

RWTH AACHEN
UNIVERSITY

Advanced Machine Learning Lecture 21

Repetition

06.02.2017

Bastian Leibe
RWTH Aachen
<http://www.vision.rwth-aachen.de/>
leibe@vision.rwth-aachen.de

Advanced Machine Learning Winter'16

RWTH AACHEN
UNIVERSITY

Announcements

- Today, I'll summarize the most important points from the lecture.
 - It is an opportunity for you to ask questions...
 - ...or get additional explanations about certain topics.
 - *So, please do ask.*
- Today's slides are intended as an index for the lecture.
 - But they are not complete, won't be sufficient as only tool.
 - Also look at the exercises - they often explain algorithms in detail.
- Exam procedure
 - **Closed-book exam**, the **core exam time will be 2h**.
 - We will send around an announcement with the exact starting times and places by email.

B. Leibe

2

Advanced Machine Learning Winter'16

RWTH AACHEN
UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN
UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN
UNIVERSITY

Recap: Regression

- Learning to predict a continuous function value
 - Given: training set $\mathbf{X} = \{x_1, \dots, x_N\}$ with target values $\mathbf{T} = \{t_1, \dots, t_N\}$.
 - ➔ Learn a continuous function $y(x)$ to predict the function value for a new input x .
- Define an error function $E(\mathbf{w})$ to optimize
 - E.g., **sum-of-squares error**
$$E(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2$$
 - Procedure: Take the derivative and set it to zero
$$\frac{\partial E(\mathbf{w})}{\partial w_j} = \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\} \frac{\partial y(x_n, \mathbf{w})}{\partial w_j} \stackrel{!}{=} 0$$

B. Leibe

5

Advanced Machine Learning Winter'16

RWTH AACHEN
UNIVERSITY

Recap: Least-Squares Regression

$$\mathbf{x}_i^T \mathbf{w} + w_0 = t_i, \quad \forall i = 1, \dots, n$$

- Setup
 - Step 1: Define $\tilde{\mathbf{x}}_i = \begin{pmatrix} \mathbf{x}_i \\ 1 \end{pmatrix}$, $\tilde{\mathbf{w}} = \begin{pmatrix} \mathbf{w} \\ w_0 \end{pmatrix}$
 - Step 2: Rewrite $\tilde{\mathbf{x}}_i^T \tilde{\mathbf{w}} = t_i, \quad \forall i = 1, \dots, n$
 - Step 3: Matrix-vector notation $\tilde{\mathbf{X}}^T \tilde{\mathbf{w}} = \mathbf{t}$ with $\tilde{\mathbf{X}} = [\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n]$ and $\mathbf{t} = [t_1, \dots, t_n]^T$
 - Step 4: Find least-squares solution $\|\tilde{\mathbf{X}}^T \tilde{\mathbf{w}} - \mathbf{t}\|^2 \rightarrow \min$
 - Solution: $\tilde{\mathbf{w}} = (\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T)^{-1} \tilde{\mathbf{X}} \mathbf{t}$

B. Leibe

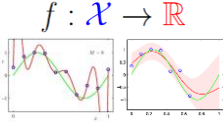
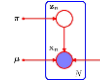
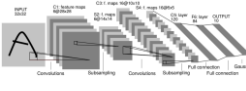
6

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Regularization

- Problem: Overfitting
 - Many parameters & little data \Rightarrow tendency to overfit to the noise
 - Side effect: The coefficient values get very large.
- Workaround: Regularization
 - Penalize large coefficient values
$$\tilde{E}(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N \{y(x_n, \mathbf{w}) - t_n\}^2 + \frac{\lambda}{2} \|\mathbf{w}\|^2$$
 - Here we've simply added a **quadratic regularizer**, which is simple to optimize
$$\|\mathbf{w}\|^2 = \mathbf{w}^T \mathbf{w} = w_0^2 + w_1^2 + \dots + w_M^2$$
 - The resulting form of the problem is called **Ridge Regression**.
 - (Note: w_0 is often omitted from the regularizer.)

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Probabilistic Regression

- First assumption:
 - Our target function values t are generated by adding noise to the ideal function estimate:
$$t = y(\mathbf{x}, \mathbf{w}) + \epsilon$$

Target function value

Regression function

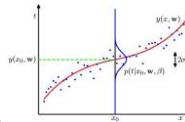
Input value

Noise

Weights or parameters
- Second assumption:
 - The noise is Gaussian distributed.
$$p(t|\mathbf{x}, \mathbf{w}, \beta) = \mathcal{N}(t|y(\mathbf{x}, \mathbf{w}), \beta^{-1})$$

Mean

Variance (β precision)



Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Probabilistic Regression

- Given
 - Training data points: $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$
 - Associated function values: $\mathbf{t} = [t_1, \dots, t_n]^T$
- Conditional likelihood (assuming i.i.d. data)

$$p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = \prod_{n=1}^N \mathcal{N}(t_n|y(\mathbf{x}_n, \mathbf{w}), \beta^{-1}) = \prod_{n=1}^N \mathcal{N}(t_n|\mathbf{w}^T \phi(\mathbf{x}_n), \beta^{-1})$$
 - \Rightarrow Maximize w.r.t. \mathbf{w}, β

Generalized linear regression function

Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Maximum Likelihood Regression

$$\nabla_{\mathbf{w}} \log p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = -\beta \sum_{n=1}^N (t_n - \mathbf{w}^T \phi(\mathbf{x}_n)) \phi(\mathbf{x}_n)$$

- Setting the gradient to zero:

$$0 = -\beta \sum_{n=1}^N (t_n - \mathbf{w}^T \phi(\mathbf{x}_n)) \phi(\mathbf{x}_n)$$

$$\Leftrightarrow \sum_{n=1}^N t_n \phi(\mathbf{x}_n) = \left[\sum_{n=1}^N \phi(\mathbf{x}_n) \phi(\mathbf{x}_n)^T \right] \mathbf{w}$$

$$\Leftrightarrow \Phi \mathbf{t} = \Phi \Phi^T \mathbf{w} \quad \Phi = [\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n)]$$

$$\Leftrightarrow \mathbf{w}_{ML} = (\Phi \Phi^T)^{-1} \Phi \mathbf{t} \quad \leftarrow \text{Same as in least-squares regression!}$$

\Rightarrow Least-squares regression is equivalent to Maximum Likelihood under the assumption of Gaussian noise.

Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Role of the Precision Parameter

- Also use ML to determine the precision parameter β :

$$\log p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = -\frac{\beta}{2} \sum_{n=1}^N \{t_n - \mathbf{w}^T \phi(\mathbf{x}_n)\}^2 + \frac{N}{2} \log \beta - \frac{N}{2} \log(2\pi)$$
- Gradient w.r.t. β :

$$\nabla_{\beta} \log p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = -\frac{1}{2} \sum_{n=1}^N \{t_n - \mathbf{w}^T \phi(\mathbf{x}_n)\}^2 + \frac{N}{2} \frac{1}{\beta}$$

$$\frac{1}{\beta_{ML}} = \frac{1}{N} \sum_{n=1}^N \{t_n - \mathbf{w}^T \phi(\mathbf{x}_n)\}^2$$

\Rightarrow The inverse of the noise precision is given by the residual variance of the target values around the regression function.

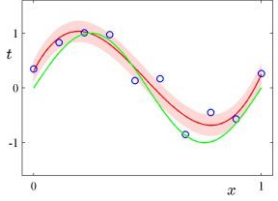
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Predictive Distribution

- Having determined the parameters \mathbf{w} and β , we can now make predictions for new values of \mathbf{x} .

$$p(t|\mathbf{X}, \mathbf{w}_{ML}, \beta_{ML}) = \mathcal{N}(t|y(\mathbf{x}, \mathbf{w}_{ML}), \beta_{ML}^{-1})$$
- This means
 - Rather than giving a point estimate, we can now also give an estimate of the estimation uncertainty.



13
B. Leibe Image source: C.M. Bishop, 2006

RWTH AACHEN UNIVERSITY

Recap: Maximum-A-Posteriori Estimation

- Introduce a prior distribution over the coefficients \mathbf{w} .
 - For simplicity, assume a zero-mean Gaussian distribution

$$p(\mathbf{w}|\alpha) = \mathcal{N}(\mathbf{w}|\mathbf{0}, \alpha^{-1}\mathbf{I}) = \left(\frac{\alpha}{2\pi}\right)^{(M+1)/2} \exp\left\{-\frac{\alpha}{2}\mathbf{w}^T\mathbf{w}\right\}$$
 - New hyperparameter α controls the distribution of model parameters.
- Express the posterior distribution over \mathbf{w} .
 - Using Bayes' theorem:

$$p(\mathbf{w}|\mathbf{X}, \mathbf{t}, \beta, \alpha) \propto p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta)p(\mathbf{w}|\alpha)$$
 - We can now determine \mathbf{w} by maximizing the posterior.
 - This technique is called maximum-a-posteriori (MAP).

14
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: MAP Solution

- Minimize the negative logarithm

$$-\log p(\mathbf{w}|\mathbf{X}, \mathbf{t}, \beta, \alpha) \propto -\log p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) - \log p(\mathbf{w}|\alpha)$$

$$-\log p(\mathbf{t}|\mathbf{X}, \mathbf{w}, \beta) = \frac{\beta}{2} \sum_{n=1}^N \{y(\mathbf{x}_n, \mathbf{w}) - t_n\}^2 + \text{const}$$

$$-\log p(\mathbf{w}|\alpha) = \frac{\alpha}{2} \mathbf{w}^T \mathbf{w} + \text{const}$$
- The MAP solution is therefore

$$\arg \min_{\mathbf{w}} \frac{\beta}{2} \sum_{n=1}^N \{y(\mathbf{x}_n, \mathbf{w}) - t_n\}^2 + \frac{\alpha}{2} \mathbf{w}^T \mathbf{w}$$

\Rightarrow Maximizing the posterior distribution is equivalent to minimizing the regularized sum-of-squares error (with $\lambda = \frac{\alpha}{\beta}$).

15
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: MAP Solution (2)

$$\nabla_{\mathbf{w}} \log p(\mathbf{w}|\mathbf{X}, \mathbf{t}, \beta, \alpha) = -\beta \sum_{n=1}^N (t_n - \mathbf{w}^T \phi(\mathbf{x}_n)) \phi(\mathbf{x}_n) + \alpha \mathbf{w}$$

- Setting the gradient to zero:

$$0 = -\beta \sum_{n=1}^N (t_n - \mathbf{w}^T \phi(\mathbf{x}_n)) \phi(\mathbf{x}_n) + \alpha \mathbf{w}$$

$$\Leftrightarrow \sum_{n=1}^N t_n \phi(\mathbf{x}_n) = \left[\sum_{n=1}^N \phi(\mathbf{x}_n) \phi(\mathbf{x}_n)^T \right] \mathbf{w} + \frac{\alpha}{\beta} \mathbf{w}$$

$$\Leftrightarrow \Phi \mathbf{t} = \left(\Phi \Phi^T + \frac{\alpha}{\beta} \mathbf{I} \right) \mathbf{w} \quad \Phi = [\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_N)]$$

$$\Leftrightarrow \mathbf{w}_{MAP} = \left(\Phi \Phi^T + \frac{\alpha}{\beta} \mathbf{I} \right)^{-1} \Phi \mathbf{t}$$

Effect of regularization:
Keeps the inverse well-conditioned

16
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Bayesian Curve Fitting

- Given
 - Training data points: $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$
 - Associated function values: $\mathbf{t} = [t_1, \dots, t_n]^T$
 - Our goal is to predict the value of t for a new point \mathbf{x} .
- Evaluate the predictive distribution

$$p(t|\mathbf{x}, \mathbf{X}, \mathbf{t}) = \int p(t|\mathbf{x}, \mathbf{w}) p(\mathbf{w}|\mathbf{X}, \mathbf{t}) d\mathbf{w}$$

What we just computed for MAP

 - Noise distribution - again assume a Gaussian here

$$p(t|\mathbf{x}, \mathbf{w}) = \mathcal{N}(t|y(\mathbf{x}, \mathbf{w}), \beta^{-1})$$
 - Assume that parameters α and β are fixed and known for now.

17
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Bayesian Curve Fitting

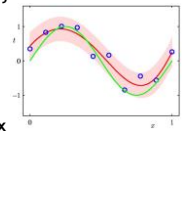
- Under those assumptions, the posterior distribution is a Gaussian and can be evaluated analytically:

$$p(t|\mathbf{x}, \mathbf{X}, \mathbf{t}) = \mathcal{N}(t|m(\mathbf{x}), s^2(\mathbf{x}))$$
 - where the mean and variance are given by

$$m(\mathbf{x}) = \beta \phi(\mathbf{x})^T \mathbf{S} \sum_{n=1}^N \phi(\mathbf{x}_n) t_n$$

$$s(\mathbf{x})^2 = \beta^{-1} + \phi(\mathbf{x})^T \mathbf{S} \phi(\mathbf{x})$$
 - and \mathbf{S} is the regularized covariance matrix

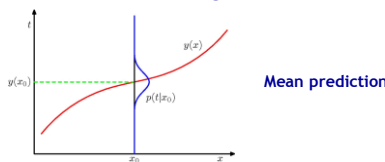
$$\mathbf{S}^{-1} = \alpha \mathbf{I} + \beta \sum_{n=1}^N \phi(\mathbf{x}_n) \phi(\mathbf{x}_n)^T$$



18
B. Leibe Image source: C.M. Bishop, 2006

RWTH AACHEN UNIVERSITY

Recap: Loss Functions for Regression



Mean prediction

- **Optimal prediction**
 - Minimize the expected loss

$$\mathbb{E}[L] = \iint L(t, y(\mathbf{x})) p(\mathbf{x}, t) d\mathbf{x} dt$$
 - Under squared loss, the optimal regression function is the mean $\mathbb{E}[t|\mathbf{x}]$ of the posterior $p(t|\mathbf{x})$ ("mean prediction").
 - For generalized linear regression function and squared loss:

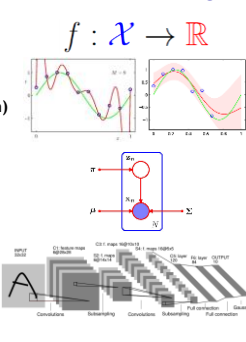
$$y(\mathbf{x}) = \int t \mathcal{N}(t|\mathbf{w}^T \phi(\mathbf{x}), \beta^{-1}) dt = \mathbf{w}^T \phi(\mathbf{x})$$

Slide adapted from Stefan Roth. B. Leibe. Image source: C.M. Bishop, 2006. 19

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- **Regression Approaches**
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- **Approximate Inference**
 - Sampling Approaches
 - MCMC
- **Deep Learning**
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.



B. Leibe

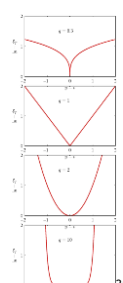
RWTH AACHEN UNIVERSITY

Recap: Loss Functions for Regression

- The squared loss is not the only possible choice
 - Poor choice when conditional distribution $p(t|\mathbf{x})$ is multimodal.
- Simple generalization: **Minkowski loss**

$$L(t, y(\mathbf{x})) = |y(\mathbf{x}) - t|^q$$
 - Expectation

$$\mathbb{E}[L_q] = \iint |y(\mathbf{x}) - t|^q p(\mathbf{x}, t) dx dt$$
- Minimum of $\mathbb{E}[L_q]$ is given by
 - **Conditional mean** for $q = 2$,
 - **Conditional median** for $q = 1$,
 - **Conditional mode** for $q = 0$.



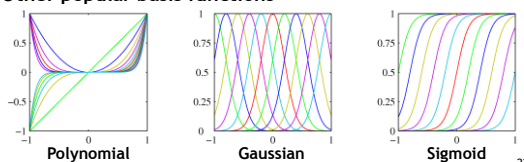
B. Leibe. 21

RWTH AACHEN UNIVERSITY

Recap: Linear Basis Function Models

- Generally, we consider models of the following form

$$y(\mathbf{x}, \mathbf{w}) = \sum_{j=0}^{M-1} w_j \phi_j(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x})$$
 - where $\phi_j(\mathbf{x})$ are known as **basis functions**.
 - In the simplest case, we use linear basis functions: $\phi_d(\mathbf{x}) = x_d$.
- Other popular basis functions
 - **Polynomial**
 - **Gaussian**
 - **Sigmoid**

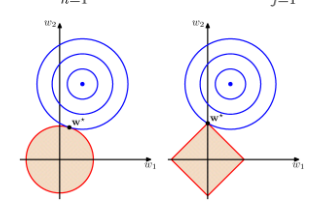


B. Leibe. 22

RWTH AACHEN UNIVERSITY

Recap: Regularized Least-Squares

- Consider more general regularization functions
 - "L_q norms": $\frac{1}{2} \sum_{n=1}^N \{t_n - \mathbf{w}^T \phi(\mathbf{x}_n)\}^2 + \frac{\lambda}{2} \sum_{j=1}^M |w_j|^q$



- Effect: **Sparsity** for $q \leq 1$.
 - Minimization tends to set many coefficients to zero

B. Leibe. Image source: C.M. Bishop, 2006. 23

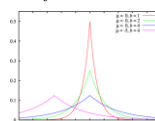
RWTH AACHEN UNIVERSITY

Recap: The Lasso

- L₁ regularization ("The Lasso")

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \frac{1}{2} \sum_{n=1}^N \{t_n - \mathbf{w}^T \phi(\mathbf{x}_n)\}^2 + \lambda \sum_{j=1}^M |w_j|$$
 - The solution will be sparse (only few coefficients non-zero)
 - The L₁ penalty makes the problem non-linear.
 - ⇒ There is no closed-form solution.
- Interpretation as Bayes Estimation
 - We can think of $|w_j|^q$ as the log-prior density for w_j .
- Prior for Lasso ($q = 1$):
 - **Laplacian distribution**

$$p(\mathbf{w}) = \frac{1}{2\tau} \exp\{-|w|/\tau\} \quad \text{with} \quad \tau = \frac{1}{\lambda}$$

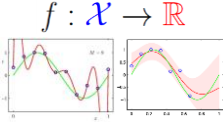
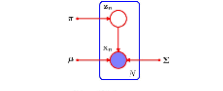
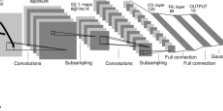


B. Leibe. Image source: Wikipedia. 24

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Kernel Ridge Regression

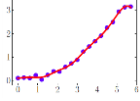
- Dual definition
 - Instead of working with w , substitute $w = \Phi^T a$ into $J(w)$ and write the result using the kernel matrix $K = \Phi \Phi^T$:

$$J(a) = \frac{1}{2} a^T K K a - a^T K t + \frac{1}{2} t^T t + \frac{\lambda}{2} a^T K a$$
 - Solving for a , we obtain

$$a = (K + \lambda I_N)^{-1} t$$
- Prediction for a new input x :
 - Writing $k(x)$ for the vector with elements $k_n(x) = k(x_n, x)$

$$y(x) = w^T \phi(x) = a^T \Phi \phi(x) = k(x)^T (K + \lambda I_N)^{-1} t$$

⇒ The dual formulation allows the solution to be entirely expressed in terms of the kernel function $k(x, x')$.



B. Leibe Image source: Christoph Lampert

RWTH AACHEN UNIVERSITY

Recap: Properties of Kernels

- Theorem
 - Let $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a positive definite kernel function. Then there exists a Hilbert Space \mathcal{H} and a mapping $\phi: \mathcal{X} \rightarrow \mathcal{H}$ such that

$$k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$$
 - where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product in \mathcal{H} .
- Translation
 - Take any set \mathcal{X} and any function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$.
 - If k is a positive definite kernel, then we can use k to learn a classifier for the elements in \mathcal{X} .
- Note
 - \mathcal{X} can be any set, e.g. $\mathcal{X} =$ "all videos on YouTube" or $\mathcal{X} =$ "all permutations of $\{1, \dots, k\}$ ", or $\mathcal{X} =$ "the internet".

Slide credit: Christoph Lampert

RWTH AACHEN UNIVERSITY

Recap: The "Kernel Trick"

Any algorithm that uses data only in the form of inner products can be kernelized.

- How to kernelize an algorithm
 - Write the algorithm only in terms of inner products.
 - Replace all inner products by kernel function evaluations.

⇒ The resulting algorithm will do the same as the linear version, but in the (hidden) feature space \mathcal{H} .

- Caveat: working in \mathcal{H} is not a guarantee for better performance. A good choice of k and model selection are important!

Slide credit: Christoph Lampert

RWTH AACHEN UNIVERSITY

Recap: How to Check if a Function is a Kernel

- Problem:
 - Checking if a given $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ fulfills the conditions for a kernel is difficult:
 - We need to prove or disprove

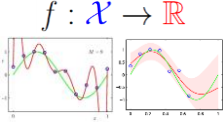
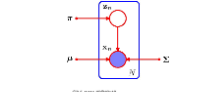
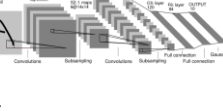
$$\sum_{i,j=1}^n t_i k(x_i, x_j) t_j \geq 0$$
 for any set $x_1, \dots, x_n \in \mathcal{X}$ and any $t \in \mathbb{R}^n$ for any $n \in \mathbb{N}$.
- Workaround:
 - It is easy to construct functions k that are positive definite kernels.

Slide credit: Christoph Lampert

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Gaussian Process

- Gaussian distribution
 - Probability distribution over scalars / vectors.
- Gaussian process (generalization of Gaussian distrib.)
 - Describes properties of functions.
 - Function: Think of a function as a long vector where each entry specifies the function value $f(x_i)$ at a particular point x_i .
 - Issue: How to deal with infinite number of points?
 - If you ask only for properties of the function at a finite number of points...
 - Then inference in Gaussian Process gives you the same answer if you ignore the infinitely many other points.
- Definition
 - A Gaussian process (GP) is a collection of random variables any finite number of which has a joint Gaussian distribution.

Slide credit: Bernt Schiele B. Leibe 31

RWTH AACHEN UNIVERSITY

Recap: Gaussian Process

- A Gaussian process is completely defined by
 - Mean function $m(x)$ and

$$m(x) = \mathbb{E}[f(x)]$$
 - Covariance function $k(x, x')$

$$k(x, x') = \mathbb{E}[(f(x) - m(x))(f(x') - m(x')))]$$
 - We write the Gaussian process (GP)

$$f(x) \sim \mathcal{GP}(m(x), k(x, x'))$$

Slide adapted from Bernt Schiele B. Leibe 32

RWTH AACHEN UNIVERSITY

Recap: GPs Define Prior over Functions

- Distribution over functions:
 - Specification of covariance function implies distribution over functions.
 - I.e. we can draw samples from the distribution of functions evaluated at a (finite) number of points.
 - Procedure
 - We choose a number of input points X_*
 - We write the corresponding covariance matrix (e.g. using SE) element-wise:

$$K(X_*, X_*)$$
 - Then we generate a random Gaussian vector with this covariance matrix:

$$f_* \sim \mathcal{N}(0, K(X_*, X_*))$$

Example of 3 functions sampled

Slide credit: Bernt Schiele B. Leibe Image source: Rasmussen & Williams, 2006 33

RWTH AACHEN UNIVERSITY

Recap: Prediction with Noise-free Observations

- Assume our observations are noise-free:

$$\{(x_n, f_n) \mid n = 1, \dots, N\}$$
 - Joint distribution of the training outputs f and test outputs f_* according to the prior:

$$\begin{bmatrix} f \\ f_* \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(X, X) & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix}\right)$$
 - Calculation of posterior corresponds to conditioning the joint Gaussian prior distribution on the observations:

$$f_* | X_*, X, f \sim \mathcal{N}(\bar{f}_*, \text{cov}[f_*]) \quad \bar{f}_* = \mathbb{E}[f_* | X, X_*, t]$$
 - with:

$$\bar{f}_* = K(X_*, X)K(X, X)^{-1}f$$

$$\text{cov}[f_*] = K(X_*, X_*) - K(X_*, X)K(X, X)^{-1}K(X, X_*)$$

Slide adapted from Bernt Schiele B. Leibe 34

RWTH AACHEN UNIVERSITY

Recap: Prediction with Noisy Observations

- Joint distribution of the observed values and the test locations under the prior:

$$\begin{bmatrix} t \\ f_* \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(X, X) + \sigma_n^2 I & K(X, X_*) \\ K(X_*, X) & K(X_*, X_*) \end{bmatrix}\right)$$
 - Calculation of posterior corresponds to conditioning the joint Gaussian prior distribution on the observations:

$$f_* | X_*, X, t \sim \mathcal{N}(\bar{f}_*, \text{cov}[f_*]) \quad \bar{f}_* = \mathbb{E}[f_* | X, X_*, t]$$
 - with:

$$\bar{f}_* = K(X_*, X) (K(X, X) + \sigma_n^2 I)^{-1} t$$

$$\text{cov}[f_*] = K(X_*, X_*) - K(X_*, X) (K(X, X) + \sigma_n^2 I)^{-1} K(X, X_*)$$

⇒ This is the key result that defines Gaussian process regression!

- Predictive distribution is Gaussian whose mean and variance depend on test points X_* and on the kernel $k(x, x')$, evaluated on X_* .

Slide adapted from Bernt Schiele B. Leibe 35

RWTH AACHEN UNIVERSITY

Recap: GP Regression Algorithm

- Very simple algorithm

input: X (inputs), y (targets), k (covariance function), σ_n^2 (noise level), x_* (test input)

- $L := \text{cholesky}(K + \sigma_n^2 I)$
- $\alpha := L^{-T} \backslash (L y)$
- $\bar{f}_* := k^T \alpha$ } predictive mean eq. (2.25)
- $v := L \backslash k_*$ } predictive variance eq. (2.26)
- $\mathbb{V}[f_*] := k(x_*, x_*) - v^T v$
- $\log p(y|X) := -\frac{1}{2} y^T \alpha - \sum_i \log L_{ii} - \frac{n}{2} \log 2\pi$ eq. (2.30)
- return: \bar{f}_* (mean), $\mathbb{V}[f_*]$ (variance), $\log p(y|X)$ (log marginal likelihood)

- Based on the following equations (Matrix inv. ↔ Cholesky fact.)

$$\bar{f}_* = k_*^T (K + \sigma_n^2 I)^{-1} t$$

$$\text{cov}[f_*] = k(x_*, x_*) - k_*^T (K + \sigma_n^2 I)^{-1} k_*$$

$$\log p(t|X) = -\frac{1}{2} t^T (K + \sigma_n^2 I)^{-1} t - \frac{1}{2} \log |K + \sigma_n^2 I| - \frac{N}{2} \log 2\pi$$

Slide adapted from Bernt Schiele B. Leibe Image source: Rasmussen & Williams, 2006 36

RWTH AACHEN UNIVERSITY

Recap: Computational Complexity

- Complexity of GP model
 - Training effort: $\mathcal{O}(N^3)$ through matrix inversion
 - Test effort: $\mathcal{O}(N^2)$ through vector-matrix multiplication
- Complexity of basis function model
 - Training effort: $\mathcal{O}(M^3)$
 - Test effort: $\mathcal{O}(M^2)$
- Discussion
 - If the number of basis functions M is smaller than the number of data points N , then the basis function model is more efficient.
 - However, advantage of GP viewpoint is that we can consider covariance functions that can only be expressed by an **infinite number of basis functions**.
 - Still, exact GP methods become infeasible for large training sets.

37

RWTH AACHEN UNIVERSITY

Recap: Bayesian Model Selection for GPs

- Goal
 - Determine/learn different parameters of Gaussian Processes
- Hierarchy of parameters
 - Lowest level
 - w - e.g. parameters of a linear model.
 - Mid-level (hyperparameters)
 - θ - e.g. controlling prior distribution of w .
 - Top level
 - Typically discrete set of model structures \mathcal{H}_i .
- Approach
 - Inference takes place one level at a time.

38

RWTH AACHEN UNIVERSITY

Recap: Model Selection at Lowest Level

- Posterior of the parameters w is given by Bayes' rule

$$p(w|t, X, \theta, \mathcal{H}_i) = \frac{p(t|X, w, \theta, \mathcal{H}_i)p(w|\theta, X, \mathcal{H}_i)}{p(t|X, \theta, \mathcal{H}_i)}$$

$$= \frac{p(t|X, w, \mathcal{H}_i)p(w|\theta, \mathcal{H}_i)}{p(t|X, \theta, \mathcal{H}_i)}$$
- with
 - $p(t|X, w, \mathcal{H}_i)$ likelihood and
 - $p(w|\theta, \mathcal{H}_i)$ prior parameters w ,
 - Denominator (normalizing constant) is independent of the parameters and is called **marginal likelihood**.
$$p(t|X, \theta, \mathcal{H}_i) = \int p(t|X, w, \mathcal{H}_i)p(w|\theta, \mathcal{H}_i)dw$$

39

RWTH AACHEN UNIVERSITY

Recap: Model Selection at Mid Level

- Posterior of parameters θ is again given by Bayes' rule

$$p(\theta|t, X, \mathcal{H}_i) = \frac{p(t|X, \theta, \mathcal{H}_i)p(\theta|X, \mathcal{H}_i)}{p(t|X, \mathcal{H}_i)}$$

$$= \frac{p(t|X, \theta, \mathcal{H}_i)p(\theta|\mathcal{H}_i)}{p(t|X, \mathcal{H}_i)}$$
- where
 - The marginal likelihood of the previous level $p(t|X, \theta, \mathcal{H}_i)$ plays the role of the likelihood of this level.
 - $p(\theta|\mathcal{H}_i)$ is the **hyperprior** (prior of the hyperparameters)
 - Denominator (normalizing constant) is given by:

$$p(t|X, \mathcal{H}_i) = \int p(t|X, \theta, \mathcal{H}_i)p(\theta|\mathcal{H}_i)d\theta$$

40

RWTH AACHEN UNIVERSITY

Recap: Model Selection at Top Level

- At the top level, we calculate the posterior of the model

$$p(\mathcal{H}_i|t, X) = \frac{p(t|X, \mathcal{H}_i)p(\mathcal{H}_i)}{p(t|X)}$$
- where
 - Again, the denominator of the previous level $p(t|X, \mathcal{H}_i)$ plays the role of the likelihood.
 - $p(\mathcal{H}_i)$ is the prior of the model structure.
 - Denominator (normalizing constant) is given by:

$$p(t|X) = \sum_i p(t|X, \mathcal{H}_i)p(\mathcal{H}_i)$$

41

RWTH AACHEN UNIVERSITY

Recap: Bayesian Model Selection

- Discussion
 - Marginal likelihood is main difference to non-Bayesian methods

$$p(t|X, \mathcal{H}_i) = \int p(t|X, \theta, \mathcal{H}_i)p(\theta|\mathcal{H}_i)d\theta$$
 - It automatically incorporates a trade-off between the model fit and the model complexity:
 - A simple model can only account for a limited range of possible sets of target values - if a simple model fits well, it obtains a high marginal likelihood.
 - A complex model can account for a large range of possible sets of target values - therefore, it can never attain a very high marginal likelihood.

42

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

Advanced Machine Learning Winter'16
B. Leibe

Recap: Sampling Idea

- Objective:
 - Evaluate expectation of a function $f(\mathbf{z})$ w.r.t. a probability distribution $p(\mathbf{z})$.

$$\mathbb{E}[f] = \int f(\mathbf{z})p(\mathbf{z})d\mathbf{z}$$
- Sampling idea
 - Draw L independent samples $\mathbf{z}^{(l)}$ with $l = 1, \dots, L$ from $p(\mathbf{z})$.
 - This allows the expectation to be approximated by a finite sum

$$\hat{f} = \frac{1}{L} \sum_{l=1}^L f(\mathbf{z}^{(l)})$$
 - As long as the samples $\mathbf{z}^{(l)}$ are drawn independently from $p(\mathbf{z})$, then

$$\|\mathbb{R}[\hat{f}] - \mathbb{R}[f]\|$$

⇒ **Unbiased estimate, independent** of the dimension of \mathbf{z} !

Advanced Machine Learning Winter'16
Slide adapted from Bernt Schiele
B. Leibe
Image source: C.M. Bishop, 2006
44

Recap: Rejection Sampling

- Assumptions
 - Sampling directly from $p(\mathbf{z})$ is difficult.
 - But we can easily evaluate $\tilde{p}(\mathbf{z})$ (up to some norm. factor Z_p):

$$p(\mathbf{z}) = \frac{1}{Z_p} \tilde{p}(\mathbf{z})$$
- Idea
 - We need some simpler distribution $q(\mathbf{z})$ (called **proposal distribution**) from which we can draw samples.
 - Choose a constant k such that: $\forall \mathbf{z} : kq(\mathbf{z}) \geq \tilde{p}(\mathbf{z})$
- Sampling procedure
 - Generate a number z_0 from $q(\mathbf{z})$.
 - Generate a number u_0 from the uniform distribution over $[0, kq(z_0)]$.
 - If $u_0 > \tilde{p}(z_0)$ reject sample, otherwise accept.

Advanced Machine Learning Winter'16
Slide adapted from Bernt Schiele
B. Leibe
Image source: C.M. Bishop, 2006
45

Recap: Sampling from a pdf

- In general, assume we are given the pdf $p(x)$ and the corresponding cumulative distribution:

$$F(x) = \int_{-\infty}^x p(z)dz$$
- To draw samples from this pdf, we can invert the cumulative distribution function:

$$u \sim \text{Uniform}(0, 1) \Rightarrow F^{-1}(u) \sim p(x)$$

Advanced Machine Learning Winter'16
Slide credit: Bernt Schiele
B. Leibe
Image source: C.M. Bishop, 2006
46

Recap: Importance Sampling

- Approach
 - Approximate expectations directly (but does not enable to draw samples from $p(\mathbf{z})$ directly).
 - Goal:
$$\mathbb{E}[f] = \int f(\mathbf{z})p(\mathbf{z})d\mathbf{z}$$
- Idea
 - Use a proposal distribution $q(\mathbf{z})$ from which it is easy to sample.
 - Express expectations in the form of a finite sum over samples $\{\mathbf{z}^{(l)}\}$ drawn from $q(\mathbf{z})$.

$$\mathbb{E}[f] = \int f(\mathbf{z})p(\mathbf{z})d\mathbf{z} = \int f(\mathbf{z}) \frac{p(\mathbf{z})}{q(\mathbf{z})} q(\mathbf{z})d\mathbf{z}$$
 - Approximate as:

$$\simeq \frac{1}{L} \sum_{l=1}^L \underbrace{\frac{p(\mathbf{z}^{(l)})}{q(\mathbf{z}^{(l)})}}_{\text{Importance weights}} f(\mathbf{z}^{(l)})$$

Advanced Machine Learning Winter'16
Slide adapted from Bernt Schiele
B. Leibe
Image source: C.M. Bishop, 2006

Recap: Sampling-Importance-Resampling

- Motivation: Avoid having to determine the constant k .
- Two stages
 - Draw L samples $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(L)}$ from $q(\mathbf{z})$.
 - Construct weights using importance weighting

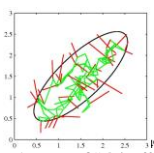
$$w_l = \frac{\tilde{r}_l}{\sum_m \tilde{r}_m} = \frac{\tilde{p}(\mathbf{z}^{(l)})}{\sum_m \tilde{p}(\mathbf{z}^{(m)})}$$
 - and draw a second set of samples $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(L)}$ with probabilities given by the weights $w^{(1)}, \dots, w^{(L)}$.
- Result
 - The resulting L samples are only approximately distributed according to $p(\mathbf{z})$, but the distribution becomes correct in the limit $L \rightarrow \infty$.

Advanced Machine Learning Winter'16
B. Leibe
48

Advanced Machine Learning Winter'16

Recap: MCMC - Markov Chain Monte Carlo

- Overview
 - Allows to sample from a large class of distributions.
 - Scales well with the dimensionality of the sample space.
- Idea
 - We maintain a record of the current state $\mathbf{z}^{(r)}$
 - The proposal distribution depends on the current state: $q(\mathbf{z} | \mathbf{z}^{(r)})$
 - The sequence of samples forms a Markov chain $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \dots$
- Approach
 - At each time step, we generate a candidate sample from the proposal distribution and accept the sample according to a criterion.
 - Different variants of MCMC for different criteria.



Slide adapted from Bernd Schiele B. Leibe Image source: C.M. Bishop, 2006

Advanced Machine Learning Winter'16

Recap: Markov Chains - Properties

- Invariant distribution
 - A distribution is said to be **invariant** (or **stationary**) w.r.t. a Markov chain if each step in the chain leaves that distribution invariant.
 - Transition probabilities:

$$T(\mathbf{z}^{(m)}, \mathbf{z}^{(m+1)}) = p(\mathbf{z}^{(m+1)} | \mathbf{z}^{(m)})$$
 - For homogeneous Markov chain, distribution $p^*(\mathbf{z})$ is invariant if:

$$p^*(\mathbf{z}) = \sum_{\mathbf{z}'} T(\mathbf{z}', \mathbf{z}) p^*(\mathbf{z}')$$
- Detailed balance
 - Sufficient (but not necessary) condition to ensure that a distribution is invariant:

$$p^*(\mathbf{z}) T(\mathbf{z}, \mathbf{z}') = p^*(\mathbf{z}') T(\mathbf{z}', \mathbf{z})$$
 - A Markov chain which respects **detailed balance** is **reversible**.

Slide credit: Bernd Schiele B. Leibe

Advanced Machine Learning Winter'16

Recap: Detailed Balance

- Detailed balance means
 - If we pick a state from the target distribution $p(\mathbf{z})$ and make a transition under T to another state, it is just as likely that we will pick \mathbf{z}_A and go from \mathbf{z}_A to \mathbf{z}_B than that we will pick \mathbf{z}_B and go from \mathbf{z}_B to \mathbf{z}_A .
 - It can easily be seen that a transition probability that satisfies detailed balance w.r.t. a particular distribution will leave that distribution invariant, because

$$\sum_{\mathbf{z}'} p^*(\mathbf{z}') T(\mathbf{z}', \mathbf{z}) = \sum_{\mathbf{z}'} p^*(\mathbf{z}) T(\mathbf{z}, \mathbf{z}') = p^*(\mathbf{z}) \sum_{\mathbf{z}'} p(\mathbf{z}' | \mathbf{z}) = p^*(\mathbf{z})$$

B. Leibe

Advanced Machine Learning Winter'16

Recap: MCMC - Metropolis Algorithm

- Metropolis algorithm [Metropolis et al., 1953]
 - Proposal distribution is symmetric: $q(\mathbf{z}_A | \mathbf{z}_B) = q(\mathbf{z}_B | \mathbf{z}_A)$
 - The new candidate sample \mathbf{z}^* is accepted with probability

$$A(\mathbf{z}^*, \mathbf{z}^{(\tau)}) = \min\left(1, \frac{\tilde{p}(\mathbf{z}^*)}{\tilde{p}(\mathbf{z}^{(\tau)})}\right)$$
 - ⇒ New candidate samples always accepted if $\tilde{p}(\mathbf{z}^*) \geq \tilde{p}(\mathbf{z}^{(\tau)})$.
 - The algorithm sometimes accepts a state with lower probability.
- Metropolis-Hastings algorithm
 - Generalization: Proposal distribution not necessarily symmetric.
 - The new candidate sample \mathbf{z}^* is accepted with probability


$$A(\mathbf{z}^*, \mathbf{z}^{(\tau)}) = \min\left(1, \frac{\tilde{p}(\mathbf{z}^*) q_k(\mathbf{z}^{(\tau)} | \mathbf{z}^*)}{\tilde{p}(\mathbf{z}^{(\tau)}) q_k(\mathbf{z}^* | \mathbf{z}^{(\tau)})}\right)$$
 - where k : labels the members of the set of considered transitions.

Slide adapted from Bernd Schiele B. Leibe

Advanced Machine Learning Winter'16

Recap: Gibbs Sampling

- Approach
 - MCMC-algorithm that is simple and widely applicable.
 - May be seen as a special case of Metropolis-Hastings.
- Idea
 - Sample variable-wise: replace z_i by a value drawn from the distribution $p(z_i | \mathbf{z}_{-i})$.
 - This means we update one coordinate at a time.
 - Repeat procedure either by cycling through all variables or by choosing the next variable.
- Properties
 - The algorithm always accepts!
 - Completely parameter free.
 - Can also be applied to subsets of variables.

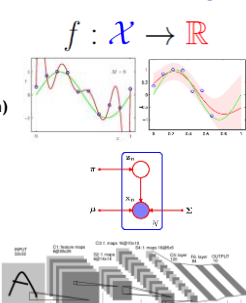


Slide adapted from Bernd Schiele B. Leibe

Advanced Machine Learning Winter'16

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.



B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Linear Discriminant Functions

- **Basic idea**
 - Directly encode decision boundary
 - Minimize misclassification probability directly.
- **Linear discriminant functions**

$$y(\mathbf{x}) = \mathbf{w}^T \mathbf{x} + w_0$$

weight vector
"bias"
(= threshold)

 - \mathbf{w}, w_0 define a hyperplane in \mathbb{R}^D .
 - If a data set can be perfectly classified by a linear discriminant, then we call it **linearly separable**.

Advanced Machine Learning Winter'16 55

Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Generalized Linear Discriminants

- **Extension with non-linear basis functions**
 - Transform vector \mathbf{x} with M nonlinear basis functions $\phi_j(\mathbf{x})$:

$$y_k(\mathbf{x}) = g \left(\sum_{j=1}^M w_{kj} \phi_j(\mathbf{x}) + w_{k0} \right)$$
 - Basis functions $\phi_j(\mathbf{x})$ allow non-linear decision boundaries.
 - Activation function $g(\cdot)$ bounds the influence of outliers.
 - Disadvantage: minimization no longer in closed form.
- **Notation**

$$y_k(\mathbf{x}) = g \left(\sum_{j=0}^M w_{kj} \phi_j(\mathbf{x}) \right) \quad \text{with } \phi_0(\mathbf{x}) = 1$$

Advanced Machine Learning Winter'16 56

Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Gradient Descent

- **Iterative minimization**
 - Start with an initial guess for the parameter values $w_{kj}^{(0)}$.
 - Move towards a (local) minimum by following the gradient.
- **Basic strategies**
 - "Batch learning"

$$w_{kj}^{(\tau+1)} = w_{kj}^{(\tau)} - \eta \left. \frac{\partial E(\mathbf{w})}{\partial w_{kj}} \right|_{\mathbf{w}^{(\tau)}}$$
 - "Sequential updating"

$$w_{kj}^{(\tau+1)} = w_{kj}^{(\tau)} - \eta \left. \frac{\partial E_n(\mathbf{w})}{\partial w_{kj}} \right|_{\mathbf{w}^{(\tau)}}$$

where $E(\mathbf{w}) = \sum_{n=1}^N E_n(\mathbf{w})$

Advanced Machine Learning Winter'16 57

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Gradient Descent

- **Example: Quadratic error function**

$$E(\mathbf{w}) = \sum_{n=1}^N (y(\mathbf{x}_n; \mathbf{w}) - t_n)^2$$
- **Sequential updating leads to delta rule (=LMS rule)**

$$w_{kj}^{(\tau+1)} = w_{kj}^{(\tau)} - \eta (y_k(\mathbf{x}_n; \mathbf{w}) - t_{kn}) \phi_j(\mathbf{x}_n)$$

$$= w_{kj}^{(\tau)} - \eta \delta_{kn} \phi_j(\mathbf{x}_n)$$
 - where

$$\delta_{kn} = y_k(\mathbf{x}_n; \mathbf{w}) - t_{kn}$$
 - Simply feed back the input data point, weighted by the classification error.

Advanced Machine Learning Winter'16 58

Slide adapted from Bernt Schiele B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Probabilistic Discriminative Models

- **Consider models of the form**

$$p(C_1 | \phi) = y(\phi) = \sigma(\mathbf{w}^T \phi)$$

with $p(C_2 | \phi) = 1 - p(C_1 | \phi)$
- **This model is called logistic regression.**
- **Properties**
 - Probabilistic interpretation
 - But discriminative method: only focus on decision hyperplane
 - Advantageous for high-dimensional spaces, requires less parameters than explicitly modeling $p(\phi | C_i)$ and $p(C_i)$.

Advanced Machine Learning Winter'16 59

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Logistic Sigmoid

- **Properties**
 - **Definition:** $\sigma(a) = \frac{1}{1 + \exp(-a)}$
 - **Inverse:** $a = \ln \left(\frac{\sigma}{1 - \sigma} \right)$ "logit" function
 - **Symmetry property:**

$$\sigma(-a) = 1 - \sigma(a)$$
 - **Derivative:**

$$\frac{d\sigma}{da} = \sigma(1 - \sigma)$$

Advanced Machine Learning Winter'16 60

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Logistic Regression

- Let's consider a data set $\{\phi_n, t_n\}$ with $n = 1, \dots, N$, where $\phi_n = \phi(\mathbf{x}_n)$ and $t_n \in \{0, 1\}$, $\mathbf{t} = (t_1, \dots, t_N)^T$.
- With $y_n = p(\mathcal{C}_1 | \phi_n)$, we can write the likelihood as

$$p(\mathbf{t} | \mathbf{w}) = \prod_{n=1}^N y_n^{t_n} \{1 - y_n\}^{1-t_n}$$
- Define the error function as the negative log-likelihood

$$E(\mathbf{w}) = -\ln p(\mathbf{t} | \mathbf{w}) = -\sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\}$$
 - This is the so-called **cross-entropy error function**.

61

RWTH AACHEN UNIVERSITY

Recap: Gradient of the Error Function

- Gradient for logistic regression

$$\nabla E(\mathbf{w}) = \sum_{n=1}^N (y_n - t_n) \phi_n$$
- This is the same result as for the Delta (=LMS) rule

$$w_{kj}^{(\tau+1)} = w_{kj}^{(\tau)} - \eta (y_k(\mathbf{x}_n; \mathbf{w}) - t_{kn}) \phi_j(\mathbf{x}_n)$$
- We can use this to derive a sequential estimation algorithm.
 - However, this will be quite slow...
 - More efficient to use 2nd-order Newton-Raphson \Rightarrow IRLS

62

RWTH AACHEN UNIVERSITY

Recap: Iteratively Reweighted Least Squares

- Result of applying Newton-Raphson to logistic regression

$$\begin{aligned} \mathbf{w}^{(\tau+1)} &= \mathbf{w}^{(\tau)} - (\Phi^T \mathbf{R} \Phi)^{-1} \Phi^T (\mathbf{y} - \mathbf{t}) \\ &= (\Phi^T \mathbf{R} \Phi)^{-1} \{ \Phi^T \mathbf{R} \Phi \mathbf{w}^{(\tau)} - \Phi^T (\mathbf{y} - \mathbf{t}) \} \\ &= (\Phi^T \mathbf{R} \Phi)^{-1} \Phi^T \mathbf{z} \end{aligned}$$
 with $\mathbf{z} = \Phi \mathbf{w}^{(\tau)} - \mathbf{R}^{-1} (\mathbf{y} - \mathbf{t})$
- Very similar form to pseudo-inverse (normal equations)
 - But now with non-constant weighing matrix \mathbf{R} (depends on \mathbf{w}).
 - Need to apply normal equations iteratively.
 - \Rightarrow **Iteratively Reweighted Least-Squares (IRLS)**

63

RWTH AACHEN UNIVERSITY

Recap: Softmax Regression

- Multi-class generalization of logistic regression
 - In logistic regression, we assumed binary labels $t_n \in \{0, 1\}$
 - Softmax generalizes this to K values in 1-of- K notation.

$$\mathbf{y}(\mathbf{x}; \mathbf{w}) = \begin{bmatrix} P(y=1|\mathbf{x}; \mathbf{w}) \\ P(y=2|\mathbf{x}; \mathbf{w}) \\ \vdots \\ P(y=K|\mathbf{x}; \mathbf{w}) \end{bmatrix} = \frac{1}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x})} \begin{bmatrix} \exp(\mathbf{w}_1^T \mathbf{x}) \\ \exp(\mathbf{w}_2^T \mathbf{x}) \\ \vdots \\ \exp(\mathbf{w}_K^T \mathbf{x}) \end{bmatrix}$$

- This uses the **softmax** function

$$\frac{\exp(a_k)}{\sum_j \exp(a_j)}$$
- Note: the resulting distribution is normalized.

64

RWTH AACHEN UNIVERSITY

Recap: Softmax Regression Cost Function

- Logistic regression
 - Alternative way of writing the cost function

$$\begin{aligned} E(\mathbf{w}) &= -\sum_{n=1}^N \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\} \\ &= -\sum_{n=1}^N \sum_{k=0}^K \{ \mathbb{I}(t_n = k) \ln P(y_n = k | \mathbf{x}_n; \mathbf{w}) \} \end{aligned}$$
- Softmax regression
 - Generalization to K classes using indicator functions.

$$E(\mathbf{w}) = -\sum_{n=1}^N \sum_{k=1}^K \left\{ \mathbb{I}(t_n = k) \ln \frac{\exp(\mathbf{w}_k^T \mathbf{x})}{\sum_{j=1}^K \exp(\mathbf{w}_j^T \mathbf{x})} \right\}$$
 - $$\nabla_{\mathbf{w}_k} E(\mathbf{w}) = -\sum_{n=1}^N \mathbb{I}(t_n = k) \ln P(y_n = k | \mathbf{x}_n; \mathbf{w})$$

65

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

66

RWTH AACHEN UNIVERSITY

Recap: Perceptrons

- One output node per class

Output layer
Weights
Input layer

- Outputs
 - Linear outputs

$$y_k(\mathbf{x}) = \sum_{i=0}^d W_{ki} x_i$$

With output nonlinearity

$$y_k(\mathbf{x}) = g\left(\sum_{i=0}^d W_{ki} x_i\right)$$

⇒ Can be used to do multidimensional linear regression or multiclass classification.

Slide adapted from Stefan Roth. B. Leibe 67

RWTH AACHEN UNIVERSITY

Recap: Non-Linear Basis Functions

- Straightforward generalization

Output layer
Weights
Feature layer
Mapping (fixed)
Input layer

- Outputs
 - Linear outputs

$$y_k(\mathbf{x}) = \sum_{i=0}^d W_{ki} \phi(x_i)$$

with output nonlinearity

$$y_k(\mathbf{x}) = g\left(\sum_{i=0}^d W_{ki} \phi(x_i)\right)$$

Advanced Machine Learning Winter'16 B. Leibe 68

RWTH AACHEN UNIVERSITY

Recap: Non-Linear Basis Functions

- Straightforward generalization

Output layer
Weights
Feature layer
Mapping (fixed)
Input layer

- Remarks
 - Perceptrons are generalized linear discriminants!
 - Everything we know about the latter can also be applied here.
 - Note: feature functions $\phi(x)$ are kept fixed, not learned!

Advanced Machine Learning Winter'16 B. Leibe 69

RWTH AACHEN UNIVERSITY

Recap: Perceptron Learning

- Process the training cases in some permutation
 - If the output unit is correct, leave the weights alone.
 - If the output unit incorrectly outputs a zero, add the input vector to the weight vector.
 - If the output unit incorrectly outputs a one, subtract the input vector from the weight vector.
- Translation

$$w_{kj}^{(\tau+1)} = w_{kj}^{(\tau)} - \eta (y_k(\mathbf{x}_n; \mathbf{w}) - t_{kn}) \phi_j(\mathbf{x}_n)$$
 - This is the Delta rule a.k.a. LMS rule!

⇒ Perceptron Learning corresponds to 1st-order (stochastic) Gradient Descent of a quadratic error function!

Slide adapted from Geoff Hinton B. Leibe 70

RWTH AACHEN UNIVERSITY

Recap: Loss Functions

- We can now also apply other loss functions
 - L_2 loss ⇒ Least-squares regression

$$L(t, y(\mathbf{x})) = \sum_n (y(\mathbf{x}_n) - t_n)^2$$
 - L_1 loss: ⇒ Median regression

$$L(t, y(\mathbf{x})) = \sum_n |y(\mathbf{x}_n) - t_n|$$
 - Cross-entropy loss ⇒ Logistic regression

$$L(t, y(\mathbf{x})) = - \sum_n \{t_n \ln y_n + (1 - t_n) \ln(1 - y_n)\}$$
 - Hinge loss ⇒ SVM classification

$$L(t, y(\mathbf{x})) = \sum_n [1 - t_n y(\mathbf{x}_n)]_+$$
 - Softmax loss ⇒ Multi-class probabilistic classification

$$L(t, y(\mathbf{x})) = - \sum_n \sum_k \left\{ \mathbb{1}(t_n = k) \ln \frac{\exp(y_k(\mathbf{x}))}{\sum_j \exp(y_j(\mathbf{x}))} \right\}$$

Advanced Machine Learning Winter'16 B. Leibe 71

RWTH AACHEN UNIVERSITY

Recap: Multi-Layer Perceptrons

- Adding more layers

Output layer
Hidden layer
Input layer

- Output

$$y_k(\mathbf{x}) = g^{(2)}\left(\sum_{i=0}^h W_{ki}^{(2)} g^{(1)}\left(\sum_{j=0}^d W_{ij}^{(1)} x_j\right)\right)$$

Slide adapted from Stefan Roth B. Leibe 72

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Learning with Hidden Units

- How can we train multi-layer networks efficiently?
 - Need an efficient way of adapting **all** weights, not just the last layer.
- Idea: Gradient Descent
 - Set up an error function

$$E(\mathbf{W}) = \sum_n L(t_n, y(\mathbf{x}_n; \mathbf{W})) + \lambda \Omega(\mathbf{W})$$
 with a loss $L(\cdot)$ and a regularizer $\Omega(\cdot)$.
 - E.g., $L(t, y(\mathbf{x}; \mathbf{W})) = \sum_n (y(\mathbf{x}_n; \mathbf{W}) - t_n)^2$ **L₂ loss**
 - $\Omega(\mathbf{W}) = \|\mathbf{W}\|_F^2$ **L₂ regularizer ("weight decay")**
 - ⇒ Update each weight $W_{ij}^{(k)}$ in the direction of the gradient $\frac{\partial E(\mathbf{W})}{\partial W_{ij}^{(k)}}$

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Gradient Descent

- Two main steps
 - Computing the gradients for each weight
 - Adjusting the weights in the direction of the gradient
- We consider those two steps separately
 - Computing the gradients: **Backpropagation**
 - Adjusting the weights: **Optimization techniques**

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Backpropagation Algorithm

- Core steps
 - Convert the discrepancy between each output and its target value into an error derivate.

$$E = \frac{1}{2} \sum_{j \in \text{output}} (t_j - y_j)^2$$

$$\frac{\partial E}{\partial y_j} = -(t_j - y_j)$$
 - Compute error derivatives in each hidden layer from error derivatives in the layer above.
 - Use error derivatives w.r.t. activities to get error derivatives w.r.t. the incoming weights

$$\frac{\partial E}{\partial y_j} \rightarrow \frac{\partial E}{\partial w_{ik}}$$

Slide adapted from Geoff Hinton

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Backpropagation Algorithm

$$\frac{\partial E}{\partial z_j} = \frac{\partial y_j}{\partial z_j} \frac{\partial E}{\partial y_j} = y_j(1 - y_j) \frac{\partial E}{\partial y_j}$$

$$\frac{\partial E}{\partial y_i} = \sum_j \frac{\partial z_j}{\partial y_i} \frac{\partial E}{\partial z_j} = \sum_j w_{ij} \frac{\partial E}{\partial z_j}$$

$$\frac{\partial E}{\partial w_{ij}} = \frac{\partial z_j}{\partial w_{ij}} \frac{\partial E}{\partial z_j} = y_i \frac{\partial E}{\partial z_j}$$

- Efficient propagation scheme
 - y_i is already known from forward pass! (Dynamic Programming)
 - ⇒ Propagate back the gradient from layer j and multiply with y_i .

Slide adapted from Geoff Hinton

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: MLP Backpropagation Algorithm

- Forward Pass


```

y(0) = x
for k = 1, ..., l do
  z(k) = W(k)y(k-1)
  y(k) = gk(z(k))
endfor
y = y(l)
E = L(t, y) + λΩ(W)
      
```
- Backward Pass


```

h ← ∂E/∂y = ∂/∂y L(t, y) + λ ∂Ω/∂y
for k = l, l-1, ..., 1 do
  h ← ∂E/∂z(k) = h ⊙ g'k(y(k))
  ∂E/∂W(k) = h y(k-1)T + λ ∂Ω/∂W(k)
  h ← ∂E/∂y(k-1) = W(k)T h
endfor
      
```
- Notes
 - For efficiency, an entire batch of data \mathbf{X} is processed at once.
 - \odot denotes the element-wise product

B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Computational Graphs

Forward-Mode Differentiation ($\frac{\partial}{\partial X}$)

Apply operator $\frac{\partial}{\partial X}$ to every node.

Reverse-Mode Differentiation ($\frac{\partial Z}{\partial}$)

Apply operator $\frac{\partial Z}{\partial}$ to every node.

- Forward differentiation needs one pass per node. Reverse-mode differentiation can compute all derivatives in one single pass.
- ⇒ Speed-up in $\mathcal{O}(\#inputs)$ compared to forward differentiation!

Slide inspired by Christopher Olah. B. Leibe. Image source: Christopher Olah, colah.github.io

RWTH AACHEN UNIVERSITY

Recap: Automatic Differentiation

- Approach for obtaining the gradients
 - Convert the network into a computational graph.
 - Each new layer/module just needs to specify how it affects the forward and backward passes.
 - Apply reverse-mode differentiation.
 - ⇒ Very general algorithm, used in today's Deep Learning packages

Advanced Machine Learning Winter'16. B. Leibe. Image source: Christopher Olah, colah.github.io

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

Advanced Machine Learning Winter'16. B. Leibe. Image source: Christopher Olah, colah.github.io

RWTH AACHEN UNIVERSITY

Recap: Data Augmentation

- Effect
 - Much larger training set
 - Robustness against expected variations
- During testing
 - When cropping was used during training, need to again apply crops to get same image size.
 - Beneficial to also apply flipping during test.
 - Applying several ColorPCA variations can bring another -1% improvement, but at a significantly increased runtime.

Augmented training data (from one original image)

Advanced Machine Learning Winter'16. B. Leibe. Image source: Lucas Bayer

RWTH AACHEN UNIVERSITY

Recap: Normalizing the Inputs

- Convergence is fastest if
 - The mean of each input variable over the training set is zero.
 - The inputs are scaled such that all have the same covariance.
 - Input variables are uncorrelated if possible.
- Advisable normalization steps (for MLPs)
 - Normalize all inputs that an input unit sees to zero-mean, unit covariance.
 - If possible, try to decorrelate them using PCA (also known as Karhunen-Loeve expansion).

Advanced Machine Learning Winter'16. B. Leibe. Image source: Yann LeCun et al., Efficient BackProp (1998)

RWTH AACHEN UNIVERSITY

Recap: Choosing the Right Learning Rate

- Convergence of Gradient Descent
 - Simple 1D example

$$W^{(\tau-1)} = W^{(\tau)} - \eta \frac{dE(W)}{dW}$$
 - What is the optimal learning rate η_{opt} ?
 - If E is quadratic, the optimal learning rate is given by the inverse of the Hessian

$$\eta_{opt} = \left(\frac{d^2 E(W^{(\tau)})}{dW^2} \right)^{-1}$$
 - Advanced optimization techniques try to approximate the Hessian by a simplified form.
 - If we exceed the optimal learning rate, bad things happen!

Advanced Machine Learning Winter'16. B. Leibe. Image source: Yann LeCun et al., Efficient BackProp (1998)

Advanced Machine Learning Winter'16

Recap: Advanced Optimization Techniques

- Momentum
 - Instead of using the gradient to change the position of the weight "particle", use it to change the velocity.
 - Effect: dampen oscillations in directions of high curvature
 - Nesterov-Momentum: Small variation in the implementation
- RMS-Prop
 - Separate learning rate for each weight: Divide the gradient by a running average of its recent magnitude.
- AdaGrad
- AdaDelta
- Adam

Some more recent techniques, work better for some problems. Try them.

B. Leibe 85
Image source: Geoff Hinton

Advanced Machine Learning Winter'16

Recap: Patience

- Saddle points dominate in high-dimensional spaces!

⇒ Learning often doesn't get stuck, you just may have to wait...

B. Leibe 86
Image source: Yoshua Bengio

Advanced Machine Learning Winter'16

Recap: Reducing the Learning Rate

- Final improvement step after convergence is reached
 - Reduce learning rate by a factor of 10.
 - Continue training for a few epochs.
 - Do this 1-3 times, then stop training.
- Effect
 - Turning down the learning rate will reduce the random fluctuations in the error due to different gradients on different minibatches.
- Be careful: Do not turn down the learning rate too soon!
 - Further progress will be much slower after that.

Slide adapted from Geoff Hinton 87
B. Leibe

Advanced Machine Learning Winter'16

Recap: Glorot Initialization

[Glorot & Bengio, '10]

- Variance of neuron activations
 - Suppose we have an input X with n components and a linear neuron with random weights W that spits out a number Y .
 - We want the variance of the input and output of a unit to be the same, therefore $n \text{Var}(W_i)$ should be 1. This means

$$\text{Var}(W_i) = \frac{1}{n} = \frac{1}{n_{in}}$$
 - Or for the backpropagated gradient

$$\text{Var}(W_i) = \frac{1}{n_{out}}$$
 - As a compromise, Glorot & Bengio propose to use

$$\text{Var}(W) = \frac{2}{n_{in} + n_{out}}$$

⇒ Randomly sample the weights with this variance. That's it.

B. Leibe 88

Advanced Machine Learning Winter'16

Recap: He Initialization

[He et al., '15]

- Extension of Glorot Initialization to ReLU units
 - Use Rectified Linear Units (ReLU)

$$g(a) = \max\{0, a\}$$
 - Effect: gradient is propagated with a constant factor

$$\frac{\partial g(a)}{\partial a} = \begin{cases} 1, & a > 0 \\ 0, & \text{else} \end{cases}$$
- Same basic idea: Output should have the input variance
 - However, the Glorot derivation was based on tanh units, linearity assumption around zero does not hold for ReLU.
 - He et al. made the derivations, proposed to use instead

$$\text{Var}(W) = \frac{2}{n_{in}}$$

B. Leibe 89

Advanced Machine Learning Winter'16

Recap: Batch Normalization

[Ioffe & Szegedy '14]

- Motivation
 - Optimization works best if all inputs of a layer are normalized.
- Idea
 - Introduce intermediate layer that centers the activations of the previous layer per minibatch.
 - I.e., perform transformations on all activations and undo those transformations when backpropagating gradients
- Effect
 - Much improved convergence

B. Leibe 90

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Dropout

[Srivastava, Hinton '12]

(a) Standard Neural Net (b) After applying dropout.

- Idea
 - Randomly switch off units during training.
 - Change network architecture for each data point, effectively training many different variants of the network.
 - When applying the trained network, multiply activations with the probability that the unit was set to zero.

⇒ Greatly improved performance

B. Leibe 91

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

This Lecture: Advanced Machine Learning

$f: \mathcal{X} \rightarrow \mathbb{R}$

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: ImageNet Challenge 2012

IMAGENET

- ImageNet
 - ~14M labeled internet images
 - 20k classes
 - Human labels via Amazon Mechanical Turk
- Challenge (ILSVRC)
 - 1.2 million training images
 - 1000 classes
 - Goal: Predict ground-truth class within top-5 responses
 - Currently one of the top benchmarks in Computer Vision

[Deng et al., CVPR'09]

B. Leibe 93

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Convolutional Neural Networks

“LeNet” architecture

- Neural network with specialized connectivity structure
 - Stack multiple stages of feature extractors
 - Higher stages compute more global, more invariant features
 - Classification layer at the end

Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, [Gradient-based learning applied to document recognition](#), Proceedings of the IEEE 86(11): 2278-2324, 1998.

Slide credit: Svetlana Lazebnik B. Leibe 94

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: CNN Structure

- Feed-forward feature extraction
 - Convolve input with learned filters
 - Non-linearity
 - Spatial pooling
 - (Normalization)
- Supervised training of convolutional filters by back-propagating classification error

Slide credit: Svetlana Lazebnik B. Leibe 95

Advanced Machine Learning Winter'16

RWTH AACHEN UNIVERSITY

Recap: Intuition of CNNs

- Convolutional net
 - Share the same parameters across different locations
 - Convolutions with learned kernels
- Learn *multiple* filters
 - E.g. 1000x1000 image
 - 100 filters
 - 10x10 filter size
 - ⇒ only 10k parameters
- Result: Response map
 - size: 1000x1000x100
 - Only memory, not params!

Slide adapted from Marc'Aurelio Ranzato B. Leibe Image source: Yann LeCun 96

Advanced Machine Learning Winter'16

Recap: Convolution Layers

Naming convention:

HEIGHT
WIDTH
DEPTH

- All Neural Net activations arranged in 3 dimensions
 - Multiple neurons all looking at the same input region, stacked in depth
 - Form a single $[1 \times 1 \times \text{depth}]$ depth column in output volume.

Slide credit: FeiFei Li, Andrei Karpathy. B. Leibe 97

Advanced Machine Learning Winter'16

Recap: Activation Maps

Activations: one filter = one depth slice (or activation map)

5x5 filters

Each activation map is a depth slice through the output volume.

Activations maps

Slide adapted from FeiFei Li, Andrei Karpathy. B. Leibe 98

Advanced Machine Learning Winter'16

Recap: Pooling Layers

Single depth slice

1	1	2	4
5	6	7	8
3	2	1	0
1	2	3	4

max pool with 2x2 filters and stride 2

6	8
3	4

- Effect:
 - Make the representation smaller without losing too much information
 - Achieve robustness to translations

Slide adapted from FeiFei Li, Andrei Karpathy. B. Leibe 99

Advanced Machine Learning Winter'16

Recap: AlexNet (2012)

- Similar framework as LeNet, but
 - Bigger model (7 hidden layers, 650k units, 60M parameters)
 - More data (10^6 images instead of 10^3)
 - GPU implementation
 - Better regularization and up-to-date tricks for training (Dropout)

A. Krizhevsky, I. Sutskever, and G. Hinton, [Image Classification with Deep Convolutional Neural Networks](#), NIPS 2012.

Image source: A. Krizhevsky, I. Sutskever and G. E. Hinton, NIPS 2012 100

Advanced Machine Learning Winter'16

Recap: VGGNet (2014/15)

- Main ideas
 - Deeper network
 - Stacked convolutional layers with smaller filters (+ nonlinearity)
 - Detailed evaluation of all components
- Results
 - Improved ILSVRC top-5 error rate to 6.7%.

ConvNet Configuration					
A	A-LRN	B	C	D	E
11 weight layers	11 weight layers	13 weight layers	16 weight layers	16 weight layers	19 weight layers
conv3-64	conv3-64 LRN	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64	conv3-64 conv3-64
conv3-128	conv3-128	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128	conv3-128 conv3-128
conv3-256	conv3-256	conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256	conv3-256 conv3-256
		maxpool	maxpool	maxpool	maxpool
conv3-512	conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512
conv3-512	conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512
		maxpool	maxpool	maxpool	maxpool
conv3-512	conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512
conv3-512	conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512	conv3-512 conv3-512
		maxpool	maxpool	maxpool	maxpool
		FC-4096	FC-4096	FC-4096	FC-4096
		FC-1000	FC-1000	FC-1000	FC-1000
		soft-max	soft-max	soft-max	soft-max

Mainly used

B. Leibe. Image source: Simonyan & Zisserman

Advanced Machine Learning Winter'16

Recap: GoogLeNet (2014)

- Ideas:
 - Learn features at multiple scales
 - Modular structure

Inception module + copies

Auxiliary classification outputs for training the lower layers (deprecated)

Convolution Pooling Softmax Other

(b) Inception module with dimension reductions

B. Leibe. Image source: Szegedy et al. 102

Recap: Residual Networks

AlexNet, 8 layers (ILSVRC 2012) VGG, 19 layers (ILSVRC 2014) ResNet, 152 layers (ILSVRC 2015)

- Core component
 - Skip connections bypassing each layer
 - Better propagation of gradients to the deeper layers
 - This makes it possible to train (much) deeper networks.

$H(x) = F(x) + x$

B. Leibe 103

Recap: Transfer Learning with CNNs

- Train on ImageNet
- If small dataset: fix all weights (treat CNN as fixed feature extractor), retrain only the classifier
- If you have a medium sized dataset, "finetune" instead: use the old weights as initialization, train the full network or only some of the higher layers.

I.e., replace the Softmax layer at the end

Retrain bigger part of the network

Slide credit: Andrej Karpathy 104

Recap: Visualizing CNNs

Image source: M. Zeiler, B. Fergus 105

Recap: Visualizing CNNs

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Slide credit: Yann LeCun 106

Recap: R-CNN for Object Detection

Regions of Interest (RoI) from a proposal method (~2k)

Input image

Warped image regions

Forward each region through ConvNet

Bbox reg SVMs Classify regions with SVMs

Slide credit: Ross Girshick 107

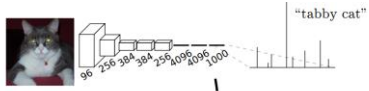
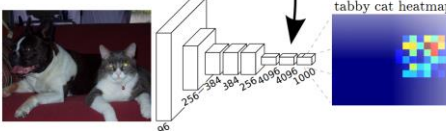
Recap: Faster R-CNN

- One network, four losses
 - Remove dependence on external region proposal algorithm.
 - Instead, infer region proposals from same CNN.
 - Feature sharing
 - Joint training

⇒ Object detection in a single pass becomes possible.

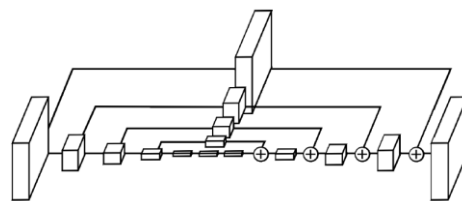
Slide credit: Ross Girshick 108

Recap: Fully Convolutional Networks

- CNN
 
- FCN
 
- Intuition
 - Think of FCNs as performing a sliding-window classification, producing a heatmap of output scores for each class

109
Image source: Long, Shelhamer, Darrell

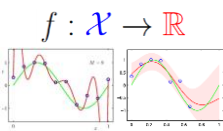
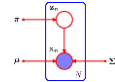
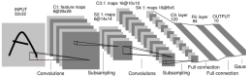
Recap: Image Segmentation Networks



- Encoder-Decoder Architecture
 - Problem: FCN output has low resolution
 - Solution: perform upsampling to get back to desired resolution
 - Use skip connections to preserve higher-resolution information

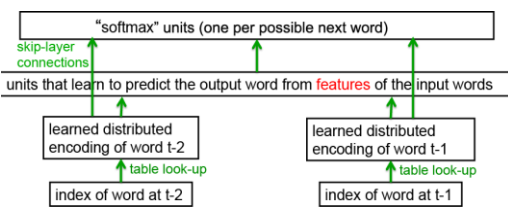
110
Image source: Newell et al

This Lecture: Advanced Machine Learning

- Regression Approaches
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- Approximate Inference
 - Sampling Approaches
 - MCMC
- Deep Learning
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

B. Leibe
113
Image source: Mikolov et al, 2013

Recap: Neural Probabilistic Language Model



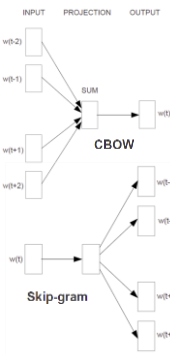
- Core idea
 - Learn a shared distributed encoding (word embedding) for the words in the vocabulary.

Y. Bengio, R. Ducharme, P. Vincent, C. Jauvin, [A Neural Probabilistic Language Model](#), in JMLR, Vol. 3, pp. 1137-1155, 2003.

Slide adapted from Geoff Hinton
B. Leibe
112
Image source: Geoff Hinton

Recap: word2vec

- Goal
 - Make it possible to learn high-quality word embeddings from huge data sets (billions of words in training set).
- Approach
 - Define two alternative learning tasks for learning the embedding:
 - "Continuous Bag of Words" (CBOW)
 - "Skip-gram"
 - Designed to require fewer parameters.

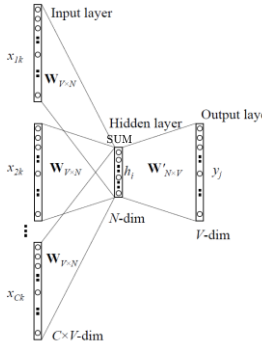


B. Leibe
113
Image source: Mikolov et al, 2013

Recap: word2vec CBOW Model

- Continuous BOW Model
 - Remove the non-linearity from the hidden layer
 - Share the projection layer for all words (their vectors are averaged)

⇒ Bag-of-Words model (order of the words does not matter anymore)



B. Leibe
114
Image source: Yin Peng, 2015

Recap: word2vec Skip-Gram Model

- **Continuous Skip-Gram Model**
 - Similar structure to CBOW
 - Instead of predicting the current word, predict words within a certain range of the current word.
 - Give less weight to the more distant words
- **Implementation**
 - Randomly choose a number $R \in [1, C]$.
 - Use R words from history and R words from the future of the current word as correct labels.
 - ⇒ $R+R$ word classifications for each input.

B. Leibe
Image source: Xin Bing, 2015

Problems with 100k-1M outputs

- **Weight matrix gets huge!**
 - Example: CBOW model
 - One-hot encoding for inputs ⇒ Input-hidden connections are just vector lookups.
 - This is not the case for the hidden-output connections!
 - State h is not one-hot, and vocabulary size is 1M.
 - ⇒ $W'_{N \times V}$ has $300 \times 1M$ entries
- **Softmax gets expensive!**
 - Need to compute normalization over 100k-1M outputs

B. Leibe
Image source: Xin Bing, 2015

Recap: Hierarchical Softmax

- **Idea**
 - Organize words in binary search tree, words are at leaves
 - Factorize probability of word w_0 as a product of node probabilities along the path.
 - Learn a linear decision function $y = v_{n(w,j)} \cdot h$ at each node to decide whether to proceed with left or right child node.
 - ⇒ Decision based on output vector of hidden units directly.

B. Leibe
Image source: Andrej Karpathy

Recap: Recurrent Neural Networks

- **Up to now**
 - Simple neural network structure: 1-to-1 mapping of inputs to outputs
- **Recurrent Neural Networks**
 - Generalize this to arbitrary mappings

B. Leibe
Image source: Andrej Karpathy

Recap: Recurrent Neural Networks (RNNs)

- **RNNs are regular NNs whose hidden units have additional connections over time.**
 - You can **unroll** them to create a network that extends over time.
 - When you do this, keep in mind that the weights for the hidden are shared between temporal layers.
- **RNNs are very powerful**
 - With enough neurons and time, they can compute anything that can be computed by your computer.

B. Leibe
Image source: Andrej Karpathy

Recap: Backpropagation Through Time (BPTT)

- **Configuration**

$$h_t = \sigma(W_{xh}x_t + W_{hh}h_{t-1} + b)$$

$$\hat{y}_t = \text{softmax}(W_{hy}h_t)$$
- **Backpropagated gradient**
 - For weight w_{ij} :
$$\frac{\partial E_t}{\partial w_{ij}} = \sum_{1 \leq k \leq t} \left(\frac{\partial E_t}{\partial h_t} \frac{\partial h_t}{\partial h_k} \frac{\partial^+ h_k}{\partial w_{ij}} \right)$$

B. Leibe
Image source: Andrej Karpathy

Advanced Machine Learning Winter'16

Recap: Backpropagation Through Time (BPTT)

- Analyzing the terms
 - For weight w_{ij} :
$$\frac{\partial E_t}{\partial w_{ij}} = \sum_{1 \leq k \leq t} \left(\frac{\partial E_t}{\partial h_t} \frac{\partial h_t}{\partial h_k} \frac{\partial^+ h_k}{\partial w_{ij}} \right)$$
 - This is the "immediate" partial derivative (with h_{k-1} as constant)

121

Advanced Machine Learning Winter'16

Recap: Backpropagation Through Time (BPTT)

- Analyzing the terms
 - For weight w_{ij} :
$$\frac{\partial E_t}{\partial w_{ij}} = \sum_{1 \leq k \leq t} \left(\frac{\partial E_t}{\partial h_t} \frac{\partial h_t}{\partial h_k} \frac{\partial^+ h_k}{\partial w_{ij}} \right)$$
 - Propagation term:
$$\frac{\partial h_t}{\partial h_k} = \prod_{t \geq i > k} \frac{\partial h_i}{\partial h_{i-1}}$$

122

Advanced Machine Learning Winter'16

Recap: Backpropagation Through Time (BPTT)

- Summary
 - Backpropagation equations

$$E = \sum_{1 \leq t \leq T} E_t$$

$$\frac{\partial E_t}{\partial w_{ij}} = \sum_{1 \leq k \leq t} \left(\frac{\partial E_t}{\partial h_t} \frac{\partial h_t}{\partial h_k} \frac{\partial^+ h_k}{\partial w_{ij}} \right)$$

$$\frac{\partial h_t}{\partial h_k} = \prod_{t \geq i > k} \frac{\partial h_i}{\partial h_{i-1}} = \prod_{t \geq i > k} \mathbf{W}_{hh}^T \text{diag}(\sigma'(\mathbf{h}_{i-1}))$$
 - Remaining issue: how to set the initial state h_0 ?
 ⇒ Learn this together with all the other parameters.

B. Leibe

123

Advanced Machine Learning Winter'16

Recap: Exploding / Vanishing Gradient Problem

- BPTT equations:

$$\frac{\partial E_t}{\partial w_{ij}} = \sum_{1 \leq k \leq t} \left(\frac{\partial E_t}{\partial h_t} \frac{\partial h_t}{\partial h_k} \frac{\partial^+ h_k}{\partial w_{ij}} \right)$$

$$\frac{\partial h_t}{\partial h_k} = \prod_{t \geq i > k} \frac{\partial h_i}{\partial h_{i-1}} = \prod_{t \geq i > k} \mathbf{W}_{hh}^T \text{diag}(\sigma'(\mathbf{h}_{i-1}))$$

$$= (\mathbf{W}_{hh}^T)^l$$

(if l goes to infinity and $l = t - k$.)
- ⇒ We are effectively taking the weight matrix to a high power.
 - The result will depend on the eigenvalues of \mathbf{W}_{hh} .
 - Largest eigenvalue > 1 ⇒ Gradients *may* explode.
 - Largest eigenvalue < 1 ⇒ Gradients *will* vanish.
 - This is very bad...

B. Leibe

124

Advanced Machine Learning Winter'16

Recap: Gradient Clipping

- Trick to handle exploding gradients
 - If the gradient is larger than a threshold, clip it to that threshold.

Algorithm 1 Pseudo-code

```

 $\hat{\mathbf{g}} \leftarrow \frac{\partial \mathcal{E}}{\partial \theta}$ 
if  $\|\hat{\mathbf{g}}\| \geq \text{threshold}$  then
 $\hat{\mathbf{g}} \leftarrow \frac{\text{threshold}}{\|\hat{\mathbf{g}}\|} \hat{\mathbf{g}}$ 
end if
    
```

- This makes a big difference in RNNs

Slide adapted from Richard Socher

B. Leibe

125

Advanced Machine Learning Winter'16

Recap: Long Short-Term Memory (LSTM)

- LSTMs
 - Inspired by the design of memory cells
 - Each module has 4 layers, interacting in a special way.

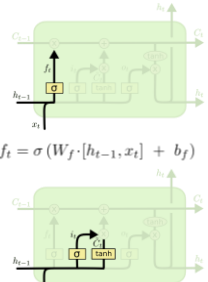
Image source: Christopher Olah - <http://olsh.github.io/posts/2015-08-11-understanding-LSTMs/>

126

Advanced Machine Learning Winter'16

Recap: Elements of LSTMs

- **Forget gate layer**
 - Look at h_{t-1} and x_t and output a number between 0 and 1 for each dimension in the cell state C_{t-1} .
 - 0: completely delete this,
 - 1: completely keep this.
- **Update gate layer**
 - Decide what information to store in the cell state.
 - Sigmoid network (input gate layer) decides which values are updated.
 - tanh layer creates a vector of new candidate values that could be added to the state.



$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$

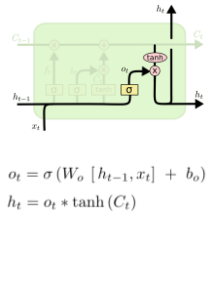
$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

Source: Christopher Olah, <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>

Advanced Machine Learning Winter'16

Recap: Elements of LSTMs

- **Output gate layer**
 - Output is a filtered version of our gate state.
 - First, apply sigmoid layer to decide what parts of the cell state to output.
 - Then, pass the cell state through a tanh (to push the values to be between -1 and 1) and multiply it with the output of the sigmoid gate.



$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o)$$

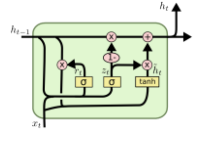
$$h_t = o_t * \tanh(C_t)$$

Source: Christopher Olah, <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>

Advanced Machine Learning Winter'16

Recap: Gated Recurrent Units (GRU)

- **Simpler model than LSTM**
 - Combines the forget and input gates into a single **update gate** z_t .
 - Similar definition for a **reset gate** r_t , but with different weights.
 - In both cases, merge the cell state and hidden state.
- **Empirical results**
 - Both LSTM and GRU can learn much longer-term dependencies than regular RNNs
 - GRU performance similar to LSTM (no clear winner yet), but fewer parameters.



$$z_t = \sigma(W_z \cdot [h_{t-1}, x_t])$$

$$r_t = \sigma(W_r \cdot [h_{t-1}, x_t])$$

$$\tilde{h}_t = \tanh(W \cdot [r_t * h_{t-1}, x_t])$$

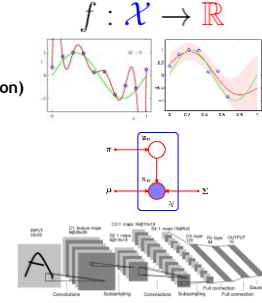
$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$

Source: Christopher Olah, <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>

Advanced Machine Learning Winter'16

This Lecture: Advanced Machine Learning

- **Regression Approaches**
 - Linear Regression
 - Regularization (Ridge, Lasso)
 - Kernels (Kernel Ridge Regression)
 - Gaussian Processes
- **Approximate Inference**
 - Sampling Approaches
 - MCMC
- **Deep Learning**
 - Linear Discriminants
 - Neural Networks
 - Backpropagation & Optimization
 - CNNs, ResNets, RNNs, Deep RL, etc.

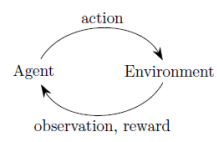


Source: B. Leibe

Advanced Machine Learning Winter'16

Recap: Reinforcement Learning

- **Motivation**
 - General purpose framework for decision making.
 - Basis: **Agent** with the capability to **interact** with its **environment**
 - Each **action** influences the agent's future **state**.
 - Success is measured by a scalar **reward** signal.
 - Goal: **select actions to maximize future rewards**.



- Formalized as a partially observable Markov decision process (POMDP)

Slide adapted from: David Silver, Sergey Levine

Advanced Machine Learning Winter'16

Recap: Reward vs. Return

- **Objective of learning**
 - We seek to maximize the **expected return** G_t as some function of the reward sequence $R_{t+1}, R_{t+2}, R_{t+3}, \dots$
 - Standard choice: **expected discounted return**

$$G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$$

where $0 \leq \gamma \leq 1$ is called the **discount rate**.

- **Difficulty**
 - We don't know which past actions caused the reward.
 - ⇒ Temporal credit assignment problem

Source: B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Policy

- **Definition**
 - A policy determines the agent's behavior
 - Map from state to action $\pi: \mathcal{S} \rightarrow \mathcal{A}$
- **Two types of policies**
 - Deterministic policy: $a = \pi(s)$
 - Stochastic policy: $\pi(a|s) = \Pr\{A_t = a | S_t = s\}$
- **Note**
 - $\pi(a|s)$ denotes the probability of taking action a when in state s .

133

RWTH AACHEN UNIVERSITY

Recap: Value Function

- **Idea**
 - Value function is a prediction of future reward
 - Used to evaluate the goodness/badness of states
 - And thus to select between actions
- **Definition**
 - The **value of a state** s under a policy π , denoted $v_\pi(s)$, is the expected return when starting in s and following π thereafter.

$$v_\pi(s) = \mathbb{E}_\pi[G_t | S_t = s] = \mathbb{E}_\pi[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s]$$
 - The **value of taking action** a in state s under a policy π , denoted $q_\pi(s, a)$, is the expected return starting from s , taking action a , and following π thereafter.

$$q_\pi(s, a) = \mathbb{E}_\pi[G_t | S_t = s, A_t = a] = \mathbb{E}_\pi[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} | S_t = s, A_t = a]$$

134

RWTH AACHEN UNIVERSITY

Recap: Optimal Value Functions

- **Bellman optimality equations**
 - For the **optimal state-value function** v_* :

$$v_*(s) = \max_{a \in \mathcal{A}(s)} q_*(s, a)$$

$$= \max_{a \in \mathcal{A}(s)} \sum_{s', r} p(s', r | s, a) [r + \gamma v_*(s')]$$
 - v_* is the unique solution to this system of nonlinear equations.
 - For the **optimal action-value function** q_* :

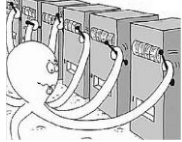
$$q_*(s, a) = \sum_{s', r} p(s', r | s, a) [r + \gamma \max_{a'} q_*(s', a')]$$
 - q_* is the unique solution to this system of nonlinear equations.
- ⇒ If the dynamics of the environment $p(s', r | s, a)$ are known, then in principle one can solve those equation systems.

135

RWTH AACHEN UNIVERSITY

Recap: Exploration-Exploitation Trade-off

- **Example: N-armed bandit problem**
 - Suppose we have the choice between N actions a_1, \dots, a_N .
 - If we knew their value functions $q_*(s, a_i)$, it would be trivial to choose the best.
 - However, we only have estimates based on our previous actions and their returns.
- **We can now**
 - **Exploit** our current knowledge
 - And choose the **greedy** action that has the highest value based on our current estimate.
 - **Explore** to gain additional knowledge
 - And choose a **non-greedy** action to improve our estimate of that action's value.



136

RWTH AACHEN UNIVERSITY

Recap: TD-Learning

- **Policy evaluation (the prediction problem)**
 - For a given policy π , compute the state-value function v_π .
- **One option: Monte-Carlo methods**
 - Play through a sequence of actions until a reward is reached, then backpropagate it to the states on the path.
$$V(S_t) \leftarrow V(S_t) + \alpha [G_t - V(S_t)]$$

Target: the actual return after time t
- **Temporal Difference Learning - TD(λ)**
 - Directly perform an update using the estimate $V(S_{t+\lambda+1})$.

$$V(S_t) \leftarrow V(S_t) + \alpha [R_{t+1} + \gamma V(S_{t+1}) - V(S_t)]$$

Target: an estimate of the return (here: TD(0))

137

RWTH AACHEN UNIVERSITY

Recap: SARSA - On-Policy TD Control

- **Idea**
 - Turn the TD idea into a control method by always updating the policy to be greedy w.r.t. the current estimate
- **Procedure**
 - Estimate $q_\pi(s, a)$ for the current policy π and for all states s and actions a .
 - TD(0) update equation

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha [R_{t+1} + \gamma Q(S_{t+1}, A_{t+1}) - Q(S_t, A_t)]$$
 - This rule is applied after every transition from a nonterminal state S_t .
 - It uses every element of the quintuple $(S_t, A_t, R_{t+1}, S_{t+1}, A_{t+1})$.

⇒ Hence, the name SARSA.

138

RWTH AACHEN UNIVERSITY

Recap: Q-Learning - Off-Policy TD Control

- **Idea**
 - Directly approximate the optimal action-value function q_* , independent of the policy being followed.
- **Procedure**
 - TD(0) update equation

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha \left[R_{t+1} + \gamma \max_a Q(S_{t+1}, a) - Q(S_t, A_t) \right]$$
 - Dramatically simplifies the analysis of the algorithm.
 - All that is required for correct convergence is that all pairs continue to be updated.

139
B. Leibe
Image source: Sutton & Barto

RWTH AACHEN UNIVERSITY

Recap: Deep Q-Learning

- **Idea**
 - Optimal Q-values should obey Bellman equation

$$Q_*(s, a) = \mathbb{E} \left[r + \gamma \max_{a'} Q(s', a') \mid s, a \right]$$
 - Treat the right-hand side $r + \gamma \max_{a'} Q(s', a', w)$ as a target
 - Minimize MSE loss by stochastic gradient descent

$$L(w) = \left(r + \gamma \max_{a'} Q(s', a', w) - Q(s, a, w) \right)^2$$
 - This converges to Q_* using a lookup table representation.
 - Unfortunately, it **diverges** using neural networks due to
 - Correlations between samples
 - Non-stationary targets

140
Slide adapted from David Silver
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Deep Q-Networks (DQN)

- **Adaptation: Experience Replay**
 - To remove correlations, build a dataset from agent's own experience

s_1, a_1, r_1, s_2
s_2, a_2, r_2, s_3
s_3, a_3, r_3, s_4
...
s_t, a_t, r_t, s_{t+1}

 $s_t, a_t, r_{t+1}, s_{t+1} \rightarrow s, a, r, s'$
 - Perform minibatch updates to samples of experience drawn at random from the pool of stored samples
 - $(s, a, r, s') \sim U(D)$ where $D = \{(s_t, a_t, r_{t+1}, s_{t+1})\}$ is the dataset
 - Advantages
 - Each experience sample is used in many updates (more efficient)
 - Avoids correlation effects when learning from consecutive samples
 - Avoids feedback loops from on-policy learning

141
Slide adapted from David Silver
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Deep Q-Networks (DQN)

- **Adaptation: Experience Replay**
 - To remove correlations, build a dataset from agent's own experience

s_1, a_1, r_1, s_2
s_2, a_2, r_2, s_3
s_3, a_3, r_3, s_4
...
s_t, a_t, r_t, s_{t+1}

 $s_t, a_t, r_{t+1}, s_{t+1} \rightarrow s, a, r, s'$
 - Sample from the dataset and apply an update

$$L(w) = \left(r + \gamma \max_{a'} Q(s', a', w^-) - Q(s, a, w) \right)^2$$
 - To deal with non-stationary parameters w^- , are held fixed.
 - Only update the target network parameters every C steps.
 - I.e., clone the network Q to generate a target network Q^- .
 - ➔ Again, this reduces oscillations to make learning more stable.

142
Slide adapted from David Silver
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Policy Gradients

- **How to make high-value actions more likely**
 - The gradient of a stochastic policy $\pi(s, \mathbf{u})$ is given by

$$\frac{\partial L(\mathbf{u})}{\partial \mathbf{u}} = \frac{\partial}{\partial \mathbf{u}} \mathbb{E}_\pi [r_1 + \gamma r_2 + \gamma^2 r_3 + \dots \mid \pi(\cdot, \mathbf{u})]$$

$$= \mathbb{E}_\pi \left[\frac{\partial \log \pi(a|s, \mathbf{u})}{\partial \mathbf{u}} Q_\pi(s, a) \right]$$
 - The gradient of a deterministic policy $a = \pi(s)$ is given by

$$\frac{\partial L(\mathbf{u})}{\partial \mathbf{u}} = \mathbb{E}_\pi \left[\frac{\partial Q_\pi(s, a)}{\partial a} \frac{\partial a}{\partial \mathbf{u}} \right]$$

if a is continuous and Q is differentiable.

143
Slide adapted from David Silver
B. Leibe

RWTH AACHEN UNIVERSITY

Recap: Deep Policy Gradients (DPG)

- **DPG is the continuous analogue of DQN**
 - **Experience replay**: build data-set from agent's experience
 - **Critic** estimates value of current policy by DQN

$$L_w(w) = \left(r + \gamma Q(s', \pi(s', \mathbf{u}^-), w^-) - Q(s, a, w) \right)^2$$
 - To deal with non-stationarity, targets \mathbf{u}^- , w^- are held fixed
 - Actor updates policy in direction that improves Q

$$\frac{\partial L_u(\mathbf{u})}{\partial \mathbf{u}} = \frac{\partial Q(s, a, w)}{\partial a} \frac{\partial a}{\partial \mathbf{u}}$$
 - In other words critic provides loss function for actor.

144
Slide credit: David Silver
B. Leibe

RWTH AACHEN UNIVERSITY

Any Questions?

So what can you do with all of this?

Advanced Machine Learning Winter'16

145

RWTH AACHEN UNIVERSITY

Robust Object Detection & Tracking

Advanced Machine Learning Winter'16

146

RWTH AACHEN UNIVERSITY

Applications for Driver Assistance Systems

Advanced Machine Learning Winter'16

147

RWTH AACHEN UNIVERSITY

Mobile Tracking in Densely Populated Settings

Advanced Machine Learning Winter'16

148
[D. Mitzel, B. Leibe, ECCV'12]

RWTH AACHEN UNIVERSITY

Articulated Multi-Person Tracking

- Multi-Person tracking
 - Recover trajectories and solve data association
- Articulated Tracking
 - Estimate detailed body pose for each tracked person

Advanced Machine Learning Winter'16

149
[Gammeter, Ess, Jaeggli, Schindler, Leibe, Van Gool, ECCV'08]


RWTH AACHEN UNIVERSITY

Semantic 2D-3D Scene Segmentation

Advanced Machine Learning Winter'16

150
B. Leibe [G. Floros, B. Leibe, CVPR'12]

Integrated 3D Point Cloud Labels



B. Leibe

151
[G. Floros, B. Leibe, CVPR'12]

Any More Questions?

Good luck for the exam!

152