Remedying Gender Bias in Open English Wordnet

John P. McCrae
Insight Centre & ADAPT Centre
University of Galway
john@mccr.ae

Haotian Zhu
University of Washington
haz060@uw.edu

Fei Xia
University of Washington
fxia@uw.edu

Al Waskow
Insight Centre
University of Galway
margaret.waskow@universityofgalway.ie

Kexin Gao
University of Washington
kexing66@uw.edu

Abstract

Open English Wordnet aims to improve and maintain a wordnet for English, based on the Princeton WordNet. In this context, we identify a number of gender biases in the existing wordnet and consider the challenges of remediating the biases in the resource. In particular, we look at structural, contextual and definitional biases in the resource and examine how changes to the structure of the wordnet and to the textual definitions can create a wordnet that more fairly represents reality. We propose a number of changes that introduce 317 new synsets as well as changing the definitions or relations of over 400 further synsets. We show that these changes reduce certain kinds of gender bias within the resource.

1 Introduction

English Wordnet, first introduced by Princeton as Princeton WordNet (Miller, 1995; Fellbaum, 2010, PWN) and more recently as an open-source project, Open English Wordnet (McCrae et al., 2019, OEWN)[1], is one of the primary resources for computational lexicography. Like many lexical resources, English Wordnet reflects some of the biases of when it was created and this has been criticised in general in lexicography (Pettini, 2021). Recently, Zhu et al. (2024) identified a number of challenges specifically with English Wordnet and in this work, we consider the challenges associated with removing or limiting these biases. To this extent, we aim to tackle three main forms of bias. Firstly, structural bias, which we interpret primarily by the inclusion of gendered role words (such as 'policeman') in the resource. We modify the resource such that gendered words are in a unique synset that is also marked as gendered by being a hyponym of the synset for 'male'[09647338-n] or 'female'[09642198-n]. Secondly, we look at the use of pronouns in definitions[2] and analysed those that represent a gender bias. We considered how to change this and decided to use the singular 'they'. Finally, we look at definitions and show that as a result of the changes we made, the definitions used for female terms are improved over PWN. These changes are all part of the 2024 release of Open English Wordnet. Finally, we note that distributional bias, as identified by Zhu et al. (2024), is still a major issue and we further identify ways in which English Wordnet under-represents women. In this paper, we only consider gender bias, but other kinds of bias exist in the resource and these techniques could be helpful in fixing these issues.

In this paper, we first consider related work in Section 2 and then provide a definition of bias in Section 3. We then introduce the Open English WordNet in Section 4 and describe our methodology for removing bias in Section 5 and the results of the work. In Section 6, we discuss some of the limitations and potential future work, before finishing with conclusions in Section 7.

2 Related Work

In this study, we focus on gender bias in English Wordnet. We first discuss a taxonomy of gender bias in human-generated text and then review previous research on gender bias in NLP research.

2.1 Taxonomy of Gender Bias

To meaningfully categorize various kinds of gender bias, Hitti et al. (2019) propose two types of gender bias in text: **structural** and **contextual** bias. **Structural** bias 'occurs when bias can be traced down from a specific grammatical construction,' including gender generalization (e.g., generic *he*) and explicit marking of sex (e.g.,

[1]'Wordnet' is the generic term for this kind of resource, 'WordNet' is a trademark of Princeton University. We use the term *English Wordnet* to encompass the wordnets released by Princeton along with those subsequently released as open-source resources

[2]We do not include the example text within definitions. These have been separated in an early version of OEWN

'*chairman*' vs. '*chairwoman*'). **Contextual** bias 'requires the learning of the association between gender marked keywords and contextual knowledge,' which includes societal bias, where societal norms reflect traditional gender roles, and behavioural bias, which is a generalization of attributes and traits onto a gendered person.

Based on Hitti et al. (2019), Doughman et al. (2021) and Doughman and Khreich (2022) provide a more fine-grained taxonomy with five types of gender bias, linking each type to possible real-world implications.

2.2 Gender Bias Study in NLP

The identification and quantification of gender bias have received increasing attention in the realm of NLP in recent years. There are various kinds of gender bias that researchers have examined: gender bias in text (Cryan et al., 2020; Li et al., 2020), in NLP systems (Zhao et al., 2018; Savoldi et al., 2021), in language models (Bordia and Bowman, 2019; Fatemi et al., 2023) and in word embeddings (Bolukbasi et al., 2016; Caliskan et al., 2017; May et al., 2019).

Another important aspect of studying gender bias lies in bias mitigation. To mitigate bias in text, Sczesny et al. (2016) explore the use of gender-fair language in overcoming gender stereotyping and discrimination. Dinan et al. (2020) propose a general-purpose technique to mitigate gender bias in human-generated dialogue utterances by leveraging data augmentation, positive-bias data collection and bias-controlled training. In this work, we focus on identifying and mitigating gender bias in human-generated text, namely, in English Wordnet.

3 Definition and Statement of Bias

We adopt the definition of gender bias as given in Doughman et al. (2021): '*an exclusionary, implicitly prejudicial, or generalized representation of a specific gender as a function of various societal stereotypes.*' To systematically understand what kinds of gender bias exist in English Wordnet, we adopt and extend the gender bias taxonomy from Hitti et al. (2019) and Doughman et al. (2021). In our study, we first consider **structural bias** and **contextual bias**. We also add two new types of bias: **distributional bias** and **definitional bias**. Table 1 lists all bias types and illustrative examples.

3.1 Structural Bias

Structural gender bias considers the association between various linguistic patterns and gender. Such bias can occur at different linguistic levels such as morphology, syntax, semantics, etc.

3.1.1 Explicit Marking of Sex (B1)

At the morphological level, explicit marking of sex appears when gender-neutral entities are denoted by gender markers such as '*-man*' and '*-woman*.' Here, the term 'gender marker' refers not to grammatical gender markers but to free morphemes such as '*-woman*' in '*needlewoman*'[3] or head nouns in compound phrases such as '*woman*' in '*slovenly woman*'. **B1** in Table 1 presents an example where '*policeman*' contains the marker '*-man*' whereas its definition denotes a gender-neutral meaning.

3.1.2 Generic *he* (B2)

We also examine the generic usage of the gendered pronoun '*he*' where the pronoun is co-indexed with a gender-neutral common noun. As shown in the example from **B2** of Table 1, the word *scientist* is gender neutral but is co-indexed with a male reflexive pronoun '*himself*'.

3.2 Contextual Bias (B3)

In the proposed gender bias taxonomy (Hitti et al., 2019), contextual bias has two subtypes: societal bias, where one gender is stereotypically assigned a social role, and behavioural bias, where certain attributes or traits associated with one gender can lead to generalized gender stereotypes. For example, in Wordnet, for the word entries '*slovenly woman*' and '*rich man*', each prescribes a specific trait to a gender, while the connotations related to the adjectival modifiers appear different.

3.3 Additional Bias

We add two gender bias types to the taxonomy:

3.3.1 Distributional Bias (B4)

Distributional bias is concerned with the uneven distribution of different genders. For example, in OEWN 2024, the number of male names (6,778) is significantly greater than that of female ones (845) as shown in Table 4.

3.3.2 Definitional Bias (B5)

The different definitions given to male and female words implicate the differentiated representation

[3] 'needleman' is attested in the Oxford English Dictionary

Type	ID	Subtype	Example
Structural Bias	**B1**	Explicit Marking of Sex	*police*man: a member of a police force
	B2	Generic *he*	*researcher*: a scientist$_i$ who devotes himself$_i$ to doing research.
Contextual Bias	**B3**	Contextual Bias	(1) *slovenly* woman vs. *rich* man (2) He made an *honest* woman of her.
Additional Bias	**B4**	Distributional Bias	For OEWN 2024, 6,778 male names and 845 female names.
	B5	Definitional Bias	*horse*man: a man skilled in equitation *horse*woman: a woman horseman

Table 1: Taxonomy with types and subtypes of gender bias and examples. In the examples, red indicates male gender; blue female; green neutral. Mentions that refer to the same person are indicated by i. Examples in B1, B2, B3 (1) and B5 are the definitions of entries from WordNet. B3 (2) is from the example sentence in the word entry 'honest woman' in Wordnet.

of men and women in lexical resources, which we denote as 'definitional bias'. As shown in **B3** in Table 1, the definition given in PWN to '*horseman*' only refers to men and is detailed, whereas '*horsewoman*', in PWN, is defined solely based on the male version: '*horseman*'[4].

4 Open English Wordnet

Open English Wordnet (McCrae et al., 2019, 2020, OEWN) is a 'fork' of Princeton WordNet (Miller, 1995; Fellbaum, 2010, PWN) that aims to further develop a wordnet for English. The project is open-source and is hosted on GitHub. It has made 5 releases of its WordNet on an annual basis since 2019. These have been released in various formats including the original Princeton WordNet database format, allowing this resource to act as a 'drop-in' replacement for Princeton WordNet, as well as in the standard formats proposed by the Global WordNet Association (McCrae et al., 2021, GWA). The scope of this project has fixed a wide range of issues from simple typos up to identifying and merging duplicate synsets and introducing new synsets. In addition, the project aims to take feedback from other wordnet projects in other languages, such as plWordNet (Maziarz et al., 2016), and incorporate changes that are relevant to English. The project aims to evolve English Wordnet along the lines set up by Princeton WordNet, and to that extent has developed guidelines that describe how the English Wordnet can be constructed, for example on criteria for inclusion of terms or novel senses. Further, the project aims to keep the resource up-to-date with modern English, not only by including neologisms but also by ensuring that the resource matches modern lexicographic practices. It was in this context that the new guidelines on the use of gendered language have been developed.

[4]This definition along with a few others was changed in OEWN 2024 to be more balanced

	Phase 1	Phase 2
Hypernym Links for Missing Gender (1a)		
Male	0	61
Female	26	58
Both	84	0
New Synset from Neutral (1b)		
Male	13	13
Female	0	5
Both (1c)	17	209
New Synset from other Gender (1d)		
Male	26	0
Female	0	34

Table 2: Summary of changes to OEWN 2024 to fix structural bias. We show the number of new hypernym links created to add missing gender information as well as the new synsets created.

5 Methodology

We consider four kinds of bias and how we can mitigate these issues. Firstly, we consider structural bias in two aspects: the explicit marking of gender in terms such as 'mailman' and the usage of the male pronoun, 'he'. We then consider definitional bias in the length of definitions and finally, we consider distributional bias in the resource.

5.1 Structural Bias - Explicit Marking of Sex

In order to examine the structural biases in English Wordnet, we examined the usage of gendered words (**B1**). The first step was to compile a list of gendered words and this was done in two stages. In **Phase 1**, we extracted all gendered words in

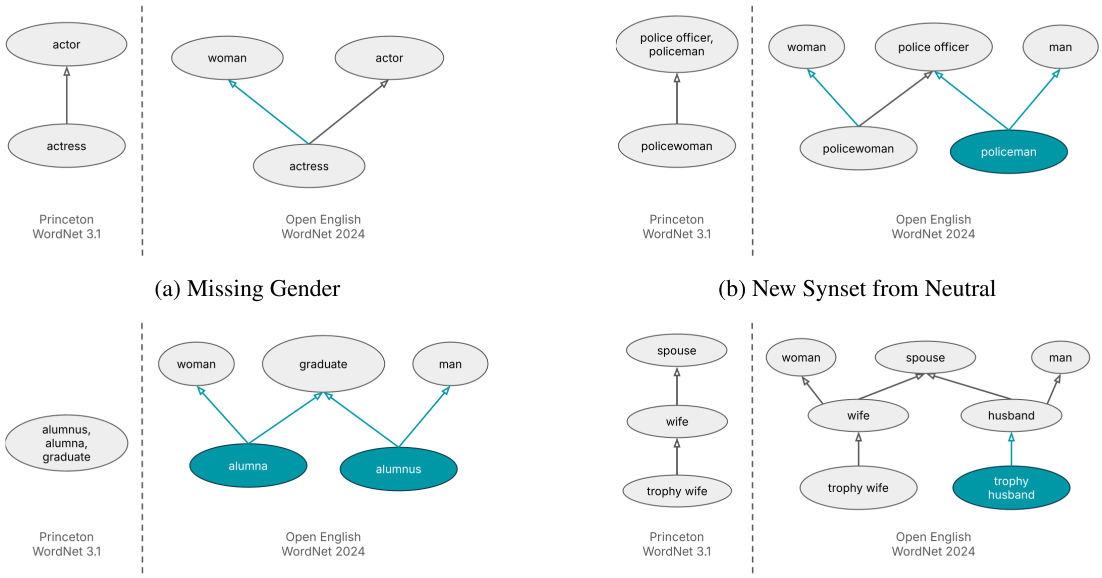

(a) Missing Gender

(b) New Synset from Neutral

(c) Both Synsets from Neutral

(d) New Synset from other Gender

Figure 1: Examples of structural biases found in English Wordnet and the proposed changes made in Open English Wordnet 2024. Changes made in Open English Wordnet are highlighted in blue.

English Wordnet based on a list in Wiktionary[5], which allowed us to identify words that have an unusual feminine morphological form (such as ‘ladette’ or ‘dominatrix’) or have no morphological clues (such as ‘bride’ or ‘wife’). For **Phase 2**, we used words that end in words known to be gendered such as ‘-man’ and ‘-woman’. This list was compiled by taking all the words in OEWN 2023 that were descended from the synset for ‘man’[10306910-n] or ‘woman’[10807146-n] with a few alterations[6]. For this second stage, we first manually classified whether the words represented a gender usage, as many terms like ‘German’[09767053-n] or ‘brahman’[09892023-n] were included. In total, 556/1092 words found this way were gendered terms. Then for each of these terms found, we examined the usage of the synset and how this can be changed to better represent gender. These can be divided into four main classes:

Missing Gender In these cases, the word was explicitly gendered but the direct hypernym was not gendered. We consider a word to be explicitly gendered if the definition contains a gendered term (such as ‘a man who’) or if all lemmas were gendered (for example by ending in ‘-man’). An example of this was ‘actress’[09787123-n], which is a hyponym of ‘actor’[09784701-n], which in this case was a gender-neutral term. For these cases, we simply added a second hypernym link to the synset for ‘woman’[10807146-n] to indicate the gender, as depicted in Figure 1a.

Synset from Neutral In this case, there was an exclusively (normally female-gendered) term and the hypernym included a term that was of the opposite gender. For example, ‘policewoman’[10468986-n] was a hyponym of a synset that included ‘police officer’ and ‘policeman’[10468557-n]. In order to fix this, we introduced a new synset that had ‘policeman’[80600405-n] as a member and was a hyponym of ‘man’[10306910-n]. In addition, both of the gendered synsets were given links to the appropriate gender term as for the previous case, as depicted in Figure 1b.

New Synsets of both Genders We also discovered that there were cases where both male-gendered and female-gendered terms occurred in the same synset or where both a male and a female variant of the term should be introduced. For example, there was a single synset with the terms ‘alumnus’, ‘alumna’ and ‘graduate’[09805779-n]. In this case, these are not true synonyms as you cannot substitute ‘alumnus’ in a sentence that uses

[5] https://en.wiktionary.org/w/index.php?title=Category:en:Female_people

[6] In particular, the words ‘ex’ and ‘cat’ were removed, although they had gendered usage.

'alumna'. As such, in this case, we introduced two new synsets ('alumnus'[80186365-n] and 'alumna'[86032704-n]) as well as links to gendered terms to arrive at the same modelling as the previous case, as in Figure 1c.

Synset from other Gender Finally, we encountered a number of terms that were clearly female-gendered but had no masculine-gendered term or vice versa. In these cases, we used corpus evidence and native-speaker intuition to decide if these terms had an equivalent and if so we introduced a new synset with the new term. In some cases, these terms were not generally linked to the synset for 'man'/'woman', but to another already gendered word. For example, we introduced 'trophy husband'[80620228-n][7] (from 'trophy wife'[10750477-n]) and linked it to 'husband'[10213586-n] as depicted in Figure 1d.

The summary of the changes is presented in Table 2 for both phases.

There were a number of challenges with this classification. Firstly, many gender-neutral terms have only become gender-neutral in recent usage, a particular example of this is 'actor' and 'actress', where the use of the female-specific term can be considered offensive (Duran, 2024). However, historically the term 'actor' only referred to males and referring to a woman as an 'actor' would have been offensive. In general, we chose to assume that terms that end in '-er'/'-or' are gender-neutral. One further exception to this was 'master', which may be used in a gender-neutral way, but also has some cases, where its use is gendered. In particular, 'schoolmaster' is marked as male in English Wordnet and in other dictionaries and would seem to be a gendered term, whereas other compounds (e.g., 'spymaster') are not gendered.

Another bias was found in that some male terms have more senses than their equivalent female terms. One particular example of this was 'viscountess'[10775729-n] defined as 'a wife or widow of a viscount' but there were two senses for 'viscount' as ' a son or younger brother or a count'[10775816-n] and 'a British peer who ranks below an earl and above a baron'[10775483-n] and, as such, a 'viscountess' should have had two senses depending on the kind of viscount the viscountess is a wife of. In general, it was noted that there were

[7]This is attested in the Cambridge Dictionary

	Male	Female
Gendered	839	157
Unbiased	37	3
Biased	181	2
Unclear	8	0
Incidental	4	2
Total	1,069	164

Table 3: Analysis of Pronoun Usage in OEWN 2023

far more male-gendered terms in previous versions of English Wordnet.

As part of the work, we introduced many new terms into OEWN 2024, and this led to some key considerations. Firstly, we must be sure that these new terms are significant in that they have documented usage, for this, we relied on looking in corpora (primarily CoCA (Davies, 2010)) and other dictionaries to find usages of these terms. This meant that in some cases, such as for 'hodman'[10199158-n], we did not introduce a female equivalent even though the job could be done by a woman and the female term can be easily derived morphosyntactically. Secondly, we also noted that the creation of terms, especially in the case of novel masculine-gendered terms, would project biases about women onto men. As such, we did not introduce some male-gendered variants of terms such as 'slovenly woman' and in most cases removed these terms, mostly due to them being compositional terms[8]. This change also eliminated some contextual bias (**B3**).

5.2 Structural Bias - Generic 'he'

Another major source of bias is the use of pronouns in definitions (**B2**), in particular, the usage of 'he'/'him'/'his'/'himself' and 'she'/'her'/'hers'/'herself'. We classified the usage into the following groups, **gendered** was usage that clearly referred to a person that identifies as male/female; **unbiased** usage, where the definition used an expression such as 'his or her' to indicate both genders; **biased** usage was for the case where a male or female pronoun was used with an ungendered noun. It should be noted that the two cases where a female pronoun was used indicated a contextual bias as they were used for the terms 'shopper'[10612003-n] and 'teacher' (in the definition

[8]So, word sense disambiguation algorithms could tag the individual words in the composition

of the verb 'shepherd'[02555865-v]), which are traditionally female professions. Finally, we saw eight cases where the reference was to a word that was potentially gendered. This included words such as 'pope', a role that can currently only be held by a man[9]. Finally, in a few cases, denoted as **incidental**, the most appropriate gender cannot be deduced from the context, primarily due to an example being given within the definition. The distribution of these is given in Table 3. We also note that there were substantially more usages of male pronouns reflecting the distributional biases in English Wordnet.

In order to fix the issues of gender bias related to the use of pronouns, there are the following strategies:

Generic 'he' The masculine pronoun has a history of usage without referring to gender (Wagner, 2003), and as such, one option would be to simply keep all the references to 'he' as is and remove other styles (such as below). This however does not reflect modern usage and so would not be appropriate.

He or she The solution already adopted in several entries in English Wordnet is to use both male and female pronouns ('he or she'). This has a number of issues, not only is it more wordy, but also the ordering of the pronouns still puts the male pronoun first. Ordering the female pronoun first would not be less biased and sounds unnatural[10]. Further, the use of 'he' or 'she' is exclusionary to non-binary people who use other pronouns.

(s)he This option only works for the plain form of the pronoun and has most of the disadvantages of the above option.

Alternation Another option would be to randomly use one of the two pronouns (or 'his/her' and 'her/his'). While this could be made unbiased so that both pronouns have equal distribution, this would not be apparent to users of the wordnet, who do not read all definitions.

Singular 'they' Singular 'they' has been used in English for a long time and is widely accepted as a gender-neutral pronoun (Balhorn, 2004). The principle disadvantage is that it is not approved of by some style guides (most notably *The Elements of Style* (Strunk and White, 1999)). However, most style guides prefer this and it avoids the disadvantages discussed above.

As such, in discussion with the community[11] of OEWN, we adopted the use of the singular 'they' for most examples. For the cases, where the gender was unclear, we rewrote the definition in a way that avoids the use of pronouns.

5.3 Definitional Bias

In English Wordnet, a potential source of biases is shorter definitions when describing females as opposed to males (**B5**). In some cases, it may be appropriate to use a shorter definition, for example, 'a female actor' as the definition of 'actress'[09787123-n] is an efficient and complete definition. In order to investigate whether female definitions in general are shorter, we needed to establish whether a particular synset referred exclusively to males, females or neutrally to both genders. Our method to do this was to consider all non-instance hyponyms of 'doer'[09786620-n] (including indirect hypernyms) and check whether they are also hyponyms of 'male'[09647338-n] or 'female'[09642198-n]. Unfortunately, this information was very incomplete in previous versions of English Wordnet and in a few cases produced synsets that were erroneously both male and female. A large number of fixes were made[12] to ensure that the hierarchy is correct, and in addition, the gender of words was inferred for existing synsets as part of the changes discussed in Section 5.1. As a result, we raised the number of gendered role words from 248 (101 male, 147 female and 8 erroneously in both[13]) in Princeton WordNet 3.1 to 1,186 in OEWN 2024. Of these, 395 synsets in PWN were gendered due to definition or lemmas but not explicitly marked as such using a hypernym. In Table 4, we present the average definition length in OEWN 2024 and PWN 3.1[14] in terms of words and characters. We see that

[9]Although there are claims of a historical female pope, this is believed to be apocryphal and this pope would have used male pronouns

[10]It sounds unnatural to the native Anglophone authors of this article. In CoCA 'his or her' has 17,285 occurrences against 472 for 'her or his' suggesting this is a widely-shared opinion

[11]https://github.com/globalwordnet/english-wordnet/issues/1058

[12]OEWN Issues: #1073, #1075, #1078-#1082

[13]According to PWN's hypernyms

[14]As PWN's gender assertions are unreliable, we decided the gender based on the OEWN 2024 synset that is aligned to

	Average Word Count	Average Characters	Number of Synsets	Number of Lemmas
Princeton WordNet 3.1				
Male Roles	9.61	51.50	309	489
Female Roles	8.08	42.83	324	523
Neutral Roles	9.52	55.11	6,023	9,683
Open English Wordnet 2024				
Male Roles	9.52	50.75	585	815
Female Roles	8.67	46.83	601	851
Neutral Roles	9.53	55.12	6,061	9,736
Male Names	13.98	91.74	2,722	6,778
Female Names	13.54	85.38	398	845

Table 4: Length of Definitions in OEWN 2024 and PWN 3.1

female definitions are generally shorter, however, OEWN has slightly longer definitions than PWN for females with about the same for male and neutral definitions. We also see that there are in both resources more female synsets than male and this is due to cases like ‘actress’, where there is no specific male term.

5.4 Distributional Bias

We also examined the definitions of named people in the resource (**B4**). Again there was no direct information for most synsets in English Wordnet about gender, except in a few cases where the named person was an instance of a gendered word like ‘queen’[10518940-n]. As such, we used a partial mapping to WikiData, where gender information is available (in particular using the P21 property)[15] and we present the analysis of the lengths and frequency of definitions based on the gender given in WordNet also in Table 4[16]. We see that while there is a huge distributional bias in the number of male figures mentioned in WordNet, the definitions are of similar length, with definitions of female figures being only slightly shorter in general.

6 Discussion

In this work, we have laid out the changes we have made to reduce gender bias in Open English Wordnet. These changes should be relevant to wordnets for other languages as well as for other lexicographic projects targeting English. In fact, such initiatives are common in lexicography, with similar initiatives as far back as the 1980s (White, 1989), and the Oxford English Dictionary was even the subject of an online petition (Pettini, 2021). Lexicographers have a dual role, both to record and reflect language as it is used, but also they can be ‘agents of change’ as reference works (Müller-Spitzer, 2023; Fuertes-Olivera and Tarp, 2022) in society. Open English Wordnet is a resource that reflects modern English with the purpose of enabling NLP applications such as word sense disambiguation, while also providing a structured organisation of words that can power psycholinguistic investigation including those that use large language models. As such, the resource needs to reflect the biases that are inherent in society, but avoid forcing more bias into its applications by, for example, forcing a sense to be disambiguated to a male synset, when the entity in the context is clearly female.

When applied to cross-linguistically, we hope that some of these fixes can be used to help with other languages. Firstly, the Open English Wordnet project will contribute these changes to the Collaborative Inter-lingual Index (Bond et al., 2016) so that they may be useful to other wordnets. In particular, this would be useful for languages that have mandatory gender for most role words (this applies to most European languages) as there are now synsets for gender-specific as well as gender-neutral terms. In this way, the German ‘Polizist’ could be linked through policeman[80600405-n] to terms such as

this PWN synset.

[15]Two named figures in English Wordnet have a non-binary identity and are excluded from this analysis

[16]We present only the results for OEWN 2024 as the total number of proper noun synsets and their definitions has not changed substantially

'policier' in French, while the feminine 'Polizistin' would be linked through 'policewoman'[10468986-n] to 'policière' in French. We note the changes made in OEWN 2024 did not remove any synsets, and added gender to a synset if its definition explicitly marked the synset as male or if all lemmas were male-gendered, which may differ from the policy of other wordnets constructed using the EXTEND methodology (Vossen, 1998). For role words, our changes created both gender-neutral and gender-specific synsets, unless we could not find attestations of the terms in English, and aimed to cover all such role words in OEWN, but some may be missed if they were not in Wiktionary and did not follow a typical derivational pattern (such as '-man'/'-woman').

Finally, we note that the change of gendered pronouns relies on a fairly unique case of English having an animate, gender-neutral pronoun, whereas most other languages either lack gendered pronouns entirely (such as 'hän' in Finnish) or gender-neutral pronouns are neologisms that are not widely used (such as 'elle' in Spanish).

7 Conclusion

This work has presented a number of changes to Open English Wordnet that have made the resource more applicable and more fair for modern usage of English. While we managed to mostly remove or limit the structural biases in English Wordnet, other biases still exist. In particular, distributional biases are still a major issue in the resource with male figures being named more than seven times as frequently as female figures and a similar frequency for the use of female pronouns versus male. It is hard to fix these in a way that reflects a world where women have traditionally been excluded from roles where they would acquire notoriety to be included in the resource. Further, OEWN has a current moratorium on the introduction or removal of proper nouns that would prevent this. We also note that this work has approached gender as a binary, which excludes many non-binary people, and there is scope for improving this in future releases.

Acknowledgements

John P. McCrae is supported by Research Ireland under Grant Number SFI/12/RC/2289_P2 Insight_2, Insight SFI Centre for Data Analytics and Grant Number 13/RC/2106_P2, ADAPT SFI Research Centre.

References

Mark Balhorn. 2004. The Rise of Epicene They. *Journal of English Linguistics*, 32(2):79–104.

Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings.

Francis Bond, Piek Vossen, John McCrae, and Christiane Fellbaum. 2016. CILI: the collaborative interlingual index. In *Proceedings of the 8th Global WordNet Conference (GWC)*, pages 50–57, Bucharest, Romania. Global Wordnet Association.

Shikha Bordia and Samuel R. Bowman. 2019. Identifying and reducing gender bias in word-level language models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 7–15, Minneapolis, Minnesota. Association for Computational Linguistics.

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

Jenna Cryan, Shiliang Tang, Xinyi Zhang, Miriam J. Metzger, Haitao Zheng, and Ben Y. Zhao. 2020. Detecting gender stereotypes: Lexicon vs. supervised learning methods. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.

Mark Davies. 2010. The Corpus of Contemporary American English as the first reliable monitor corpus of English. *Literary and Linguistic Computing*, 25(4):447–464.

Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. 2020. Queens are powerful too: Mitigating gender bias in dialogue generation.

Jad Doughman and Wael Khreich. 2022. Gender bias in text: Labeled datasets and lexicons. *CoRR*, abs/2201.08675.

Jad Doughman, Wael Khreich, Maya El Gharib, Maha Wiss, and Zahraa Berjawi. 2021. Gender bias in text: Origin, taxonomy, and implications. In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 34–44, Online. Association for Computational Linguistics.

Crystal Duran. 2024. Actor vs. actress. *Backstage*.

Zahra Fatemi, Chen Xing, Wenhao Liu, and Caiming Xiong. 2023. Improving gender fairness of pre-trained language models without catastrophic forgetting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1249–1262, Toronto, Canada. Association for Computational Linguistics.

Christiane Fellbaum. 2010. *WordNet*, pages 231–243. Springer Netherlands, Dordrecht.

Pedro A Fuertes-Olivera and Sven Tarp. 2022. Critical lexicography at work: reflections and proposals for eliminating gender bias in general dictionaries of Spanish. *Lexikos*, 32(2):105–132.

Yasmeen Hitti, Eunbee Jang, Ines Moreno, and Carolyne Pelletier. 2019. Proposed taxonomy for gender bias in text; a filtering methodology for the gender generalization subtype. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 8–17, Florence, Italy. Association for Computational Linguistics.

Lucy Li, Demszky Dorottya, Bromley Patricia, and Jurafsky Dan. 2020. Content analysis of textbooks via natural language processing: Findings on gender, race, and ethnicity in texas u.s. history textbooks. *AERA Open*, 6.

Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. *ArXiv*, abs/1903.10561.

Marek Maziarz, Maciej Piasecki, Ewa Rudnicka, Stan Szpakowicz, and Paweł Kędzia. 2016. plWordNet 3.0 – a comprehensive lexical-semantic resource. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2259–2268, Osaka, Japan. The COLING 2016 Organizing Committee.

John P. McCrae, Michael Wayne Goodman, Francis Bond, Alexandre Rademaker, Ewa Rudnicka, and Luis Morgado Da Costa. 2021. The GlobalWordNet formats: Updates for 2020. In *Proceedings of the 11th Global Wordnet Conference*, pages 91–99, University of South Africa (UNISA). Global Wordnet Association.

John P. McCrae, Alexandre Rademaker, Francis Bond, Ewa Rudnicka, and Christiane Fellbaum. 2019. English WordNet 2019 – an open-source WordNet for English. In *Proceedings of the 10th Global Wordnet Conference*, pages 245–252, Wroclaw, Poland. Global Wordnet Association.

John Philip McCrae, Alexandre Rademaker, Ewa Rudnicka, and Francis Bond. 2020. English WordNet 2020: Improving and extending a WordNet for English using an open-source methodology. In *Proceedings of the LREC 2020 Workshop on Multimodal Wordnets (MMW2020)*, pages 14–19, Marseille, France. The European Language Resources Association (ELRA).

George A. Miller. 1995. Wordnet: a lexical database for english. *Commun. ACM*, 38(11):39–41.

Carolin Müller-Spitzer. 2023. Gender stereotypes in dictionaries: the challenge of reconciling usage-based lexicography with the role of dictionaries as social agents. *Lexikos*, 33(2):79–94.

Silvia Pettini. 2021. "One is a woman, so that's encouraging too". the representation of social gender in "powered by Oxford" online lexicography. *Lingue e Linguaggi*, (44):275–295.

Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. Gender bias in machine translation. *Transactions of the Association for Computational Linguistics*, 9:845–874.

Sabine Sczesny, Magda Formanowicz, and Franziska Moser. 2016. Can gender-fair language reduce gender stereotyping and discrimination? *Frontiers in Psychology*, 7.

William Strunk and E. B. White. 1999. *The Elements of style*. Allyn and Bacon : Longman, Boston (USA).

Piek Vossen. 1998. *Introduction to EuroWordNet*, pages 1–17. Springer Netherlands, Dordrecht.

Susanne Wagner. 2003. *Gender in English pronouns: Myth and reality*. Ph.D. thesis, Freiburg (Breisgau), Univ., Diss., 2003.

Linda White. 1989. Feminism and lexicography: Dealing with sexist language in a bilingual dictionary. *Frontiers: A Journal of Women Studies*, pages 61–64.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.

Haotian Zhu, Kexin Gao, Fei Xia, and Mari Ostendorf. 2024. Disagreeable, slovenly, honest and un-named women? investigating gender bias in English educational resources by extending existing gender bias taxonomies. In *Proceedings of the 5th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 219–236, Bangkok, Thailand. Association for Computational Linguistics.