Seminar Natural Language Generation (NLG) — Part 2

Overview of Seminar Topics

Henning Wachsmuth

https://ai.uni-hannover.de

Natural Language Processing

Leibniz Universität Hannover

Assignment of seminar topics

- **This talk**
  - Overview of all 15 possible seminar topics, presented by the respective advisors
  - For each topic, two articles are given that provide the basis of the topic

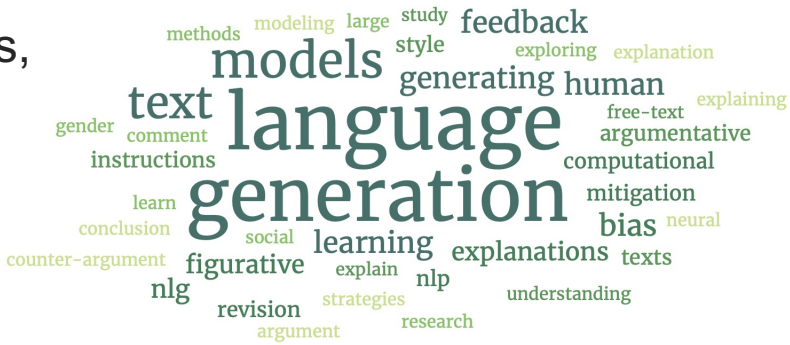

- **Concept behind**
  - Each seminar participant will be assigned one topic
  - The two articles (+ *at least* one further relevant article to be found by you) should be discussed in the talk and in the article
  - You can choose topic preferences, we then assign topics

- **Your task**
  - Inform yourself about the topics and articles in this presentation
  - Choose 3 topics with preferences
  - Until Saturday, April 29, 23:59 GTM+2. Send e-mail with preferences

Choosing preferences: eMail and subsequent process

- **Your eMail**
  - Recipient. h.wachsmuth@ai.uni-hannover.de
  - Subject. ”[nlg] Topic preferences“
  - Content. Your name and matriculation number + 3 topic preferences
  - Example. On the right, you see how the content of your eMail could look like

Name:
Timon Ziegenbein

Matriculation number:
1234567

Topic preferences:
1) A4. Text Revision Generation
2) E1. Feedback Comment Generation
3) A1. Conclusion Generation

- **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
    You can get a rough idea of the schedule from the ordering on the next slides.
  - Topic assignment and schedule will be announced until the next session

Overview of topics

Seminar topics organized into four areas

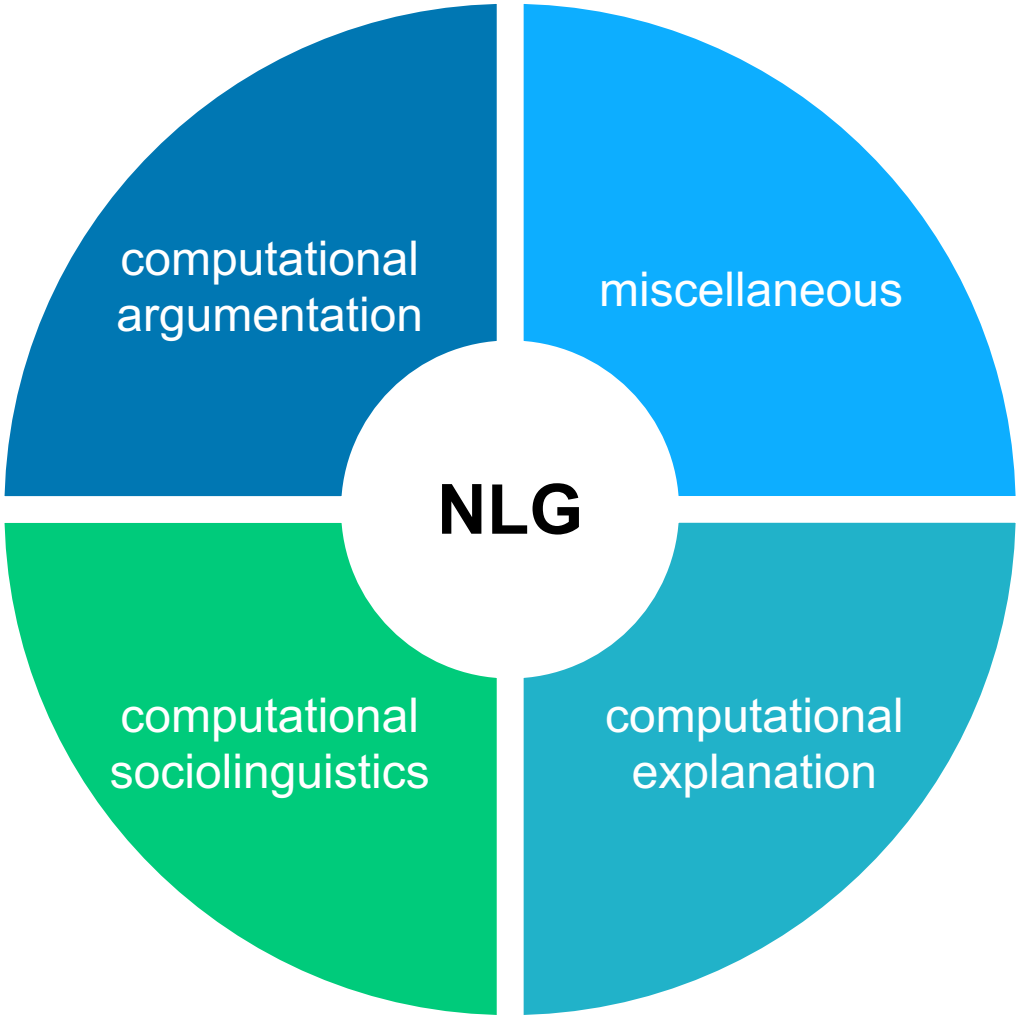

Topics for the seminar talks and articles (1 out of 2)

- **Computational argumentation**
  - A1. Conclusion Generation — Milad
  - A2. Audience-Aware Argument Generation — Milad
  - A3. Counter-Argument Generation — Milad
  - A4. Text Revision Generation — Gabriella

- **Computational sociolinguistics**
  - S1. Social Bias Evaluation of Large Language Models — Max
  - S2. Interpretable Language Modeling with Minimal Editing — Timon
  - S3. Gender Bias Mitigation — Maja

More on next slide…

Topics for the seminar talks and articles (2 out of 2)

- **Computational explanation**

E1.	Generating Free-Text Explanations	Maja
E2.	Feedback Comment Generation	Maja
E3.	Shapley Values for Explanation Generation	Meghdut
E4.	Explaining Figurative Language	Meghdut

- **Miscellaneous**

M1.	Metaphor Generation	Meghdut
M2.	Selection Strategies for Few-Shot Learning	Yamen
M3.	Human-aligned Text Summarization	Timon
M4.	Teaching Language Models to Follow Instructions	Timon

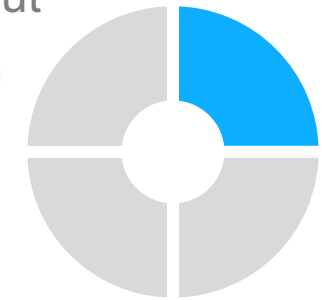

Topics and Literature

Topics supervised by Milad

A1. Conclusion Generation

- Shahbaz Syed, Khalid Al Khatib, Milad Alshomary, Henning Wachsmuth, and Martin Potthast (2021). Generating Informative Conclusions for Argumentative Texts. https://aclanthology.org/2021.findings-acl.306/
- Philipp Heinisch, Anette Frank, Juri Opitz, and Philipp Cimiano (2022). Strategies for Framing Argumentative Conclusion Generation. https://aclanthology.org/2022.inlg-main.20/

A2. Audience-Aware Argument Generation

- Milad Alshomary, Wei-Fan Chen, Timon Gurcke, Henning Wachsmuth (2021). Belief-based Generation of Argumentative Claims. https://aclanthology.org/2021.eacl-main.17/
- Milad Alshomary, Roxanne El Baff, Timon Gurcke, Henning Wachsmuth (2022). The Moral Debater: A Study on the Computational Generation of Morally Framed Arguments. https://aclanthology.org/2022.acl-long.601/

A3. Counter-Argument Generation

- Milad Alshomary, Shahbaz Syed, Arkajit Dhar, Martin Potthast, and Henning Wachsmuth (2021). Argument Undermining: Counter-Argument Generation by Attacking Weak Premises. https://aclanthology.org/2021.findings-acl.159/
- Xinyu Hua and Lu Wang (2019). Sentence-level content planning and style specification for neural text generation. https://aclanthology.org/D19-1055/

Topics supervised by Maja

S3. Gender Bias Mitigation

- Xinyao Ma, Maarten Sap, Hannah Rashkin, Yejin Choi (2020). PowerTransformer: Unsupervised Controllable Revision for Biased Language Correction. https://aclanthology.org/2020.emnlp-main.602/
- Maja Stahl, Maximilian Spliethöver, Henning Wachsmuth (2022). To Prefer or to Choose? Generating Agency and Power Counterfactuals Jointly for Gender Bias Mitigation. https://aclanthology.org/2022.nlpcss-1.6/

E1. Generating Free-Text Explanations

- Sarah Wiegreffe, Jack Hessel, Swabha Swayamdipta, Mark Riedl, Yejin Choi (2022). Reframing Human-AI Collaboration for Generating Free-Text Explanations. https://aclanthology.org/2022.naacl-main.47/
- Swarnadeep Saha, Peter Hase, Nazneen Rajani, Mohit Bansal (2022). Are Hard Examples also Harder to Explain? A Study with Human and Model-Generated Explanations. https://aclanthology.org/2022.emnlp-main.137/

E2. Feedback Comment Generation

- Ryo Nagata, Kentaro Inui, Shinaichiro Ishikawa (2020). Creating Corpora for Research in Feedback Comment Generation. https://aclanthology.org/2020.lrec-1.42/
- Kazuaki Hanawa, Ryo Nagata, Kentaro Inui (2021). Exploring Methods for Generating Feedback Comments for Writing Learning. https://aclanthology.org/2021.emnlp-main.766/

Topics supervised by Meghdut

E3. Shapley Values for Explanation Generation

- Edoardo Mosca, Defne Demirtürk, Luca Mülln, Fabio Raffagnato, Georg Groh (2022). GrammarSHAP: An Efficient Model-Agnostic and Structure-Aware NLP Explainer. https://aclanthology.org/2022.lnls-1.2/
- Edoardo Mosca, Ferenc Szigeti, Stella Tragianni, Daniel Gallagher, Georg Groh (2022). SHAP-Based Explanation Methods: A Review for NLP Interpretability. https://aclanthology.org/2022.coling-1.406/

E4. Explaining Figurative Language

- Mohadeseh Bastan, Yash Kumar Lal (2022). SBU Figures It Out: Models Explain Figurative Language. https://aclanthology.org/2022.flp-1.20/
- Tuhin Chakrabarty, Arkadiy Saakyan, Debanjan Ghosh, Smaranda Muresan (2022). Figurative Language Understanding through Textual Explanations. https://aclanthology.org/2022.emnlp-main.481/

M1. Metaphor Generation

- Kevin Stowe, Nils Beck, Iryna Gurevych (2021). Exploring Metaphoric Paraphrase Generation. https://aclanthology.org/2021.conll-1.26/
- Kevin Stowe, Prasetya Utama, Iryna Gurevych (2022). IMPLI: Investigating NLI Modelsâ□□ Performance on Figurative Language. https://aclanthology.org/2022.acl-long.369/

Topics supervised by Timon

S2. Interpretable Language Modeling with Minimal Editing

- Alexis Ross, Ana Marasovic, Matthew E. Peters (2021). Explaining NLP Models via Minimal Contrastive Editing (MICE). https://aclanthology.org/2021.findings-acl.336/
- Ewoenam Kwaku Tokpo, Toon Calders (2022). Text Style Transfer for Bias Mitigation using Masked Language Modeling. https://aclanthology.org/2022.naacl-srw.21/

M3. Human Aligned Text Summarization

- Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul Christiano, Geoffrey Irving (2020). Fine-Tuning Language Models from Human Preferences. https://arxiv.org/abs/1909.08593
- Nisan Stiennonr, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, Paul Christiano (2022). Learning to summarize from human feedback. https://arxiv.org/abs/2009.01325

M4. Teaching Language Models to Follow Instructions

- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin et al. (2022). Training language models to follow instructions with human feedback. https://arxiv.org/abs/2203.02155
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, Hannaneh Hajishirzi. SELF-INSTRUCT: Aligning Language Model with Self Generated Instructions. https://arxiv.org/abs/2212.10560

Topics supervised by Gabriella, Max, and Yamen

A4. Text Revision Generation

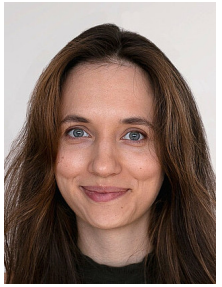

- Gabriella Skitalinskaya, Maximilian Spliethöver, Henning Wachsmuth (2022). Claim Optimization in Computational Argumentation. https://arxiv.org/abs/2212.08913
- Wanyu Du, Vipul Raheja, Dhruv Kumar, Zae Myung Kim, Melissa Lopez, Dongyeop Kang (2022). Understanding Iterative Revision from Human-Written Text. https://aclanthology.org/2022.acl-long.250/

S1. Social Bias Evaluation of Large Language Models

- Debora Nozza, Federico Bianchi, Dirk Hovy (2022). Pipelines for Social Bias Testing of Large Language Models. https://aclanthology.org/2022.bigscience-1.6/
- Rochelle Choenni, Ekaterina Shutova, Robert von Rooij (2021). Stepmothers are mean and academics are pretentious: What do pretrained language models learn about you? https://aclanthology.org/2021.emnlp-main.111/

M2. Selection Strategies for Few-Shot Learning

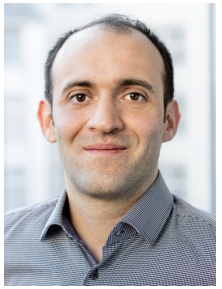

- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, Pontus Stenetorp (2022). Fantastically Ordered Prompts and Where to Find Them: Overcoming Few-Shot Prompt Order Sensitivity. https://aclanthology.org/2022.acl-long.556/
- Ohad Rubin, Jonathan Herzig, Jonathan Berant. Learning To Retrieve Prompts for In-Context Learning. https://aclanthology.org/2022.naacl-main.191/

Sum-up

Conclusion

- **Seminar topics**
  - 14 candidate topics related to four research areas
  - Each of you will be assigned one of these topics
  - Given + further literature form the basis of talk and article

- **Topic assignment**
  - You choose topic preferences, we assign topics
  - Inform yourself about the topics of the given literature
  - Send me your topic preferences by Saturday this week!

- **Next up**
  - Topic assignment will be presented until the next session
  - Talk preparation starts then
  - Basics of scientific presentation in the weak afterwards