Personal Noun Detection for German

Carla Sökefeld
Universität Hamburg
carla.soekefeld@uni-hamburg.de

Melanie Andresen
University of Stuttgart
melanie.andresen@uni-stuttgart.de

Johanna Binnewitt
Bundesinstitut für Berufsbildung
Johanna.Binnewitt@bibb.de

Heike Zinsmeister
Universität Hamburg
heike.zinsmeister@uni-hamburg.de

Abstract

Common nouns denoting human beings such as *teacher* or *visitor*—henceforth **personal nouns**—play an important role in manifesting gender and gender stereotypes in texts, especially for languages with grammatical gender like German. Automatically detecting and extracting personal nouns can thus be of interest for a wide range of different tasks such as minimizing gender bias in language models and researching gender stereotypes or gender-fair language. However, personal noun detection is complicated by the morphological heterogeneity and ambiguity of personal and non-personal nouns, which restrict lexicon-based approaches. In this paper, we introduce the new task of personal noun detection and present a classifier that detects personal nouns in German, created by fine-tuning a BERT-based transformer model. Although some phenomena like ambiguity and metalinguistic uses are still problematic, the model is able to classify personal nouns with robust performance (f1-score: 0.94).

1 Motivation

Following Elmiger (2018), personal nouns are common nouns denoting humans such as kinship terms (*daughter*) or occupational titles (*teacher*). They form a segment of the animacy hierarchy (Silverstein, 1976), which is widely used in language typology, see Figure 1. Personal nouns correspond to the segment characterized as [–proper, +human], between proper names and common nouns denoting non-human living beings.

Identifying personal nouns is not only motivated by typological interests. In German, a language with a tri-partite grammatical gender system (masculine, feminine, neuter), there are morphological means to express the gender of the persons referred to, which leads to a congruent interpretation of grammatical form and human gender (such as *mother* and *father* or *actor* and *actress* in English). In recent years, there has been a vigorous debate in Germany whether to consequently disambiguate personal nouns in concordance with the gender of their referents (Kunkel-Razum, 2020). The actual implementation in texts varies between using masculine forms as the traditional 'general' expression (e.g. *die Zuschauer* ['the spectators, masculine']), explicit markings of feminized (e.g. *die Zuschauerinnen* ['the spectators, feminine']) and gender-diverse forms with a special character and feminine suffix (e.g. *die Zuschauer:innen* ['the spectators, gender-diverse']), or using neutral forms (e.g. *die Zuschauenden* ['the spectators, neutralized for gender']).

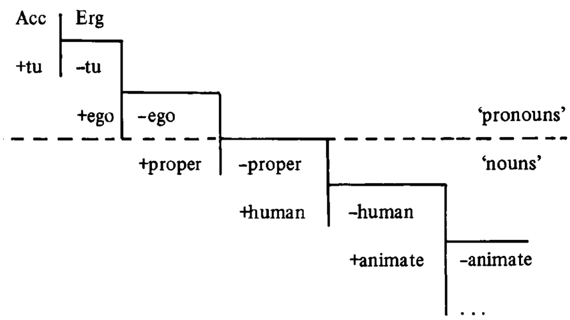

Figure 1: Personal nouns form the segment [–proper, +human]. The animacy hierarchy of Silverstein (1976) was originally introduced for typological analyses of 'accusative' vs. 'ergative' case-marking splits.

As a result, personal nouns are a crucial part of expressing gender in German texts, and thus also a crucial part of manifesting gender stereotypes in texts. The detection of personal nouns is useful for analyzing these stereotypes from various perspectives.

Gender bias in large language models or their training data has become an active research field

Proceedings of the 19th Joint ACL - ISO Workshop on Interoperable Semantic Annotation, pages 33–39
June 20, 2023. ©2023 Association for Computational Linguistics

in NLP.[1] There are methods of detecting gender bias in word embeddings such as the Word Embedding Association Test (WEAT) (Caliskan et al., 2017). One method of balancing gender in the training data, for example, is 'Counterfactual Data Augmentation' (Lu et al., 2019) which is based on adding synthetic sentences to the training corpus that are created by means of a bidirectional lexicon of gendered words such as *actor:actress*. In languages like German, such a lexicon would need to include all personal nouns, because German uses lexical and morphological means very productively to create their feminized or neuter forms.

From a linguistic standpoint, recent developments of gender-fair language in German (Kunkel-Razum, 2020) have led to increasing interest in the forms and use of personal nouns, e.g. regarding frequencies of feminized forms (*Student* ['student, masculine'] > *Studentin* ['student, feminine']) or neutral forms derived from a verbal participle form (*Studenten* ['students, masculine'] > *Studierende* ['people who study, plural, neutralized for gender']). Newer overtly gender-inclusive forms employ e.g. an asterisk (*Wähler*innen* ['voters, female plural suffix']) or a colon (*Bürger:innen* ['citizens, female plural suffix']) to explicitly include not only women but people of all non-binary genders. The problem with researching these phenomena in a quantitative way is that it has not been possible to gauge the basic population of personal nouns in a given corpus in order to put frequencies of e.g. forms with an asterisk into perspective, for instance to approximate whether such forms are getting more frequent.

This is due to personal nouns being a heterogeneous class in German that includes the products of many different word formation processes. Derivational suffixes for personal nouns, for example, include *-er* (*Lehrer* 'teacher'), *-ung* (*Leitung* 'leader/manager') and *-ling* (*Lehrling* 'apprentice'). This heterogeneity is further complicated by some personal nouns being ambiguous with non-personal nouns, e.g. *Leitung* 'leader, manager' vs. *Leitung* 'wire, pipeline', restricting the use of word-list based approaches like Kokkinakis et al. (2015) for Swedish vocational terms. Furthermore, other nouns that do not refer to a human contain these suffixes as well (e.g. *Gräber* 'graves', *Fälschung* 'forgery', *Frischling* 'shoat'), leading to false positives when querying a corpus for these word formation patterns. Thus, it is not possible to identify all personal nouns in a corpus with a regular expression without extensive manual correction. Instead, machine learning-based token classification could be the way to go.

To test the feasibility of such a semantic annotation, we have fine-tuned a pre-trained language model on manually annotated data to automatically identify personal nouns in a corpus. We discuss problems of the annotation and perform a qualitative error analysis on the results. The classifier model is freely available.[2]

While our work focuses on German, research on gender-fair language has been conducted for other gendered languages as well (see Robiche 2018 for French and Verelst 2022 for Dutch). Thus, a classifier that is able to detect personal nouns could also be fruitful for research in other languages.

2 Previous work on personal nouns

Quantitative work on personal nouns in German so far has either looked at pre-chosen lexemes where it is possible to extract all forms of the whole paradigm and thus know the basic population (e.g. Elmiger et al. 2017; Adler and Plewnia 2019), or has resorted to manually identify personal nouns in a corpus (e.g. Ivanov et al. 2018, Acke 2019, Müller-Spitzer et al. 2022). Elmiger (2018, 184) defines personal nouns as "nominal expressions that are used [...] to refer to human beings". While this definition might seem straightforward, it is often difficult to determine if a noun is indeed used to refer to human beings. Some problems identified in Elmiger et al. (2017) include ambiguous nouns and collective nouns that can be used either in a personalized way (1-a) or instead referring to an institution or organization (1-b) (Elmiger et al., 2017, 195-197) .

(1) a. Most Democrats voted in favor of the motion.
b. The Democrats lost votes to the Republicans.

These issues, especially ambiguity, lead to the problem that even to query specific lexemes in a corpus will yield false positives, for example for nouns such as *Berliner* that can be used both as an adjec-

[1] See, e.g., the workshop series on Gender Bias in Natural Language Processing (https://genderbiasnlp.talp.cat/) and their proceedings on https://aclanthology.org/.

[2] https://huggingface.co/CarlaSoe/personal-noun-detection-german-bert.

tive and as a noun, and which has a personal and non-personal meaning as a noun on top of that, see the examples in (2).

(2) a. Er ist ein *Berliner* Bäcker.
 'He is a *Berlin* baker'
 "He is a baker from *Berlin*."
 b. Er ist ein *Berliner*.
 "He is a *Berliner*." (Berlin native)
 c. Er isst einen *Berliner*.
 "He eats a *donut*."

While POS annotation can help to distinguish the adjectival from the nominal use, it does not help to distinguish between the latter two nominal usages.

In the context of digital humanities, Flüh and Schumacher (2021) trained a classifier to extract and assign gender roles in German literary texts, targeting personal nouns as well as proper names of literary characters. While the task of automatically detecting personal nouns is similar to Named-Entity Recognition as it is a token classification task, it differs insofar as the tokens to be detected are crucially not named entities—proper names are not part of the semantic class of personal nouns.

3 The personal noun detection task

Objective. The objective of the detection task is a binary classification of all tokens in a corpus as either personal noun (PERS_N) or other (O).

Conceptually, a personal noun is a token t in a text that meets the following criteria: (i) the lexical-semantic class of t is [–proper, +human]; (ii) t is used in a context that refers to a person or a group of humans; (iii) the part of speech of t is common noun.

Following Elmiger's (2018) approach to personal nouns, the detection task targets all noun tokens denoting humans regardless of their referential context, including generic, non-generic, and predicative contexts. Metonymic uses of nouns—such as referring to an institution or an organization instead of referring to a person—are labeled "other" (O) according to criterion (ii). For example, *Gewinnerin* ('winner, female') in example (3), which refers to the Green Party, is labeled O.

(3) [. . .] die Grüne Partei der Schweiz (GPS) [ist] die große *Gewinnerin* [. . .]
 "[. . .] the Green Party of Switzerland (GPS) [is] the big *winner* [. . .]"

Furthermore, the task excludes personal noun instances that occur as the first part of a compound noun such as *Bauern* in example (4).[3] Because the token *Bauern-Proteste* refers to the event of 'protest', and only the subtoken *Bauern* fulfills the definition of a personal noun, it is disregarded for the annotation.

(4) Die größten *Bauern*-Proteste gab es in Bonn.
 "The biggest *farmers*' protests took place in Bonn."

Proper names are, by definition, not personal nouns and are labeled "other" (O).

We would like to point out that the task operates on the token level, instead of the phrasal level, because our research interest are forms of gender-fair language in German. This is essentially expressed on the lexical level even if it requires contextual and referential disambiguation. The personal noun detection task is therefore different from, e.g., the task of (phrasal) markable detection in coreference resolution.

3.1 Data

We use the corpus from Sökefeld (2021) which consists of roughly 130,000 tokens from two different text types (newspaper and blog). The news subcorpus was compiled by selecting twelve articles each from the politics section of seven German online news outlets.[4] For the blog subcorpus, posts from the blogging platform `wordpress.com` were selected that had been tagged either as "Alltag" ('everyday life') or as "Tagebuch" ('diary') in order to capture more colloquial language use.

For the new task of personal noun detection, we enriched the corpus with additional annotations (see section 3.2).

Because of copyright issues, it is not possible to share the corpus, but metadata with links to the articles and blog posts is provided with the classifier model (see Section 3.3).

3.2 Annotation

In the initial corpus, only personal nouns that refer to a person or people of more than one gender as in

[3] Compounding is very productive in German and results either in merged words without a space or in hyphenated compounds.

[4] *Bild*, *Frankfurter Rundschau*, *Neues Deutschland*, *Süddeutsche Zeitung*, *taz. die tageszeitung*, *Die Welt*, and *Die Zeit*.

Data	Tokens	Types
Training	3,342	1,331
Test	384	289

Table 1: PERS_N types and tokens in training and test set.

example (5)[5], or where the gender of the referent(s) is unclear as in example (6), had been annotated manually.

(5) Die meisten *Migranten* zogen weg, nur fünf Familien blieben.
"Most *migrants* moved away, only five families stayed."

(6) Am besten holt ihr noch ein Familienmitglied oder eine/n gute/n *Freund/in* ins Boot.
"It would be best if you got a family member or a good *friend* on board."

For the personal noun detection task, the original corpus was enriched and all personal nouns with a gender-specific referent (either a male or female individual, or a group of only male or female people) were annotated with a semi-automatic approach. This was conducted in four steps: First, a list of word forms was derived from Sökefeld's (2021) annotations; Second, the list was applied to automatically tag all additional gender-specific instances of these word forms in the corpus; Third, the resulting annotations were manually corrected and, fourth, additional personal noun tokens that had not been included in the earlier list of word forms were annotated in the correction process. The second step yielded many false-positive labels for ambiguous word forms such as *Deutsche* 'German', which can either be used as a personal noun or as an adjective, or *Alter* 'old person'; 'age', suggesting that an approach of matching a list of previously discerned personal nouns to a corpus would not yield sufficient accuracy. All manual annotation was carried out by one annotator.

All in all, the label PERS_N was not very prevalent in the data. There were only 3,726 tokens (roughly 3%, 1,441 different types) labeled as PERS_N compared to 126,459 "other" tokens.

3.3 Training

We split the sentence-wise annotated corpus in 10% test data and 90% training data for fine-tuning a

[5]Target words are italicized in the examples.

Label	Precision	Recall	f1-Score	Support
O	1.00	1.00	1.00	12,495
PERS_N	0.94	0.93	0.94	384
PERS_N$_{OOV}$	1.00	0.88	0.93	113

Table 2: Results of the fine-tuned model on the test set, with scores for overall PERS_N-types and OOV-PERS_N-types.

token classifier[6] based on the pre-trained language model *bert-base-german-cased*[7] for the new task of personal noun detection.

Since the personal noun annotation was performed on the token level, we applied the transformer tokenizer on already tokenized sentences. We used the default hyperparameters for training[8] and evaluated the model on token level on the remaining 10% of the corpus (with 384 tokens (289 types) marked as PERS_N). Of the personal nouns in the test set, 110 types were out-of-vocabulary in the sense of not being present in the training set (although they might be present in the pre-trained language model). Table 1 shows the distribution of personal noun types and tokens in the training and test set.

The fine-tuned model and information on the corpus (metadata and URLs to the original texts) are provided on Huggingface.[9]

4 Results and discussion

The results of the fine-tuned model's performance on the test data are shown in Table 2. The results were quite good for both recall and precision, particularly considering the small amount of data and the low frequency of the target category in this data. Performance on out-of-vocabulary types (see Section 3.3) was similar to the overall results, but with a higher precision and a lower recall.

Overall, there were 22 cases of false positives (see (7) for an example) and 27 cases of false negatives (see (8) for an example) in the test data. Ex-

[6]By following the tutorial on `https://huggingface.co/course/chapter7/2?fw=pt` (last used May 8th 2023).

[7]`https://huggingface.co/bert-base-german-cased` (last used May 8th 2023)

[8]As specified in the huggingface tutorial, see footnote 6: Number of training epochs: 3; learning rate: $2e^{-5}$; weight decay: 0.01.

[9]`https://huggingface.co/CarlaSoe/personal-noun-detection-german-bert/tree/main`.

ample (7) showcases an interesting example of a false positive that could be considered a peripheral, non-prototypical personal noun, as a generation is made up of people. The model's classification of this token showcases that some categorization decisions are not as clear-cut as they may seem on the surface.

(7) Von Generation zu *Generation* schwand das Wissen um den Ursprung des Wohlstands der Familie.
"The knowledge about the origin of the family's wealth faded from generation to *generation*."

(8) Du bist ein elender *Heuchler*.
"You are a wretched *hypocrite*."

The personal noun *Heuchler* in example (8) was not detected by the model as such. This could be due to its relative infrequency.[10] It was also not part of the training data for the fine-tuning.

On closer inspection, though, the false negatives and positives in some cases revealed not a mistake of the model, but an error in the manual annotation. These included errors from the automatic annotation that were not caught and corrected during the manual correction, such as *Deutschen* being categorized as a personal noun in example (9). These oversights stress the importance of using more than one annotator when manually labeling data, so that errors like this can be avoided.

(9) Ähnlich äußerte sich der Präsident des *Deutschen* Städtetags [. . .]
"The president of the *German* Association of Cities expressed himself similarly [. . .]"

Apart from looking at the model's performance on the test data, we also tested instances of challenging phenomena as identified by Elmiger (2018) that make distinguishing between personal nouns and other words difficult.

First of all, ambiguity can pose a problem. We tested the two word forms *Berliner* and *Hamburger* that can both be used as an adjective and as a noun, as well as having both a personal noun usage and a 'food' meaning. Both word forms were correctly not classified as a personal noun in their adjectival usage, but *Berliner* as a noun was labeled a personal noun in both the 'food' usage and the 'person from Berlin' usage. For *Hamburger*, on the other hand, the model correctly only labeled the usage as a personal noun as such.

Secondly, personalized and institutional usages of collective nouns were tested with the word forms *Polizei* 'police' and *Menge* 'amount', 'crowd'. For both word forms, the model managed to correctly label the personalized usage as a personal noun in example (10-a), and not label the impersonal usage in example (10-b).

(10) a. Die *Polizei* schoss auf Demonstrant:innen.
"The *police* shot at protestors."
b. Die *Polizei* ist Teil der Exekutive.
"The *police* is part of the executive."

Proper names could also pose a problem for the classification, as a lot of last names are derived from personal nouns but should not be detected by the model. In fact, the model was able to differentiate correctly between personal noun, in example (11-a), and proper name usage, in example (11-b), for *Schneider* ('tailor'), but it did not detect *Müller* ('miller') as a personal noun in example (11-c), which is the most common family name in Germany,[11] but the occupation has become rare, so that *Müller* only appears as a last name in the training data and not in its personal noun usage.

(11) a. Ich bringe ein Hemd zum *Schneider*.
"I bring a shirt to the *tailor*."
b. Frau *Schneider* sitzt auf einer Bank.
"Ms *Schneider* is sitting on a bench."
c. Ich bringe das Getreide zum *Müller*.
"I bring the grain to the *miller*."

Finally, we tested how the model responds to metalinguistic uses of personal nouns. The model labeled the word forms of *Frau* and *Mann* in their metalinguistic uses in the examples in (12) as personal nouns.

(12) a. *Frauen* ist der Plural von *Frau*.
"*Women* is the plural of *woman*."
b. Das Wort *Mann* ist ein Nomen.
"The word *man* is a noun."

(13) "*Frauen*" ist der Plural von "*Frau*".
"'*Women*' is the plural of '*woman*'."

[10] See https://corpora.uni-leipzig.de/de/res?corpusId=deu_news_2022&word=Heuchler (last used May 8th 2023) for frequency information.

[11] For a list of common family names in Germany see https://de.wiktionary.org/wiki/Verzeichnis:Deutsch/Namen/die_hufigsten_Nachnamen_Deutschlands (last used May 8th 2023)

Interestingly, when adding quotation marks to the sentence in (12-a) as in example (13), the model only labeled *Frau* as a personal noun, but not *Frauen*. For the sentence in example (12-b), though, it did not make a difference whether *Mann* was set in quotation marks or not.

Another challenge for the study of gender-fair language is that new forms keep evolving. Testing the new colon form (e.g. *Schüler:innen* 'students') that became popular only after the training corpus was compiled in 2019, the model still labeled the token *Demonstrant:innen* in example (14) as a personal noun. This shows that it could be useful for identifying new strategies of gender-fair language emerging in the future as well.

(14) Die Polizei schoss auf *Demonstrant:innen*.
"The police shot at *protestors*."

5 Conclusion

Personal nouns, the semantic class of common nouns denoting humans, are of great importance in the context of current discussions and developments in research on gender-fair language and language use in linguistics and digital humanities, as well as gender-fair NLP. In order to facilitate quantitative research, we defined the task of personal noun detection and fine-tuned a pre-trained language model for the detection of personal nouns in German.

The fine-tuning yielded surprisingly good results (f1-score: 0.94), considering the small amount of training data and the fact that the actual tokens of interest were not very prevalent. Further training on more diverse data including other text types, for example literary texts, which probably contain a range of different personal nouns not covered in news writing or personal blog posts, could improve the results even more. New training data could also include specifically selected sentences containing some of the more difficult to distinguish words as discussed in Section 4, like ambiguous words, proper names, and metalinguistic usages.

So far, the classifier only detects personal nouns but does not give any additional information on them. Ideally, a future version of the model would further enrich this classification. An initial expansion could be to detect grammatical gender. Much less trivial, but desirable, would be to implement a further classification of the type of reference, as qualitative research has shown that gender-fair forms tend to be used more frequently in cases of non-generic reference (Pettersson 2011, Sökefeld 2021). Incorporating a distinction between generic (15) and non-generic (16) use (see Friedrich and Pinkal 2015) into the classifier would make it possible to test whether this holds true on a larger scale.

(15) Kein *Bauarbeiter* hält bis 69 durch.
"No *construction worker* will manage to keep it up until the age of 69."

(16) Als Reaktion sprangen *Schüler*innen* und *Studierende* zunächst über die Drehkreuze an den Zugängen zu den Bahnsteigen.
"As a reaction, *pupils* and *students* initially jumped the barriers at the entry to the platform."

Similarly, whether a personal noun refers gender-specifically (e.g. masculine *Lehrer* referring to only male teachers) or gender-independently (e.g. masculine *Lehrer* referring to a mixed-gender group of teachers) is necessary information in order to quantify the amount of masculine personal nouns used to refer to gender-diverse groups.

Training the language model to classify personal nouns in these three categories would thus be a next step.

6 Ethics statement

We are aware that the corpus we used as training data contains texts that potentially include gender stereotypes. A possible application of our classifier could be to identify such stereotypical depictions.

References

Hanna Acke. 2019. Sprachwandel durch feministische Sprachkritik: Geschlechtergerechter Sprachgebrauch an den Berliner Universitäten. *Zeitschrift für Literaturwissenschaft und Linguistik*, 49(2):303–320.

Astrid Adler and Albrecht Plewnia. 2019. Die Macht der großen Zahlen. Aktuelle Spracheinstellungen in Deutschland. In Ludwig M. Eichinger, editor, *Neues vom Heutigen Deutsch: Empirisch - Methodisch - Theoretisch*, in Jahrbuch des Instituts für Deutsche Sprache. De Gruyter, Boston, MA.

Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.

Daniel Elmiger. 2018. French anthroponyms as a heterogeneous category: Is there such a thing as personal nouns? *International Journal of Language and Culture*, 5(2):184–202.

Daniel Elmiger, Eva Schaeffer-Lacroix, and Verena Tunger. 2017. Geschlechtergerechte Sprache in Schweizer Behördentexten. Möglichkeiten und Grenzen einer mehrsprachigen Umsetzung. In Constanze Spieß and Martin Reisigl, editors, *Sprache und Geschlecht. Band 1: Sprachpolitiken und Grammatik*, volume 90 of *Osnabrücker Beiträge zur Sprachtheorie*, pages 61–90. Duisburg.

Marie Flüh and Mareike Schumacher. 2021. Digitale diachrone Korpusanalyse am Beispiel des Projekts "m*w – Gender Stereotype in der Literatur". In Gisela Mettele, Marint Prell, and Pia Marzell, editors, *Digital humanities and gender history*. Jena.

Annemarie Friedrich and Manfred Pinkal. 2015. Discourse-sensitive Automatic Identification of Generic Expressions. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1272–1281, Beijing, China. Association for Computational Linguistics.

Christine Ivanov, Maria B. Lange, and Tabea Tiemeyer. 2018. Geschlechtergerechte Personenbezeichnungen in deutscher Wissenschaftssprache. Von frühen feministischen Vorschlägen für geschlechtergerechte Sprache zu deren Umsetzung in wissenschaftlichen Abstracts. *Suvremena lingvistika*, 44(86):261–290.

Dimitrios Kokkinakis, Ann Ighe, and Mats Malm. 2015. Gender-Based Vocation Identification in Swedish 19th Century Prose Fiction using Linguistic Patterns, NER and CRF Learning. In *Proceedings of the Fourth Workshop on Computational Linguistics for Literature*, pages 89–97, Denver, Colorado, USA. Association for Computational Linguistics.

Kathrin Kunkel-Razum. 2020. Gender-Neutral German: Asterisk or Underscore? *Germany #NOFilter*. Online. Goethe-Institut.

Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. 2019. Gender bias in neural natural language processing. *CoRR*, abs/1807.11714.

Carolin Müller-Spitzer, Jan Oliver Rüdiger, and Sascha Wolfer. 2022. Olaf Scholz gendert. Eine Analyse von Personenbezeichnungen in Weihnachts- und Neujahrsansprachen. *Linguistische Werkstattberichte*. Online.

Magnus Pettersson. 2011. *Geschlechtsübergreifende Personenbezeichnungen: eine Referenz- und Relevanzanalyse an Texten*. Vol. 11 in Europäische Studien zur Textlinguistik. Narr Verlag, Tübingen.

Lou Robiche. 2018. *Pratiques sociolinguistiques francophones de féminisation et de dégenrement*. Espaces discursifs. L'Harmattan, Paris.

Michael Silverstein. 1976. Hierarchy of features and ergativity. In R.M.W. Dixon, editor, *Grammatical categories in Australian languages*, pages 112–171. Australian Institute of Aboriginal Studies, Canberra. Re-published in: Pieter Muysken and Henk van Riemsdijk, editors, 1986. *Features and Projections*, pages 163–232. Foris, Dordrecht.

Carla Sökefeld. 2021. Gender(un)gerechte Personenbezeichnungen: derzeitiger Sprachgebrauch, Einflussfaktoren auf die Sprachwahl und diachrone Entwicklung. *Sprachwissenschaft*, 46(1):111–141.

Natalie Verelst. 2022. Movierung und generisches Maskulinum aus kontrastiver Perspektive. Talk presented at *Die Movierung. Formen – Funktionen – Bewertungen*, University of Passau.