Representation
Learning 2016

Representation Learning 2016

Course Introduction

October 17, 2016

Sebastian Stober <sstober@uni-potsdam.de>

Introductions

A Brief Intro to

Representation Learning

My 1st ML Project

- A coffee mug though the “eyes” of a road-sign detector in 2003:

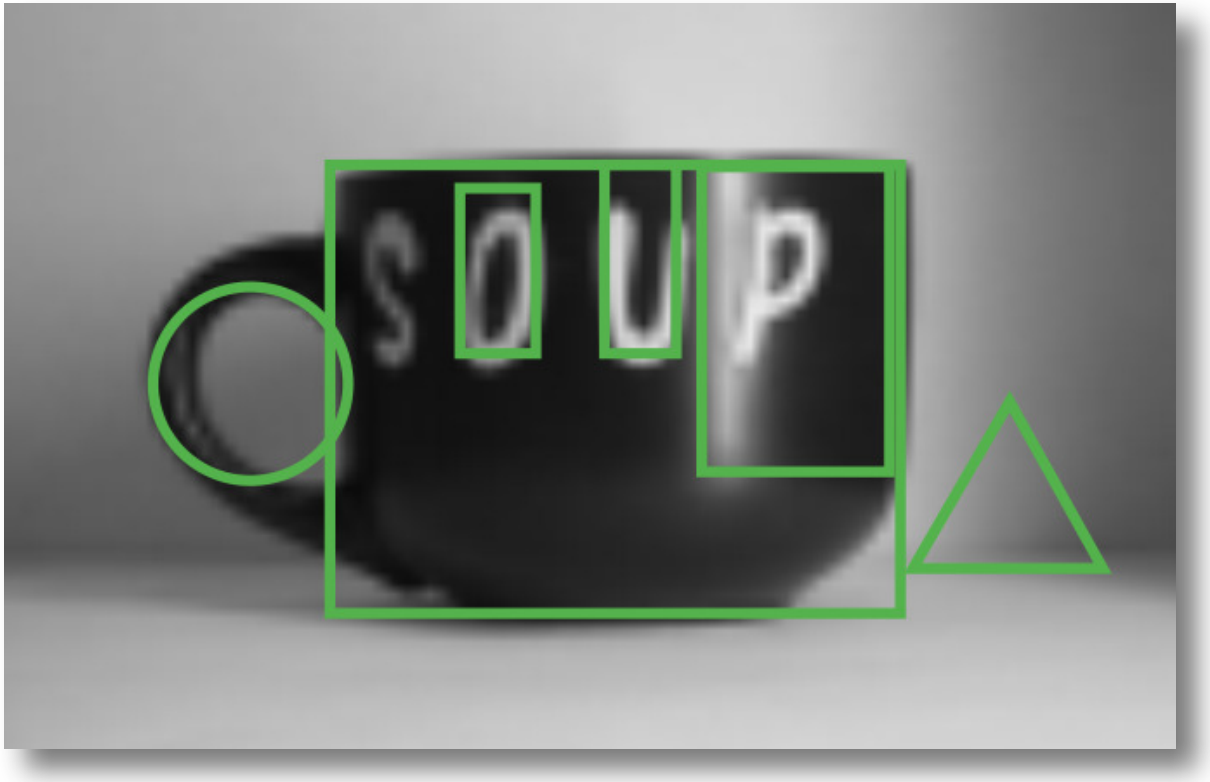

My 1st ML Project

- A coffee mug though the “eyes” of a road-sign detector in 2003:

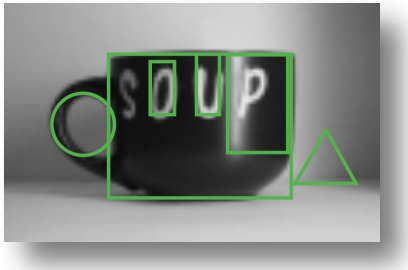

representation
to be used by classifier

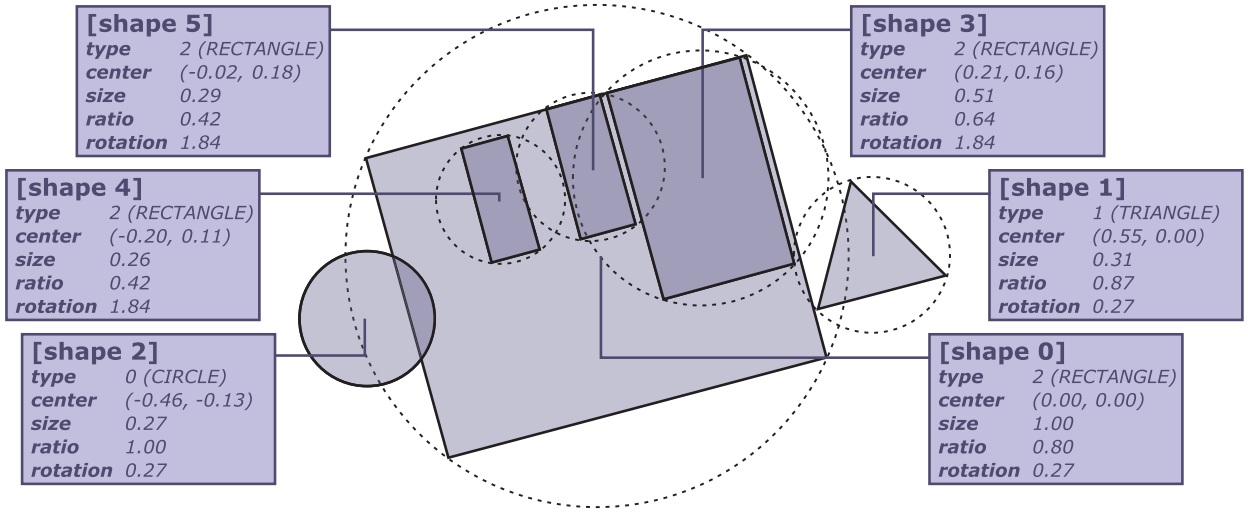

Typical Machine Learning Workflow (for Classification)

make use of domain knowledge from experts

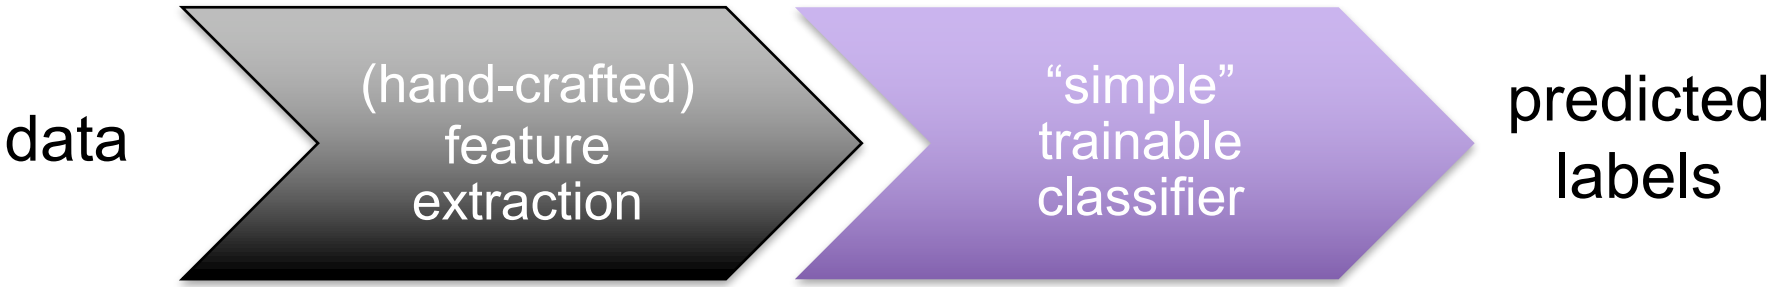

train with labeled data (supervised)

Typical Deep Learning Workflow (for Classification)

make use of abundant data and (GPU) compute power

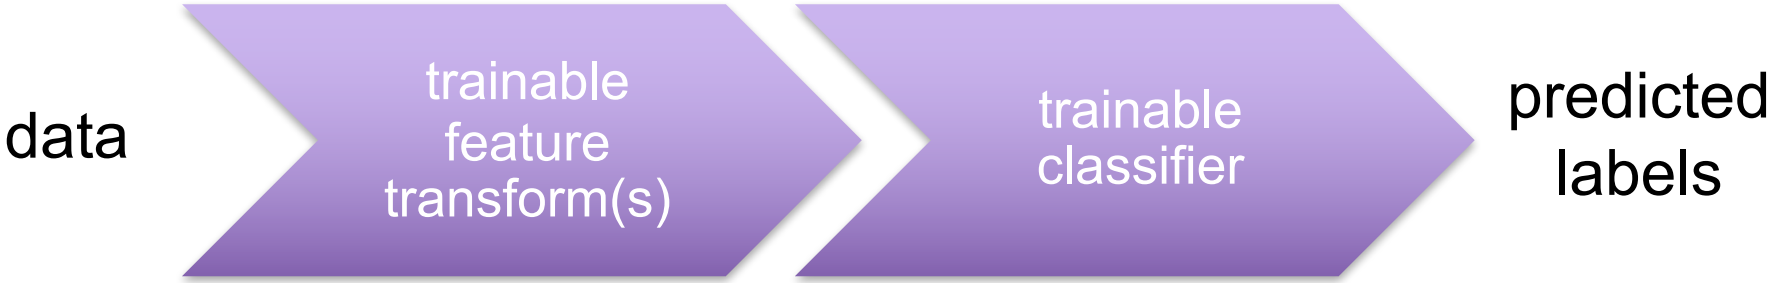

e.g., “pre-train” with unlabeled data (unsupervised)

train with labeled data (supervised)

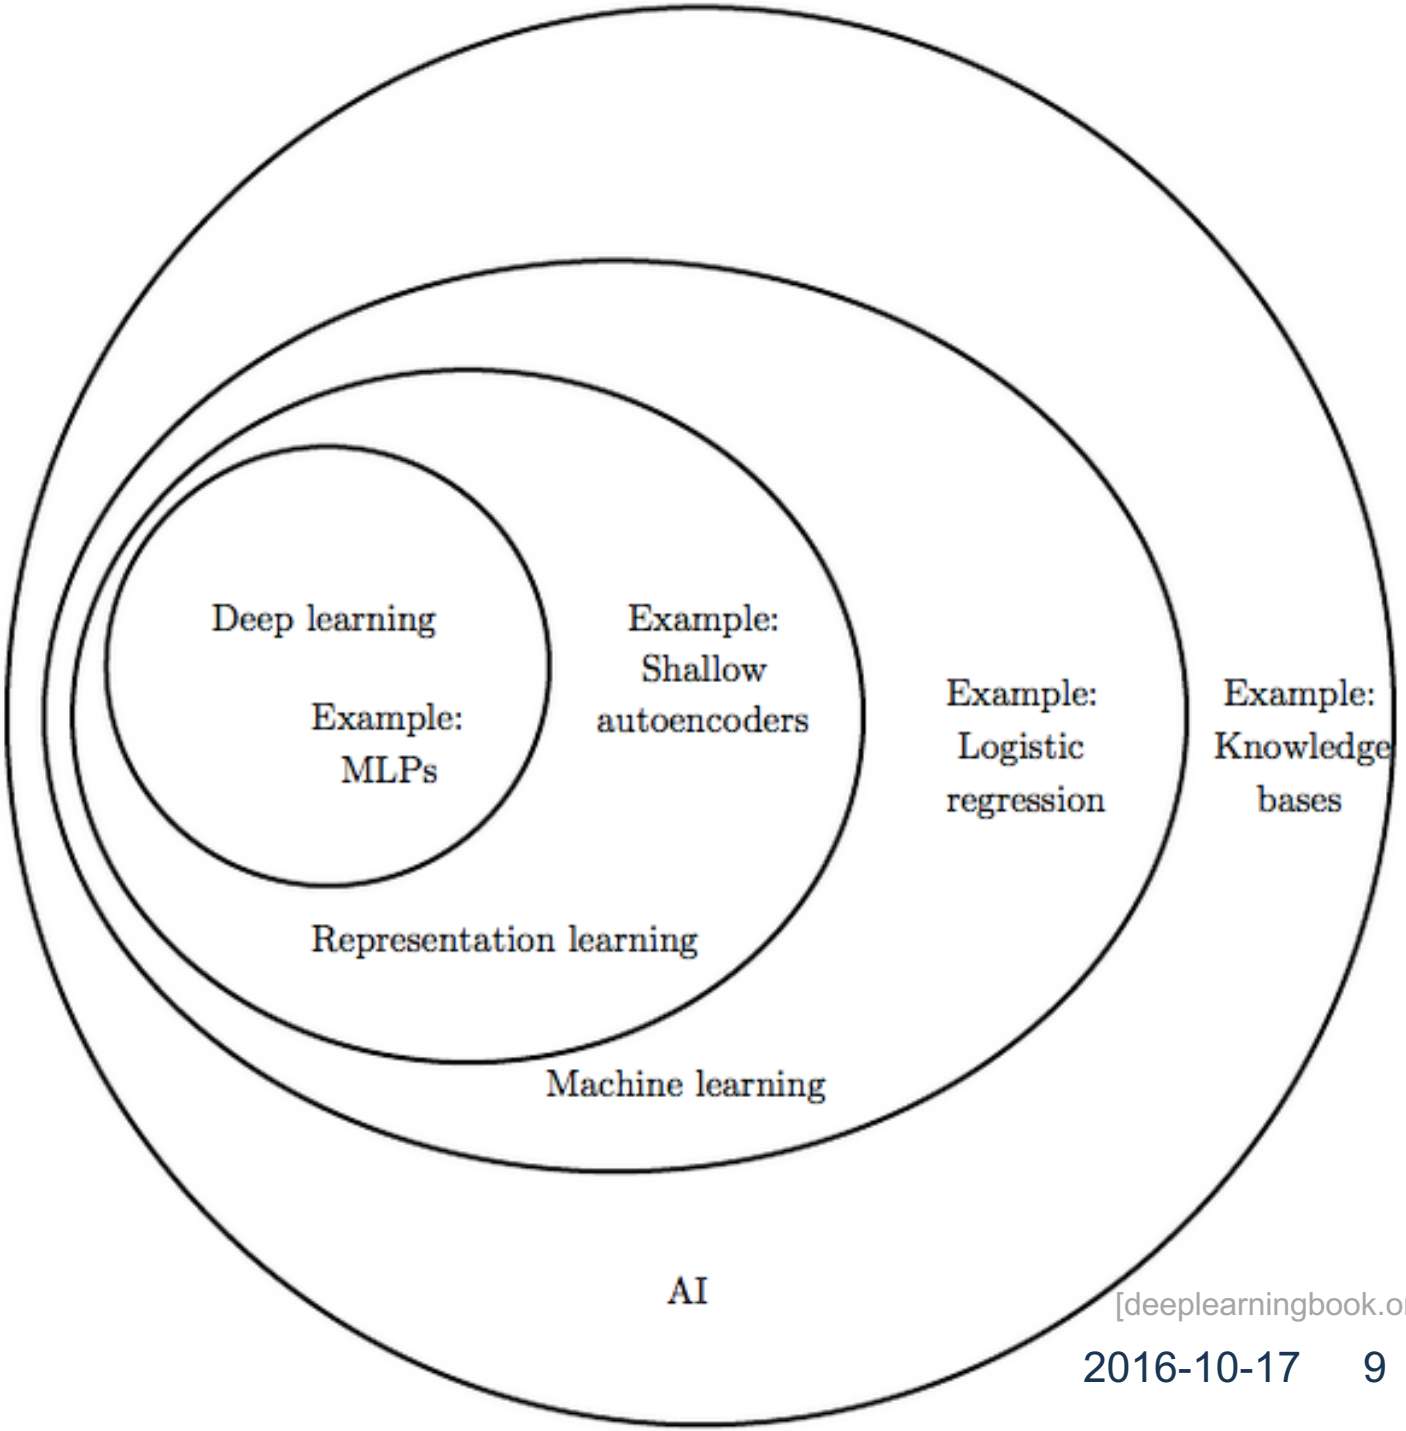

[deeplearningbook.org]

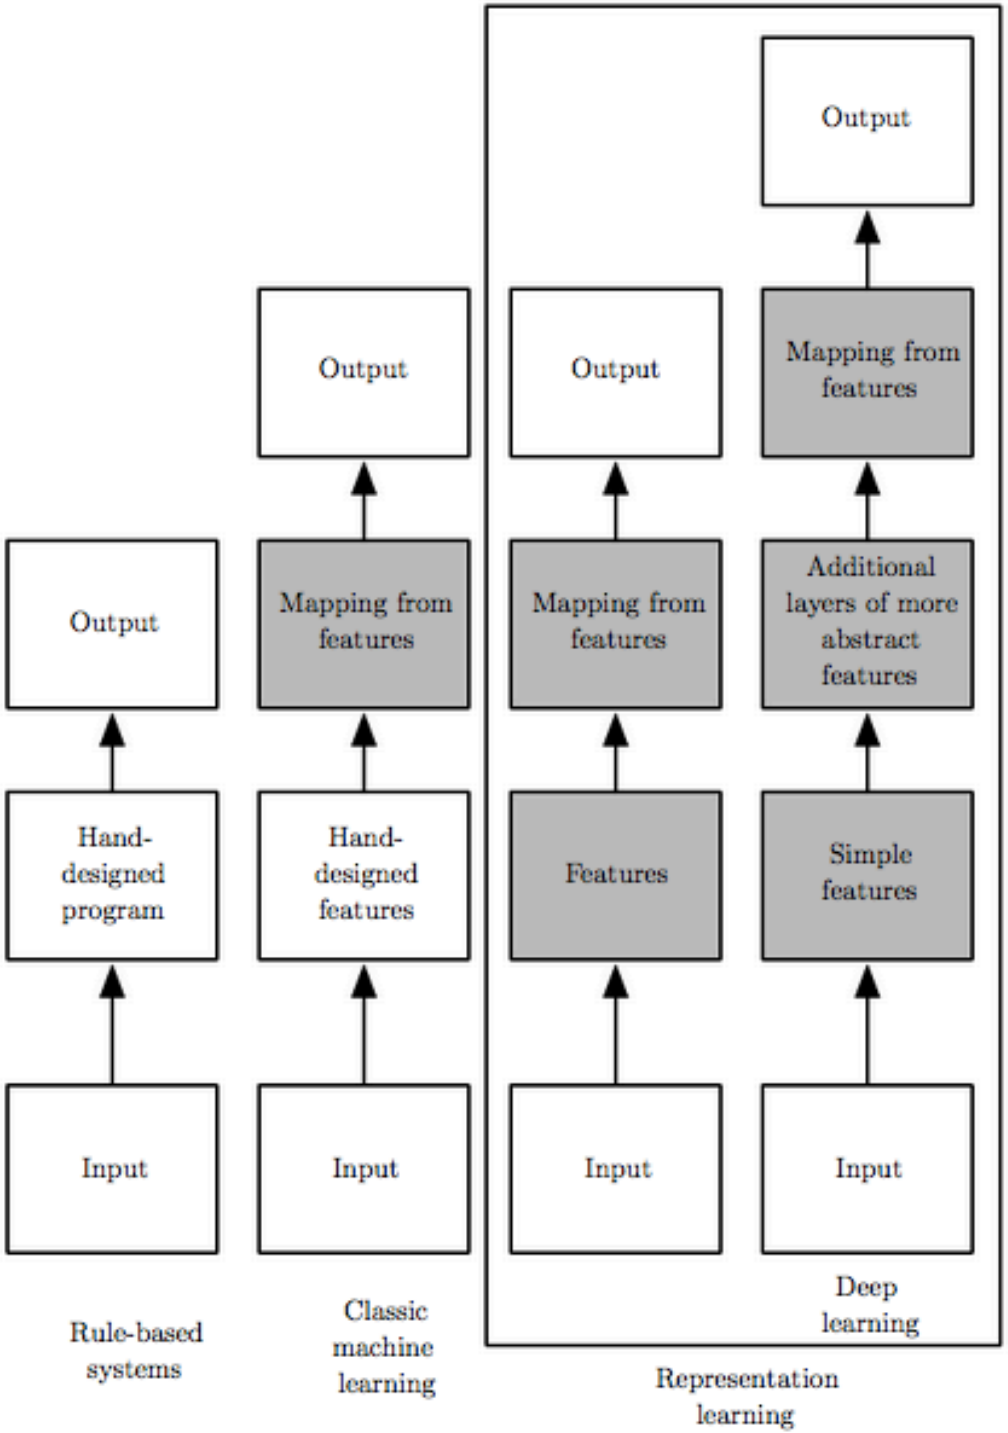

[deeplearningbook.org]

The Promise of Deep Learning

- learn suitable feature representations along with the actual learning task
- using a general-purpose learning procedure

An Example Deep Net

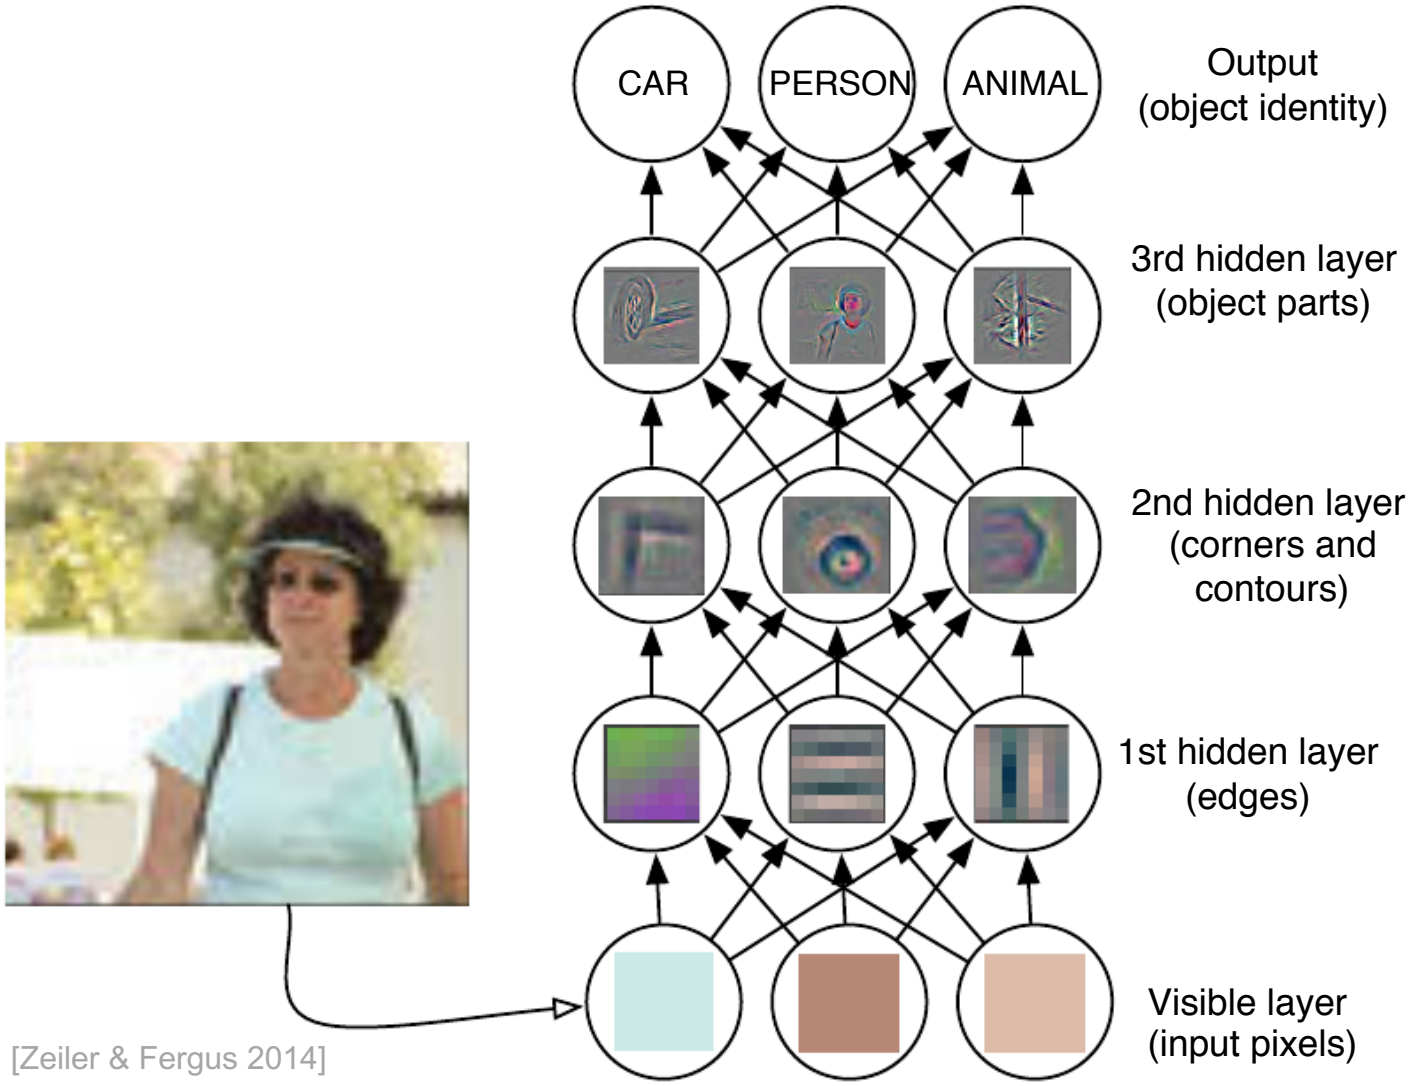

[Zeiler & Fergus 2014]

Course Rationale & Design

Learning Goals

- Think – Pair – Share
  1. Think about your personal learning goals for this course!
  2. Discuss with your neighbor & create a ranking!
  3. Name your most important one!

Learning Goals

- in-depth theory (math) and background (why and how it works)
- hands-on experience with frameworks and design decisions, “craftsmanship”
- unsupervised RL approaches
- “big picture” – applications to various topics / fields
- understand inner workings
- hierarchical feature representations for language
- working system (portfolio)

Overall Learning Goals

- At the end of the course, you are able to …
  - **confidently** apply RL techniques to develop a solution for a given problem
  - follow recent RL publications and **critically** assess their contributions
  - formulate **hypotheses** and design & conduct RL experiments to **validate** them
  - **document** progress & design decisions for **reproducibility** and transparency

Please add your goals!

Disclaimer

this course may not be suitable for …

- mere credit collectors
- passive attendees
- remote students
- the lighthearted ;-)

end of withdrawal period:
November 20, 2016

Course Design

Preparation	In Class	Course Project
(session summary)	"last episode on RL" (3-min summary)	weekly warm-up assignments (until Nov. 20)
reading • book chapters • papers • blogs	literature discussion / Q&A small-group activity (25-60 min)	working in teams • up to 4 teams • scrum-style weekly sprints
forum discussions		forum discussions
weekly blog posts • new insights • hints, tricks & hacks • open questions	project / assignments discussion / Q&A (25-60 min)	team progress blog
3-6h per week →	**150h total**	← 3-6h per week

grading: *oral exam (20min)

Topics (Tentative)

- MLPs, Gradient Descent & Backpropagation
- Autoencoders
- Convolutional Neural Networks
- Recurrent/Recursive Neural Networks
- Visualization & Sonification
- Regularization Techniques
- Advanced Regularization Techniques
- Introspection & Inception
- Optimization Techniques
- Advanced Training Strategies
- Reinforcement Learning

Online Tools

- Campus.UP (workspace “RL2016”)
  - blogs (course / personal / team)
  - forum
  - wiki
  - microblog / messaging (?)
- Slack channel (?)
- GPU compute environment (medusa)
  - shell access & jupyterhub for notebooks

GPU Compute Server

- 4 Maxwell Geforce Titan X GPUs
- 128 GB RAM
- 12 CPU cores
- jupyterhub server for notebooks

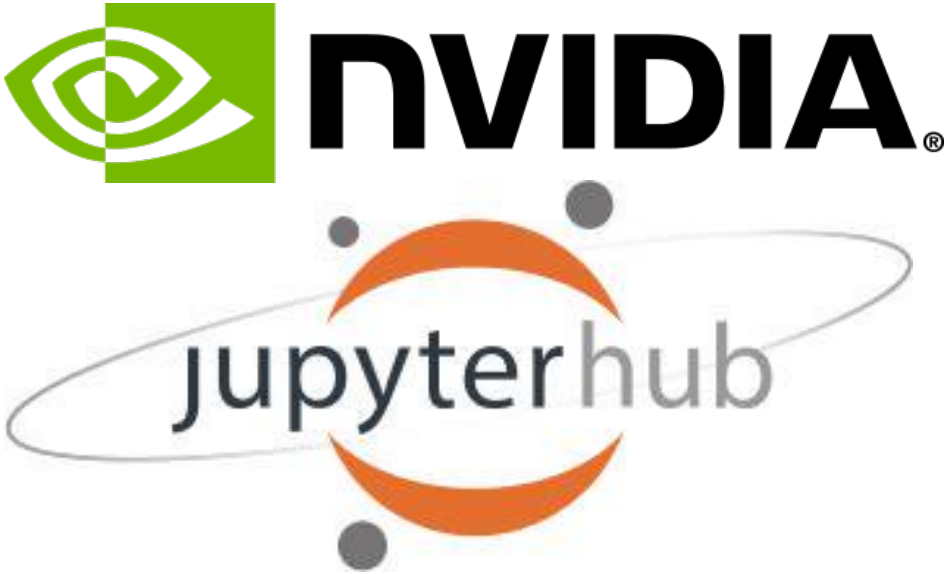

https://jupyter.org/

… more info in assignment #1

Jupyterhub

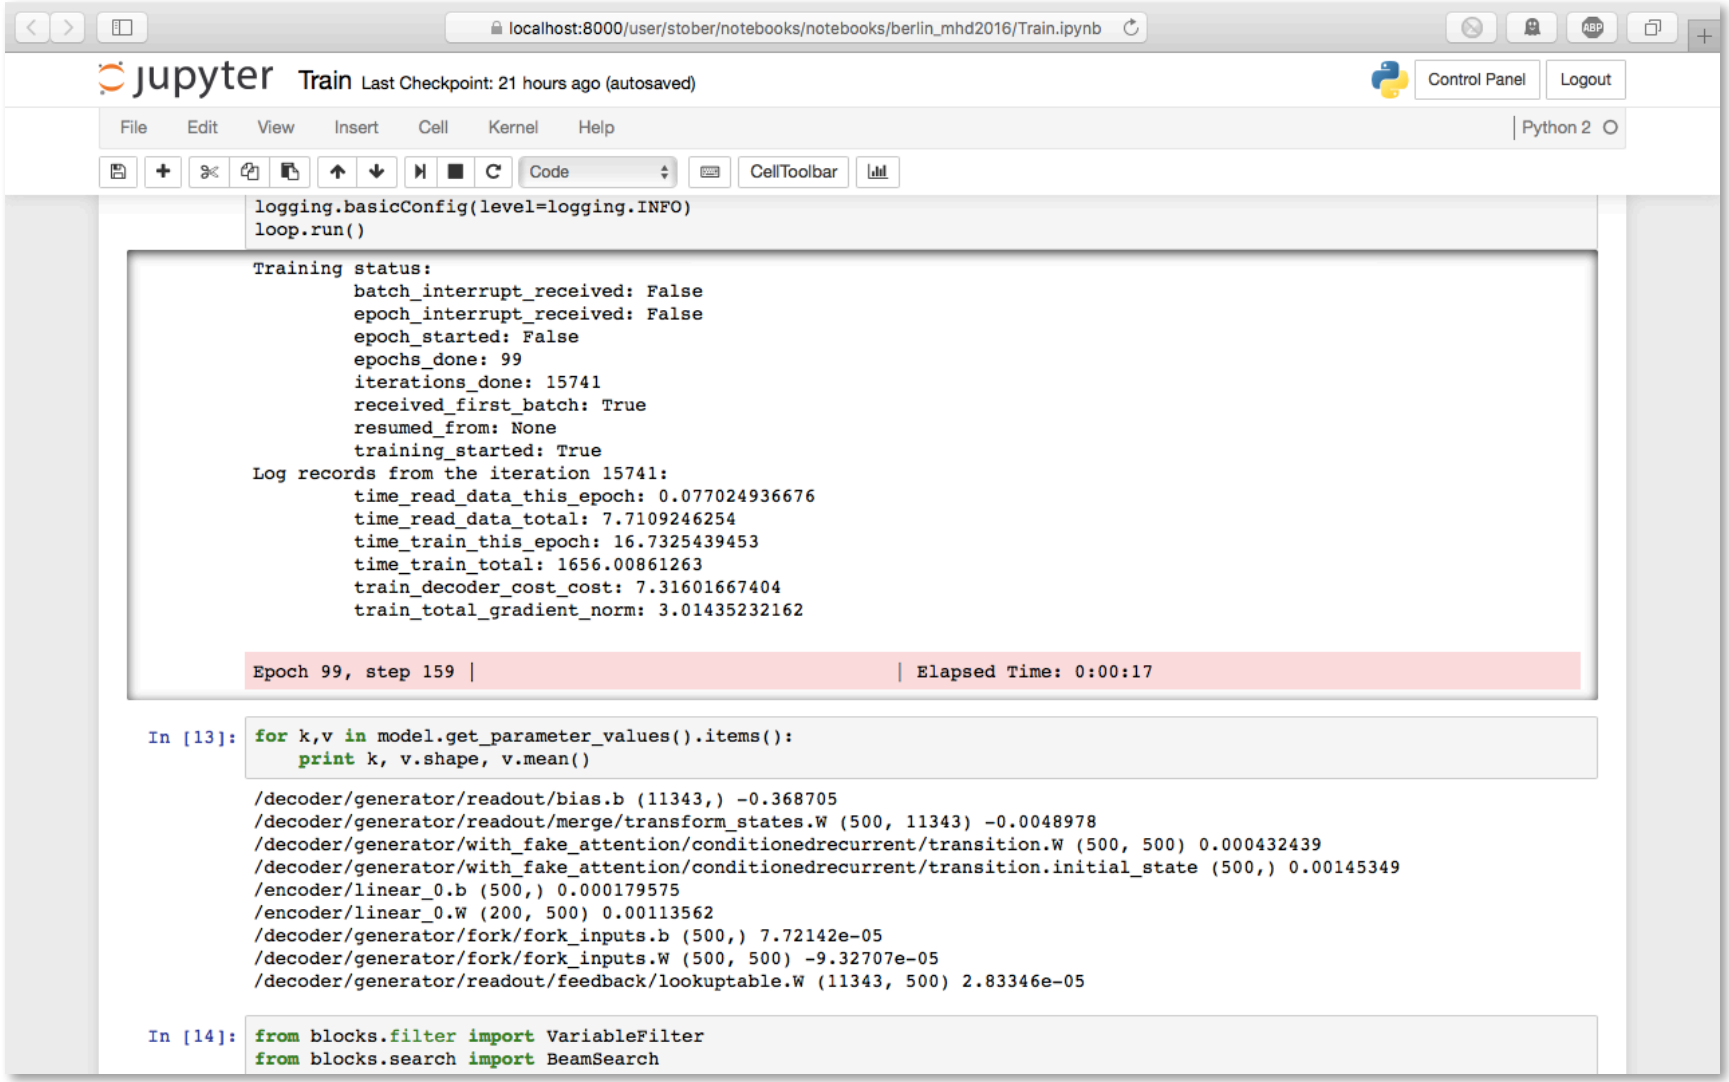
localhost:8000/user/stober/notebooks/notebooks/berlin_mhd2016/Train.ipynb
jupyter Train Last Checkpoint: 21 hours ago (autosaved)
Control Panel
Logout
File Edit View Insert Cell Kernel Help
Python 2
Code
CellToolbar
logging.basicConfig(level=logging.INFO)
loop.run()
Training status:
batch_interrupt_received: False
epoch_interrupt_received: False
epoch_started: False
epochs_done: 99
iterations_done: 15741
received_first_batch: True
resumed_from: None
training_started: True
Log records from the iteration 15741:
time_read_data_this_epoch: 0.077024936676
time_read_data_total: 7.7109246254
time_train_this_epoch: 16.7325439453
time_train_total: 1656.00861263
train_decoder_cost_cost: 7.31601667404
train_total_gradient_norm: 3.01435232162
Epoch 99, step 159 | | Elapsed Time: 0:00:17
In [13]:
for k,v in model.get_parameter_values().items():
print k, v.shape, v.mean()
/decoder/generator/readout/bias.b (11343,) -0.368705
/decoder/generator/readout/merge/transform_states.W (500, 11343) -0.0048978
/decoder/generator/with_fake_attention/conditionedrecurrent/transition.W (500, 500) 0.000432439
/decoder/generator/with_fake_attention/conditionedrecurrent/transition.initial_state (500,) 0.00145349
/encoder/linear_0.b (500,) 0.000179575
/encoder/linear_0.W (200, 500) 0.00113562
/decoder/generator/fork/fork_inputs.b (500,) 7.72142e-05
/decoder/generator/fork/fork_inputs.W (500, 500) -9.32707e-05
/decoder/generator/readout/feedback/lookuptable.W (11343, 500) 2.83346e-05
In [14]:
from blocks.filter import VariableFilter
from blocks.search import BeamSearch

Campus.UP Workspace

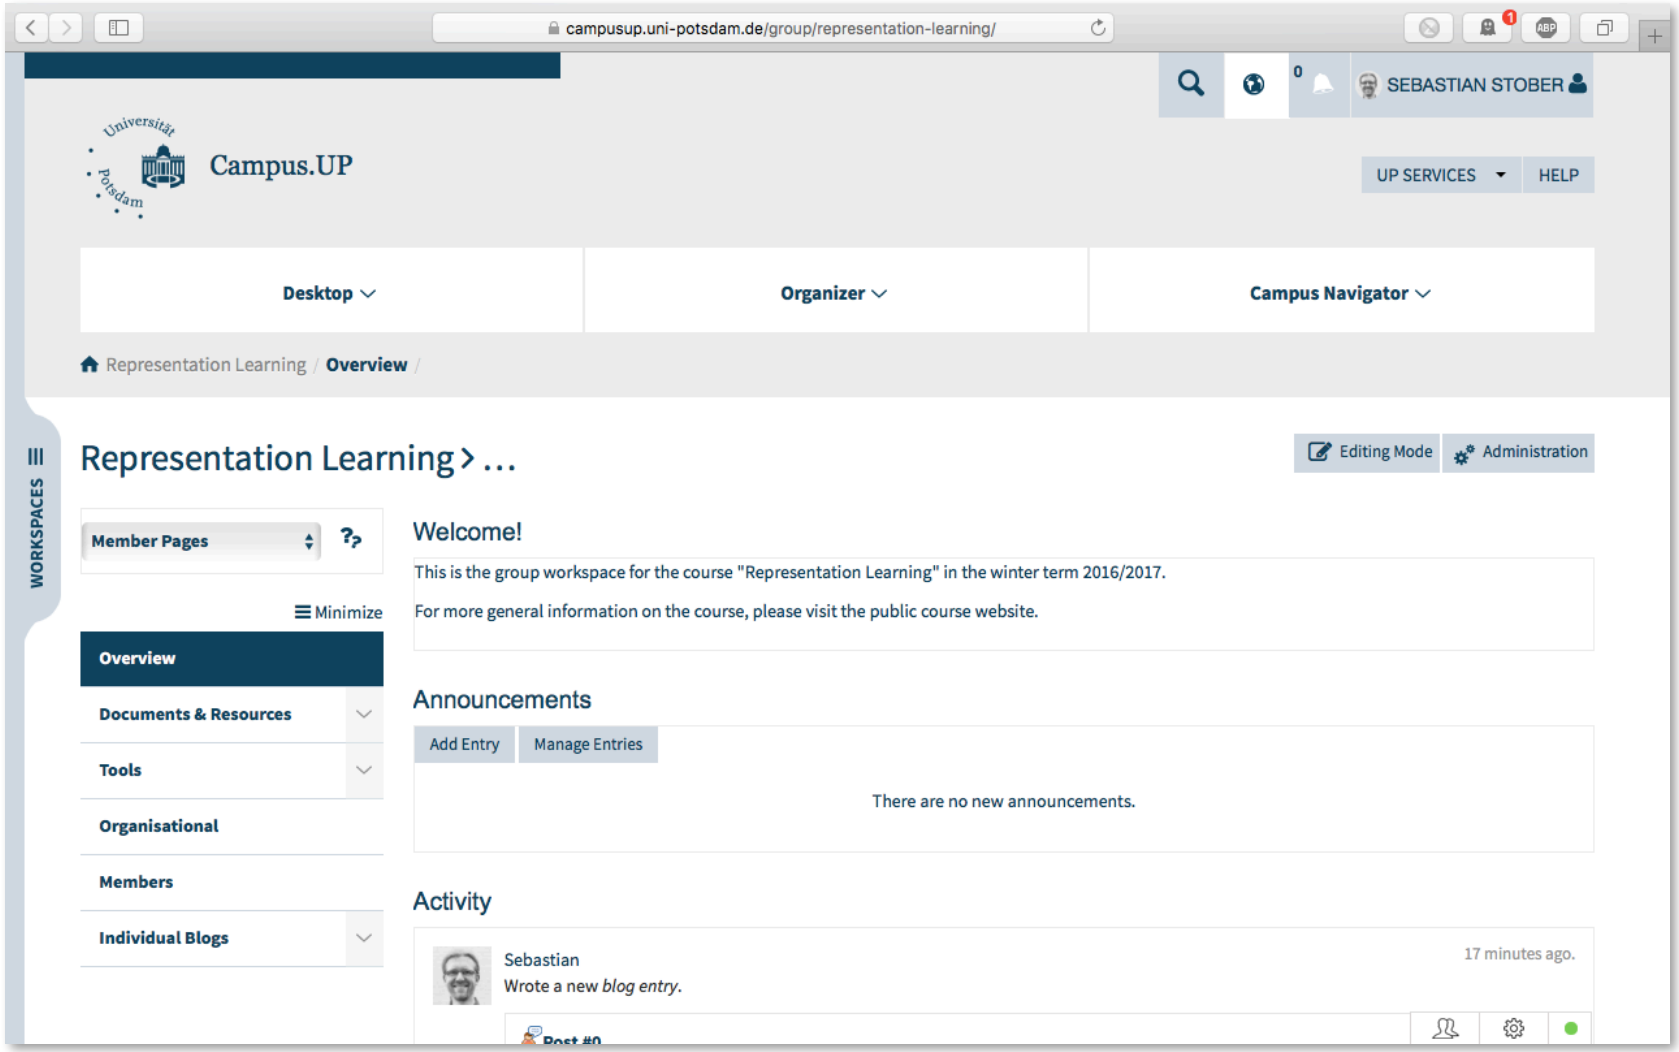

Session Summaries

- rotating job!
  (one session per person, assignment by poll)

- short summary blog post (in course blog)
  + 3-min intro recap at next session
  - key topics
  - results of the discussion
  - optional photos

Your Personal Blog

- document your learning / project progress
  - one post per week
  - share your experiences!
  - visible only to course participants

- examples:
  - https://deeprandommumbling.wordpress.com/
  - http://bartvanmerrienboer.nl/#blog
- guidelines:
  - https://www2.uwstout.edu/content/profdev/rubrics/blogrubric.html

Team Blogs

- weekly progress reports for course project
  - similar to scrum
  - compare original goals with outcomes
    - What has worked well?
    - What did not work / had to be changed?
  - outline plan for next week
    - What would you like to try / investigate next?
- can be written up by one designated team member or in turns

Open Questions (Blog)

- guide for what is covered in class
  deadline: Monday morning 7am

- do not hesitate to post questions!
  (If you got one, you are probably not the only one!)

- post a comment if you know the answer

Contribute!

- ask – in your blog and the forum
- comment / like / rate
- answer
- document – in your blog and the wiki
  - hints, tricks & hacks
- recommend
  - additional readings (papers, blogs, etc.)
- give (constructive) feedback

Course Project

Course Project

- vision: speech-base interaction

real systems:

- Siri (Apple)
- Alexa (Amazon)
- Cortana (Microsoft)
- Google Speech
- Skype Translate
- …

fictional characters:

- J.A.R.V.I.S. (Iron Man)
- Samantha (Her)
- Jane (Ender's Game)
- ...

Course Project

- Automatic Speech Recognition (ASR)
  - state-of-the-art systems use deep nets

- large-scale dataset (cc-by 4.0):
  1000h corpus of read English speech:
  "LibriSpeech: an ASR corpus based on public domain audio books", V. Panayotov, G. Chen, D. Povey and S. Khudanpur, ICASSP 2015.
  http://www.openslr.org/12/

Course Project

- collaborative effort
- “coopetition” (cooperative + competition)
- up to 4 teams:
  - form after course withdrawal deadline (Nov. 20)
  - self-organized (heterogeneous if possible)
  - scrum-like approach
  - focus on different aspects / strategies / tools

Course Project

- available deep learning frameworks:
  - Theano (+ Blocks&Fuel or Keras)
  - Tensorflow (+ Keras)

Course Project

- optional: collaborative overview-paper

http://www.aes.org/conferences/2017/semantic/

- manuscript deadline: January 22, 2017