

Conference Report: KONVENS 2024*

Roman Klinger

2024-10-12

KONVENS

In September 2024, I participated in the [KONVENS](#) – the “Konferenz zur Verarbeitung natürlicher Sprache” (Conference on Natural Language Processing) in Vienna.

KONVENS is the computational linguistics and natural language processing conference in the German speaking countries. Various countries have such more local CL/NLP conferences, complementing the large and global conferences by the [ACL](#), the [COLING](#), or [LREC](#), which have different foci, but are always very international. (there are also many other venues, like machine learning, language models, and AI focused events, but given that KONVENS is CL/NLP, I only contrast it to this field here).

Other examples for established more regional conferences are the [NoDaLiDa](#) (Nordic Conference on Computational Linguistics), [Clic-it](#) (Italy), or the [CLIN](#) (Netherlands).

You may ask: Why would I go to such a regional conference? (and by the way, all of these conferences are international these days, and the language spoken there is English, but the focus is a bit more regional)

I think there are a couple of reasons:

- There are papers that fit better to KONVENS than to larger, global venues. In NLP, we mostly publish at conferences. Regional conferences also publish proceedings as larger venues do, which typically also go into the [ACL Anthology](#) the main paper repository in the field (and all open access). The reputation of these regional conferences is lower than [EMNLP](#) or [ACL](#), but, as with focus workshops, there are papers which find a more interested audience here. For example, if you work on the German language, it's more likely that you find German speaking people at KONVENS.
- You don't need travel so far. Sure, UAE or Miami might be nice for some, but for others, traveling there is not an option. Be it visa issues or not feeling comfortable with the legal situation in a place ([some readers might find this a euphemism, it can be pretty bad for some people in some countries](#)), or they are hesitant to travel far, by plane.

*Published as blog post at <https://www.romanklinger.de/blog/2024-10-12-konvens/>

- Sometimes there is no funding available to go to a distant conference. With KONVENS and other regional venues, also papers that have been written based on a, for instance, Master's thesis, where the main author might not have an affiliation at the moment the conference takes place, could be published.
- Networking. It's so much easier to enter a new field in smaller conferences than in bigger ones, and you meet people who are typically geographically closer to you. This makes it easier to collaborate, based on discussions that may take place at the conference. Networking is the main reason I participate in these conferences.

KONVENS 2024

KONVENS 2024 took place in Vienna, and has been organized not only by the German Society for Computational Linguistics ([GSCL](#)) but jointly with the Austrian Research Institute for Artificial Intelligence ([OFAI](#)) and the Austrian Society for Artificial Intelligence ([ASAI](#)). The main local organizer has been the [University of Vienna](#).

The conference received 57 submissions and accepted 39 papers. During the conference, there were 30 poster presentations and 9 oral presentations. Most papers came from Germany (70), Austria (20), and Switzerland (14). Authors from other countries contributed 7 more papers. In addition, there were three invited talks ([Leonie Weissweiler](#) who is a postdoc at UT Austin; [Sebastian Schuster](#) from UC London; [Jana Diesner](#) from TU Munich). The conference was complemented by a set of (partially as large as the main conference) workshops: [GermEval Shared Task 2: Statement Segmentation in German Easy Language \(StaGE\)](#), [Workshop on Linguistic Insights from and for Multimodal Language Processing \(LIMO\)](#), and [Workshop on Computational Linguistics for Political Text Analysis \(CPSS\)](#), and [GERMS-DETECT Sexism Detection in German Online News Fora \(GERMS-DETECT\)](#).

My Favorite Contributions

All of the invited talks were awesome. I'd like to point out the presentation by Sebastian Schuster (because I found it most relevant for my own work), who explained limitations of large language models based on inference tasks that are easy for humans and difficult for machines. The main paper his talk was based on is [Kim and Schuster \(2023\)](#), which also won a best paper award at the recent ACL in Toronto 2023. The task is to follow a description how entities are moved from one box to another, and the model needs to say in which box which entity is.

The whole [proceedings of KONVENS](#) are available in the ACL Anthology.

Under the assumption that you might be reading this because you have similar research interests as I do, I'd like to point out papers, that I personally found particularly interesting and relevant (for my work).

- [Hellwig et al. \(2024\)](#) report on a German restaurant review dataset, annotated for aspect-based sentiment analysis. There are a couple of German sentiment corpora (for instance our own corpora USAGE [Klinger and Cimiano \(2014\)](#) and SCARE [Sänger et al. \(2016\)](#)), but in contrast to English, there is not a lot, and

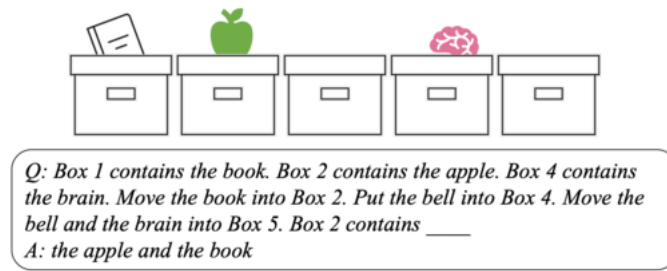


Figure 1: Entity Tracking Task, presented by [Kim and Schuster \(2023\)](#) as a keynote

the restaurant domain did, as far as I know, not receive any attention yet. The resource consists of more than 3000 manually annotated reviews.

- Language models are often used for text classification now, and offer themselves as a training data efficient method, via prompting. [Kluge and Kähler \(2024\)](#) present experiments on indexing medical book titles via prompting. The authors work German National Library, so I assume that this paper reports not only on a purely academic work, but on something that has practical relevance for their direct environment. Subject indexing is an interesting and challenging task, sometimes considered extreme classification, because you need to decide for many labels which are fitting. While the paper does not provide statistics on the inventory of possible labels used here, I assume that the set is large.
- [Petersen-Frey and Biemann \(2024\)](#) present a method on quotation and attribution – the task is to detect speech in written text and attribute it to the speaker (“Roman said ‘this is true’”). We worked on speaker and quotation identification a while ago ([Scheible, Klinger, and Padó \(2016\)](#)) and my former collaborators continued to contribute to the topic (e.g., [Papay and Padó \(2020\)](#)). Petersen-Frey and Biemann (2024) approach the task in a structured prediction framework.
- While a lot of efforts go into mitigating gender bias in representations (see [Sun et al. \(2019\)](#) for a survey), [Gross et al. \(2024\)](#) take a different approach: they induce gender bias in language models to then be able to study the effects in a controlled environment.
- With the increasing popularity of populist parties, some research goes into analysing the language of populists in contrast to other political parties. While we know that populists use particular rhetoric strategies (to convince people without having actually good arguments) more frequently than other parties, there is not too much work on the language complexity. [Zanotto, Frassinelli, and Butt \(2024\)](#) investigate the hypothesis that populists use simpler language (for instance to have a larger outreach). They do, however, not find any significant effects, but confirm the more frequent use of persuasion tactics.

Awards

I cannot write this blog post without mentioning that my Ph.D. student Enrica Troiano won the award of the GSCL for the best thesis in the years 2023. Her thesis is on bringing event analysis and emotion analysis together. In contrast to the various, sometimes a bit decontextualized papers we wrote, the thesis is really a nice aggregation

of the work, and worth reading (Troiano (2023)).



Figure 2: Enrica Troiano won the GSCL PhD award

Venue and Place

The conference took place in Vienna - a city I should have visited more often already. Now, from my new workplace in Bamberg, Vienna is reachable with a short train trip (4 hours from Nurnberg). Of course, I brought my bicycle, so I could commute from the hotel to the conference venue by bike. The bike infrastructure in Vienna is very good and protects cyclists well from cars, and it's much faster to get around than by public transportation. Unfortunately, there was a storm and rain warning, so towards the end of the conference, cycling around became a bit challenging. Actually, when I traveled back, I took one of the last three trains that made it to Germany, before the track was shut down for a couple of days. Read more about this storm [here](#).

The conference itself took place at the University of Vienna, in a pretty modern lecture hall. The poster sessions were right in the lobby, so no long commutes between places for various parts of the program.

The social event was a small walk through the vineyards and dinner in a beergarden. I prefer more vegetarian-friendly places and non-seated dinners at conferences, but the place was very nice. I also really liked to combination of a short walk with the dinner.



Figure 3: Social Event

Bibliography

- Gross, Stephanie, Brigitte Krenn, Craig Lincoln, and Lena Holzwarth. 2024. "Analysing Effects of Inducing Gender Bias in Language Models." In *Proceedings of the 20th Conference on Natural Language Processing (KONVENS 2024)*, edited by Pedro Henrique Luz de Araujo, Andreas Baumann, Dagmar Gromann, Brigitte Krenn, Benjamin Roth, and Michael Wiegand, 222–30. Vienna, Austria: Association for Computational Linguistics. <https://aclanthology.org/2024.konvens-main.24>.
- Hellwig, Nils Constantin, Jakob Fehle, Markus Bink, and Christian Wolff. 2024. "GER-restaurant: A German Dataset of Annotated Restaurant Reviews for Aspect-Based Sentiment Analysis." In *Proceedings of the 20th Conference on Natural Language Processing (KONVENS 2024)*, edited by Pedro Henrique Luz de Araujo, Andreas Baumann, Dagmar Gromann, Brigitte Krenn, Benjamin Roth, and Michael Wiegand, 123–33. Vienna, Austria: Association for Computational Linguistics. <https://aclanthology.org/2024.konvens-main.14>.
- Kim, Najoung, and Sebastian Schuster. 2023. "Entity Tracking in Language Models." In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, edited by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, 3835–55. Toronto, Canada: Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.213>.
- Klinger, Roman, and Philipp Cimiano. 2014. "The USAGE Review Corpus for Fine Grained Multi Lingual Opinion Analysis." In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, edited by Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Hrafn Loftsson, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk, and Stelios Piperidis, 2211–18. Reykjavik, Iceland: European Language Resources Association (ELRA). http://www.lrec-conf.org/proceedings/lrec2014/pdf/85_Paper.pdf.
- Kluge, Lisa, and Maximilian Kähler. 2024. "Few-Shot Prompting for Subject Indexing of German Medical Book Titles." In *Proceedings of the 20th Conference on Natural Language Processing (KONVENS 2024)*, edited by Pedro Henrique Luz de Araujo, Andreas Baumann, Dagmar Gromann, Brigitte Krenn, Benjamin Roth, and Michael Wiegand, 141–48. Vienna, Austria: Association for Computational Linguistics. <https://aclanthology.org/2024.konvens-main.16>.
- Papay, Sean, and Sebastian Padó. 2020. "RiQuA: A Corpus of Rich Quotation Annotation for English Literary Text." In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, edited by Nicoletta Calzolari, Frédéric Béchet,

- Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, et al., 835–41. Marseille, France: European Language Resources Association. <https://aclanthology.org/2020.lrec-1.104>.
- Petersen-Frey, Fynn, and Chris Biemann. 2024. “Fine-Grained Quotation Detection and Attribution in German News Articles.” In *Proceedings of the 20th Conference on Natural Language Processing (KONVENS 2024)*, edited by Pedro Henrique Luz de Araujo, Andreas Baumann, Dagmar Gromann, Brigitte Krenn, Benjamin Roth, and Michael Wiegand, 196–208. Vienna, Austria: Association for Computational Linguistics. <https://aclanthology.org/2024.konvens-main.22>.
- Sänger, Mario, Ulf Leser, Steffen Kemmerer, Peter Adolphs, and Roman Klinger. 2016. “SCARE — the Sentiment Corpus of App Reviews with Fine-Grained Annotations in German.” In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, edited by Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, et al., 1114–21. Portorož, Slovenia: European Language Resources Association (ELRA). <https://aclanthology.org/L16-1178>.
- Scheible, Christian, Roman Klinger, and Sebastian Padó. 2016. “Model Architectures for Quotation Detection.” In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, edited by Katrin Erk and Noah A. Smith, 1736–45. Berlin, Germany: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P16-1164>.
- Sun, Tony, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. “Mitigating Gender Bias in Natural Language Processing: Literature Review.” In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, edited by Anna Korhonen, David Traum, and Lluís Màrquez, 1630–40. Florence, Italy: Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1159>.
- Troiano, Enrica. 2023. “Where Are Emotions in Text? A Human-Based and Computational Investigation of Emotion Recognition and Generation.” PhD thesis, University of Stuttgart. <https://elib.uni-stuttgart.de/handle/11682/13671>.
- Zanotto, Sergio E., Diego Frassinelli, and Miriam Butt. 2024. “Language Complexity in Populist Rhetoric.” In *Proceedings of the 4th Workshop on Computational Linguistics for the Political and Social Sciences: Long and Short Papers*, edited by Christopher Klamm, Gabriella Lapesa, Simone Paolo Ponzetto, Ines Rehbein, and Indira Sen, 61–80. Vienna, Austria: Association for Computational Linguistics. <https://aclanthology.org/2024.cpss-1.5>.

[[Download this post as PDF](#)]