Beyond Deep Learning: Uncertainty Aware Models

Christian Tomani, Yuesong Shen

Technical University of Munich

Chair of Computer Vision and Artificial Intelligence

Garching, April 22nd, 2020

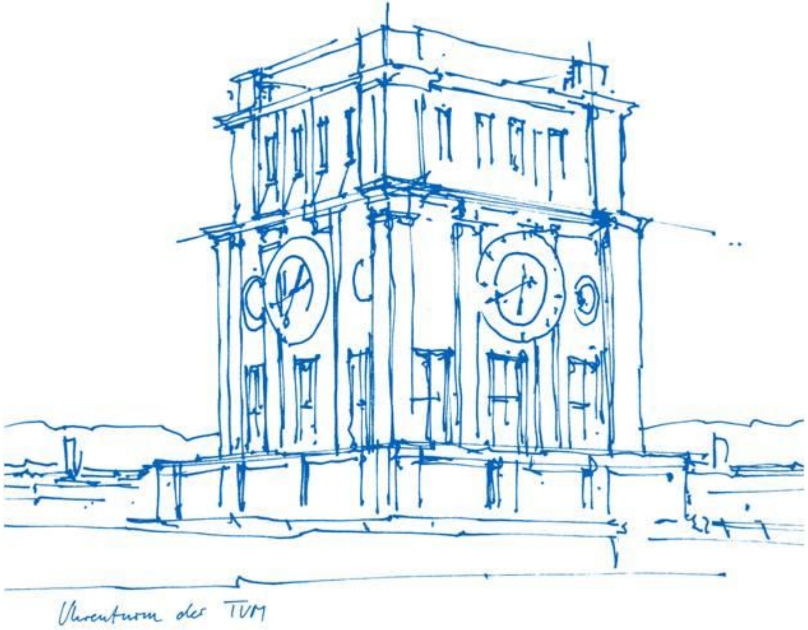

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Agenda

- Why do we need uncertainty aware models?
- What are the topics we will cover?
  - Project no.1: Folding of Proteins
  - Project no.2: Uncertainty aware Differentiable Neural Computers
  - Project no.3: Time Series Analysis on real world datasets
  - Project no.4: Interpretability of neural networks
  - Project no.5: Hierarchical graphical models
  - Project no.6: Sum-product network
  - Project no.7: Neural network as Gaussian process
- How is the course organized?
- How to use our resources?

Why do we need uncertainty aware models?

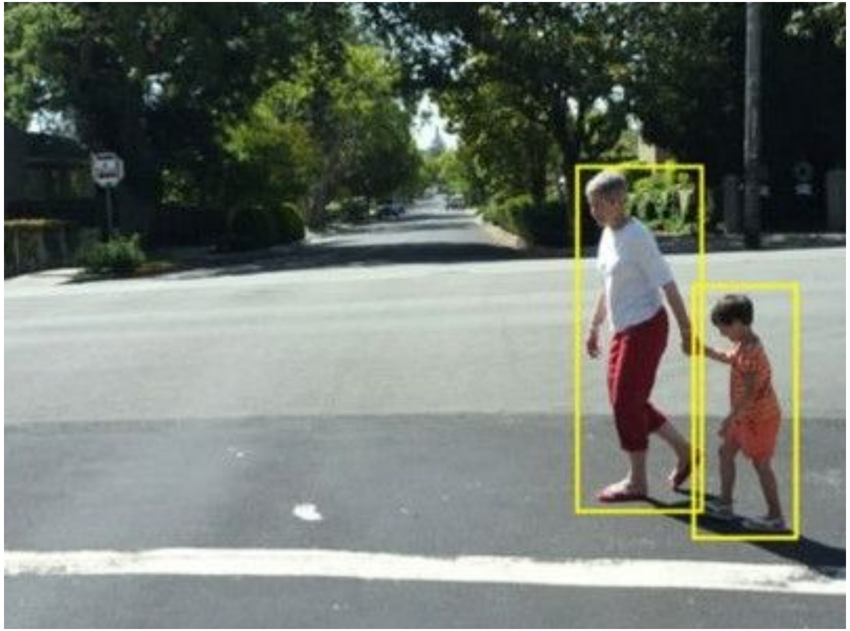

https://www.paperswithcode.com, 03/01/2020

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

The issue with Deep Learning - Can we trust the model?

Setup

LeNet-5 Model with weight decay

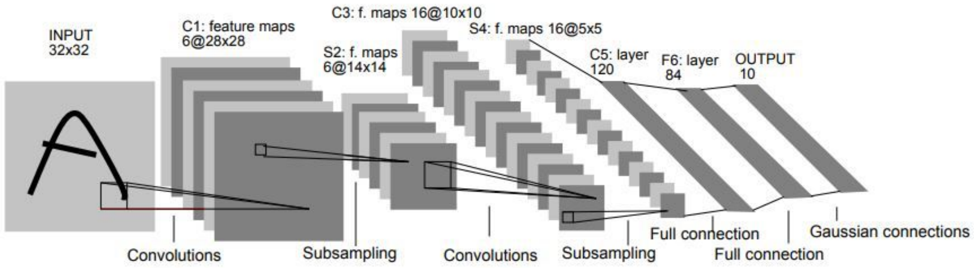

MNIST Dataset

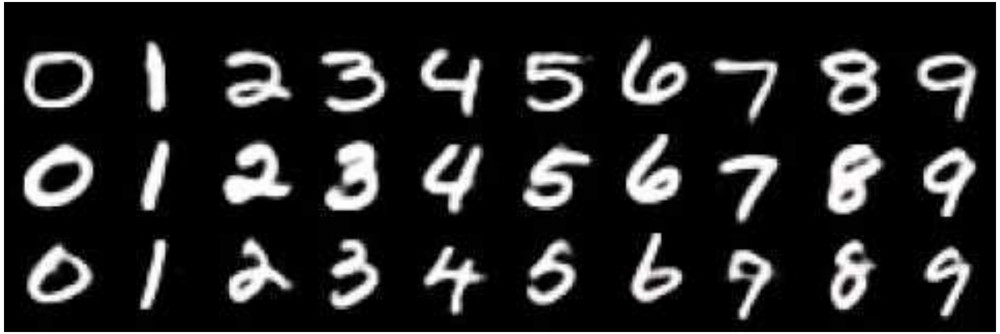

LeCun et al. - Gradient Based Learning Applied to Document Recognition, 1998
https://github.com/cazala/mnist

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

The issue with Deep Learning - Can we trust the model?

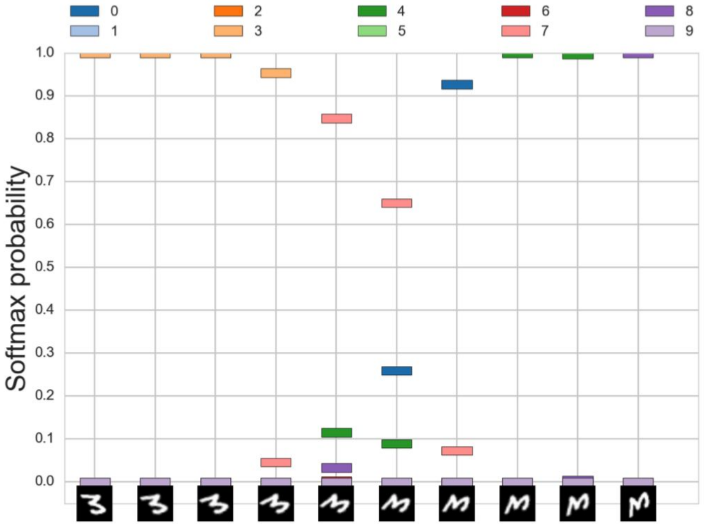

Vanilla LeNet-5 Model on MNIST

- Model is unreliable and not calibrated
- Gives totally wrong but highly confident predictions if data is perturbed
- wrong predictions cannot be distinguished from correct ones

→ Softmax Probability ≠ Confidence

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

The issue with Deep Learning - Can we trust the model?

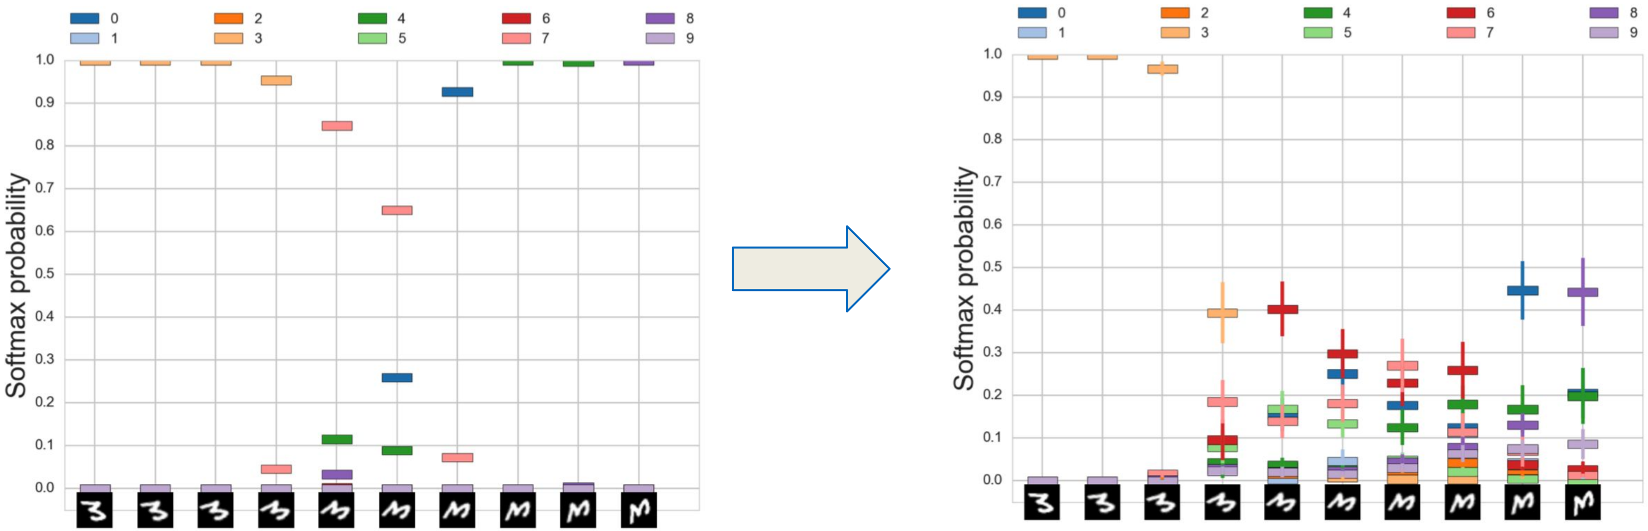

Christos Louizos - Multiplicative Normalizing Flows for Variational Bayesian Neural Networks, 2017

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

What are the topics we will cover?

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.1: Folding of Proteins

Understand and implement AlphaFold with ProteinNet as well as improve uncertainty awareness:

- AlphaFold is the state of the art model for predicting secondary and tertiary structures of proteins
- ProteinNet is a big scale real world dataset of amino acid sequences
- Goal:
    - Predict the 3D-structure of proteins based solely on the DNA sequence
    - Increase uncertainty awareness of predictions through implementing methods such as MC-Dropout and Deep Ensembles.

Some initial literature:

- https://www.nature.com/articles/s41586-019-1923-7
- https://github.com/deepmind/deepmind- research/tree/master/alphafold_casp13
- https://bmcbioinformatics.biomedcentral.com/articles/10.1186/s12859-019-2932-0

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.1: Folding of Proteins

Alpha Fold:

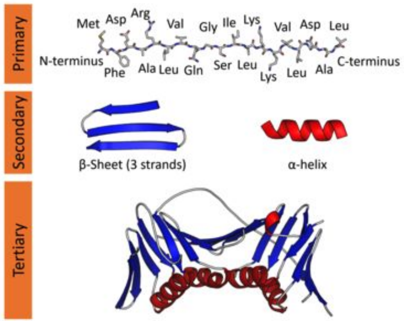

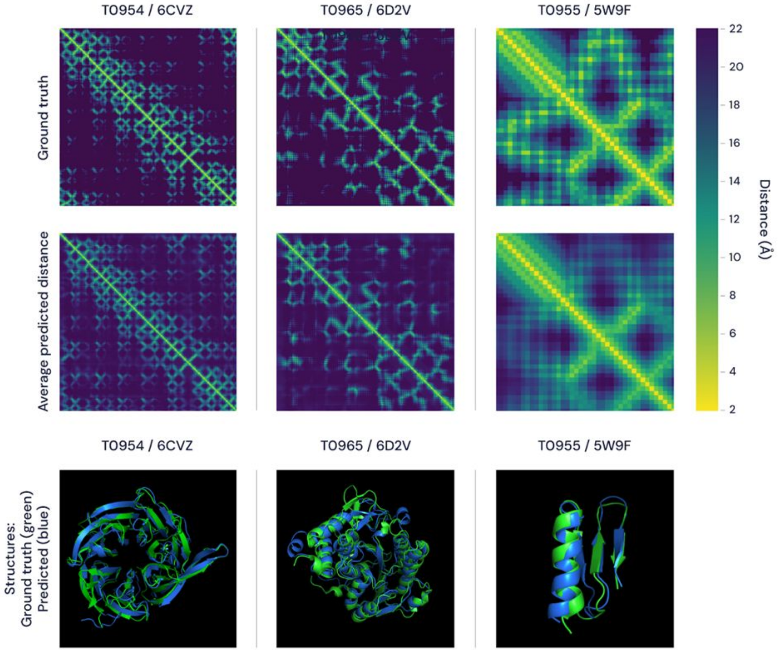

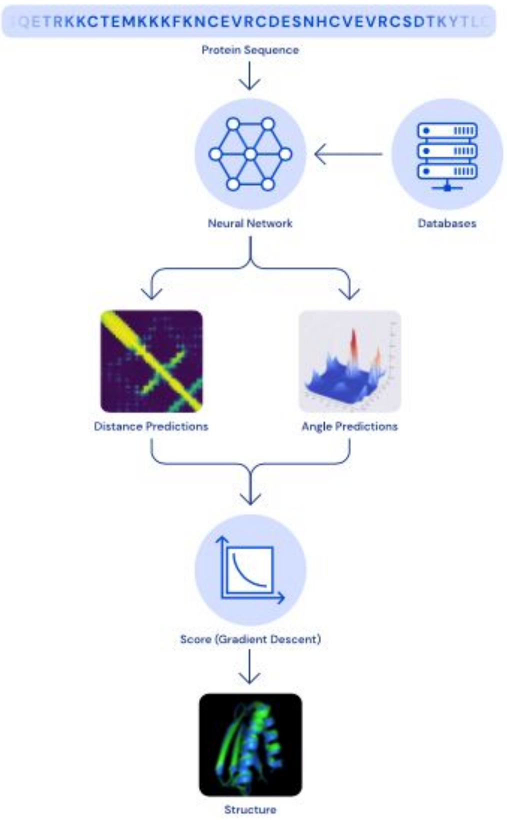

https://deepmind.com/blog/article/AlphaFold-Using-AI-for-scientific-discovery

Project details

Project no.2: Uncertainty aware Differentiable Neural Computers

Understand and implement the Differentiable Neural Computer (DNC) and the bAbI dataset in the context of uncertainty aware predictions:

- DNC is a fully differentiable and trainable model with memory as well as attention
- bAbI is algorithmically generated question answering dataset containing 20 different tasks.
- Goal:
    - Train and evaluate the DNC on datasets (e.g., bAbI)
    - Increase uncertainty awareness of predictions through implementing methods such as MC-Dropout and Temperature Scaling.

Some initial literature:

- https://www.nature.com/articles/nature20101
- https://github.com/deepmind/dnc
- Weston, J., Bordes, A., Chopra, S. & Mikolov, T. Towards AI-complete question answering: a set of prerequisite toy tasks. Preprint at http://arxiv.org/ abs/1502.05698 (2015).

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.2: Uncertainty aware Differentiable Neural Computer

Differentiable Neural Computer:

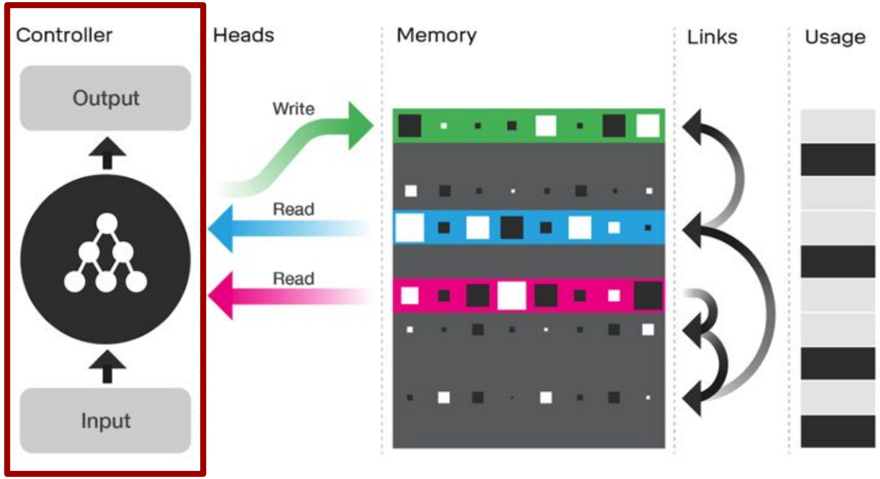

https://deepmind.com/blog/article/differentiable-neural-computers

Project details

Project no.2: Uncertainty aware Differentiable Neural Computer

Dataset:

```
1 Mary moved to the bathroom.
2 John went to the hallway.
3 Where is Mary?        bathroom        1
4 Daniel went back to the hallway.
5 Sandra moved to the garden.
6 Where is Daniel?      hallway 4
```

```
1 Mary moved to the bathroom.
2 Sandra journeyed to the bedroom.
3 John went to the kitchen.
4 Mary took the football there.
5 How many objects is Mary carrying?    one
```

Some initial literature:

- https://www.nature.com/articles/nature20101
- https://github.com/deepmind/dnc
- Weston, J., Bordes, A., Chopra, S. & Mikolov, T. Towards AI-complete question answering: a set of prerequisite toy tasks. Preprint at http://arxiv.org/ abs/1502.05698 (2015).

Project details

Project no.3: Time Series Analysis on real world datasets

Understand and implement different types of RNNs on various time series real world datasets.

- Various Datasets from different domains (e.g., medicine, personal assistant keywords, text data)
- Work with Recurrent Neural Networks, Long-Short-Term-Memory Networks, GRUs
- Goal:
  - Preprocess datasets and train classifiers
  - Implement EDL (Evidential Deep Learning), MC-Dropout and Temperature Scaling in order to increase uncertainty awareness of predictions.

Some initial literature:

- https://www.bioinf.jku.at/publications/older/2604.pdf
- https://arxiv.org/abs/1706.04599

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.4: Interpretability of neural networks

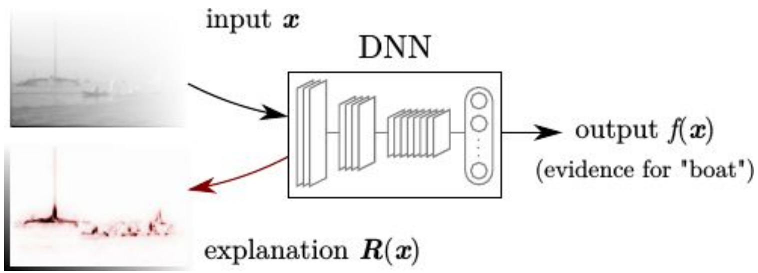

From "Methods for Interpreting and Understanding Deep Neural Networks"

Aim: obtain insights on trained NNs and decisions.

Task: Understand and implement NN interpretation techniques:

- e.g., Deconvolution, Layer-wise Relevance Propagation, etc.
- Categorization, comparison

Going further: risk of misinterpretation, effects of uncertainty aware methods ...

Some starting points for literature search:

- Methods for Interpreting and Understanding Deep Neural Networks
- https://github.com/1202kbs/Understanding-NN
- http://www.heatmapping.org/ [MICCAI'18 talk]

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.5: Hierarchical graphical models

Aim: probabilistic generative modeling for generative / unsupervised tasks

Task: understand and implement hierarchical GM structures and learning algorithms:

- Deep belief network: contrastive divergence, annealed importance sampling …
- More models: deep Boltzmann machine, sigmoid belief network …
- More methods: mean field, Gibbs sampling, wake-sleep algorithm ...

Some starting points for literature search:

- DBN, DBM, SBN
- Guide on contrastive divergence AIS for DBN
- Chapter 20 of deep learning book

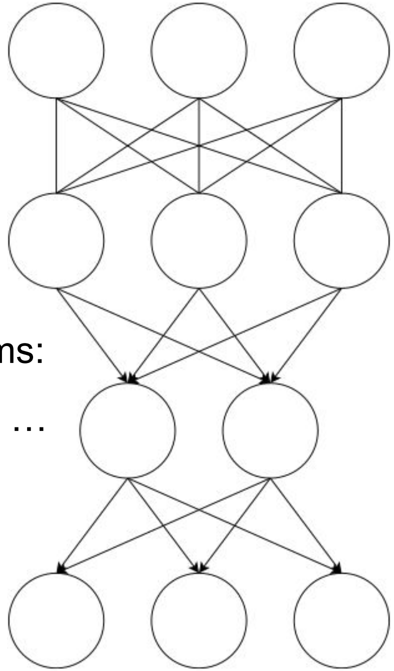

Illustration of a deep belief network

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Project details

Project no.6: Sum-product network

Aim: tractable probabilistic generative modeling

Task: understand and implement SPN

- generative / discriminative learning, node representation

Further directions: Conv SPN, conditional SPN, structure learning ...

Some starting points for literature search:

- [Sum-Product Networks: A New Deep Architecture]
- https://github.com/arranger1044/awesome-spn
- [A PyTorch implementation in SPFlow]

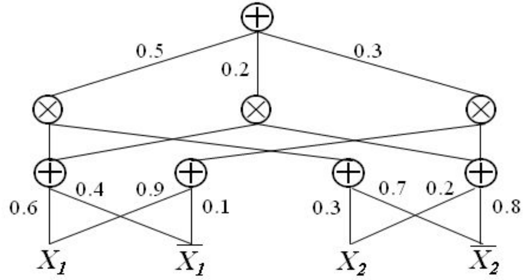

From: Sum-Product Networks: A New Deep Architecture

Project details

Project no.7: Neural network as Gaussian process

Aim: investigate the link between infinite width NN and Gaussian process

Task: understand and implement NN as GP to reproduce results from papers.

- deep neural networks as Gaussian processes, Neural tangent kernel

Further directions: deep GP, signal propagation, scalable GP methods ...

Some starting points for literature search:

- see above, plus
- Approximate Inference Turns Deep Networks into Gaussian Processes
- Ultra-Wide Deep Nets and the Neural Tangent Kernel (NTK)

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Your tutors

Christian Tomani

PhD Student

Prefers TensorFlow

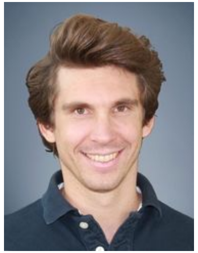

Supervises the following projects:

- 1. Folding of Proteins
- 2. Uncertainty aware Differentiable Neural Computers
- 3. Time Series Analysis on real world datasets

Yuesong Shen

PhD Student

Prefers PyTorch

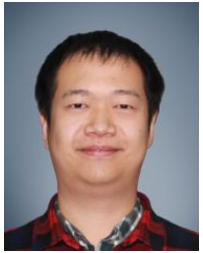

Supervises 3 of the following projects:

- 4. Interpretability of NN
- 5. Hierarchical graphical models
- 6. Sum-product network
- 7. NN as Gaussian process

Projects preferences

Information about project matching

The email should be sent to bdluam-ss20@vision.in.tum.de **latest April 26th** with the subject: “[Project Matching] <Firstname> <Lastname>” and contain the filled information form (template on course website, rename to “firstname_lastname.xlsx”).

Fill in:

- Immatriculation No.
- First name
- Last Name
- e-mail
- rank all projects from 1 (lowest preference) to 7 (highest preference)
- comment section

Chair of Computer Vision & Artificial Intelligence
Department of Informatics
Technical University of Munich

Course organization

Course website: https://vision.in.tum.de/teaching/ss2020/bdluam_ss2020
Material: https://vision.in.tum.de/teaching/ss2020/bdluam_ss2020/materials password: `uncertain_password`
Course e-mail: bdluam-ss20@vision.in.tum.de

- Course structure:
  - Approx. 2 students per team
  - Weekly progress reports (e-mail: progress, questions; if necessary attach pdf for results)
  - If needed weekly team meetings with the tutor
- 4 joint meetings along the semester:
  - Kickoff meeting: project intro & matching (April 22nd)
  - 2 progress meetings: project progress check points & exchange ideas (May 27th & June 24th)
  - Final presentations: final project delivery and presentation (July 22nd)

Course organization

Course website: https://vision.in.tum.de/teaching/ss2020/bdluam_ss2020
Material: https://vision.in.tum.de/teaching/ss2020/bdluam_ss2020/materials password: `uncertain_password`
Course e-mail: bdluam-ss20@vision.in.tum.de

- Evaluation:
  - Overall work done and presentations (in particular final presentation)
  - Code repository - running code and well documented (gitlab9.in.tum.de)
  - Final Report (continuous effort due to weekly reports)
- Next steps:
  - Students will get notified at the beginning of next week about their project
  - First team meetings will take place starting next week

Thank you! Questions?

Access the presentation:

https://vision.in.tum.de/teaching/ss2020/bdluam_ss2020/materials

password: uncertain_password