


Introduction to Deep Learning

CNN papers

Andreas Krug, M.Sc.

ankrug@uni-potsdam.de

12. November 2018



Group exercise

CNN papers

CNN papers

1

Kalchbrenner et al.
CNN sentence
modelling

3

Gatys et al.
image
style transfer

5

Jansson et al.
singing voice
separation with U-Net

2

Pons et al.
music
audio tagging

4

Dong et al.
image
super resolution

6

Redmon et al.
YOLO real-time
object detection



1. Expert presents
for example including:
 - learning problem & specific challenges
 - network design
 - main innovations
 - main experimental findings
 - weak spots?
 - performance & hardware requirements?
2. Group asks questions
3. After 9': Next expert's turn

Programming tf.Estimator

open discussion

Feeding data to Estimator

```
train_img = np.load(directory + "mnist_train_imgs.npy")
train_lab = np.load(directory + "mnist_train_lbls.npy")

def train_input_fn(features, labels, batch_size):
    # convert input to dataset
    dataset = tf.data.Dataset.from_tensor_slices({"img": features, labels})
    # shuffle, repeat, and batch the examples
    dataset = dataset.shuffle(1000).repeat().batch(batch_size)
    return dataset.make_one_shot_iterator().get_next()
...

mnist_classifier.train(
    input_fn=train_input_fn(train_img, train_lab, 100),
    steps=20000,
    hooks=[logging_hook])
```

TypeError: <BatchDataset shapes: ({img: (?, 784)}, (?,)),
types: ({img: tf.uint8}, tf.uint8)> is not a callable object

input_fn = **lambda**: train_input_fn(train_img, train_lab, 100),

Assignments until next week

- Responsible for recap: Alexander & Ankur
- Reading:
Recurrent/Recursive Neural Networks part I
- (last mandatory) Programming exercise:
TF object detection API

Slides & assignments on: https://mlcogup.github.io/idl_ws18/