Beyond Deep Learning: Selected Topics

Christian Tomani, Yuesong Shen

Technical University of Munich

Chair of Computer Vision and Artificial Intelligence

Garching, April 27th, 2022

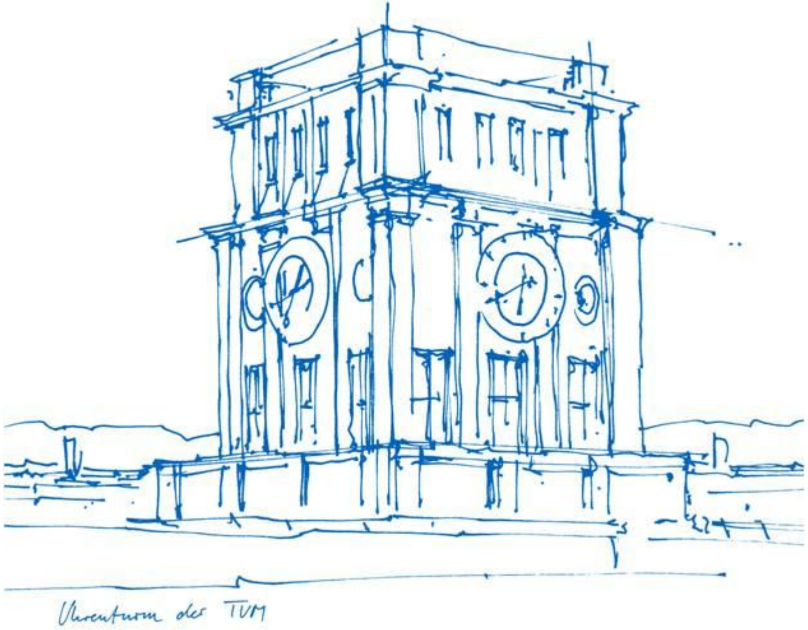

Agenda

- What are the topics we will cover?
    - Uncertainty awareness and calibration
    - Time series & sequence analysis
    - Layer and Architecture Designs
    - Alternatives to Neural Networks
- Project matching?
- How is the course organized?

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

Your tutors

Christian Tomani

PhD Student

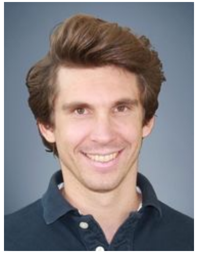

Supervises the following topics:

- 1. Calibration via post-processing methods
- 2. Uncertainty inducing methods
- 3. Self Supervised Learning with Masks
- 7. Transformer for Time Series
- 8. Time Series Analysis - IndRNN & Transformer
- 9. Sequence Analysis - Protein Folding

Yuesong Shen

PhD Student

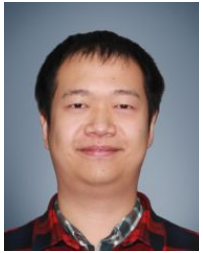

Supervises the following topics:

- 4: Uncertainty via stochastic variational inference
- 5: Uncertainty via MCMC
- 6: Uncertainty via ensemble learning
- 10: Self-supervised representation learning
- 11: Learning in vision beyond CNNs
- 12: Sum-product network: an NN alternative?

Topic 1: Calibration via post-processing methods

Introduction:

Calibration denotes that the probability of a system output for an event should reflect the true frequency of that event. Calibration can be increased via post-processing outputs of neural networks.

Goal:

- Understand how calibration is calculated and what issues arise when estimating it.
- Compare different calibration inducing methods (scaling-binning calibrator, TS, PS, ODIR).

Literature:

- Verified Uncertainty Calibration, Kumar et al., 2020
- Beyond temperature scaling: Obtaining well-calibrated multiclass probabilities with Dirichlet calibration, Kull et al
- On Calibration of Modern Neural Networks, Guo et al, 2017
- Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods, Platt et al., 1999

Topic 2: Uncertainty inducing methods

Introduction:

Calibration denotes that the probability of a system output for an event should reflect the true frequency of that event.

Goal:

- Understand calibration and describe basic calibration inducing methods: Intra Order-Preserving methods, multi-class uncertainty calibration via mutual information maximization-based binning, calibration of neural networks using splines.
- Compare these 3 methods to each other.

Literature:

- Intra Order-Preserving Functions for Calibration of Multiclass Neural Networks, Rahimi et al., 2020
- Multiclass Uncertainty Calibration via Mutual Information Maximization-based Binning, Patel et al., 2021
- Calibration of Neural Networks using Splines, Gupta et al., 2021

Topic 3: Self Supervised Learning with Masks

Introduction:

While the general idea of self-supervised learning is identical across modalities, the actual algorithms and objectives differ widely because they were developed with a single modality in mind.

Goal:

- Understand self-supervised learning with tokens and masks in images and sequences
- Compare different methods

Literature:

- data2vec: A General Framework for Self-supervised Learning in Speech, Vision and Language, Baevski et al., 2022
- BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, Devlin et al., 2019
- BEIT: BERT Pre-Training of Image Transformers, Bao et al. 2021

Topic 4: Uncertainty via stochastic variational inference

Introduction:

A standard approach to perform scalable inference and learning generic probabilistic models (including Bayesian NNs) is to employ SVI (check out probabilistic programming libraries e.g., Pyro / Stan ...).

Goal:

- Understand and summarize the SVI approach
- How to make SVI work with BNN in practice? What are the possible extensions?

Literature:

- Weight Uncertainty in Neural Networks, Blundell et al., 2015
- Auto-Encoding Variational Bayes, Kingma and Welling, 2013
- Radial Bayesian Neural Networks: Beyond Discrete Support In Large-Scale Bayesian Deep Learning, Farquhar et al, 2021
- Flipout: Efficient Pseudo-Independent Weight Perturbations on Mini-Batches, Wen et al, 2020

Topic 5: Uncertainty via MCMC

Introduction:

An alternative line of work uses Markov chain Monte Carlo to obtain stochastic estimations of the posterior distribution for Bayesian deep learning.

Goal:

- Understand and summarize the MCMC approaches
- How do they apply to Bayesian neural networks in practice? What are the challenges?

Literature:

- (Intro on sampling: Chapter 27 of “Bayesian Reasoning and Machine Learning”, Barber, 2011)
- A Conceptual Introduction to Hamiltonian Monte Carlo, Betancourt, 2018
- A Complete Recipe for Stochastic Gradient MCMC, Ma et al., 2015
- Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning, zhang et al, 2020

Topic 6: Uncertainty via ensemble learning

Introduction:

Another line of approach to introduce uncertainty is via the idea of ensembles: training multiple versions of the model and aggregate their predictions. This is a simple idea that works surprisingly well in practice.

Goal:

- Search for, summarize and compare related work such as SWA, SWAG, deep ensemble …
- Compare with dropout-based uncertainty methods: MC-dropout, variational/concrete dropout ...
- Compare with traditional Bayesian model averaging approach.

Literature:

- A Simple Baseline for Bayesian Uncertainty in Deep Learning, Maddox et al, 2019
- Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles, Lakshminarayanan et al, 2017

Topic 7: Transformer for Time Series

Introduction:

Time series data are prevalent in many scientific and engineering disciplines. Transformer-based machine learning models can also be used to forecast time series data.

Goal:

- Understand transformer models and attention based methods
- Compare the models to each other and show pros and cons.

Literature:

- Deep Transformer Models for Time Series Forecasting: The Influenza Prevalence Case, Wu et al., 2020
- Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting, Zhou et al, 2021
- Attend and diagnose: Clinical time series analysis using attention models, Song et al., 2018

Topic 8: Time Series Analysis - IndRNN & Transformer

Introduction:

Independently RNNs are models where neurons in the same layer are independent of each other.

Transformers are a form of time series models that rely on attention mechanisms without recurrences.

Goal:

- Understand the methods and show the differences: IndRNN, Transformer, LSTMs, RNNs.
- Compare the IndRNN and the Transformer models to all benchmark models used in the papers.
- Look out and discuss new models/methods based on transformers.

Literature:

- Attention Is All You Need, Vaswani et al., 2017
- Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context, Dai et al., 2019
- Independently Recurrent Neural Network (IndRNN): Building A Longer and Deeper RNN, Li et al., 2018
- Long short-term memory, Hochreiter et al, 1997

Topic 9: Sequence Analysis - Protein Folding

Introduction:

Protein folding is a task in bioinformatics where amino-acid sequences are fed to a model with the goal of predicting the 3D structure of proteins.

Goal:

- Understand the following methods: AlphaFold, AlphaFold2, RaptorX, TripletRes.
- Compare the 4 methods based on their inputs, architectures and results.

Literature:

- Improved protein structure prediction using potentials from deep learning, Senior et al., 2020
- Highly accurate protein structure prediction with AlphaFold, Jumper et al. 2021
- Analysis of distance-based protein structure prediction by deep learning in CASP13, Xu et al., 2019
- Ensembling multiple raw coevolutionary features with deep residual neural networks for contact-map prediction in CASP13, Li et al, 2019

Topic 10: Self-supervised representation learning

Introduction:

Traditional deep learning methods have been mainly focused on supervised learning problems which rely on ground-truths. More recent studies try to overcome this limitation and use neural networks to solve learning tasks in a self-supervised manner.

Goal:

- Survey different self-supervised deep learning methods
- Summarize and compare components that make self-supervised deep learning work

Literature:

- A Simple Framework for Contrastive Learning of Visual Representations, Chen et al., 2020
- Exploring Simple Siamese Representation Learning, Chen and He, 2020
- Barlow Twins: Self-Supervised Learning via Redundancy Reduction, Zbontar et al., 2021

Topic 11: Learning in vision beyond CNNs

Introduction:

Since the breakthrough of AlexNet in 2012, CNNs have been the predominant choice for vision tasks.

Recent studies consider the idea of attention and its combination with convolutional models.

Goal:

- Search for and summarize related work on attention (+ convolution) models for vision tasks
- What are the strong/weak points of attention/convolution? Can we get the best of both worlds?

Literature:

- Attention Augmented Convolutional Networks, Bello et al, 2019
- An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, Dosvotskiy et al, 2020
- Involution: Inverting the Inherence of Convolution for Visual Recognition, Li et al, 2021

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

Topic 12: Sum-product network: an NN alternative?

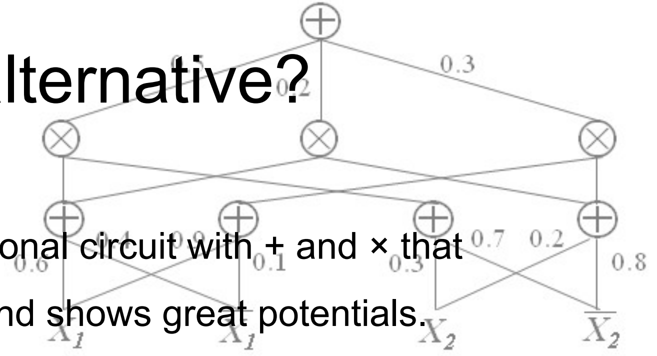

Introduction:

SPN is an interesting alternative of NN, which is a hierarchical computational circuit with + and × that encode a joint probabilistic distribution. SPN is also quite easy to train, and shows great potentials.

Goal:

- Understand and summarize SPN structures, inference and learning methods.
- Compare to NNs: what are the advantages/disadvantages of SPN?

Literature:

- Sum-Product Networks: A New Deep Architecture, Poon and Domingos, 2011
- Sum-product networks: A survey, París et al, 2020
- https://github.com/arranger1044/awesome-spn

Original Image published in:
“Sum-Product Networks: A New Deep Architecture, Poon and Domingos, 2011”

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

Projects preferences

Information about project matching

The email should be sent to bdlstnc-ss22@vision.in.tum.de **latest May 4th** with the subject:
“[Project_Matching] <Firstname> <Lastname>” and contain the filled information form (template on course website, rename to “firstname_lastname.xlsx”, do not transform it to another format).

Fill in:

- Immatriculation No.
- First name
- Last Name
- e-mail
- rank all projects (=topics) from 1 (lowest preference) to 12 (highest preference)
- comment section

Indicate your preference for each topic from 1 (**lowest** preference) to 12 (**highest** preference). Use each number (=preference) only once!											
Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6	Topic 7	Topic 8	Topic 9	Topic 10	Topic 11	Topic 12

Course organization

Course website: https://vision.in.tum.de/teaching/ss2022/bdlstnc_ss2022
Course e-mail: bdlstnc-ss22@vision.in.tum.de

- Course structure:
  - 1 student per topic
  - When needed you can contact your tutor and schedule a personal meeting
  - Highly recommended: meet with your tutor to discuss about the draft before presentation

Course organization

Course website: https://vision.in.tum.de/teaching/ss2022/bdlstnc_ss2022
Course e-mail: bdlstnc-ss22@vision.in.tum.de

- Evaluation:
  - Presentation (20-25 min) and Report (10-15 pages)
    - Draft: June 22nd
    - Oral Presentations: July 5th from 1pm-4pm and July 6th 9am-4pm
    - Final Version: July 17th
  - Participation
- Next steps:
  - Students will get notified at the end of next week about their project

Guidelines about Presentation & Report

- Presentation:
  - Duration: 20-25 min
  - Content: Sufficient materials, Introduction, Method description, Results, etc.
  - Layout: Consistent visuals, proper graphs and diagrams
  - Presentation Style: Clear, easy to follow, engaging for the audience
  - Proper citations
- Report:
  - Length: 10-15 pages
  - Scientific layout, paper style (Introduction, Methods, Results, Conclusion, etc.)
  - Clear Explanation, consistent notations, technical details
  - Proper citations

Thank you! Questions?