

# Preliminary Meeting Seminar Ethics in NLP SS2023


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

Tobias Eder, M.A. M.Sc.,

Prof. Dr. Georg Groh

Research Group Social Computing,  
School of Computation, Information and Technology,  
Technical University of Munich

01.02.2023



*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.



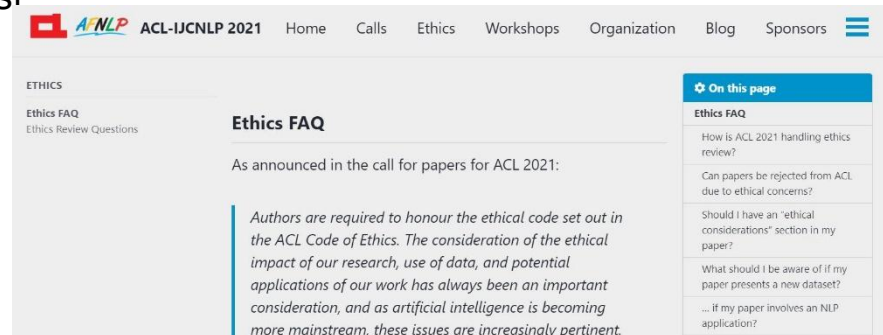
# About this course

## 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.

The image is a screenshot of the ACL-IJCNLP 2021 website, specifically the Ethics FAQ page. The navigation bar at the top includes links for Home, Calls, Ethics, Workshops, Organization, Blog, and Sponsors. The main content area is titled 'Ethics FAQ' and contains the following text: 'As announced in the call for papers for ACL 2021: *Authors are required to honour the ethical code set out in the ACL Code of Ethics. The consideration of the ethical impact of our research, use of data, and potential applications of our work has always been an important consideration, and as artificial intelligence is becoming more mainstream, these issues are increasingly pertinent.*' On the right side, there is a 'On this page' sidebar with a list of questions: 'How is ACL 2021 handling ethics review?', 'Can papers be rejected from ACL due to ethical concerns?', 'Should I have an "ethical considerations" section in my paper?', 'What should I be aware of if my paper presents a new dataset?', and '... if my paper involves an NLP application?'.

# Semester plan

- At the start of the semester (end of April) there will be an introductory sessions to cover basics.
- After the intro sessions individual topics will be distributed.
- The seminar will take place on a bi-weekly schedule starting in June.
- You will prepare a (25-minute) presentation for your topic.
- After each presentation we will take an additional 20 minutes for discussions.

## Deliverables:

- 25-minute presentation.
- ~8-page seminar paper (Deadline end of August) 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?

## NLP for manipulation and misinformation

- How can NLP be used to foster and/or fight misinformation?
- How does the state of the art of these systems look like?

## Automatic text generation and language models

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



# Topic examples

## Explainability for NLP

- How can NLP systems be made transparent?
- What methods for explainability are there and how do they work?

## Robustness of NLP methods

- 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 **14 Feb**, fill out the registration form:  
<https://wiki.tum.de/display/socialcomputing/E4N%3A+Ethics+in+Natural+Language+Processing+SS2023>
- This will be the basis for us doing the matching from our side.
- Until **14 Feb**, you also have to register for the course on the matching system!
- Around the **23 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?