Supporting Gender-Neutral Writing in German

THEODOR DIESNER-MAYER and NIELS SEIDEL, FernUniversität in Hagen, Germany

The avoidance of the generic masculine is an important part of a gender-neutral use of the German language. This paper presents a rule-based natural language processing system that identifies occurrences of the generic masculine and suggests gender-neutral corrections. The system was evaluated with 238 labeled news articles from the TIGER corpus. About 88 % of the occurrences were correctly identified. Grammatically well-formed suggestions could be generated for about 94 % of the correctly identified occurrences. A web application was developed to support gender-neutral writing by suggesting corrections for nouns and pronouns written in the generic masculine.

CCS Concepts: • **Computing methodologies** → **Natural language processing**.

Additional Key Words and Phrases: natural language processing, writing support, gender

ACM Reference Format:
Theodor Diesner-Mayer and Niels Seidel. 2022. Supporting Gender-Neutral Writing in German. In *Mensch und Computer 2022 (MuC '22), September 4–7, 2022, Darmstadt, Germany*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3543758.3547566

1 INTRODUCTION

Gender-neutral language is language use that aims at equal treatment and gender equality in spoken and written language. In such language use, gender bias effects are to be avoided [4]. In this context, such formulations should not be used that lead to false or misleading gender-related representations. In the German language, such a *gender bias* arises in particular from the use of the generic masculine, according to advocates of gender-neutral language.
Natural languages differ, among other things, in the use of grammatical gender. There are languages, like English, which do not know a grammatical gender. The German language belongs to the so-called *grammatical gender languages* (GGL), which are characterized by the fact that nouns and pronouns are assigned a grammatical gender, the so-called genus (e.g. *der* Stuhl/the chair). If the noun or pronoun refers to a person or a group of persons, in GGL the genus of the noun or pronoun usually corresponds to the gender of the person or group of persons mentioned [6, p. 269].
If a pronoun or noun refers to an indefinite person or a mixed-gender group, it is common in German to use a noun with a masculine gender (e.g., in the sentence: Die *Bürger* gehen wählen.). This is called the generic masculine (German: Generisches Maskulinum). In the generic masculine, the masculine form of the person's name is always used. Even if the group meant consists of persons of different genders. For example, the headline "Bäcker erzielen große Gewinne" ("Bakers achieve high profits") probably refers not only to male bakers, but will mean persons of all genders.
Advocates of gender-neutral language see in the use of the generic masculine the problem that the non-male persons are left behind in the imagination of readers and listeners [7]. The consequence of this, is that existing prejudices in society about role models of the genders are maintained or even strengthened.

This paper presents a system that is able to detect usage of generic masculine in text. The system is also able to suggest gender-neutral alternative formulations for the user's given text. To achieve these results, the system uses

Manuscript submitted to ACM

natural language processing techniques, mainly based on static grammatical rules, to achieve this goal. Given this functionality, the system aims to support and remind users of the (correct) usage of gender-neutral language. The system should work for a wide range of texts and topics. Correction suggestions are created for a single sentence, though a longer text may be given as input for the analysis beforehand and the context of multiple sentences are taken into account for the detection. In the future, such systems could be part of common word processing programs such as Libre Office, in addition to existing spell and grammar checking tools.

In a more formal way: The system takes the text of a user as input parameter. Given this, its goal is to highlight occurrences of words, which are formulated as generic masculine. For each of these occurrences, it also generates two correction suggestions, of which the user can choose. It suggests the usage of the pair form (mentioning the masculine and feminine of the word) or the usage of special gender separators such as the gender asterisk (such as the gender star). The gender star is an (unofficial) typographic style, which means to include all genders when referring to a person or group of persons in the German language. For example, given the sentence "Der Student benutzt den Toaster." ("The student uses the toaster."), the system determines that "Student" is a usage of the generic masculine, while "Toaster" is not (though "toaster" has a masculine grammatical gender). As correction, for example, the sentence "Der*Die Student*in benutzt den Toaster" is suggested.

2 RELATED WORKS

There are only a few related works about such assistance systems for gender-neutral language for the German language. Especially for the problem at hand, the avoidance of the generic masculine, only one publication with a similar goal was found.

Carl et al. [2] presented in their paper "Gendercheck Editor" the methodology behind their developed tool, the "Gendercheck Editor". The tool is based on a rule-based procedure that uses the grammatical structure of a sentence to decide whether an occurrence of a noun or pronoun has been formulated in a gender-neutral way or not. In this respect, it is precisely the recognition of the generic masculine. This tool informs users when they find a potentially gender-inappropriate phrase and provides information on how to manually correct the phrase [2]. It should be taken into account that the detection of the Gendercheck Editor is mainly aimed at the detection of phrases in administrative and legal texts. There is no automatic correction of the found problematic phrases. The GenderCheck editor is able to find three classes of errors: (i) The denotations of persons in the generic masculine should be detected (e.g. "Der *Beamte* muss den Anforderungen genügen." / "The official must meet the requirements."). (ii) Male relative pronouns, possessive pronouns, and personal pronouns that refer to a person whose gender is not known, shall be recognized (e.g. "*Jeder* muss seinen Beitrag leisten." / "Everyone has to do their part."). (iii) Inappropriate uses of a masculine denotation for, clearly identifiable, female persons should be recognized (e.g. "Ihr *Ansprechpartner* ist Frau Müller." / "Your contact person is Mrs. Müller.").

The recognition is rule-based using the morphology of the sentences, which is determined by the so-called MPRO system [5]. For example, a masculine noun is not marked if the noun is followed by a proper noun ("Der Beamte Müller soll kommen." / "The official Müller is supposed to come."). The rules are only applied to the sentences. Consequently, an occurrence of a non-gender-neutral phrase can only be found if the sentence itself, but not the surrounding sentences, indicates such a phrase. For the sentence like "Einem *Bewerber* ist Urlaub zu gewähren." ("Leave shall be granted to an applicant."), a warning can therefore be given. In the sentence "Ein Anspruch auf Fortzahlung *seiner* Bezüge besteht nicht." ("He shall not be entitled to continued payment of his remuneration."), it cannot be recognized, that the possessive pronoun *seiner* (*his*) must also be changed since the sentence boundary is exceeded.

Since the Gendercheck Editor is not publicly available, no qualitative comparisons could be made with the system presented here. While the works share similar goals, there are some differences in the general goals: (i) The presented system should be able to be used for different types of texts. Utilizing a publically available, crowdsourced word database, more types of text should be supported. (ii) The presented system aims for a better detection accuracy by analyzing surrounding sentences. (iii) The presented systems aims to generate alternative formulations as suggestions, which can be applied by a user.

For the English language, only a few researchers addressed the issue of gender-neutral language so far. [9] and [8] used simple rules for the grammatical conversion of pronouns (*he, she*) into a form in which all genders are mentioned (*he/she/they*). Furthermore, [9] enabled the transformation of gendered nouns (e.g. *actress*, *chairman*) to its corresponding gender-neutral form (*actor*, *chair*). The results of these two rule-based methods have been used to train Transformer Models that are applicable to translate any English text into a gender-neutral form [8, 9].

3 IMPLEMENTATION TO SUPPORT GENDER-NEUTRAL WRITING

The system implements the detection and creation of the corrections in three steps: identification of masculine nouns and pronouns (step 1), checking for forms of the generic masculine (step 2), and creation of suggested corrections (step 3). For both, the second and the third step, a comprehensive database of German words and their inflected forms was necessary. To make the system accessible for users, a web application was developed, which provides convenient graphical user interface. The source code of the system and the web application presented in the following is publically available at https://github.com/theodm/gender-assistenz.

3.1 Step 1: Search for nouns and pronouns in masculine form

Only masculine nouns and pronouns can be phrased in the generic masculine. The first step is therefore to filter the text for such occurrences. For this search, the text is first tokenized into individual words and punctuation marks (Tokenization). Using morphological analysis, only words whose grammatical gender is masculine are selected. Using part-of-speech tagging (PoS), the set of words is filtered so that only grammatical masculine nouns and pronouns (including substantive and adjectival pronouns) remain to be tagged appropriatly. In summary, after the first step only occurrences of nouns and pronouns are left, whose grammatical gender is masculine. Not yet established for these nouns and pronouns is, whether they're occurrences of *generic* masculine. Since the building blocks of the first step (Tokenization, PoS-Tagging, morphological analysis) are quite elaborated and proven techniques, one can assume that the results of this step are quite accurate.

3.2 Step 2: Checking for the generic masculine

A masculine noun or pronoun is generic masculine, if it refers to a female person or a mixed-gender group. The system assumes, that all occurrences of the first step are generic masculine, unless one of the three following cases applies.

Only personal designations. An occurrence of a noun or pronoun is not generic masculine and shouldn't be phrased in a gender-neutral manner if it doesn't refer to a person. In the sentence "Der Toaster ist ein Gerät" ("The toaster is an appliance") for example, the masculine noun "Toaster" does not refer to a person. The system assumes that a noun does not refer to a person, if a feminine declination of the noun does not exist. For pronouns, coreference resolution is applied to determine the corresponding noun. If this noun does not refer to a person, then neither the pronoun does.

Exclude gender-neutral forms. An occurrence of a noun or pronoun doesn't need to be corrected by the system, if a gender-neutral phrasing has already been used. The system makes this assumption if both, the feminine and the masculine noun or pronouns, are mentioned in the same coordinating conjunction. Other forms of gender-neutral phrasing, for example the usage of the gender-star are not recognized. A coordinating conjunction joins two equivalent phrases or clauses, for example, with the help of commas, "und" ("and") or "oder" ("or") [3]. Such conjunctions can be recognized employing dependency analysis. The nouns and pronouns of a coordinating conjunction are connected by the edges *cd* (conjunct) and *cj* (coordinating conjunction). Starting from the masculine noun or pronoun, this coordination can now be iteratively run through. If the feminine form of the noun/pronoun is found there, then the occurrence does not have to be rephrased.

Nouns/pronouns related to male persons. An occurrence of a noun or pronoun should also not be gender-neutral, if it does refer to a male person or male group. The system assumes such a case, if the noun/pronoun refers to a proper name. In the sentence "Er heißt Peter." ("His name is Peter.") the word "Er" refers to the proper noun "Peter" and as such won't be corrected. Dependency Analysis and Coreference Analysis was used to detect if a noun or pronoun refers to a proper noun. Dependency Analysis was used for inter-sentence Coreference Analysis by detecting certain sentence structures (e.g. Close Appositions, Loose Appositions, Copula Constructions). Coreference Analysis is a NLP-technique to detect if multiple words refer to the same entity or in this case to the same person. In cases where a masculine-declined noun is used for a female person or refers to a group of people of different genders, the assumption is incorrect, e.g. "Die *Studenten* Clara und Marc gehen ein Eis essen." ("The *students* Clara and Marc are going out for ice cream.").

3.3 Step 3: Generating suggestions

After the second step, ideally only occurrences of generic masculine, which can be corrected, are left. For each occurrence, two suggestions are generated. The transformation of occurrences of nouns in plural form are quite simple: Only the noun itself must be replaced with its gender-neutral counter-part. For example, the suggestions for the sentence "Die Studenten gehen ein Eis essen." ("The students go out for ice cream.") are *Die Student*innen gehen ein Eis essen.* (gender-star) and *Die Studentinnen und Studenten gehen ein Eis essen.* (both forms).

In the singular case, multiple words beside the noun / pronoun in question need to change. In order for the sentence to remain grammatically correct when a correction proposal is made, the genus congruence must be preserved. The genus of the so-called attributive adjectives, pronouns and articles depends on the genus of the noun (so-called genus congruence). Therefore, also the corresponding words of the noun / pronoun need to be changed. The corresponding words are found (again) by static grammatical rules with dependency analysis and morphology analysis. The needed forms of adjectives, pronouns and articles are determined statically. For example, the sentence "Der Student geht ein Eis essen." ("The student goes for an ice cream.") yields the suggestions: "Der*Die Student*in geht ein Eis essen." (gender-star) or "Student*innen gehen ein Eis essen." (pluralized gender-star).

3.4 Creation of a dictionary for declined nouns and conjugated verbs

For the different steps of the system, it is necessary to know the different inflection forms for nouns and verbs. For this purpose, a database was created, which contains the different words and their inflections. The data was extracted from the German Wiktionary project[1]. Since the German language allows the creation of compound nouns, which consist of

[1]See: https://dumps.wikimedia.org/ (retrieved 11/20/2021).

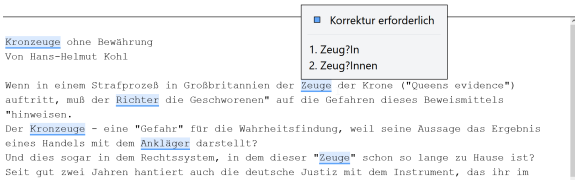

Fig. 1. The web application shows the detected occurrences and suggests ways to formulate them in a gender-neutral way.

multiple nouns and the Wiktionary project doesn't contain all possible compound nouns, a method was developed to infer the inflection forms by grammatical rules and the available noun data.

3.5 Web application

A web application was developed that supports users in checking and improving their writings to become more gender-neutral. The system for the detection of occurrences and correction of generic masculine (steps 1-3) are implemented as a server-side-application in the programming language Python. Using an HTTP-API, any given string can be analyzed. On the client side application, a text field allows the user to enter his text to be analyzed. The found occurrences are highlighted (blue occurrences in figure 1) and via mouse-over corrections are suggested. The changes made by applying a correction can be previewed, since multiple words of the sentence can change. For debugging purposes, the nouns and pronouns, which don't need to be corrected, are also highlighted (gray occurrences in figure 1).

4 EVALUATION

Dataset and Preprocessing. The corpus used to analyze the performance of the system is part of the freely available TIGER corpus [1], whose texts consist of newspaper articles from the newspaper Frankfurter Rundschau from the years 1995–1997. The journalistic texts from politics and economics from this period were not written in a gender-neutral language and, in contrast to more recent texts, contain many occurrences of the generic masculine. The corpus contains 238 articles, 3,765 words, and 79,445 tokens. For determining the performance, all nouns and pronouns were manually marked (7,516 in total) by the first author of this paper to see if it was an occurrence of generic masculine (958) or not (6,556).

Analysis. The system was first tasked to identify nouns or pronouns contained generic masculine. When comparing the occurrences identified by the system with the manually marked occurrences, a distinction was made whether they matched (true positive), an unmarked occurrence was identified to be generic masculine (false positive), or, a marked position was not identified to be generic masculine (true negative). In a second step, the system was tasked to generate corrections for the found occurrences. The generated correction suggestions were manually checked by a human for their grammatical correctness.

Results. As a result of the first task, 88.56 % of the marked occurrences were correctly detected by the system (true positive). A total of 11.44 % of the detected occurrences were incorrect (false positive). 1.42 % of the markers were incorrectly not recognized as an occurrence generic masculine (false negative). 10.02 % of the marked occurrences with generic masculine were not detected (false positive).

The first task of the system detected 1604 occurrences of generic masculine. 851 occurrences were correctly identified (true positive). 753 occurrences were incorrectly identified as being occurrences of generic masculine. For these 1604

occurrences, the system generated two suggestions to reword the sentence in a gender-neutral way. In 79.30 % of these occurrences both generated suggestions were grammatically correct, in 12.09 % only one of both suggestions was correct; in 8,60 % of the cases the correction suggestion led to incorrect grammar. Considering only the occurrences that were correctly identified as such when checking for the generic masculine, for 851 occurrences 94.24 % resulted in correct correction suggestions, 2.82 % in partially correct correction suggestions, and 2.94 % in incorrect correction suggestions.

Discussion. The detection of the generic masculine in nouns and pronouns was particularly successful in case the occurrence was a noun in plural form. Singular nouns were more likely to be misclassified as generic masculine when it was not. The detection in pronouns was mixed.

To recap, the system mainly works on the following assumptions: (1) every noun or pronoun with a grammatical masculine gender is also generic masculine unless (2) no feminine declination of the noun exists [Only personal designations] (3) or it refers to a proper name [Relation to male persons]. The false positive detection of generic masculine was often caused by insufficient coreference resolution used by assumptions (2) and (3). In many cases, the system was not able to detect relations between a noun or pronoun and a person that was mentioned in the text already. One solution would be to improve the coreference resolution method.

Another problem are male persons that have not been explicitly introduced in the text by a proper noun. Since the assumption is based on the existence of a proper noun such occurrences cannot be detected. A possible solution would be to include certain tendencies of the found occurrences (is it in singular form, does a specific word follow, etc.). These tendencies could be identified with machine learning approaches to avoid false positives and false negatives.

The generation of suggestions for gender-neutral plural nouns was particularly successful. The results for the singular nouns and pronouns were mixed. Especially in complex sentence structures, the generation of suggestions fails to change the flection of corresponding words needed to create a new grammatically correct sentence. This errors is traceable to missing grammatical rules and incorrect results of the dependency analysis. The authors conclude that an enhanced rule-based approach can yield even better results for the intended purpose. Since the quality of the suggestions heavily depends on the detection of occurrences of generic masculine, the detection has to be improved first.

5 SUMMARY AND OUTLOOK

In this paper, we present a rule-based natural language processing system that identifies occurrences of the generic masculine and makes suggestions for gender-neutral corrections. The system was evaluated with 238 labeled news articles from the TIGER corpus. About 88 % of the occurrences were identified correctly. Grammatically well-formed suggestions were generated for about 94 % of the correctly identified occurrences. In a web-based text editor, the systems support gender-neutral writing by suggesting a correction for nouns and pronouns written in the generic masculine. In future work, the results of the rule-based approach could be used as a baseline for transformer models. Considering human factors, the acceptance and User Experience of the writing tool as well as the readability of the resulting gender-neutral text should be investigated.

ACKNOWLEDGMENTS

This research was supported by the Research Cluster "Digitalization, Diversity and Lifelong Learning – Consequences for Higher Education" (D^2L^2) of the FernUniversität in Hagen, Germany.

REFERENCES

[1] Sabine Brants, Stefanie Dipper, Peter Eisenberg, Silvia Hansen-Schirra, Esther König, Wolfgang Lezius, Christian Rohrer, George Smith, and Hans Uszkoreit. 2004. TIGER: Linguistic Interpretation of a German Corpus. *Research on Language and Computation* 2, 4 (2004), 597–620. https://doi.org/10.1007/s11168-004-7431-3

[2] M. Carl, S. Garnier, J. Haller, A. Altmayer, and B. Miemietz. 2004. Controlling Gender Equality with Shallow NLP Techniques. In *Proceedings of the 20th International Conference on Computational Linguistics* (Geneva, Switzerland) *(COLING '04)*. Association for Computational Linguistics, USA, 820–es. https://doi.org/10.3115/1220355.1220473

[3] Elke Hentschel and Harald Weydt. 2005. *Handbuch der deutschen Grammatik. 3., völl. neu bearb. Aufl.* (3 ed.). De Gruyter, Berlin. https://doi.org/10.37307/j.2198-2430.2005.03.13

[4] Jutta Kühl. 2010. Geschlechtsbezogener Verzerrungseffekt (Gender Bias). , 4 pages. http://www.genderkompetenz.info/w/files/gkompzpdf/gender_bias.pdf

[5] H.-D. Maas. 1996. MPRO – Ein System zur Analyse und Synthese deutscher Wörter. In *inguistische Verifikation, Sprache und Information. Dokumentation zur Ersten Morpholympics 1994.*, Hausser R. (Ed.). Niemeyer, Tübingen, 141–166.

[6] Jennifer Prewitt-Freilino, T. Andrew Caswell, and Emmi Laakso. 2011. The Gendering of Language: A Comparison of Gender Equality in Countries with Gendered, Natural Gender, and Genderless Languages. *Sex Roles* 66 (02 2011). https://doi.org/10.1007/s11199-011-0083-5

[7] Luise F. Pusch. 1987. Totale Feminisierung. In *Frau ohne Herz: feministische Lesbenzeitschrift.* Frauenzentrum Zürich, Winterthur, 23–23.

[8] Tony Sun, Kellie Webster, Apu Shah, William Yang Wang, and Melvin Johnson. 2021. They, Them, Theirs: Rewriting with Gender-Neutral English. arXiv:2102.06788 http://arxiv.org/abs/2102.06788

[9] Eva Vanmassenhove, Chris Emmery, and Dimitar Shterionov. 2021. NeuTral Rewriter: A Rule-Based and Neural Approach to Automatic Rewriting into Gender Neutral Alternatives. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 8940–8948. https://aclanthology.org/2021.emnlp-main.704