


Preliminary Meeting Seminar Ethics in NLP SS2021

Master-Seminar – E4N: Ethics in Natural Language Processing (IN2107)

Tobias Eder, M.A. M.Sc.,
Prof. Dr. Georg Groh

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

05.02.2021



TUM Uhrenturm

Outline


1. About this course
2. Your profile
3. Topic examples
4. Applying for a spot

About this course

E4N (Ethics for Nerds) is an existing seminar discussing topics from Computer Science from an ethical point of view.

This seminar is a more specialized version, dealing with problems from the area of Natural Language Processing (NLP) specifically.

The course tries to bridge the gap between state-of-the-art NLP research and its impact on society.



About this course

Why now?

- Ethical discussions around AI and NLP have been increasing over the past few years.
- The impact of NLP systems is felt much more broadly in society already.
- NLP is a technology that normal end users are most likely to be in contact with every single day.

Recently all major NLP conferences have started employing ethics guidelines and ethics reviews. Practitioners are expected to discuss their contributions from an ethical perspective.



Semester plan

- At the start of the semester there will be 1-2 introductory sessions to cover basics.
- After the intro sessions individual topics will be distributed.
- You will prepare a (30-minute) presentation for your topic.
- After each presentation we will take an additional 30 minutes for discussions.
- Presentations will be held towards the end of the semester (late June – early July) and in block.

Deliverables:

- 30-minute presentation towards the end of the semester.
- ~8-page seminar paper (Deadline in the lecture free period in August t.b.d.) on your topic.
- Mandatory attendance for all presentations and discussion sessions.

Your profile

Minimum:

- Master's level student of any discipline
- Basic knowledge about Natural Language Processing

Important:

- Basic knowledge about machine learning and deep learning
- Basic proficiency in at least one programming language

Optimal:

- Knowledge about current (ethical) issues in NLP research
- Some experience tackling practical machine learning problems

Topic examples

Data sources, labeling and bias


- Where does our NLP data come from?
- How do the processes to collect and label data affect systems?

Representation learning and semantic information

- What information can we gain from learned embeddings that are not obvious?
- How much simplification is going on when we reduce dimensionality?

Automatic text generation and language models

- What problems can we solve with models like GPT-3?
- What problems do they create?



Topic examples

Privacy and transparency in NLP

- How can NLP be designed with privacy in mind?
- What about transparency issues and right-to-explain?

Robustness and probing NLP models

- How can models be deliberately abused in harmful ways?
- Can we learn sensitive information from NLP models in unforeseen ways?

Environmental impact of NLP

- How bad is an environmental footprint of developing and deploying NLP systems?
- Can we weigh costs and benefits? Is it even really an issue?



Applying for a spot

- Until **16 Feb**, fill out the registration form:
<https://wiki.tum.de/display/socialcomputing/Application+E4N%3A+Ethics+in+Natural+Language+Processing+SS2021>
- This will be the basis for us doing the matching from our side.
- Until **16 Feb**, you also have to register for the course on the matching system!
- Around the **25 Feb**, you are (probably) notified by the matching system about the status of participation.
- If you get assigned to the course you will be then sent an E-Mail with further steps.

Questions?