Seminar Natural Language Processing (NLP) — Part 1

Introduction to NLP and Computational Sociolinguistics

Henning Wachsmuth

https://ai.uni-hannover.de

Natural Language Processing

Leibniz Universität Hannover

Outline

- **Motivation**

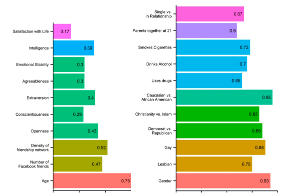

(Kosingki et al., 2013)

- **Natural language processing (NLP)**

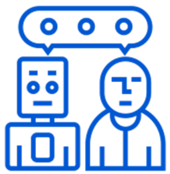

https://thenounproject.com

- **Computational sociolinguistics (CSL)**

https://thenounproject.com

- **CSL topics in this seminar**

- **Conclusion**

https://pixabay.com

Motivation

Example: Predictiveness of likes (Kosinski et al., 2013)

- **What social media likes reveal**

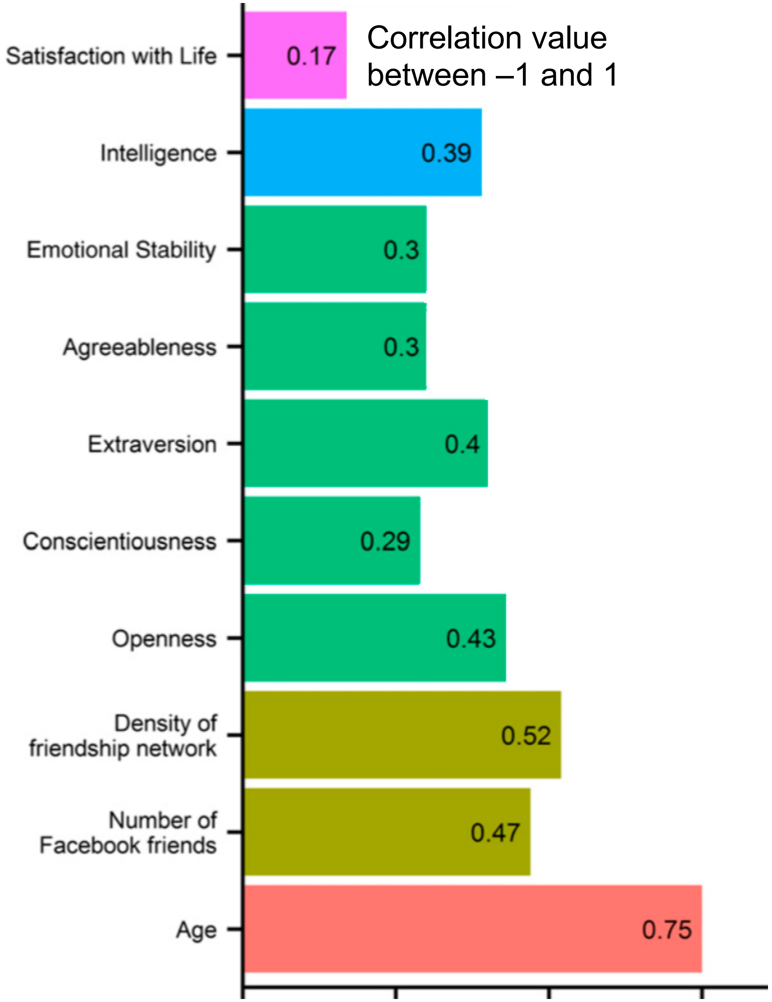

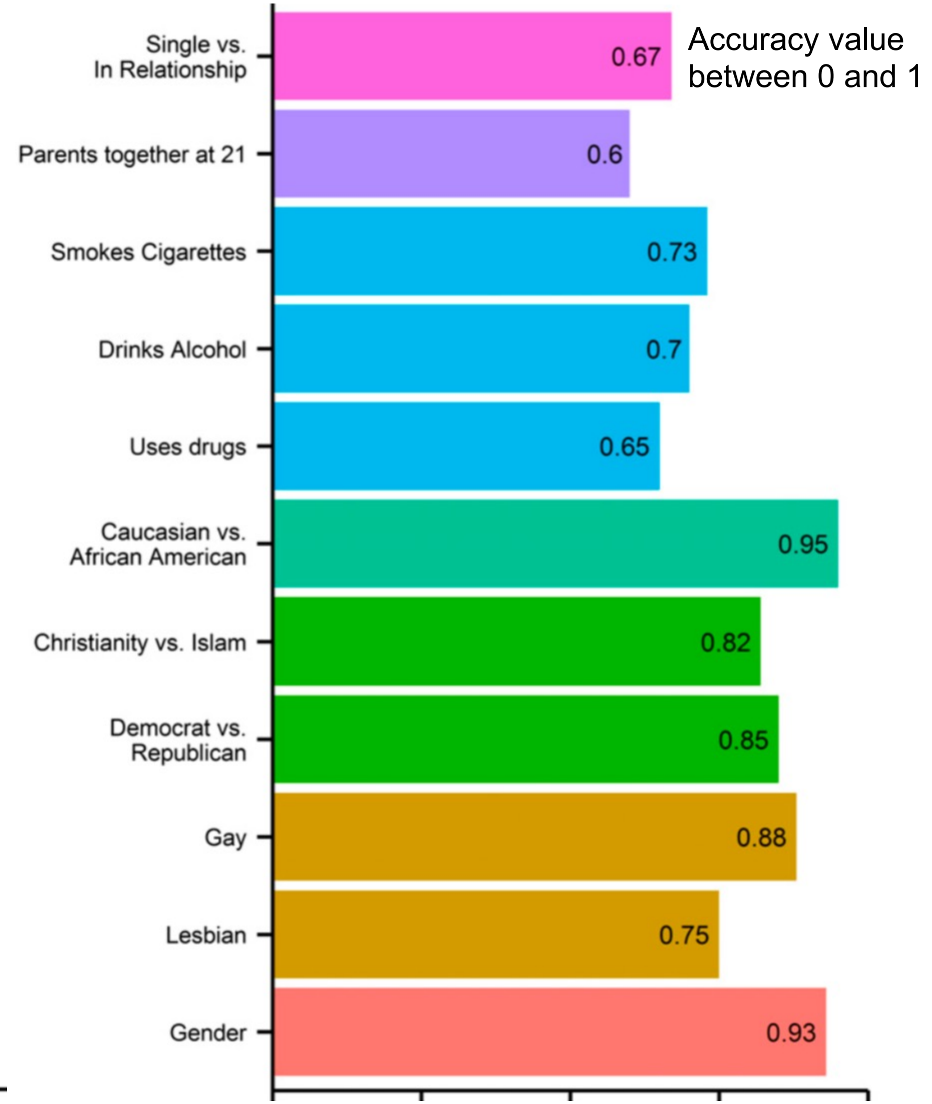

Example: Ethnicity-related police behavior (Voigt et al., 2017)

- **Language of US police officers toward black and white car drivers**

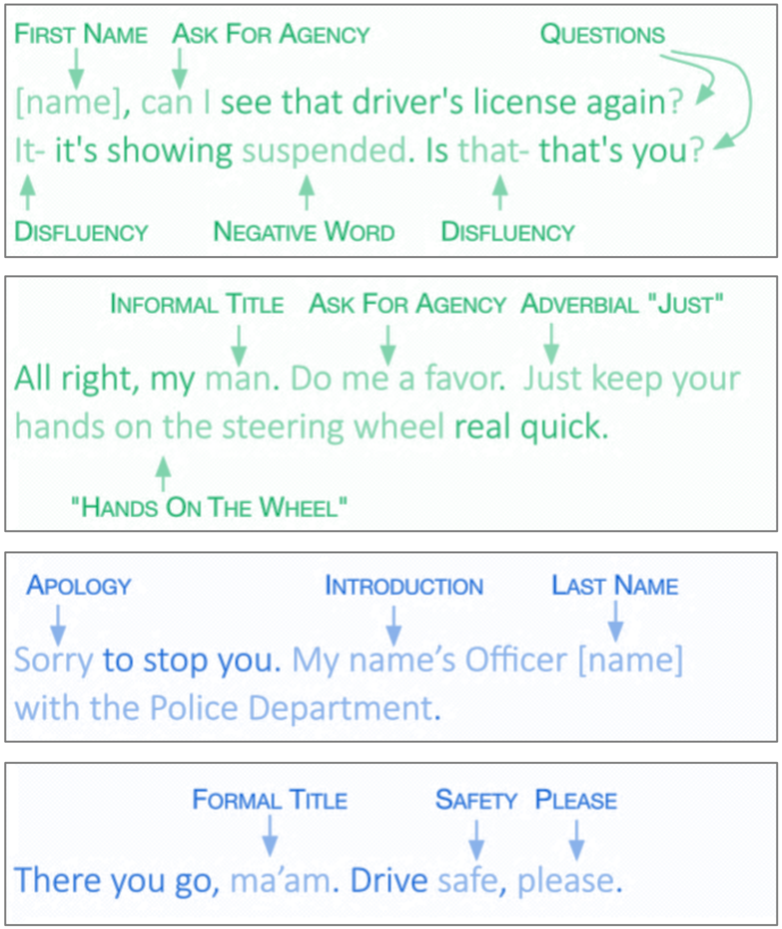

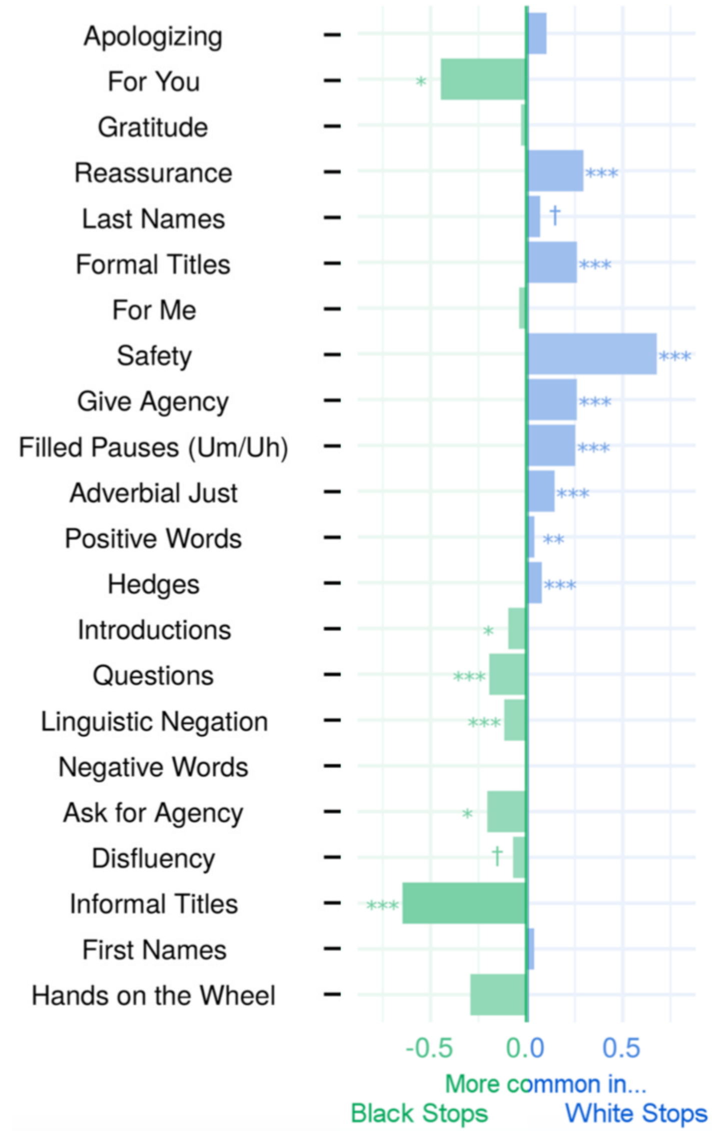

Example: Social bias in language models (Holtermann et al., 2022)

- **Social bias in language models**
  before fine-tuning and after fine-tuning on subjective language

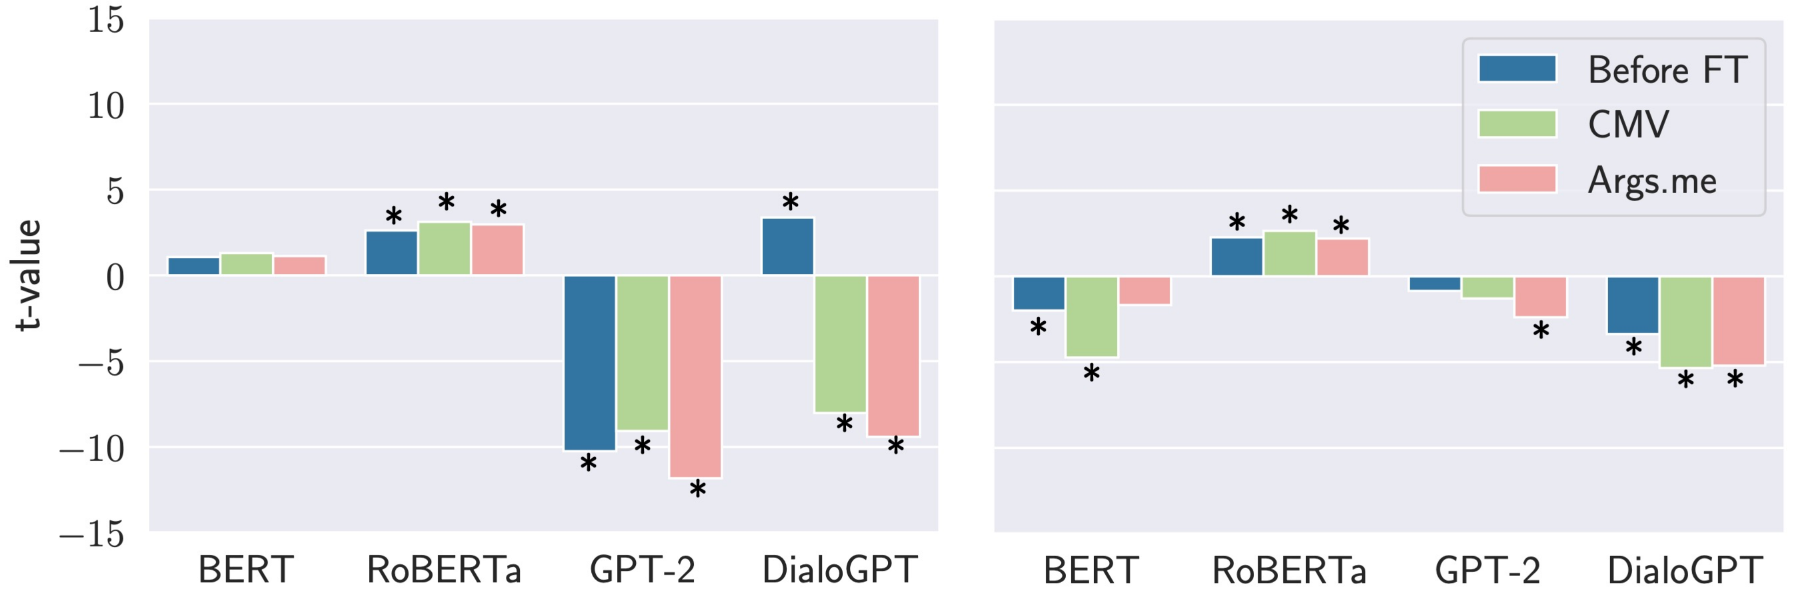

(a) LMB for Queerphobia (b) LMB for Islamophobia

(the longer the bar, the higher the bias)

Natural Language Processing (NLP)

Natural language processing (NLP)

- **Natural language processing (NLP)** (Tsujii, 2011)
  - Computational methods for understanding and generating text (or speech)
  - Subfield of AI and one part of computational linguistics
  - Applications in data science and human-AI interaction

- **Computational linguistics**
  - Intersection of computer science and linguistics
  - Models to explain linguistic phenomena, based on knowledge and statistics
  - Methods for tackling language tasks automatically

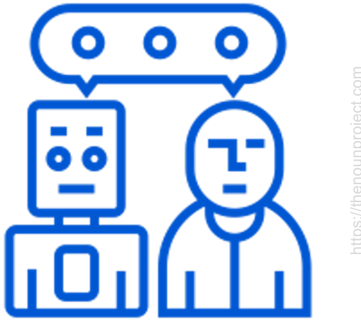

- **NLP tasks**
  - NLP deals with inference tasks for analysis and synthesis
  - Analysis. Inference of new information from given text
    Also referred to as *natural language understanding (NLU)*
  - Synthesis. Inference of new text from given information
    Also referred to as *natural language generation (NLG)*

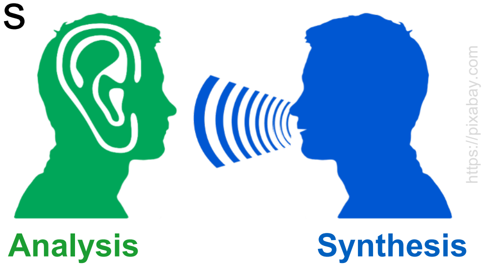

NLP examples: Traditional approach

- **Example: Argument mining**
  - Identifying and classifying arguments in natural language text

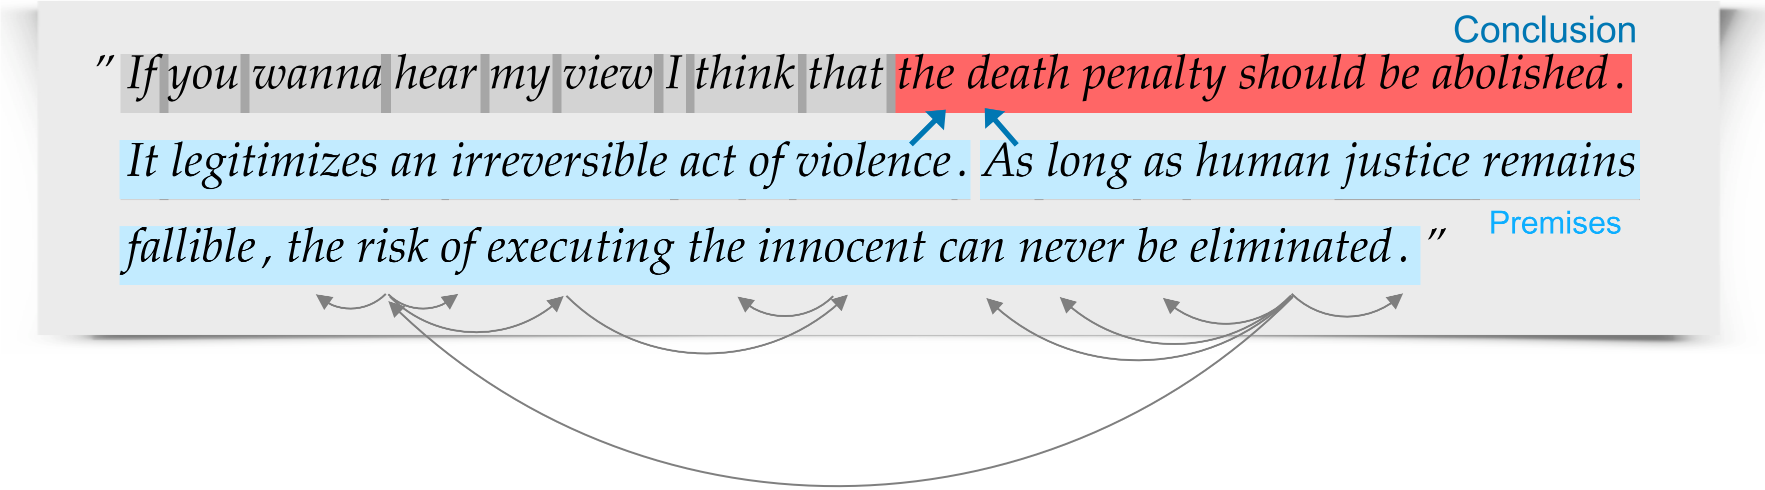

- **Traditional NLP approach**
  - Pipeline with several methods for specific tasks
  - Methods use hand-written rules and/or learned statistical patterns

NLP examples: Modern approach

- **Example: Language modeling**
  - Extending a given text word by word until a suitable ending is reached.

H In one short sentence: What is natural language processing?

Natural Language Processing (NLP) is a field of computer science and artificial intelligence that deals with the interaction between computers and humans through natural language.

- **Modern NLP approach**
  - Pretrain general language model on huge amounts of text examples
  - Fine-tune model to answer prompts aligned with human preferences

NLP and machine learning

- **Machine learning**
  - A method learns from experience wrt. a task and a performance measure, if its performance on the task increases with the experience. (Mitchell, 1997)
  - The goal is to learn a model *y* that approximates an unknown function γ

γ. A human expert on arguments

y. An NLP method for argument mining

- **Main types of machine learning**
  - Supervised. Derive *y* from data labeled with correct outputs; *y* can then predict output for other inputs
  - Unsupervised. Derive *y* from unlabeled data only; y describes data organization and association

Labeled data. Texts with known arguments

Unlabeled data. Argumentative texts

- **Machine learning in NLP**
  - NLP methods often use machine learning to infer outputs
  - Vice versa, the output of NLP may serve as input to machine learning

NLP techniques and tasks

- **Common NLP techniques**
  - Rule-based. Inference based on manually encoded expert knowledge
    Knowledge includes decision rules, lexicons, regular expressions, grammars, ...
  - Statistical. Inference based on statistical patterns in defined text features
    Features are encoded manually or semi-automatically.
  - Neural. Inference based on statistical patterns in self-learned functions
    Functions of arbitrary complexity may be approximated.

- **Common NLP tasks**
  - Similarity measures. The similarity of two instances is quantified.
  - Clustering. A set of instances is grouped into not-predefined classes.
  - Classification. Each instance is assigned a class label.
  - Regression. Each instance is assigned a numeric value.
  - Sequence labeling. A sequence of (interdependent) instances is classified.
  - Language modeling. A given input text is extended token by token.

Development and evaluation

- **Development**
  - NLP methods tackle language-related analysis and synthesis tasks.
  - To this end, they operationalize expert rules and/or statistical patterns.
  - Rules and patterns are derived from analyses of training data.

- **Evaluation**
  - The output of NLP methods is rarely free of errors due to the ambiguity of language.
  - Thus, they are evaluated empirically on test data.
  - The *effectiveness* of the methods is quantified in terms of metrics, such as accuracy.

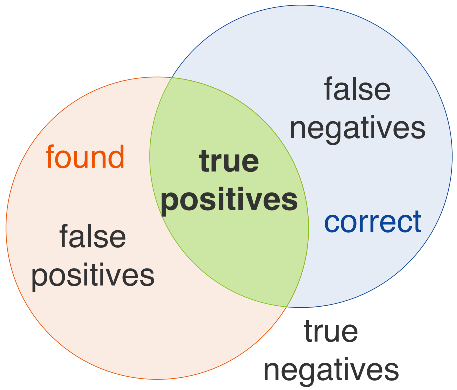

- **Comparison**
  - In general, it is unclear how good a measured effectiveness in a given task is.
  - New methods are thus compared to other methods, so called *baselines*.

Computational Sociolinguistics (CSL)

An interdisciplinary research area

- **Two views of computational sociolinguistics (CSL)**
  - The intersection of computational linguistics and sociolinguistics
  - Computational social science on language data

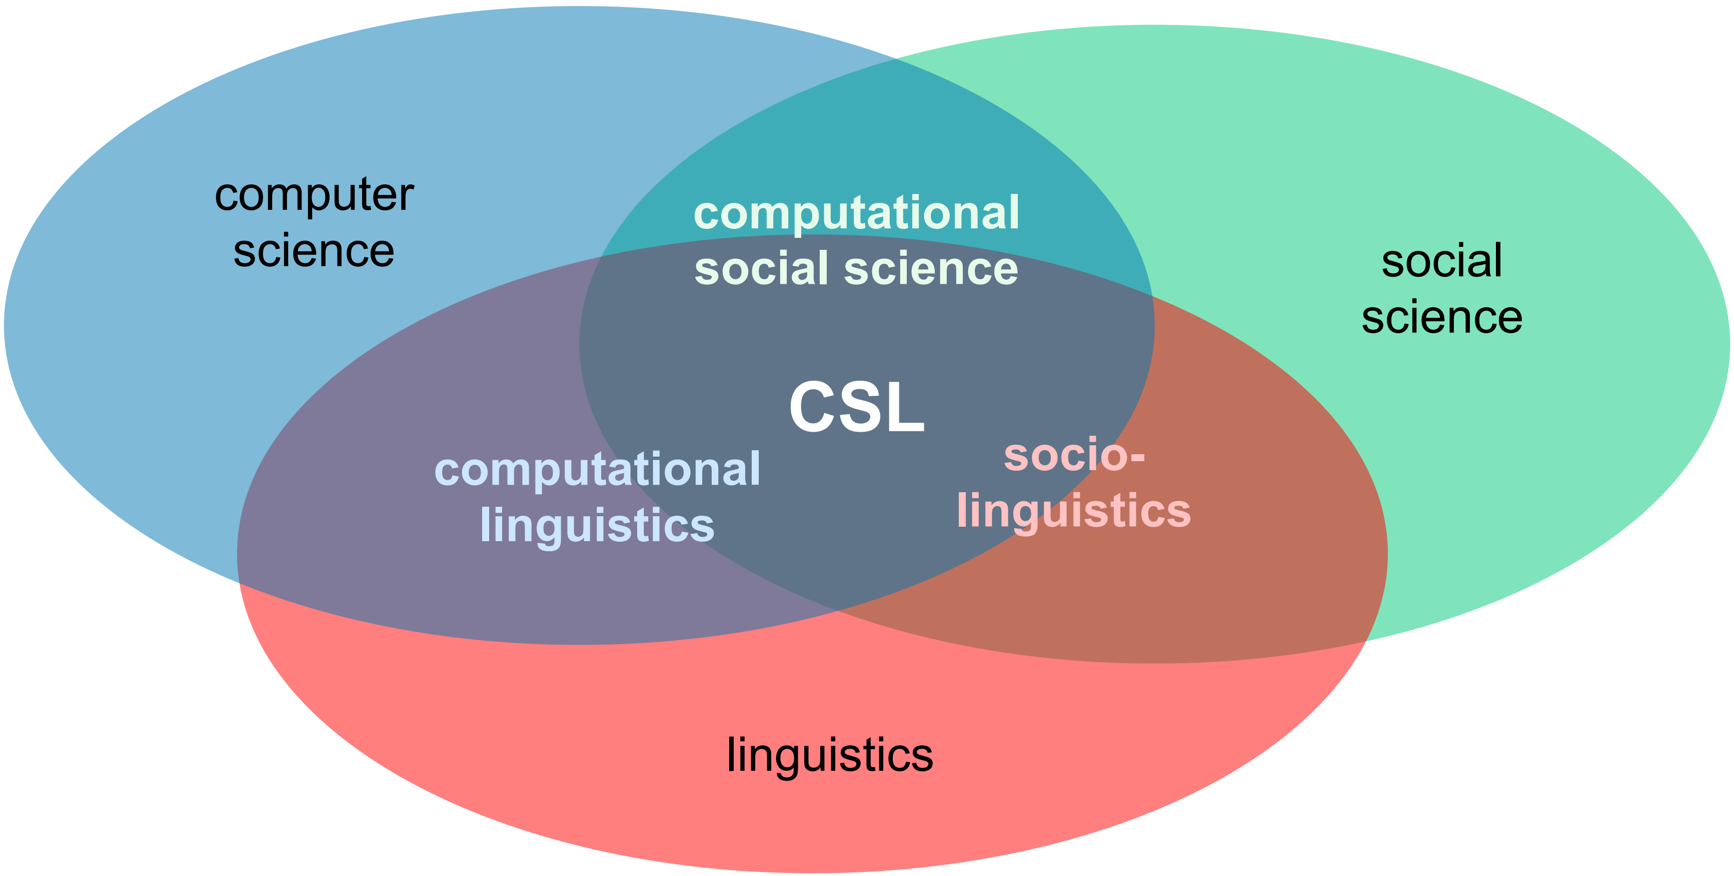

Computational linguistics based on Tsujii (2011)

- **Computational linguistics (CL)**
  - Intersection of computer science and linguistics
  - Models to explain linguistic phenomena, based on knowledge and statistics
  - Methods for tackling language tasks automatically

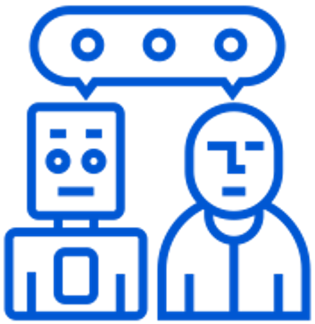

https://thenounproject.com

- **Language as an empirical phenomenon**
  - Language and its use follow common statistical patterns
  - Syntactic, semantic, and pragmatic aspects of language interact
  - Analyzing language requires real-world language samples

- **Goals of research**
  - Creativity. Novelty of developed models and methods
  - Accuracy. Effectiveness in tackling tasks
  - Empirical research is often seen as stronger than theory

Sociolinguistics based on Nguyen et al. (2016)

- **Sociolinguistics (SL)**
  - Studies the mutual interaction of society and language
  - Relations between social variables and language use
  - Language variation across social groups, social contexts, and communicative situations

https://thenounproject.com

- **Language as a social phenomenon**
  - Social identity of speakers and listeners inherently connected to language use
  - People can choose how to use language to achieve their goals
  - Analyzing language often requires to consider the people

- **Goals of research**
  - Validity. Extent to which research design isolates an issue to be studied
  - Reliability. Reproducibility of a result
  - Empirical research is seen as a means to support theory

Computational social science

- **Computational social science (CSS)**
  - Studies questions from the social science through empirical data analysis
  - Insights into social phenomena and dynamics (primary)
  - Technologies to support social context (secondary)

- **Data from social contexts**
  - Sociocultural key indicators
  - Social network structures
  - Online activities
  - News and social media texts

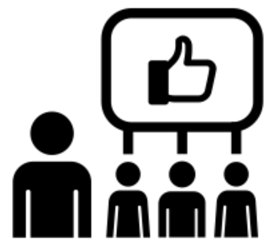

- **Methods from computer science**
  - Statistical correlation analyses
  - Data mining
  - Natural language processing

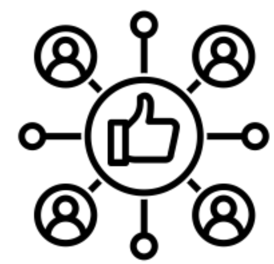

Computational sociolinguistics based on Nguyen et al. (2016)

- **Computational sociolinguistics (CSL)**
  - Studies relations between language and society computationally based on data
  - Questions emerging from theory in sociolinguistics
  - Methods from computational linguistics

- **NLP in the context of CSL**
  - Data. Natural language texts, along with sociocultural meta-information
  - Methods. Primarily analysis (classification, regression, clustering, …), but also text generation may be involved
  - Applications. Tools with social dimensions (chatbots, writing support, …)

- **Mutual impact of involved fields**
  - SL → CL. Build more robust and well-grounded computational methods
  - CL → SL. Refine theoretical models, better understand social dynamics

CSL topics in this seminar

NLP and CSL in this seminar

- **General frame**
  - Basics of NLP for computational sociolinguistics
  - State-of-the-art NLP research in this area
  - Required basics of NLP to be acquired rather than taught

- **Covered topics**
  - We take a broad view on CSL here
  - 12 topics related to research of NLP group
  - Overview is given in the following

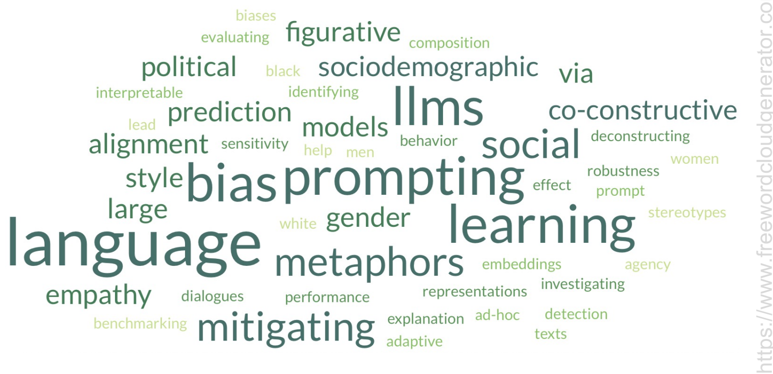

- **Concept behind**
  - Each seminar participant will be assigned one topic
  - Per topic, one paper is given as the basis for talk and article
  - Further literature (basics, related work, …) should be added where useful
  - You have to submit topic preferences, we then assign topics (details below)

Seminar topics by language type and advisor

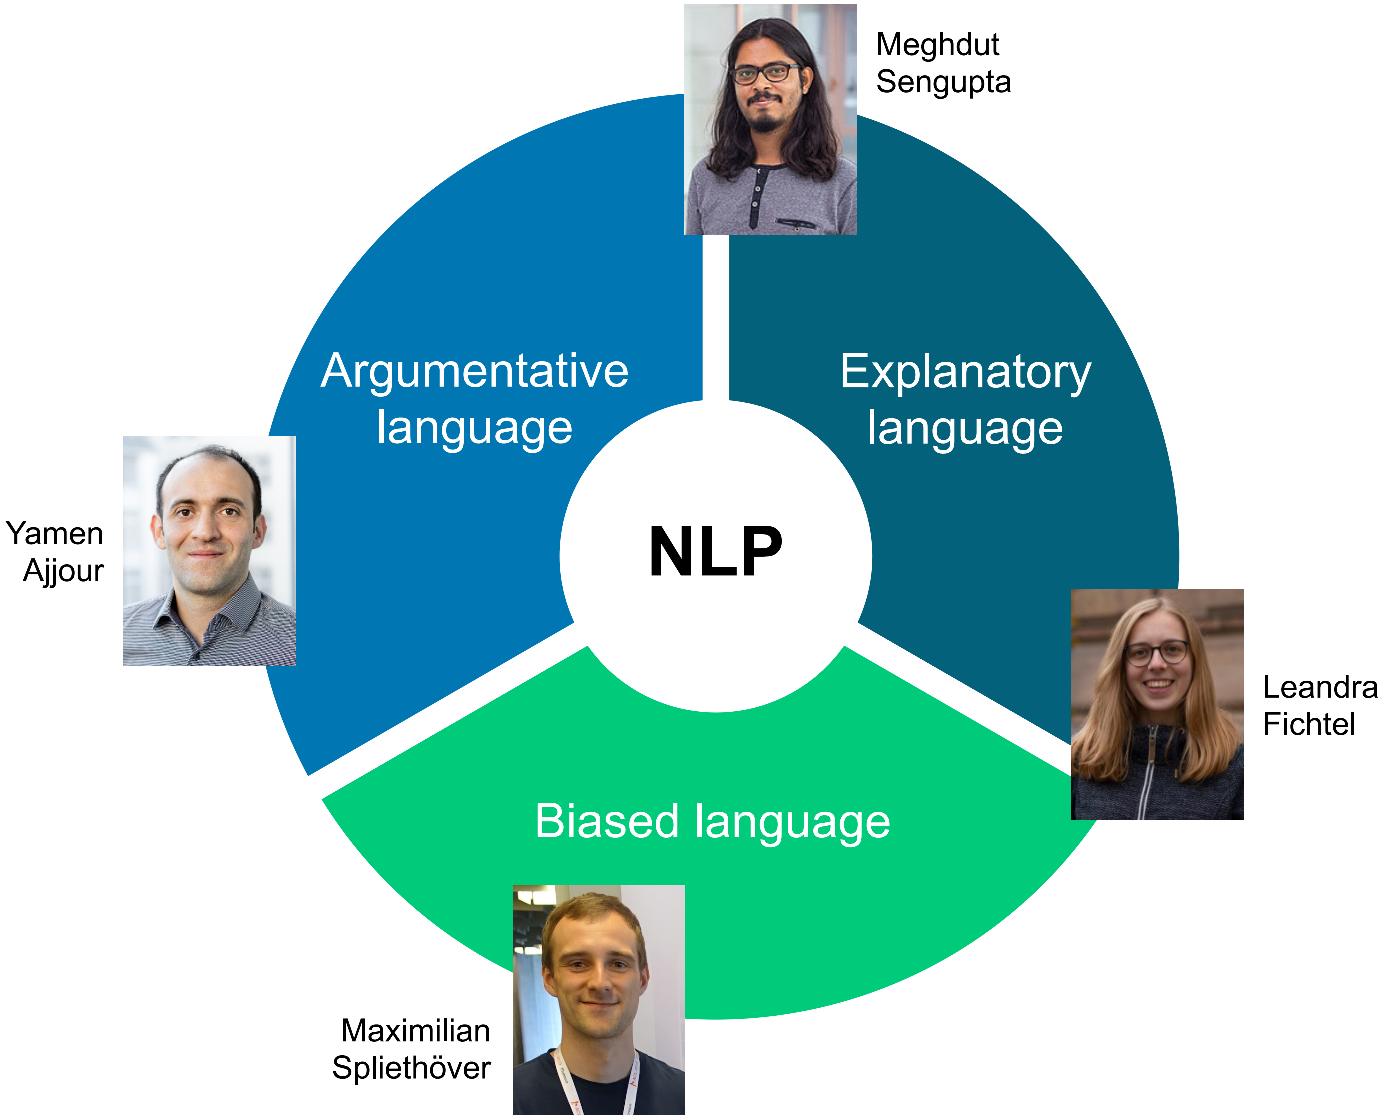

Types of language covered in the seminar

- **Biased language**
  - Language reflecting positive/negative views of social groups
  - Language reflecting stereotypes about the groups
  - Values and political views expressed through language

- **Explanatory language**
  - Explanations given to describe or justify phenomena
  - Linguistic devices used in explanations
  - Explanatory dialogues between two or more people

- **Argumentative language**
  - Online discussions on controversial issues
  - Claims and arguments that express viewpoints
  - Scientific articles that present new knowledge

Topics for the seminar talks and articles

- **Max**
  - X1. Identifying Social bias in Texts
  - X2. Evaluating Stereotypes in LLMs
  - X3. Mitigating Gender Bias through Alignment

- **Leandra**
  - L1. Sociodemographic Prompting
  - L2. Learning Style Representations
  - L3. Co-constructive LLMs

- **Meghdut**
  - M1. Multimodal Figurative Language
  - M2. Metaphor Interpretation
  - M3. Political Persuasion via Metaphors

- **Yamen**
  - Y1. Fallacious Argumentation in Science
  - Y2. Perspectivist Method For Empathy Prediction
  - Y3. Task-specific Fine-tuning

Topics advised by Max

X1. Identifying Social bias in Texts

- Maximilian Spliethöver, Tim Knebler, Fabian Fumagalli, Maximilian Muschalik, Barbara Hammer, Eyke Hüllermeier, Henning Wachsmuth. 2025. Adaptive Prompting: Ad-hoc Prompt Composition for Social Bias Detection.
  https://aclanthology.org/2025.naacl-long.122/

X2. Evaluating Stereotypes in LLMs

- Yixin Wan, and Kai-Wei Chang. 2025. White Men Lead, Black Women Help? Benchmarking and Mitigating Language Agency Social Biases in LLMs.
  https://aclanthology.org/2025.acl-long.445/

X3. Mitigating Gender Bias through Alignment

- Tao Zhang, Ziqian Zeng, YuxiangXiao YuxiangXiao, Huiping Zhuang, Cen Chen, James R. Foulds, and Shimei Pan. 2025. GenderAlign: An Alignment Dataset for Mitigating Gender Bias in Large Language Models.
  https://aclanthology.org/2025.acl-long.553/

Topics advised by Leandra

L1. Sociodemographic Prompting

- Tilman Beck, Hendrik Schuff, Anne Lauscher, Iryna Gurevych. 2024. Sensitivity, Performance, Robustness: Deconstructing the Effect of Sociodemographic Prompting. https://aclanthology.org/2024.eacl-long.159

L2. Learning Style Representations

- Ajay Patel, Delip Rao, Ansh Kothary, Kathleen McKeown, Chris Callison-Burch. 2023. Learning Interpretable Style Embeddings via Prompting LLMs. https://aclanthology.org/2023.findings-emnlp.1020

L3. Co-constructive LLMs

- Leandra Fichtel, Maximilian Spliethöver, Eyke Hüllermeier, Patricia Jimenez, Nils Klowait, Stefan Kopp, Axel-Cyrille Ngonga Ngomo, Amelie Robrecht, Ingrid Scharlau, Lutz Terfloth, Anna-Lisa Vollmer, Henning Wachsmuth. 2025. Investigating Co-Constructive Behavior of Large Language Models in Explanation Dialogues. https://aclanthology.org/2025.sigdial-1.1/

Topics advised by Meghdut

M1. Multimodal Figurative Language

- Arkadiy Saakyan, Shreyas Kulkarni, Tuhin Chakrabarty, and Smaranda Muresan. 2024. V-FLUTE: Visual Figurative Language Understanding with Textual Explanations. https://arxiv.org/pdf/2405.01474

M2. Metaphor Interpretation

- Meghdut Sengupta, Milad Alshomary, Ingrid Scharlau, and Henning Wachsmuth. 2023. Modeling Highlighting of Metaphors in Multitask Contrastive Learning Paradigms. https://aclanthology.org/2023.findings-emnlp.308

M3. Political Persuasion via Metaphors

- Meghdut Sengupta, Roxanne El Baff, Milad Alshomary, and Henning Wachsmuth. 2024. Analyzing the Use of Metaphors in News Editorials for Political Framing. https://aclanthology.org/2024.naacl-long.199

Topics advised by Yamen

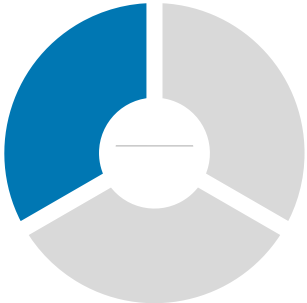

Y1. Fallacious Argumentation in Science

- Max Glockner, Yufang Hou, Perslav Nakov, Iryna Gurevych. 2025. Grounding Fallacies Misrepresenting Scientific Publications in Evidence.
  https://aclanthology.org/2025.naacl-long.491

Y2. Perspectivist Method For Empathy Prediction

- Francine Chen, Scott Carter, Tatiana Lau, Nayeli Bravo, Sumanta Bhattacharrya, Kate Sieck and Charlene Wu. 2025. Empathy Prediction from Diverse Perspectives.
  https://aclanthology.org/2025.acl-long.439

Y3. Task-specific Fine-tuning

- David Schulte, Felix Hamborg, Alan Akbik. 2024. Less is More: Parameter-Efficient Selection of Intermediate Tasks for Transfer Learning.
  https://aclanthology.org/2024.emnlp-main.529

Conclusion

Conclusion

- **NLP and computational sociolinguistics (CSL)**
  - NLP studies methods for analyzing and synthesizing language
  - CSL focuses on social aspects interacting with language
  - Goal is to understand relations between language and society

https://thenounproject.com

- **This seminar**
  - State-of-the-art NLP research on CSL
  - Focus on biased, explanatory, and argumentative language
  - Close connection to research in the NLP Group

https://www. svgsilh.com

- **Your task**
  - Inform yourself about the topics and papers in this presentation
  - Choose 3 topics with preferences
  - Until Monday, October 20, 23:59 UTC+2. Send me preferences

    Both direct e-mail and Stud.IP message are fine.

https://pixabay.com

Conclusion: Your e-mail and subsequent process

- **Your e-mail**
  - Recipient. h.wachsmuth@ai.uni-hannover.de
  - Subject. ”[nlp] Topic preferences“
  - Content. Your name, matriculation number, and 3 topic preferences
    Notice: Less than 3 does not increase chances to get one of them
  - Example. On the right, you see how the content of your e-mail could look like

Name:
Jane Doe

Matriculation number:
12345678

Topic preferences:
1) L2. Learning Style Representations
2) M1. Multimodal Figurative Language
3) L1. Sociodemographic Prompting

- **Subsequent process**
  - We will assign topics based on preferences, special reasons, and randomly
  - If you don‘t send your e-mail in time, you will *not* be assigned any topic
  - The final schedule will be decided based on the topic assignment
  - Assignment and schedule will be announced next week

References

- **Holtermann et al. (2022).** Carolin Holtermann, Anne Lauscher, and Simone Ponzetto. Fair and Argumentative Language Modeling for Computational Argumentation. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 7841–7861, 2022.
- **Kosinski et al. (2013).** Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable from digital records of human behavior. Proceedings of the National Academy of Sciences, 110(15):5802–5805, 2013.
- **Nguyen et al. (2016).** Dong Nguyen, A. Seza Doğruöz, Carolyn P. Rosé, Franciska de Jong. Computational Sociolinguistics: A Survey. Computational Linguistics 42(3): 537–593, 2016.
- **Tsujii (2011).** Jun'ichi Tsujii. Computational Linguistics and Natural Language Processing. In Proceedings of the 12th International Conference on Computational linguistics and Intelligent Text Processing - Volume Part I, pages 52–67, 2011.
- **Voigt et al. (2017).** Rob Voigt, Nicholas P. Camp, Vinodkumar Prabhakaran, William L. Hamilton, Rebecca C. Hetey, Camilla M. Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L. Eberhardt. Language from police body camera footage shows racial disparities in officer respect. Proceedings of the National Academy of Sciences, 114(25):6521–6526, 2017.