

Objectifying Women? A Syntactic Bias in French and English Corpora.

Yanis da Cunha, Anne Abeillé

Laboratoire de Linguistique Formelle, Université Paris Cité
8, Rue Albert Einstein 75013 Paris
yanis.dc@gmail.com, anne.abeille@u-paris.fr

Abstract

Gender biases in syntax have been documented for languages with grammatical gender for cases where mixed-gender coordination structures take masculine agreement, or with male-first preference in the ordering of pairs (Adam and Eve). On the basis of various annotated corpora spanning different genres (fiction, newspapers, speech and web), we show another syntactic gender bias: masculine pronouns are more often subjects than feminine pronouns, in both English and French. We find the same bias towards masculine subjects for French human nouns, which then refer to males and females. Comparing the subject of passive verbs and the object of active verbs, we show that this syntactic function bias is not reducible to a bias in semantic role assignment since it is also found with non-agentive subjects. For French fiction, we also found that the masculine syntactic function bias is larger in text written by male authors – female authors seem to be unbiased. We finally discuss two principles as possible explanations, ‘Like Me’ and ‘Easy first’, and examine the effect of the discourse tendency for men being agents and topics. We conclude by addressing the impact of such biases in language technologies.

Keywords: corpus, gender bias, syntactic function, French, English, syntax, treebank

1. Introduction

Gender biases have been documented at various levels of grammar in various languages. Among others, there are biases in favor of the masculine in agreement, where coordinations of mixed genders generally trigger masculine agreement across languages (Corbett, 1983). Despite the possibility of closest conjunct agreement (An and Abeillé, 2021), masculine controllers therefore have a privileged status in agreement patterns. In word order, men generally appear before women in binomials in English (Mollin, 2013; Mollin, 2014) and French (mari et femme ‘husband and wife’, frères et sœurs ‘brothers and sisters’, Abeillé et al. (2018)). However, some reversals are attested (aunts and uncles, mother and father, (Goldberg and Lee, 2021)). Experiments on English showed that a men-first bias can also occur in sentence production (Brough et al., 2020). For semantic roles, psycholinguistic experiments on French and German showed that it is more expected for men to be agents than for women (Esaurova and Von Stockhausen, 2015). There is thus converging evidence, from experiments and corpora, that gender stereotypes and biases can affect linguistic patterns.

This paper aims to shed light on another type of gender bias, which affects syntactic function: men are more likely to be a syntactic subject than women. Such a bias has been noticed in examples used linguistics papers, both in English (Kotek et al., 2021; Cépeda et al., 2021) and in French (Richy and Burnett, 2020). For instance, in the linguistic examples of the French journal *Langue Française* (1969-1971 and 2008-1017), Richy and Burnett (2020) show that women represent 12% of subjects and 30% of objects, while men represent 88% of subjects but 70% of objects. This difference

was significant, and year of publication and author gender did not play a role, suggesting that this gender bias is stable across time and authors. However, it can be asked whether this bias is specific to linguists’ usage or whether it is a more general trend.

Such an effect of gender is reminiscent of the effect of animacy, definiteness, person or pronominality on syntactic functions. Studies from formalist (Aissen, 1999; Aissen, 2003; Jelinek and Carnie, 2003), typological (Haspelmath, 2021) and psycholinguistic (MacDonald, 2013; Lamers and De Swart, 2011) perspectives have shown that function coding is driven by hierarchically ordered information generally characterized as ‘prominence features’. Such features can be represented in the form of scales (exemplified in 1), where $>$ means ‘more prominent than’.

- (1)
 - a. Animate $>$ Inanimate
 - b. Definite $>$ Indefinite
 - c. Pronoun $>$ Noun
- (2) Subject $>$ Object

Scales like these formalize the fact that prominent referents are more likely to occupy more prominent functions (subjects) and less prominent referents tend to occupy less prominent functions (objects). Thus prominence scales in (1) tend to align with the syntactic function scale in (2). We will refer to animate, definite and pronominal subjects as ‘aligned configurations’ (prominent referents with a prominent function), while inanimate, indefinite and nominal subjects would be ‘unaligned configurations’. This general effect shows up in two ways across languages (Bresnan et al., 2001). On one hand, prominence scales can induce strong

grammaticality contrasts, bringing into play differential argument coding or obligatory voice alternations. For example, in Spanish or Hindi, animate objects have to be coded with an extra case marker, because animate objects do not represent an aligned configuration. On the other hand, prominence scales can induce production and processing preferences, making aligned configurations easier to predict and more frequent in corpora. Therefore, it has been noted in various languages that animate patients are more likely to be used as passive subjects, to favor an aligned animate-subject configuration, and avoid an unaligned animate-object configuration in active voice (for animacy effects in active/passive alternation see Tanaka et al. (2011) for Japanese, Hundt et al. (2021) for English, Thuilier et al. (2021; da Cunha and Abeillé (2020) for French).

Finding syntactic function gender biases in corpora would thus provide evidence for integrating gender information as a prominence feature, as suggested by Esaulova and Von Stockhausen (2015). This would have consequences for psychological and typological studies, where gender would have to be taken into account for its possible effects on syntactic patterns, but it would also highlight the importance of gender biases in language for language technologies. (Wisniewski et al., 2021; Sun et al., 2019; Costa-jussà, 2019; Brown et al., 2020).

Our first goal is thus to replicate the findings that men are more likely to be subjects in linguistic examples (Cépeda et al., 2021; Richy and Burnett, 2020; Kotek et al., 2021) for more genres : newspapers, fiction, speech and web language. In a second step, we will take into account semantic roles, topicality and author gender to explore possible explanations for a gender bias in syntactic function. We will then discuss on the possible integration of gender among prominence features.

2. Methodology

We aim to detect and compare syntactic function gender biases across languages and genres. We selected corpora both in French and English. In French, we used the French TreeBank (FTB) for the journalistic genre (Abeillé et al., 2019), using a version annotated for expletive subjects (Candito et al., 2014). For spontaneous speech, we used three corpora from the Orféo project (Benzitoun and Debaisieux, 2020), namely the CFFP, the CRFP and the C-Oral-Rom. For fiction, we selected novels from contemporary Frantext (ATILF, 2022) and for web French we used FrWac (Baroni et al., 2009). For English, we used the Universal Dependencies (UD) corpora annotated for genre. We selected the English Web Treebank (EWT) (Silveira et al., 2014), the Georgetown University Multilayer (GUM) corpus (Zeldes, 2017) (which contains various genres, among others : fiction, news, conversation, interviews...), the LinES corpus (Ahrenberg, 2015), from which we only kept literature, and finally the English portion of the Parallel Universal Dependencies (PUD)

corpus (McDonald et al., 2013), from which we only kept news. The table 1 summarizes some information about the selected corpora.

Language	Genre	Corpus	Period
French	Newspapers	FTB	1900-1993
	Speech	Orféo	1994-2012
	Fiction	Frantext	1980-2021
	Web	FrWac	2010
English	Web	EWT	1999-2011
	Varied	GUM	2000-2020
	Fiction	Lines	1899-1998
	News	PUD	—

Table 1: Selected corpora

We extracted all subjects and objects from dependency-annotated corpora (FTB, Orféo corpora, UD English corpora). For French, we kept singular nouns and clitic pronouns (*il* ‘he/it’, *elle* ‘she/it’, *le* ‘him/it’, *la* ‘her/it’), for English just singular pronouns (*he*, *him*, *she*, *her*, *it*). For FrWac and Frantext, which have no dependency annotation, we took sequences defined as : *no preposition + determiner + noun + conjugated verb + determiner + noun*. The *no preposition* condition filters out examples such as (3), where a preverbal noun (here, SG) is not a subject. This allows us to assume that preverbal nouns are subjects and postverbal ones are objects, as in (4). From Frantext and FrWac, we also took a sample of singular clitic pronouns (*il*, *elle*, *le*, *la*), whose form already indicates their syntactic function.

- (3) [...] le président de la SG [Société Générale] écarte l’idée d’un rapprochement avec BNP [Banque Nationale de Paris] Paribas (FrWac, efinancialcareers.fr)
‘The president(MASC) of the SG(FEM) rules out the idea(FEM) of a merger with BNP Paribas’
- (4) a. Votre fils apprendra la voltige (CHANDER-NAGOR Françoise, L’Enfant des Lumières, 1995, Frantext)
‘Your son(MASC) will learn aerobatics(FEM)’
b. La confédération assure le cadre permanent de discussion et d’action [...] (FrWac, gauchepopulaire.fr)
‘The confederation(FEM) provides the permanent framework(MASC) for discussion and action’

The FTB annotation allows us to filter out expletive subjects *il* and predicative complements. For other French corpora, we removed the most frequent impersonal and predicative verbs according to the FTB : *falloir* ‘to be necessary’, *être* ‘to be’, *rester* ‘to remain’, *devenir* ‘to become’, *sembler* ‘to seem’, *paraître* ‘to look like’. We only kept singular nouns, to avoid mixed-gender and generic forms. To do so, we removed lemmas whose token contains an additional *-s*,

which is a plural marker in French. With regard to gender annotation, the situation is different for nouns and pronouns. French and English pronouns provide grammatical and social gender information respectively. For French nouns, the FTB is already annotated for grammatical gender. For the other corpora, we annotated grammatical gender for all nouns using information available in Flexique (Bonami et al., 2014), a French dictionary which provides grammatical gender information for 31 000 nouns. Finally, we annotated animacy (human, animate, inanimate) using an animacy-annotated version of Flexique (Bonami, p.c.). The table 2 shows our annotated data set. For French nouns, 73% of the whole data set has been annotated for grammatical gender and 70% for animacy (human vs inanimate nouns). Only the annotated data is reported there. It can be seen that English has much fewer data points than French, but excepts for FrWac (web) and Frantext (fiction), this is due to corpus size.

3. Results

3.1. Syntactic Function Bias across Genres

We first report results for English, in figure 1. Masculine bias can be seen in two ways. First, it appears that masculine pronouns are always more frequent than feminine ones, independently of syntactic function. For example, fiction contains 112 masculine pronouns but only 57 feminine ones. This imbalance is consistent across genres, but less strong in speech. Secondly, aside from being rarer, feminine pronouns also appear more often as objects than masculine one. This can be seen by the height of the orange areas. We can also see that within objects 'it', the inanimate pronoun, is the most frequent, followed by feminine pronouns and finally masculine pronouns. Masculine pronouns are thus more often subjects (height of the blue areas). The less biased genre seems to be speech, where feminine and masculine pronouns are almost equally frequent, and where there does not seem to be a syntactic function bias. We thus see a tendency for masculine pronouns to be subjects across genres, generalizing previous results found in linguistic examples to the whole of the English language (Kotek et al., 2021; Cépeda et al., 2021). We can compare these results with those for French clitic pronouns, in figure 2. We find again that masculine pronouns are more frequent than feminine ones, but the bias for masculine pronouns to be subjects does not appear as clearly. All genres show slightly more masculine pronouns as subjects, except for newspapers where the bias is reversed. The main problem here is that contrary to English pronouns, French pronouns do not reflect social gender but grammatical gender. As a consequence, French pronouns are not specified for animacy, they may either refer to humans or to inanimates, as in (5). If social gender plays a role in syntactic function assignment, it would do so only for humans, where grammatical gender is interpreted as social gender in most cases (Richy and Burnett, 2021).

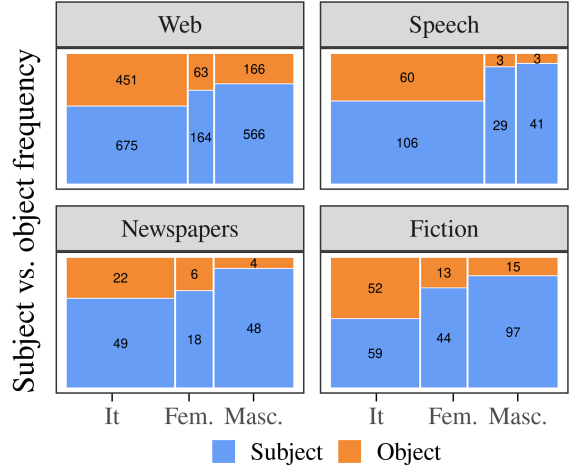


Figure 1: Gender and function frequencies for English personal pronouns




Figure 2: Gender and function frequencies for French clitic pronouns

- (5) Mais je veux dire la gestion_i de la ville est relativement bonne [...] Elle_i correspond au type de population qui réside à Montreuil.
 'But I mean, the management(FEM)_i of the city is relatively good. It(FEM)_i corresponds to the type of population that lives in Montreuil'

To reduce noise, we can look at the subject bias for nouns, for which we have animacy and grammatical gender information. Figure 3 shows the frequency in subject function by gender and animacy of French nouns. Error bars indicate standard error. We can now see a difference between inanimate and human nouns. On one hand, feminine and masculine inanimates do not differ in their frequency as subject (only fiction texts show a slightly greater frequency for feminine noun subjects). On the other hand masculine humans, ie. men, show a bias toward subject function and are thus more often subjects than feminine humans, ie. women. This difference between masculine and femi-

	Fiction	Newspapers	Web	Speech	Total
English pronouns	280	147	2085	242	2754
French pronouns	9380	2329	53093	23514	88316
French nouns	20444	16489	115059	21862	173854
Total	30104	18965	170237	45618	264924

Table 2: Composition of the studied sample

nine human nouns can be seen for each genre except for fiction, where the bias is less strong. This result indicates that gender in French does indeed have a different impact on syntactic function use for inanimate and human nouns, since it matters for the latter but not the former (as supported by Richy and Burnett (2021) among others). As inanimate nouns are not biased for syntactic function, it corroborates the idea that their grammatical gender is not interpreted in the same way as the grammatical gender of human nouns. However we do not claim that grammatical gender of inanimate nouns could not be interpreted at all (see Williams et al. (2021 03 17) for a discussion). For what concerns syntactic function, it is clear that French human nouns show the same pattern as English pronouns, where masculines are more frequently subjects, which is evidence for a similar impact of social gender.

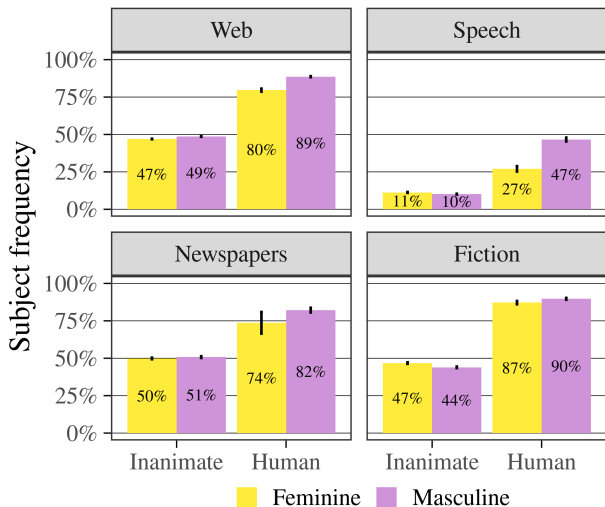


Figure 3: Subject frequency of French singular nouns according to grammatical gender and animacy

There appears to be a bias for men being subjects more frequently than women, both in French (human nouns) and in English (human pronouns). We can now compare these two languages across genres. Figure 4 shows the strength of masculines-as-subjects bias in the four genres we studied. Our bias measure corresponds to the difference between masculine subject frequency and feminine subject frequency. As a consequence, the greater this difference is, the more men are found as subjects compared to women. For example, English fiction has a bias of 10 points. So, masculine subject

frequency in English fiction (86%) is 10 points higher than feminine subject frequency (76%). We established this measure to allow easier comparisons of gender biases in syntactic functions between genres and languages.

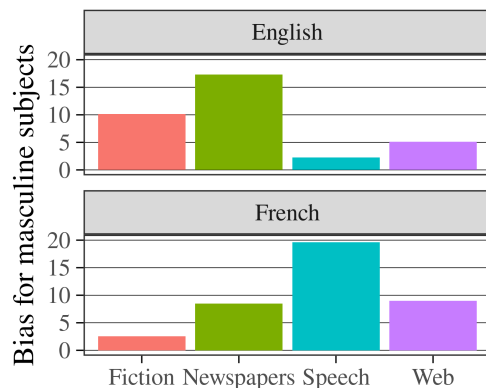


Figure 4: Bias for masculine subjects across genres in English (personal pronouns) and French (human nouns). The bias measure indicates to what extent masculines are more frequently subjects than feminines.

We show that English narrative genres (fiction and newspapers) have a rather strong bias for masculine subjects (10 points or more). Interactive genres (speech, web) are less biased (less than 5 points). In French such a generalization seems not to hold : newspapers and web both show a moderate bias (about 8 points) and speech shows the strongest one (20 points). This time, fiction shows no bias (just 2.5 points). Although there seems to be a bias across genres, we do not see a clear link between the type of genre and the strength of its masculine subject bias. A general conclusion we can draw is that the type of function bias noticed in linguistic examples in French and English papers (Richy and Burnett, 2020; Kotek et al., 2021; Cépéda et al., 2021) is not genre-specific but reflects a general trend in other genres both in English and French.

Through this method, we can ask whether this bias for masculine subjects is due to syntax and/or to other factors such as semantic role and discourse. Indeed, subjects are prototypical agents and topics while objects are prototypical patients, which could explain the syntactic biases we observe. To investigate this, we looked at semantic roles and pronominalization rate of mascu-

line and feminine subjects and objects.

3.2. Syntactic Bias and Semantic Roles

We aim to investigate whether gender biases remain after taking semantic roles into account. To probe this, we compare subjects of passive verbs and objects of active verbs (which bear the same patient-like roles) with subjects of active verbs (which are more agent-like). Assuming that active objects and passive subjects may both bear the same patient-like roles, observing gender to have a differential effect on the two syntactic functions conditional on the same semantic role would indicate that the gender bias goes beyond semantics. In the case of the active/passive alternation, the difference between subjects and objects is indeed more closely related to syntax, and information structure, than to semantics.

For this part of the study, we only consider at corpora annotated for passive : FTB and English UD corpora. As we lack data for human feminine nouns in French (only 6 passive subjects), we report all data for French nouns and pronouns, including data points for which animacy was not provided by Flexique. We may take animacy into account for future research.

Figure 5 summarizes our results. It shows that both in English and French, objects are more often feminine and active subjects are more often masculine. Passive subjects, which share syntactic properties with active subjects and semantic properties with objects, are an in-between case : they are more often masculine than objects, but less often than active subjects. The difference between objects and subjects in general echoes the differences seen in the previous section (3.1). It is to be noted that English pronoun gender represents social gender (men vs. women), while for French pronouns and nouns, gender is grammatical gender, which is correlated but not equivalent to social gender. We hypothesize that the difference observed in French is due to social gender, but we leave the testing of this hypothesis for future work.

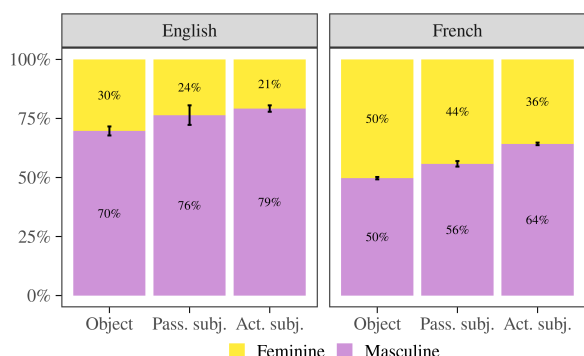


Figure 5: Proportions of masculine and feminine referents according to their syntactic function for English pronouns and French FTB pronouns and nouns, comparing transitive active and intransitive passive.

These results bring new evidence for a gender bias in syntactic functions which is not completely reducible to semantic roles. Indeed, even if objects and passive subjects both bear patient-like roles, we still observe a bias for masculine subjects, suggesting that the syntactic function bias is not due to semantics only. We showed that there is a superadditive effect between semantics and syntax: active subjects are even more often masculine than passive subjects. So there is also a bias for masculines to be agents. This result is consistent with previous literature based on linguistic examples (Kotek et al., 2021; Richy and Burnett, 2020; Cépeda et al., 2021), and we now show that it holds in other genres.

We thus showed that with constant semantic role, a bias for masculine subjects still appears. Gender biases we observed in syntactic functions cannot be explained only by a discourse tendency for men to be agents more frequently than women.

3.3. Syntactic Bias and Topicality

Another factor we now have to explore is topicality, which could also explain the observed pattern. Topicality can be assessed in various ways. We adopt here a definition in terms of topic-worthiness (Dalrymple and Nikolaeva, 2011) or Topic Accessibility Scale (Lambrecht, 1996), that is to say the likelihood of being a good topic candidate. One of the criteria for topicality is being a pronoun, since pronouns encode active referents in the discourse universe. In linguistic examples, it has been found that men are more often referred to by pronouns than women (Richy and Burnett, 2021; Cépeda et al., 2021; Kotek et al., 2021). We thus looked at the pronominalization rate of masculine and feminine referents across genres. Figure 6 presents our results.

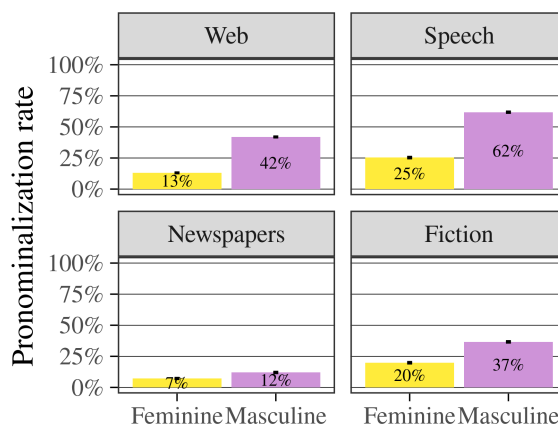


Figure 6: Pronominalization rate for masculine and feminine referents in French corpora.

We show that in the four genres under consideration, feminine referents are less often coded as pronouns than masculine ones. The less biased genre are newspapers, but French newspapers in general use fewer pro-

nouns (Poiret and Liu, 2020). On the contrary, speech shows the greatest difference, because spoken language uses more pronouns (*Ibid.*). The main consequence of this imbalance would be that masculine referents are more often topics than feminine referents. Yet, as subjects are canonical topics (Lambrecht, 1996; Givón, 1983), it would imply that subjects are more often masculine, and that the syntactic function gender bias may be reduced to this. We discuss such a hypothesis in the last section (4)

3.4. Syntactic Bias and Speaker Gender

Finally we investigate whether speaker gender plays a role: is masculine subject bias a male speaker tendency, like a *‘Me-First’ principle* (Cooper and Ross, 1975) or a *‘Like Me’ effect* (Brough et al., 2020) ? or is it a more general bias shared by male and female speakers ? For English linguistic examples, (Kotek et al., 2021) showed that author gender plays a role in gender biases, but not for French linguistic examples (Richy and Burnett, 2020). A *‘Like Me’ effect* in gender biases in syntax would thus constitute another type of explanation for the observed pattern.

Here, we only look at the data from the Frantext corpus (French, fiction), which is annotated for speaker gender. We aim to see whether the syntactic bias for masculine subjects highlighted until now depends on the gender of the speaker. Figure 7 reports our results.

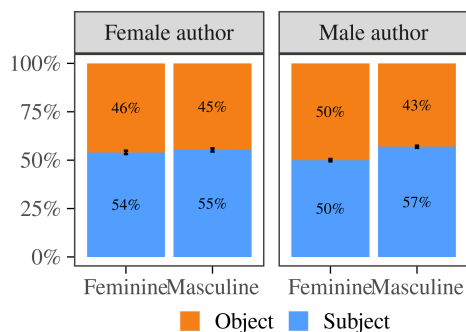


Figure 7: Proportions of subjects and objects according to speaker and noun gender in French fiction.

Whereas female authors do not show a difference between masculine and feminine nouns for subject and object frequencies (more or less 55% for both), male authors do show a difference. For them, masculine nouns are more often subjects (57%) than feminine ones (50%). Thus, a syntactic bias for masculine subjects seems to hold only for male authors here.

We analyze these data to test significance of the interaction between author and noun gender with a logistic regression model (package *lme4* on R (Bates et al., 2014)). We use function as predicted variable ($Subject = 1$, $Object = 0$), author and noun gender as predictors (which we normalized), with their interaction, and noun lemma and author as random variables, with ran-

dom intercepts only. Table 3 presents our results. We find a significant effect of noun gender ($E = 0.11$; $SE = 0.03$; $p < 0.001$) : masculine nouns are more likely to be subjects than feminine ones. We do not find an effect for author gender ($p > 0.05$) but there is a significant interaction between author and noun gender ($E = 0.04$; $SE = 0.02$; $p < 0.05$). The effect of noun gender thus significantly interacts with author gender : noun gender only matters for male authors, who use masculine subjects more frequently.

One interesting consequence of this result is that it partially corresponds to a *‘Like-Me’ effect*. Indeed men do tend to use masculines as subjects. But why don’t women use more feminines as subjects ? It would be interesting to study this type of interaction between speaker gender and gender syntactic biases for other languages and genres, taking into account animacy.

4. Discussion & conclusion

We found a gender bias in syntactic functions in both English and French across different genres: female referents (French human nouns and English pronouns) are less likely to be subjects than male referents. In French, we showed how this bias interacts with animacy, since grammatical gender has an effect only for human referents. We saw that the strength of the masculine bias for subjects is not clearly linked to genre characteristics (narrative, interactive etc).

We also explore two possible explanations for this bias : if men are more often subjects, it would come from other properties of subjects, like being canonical agents and topics. Discourse tendencies for men to be agents and topics would then be a source for syntactic biases. We showed that, although masculine referents are indeed more often agents, the syntactic bias goes beyond semantics, since it holds even when semantic roles are kept constant. If one considers only patientive referents (objects and passive subjects), a bias towards masculine subjects remains. For topicality, we found that feminine referents are indeed less referred to via pronouns. As pronouns encode active referents, they are more topical, and thus more often found as subjects, the canonical topics. Controlling for available topics in a text would be useful to corroborate this hypothesis, in a similar way to what Huet et al. (2013) did for French newspapers. Huet et al. (2013) showed that in the French journal *Le Monde* (from which the FTB, used in our study, was extracted), in 1985 (five years before the FTB), only 10% of the articles mentioned women, while 50% of them mentioned men (Huet et al., 2013). If most human referents in a text are men, it is not surprising to find them more often as subjects, since human subjects are canonical topics/agents. Nevertheless, we observe the same type of bias in other genres, including speech, which may not have the same referents as newspapers.

Finally, we investigate the possibility that gender biases are due to a kind of *‘Me-First’ principle* or *‘Like Me’*

	Estimate	Std. deviation	<i>z</i> -value	<i>p</i> -value
Intercept	0.18	0.03	5.59	< 0.001
Masculine vs. feminin noun	0.11	0.03	3.82	0.00014
Male vs. female author	-0.03	0.02	-1.33	0.18282
Interaction	0.05	0.02	2.93	0.00334

Table 3: Logistic regression modeling syntactic function with the interaction between speaker and noun gender (number of data points = 21 995).

effect (Cooper and Ross, 1975; Brough et al., 2020), which makes speakers produce/process referents they identify with more easily. In French fiction, which was the least biased genre in our cross-genre comparison (Figure 4), female and male authors behave differently. Indeed, only male authors exhibit a bias for masculine subjects, which supports a general idea of a ‘*Like Me*’ *effect* (Brough et al., 2020). However, women showed a rather unbiased usage in our data, casting doubt on this conclusion.

More generally, among semantic roles, topicality or ‘*Like Me*’ effects, disentangling syntactic biases from other kinds of gender bias will be necessary to find explanations for them.

Therefore our work extends literature on gender biases in syntax, showing that it holds across genres. It also opens the question of whether the discourse tendency of masculine subjects and feminine objects could be formalized into a gender prominence scale like that of animacy, definiteness or person prominence. It would then be take the form of the following scale (6), which tend to be aligned with the syntactic function scale.

- (6) Masculine > feminine
Subject > object

The gender bias we found thus seems comparable to other preferences in function assignment. These preferences can be summarized by the *Easy first* principle, which states that referents “important or conceptually salient to the speaker” and “more easily retrieved from memory” tend to appear earlier or as subjects in a sentence (MacDonald, 2013, 3). Investigating whether male referents can be considered as “easier” for some speakers (taking speaker gender into account), would lead to a better understanding of the gender biases we found.

We finally point out that it’s important to detect gender biases of this kind since they have an impact on language technologies : they can be learned by neural models (Brown et al., 2020) and can yield to biases in NLP tasks such as automatic translation (Wisniewski et al., 2021; Sun et al., 2019; Costa-jussà, 2019).

5. References

Abeillé, A., An, A., and Shiraishi, A. (2018). L’accord de proximité du déterminant en français. *Discours*.

Revue de linguistique, psycholinguistique et informatique. A journal of linguistics, psycholinguistics and computational linguistics, (22).

Abeillé, A., Clément, L., and Liégeois, L. (2019). Un corpus arboré pour le français: le french treebank. *Traitement Automatique des Langues*, 60(2):19–43.

Ahrenberg, L. (2015). Converting an english-swedish parallel treebank to universal dependencies. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 10–19.

Aissen, J. (1999). Markedness and subject choice in optimality theory. *Natural Language & Linguistic Theory*, 17(4):673–711.

Aissen, J. (2003). Differential object marking: Iconicity vs. economy. *Natural Language & Linguistic Theory*, 21(3):435–483.

An, A. and Abeillé, A. (2021). Closest conjunct agreement with attributive adjectives. *Journal of French Language Studies*, pages 1–28. Publisher: Cambridge University Press.

ATILF. (2022). Base textuelle frantext (online). Accessed: 2022-04-05.

Baroni, M., Bernardini, S., Ferraresi, A., and Zanchetta, E. (2009). The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, 43(3):209–226.

Bates, D., Mächler, M., Bolker, B., and Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.

Benzitoun, C. and Debaisieux, J.-M. (2020). Orféo: un corpus et une plateforme pour l’étude du français contemporain.

Bonami, O., Caron, G., and Plancq, C. (2014). Construction d’un lexique flexionnel phonétisé libre du français. In *SHS Web of Conferences*, volume 8, pages 2583–2596. EDP Sciences.

Bresnan, J., Dingare, S., and Manning, C. D. (2001). Soft constraints mirror hard constraints: Voice and person in english and lummi. In *Proceedings of the LFG01 Conference*, pages 13–32. Stanford: CSLI Publications.

Brough, J., Branigan, H., Harris, L., and Rabagliati, H. (2020). The influence of race and gender on perspective-taking during language production.

Brown, T. B., Mann, B., Ryder, N., Subbiah, M.,

- Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., and Askell, A. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Candito, M., Perrier, G., Guillaume, B., Ribeyre, C., Fort, K., Seddah, D., and de La Clergerie, E. V. (2014). Deep syntax annotation of the sequoia french treebank. In *International Conference on Language Resources and Evaluation (LREC)*.
- Cooper, W. E. and Ross, J. R. (1975). World order. *Papers from the parasession on functionalism*, pages 63–111.
- Corbett, G. (1983). Resolution rules: agreement in person, number, and gender. *Order, concord and constituency*, pages 175–206. Publisher: Foris Publications Dordrecht.
- Costa-jussà, M. R. (2019). An analysis of gender bias studies in natural language processing. *Nature Machine Intelligence*, 1(11):495–496. Publisher: Nature Publishing Group.
- Cépeda, P., Kotek, H., Pabst, K., and Syrett, K. (2021). Gender bias in linguistics textbooks: Has anything changed since macaulay & brice 1997? *Language*. Publisher: Linguistic Society of America.
- da Cunha, Y. and Abeillé, A. (2020). L’alternance actif/passif en français: une étude statistique sur corpus écrit. *Discours. Revue de linguistique, psycholinguistique et informatique. A journal of linguistics, psycholinguistics and computational linguistics*, (27).
- Dalrymple, M. and Nikolaeva, I. (2011). *Objects and information structure*. Number 131. Cambridge University Press.
- Esaulova, Y. and Von Stockhausen, L. (2015). Cross-linguistic evidence for gender as a prominence feature. *Frontiers in Psychology*, 6. Publisher: Frontiers.
- Givón, T. (1983). *Topic continuity in discourse*. Amsterdam: John Benjamins.
- Goldberg, A. E. and Lee, C. (2021). Accessibility and historical change: An emergent cluster led uncles and aunts to become aunts and uncles. *Frontiers in psychology*, 12:1418. Publisher: Frontiers.
- Haspelmath, M. (2021). Role-reference associations and the explanation of argument coding splits. *Linguistics*, 59(1):123–174.
- Huet, T., Biega, J., and Suchanek, F. M. (2013). Mining history with le monde. In *Proceedings of the 2013 workshop on Automated knowledge base construction - AKBC '13*, pages 49–54. ACM Press.
- Hundt, M., Röthlisberger, M., and Seoane, E. (2021). Predicting voice alternation across academic englishes. *Corpus Linguistics and Linguistic Theory*, 17(1):189–222.
- Jelinek, E. and Carnie, A. (2003). Argument hierarchies and the mapping principle. In *In Festschrift for Jelinek*.
- Kotek, H., Dockum, R., Babinski, S., and Geissler, C. (2021). Gender bias and stereotypes in linguistic example sentences. *Language*. Publisher: Linguistic Society of America.
- Lambrecht, K. (1996). *Information structure and sentence form: Topic, focus, and the mental representations of discourse referents*, volume 71. Cambridge university press.
- Lamers, M. and De Swart, P. (2011). *Case, word order and prominence: Interacting cues in language production and comprehension*, volume 40. Springer Science & Business Media.
- MacDonald, M. C. (2013). How language production shapes language form and comprehension. *Frontiers in psychology*, 4:226.
- McDonald, R., Nivre, J., Quirnbach-Brundage, Y., Goldberg, Y., Das, D., Ganchev, K., Hall, K., Petrov, S., Zhang, H., Täckström, O., Bedini, C., Bertomeu Castelló, N., and Lee, J. (2013). Universal Dependency annotation for multilingual parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97, Sofia, Bulgaria, August. Association for Computational Linguistics.
- Mollin, S. (2013). Pathways of change in the diachronic development of binomial reversibility in late modern american english. *Journal of English Linguistics*, 41(2):168–203. Publisher: Sage Publications Sage CA: Los Angeles, CA.
- Mollin, S. (2014). *The (ir) reversibility of English binomials: Corpus, constraints, developments*, volume 64. John Benjamins Publishing Company.
- Poiret, R. and Liu, H. (2020). Some quantitative aspects of written and spoken french based on syntactically annotated corpora. *Journal of French Language Studies*, 30(3):355–380.
- Richy, C. and Burnett, H. (2020). Jean does the dishes while marie fixes the car: a qualitative and quantitative study of social gender in french syntax articles. *Journal of French Language Studies*, 30(1):47–72. Publisher: Cambridge University Press.
- Richy, C. and Burnett, H. (2021). Démêler les effets des stéréotypes et le genre grammatical dans le biais masculin: une approche expérimentale. *GLAD! Revue sur le langage, le genre, les sexualités*, (10).
- Silveira, N., Dozat, T., de Marneffe, M.-C., Bowman, S., Connor, M., Bauer, J., and Manning, C. D. (2014). A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.
- Sun, T., Gaut, A., Tang, S., Huang, Y., ElSherief, M., Zhao, J., Mirza, D., Belding, E., Chang, K.-W., and Wang, W. Y. (2019). Mitigating gender bias in natural language processing: Literature review. *arXiv preprint arXiv:1906.08976*.
- Tanaka, M. N., Branigan, H. P., McLean, J. F., and Pickering, M. J. (2011). Conceptual influences on word order and voice in sentence production: Evi-

- dence from japanese. *Journal of Memory and Language*, 65(3):318–330.
- Thuilier, J., Grant, M., Crabbé, B., and Abeillé, A. (2021). Word order in french: the role of animacy. *Glossa: a journal of general linguistics*, 6(1).
- Williams, A., Cotterell, R., Wolf-Sonkin, L., Blasi, D., and Wallach, H. (2021-03-17). On the relationships between the grammatical genders of inanimate nouns and their co-occurring adjectives and verbs. *Transactions of the Association for Computational Linguistics*, 9:139–159.
- Wisniewski, G., Zhu, L., Ballier, N., and Yvon, F. (2021). Biais de genre dans un système de traduction automatique neuronale : une étude préliminaire. In *Actes de la 28e Conférence sur le Traitement Automatique des Langues Naturelles. Volume 1: conférence principale*, pages 11–25.
- Zeldes, A. (2017). The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612.