Markers of translator gender: do they really matter?[1,2]

Miriam Shlesinger, Moshe Koppel, Noam Ordan and Brenda Malkiel

The summers spent with Arnt and colleagues from the Copenhagen Business School at the retreat in Skagen were as close to heaven as I'll ever get. In his enviably laid-back way, Arnt managed to mold us into an irrepressible think tank. We wrote articles, he produced grant proposals, and all of us reveled in long walks along the magnificent beaches of that picturesque town on the tip of Jutland. Our explorations of the "migratory dune," our visits to the "sunken church," and the wonderful evenings spent lolling in the spacious parlor of the mansion that became our home during those magical weeks – are etched in my memory and in my heart. Back in Copenhagen, it was Arnt who introduced me to Karen Blixen's home, and all the stories that go with it, and Arnt – never rushed, never appearing pressured – who took me to see the opera building and other highlights of Copenhagen architecture. It was also Arnt who set up CRITT and so many other cornerstones of collaborative research at CBS. And on an even more personal level, it was Arnt who did me the unforgettable honor of co-editing and publishing Interpreting Studies and Beyond – a very special volume of the Copenhagen Studies in Language. I feel privileged to have been given this opportunity to express my appreciation, in some small way, for the gift of having Arnt as my friend. (And I am grateful to my co-authors, who have not had the good fortune of meeting Arnt, for agreeing to dedicate our first collaborative article to Arnt Lykke Jakobsen.)

Miriam Shlesinger

[1] This research was partially supported by the Israel Science Foundation, grant no. 1180/06.

[2] We thank our colleague Jonathan Fine for his review of this article and for his valuable suggestions.

Abstract

Given the impressive record of machine learning in telling male- from female-authored texts in various genres, we asked whether the computer could also be "taught" to tell male- from female-translated texts. Our corpus, downloaded from the website of Words Without Borders, consisted of 273 samples of literary prose translated into English from a variety of languages. We found that despite its ability to isolate particular features of male- vs. female-translated texts, the computer could not be trained to accurately predict the gender of the translator. We see the difference between our results and those for original texts as highlighting the limitations of the classical social-science methodologies; i.e. notwithstanding the successful application of methods for isolating discrete features of male-translated vs. female-translated texts, these features were found to have little or no predictive value when tested in a cross-validation experiment. In other words, the same cross-validation approach that has been shown to be highly predictive in the case of author-gender attribution has proven unreliable for translator-gender attribution. We explore the implications of these results, both with regard to the competing methodologies and in terms of their implications for Translation Studies.

1. Introduction

In the Book of Judges we are told of a battle between two peoples, the Gileadites and Ephraimites. When the defeated Ephraimites tried to cross the Jordan River in an effort to return to their own territory, the Gileadites blocked the passage. Being unable to spot an Ephraimite *by appearance* – for all practical purposes, the Ephraimites resembled members of the neighboring tribe of Menashe – the Gileadites devised a method of identifying them *by their pronunciation*: capitalizing on the Ephraimites' well-known inability to distinguish between [s] and [š]. Every person who wished to cross the river was required to pronounce the telltale [šɪbələθ], and whoever pronounced it [sɪbələθ] met his demise. The biblical story (Judges, 12) is among the earliest descriptions of a categorization task. It includes two important concepts – *classes* and *features* – with the

Ephraimites and non-Ephraimites as classes and the minimal pair [s] and [š] a *phonetic feature*.

In many ways, the story of the Ephraimites is reminiscent of contemporary work on *text categorization*, in which one aims to assign an anonymous text to a particular class. Early work in authorship attribution used a variety of statistical methods to identify what are known as stylistic discriminators – characteristics that were relatively consistent in the oeuvre of a particular author, but varied from one author to the next. Mathematicians and linguists saw the potential of machine learning for extending the scope of text categorization, and attempted to bridge between early work in stylometrics and contemporary computer-based methodologies. One of the main by-products of this interface, the study of author attribution, is a productive field of inquiry in computer science to this day (see e.g. Joula 2008, Koppel *et al.* 2008, Stamatatos 2009). Indeed, the growing popularity of machine-learning techniques at the turn of the millennium has allowed for a sophisticated, robust and accurate approach to this endeavor. As in the case of our opening example, typical features of language, such as the distribution patterns of function words or content words, n-grams, lexical categories (technically part-of-speech tags), and even character n-grams have been used to categorize items by gender, age group, personality, mother tongue, etc.

There are two obvious advantages to automated text categorization over the standard (i.e. social science) statistical methods: (a) it is incomparably faster; and (b) it is bias-free, as evidenced by the fact that its analyses are completely reproducible. Among other applications, it can be used to examine problems of interest to scholars in the humanities, and can be tailored to a broad range of languages, text types, and research questions. Thus, for example, Koppel *et al.* (2006) were able to identify the most likely author of a document of unknown provenance and Strous *et al.* (forthcoming) succeeded in discriminating between texts produced by those diagnosed with schizophrenia and those produced by non-psychiatrically ill individuals.

In the present study, we set out to explore another categorization task – the classification of translated texts, focusing on *translator-gender* attribution. We will first consider the matter from a standard social-science perspective. We will base our study on a corpus of documents some of which were translated by males and some of which were translated by females. We will determine which, if any, linguistic features are used significantly differently by male and female translators and will attempt to explain any such differences. We will then consider the matter from the text categorization perspective to determine if the differences we have found have predictive value.

2. The standard approach: identifying indicators of translator gender

Gender is one of the major preoccupations of contemporary social sciences and humanities. Early research into differences between male and female discourse concentrated on socio-psychological aspects, and on the stereotypization and power differential of gender roles, as represented in language (Lakoff 1975). Studies often included discussion of conversational cultures, and of the ways in which disparate discourse patterns – Tannen (1990) goes so far as to speak of a *genderlect* – may lead to misunderstandings. Along similar lines, linguists have examined phonological, lexical, syntactic and pragmatic differences as well as turn-taking patterns (Muchnik 1997), whether in oral discourse (Holmes 1990; Labov 1990), informal writing (Mulac *et al.* 1990; Mulac & Lundell 1994), texting (Herring 1996), or formal written texts (Argamon *et al.* 2007).

Among the findings reported in these works (see Koppel *et al.* 2002 and Argamon *et al.* 2003) are that males used more determiners and more cardinal numbers, and were more prone to “specify” the things they wrote about. They were also more “informational”, as demonstrated by their greater use of such “indicators” as post-head modification, as in “garden *of* roses”. Women, on the other hand, relied heavily on pronouns, especially first- and second-person singular. Their frequent use of “I” has been interpreted as a way for the writer to introduce herself into the text, and to render it more personal; and the frequent use of “you” has been seen as a form of “involvedness” (Biber 1995).

Inspired by the body of evidence indicating that gender manifests itself in original texts, Elraz (2004) set out to investigate translations, using the findings of Koppel *et al.* (2002) as a point of departure. Her study of translations of the same text by twenty men and twenty women, revealed three general tendencies: (1) male translators use more questions in their translations; (2) female translators choose more-specific color terms; and (3) female translators are more explicit.[3] [4] Saldanha (2003) and Leonardi (2007) also found that men and women are likely to translate differently. The former, using a corpus that was similar, in principle, to our own (i.e. narratives translated from different languages into English), found differences in markers of adherence (or non-adherence) to standard forms, including such features as split infinitives and long sentences. The latter aimed at establishing a comparative framework for the contrastive analysis of male and female translation strategies. Using a critical contrastive text linguistics (CCTL) paradigm, she provided tentative indications of ideologically driven shifts in the translation process as a result of gender differences.

The corpus used in the present study was taken in its entirety from the website of Words Without Borders [wordswithoutborders.org] in February 2009. It consisted of 273 samples of literary prose translated into English from over 30 languages, among them Arabic, Chinese, French, German, Greek, Hebrew, Italian, Japanese, Korean, Portuguese, Russian, and Spanish, with a total of 908,000 tokens.[5] Poetry, graphic novels, drama, essays and interviews were excluded, as were second-hand translations, translator's notes, biographical information about the author, and footnotes. The website specifies the gender of both the author and the translator. For both the male-translated documents and the female-translated documents included in our corpus, about two thirds were authored by males.

In Table 1 we present the function words which proved to be significant at $p<0.05$. Words which were relevant as classifiers only for

---

[3] Elraz found clear differences between male and female interpreters as well, but the number of participants – 2 males and 2 females – did not allow for any clear conclusions.

[4] Elraz – citing Vásquez-Ayora (1977) – draws a clear distinction between explicitation, on the one hand, and addition or amplification, on the other, with the former being a broader, more "generic" term relating to *the need to convey meaning* whereas the latter refers simply to the grammatical constraints of the target language. Her implication, therefore, is that translations by women reflect a greater concern for the transfer of meaning.

[5] To quote the website: "Words Without Borders (WWB) opens doors to international exchange through translation, publication, and promotion of the world's best writing. WWB publishes selected prose and poetry on the web […]. Monthly issues of its online magazine feature new selections of contemporary world literature, most of which would never have been accessible to English-speaking readers without WWB."

male-authored texts are marked with one plus-sign (+). Words relevant only for female-authored texts are marked with a double plus-sign (++). The items within each group are sorted according to their frequency in the texts in descending order.

Table 1. Function words with significant differences (p<0.05) in usage between male and female translators.

Male translators	**Female translators**
and	*how*
now	*its*
off	*first*
however	*say*
today	*each*
near	*anything*
till	*until*
nevertheless	*enough*
everybody	*probably*
might+	*isn't*
is++	*never*+
are++	*than*+
where++	*can't*+
has++	*he's*+
should++	*either*+
getting++	*dear*+
done++	*was*++
	had++
	could++
	been++
	someone++
	couldn't++
	whose++
	sometimes++
	she's++
	beneath++

We also considered a set of features taken from Halliday's Systemic Functional Linguistics (SFL), where function words are represented in a hierarchy and each item or set of items is checked relative to other items or sets of items in the hierarchy (Halliday and Matthiessen 2003). Using this method we were able to discern seven more sets of features which were distributed differently in texts translated by men vs. texts translated by women (see Table 2).

Some of the features in the male column are apparently related to the “indicators” mentioned above and to men’s more frequent use of concrete time- and space-bound references, as in the case of ‘now’, ‘today’ and ‘near’. As for the use of *one* in male translations vs. *each* in female translations, these might relate to the more personal tone in female writing. Another contrast is between male use of the plural pronouns *we*, *us*, *our*, *ours*, *ourselves* and the female use of *I’m*.

Table 2. SFL features with significant differences ($p<0.05$) in usage between male and female translators.

Male translators	**Female translators**
{and, or, but, yet, however} relative to the set of other conjunctions	{each} relative to {someone, anyone, each, anybody, one, whoever, whomever}
{one} relative to {someone, anyone, each, anybody, one, whoever, whomever}	{I’m} relative to { is, isn’t, am, ain’t, I’m, are, you’re, aren’t, he’s, she’s, they’re}
{we, us, our, ours, ourselves} relative to the set of second and third person pronouns	
{is} relative to { is, isn’t, am, ain’t, I’m, are, you’re, aren’t, he’s, she’s, they’re}	
{as} relative to {like, unlike, than, as}	

In addition, we checked two more features which have proven useful in corpus-based translation studies, namely the type-token ratio (TTR) and mean sentence length (MSL); see Table 3.

Table 3. Mean sentence length and TTR by translator gender

	Male translators	**Female translators**
MSL	19.18	16.93
TTR	0.052	0.048

The mean sentence length in translations by men in our corpus was found to be much higher than in those by women: 19.18 words per sentence as opposed to 16.93. This complements the finding that men make greater use of the logical connectives *and*, *or*, *but*, *yet*, *however*. We also found that the TTR in men's translation was higher than in women's: 0.052 vs. 0.048, though this difference is not significant.

To summarize then, we have found a number of indicators of translator gender. Taken together, these markers seem to suggest that there are distinctive male and female translator styles.

3. The text categorization approach

We now ask a simple question: how meaningful are the differences that we have found above? We propose a measure of the meaningfulness of such differences based on their predictive value. Specifically, we ask whether such differences are adequate for allowing us to correctly determine the gender of the translator of a previously unseen document. To do so, we introduce in some detail the methodology and testing protocols now commonly used in text categorization studies.

Figure 1, below, presents the basic architecture of a text-categorization system. Here we are given examples of two classes of documents, Class A and Class B.

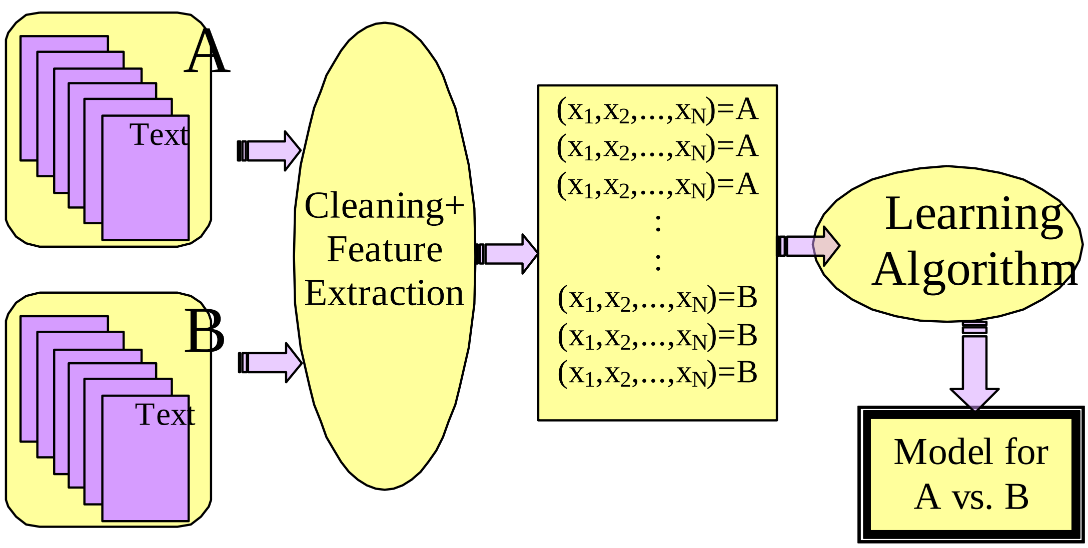

Figure 1. Architecture of a text categorization system

The first step, document representation, involves defining a set of textual features that might potentially be useful for categorizing the texts and then representing each text as a vector. A vector can be thought of as a long sequence of entries, with each representing the frequency of a particular feature in the text. If needed, this architecture can be expanded to accommodate three or more classes.

Part of the art of this methodology lies in selecting the feature types to be considered in converting documents to numerical vector representations. One commonly used feature set consists of function words, which are exceptionally useful for two very different reasons: (a) their frequencies are unlikely to be affected by subject matter, and (b) it is doubtful that people can consciously control their use of function words (Chung & Pennebaker 2007). Studies of the frequencies of function words in English texts generally focus on a few hundred items, such as determiners, pronouns, prepositions, auxiliary and modal verbs, and conjunctions. Although numbers and interjections are not generally considered to be function words, they too are often included, since they too are independent of the subject matter. Interestingly, different lists of function words enjoy roughly similar rates of success in authorship-attribution tasks (Koppel *et al.* 2008).

Once documents have been represented as vectors, a number of learning algorithms can be used to construct models capable of distinguishing vectors representing documents in Class A from those representing documents in Class B. Here we used a linear separator, in which the number of points assigned to each feature serves as an indicator of whether this specific feature provides a reliable tool for telling one class from the other. Since the precise number of points assigned to each class for a given feature is determined automatically by the learning algorithm, based on the training documents, no biases of any kind enter the process. (Contrast this with some sociolinguistic research on gender-specific variation which has been accused of ideological bias (cf. Wodak and Benke 1996: 128).) Typically, the features that are assigned a higher number of points for a particular class are the most frequent ones in documents in that class, relative to documents in the other classes. Once the point values have been determined, a new document is classified by scanning it and counting the points it contains for each class. The class assigned the highest number of points for the particular document is the one to which that document is assigned.

K-fold cross-validation is used to assess the reliability of the system. Ten-fold cross-validation is a common choice. It is illustrated in Figure 2:

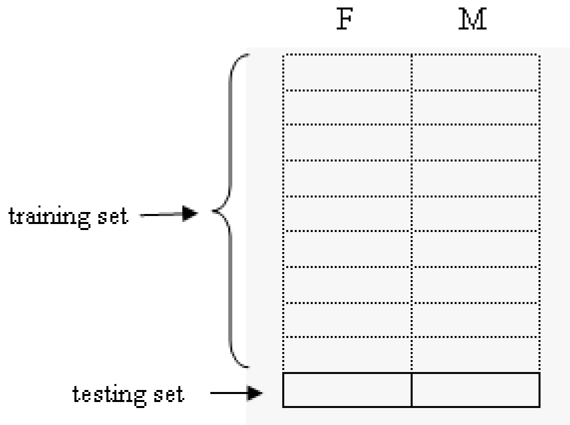

Figure 2. 10-fold cross-validation

Vertically, each column represents the categories – translations by females and translations by males. Horizontally, there are 10 folds, each of which contains a sample of each sub-corpus. The computer is trained on 90% of the corpus (9 folds) in what is called the *training set*. In this process the optimal features for distinguishing each of the sub-corpora are selected, and each feature is given a certain weight, relative to the other features. The features are represented by means of vectors, as explained earlier. The tenth fold, *the testing set*, is left out for testing. The learning algorithm is now applied to this fold and serves as an indication of the success rate of the classification task at hand. As the name ten-fold cross-validation implies, the procedure involves running this validation test ten times, in each of which one fold is kept out for testing and the other nine are used as the training set.

The last decade has seen an explosion of research in automated text categorization. Texts may be classified by author, by topic, by period, or by any other relevant criterion, including an anonymous author’s native language (Koppel *et al.* 2005) and ontological status – original versus translation (Baroni and Bernadini 2006). Provided enough texts are available and the individual texts are sufficiently long, virtually any type of text may be classified, using machine-learning techniques such as the one described above. The second author and his colleagues have performed classification tasks on sections of the British National Corpus and the International Corpus of Learner English as well as on e-mail messages, blogs, nineteenth-century novels, rabbinic literature, and commissioned essays.

Linguists and computer scientists working in the field of author attribution have also used machine-learning techniques to discriminate between texts written by men and those written by women. Recent publications in this direction include Koppel *et al.* (2002), and Argamon *et al.* (2003). Among other findings, the authors showed that based solely on function words and parts of speech, they could predict which texts had been written by men and which by women with accuracy of about 80% in ten-fold cross-validation tests. In related work (Argamon *et al.* 2007), similar features were used to determine an author's gender, age, mother tongue, and personality.

4. Findings

Our research was designed to investigate whether the techniques used in the studies described above might be applied to translations as well. In other words, our objective in this paper has been to apply text categorization methods for determining translator gender, and to see whether a computer could be taught to discriminate between male- and female-translated texts. In terms of Figure 1, female translators would be Class A and male translators Class B. We used ten-fold cross-validation tests on function-word frequencies, treating each of the 273 translations as a single example.

Surprisingly, we found that while the computer was successful in discerning many differences between the translated-by-males and translated-by-females texts, including close to 50 differences significant at $p<.05$, it was not able to distinguish between the two categories as such, i.e., it was not able to predict translator gender. Specifically, ten-fold cross-validation yielded an accuracy rate of only slightly above 50%, that is, approximately the same as random guessing. This means that although each of the indicators of translator gender was in itself significantly more frequent in translations by one gender or the other, these indicators proved insufficient for reliably distinguishing between translations by males and by females when using the more advanced and comprehensive methods currently used for author attribution and related tasks.

While it might be the case that some other feature types might have led to better results, we note that many other text categorization problems relating to writing style have been successfully solved using the very feature types used here, as in the studies cited above. In fact, when we used ten-fold cross-validation on the selfsame corpus to test for *author* gender, we were able to distinguish between male and female authors with an accuracy of 68%. It is worth noting that this accuracy is, on the one hand, substantially better than that obtained for translator gender, but, on the other hand, somewhat worse than that reported for author gender on comparable corpora (Koppel *et al.* 2002). The first of these observations simply reflects the fact that authors have more control over a text than translators do. Thus, we find in this corpus, as in others studied previously and cited above (e.g. Koppel *et al.* 2002), that female authors use significantly more singular pronouns – especially female pronouns – than male authors, while male authors use significantly more numbers than female authors. Clearly usage of telltale features such as these is controlled primarily by authors rather than by translators. The somewhat diminished accuracy on author gender compared to that obtained in previous studies is likely tied to the limited size of our corpus in terms of both the number of documents and their length. The limited quantity of training data yields a slightly less robust classifier, as is usually the case. However, it is also possible that the fact that the documents passed through translation after leaving the authors' hands resulted in the attenuation of some of the differences that are found in source texts. Thus, although the evidence is quite weak, it might be the case that the diminished accuracy as compared to previous studies is an example of the phenomenon of "leveling" (Shlesinger 1990, Baker 1996, Laviosa-Braithwaite 1996) or of standardization (Toury 1995).

5. Discussion

From the perspective of Translation Studies, once it was established that translations effectively form a distinctive textual system we had to "refine our ideas and learn to set more specific and local agendas" (Baker 2004: 29), with the influence of gender on translations figuring as one such specific agenda – one which was the focus of the present study. As Baker (2004) points out, the models available for systematic analysis do not cover the full range of features that are proving to be of interest to translation scholars, in terms of analyzing differences between translated and non-translated text, with translations being treated here as a kind of genre or text type. Based on previous studies in both computer science and Translation Studies, we had expected ten-fold cross-validation to bring us closer to devising such a model, and discriminating between translations by men and by women, but found that there is no reliable model available that will allow us to examine and account for the uneven distribution of these features across categories. Thus, although we – like Saldanha (2003), Elraz (2004) and Leonardi (2007) – did find certain features to be significantly more common in the translations by men or by women, these findings fell short of providing a profile that might serve for making reliable predictions about translator gender.[6]

Our data appear to indicate that whatever the differences between male- and female-translated documents, they are less robust than those found in original writing – so much so that even when we controlled for author gender (i.e. even when comparing male-translated vs. female-translated texts written solely by males or solely by females), our attempts to reproduce a more complete profile of the gender distinction failed.

In our view, the primary importance of these findings lies in their methodological implications. The failure of cross validation to produce a reliable means of predicting translator gender may call into question the significance of traditional statistical methods, in which individual phenomena are selected for testing, and isolated instances of significant differences are assumed to add up to a means of distinguishing male- from female-translated documents. More broadly, if indeed the use of cross-

[6] Some of our findings, moreover, run counter to theirs; e.g. Saldanha (2003) found that females use longer sentences than males, whereas our findings point in the other direction. Whether the different results derive from the methodologies used or from the corpora is not clear. To resolve the issue we would presumably need to enlarge and refine our corpora as well as our methods and to include not only lexical features but others (e.g. word order, morphology, indicators of lexical variety, subtle stylistic markers etc.) as well.

validation provides a more powerful and fully replicable means of discerning differences (and similarities) between subcorpora, these findings may imply the need to revisit conclusions based solely on traditional (mostly social science) statistical methods, since the automated text categorization methods used here call into question the predictive value of such features. In the study reported here, many features were found to distinguish male- from female-authored texts, including close to 50 differences significant at $p<.05$, but these proved to have little or no predictive value. We therefore question the effectiveness of the standard methods of categorization and suggest that automatic text categorization may provide more rigorous and reliable, as well as replicable, results.[7]

6. Conclusions

Both the biblical story and modern author attribution studies present us with ways of looking at the interface between *classes* and *features*, between the status and social situation of the pre-given social categories and the ostensibly neutral, indeed technical, elements used to distinguish between them. The ethnic identity of the Ephraimites was a *social fact*, a structural category imposed on individuals by society, a politicized category, with political consequences. Although the consequences are not a direct result of the means used to distinguish between them and non-Ephraimites, the means of categorization proved effective, and may well teach us something about classes and about identities. The means used in the present study fell short of providing such a profile. Fortunately, the stakes are less fateful.

[7] The present study focused on translation as a *product*. Triangulation with process-oriented methodologies may be useful as well in homing in on differences in micro-level decisions of male and female translators, as revealed in the introspective meta-translational discourse (e.g. think-aloud protocols), online tracking (through such programs as Translog), or even eye-tracking (see Göpferich *et al.* 2008). We hope to be able to revisit process-oriented studies from the translator-attribution perspective, and to pinpoint junctures at which the translator's decision-making process may show consistent gender-specific patterns.

References

Argamon, S., Koppel, M., Fine, J., & Shimoni, A. 2003. Gender, genre, and writing style in formal written texts. *Text* 23 (3): 401–412.

Argamon, S., Koppel, M., Pennebaker, J. & Schler, J. 2007. Automatically profiling the author of an anonymous text. *Communications of the ACM*: 119–123.

Baker, M. 1996. Corpus-based Translation Studies: the challenges that lie ahead. In H. L. Somers (ed.) *Terminology, LSP and Translation: Studies in Language Engineering in Honour of Juan C. Sager*. Amsterdam/ Philadelphia: John Benjamins. 175–186.

Baker, M. 2004. The treatment of variation in corpus-based translation studies. In A. Kruger (ed.). *Language Matters: Studies in the Languages of Africa* 35 (1): 28–38.

Baroni, M. & Bernardini, S. 2006. A new approach to the study of translationese: machine-learning the difference between original & translated text. *Literary and Linguistic Computing* 21 (3): 259–274.

Biber, D. 1995. *Dimensions of Register Variation: A Cross-linguistic Comparison.* Cambridge: Cambridge University Press.

Chung, C. & Pennebaker, J. 2007. The psychological function of function words. In K. Fiedler (ed.). *Social Communication*. New York: Psychology Press. 243–259.

Elraz, I. 2004. *His vs. Hers: Does Gender Shape One's Translation?* M.A. thesis. Department of Translation and Interpreting Studies. Bar-Ilan University. (Unpublished). http://www.biu.ac.il/HU/tr/stud-pub/tr-pub/takzir/Elraz,%20Inbal%20.htm.

Göpferich, S., Jakobsen, A. L. & Mees, I. M. (eds). 2008. *Looking at Eyes: Eye-Tracking Studies of Reading and Translation Processing*. (Copenhagen Studies in Language 36). Copenhagen: Samfundslitteratur.

Halliday, M. A. K. and Matthiessen, C. M. I. M. 2003. *An Introduction to Functional Grammar*. Hodder Arnold.

Herring, S. 1996. Two variants of an electronic message schema. In S. Herring (ed.). *Computer-Mediated Communication: Linguistic, Social & Cross-Cultural Perspectives*. Amsterdam: John Benjamins. 81–106.

Holmes, J. 1990. Hedges and boosters in women's and men's speech. *Language and Communication* 10 (3): 185–205.

Joula, P. 2008. Author attribution. *Foundations and Trends in Information Retrieval* 1 (3): 233–334.

Koppel, M., Argamon, S. & Shimoni, A. 2002. Automatically categorizing written texts by author gender. *Literary and Linguistic Computing* 17 (4): 401–412.

Koppel, M., Mughaz, D. & Akiva, N. 2006. New methods for attribution of rabbinic literature. *Hebrew Linguistics: A Journal for Hebrew Descriptive, Computational and Applied Linguistics* 57: 5–18.

Koppel, M., Schler, J. & Argamon, S. 2008. Computational methods in authorship attribution. *JASIST* 60 (1): 9–26.

Koppel, M., Schler, J., & Zigdon, K. 2005. Determining an author's native language by mining a text for errors. *Proceedings of KDD, Chicago, Illinois*. 624–628.

Labov, W. 1990. The intersection of sex and social class in the course of linguistic change. *Language Variation and Change* 2: 205–254.

Lakoff, R. T. 1975. *Language and Women's Place.* New York: Harper Colophon Books.

Laviosa-Braithwaite, S. 1996. *The English Comparable Corpus (ECC): A Resource and a Methodology for the Empirical Study of Translation.* Unpublished PhD thesis, UMIST, Manchester.

Leonardi, V. 2007. *Gender & Ideology in Translation: Do Women and Men Translate Differently? A Contrastive Analysis from Italian into English.* Bern: Peter Lang.

Muchnik, M. 1997. Men vs. women – Different communication patterns. *Hebrew Linguistics* 41–42: 79–86.

Mulac, A. & Lundell, T. L. 1994. Effects of gender-linked language differences in adults' written discourse: multivariate tests of language effects. *Language and Communication* 14: 299–309.

Mulac, A., Studley, L.B. & Blau, S. 1990. The gender-linked language effect in primary and secondary students' impromptu essays. *Sex Roles* 23 (9/10): 439–469.

Saldanha, G. 2003. Investigating gender-related linguistic features in translation. In J. Santaemilia (ed.). *Género, lenguaje y traducción.* Valencia: Universitat de València. 420–432.

Shlesinger, M. 1990. 5. 1990. Factors affecting the applicability of the oral-literate continuum to interpretation research. *Hebrew Linguistics* 28–30: 49–56.

Stamatatos, E. 2009. A survey of modern authorship attribution methods. *JASIST* 60 (3): 538–556.

Strous, R., Koppel, M., Fine, J., Nahaliel, S., Shaked, G. & Zivotofsky, A. Forthcoming. Automated characterization and identification of schizophrenia in writing. *Journal of Nervous & Mental Disorders*.

Tannen, D. 1990. *You Just Don't Understand: Women and Men in Conversation.* New York: Balantine Books.

Toury, G. 1995. *Descriptive Translation Studies and Beyond.* Amsterdam/ Philadelphia: John Benjamins.

Vásquez-Ayora, Gerardo. 1977. *Introducción a la Traductología*. Washington D.C: Georgetown University Press.

Wodak, R. & Benke, G. 1996. Gender as a sociolinguistic variable: new perspectives on variation studies. In F. Coulmas (ed.). *The Handbook of Sociolinguistics.* Oxford: Blackwell. 127–150.