



Introduction to
Deep Learning

Course Introduction

April 9, 2018

Sebastian Stober <sstober@uni-potsdam.de>
Jens Johannsmeier <johannsmeier@uni-potsdam.de>

Introductions

“Living Stats”

- How far away from home are you?
- How long have you studied already?
- What did you study?
- How experienced are you in ML?
- How experienced are you in DL?

My Experience

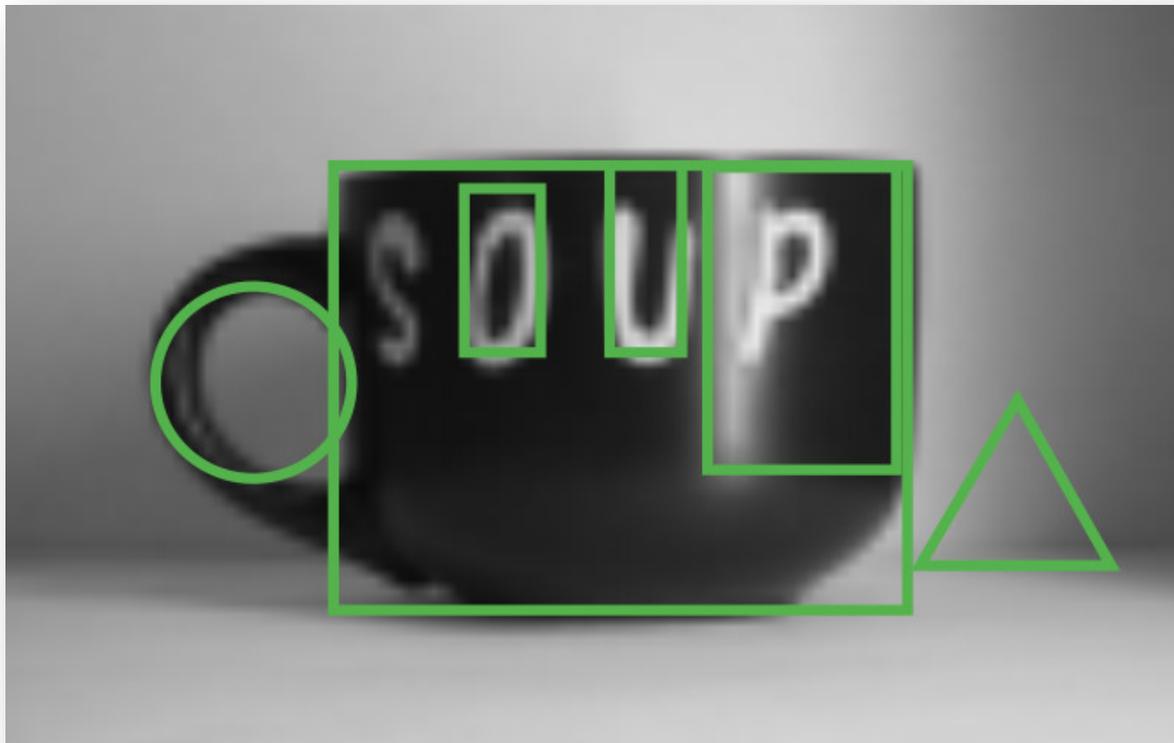
1. ML / artificial neural nets > 16 years
2. deep nets > 4 years
3. generative nets > 1 year

A Brief Intro to

Deep Learning

My 1st ML Project

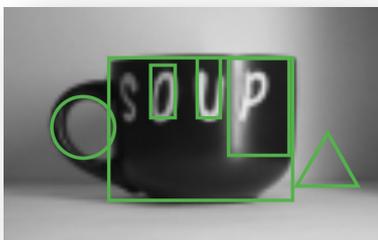
- A coffee mug through the “eyes” of a road-sign detector in 2003:



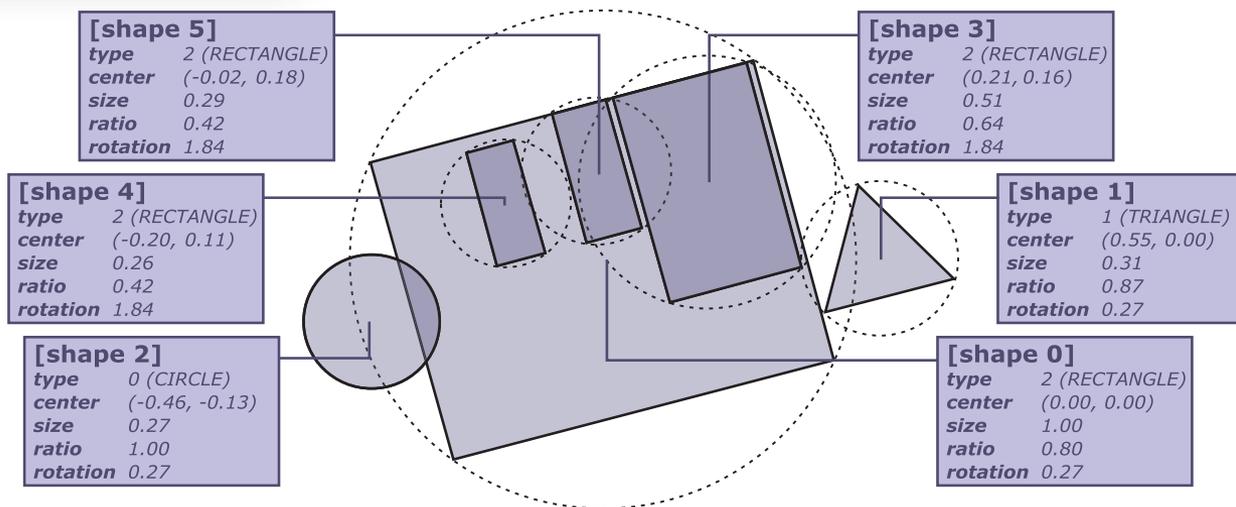


My 1st ML Project

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

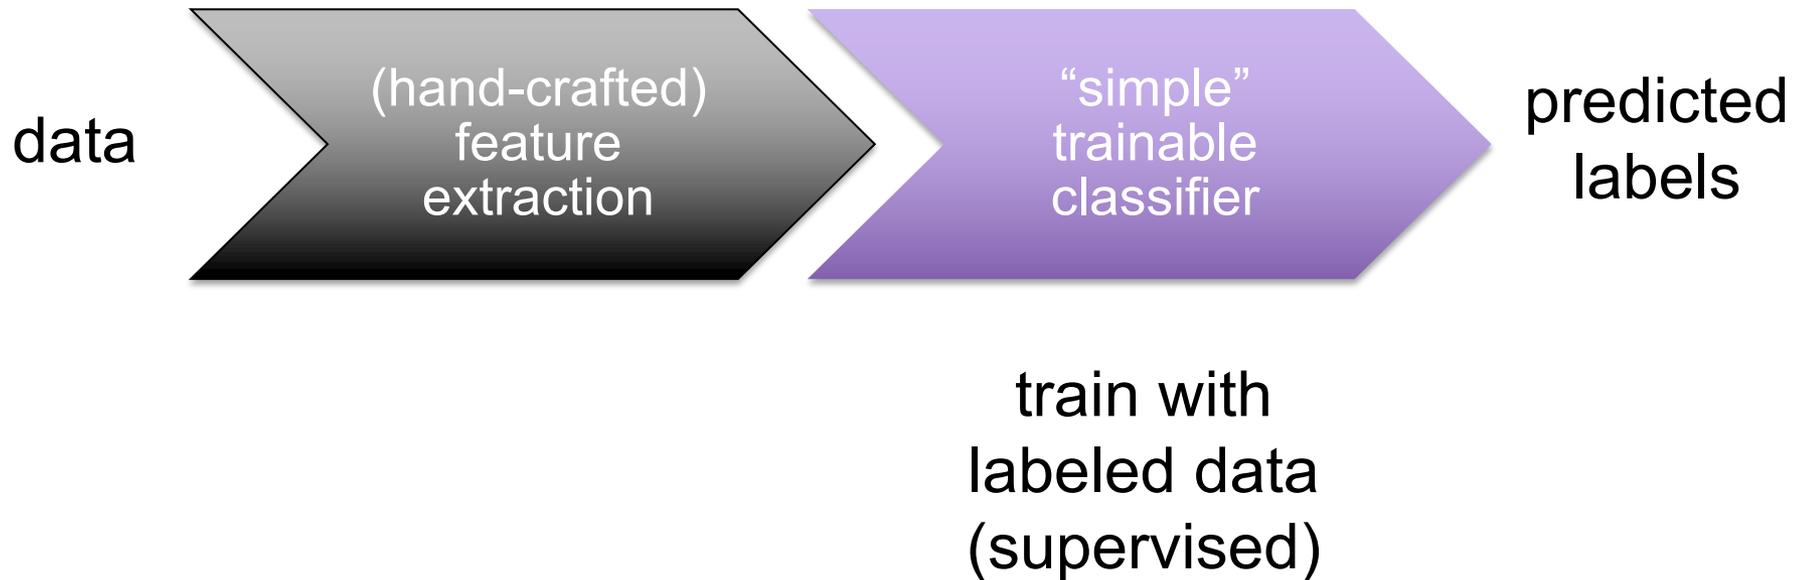


representation
to be used by classifier



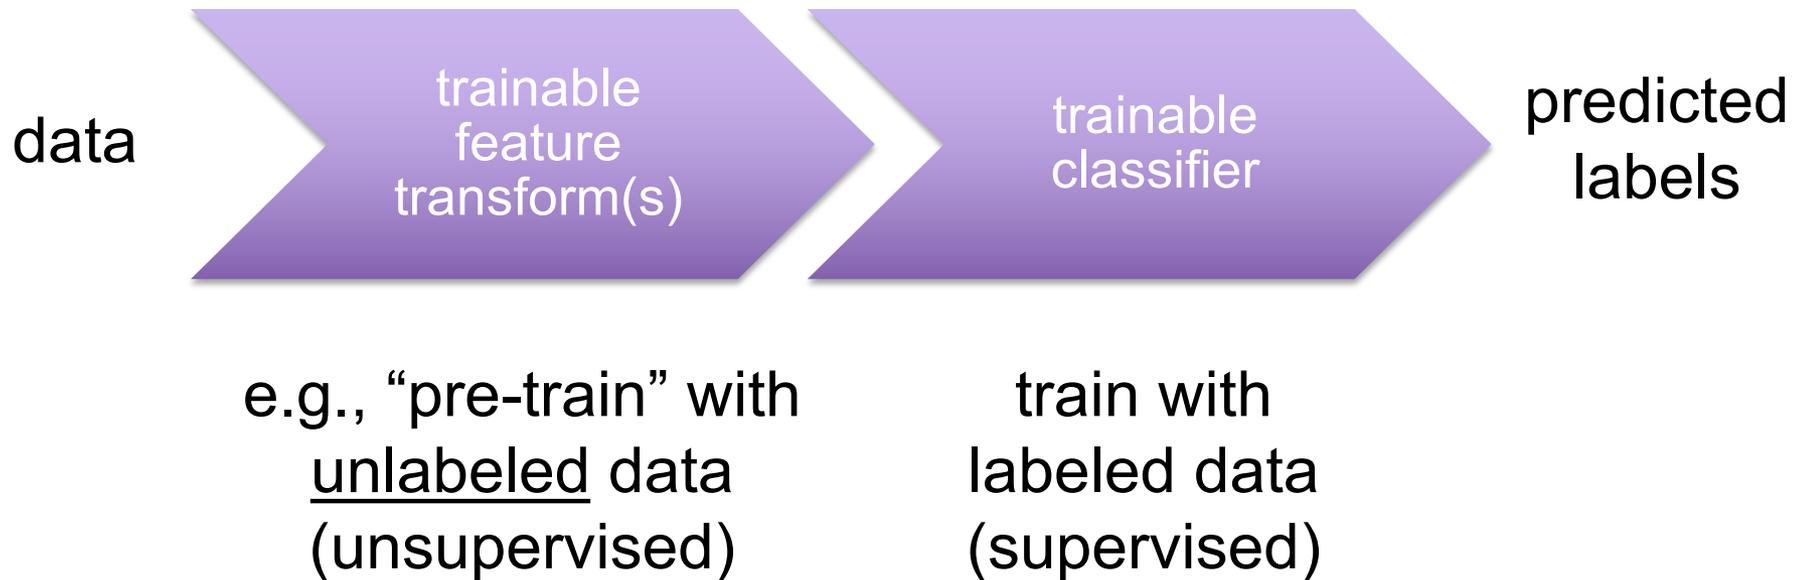
Typical Machine Learning Workflow (for Classification)

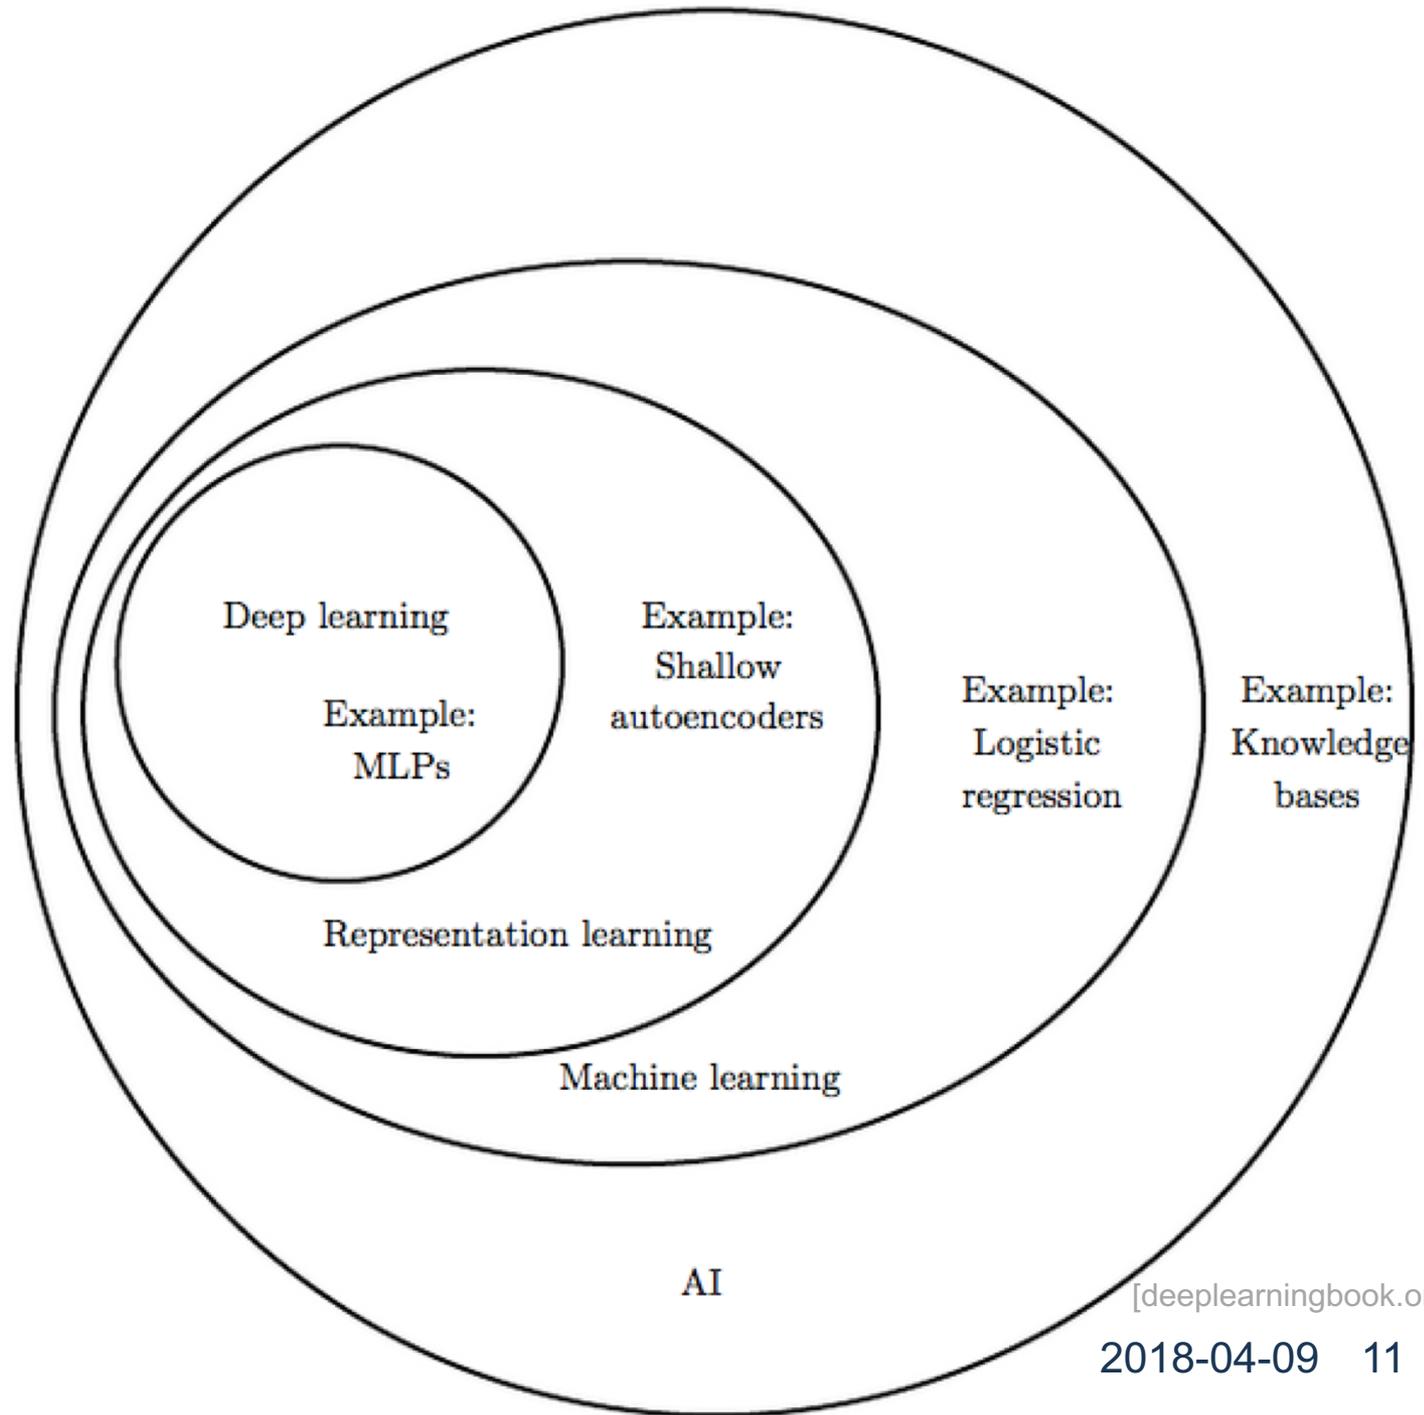
make use of domain
knowledge from experts

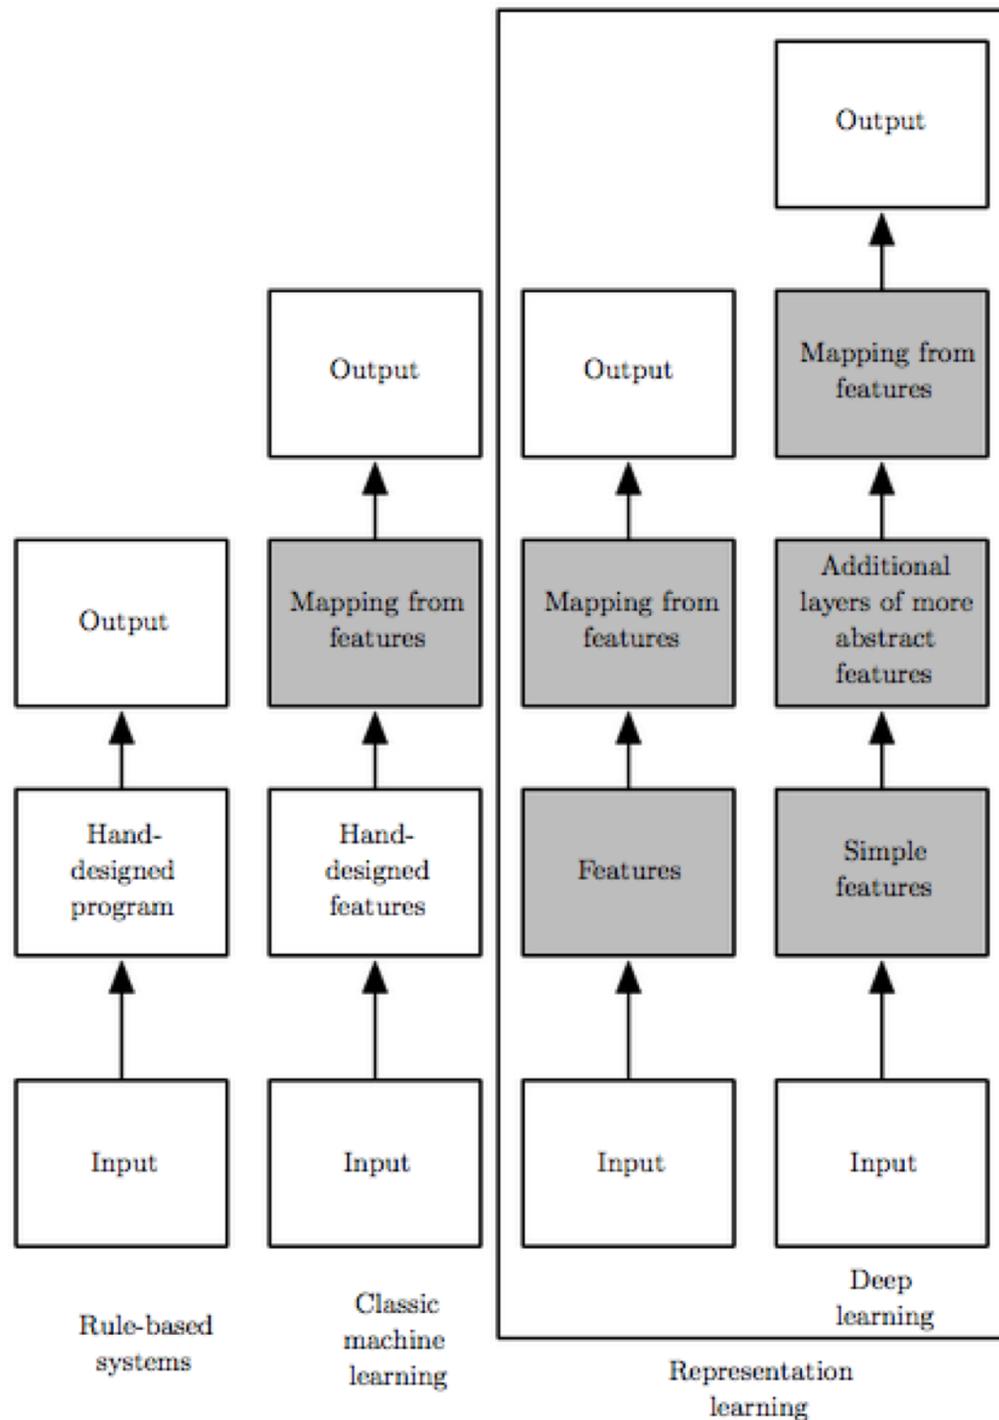


Typical Deep Learning Workflow (for Classification)

make use of abundant data
and (GPU) compute power





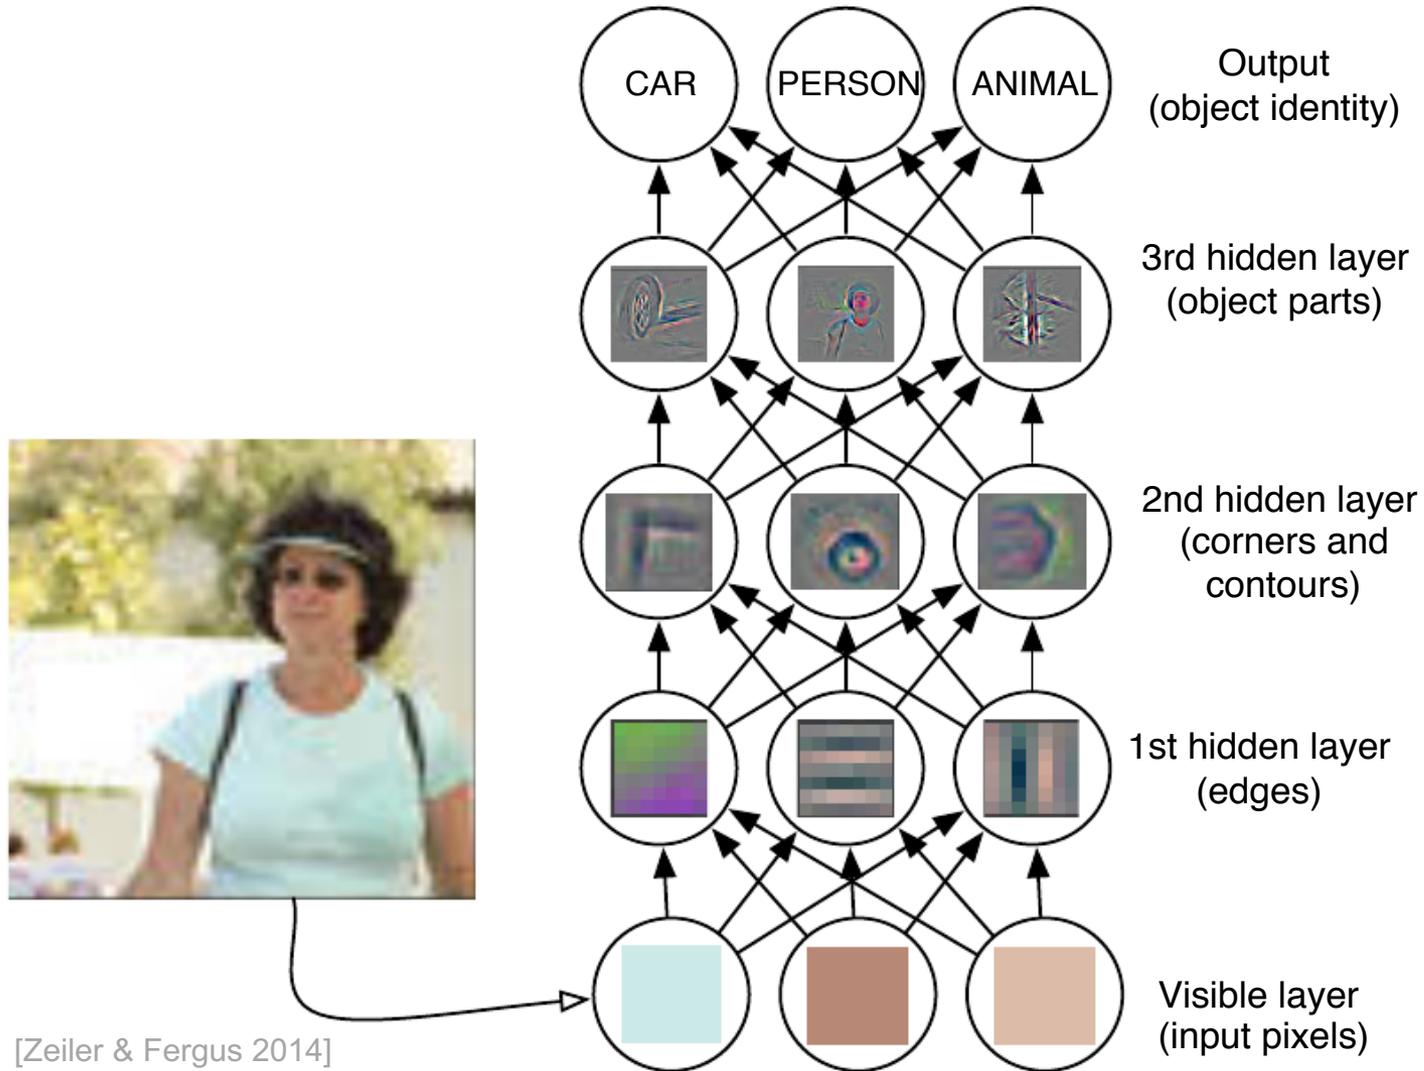


[deeplearningbook.org]

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



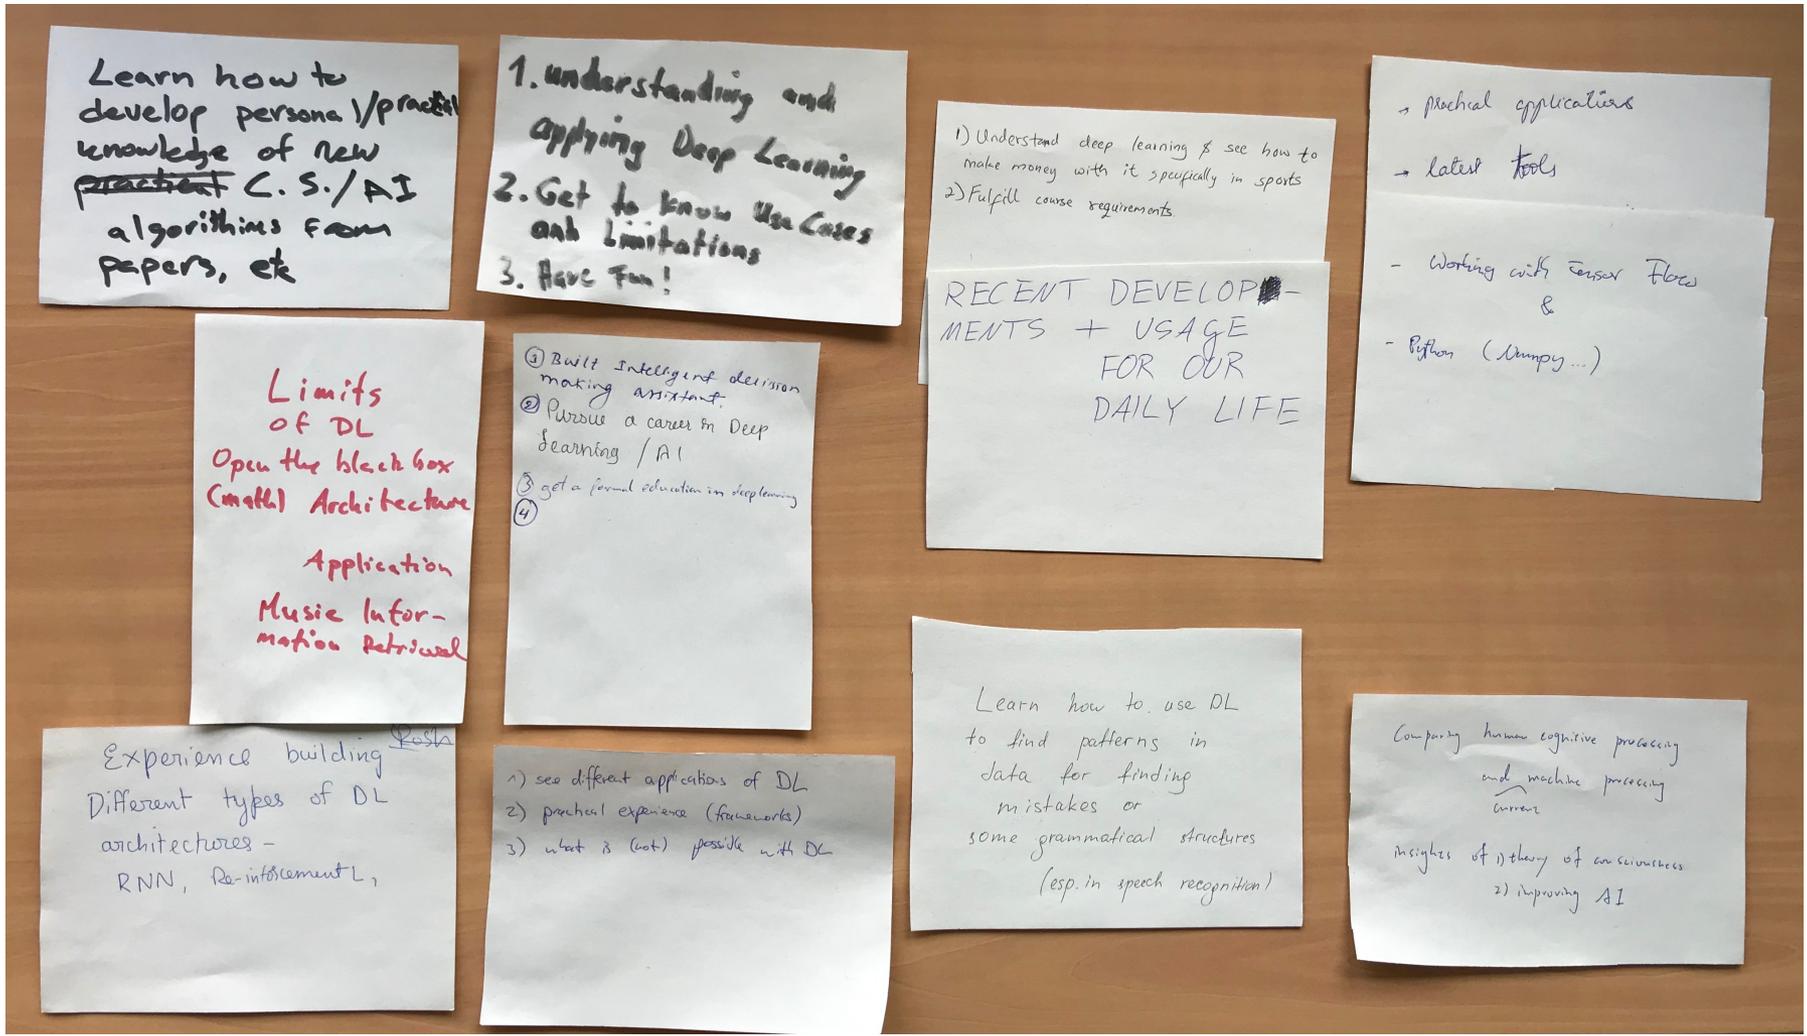
Course Rationale & Design

Learning Goals

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



Learning Goals



Overall Learning Goals

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

Please add your goals!

Disclaimer



HERE BE DRAGONS!

WELL, NOT REALLY. WE WERE JUST
TOO LAZY TO LIST THE HAZARDS, AND
THIS IS MORE ACCURATE THAN NOTHING AT ALL

this course may not be
suitable for ...

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

*end of withdrawal period:
May 20, 2018*

No Free Lunch!

This is **not** how
you will learn...

you will need to
participate



The Nuremberg Funnel (1647)

Course Design

Preparation

(session summary)

reading

- book chapters
- papers
- blogs

channel discussions

weekly blog posts

- new insights
- hints, tricks & hacks
- open questions

3-6h per week

In Class

“last episode on IDL”
(3-min summary)

literature
discussion / Q&A
small-group activity
(25-60 min)

project / assignments
discussion / Q&A
(25-60 min)

150h total

Practical Work

weekly exercise
assignments
(until week 7 or 8)

channel discussions

course project
working in teams

- 2-4 team members
- scrum-style
weekly sprints
- team progress blog

3-6h per week

grading: *oral exam (20min)

Topics (Tentative)

- MLPs, Gradient Descent & Backpropagation
- Convolutional Neural Networks
- Recurrent/Recursive Neural Networks
- Auto-Encoders
- Regularization Techniques
- Advanced Regularization Techniques
- Introspection & Inception
- Optimization Techniques
- Advanced Training Strategies

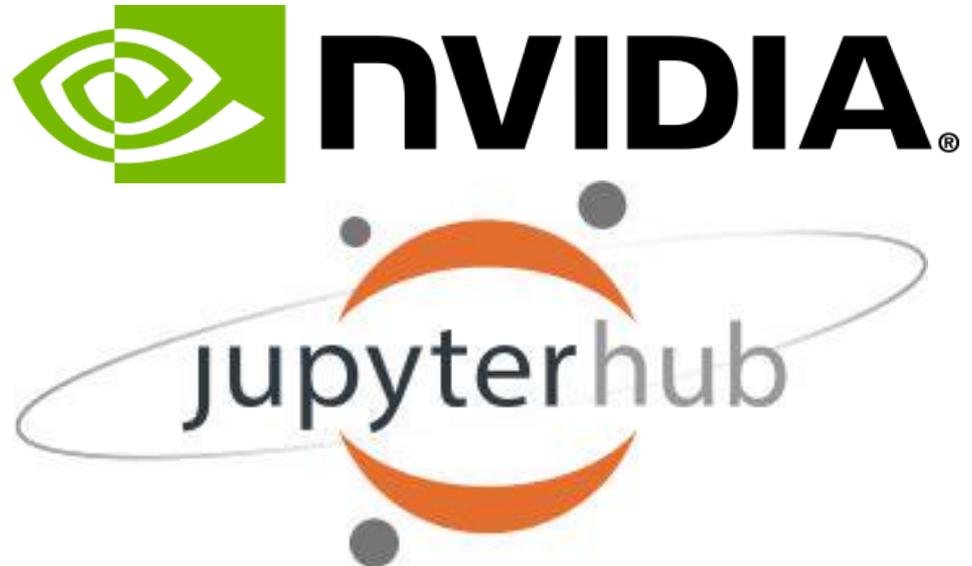
Online Tools

- Mattermost (~Slack) channel
+ Campus.UP workspace
 - channels / blogs (course / personal / team)
 - forum
 - wiki
 - messaging
- GPU compute environment
 - shell access & jupyterhub for notebooks

GPU Compute Servers

10 GPU Compute Servers

- 8 Pascal Geforce 1080 Ti GPUs
- 256 GB RAM
- 24 CPU cores
- jupyterhub server for notebooks
- fully dockerized



<https://jupyter.org/>



Jupyterhub

The screenshot displays a Jupyter Notebook running on a local host. The browser address bar shows the URL: localhost:8000/user/stober/notebooks/notebooks/berlin_mhd2016/Train.ipynb. The notebook interface includes a menu bar (File, Edit, View, Insert, Cell, Kernel, Help) and a toolbar with various icons for file operations and execution. The main content area shows the following code and output:

```
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
```



Mattermost Channel

The screenshot shows the Mattermost interface for a channel named 'Town Square'. On the left is a dark sidebar with a channel list. The main area displays a welcome message and a history of system messages indicating new members joining the channel.

Channel Sidebar:

- ID IDL 20... (selected)
- ML MLCog
- UP UPrac...
- + (add more)
- PUBLIC CHANNELS
 - Deep Learning News
 - Off-Topic
 - Q & A
 - Town Square (selected)
 - More...
- PRIVATE CHANNELS +
- DIRECT MESSAGES +
 - ankrug
 - More...
- Switch Channels - ⌘K

Channel Header:

- ☆ Town Square ▾
- Add a channel description
- 3 👤 (members)
- 🌟 (star)
- 🔍 Search
- @ (mention)
- 📄 (posts)

Channel Content:

Beginning of Town Square

Welcome to Town Square!

Post messages here that you want everyone to see. Everyone automatically becomes a permanent member of this channel when they join the team.

[+ Invite others to this team](#) [🔧 Set a Header](#)

Tue, Sep 26, 2017

System 23:17
sstober has joined the channel.

Wed, Sep 27, 2017

System 18:24
jens has joined the channel.

jens 18:25 🙌

Thu, Sep 28, 2017

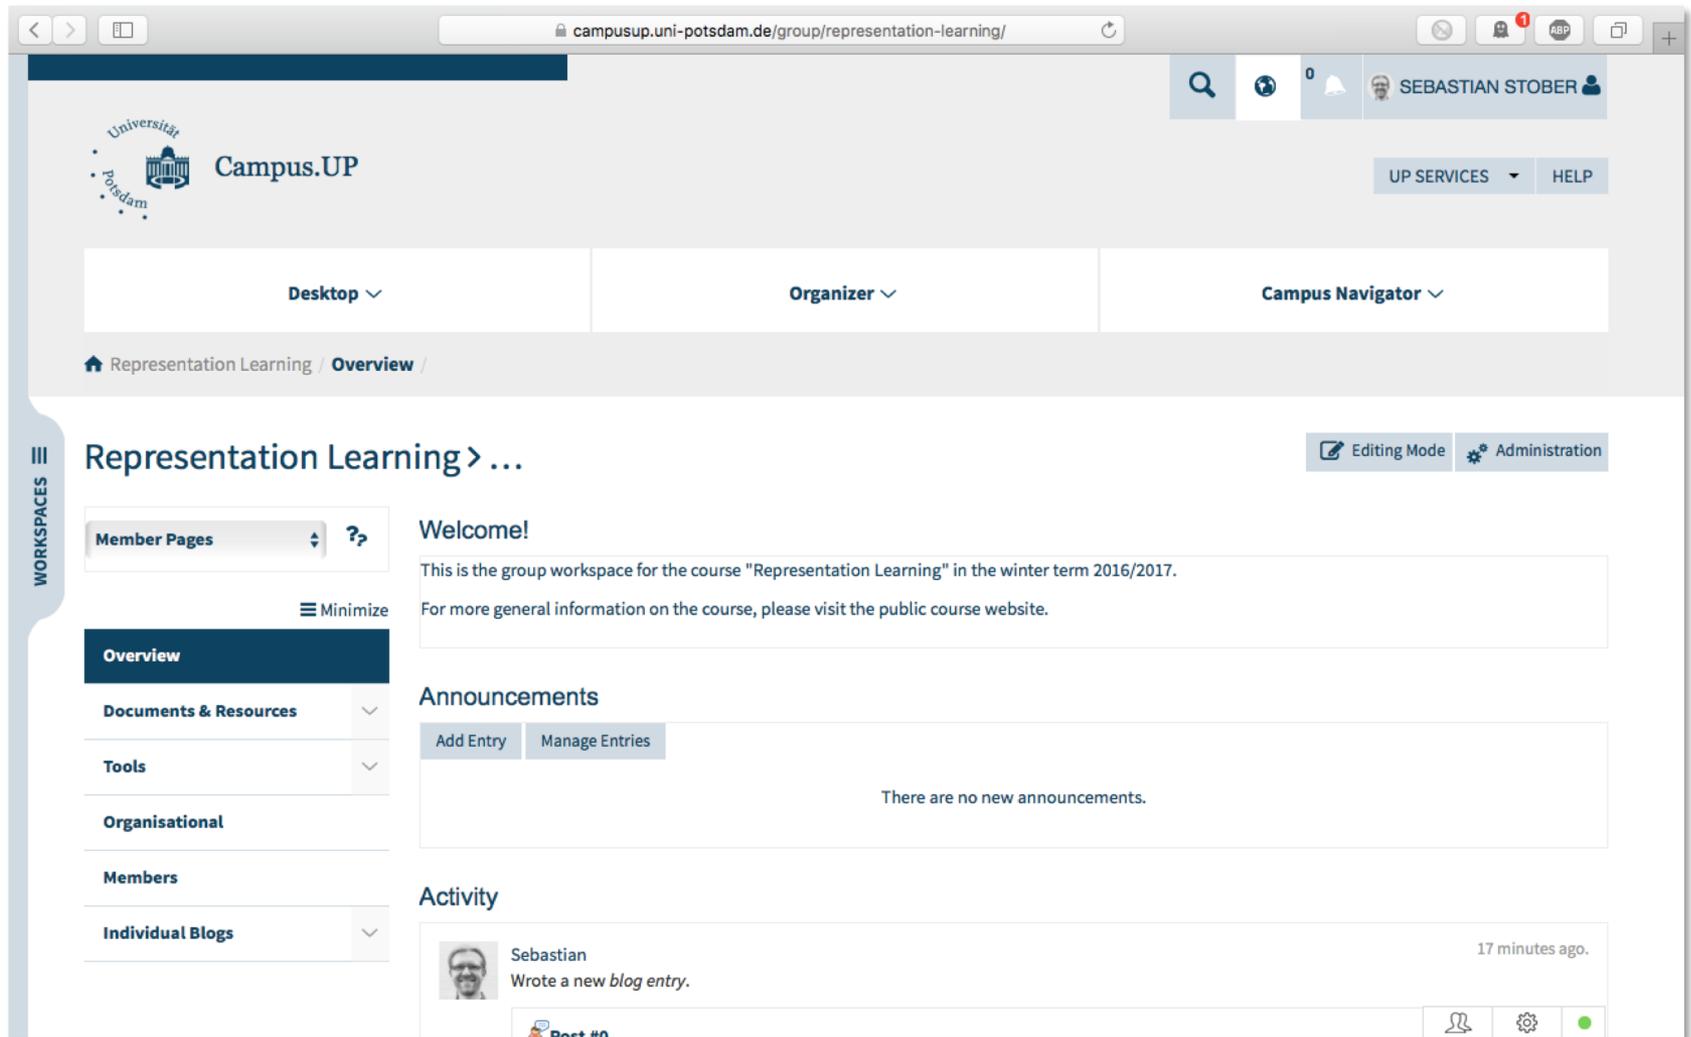
System 11:19
ilia has joined the channel.

ilia 11:19 🙌

Write a message... [📎](#) [😊](#)

[Help](#)

Campus.UP Workspace



The screenshot shows a web browser window with the URL `campusup.uni-potsdam.de/group/representation-learning/`. The interface features a top navigation bar with the University of Potsdam logo, the text "Campus.UP", and user information for "SEBASTIAN STOBER". Below this is a secondary navigation bar with "UP SERVICES" and "HELP" buttons. The main content area is divided into three sections: "Desktop", "Organizer", and "Campus Navigator". A breadcrumb trail shows the current location: "Representation Learning / Overview".

On the left side, there is a "WORKSPACES" sidebar with a menu containing "Member Pages", "Overview" (selected), "Documents & Resources", "Tools", "Organisational", "Members", and "Individual Blogs".

The main content area displays a "Welcome!" message: "This is the group workspace for the course 'Representation Learning' in the winter term 2016/2017. For more general information on the course, please visit the public course website." Below this is an "Announcements" section with "Add Entry" and "Manage Entries" buttons, and a message stating "There are no new announcements." The "Activity" section shows a recent post by "Sebastian" who "Wrote a new blog entry." 17 minutes ago.

At the top right of the workspace, there are buttons for "Editing Mode" and "Administration".

Session Summaries

last episode on
“Introduction to Deep Learning”

...

- rotating job!
(2 persons per session,
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 Channel

- document your learning / project progress
 - one post per week
 - share your experiences!
 - visible only to course participants
- examples:
 - <https://deerandommumbling.wordpress.com/>
 - <http://bartvanmerrienboer.nl/#blog>
- guidelines:
 - <https://www2.uwstout.edu/content/profdev/rubrics/blogrubric.html>

Q & A Channel

- guide for what is covered in class
deadline: Friday noon (12:00)
- 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 channel and Q&A
- comment / like / rate
- answer
- document
 - hints, tricks & hacks
- recommend
 - additional readings (papers, blogs, etc.)
- give (constructive) feedback

Important Change

- programming exercises are mandatory!
- solutions need to be submitted by Sunday 18:00 (day before lecture)
- you should be able to present your solution in course (random selection)
- 1 missed exercise can be excused
- more than 1 missed exercise leads to exclusion from course by May 20!

Course Project

Course Project

- collaborative effort
- “coopetition” (cooperative + competition)
- multiple teams of 2-4 students:
 - discuss ideas and form team
 - self-organized (heterogeneous if possible)
 - scrum-like approach
 - focus on different goals / aspects / strategies

Team Channels

- 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