

# Seminar Advanced NLP Intro Session

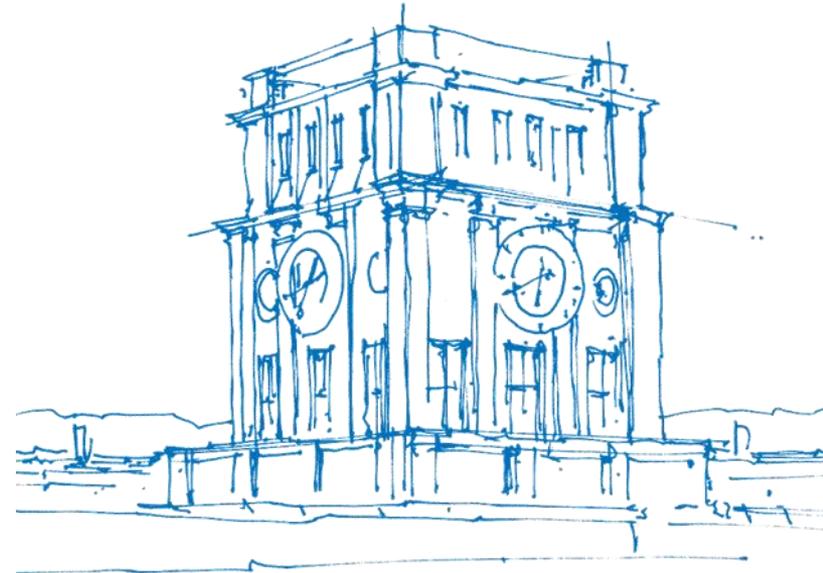
Master-Seminar – Advanced Techniques in Natural Language Processing (IN2107, IN4428)

Tobias Eder, M.A. M.Sc.

Prof. Dr. Georg Groh

Research Group Social Computing, Department of Informatics,  
Technical University of Munich

22.10.2021



*TUM Uhrenturm*

# Contents

- **Developments in NLP**
- This Seminar (Dates, Expectations, Guidelines)
- Topics
- Next Steps

# Why this seminar?

- NLP is a **fast-evolving field**
- There are a lot of **different facets** to NLP research
- Two main goals:
  - Establish **familiarity** with current NLP research
  - Look at **impactful scientific papers** to understand what makes them impactful.

# Contents

- Developments in NLP
- This Seminar (Dates, Expectations, Guidelines)
- Topics
- Next Steps

# Seminar Schedule

- Topic selection by next week ([mail](#))
- Four / Six seminar sessions in December and January
- Seminar paper OR (tbd) blogpost style submission due by February 28<sup>th</sup>

# Topic sessions

Seminar dates **choice A**:

- December 3<sup>rd</sup> & December 17<sup>th</sup>
- January 14<sup>th</sup> & January 28<sup>th</sup>
- Time: 10 am to 12 pm on Fridays
- **3 topics** per meeting

# Topic sessions

Seminar dates **choice B**:

- December 3<sup>rd</sup> & December 17<sup>th</sup>
- January 14<sup>th</sup> & January 21<sup>st</sup> & January 28<sup>th</sup> & February 4<sup>th</sup>
- Time: 10 am to 12 pm on Fridays
- **2 topics** per meeting

# Sessions

- Everyone must read the [main paper](#) for each topic
- You are encouraged to [take notes and prepare questions](#) on things that are unclear to you.
- There will be one [expert](#) on each topic who has prepared a presentation on the topic and has done additional background reading.

# Presentations

- **25-minute length presentation** on your chosen topic
- Technical presentation on the contents of the main paper, including additional background info (which subfield of NLP or otherwise, extended related work, preliminaries to better understand the paper)
- **15–20-minute discussion** on the topic after each presentation

# Seminar paper

Due date by **February 28<sup>th</sup>** (Monday) “anywhere on earth”

- ~**8 pages** long (template see webpage)
- On your presentation topic in more detail
- Should include **discussion on applications** and **impact** of the work (!)
- Bonus: Can include ideas for future work

# Contents

- Developments in NLP
- This Seminar (Dates, Expectations, Guidelines)
- **Topics**
- Next Steps

# Topics

## Topic 1: Renewing modern Deep NLP

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin. Attention Is All You Need. NeurIPS 2017.

<https://arxiv.org/abs/1706.03762>

## Topic 2: (Perceived) Transparency of NLP Models

Sarthak Jain, Byron C. Wallace. Attention is not Explanation. NAACL 2019.

<https://aclanthology.org/N19-1357/>

# Topics

## Topic 3: Complex Representations

Jacob Devlin, Ming-Wei Chang, Kenton Lee and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL 2019.

<https://aclanthology.org/N19-1423/>

## Topic 4: Benefits of the Compute Arms Race

Stephen Merity (2019). Single Headed Attention RNN: Stop Thinking With Your Head. arXiv.

<https://arxiv.org/abs/1911.11423>

# Topics

## Topic 5: Intrinsic Knowledge through Language

Adam Roberts, Colin Raffel, Noam Shazeer. How Much Knowledge Can You Pack Into the Parameters of a Language Model? EMNLP 2020.

<https://aclanthology.org/2020.emnlp-main.437/>

## Topic 6: Probabilistic Assumptions of NLP

Clara Meister, Tim Vieira, Ryan Cotterell (2020). If beam search is the answer, what was the question? EMNLP 2020

<https://aclanthology.org/2020.emnlp-main.170/>

# Topics

## Topic 7: Size Matters Most?

Tom B. Brown, Benjamin Mann, [Open AI Team] (2020). Language Models are Few-Shot Learners. NeurIPS 2020.

<https://papers.nips.cc/paper/2020/hash/1457c0d6bfcb4967418bf8ac142f64a-Abstract.html>

## Topic 8: Grounding Language Processing

Yonatan Bisk, Ari Holtzman et al. (2020). Experience Grounds Language. EMNLP 2020

<https://aclanthology.org/2020.emnlp-main.703/>

# Topics

## Topic 9: Learning to Learn

Farhad Nooralahzadeh, Giannis Bekoulis, Johannes Bjerva, Isabelle Augenstein (2020). Zero-Shot Cross-Lingual Transfer with Meta Learning. EMNLP 2020.

<https://aclanthology.org/2020.emnlp-main.368/>

## Topic 10: Adapting Models to the Task

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey and Noah A. Smith. Don't Stop Pretraining: Adapt Language Models to Domains and Tasks. ACL 2020.

<https://aclanthology.org/2020.acl-main.740/>

# Topics

## Topic 11: Beyond Transformers

Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, Joao Carreira (2021). Perceiver: General Perception with Iterative Attention. arXiv 2021.  
<https://arxiv.org/abs/2103.03206>

## Topic 12: Learning without Forgetting

Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, Iryna Gurevych (2021). AdapterFusion: Non-Destructive Task Composition for Transfer Learning.  
<https://aclanthology.org/2021.eacl-main.39/>

# Contents

- Developments in NLP
- This Seminar (Dates, Expectations, Guidelines)
- Topics
- Next Steps

# Next Steps

- Think about which topics you're most interested in
- Your own topic ideas are also welcome
- Selection of topics: Send me a mail until **Oct 29<sup>th</sup>** stating **at least 3 preferences** and if you prefer 4 or 6 sessions.
- Topics and schedule will be announced on November 2<sup>nd</sup> (mail and website)

# Questions?