrization function:
$$\mathbf{h}_\phi(\boldsymbol{\xi}, \mathbf{x}) = \boldsymbol{\mu}_\phi(\mathbf{x}) + \boldsymbol{\sigma}_\phi(\mathbf{x}) \circ \boldsymbol{\xi}$$

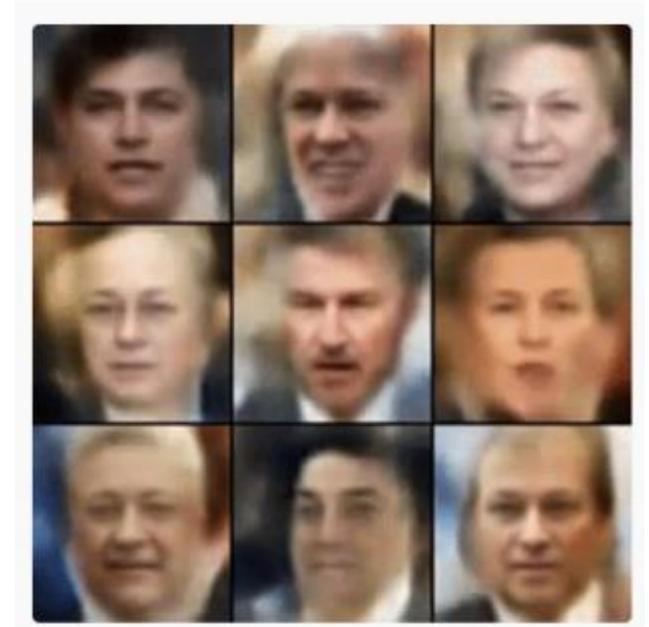
- Reparametrized lower bound:

$$\mathcal{L}(q, p) = \frac{1}{N} \sum_i \int p(\boldsymbol{\xi}) \left(\underbrace{\log p_\varphi(\mathbf{x}_i|\mathbf{h}(\boldsymbol{\xi}, \mathbf{x}))}_{\text{reconstruction}} + \underbrace{\log p_\varphi(\mathbf{h}(\boldsymbol{\xi}, \mathbf{x})) - \log q_\phi(\mathbf{h}(\boldsymbol{\xi}, \mathbf{x})|\mathbf{x})}_{\text{KL-term}} \right)$$

Variational Auto-Encoders

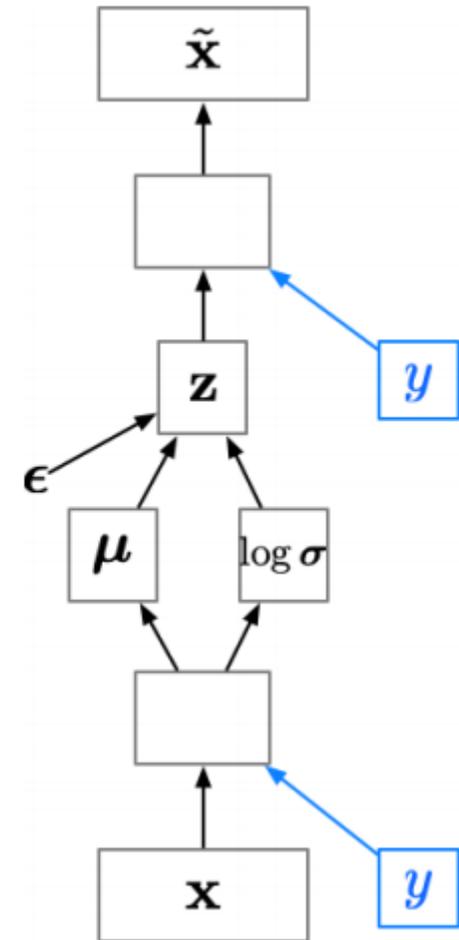
Faces produced by variational auto-encoders

- In comparison to other generative models (such as GANs), VAEs produce rather blurry images
 - More specific methods (e.g. hierarchical VAEs) can achieve similar performance to GANs
 - Most likely cause: Maximum Likelihood objective of VAEs
- In short, a VAE is like an autoencoder, except that it's also a generative model
 - defines a distribution $p(x)$



Class conditional VAEs

- So far, we haven't used the labels y . A **class-conditional VAE** provides the labels to both the encoder and the decoder.
- Since the latent code \mathbf{z} no longer has to model the image category, it can focus on modeling the stylistic features.
- If we're lucky, this lets us disentangle style and content. (Note: disentanglement is still a dark art.)
 - See Kingma et al., “Semi-supervised learning with deep generative models.”



Class-conditional VAE

- By varying two **latent dimensions** (i.e. dimensions of z) while holding y fixed, we can visualize the **latent space**.



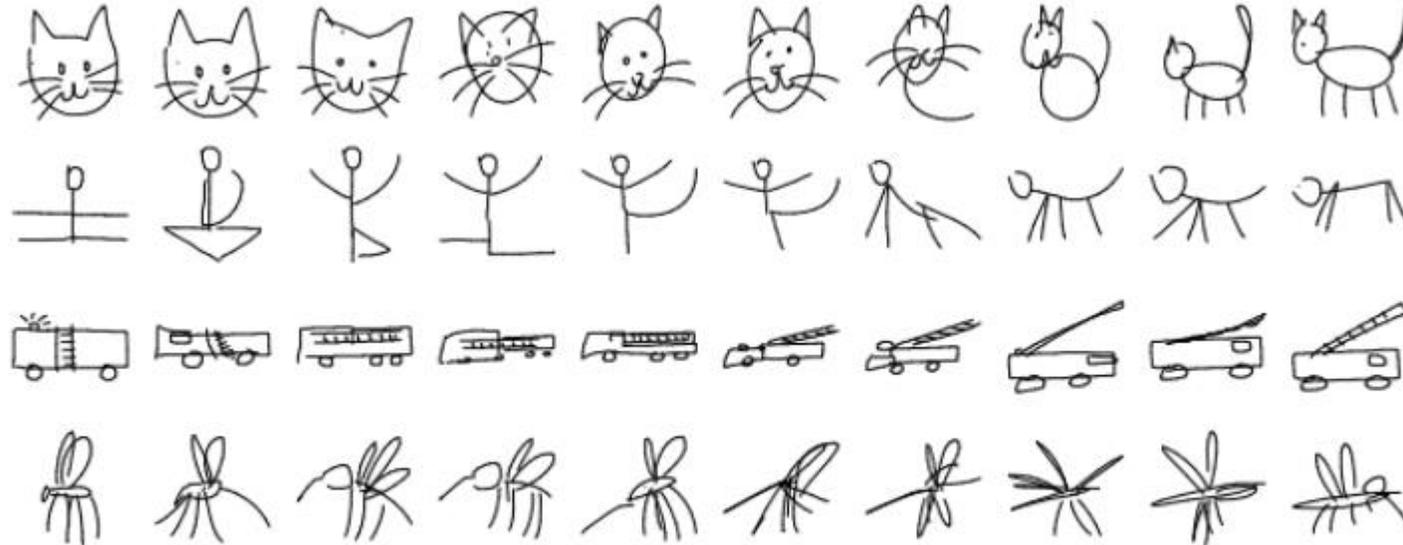
Class-conditional VAE

- By varying the label y while holding z fixed, we can solve image analogies



Latent Space Interpolations

- You can often get interesting results by interpolating between two vectors in the latent space:



Ha and Eck, "A neural representation of sketch drawings"

Wrap-Up Variational Bayes

VB is a method for optimizing the **data likelihood of latent variable models**

- It introduces a variational distribution $q_i(\mathbf{z})$ over the latent variable
 - Variational distribution should approximate posterior $q_i(\mathbf{z}) \approx p(\mathbf{z}|\mathbf{x}_i)$
- ... and decomposes the marginal likelihood in a lower bound and a KL-term
- **The lower bound is in general easier to optimize than the marginal log-likelihood:**
 - ✓ The integral has moved outside the log
 - ✓ More direct sampling in latent space by sampling from approximate posterior $q_i(\mathbf{z})$ instead of prior $p(\mathbf{z})$
 - ✗ No guarantee that we also improve marginal loglikelihood (except in the special case of EM)
- Expectation Maximization is a special case where posterior can be computed analytically
- Most prominent application of VB is the Variational Auto-Encoder

Intermediate Lecture Wrap-Up – Algorithms

Chapter 1: Classical Supervised Learning

- ✓ Linear Regression,
- ✓ Ridge Regression,
- ✓ k-NN,
- ✓ Trees and Forests

Chapter 2: Kernel Methods

- ✓ Kernel-Regression
- ✓ Support Vector Machines

Chapter 3: Bayesian Learning

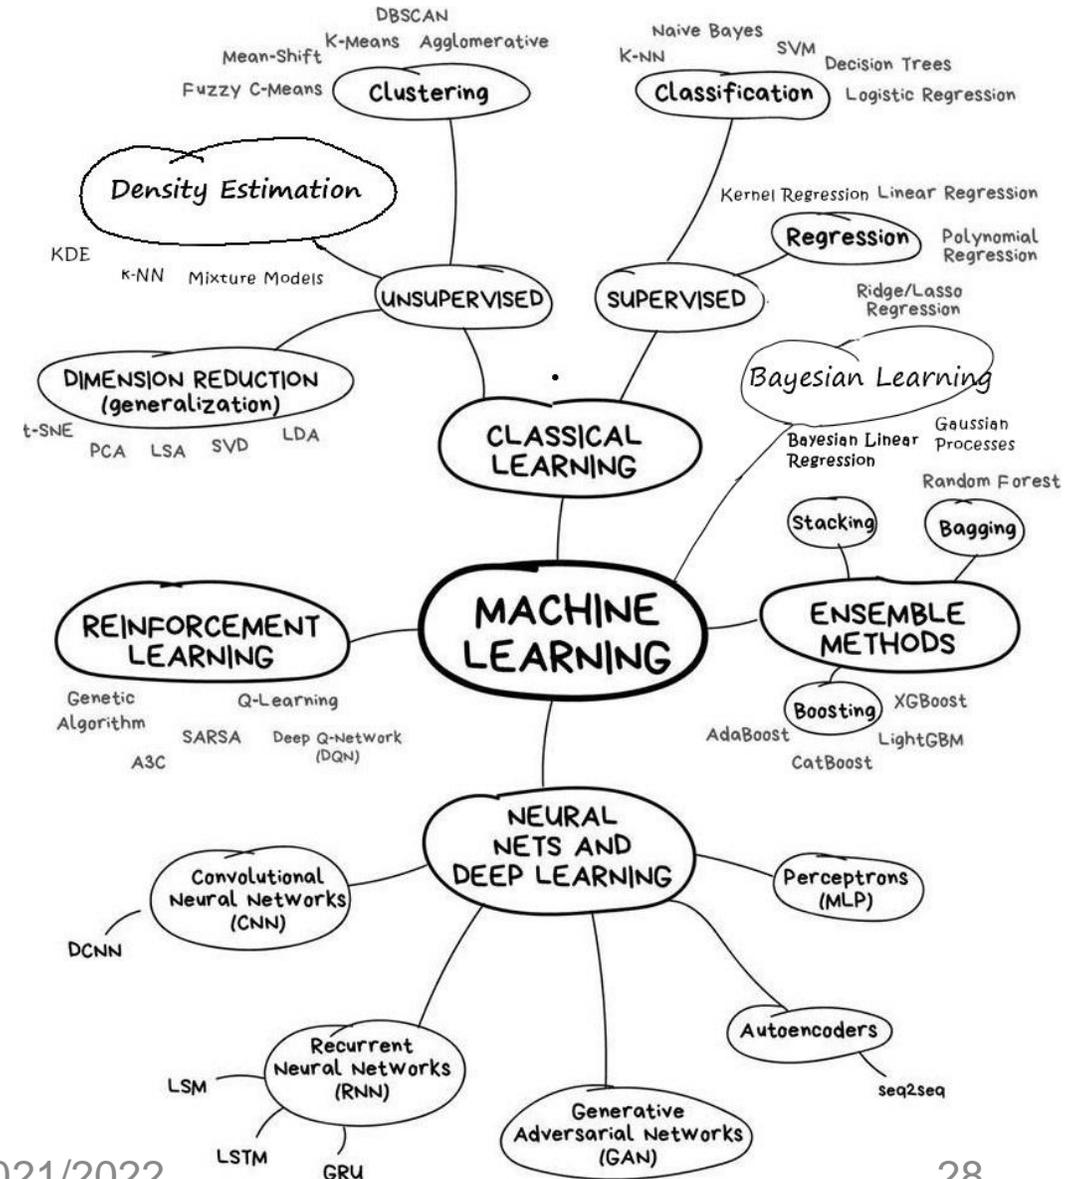
- ✓ Bayesian Linear Regression
- ✓ Gaussian Processes

Chapter 4: Neural Networks

- ✓ Backpropagation
- ✓ MLPs, CNNs, LSTMs

Chapter 5: Unsupervised Learning

- ✓ PCA
- ✓ K-means
- ✓ Expectation-Maximization
- ✓ Variational Auto Encoders



Intermediate Lecture Wrap-Up – Basics

Chapter 1: Classical Supervised Learning

- ✓ Matrix/Vector Calculus
- ✓ Probability Theory, Maximum Likelihood
- ✓ Gradient Descent

Chapter 2: Kernel Methods

- ✓ Sub-gradients
- ✓ Constraint Optimization

Chapter 3: Bayesian Learning

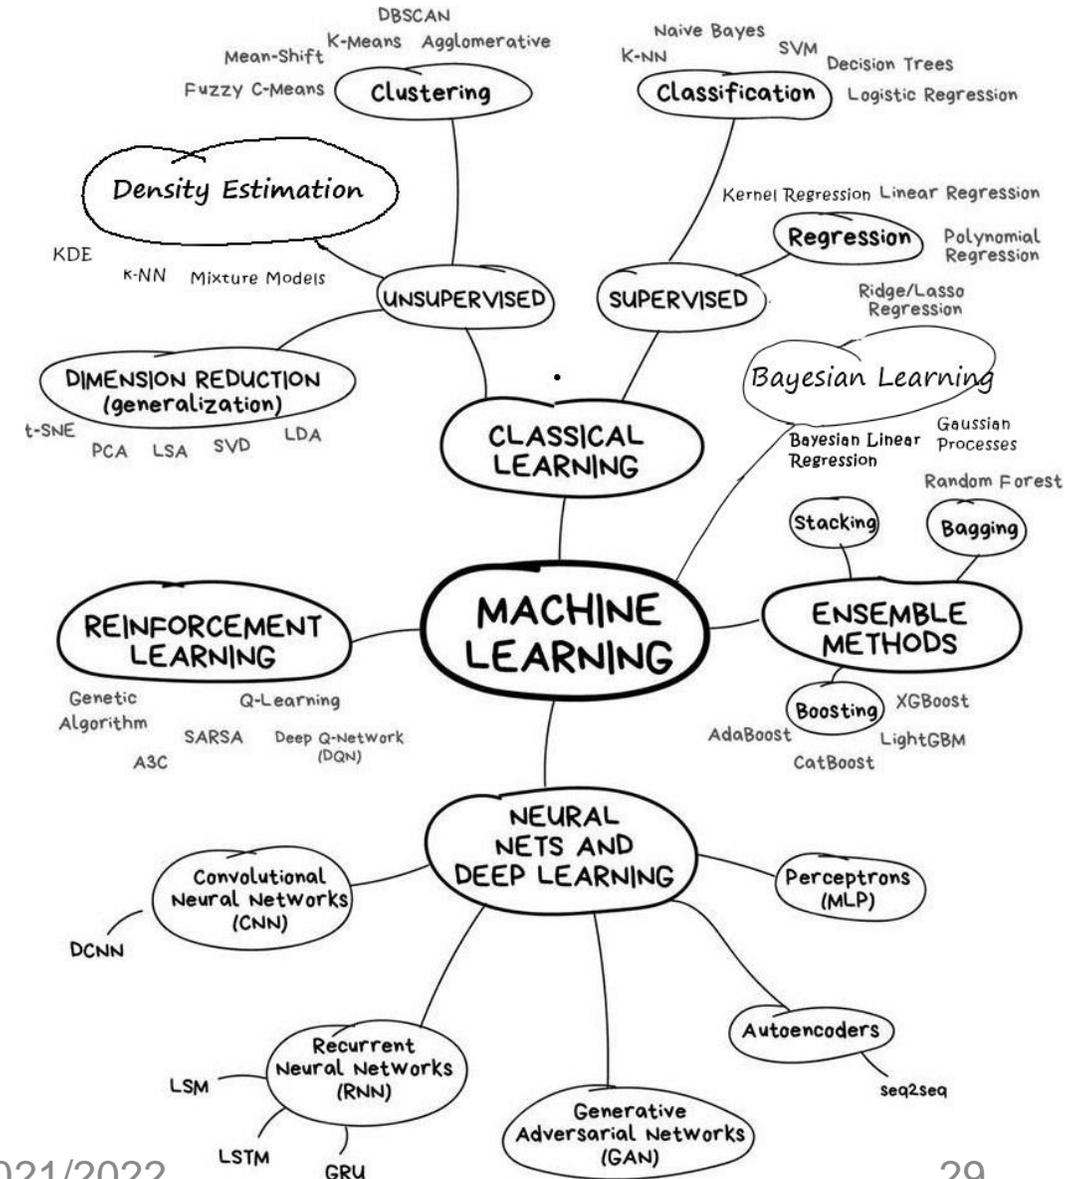
- ✓ “Completing the Square”
- ✓ Gaussian Conditioning

Chapter 4: Neural Networks

- ✓ Multivariate chain rule

Chapter 5: Unsupervised Learning

- ✓ KL-divergences
- ✓ Reparametrization trick



The ML algorithm “coordinate system”

Most ML algorithms can be grouped along 3 axis:

- **Representation:** What is the underlying representation of our model?
- **Loss function:** How do we define what is a good and what is a poor model?
- **Optimization:** How do we optimize?

... of course more axis exists, e.g. Regularization

Intermediate Lecture Wrap-Up – Representations

Chapter 1: Classical Supervised Learning

- ✓ Features / Basis Functions: Linear (Ridge) Regression, Logistic Regression
- ✓ Instances: k-NN
- ✓ Trees: CART
- ✓ Ensembles: Forests

Chapter 2: Kernel Methods

- ✓ Kernels: SVM and Kernel Regression

Chapter 3: Bayesian Learning

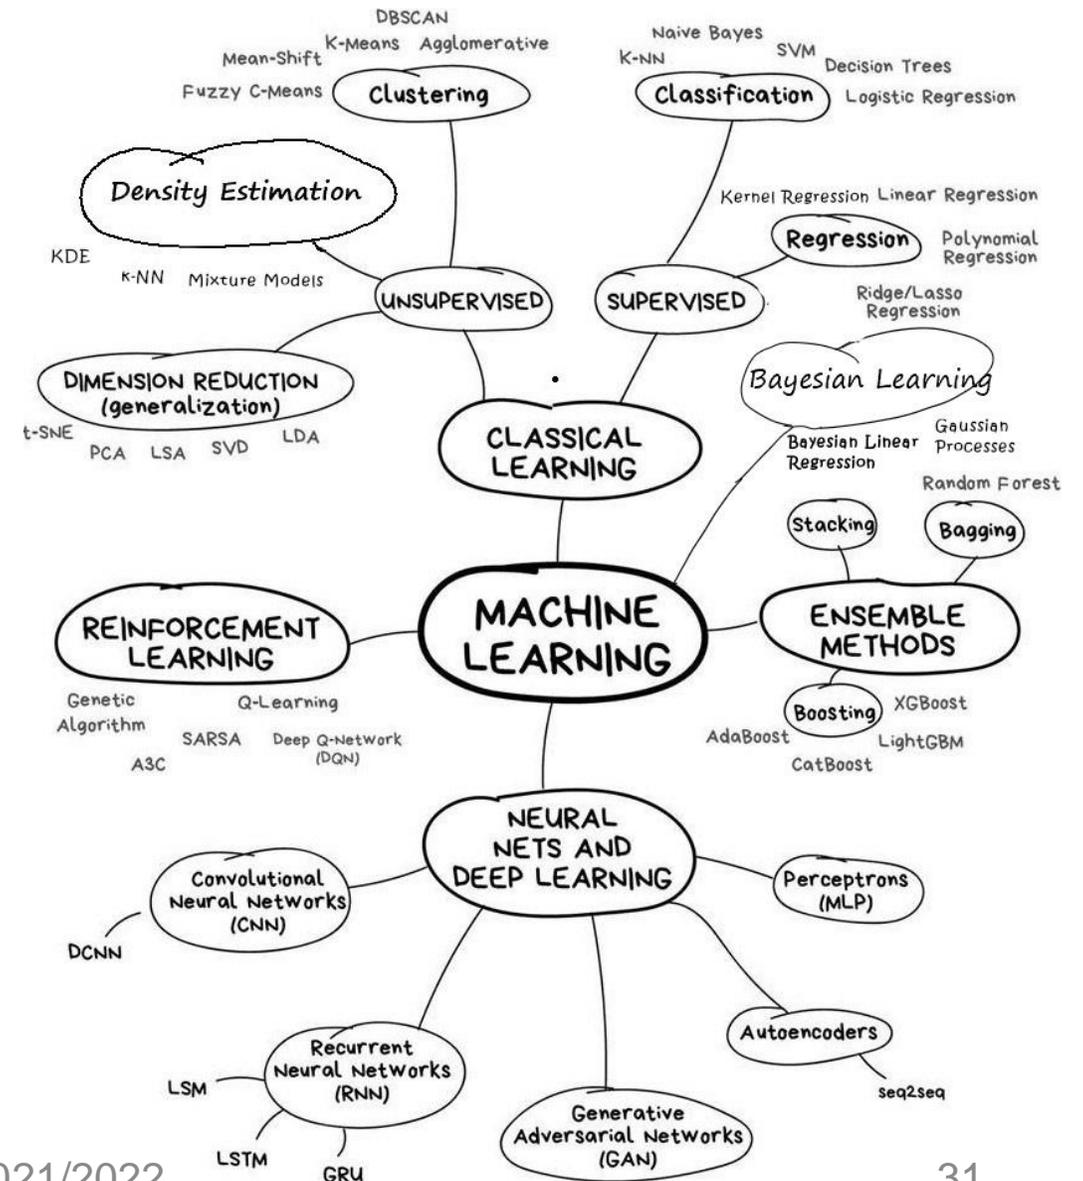
- ✓ Features: Bayesian Linear Regression
- ✓ Kernels: Gaussian Processes

Chapter 4: Neural Networks

- ✓ Feed-forward Neural Networks, CNNs, Recurrent Neural Networks, LSTMs, GRU: Backprop

Chapter 5: Unsupervised Learning

- ✓ Cluster centroids: k-means
- ✓ Linear subspaces: PCA
- ✓ Mixture Models: Expectation Maximization
- ✓ (Variational) Auto-encoders: Variational Bayes



Intermediate Lecture Wrap-Up – Loss Functions

Chapter 1: Classical Supervised Learning

- ✓ Mean/Summed Squared error (SSE): Linear Regression
- ✓ Gaussian Log-Likelihood: Probabilistic linear Regression
- ✓ Binary Cross Entropy Likelihood: Logistic Regression
- ✓ Soft-Max Likelihood: Multi-class classification

Chapter 2: Kernel Methods

- ✓ SSE: Kernel Regression
- ✓ Maximum Margin or Hinge Loss: SVM

Chapter 3: Bayesian Learning

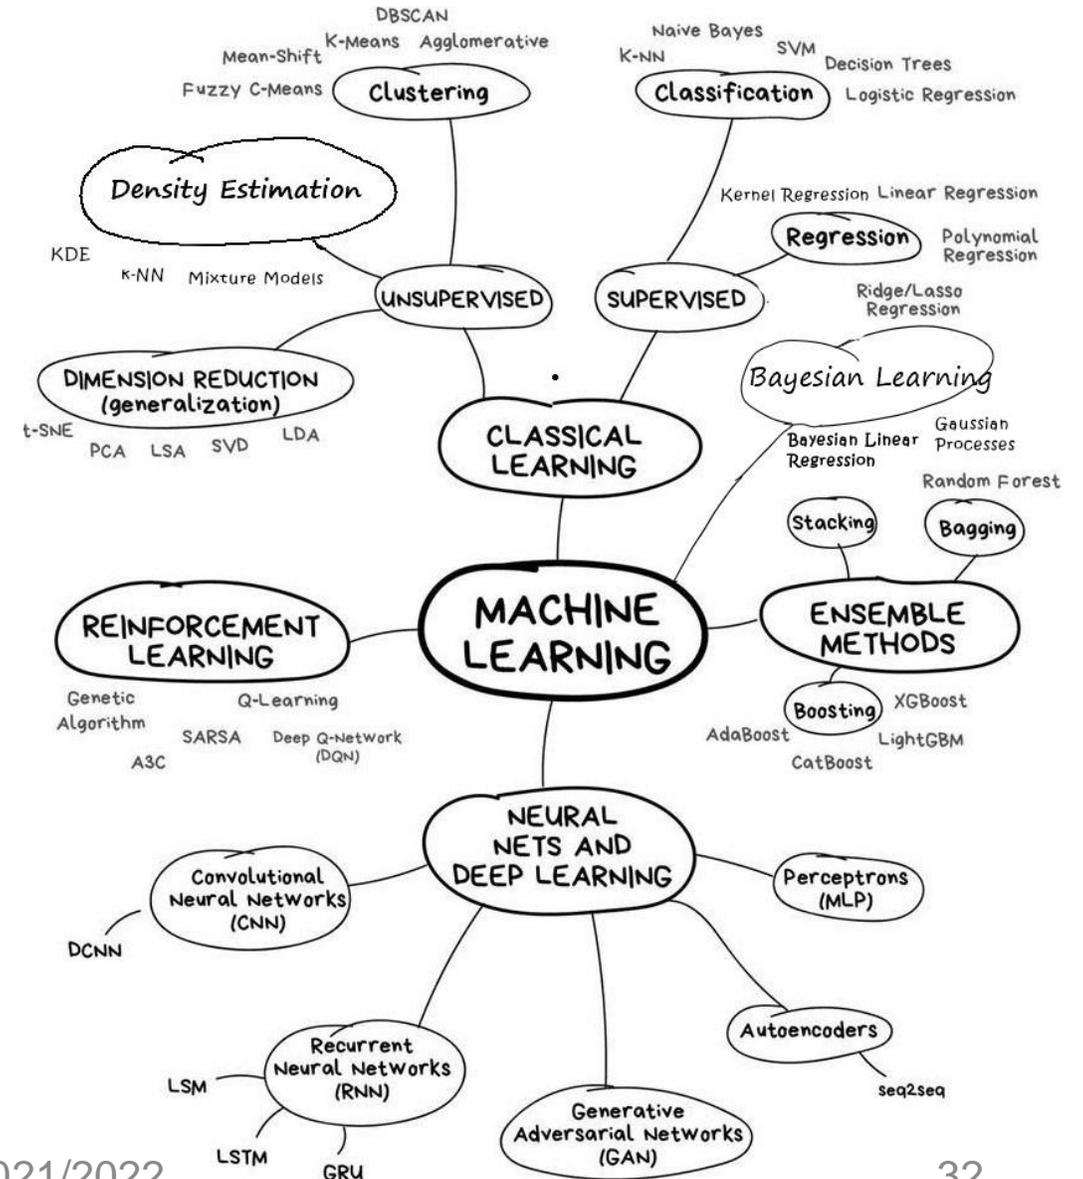
- ✓ Maximum a-posteriori solution: Probabilistic ridge regression

Chapter 4: Neural Networks

- Most of that above...

Chapter 5: Unsupervised Learning

- ✓ Reconstruction Loss: PCA, k-means
- ✓ Marginal Log-likelihoods: EM
- ✓ Evidence Lower Bound (ELBO): Variational Bayes



Intermediate Lecture Wrap-Up – Optimization Methods

Chapter 1: Classical Supervised Learning

- ✓ Least Squares Solution: Linear Regression
- ✓ Gradient Descent: Logistic Regression

Chapter 2: Kernel Methods

- ✓ Least Squares Solution: Kernel Regression
- ✓ Sub-Gradients: SVM
- ✓ Lagrangian Optimization: SVMs

Chapter 3: Bayesian Learning

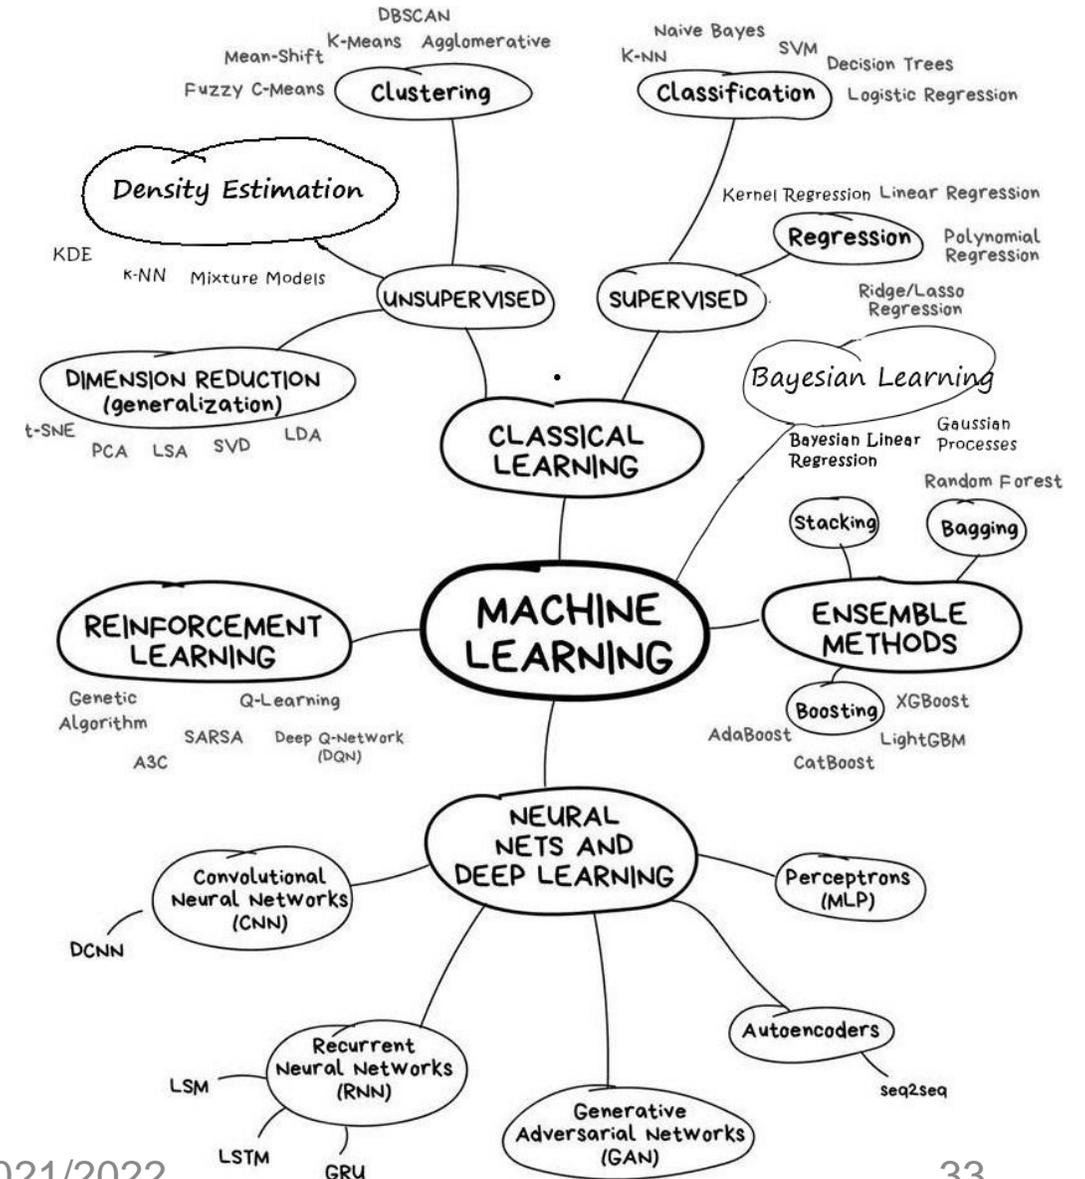
- ✓ Posterior approximation

Chapter 4: Neural Networks

- ✓ More specialized gradient descent methods
- ✓ Adam, 2nd order methods

Chapter 5: Unsupervised Learning

- ✓ Expectation-Maximization
- ✓ Variational Bayes



Where to go from now?

Other ML lectures:

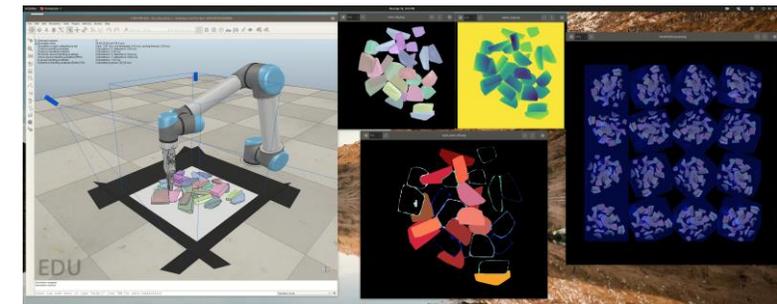
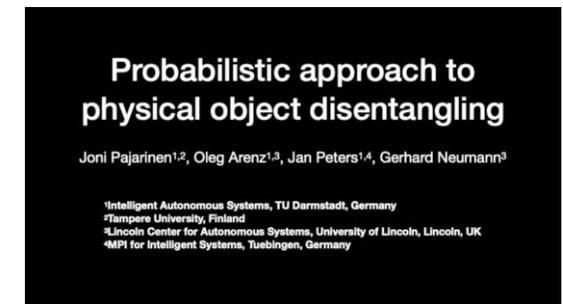
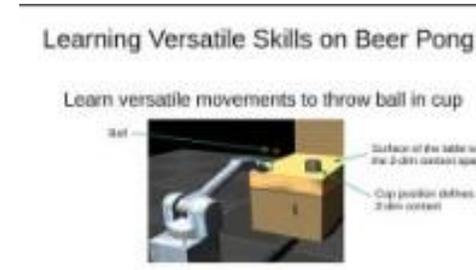
- WS: Reinforcement Learning, **Me**
- SS: Deep Learning and Neural Networks: Prof. Waibel
- SS: Deep Learning for Computer Vision: Prof. Stiefelhagen
- WS: Optimization Methods for Machine Learning and Engineering
- SS: Cognitive Systems, Prof. Waibel and **Me**
- SS: Pattern Recognition, Prof. Beyerer
- SS: Maschinelles Lernen in den Materialwissenschaften, Prof. Friedrich
- SS: Maschinelles Lernen in der Computersicherheit, Prof. Wressnegger

A bit of self-advertisement



What else do we offer?

- **Hot research topics:** Robot Reinforcement Learning, Deep Learning, Imitation Learning, Robotics, Human-Robot Collaboration, Variational Inference
- **Projektpraktikum and Seminar**
 - Work on your own research topic together with your supervisor
 - Get to know latest state of the art algorithms
 - Get experience in doing top-nodge research
- **Praxis der Forschung**
 - 2 semester, 24 ECTS intensive research project
- Interested in a **Master-Thesis or Bachelor Thesis?**
 - Have a look at <https://alr.anthropomatik.kit.edu/>
 - Use real robots (Franka Panda arms)
 - High success-rate of turning your thesis into a paper!
- **Hiwi Positions:**
 - Use robots, cameras, physics simulation, benchmark algorithms etc...



The end

Announcement: Fragestunde, 18.02 16:00

