Machine Learning

– winter term 2016/17 –

Course Mechanics

Prof. Adrian Ulges
Masters “Computer Science”
DCSM Department
University of Applied Sciences RheinMain

Outline

Food for Thought

Wikification

- Demo should enrich **some** kind of text with Wikipedia
- Live demo!?
- Remember: Choose Wikipedia Category and detect+link key terms
- Training data: Wikipedia (**text + link structure**)
- focus on certain categories?

Gaze Tracking

- steps: camera calibration, landmark detection, normalization, angle estimation
- hardware: webcam? Maybe limit head pose?
- Training phase for user?
- Live Demo!?

Food for Thought

Fish on Kaggle?

- in general, why not?
- Condition: Stick to your **own** approach
- Condition: Live Demo (?)

Outline

1. Projects!

2. 27.10.

3. Introduction

Organisatorisches

Prüfungsanmeldung

- seeehr analog $\rightarrow$ Zettel unterschreiben!
- heute (?)

Pause heute

- im zweiten Block (bitte dran erinnern / laut gähnen o.ä.)

http://video.cs.hs-rm.de (AMIGO Video-Plattform)

- Empfohlen: Firefox (mobiler Zugriff möglich).

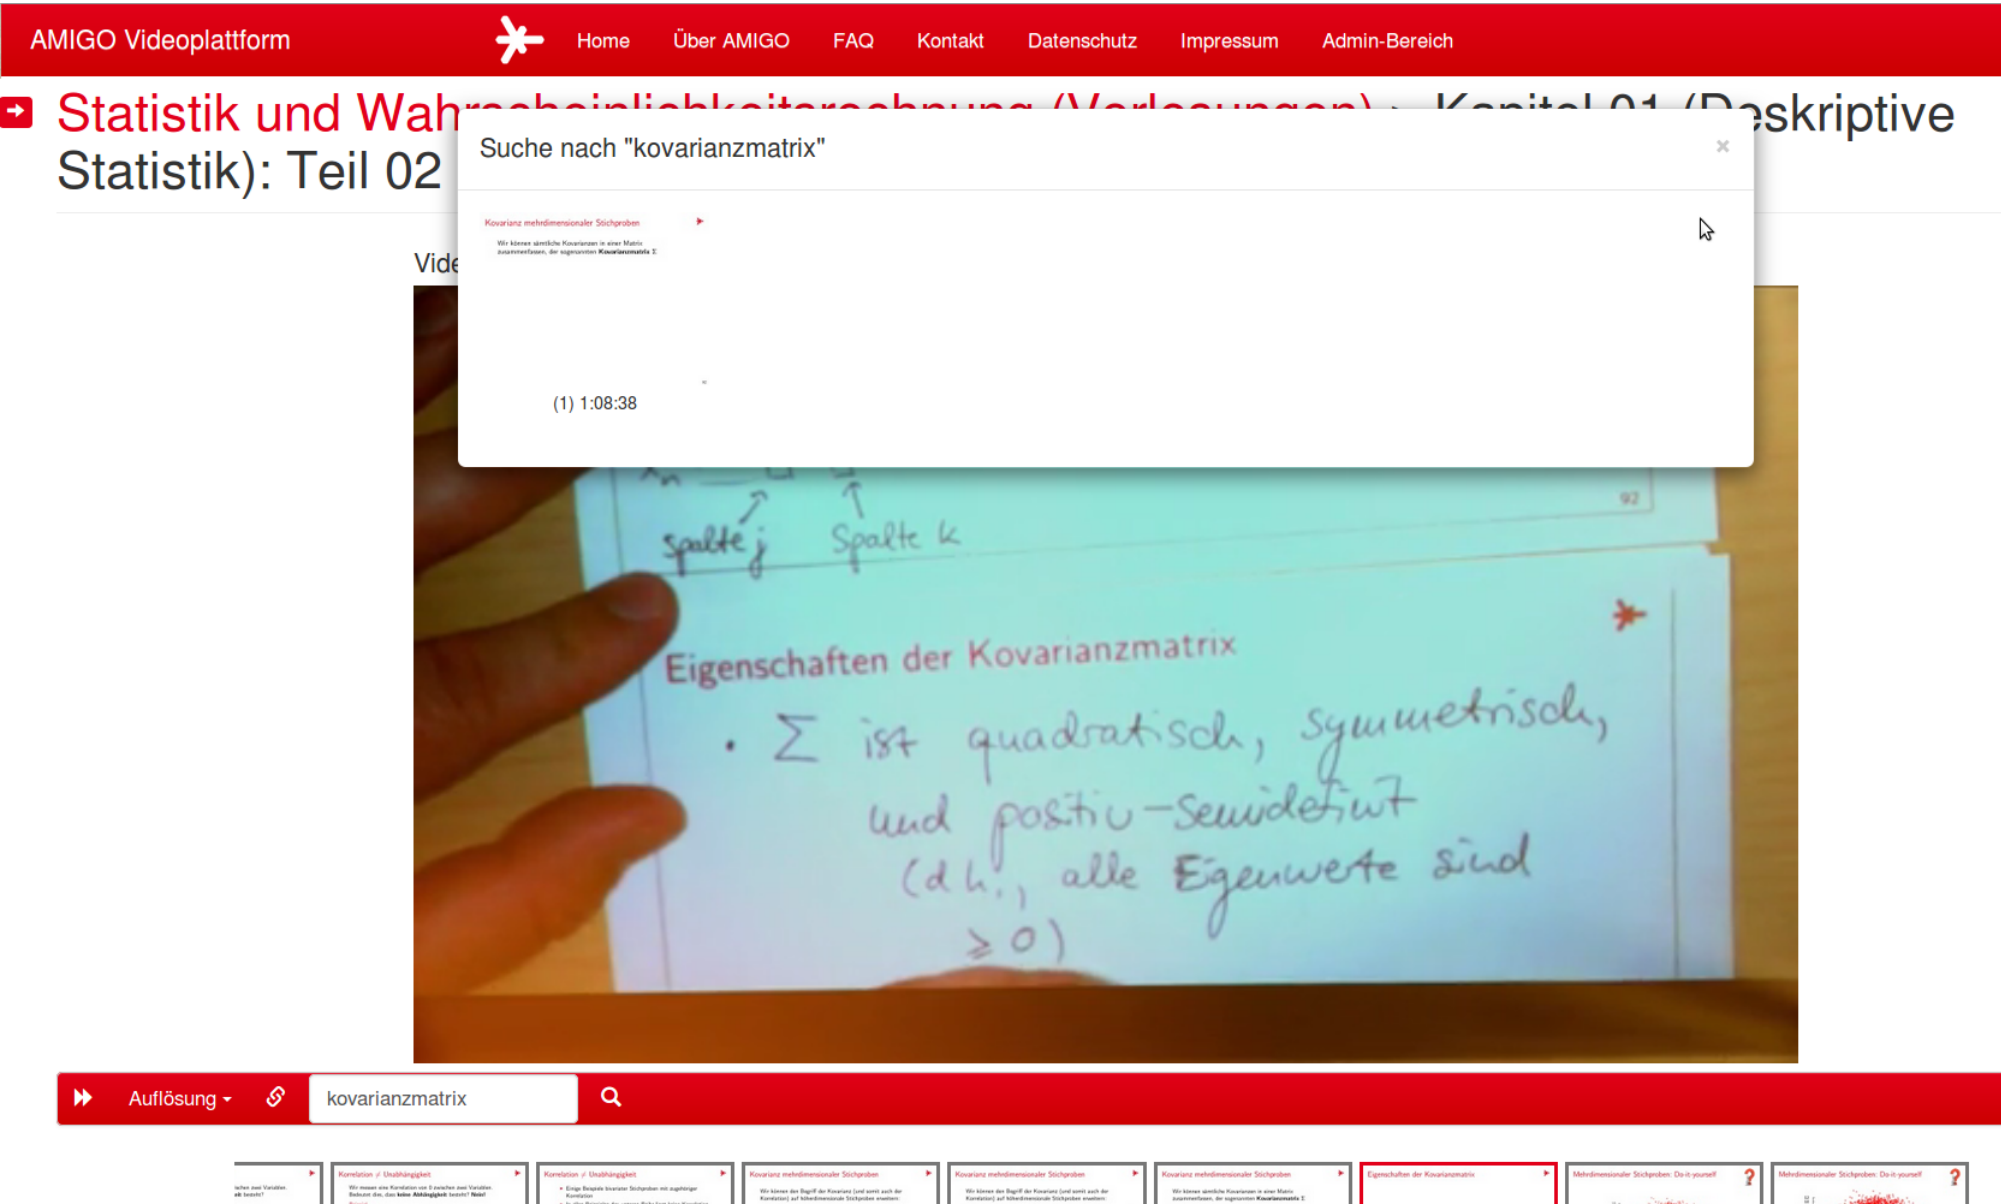

Outline

Videos

- Der Plan ist, diese Vorlesung zu **filmen**
- Der Plan war, die Vorlesungen online zu stellen
- Wie stehen Sie dazu?
- **Problem**: Art 52a UrhG

Course Language: English

But Why?

- ML terminology is mostly English anyway
- For the learning experience

But How?

- Lecturer speaks English
- Students speak whatever they like *(English appreciated)*
- Fachgespräch in German *(English terminology where it feels natural)*
- Coursework in German/English
- Presentations in X (?)
- Vocabulary issues ...? *(speak up!)*

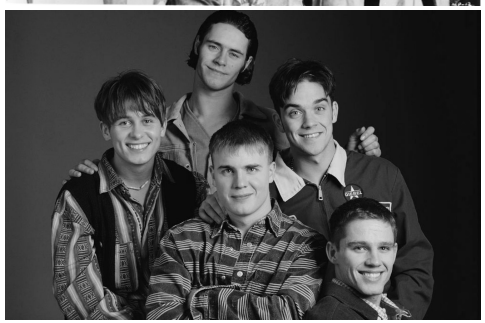

Course Process

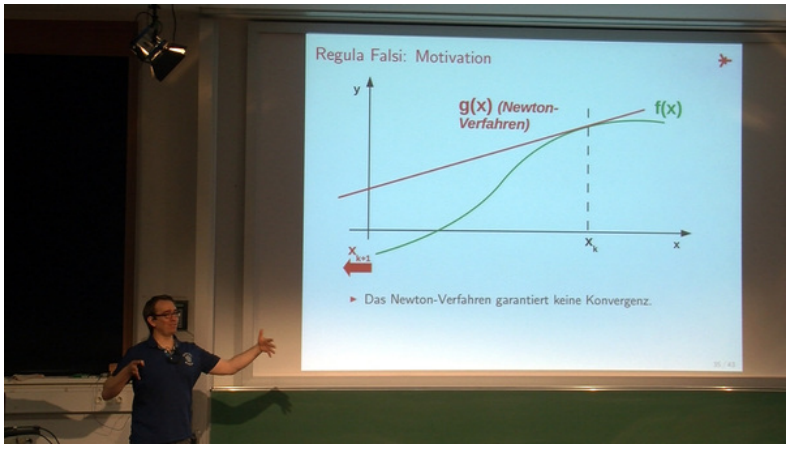

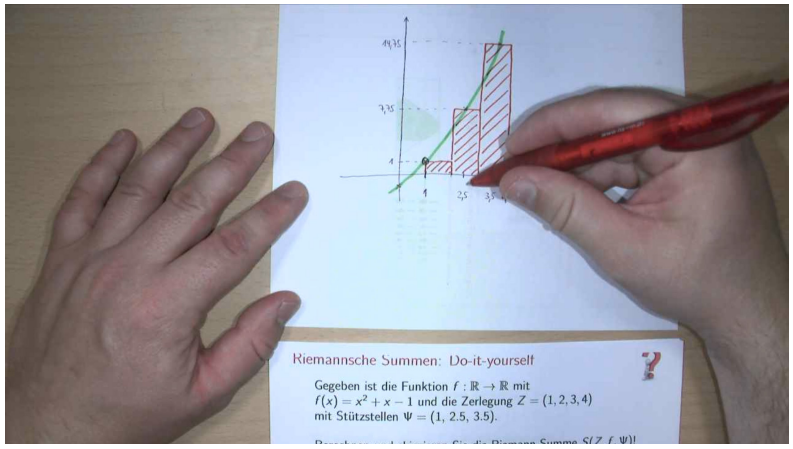

Part 01 (7 Weeks)

- **2 blocks** of Lecture-style presentation
- Some interactive parts ("Do-it-yourselfs"), lots of questions
- Writing on some slides $\rightarrow$ Do print slides beforehand!
- Weekly **coursework**

Part 02 (7 Weeks)

- **Project**: Solving a larger-scale machine learning problem

Time Table (tentative)

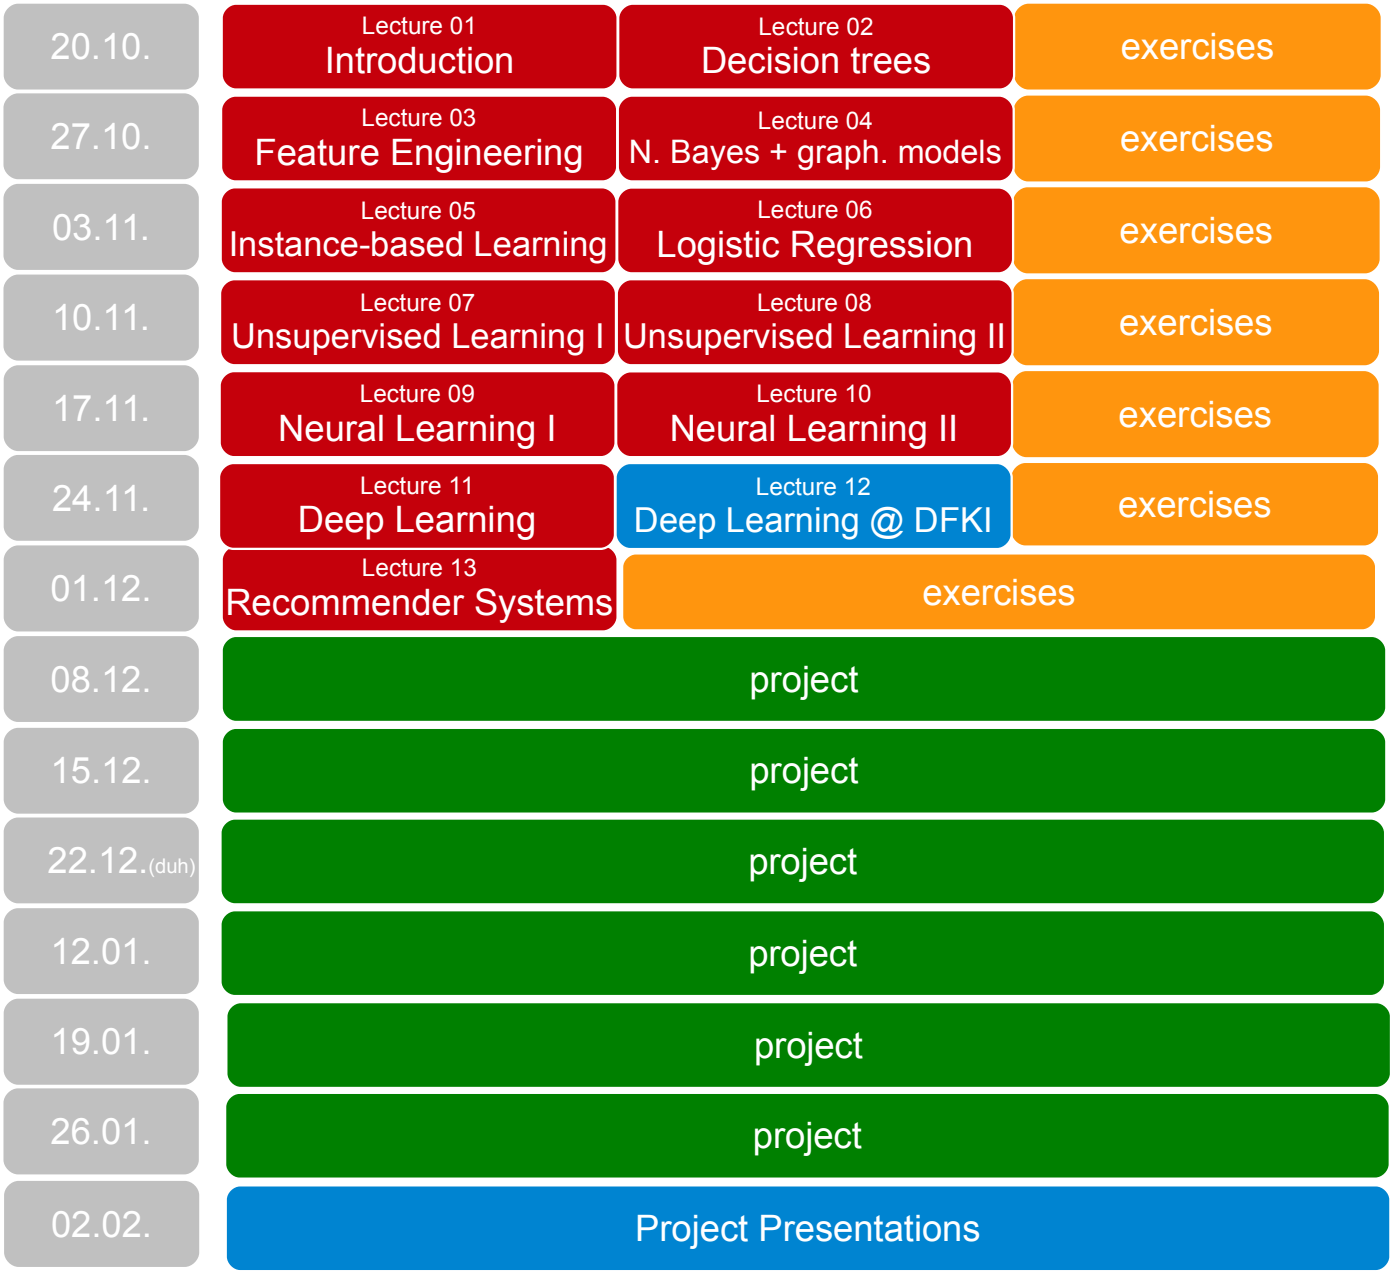

Learning Objectives

After this course, you should ...

... have an algorithmic understanding of the **most common machine learning techniques** (including 9 of the *Top 10 Algorithms in Data Mining*[1])

... be able to **assess** the benefits and shortcomings of ML algorithms

... be able to **apply** ML algorithms using state-of-the-art technology *(Python)*

... have some understanding of the **"dark art" aspects** of ML

... have gone through an **experimental development cycle** of an ML system.

[1] Wu et al., *Top 10 Algorithms in Data Mining, Knowl. Inf. Syst., 2008.*

Ressources

Primary Read

- S. Marsland, *Machine Learning: An Algorithmic Perspective* (Chapman & Hall/CRC).
- Several copies available in the library

Other Good Reads

- Bishop: Pattern Recognition and Machine Learning (Springer)
- Duda Hart Stork: Pattern Classification (Wiley Interscience)
- Specific resources per chapter (later)
- Software: R, (matlab), numpy, scipy, matplotlib, scikit-learn

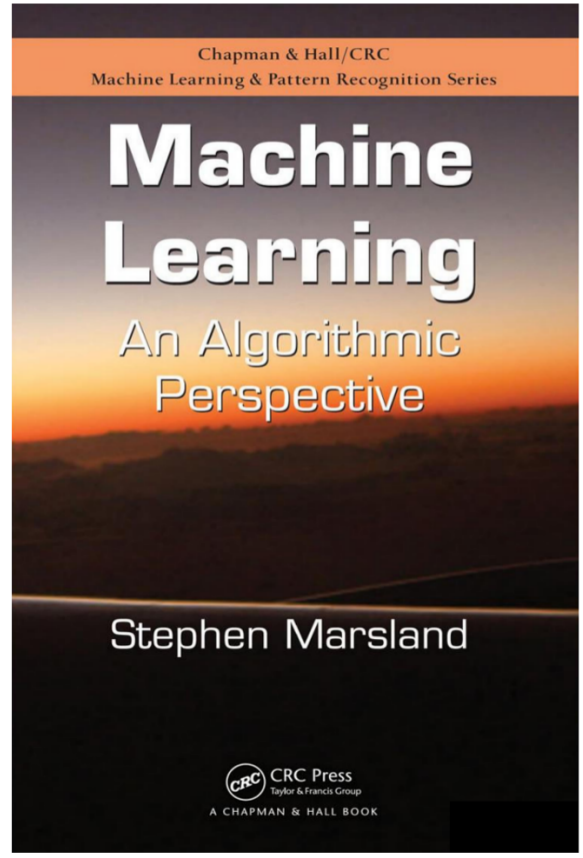

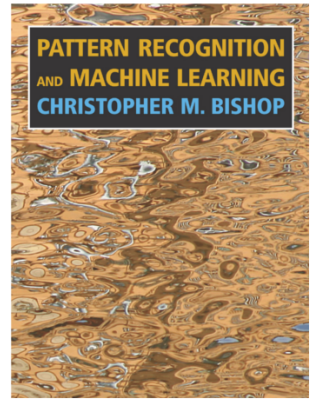

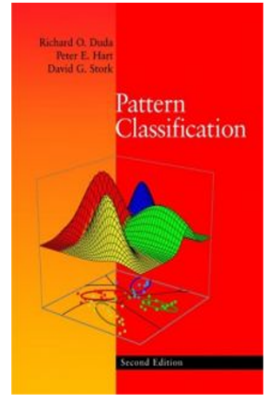

There's more...

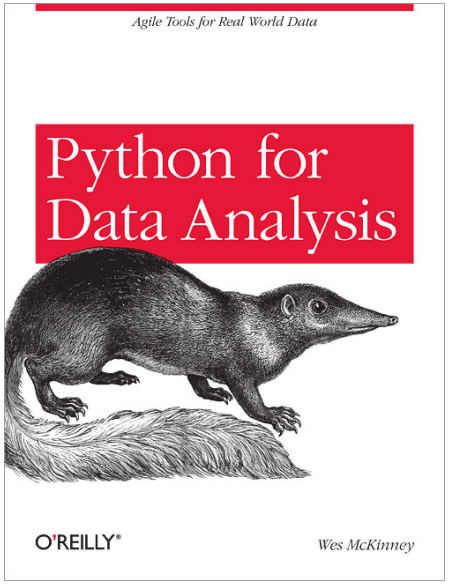
Agile Tools for Real World Data
Python for Data Analysis
O'REILLY
Wes McKinney

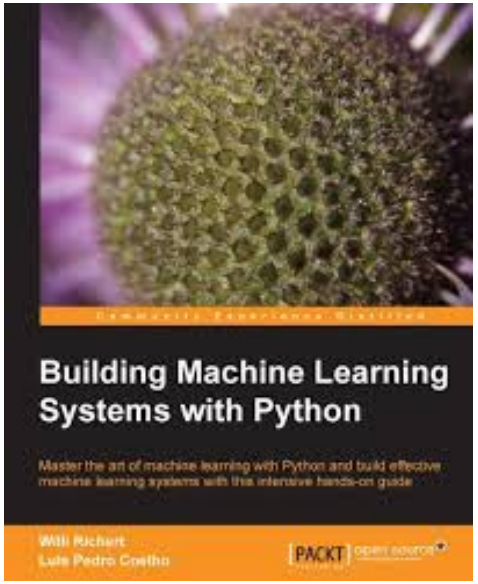
Building Machine Learning Systems with Python
Willi Richert
Luis Pedro Coelho
PACKT

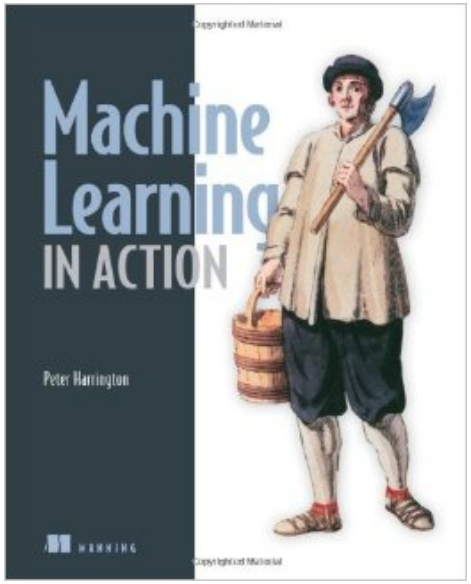
Machine Learning IN ACTION
Peter Harrington
MANNING

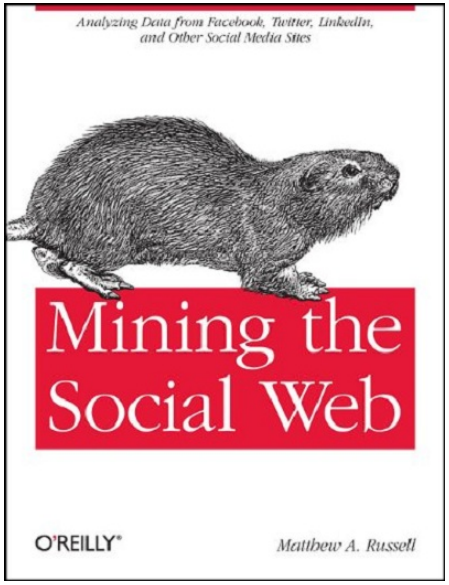
Analyzing Data from Facebook, Twitter, LinkedIn, and Other Social Media Sites
Mining the Social Web
O'REILLY
Matthew A. Russell

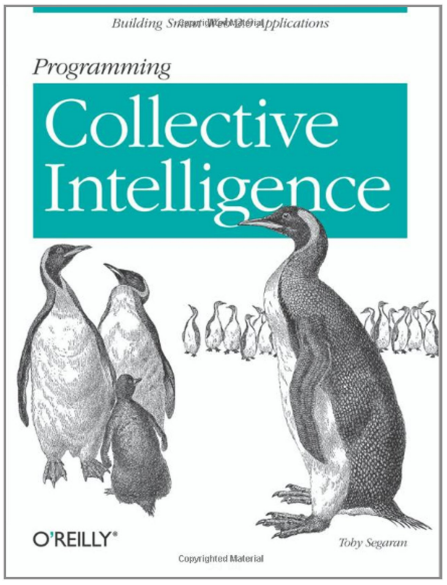
Programming
Collective Intelligence
O'REILLY
Toby Segaran

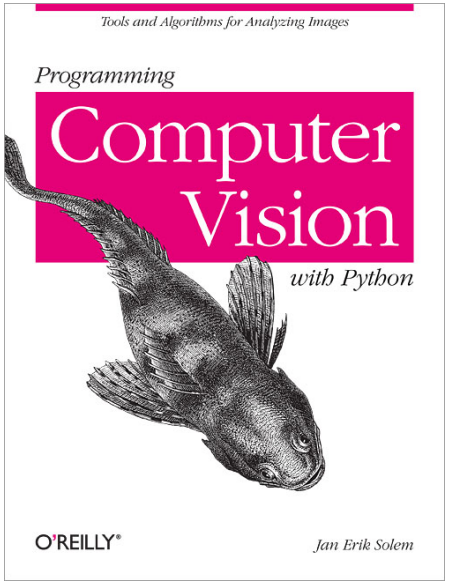
Tools and Algorithms for Analyzing Images
Programming
Computer Vision
with Python
O'REILLY
Jan Erik Solem

Python: Why?

- OO scripting language with functional elements
- easy to learn *("pseudo-code that runs")*, powerful data structures + libraries
- widely used: most popular general-purpose scripting language *(rank 5 in the TIOBE Top Ten[2])*
- very popular for data analysis

What Analytics, Big Data, Data mining, Data Science software you used in the past 12 months for a real project? [2759 voters]

KDnuggets™

Legend: Red: Free/Open Source tools
Green: Commercial tools
Fuchsia: Hadoop/Big Data tools

	% users in 2015	% users in 2014	% users in 2013
R (1293), 3.6% alone	46.9%	38.5%	37.4%
RapidMiner (870), 13.7% alone	31.5%	44.2%	39.2%
SQL (853), 0% alone	30.9%	25.3%	na
Python (837), 0% alone	30.3%	19.5%	13.3%
Excel (631), 0% alone	22.9%	25.8%	28.0%
KNIME (553), 6.7% alone	20%	15.0%	5.9%
Hadoop (507), 0% alone	18.4%	12.7%	9.3%
Tableau (341), 0% alone	12.4%	9.1%	6.3%
SAS base (313), 0.6% alone	11.3%	10.9%	10.7%
Spark (311), 0% alone	11.3%	2.6%	na
Weka (310), 0% alone	11.2%	17.0%	14.3%

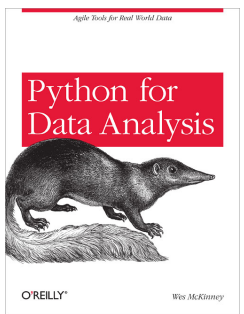

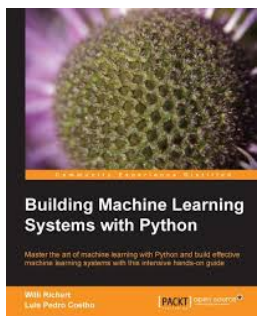

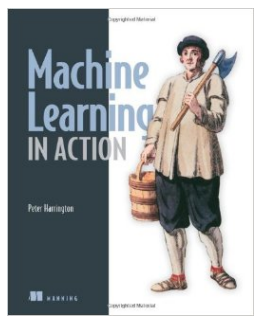

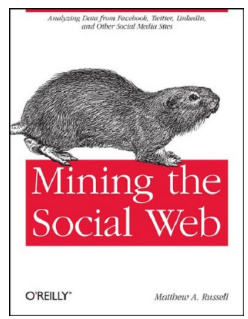

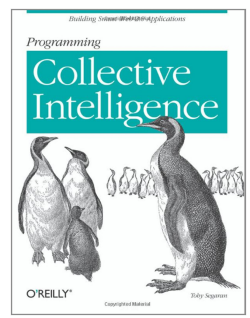

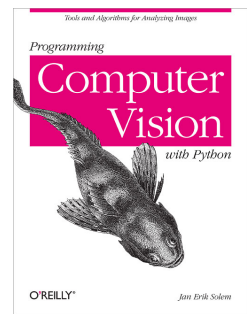

[2] http://www.tiobe.com/index.php/content/paperinfo/tpci/index.html

The Python ML Ecosystem

Numpy

- working with vectors + matrices
- array operations for compact code
- numerical algorithms wrapped

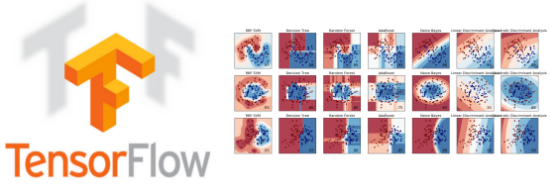

Scikit-Learn (sklearn)

- machine learning algorithms!!!

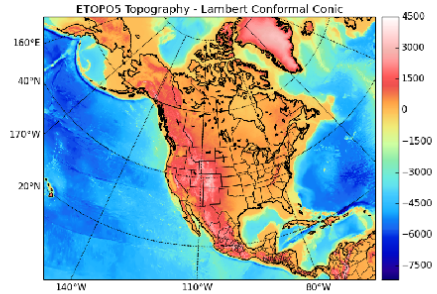

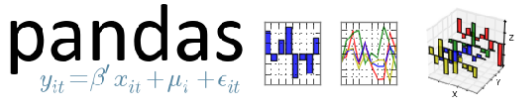

Tensorflow

- neural networks + deep learning
- scalable, portable, GPU support
- Alternatives: Caffe, Theano

Pandas

- 'excel perspective'
- feature engineering

Matplotlib

- visualization

Jupyter Notebook

- run experiments in a Wiki
- demo

Projects

You Will...

- ... work in teams of 2 *(for the whole course, actually)*
- ... solve a machine learning challenge
- ... present your results in a talk *(last week)*
- ... honor the **ground rule** *(do not train on the testing data)*

Challenge 1: Gaze Tracking

- Build a system that detects where on a computer screen the **user is looking**
- eye detection $\rightarrow$ pupil tracking $\rightarrow$ estimating gaze direction
- optionally: adapt to a specific user
- build a demo system
- Technology: Deep CNNs, Tensorflow

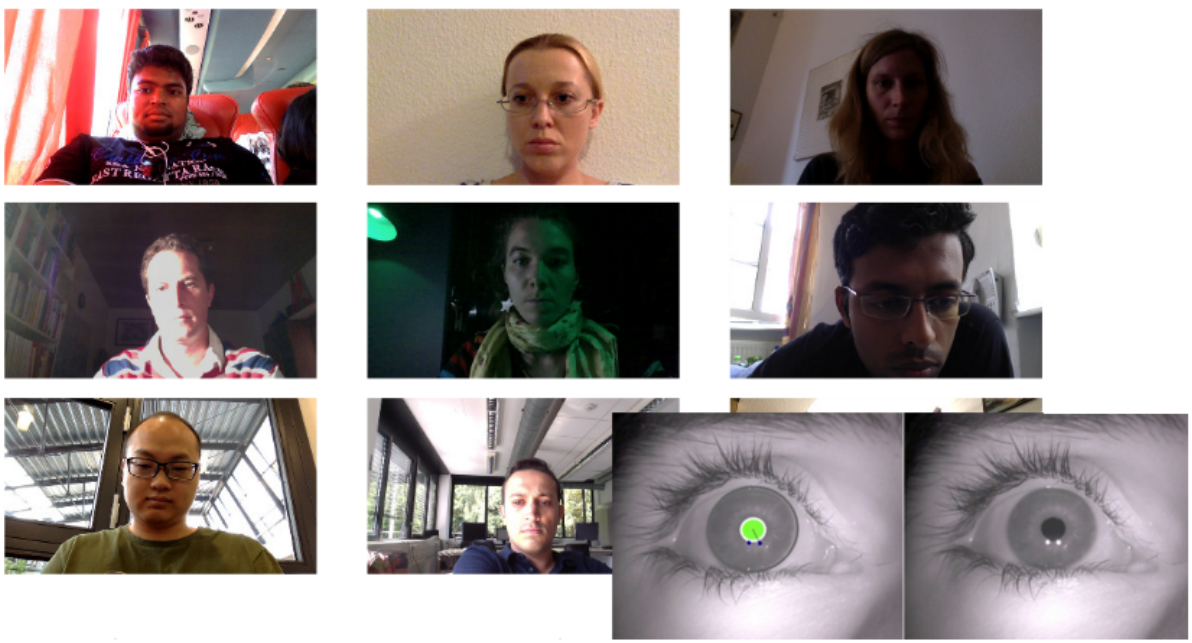

Challenge 2: Wikification

- Vision: Enrich any piece of text with **Wikipedia**
- Step 1: Sort text into the Wikipedia category tree
- Step 2: Detect keywords and link them to Wikipedia articles
- build a demo system
- Training data: Wikipedia (**text + link structure**)
- Technology: free of choice (NLTK recommended)

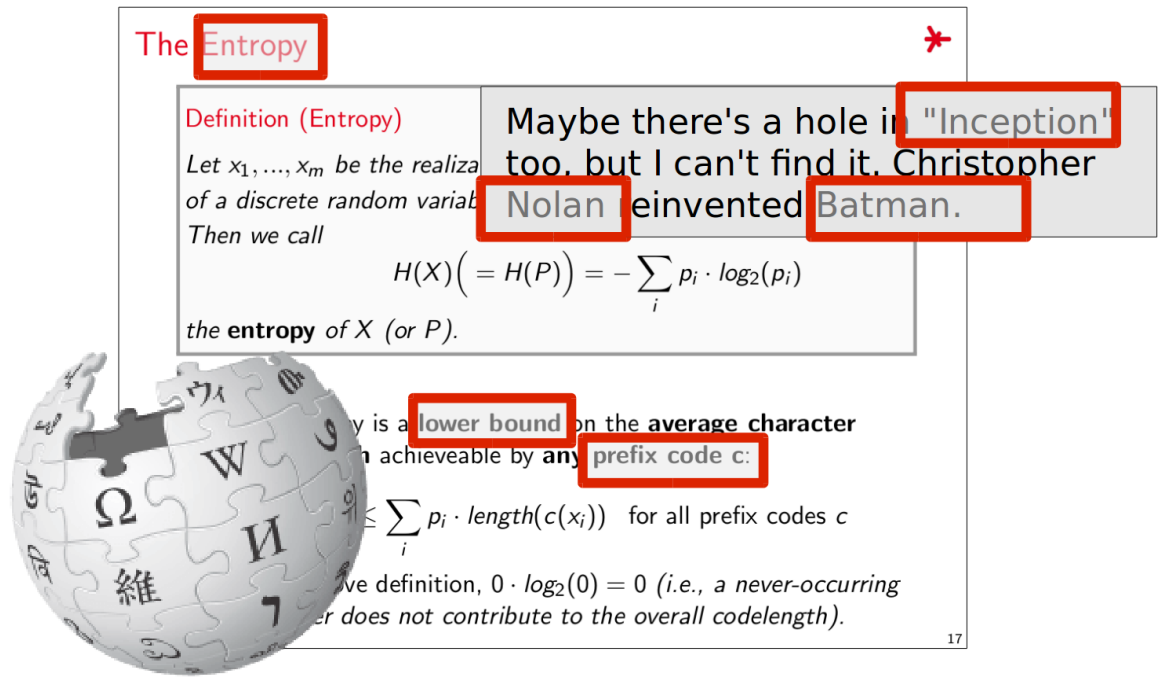

Grading

- Grading for this course will adhere to the rules of **Praktische Tätigkeit und Fachgespräch**
- Remember to sign in early on QIS!
- Attendance to "Praktikum" will be recorded *(do not miss more than 3!)*
- **Final grade**: 30% weekly coursework, 40% project, 30% Fachgespräch

1. Weekly Coursework

- Before the project phase, there will be weekly (mostly programming) exercises
- Grading based on brief "weeklies"
  - general progress
  - brief demo / results
  - code walk-through

Grading (cont'd)

2. Project

- Main deliverable: **project presentation** *(experiments+reading done, motivation for investigative path, scientific standard)*
- Do record results **while** working on the project!
- I will also check code *(should be able to install and run)*
- Please create a Gitlab repo

3. Fachgespräch

- Will use your project work as basis, but also include questions from lecture/exercises
- Prepare like for an **oral exam**!