

Bleaching Text: Abstract Features for Cross-lingual Gender Prediction

Rob van der Goot[♡] Nikola Ljubešić[♣] Ian Matroos[♡] Malvina Nissim[♡] Barbara Plank^{♡♣}

[♡] Center for Language and Cognition, University of Groningen, The Netherlands

[♣] Department of Knowledge Technologies, Jozef Stefan Institute, Ljubljana, Slovenia

[♣] IT University of Copenhagen, Copenhagen, Denmark

{r.van.der.goot,i.matroos,m.nissim}@rug.nl,nljubesi@gmail.com,bplank@itu.dk

Abstract

Gender prediction has typically focused on lexical and social network features, yielding good performance, but making systems highly language-, topic-, and platform-dependent. Cross-lingual embeddings circumvent some of these limitations, but capture gender-specific style less. We propose an alternative: *bleaching text*, i.e., transforming lexical strings into more abstract features. This study provides evidence that such features allow for better transfer across languages. Moreover, we present a first study on the ability of humans to perform *cross-lingual* gender prediction. We find that human predictive power proves similar to that of our bleached models, and both perform better than lexical models.

1 Introduction

Author profiling is the task of discovering latent user attributes disclosed through text, such as gender, age, personality, income, location and occupation (Rao et al., 2010; Burger et al., 2011; Feng et al., 2012; Jurgens, 2013; Bamman et al., 2014; Plank and Hovy, 2015; Flekova et al., 2016). It is of interest to several applications including personalized machine translation, forensics, and marketing (Mirkin et al., 2015; Rangel et al., 2015).

Early approaches to gender prediction (Koppel et al., 2002; Schler et al., 2006, e.g.) are inspired by pioneering work on authorship attribution (Mosteller and Wallace, 1964). Such stylometric models typically rely on carefully hand-selected sets of content-independent features to capture style beyond topic. Recently, *open vocabulary* approaches (Schwartz et al., 2013), where the entire linguistic production of an author is used, yielded substantial performance gains in on-

line user-attribute prediction (Nguyen et al., 2014; Preoțiuc-Pietro et al., 2015; Emmerly et al., 2017). Indeed, the best performing gender prediction models exploit chiefly lexical information (Rangel et al., 2017; Basile et al., 2017).

Relying heavily on the lexicon though has its limitations, as it results in models with limited portability. Moreover, performance might be overly optimistic due to topic bias (Sarawgi et al., 2011). Recent work on cross-lingual author profiling has proposed the use of solely language-independent features (Ljubešić et al., 2017), e.g., specific textual elements (percentage of emojis, URLs, etc) and users' meta-data/network (number of followers, etc), but this information is not always available.

We propose a novel approach where the actual text is still used, but bleached out and transformed into more abstract, and potentially better transferable features. One could view this as a method in between the open vocabulary strategy and the stylometric approach. It has the advantage of fading out content in favor of more shallow patterns still based on the original text, without introducing additional processing such as part-of-speech tagging. In particular, we investigate to what extent gender prediction can rely on generic non-lexical features (**RQ1**), and how predictive such models are when transferred to other languages (**RQ2**). We also glean insights from human judgments, and investigate how well people can perform cross-lingual gender prediction (**RQ3**). We focus on gender prediction for Twitter, motivated by data availability.

Contributions In this work i) we are the first to study cross-lingual gender prediction without relying on users' meta-data; ii) we propose a novel simple abstract feature representation which is surprisingly effective; and iii) we gauge human ability to perform cross-lingual gender detection, an angle of analysis which has not been studied thus far.

| | | | | | | | |
|------------------|----|-----|----|---------|-----|--------|------|
| Original | a | bag | of | Doritos | for | lunch! | 🍌🍌🍌🍌 |
| Frequency | 4 | 2 | 4 | 0 | 4 | 1 | 0 |
| Length | 01 | 03 | 02 | 07 | 03 | 06 | 04 |
| PunctC | W | W | W | W | W | W! | 🍌🍌🍌🍌 |
| PunctA | W | W | W | W | W | WP | JJJJ |
| Shape | L | LL | LL | ULL | LL | LLX | XX |
| Vowels | V | CVC | VC | CVCVCVC | CVC | CVCCCO | OOOO |

Table 1: Abstract features example transformation.

2 Profiling with Abstract Features

Can we recover the gender of an author from bleached text, i.e., transformed text where the raw lexical strings are converted into abstract features? We investigate this question by building a series of predictive models to infer the gender of a Twitter user, in absence of additional user-specific metadata. Our approach can be seen as taking advantage of elements from a data-driven open-vocabulary approach, while trying to capture gender-specific style in text beyond topic.

To represent utterances in a more language agnostic way, we propose to simply transform the text into alternative textual representations, which deviate from the lexical form to allow for abstraction. We propose the following transformations, exemplified in Table 1. They are mostly motivated by intuition and inspired by prior work, like the use of shape features from NER and parsing (Petrov and Klein, 2007; Schnabel and Schütze, 2014; Plank et al., 2016; Limsopatham and Collier, 2016):

- **Frequency** Each word is presented as its binned frequency in the training data; bins are sized by orders of magnitude.
- **Length** Number of characters (prefixed by 0 to avoid collision with the next transformation).
- **PunctC** Merges all consecutive alphanumeric characters to one ‘W’ and leaves all other characters as they are (C for conservative).
- **PunctA** Generalization of PunctC (A for aggressive), converting different types of punctuation to classes: emoticons¹ to ‘E’ and emojis² to ‘J’, other punctuation to ‘P’.
- **Shape** Transforms uppercase characters to ‘U’, lowercase characters to ‘L’, digits to ‘D’ and all other characters to ‘X’. Repetitions

¹Using the NLTK tokenizer http://www.nltk.org/_modules/nltk/tokenize/casual.html

²<https://pypi.python.org/pypi/emoji/>

of transformed characters are condensed to a maximum of 2 for greater generalization.

- **Vowel-Consonant** To approximate vowels, while being able to generalize over (Indo-European) languages, we convert any of the ‘aeiou’ characters to ‘V’, other alphabetic character to ‘C’, and all other characters to ‘O’.
- **AllAbs** A combination (concatenation) of all previously described features.

3 Experiments

In order to test whether abstract features are effective and transfer across languages, we set up experiments for gender prediction comparing lexicalized and bleached models for both in- and cross-language experiments. We compare them to a model using multilingual embeddings (Ruder, 2017). Finally, we elicit human judgments both within language and across language. The latter is to check whether a person with no prior knowledge of (the lexicon of) a given language can predict the gender of a user, and how that compares to an in-language setup and the machine. If humans can predict gender cross-lingually, they are likely to rely on aspects beyond lexical information.

Data We obtain data from the TWISTY corpus (Verhoeven et al., 2016), a multi-lingual collection of Twitter users, for the languages with 500+ users, namely Dutch, French, Portuguese, and Spanish. We complement them with English, using data from a predecessor of TWISTY (Plank and Hovy, 2015). All datasets contain manually annotated gender information. To simplify interpretation for the cross-language experiments, we balance gender in all datasets by downsampling to the minority class. The datasets’ final sizes are given in Table 2. We use 200 tweets per user, as done by previous work (Verhoeven et al., 2016). We leave the data untokenized to exclude any language-dependent processing, because original tokenization could preserve some signal. Apart from mapping usernames to ‘USER’ and urls to ‘URL’ we do not perform any further data pre-processing.

3.1 Lexical vs Bleached Models

We use the `scikit-learn` (Pedregosa et al., 2011) implementation of a linear SVM with default parameters (e.g., L2 regularization). We use 10-fold cross validation for all in-language experiments. For the cross-lingual experiments, we train

| TEST | USERS | IN-LANGUAGE | | CROSS-LANGUAGE | | | | |
|------|-------|-------------|----------|----------------|---------|--------|---------|---------|
| | | LEXICAL | ABSTRACT | LEX AVG | LEX ALL | EMBEDS | ABS AVG | ABS ALL |
| EN | 850 | 69.3 | 66.1 | 51.8 | 50.5 | 61.6 | 55.3 | 59.8 |
| NL | 894 | 81.3 | 71.8 | 52.3 | 50.0 | 56.8 | 59.5 | 69.2 |
| FR | 1,008 | 80.8 | 68.3 | 53.4 | 53.8 | 50.0 | 58.7 | 65.4 |
| PT | 3,066 | 86.0 | 68.1 | 55.3 | 63.8 | 59.5 | 59.3 | 58.9 |
| ES | 8,112 | 85.3 | 69.8 | 55.6 | 63.5 | 71.3 | 56.6 | 66.0 |

Table 2: Number of users per language and results for gender prediction (accuracy). IN-LANGUAGE: 10-fold cross-validation. CROSS-LANGUAGE: Testing on all test data in two setups: averages over single source models (AVG) or training a single model on all languages except the target (ALL). Comparison of lexical n-gram models (LEX), bleached models (ABS) and multilingual embeddings model (EMBEDS).

on all available source language data and test on all target language data.

For the lexicalized experiments, we adopt the features from the best performing system at the latest PAN evaluation campaign³ (Basile et al., 2017) (word 1-2 grams and character 3-6 grams).

For the multilingual embeddings model we use the mean embedding representation from the system of (Plank, 2017) and add max, std and coverage features. We create multilingual embeddings by projecting monolingual embeddings to a single multilingual space for all five languages using a recently proposed SVD-based projection method with a pseudo-dictionary (Smith et al., 2017). The monolingual embeddings are trained on large amounts of in-house Twitter data (as much data as we had access to, i.e., ranging from 30M tweets for French to 1,500M tweets in Dutch, with a word type coverage between 63 and 77%). This results in an embedding space with a vocabulary size of 16M word types. All code is available at <https://github.com/bplank/bleaching-text>.

For the bleached experiments, we ran models with each feature set separately. In this paper, we report results for the model where all features are combined, as it proved to be the most robust across languages. We tuned the n -gram size of this model through in-language cross-validation, finding that $n = 5$ performs best.

When testing across languages, we report accuracy for two setups: average accuracy over each single-language model (AVG), and accuracy obtained when training on the concatenation of all languages but the target one (ALL). The latter setting is also used for the embeddings model. We report accuracy for all experiments.

³<http://pan.webis.de>

| | Test → | EN | NL | FR | PT | ES |
|-------|--------|------|------|------|------|------|
| Train | EN | | 52.8 | 48.0 | 51.6 | 50.4 |
| | NL | 51.1 | | 50.3 | 50.0 | 50.2 |
| | FR | 55.2 | 50.0 | | 58.3 | 57.1 |
| | PT | 50.2 | 56.4 | 59.6 | | 64.8 |
| | ES | 50.8 | 50.1 | 55.6 | 61.2 | |
| | Avg | 51.8 | 52.3 | 53.4 | 55.3 | 55.6 |

Table 3: Pair-wise results for lexicalized models.

Results and Analysis Table 2 shows results for both the cross-language and in-language experiments in the lexical and abstract-feature setting.

Within language, the lexical features unsurprisingly work the best, achieving an average accuracy of 80.5% over all languages. The abstract features lose some information and score on average 11.8% lower, still beating the majority baseline (50%) by a large margin (68.7%). If we go across language, the lexical approaches break down (overall to 53.7% for LEX AVG/56.3% for ALL), except for Portuguese and Spanish, thanks to their similarities (see Table 3 for pair-wise results). The closely-related-language effect is also observed when training on all languages, as scores go up when the classifier has access to the related language. The same holds for the multilingual embeddings model. On average it reaches an accuracy of 59.8%.

The closeness effect for Portuguese and Spanish can also be observed in language-to-language experiments, where scores for ES \rightarrow PT and PT \rightarrow ES are the highest. Results for the lexical models are generally lower on English, which might be due to smaller amounts of data (see first column in Table 2 providing number of users per language).

The abstract features fare surprisingly well and

| | Male | Female |
|----|----------------|-----------------|
| 1 | W W W W "W" | USER E W W W |
| 2 | W W W W ? | 3 5 1 5 2 |
| 3 | 2 5 0 5 2 | W W W W ♥ |
| 4 | 5 4 4 5 4 | E W W W W |
| 5 | W W, W W W? | LL LL LL LL LX |
| 6 | 4 4 2 1 4 | LL LL LL LL LUU |
| 7 | PP W W W W | W W W W *_* |
| 8 | 5 5 2 2 5 | W W W W JJJ |
| 9 | 02 02 05 02 06 | W W W W &W;W |
| 10 | 5 0 5 5 2 | J W W W W |

Table 4: Ten most predictive features of the ABS model across all five languages. Features are ranked by how often they were in the top-ranked features for each language. Those prefixed with 0 (line 9) are length features. The prefix is used to avoid clashes with the frequency features.

work a lot better across languages. The performance is on average 6% higher across all languages (57.9% for AVG, 63.9% for ALL) in comparison to their lexicalized counterparts, where ABS ALL results in the overall best model. For Spanish, the multilingual embedding model clearly outperforms ABS. However, the approach requires large Twitter-specific embeddings.⁴

For our ABS model, if we investigate predictive features over all languages, cf. Table 4, we can see that the use of an emoji (like ♥) and shape-based features are predictive of female users. Quotes, question marks and length features, for example, appear to be more predictive of male users.

3.2 Human Evaluation

We experimented with three different conditions, one within language and two across language. For the latter, we set up an experiment where native speakers of Dutch were presented with tweets written in Portuguese and were asked to guess the poster’s gender. In the other experiment, we asked speakers of French to identify the gender of the writer when reading Dutch tweets. In both cases, the participants declared to have no prior knowledge of the target language. For the in-language experiment, we asked Dutch speakers to identify the gender of a user writing Dutch tweets. The

⁴We tested the approach with more generic (from Wikipedia) but smaller (in terms of vocabulary size) Polyglot embeddings resulting in inferior multilingual embeddings for our task.

| tweets/user: | Human | Mach. LEX | Mach. ABS |
|---------------------|-------|-----------|-----------|
| | 20 | 20 | 200 |
| NL _l →NL | 70.5 | 69.0 | 81.0 |
| NL _l →PT | 58.7 | 49.5 | 50.5 |
| FR _l →NL | 60.3 | 50.0 | 50.0 |

Table 5: Accuracy human versus machine.

Dutch speakers who participated in the two experiments are distinct individuals. Participants were informed of the experiment’s goal. Their identity is anonymized in the data.

We selected a random sample of 200 users from the Dutch and Portuguese data, preserving a 50/50 gender distribution. Each user was represented by twenty tweets. The answer key (F/M) order was randomized. For each of the three experiments we had six judges, balanced for gender, and obtained three annotations per target user.

Results and Analysis Inter-annotator agreement for the tasks was measured via Fleiss kappa ($n = 3$, $N = 200$), and was higher for the in-language experiment ($K = 0.40$) than for the cross-language tasks (NL_l→PT: $K = 0.25$; FR_l→NL: $K = 0.28$). Table 5 shows accuracy against the gold labels, comparing humans (average accuracy over three annotators) to lexical and bleached models on the exact same subset of 200 users. Systems were tested under two different conditions regarding the number of tweets per user for the target language: machine and human saw the exact same twenty tweets, or the full set of tweets (200) per user, as done during training (Section 3.1).

First of all, our results indicate that in-language performance of humans is 70.5%, which is quite in line with the findings of Flekova et al. (2016), who report an accuracy of 75% on English. Within language, lexicalized models are superior to humans if exposed to enough information (200 tweets setup). One explanation for this might lie in an observation by Flekova et al. (2016), according to which people tend to rely too much on stereotypical lexical indicators when assigning gender to the poster of a tweet, while machines model less evident patterns. Lexicalized models are also superior to the bleached ones, as already seen on the full datasets (Table 2).

We can also observe that the amount of information available to represent a user influences system’s performance. Training on 200 tweets per

user, but testing on 20 tweets only, decreases performance by 12 percentage points. This is likely due to the fact that inputs are sparser, especially since the bleached model is trained on 5-grams.⁵ The bleached model, when given 200 tweets per user, yields a performance that is slightly higher than human accuracy.

In the cross-language setting, the picture is very different. Here, human performance is superior to the lexicalized models, independently of the amount of tweets per user at testing time. This seems to indicate that if humans cannot rely on the lexicon, they might be exploiting some other signal when guessing the gender of a user who tweets in a language unknown to them. Interestingly, the bleached models, which rely on non-lexical features, not only outperform the lexicalized ones in the cross-language experiments, but also neatly match the human scores.

4 Related Work

Most existing work on gender prediction exploits shallow lexical information based on the linguistic production of the users. Few studies investigate deeper syntactic information (Koppel et al., 2002; Feng et al., 2012) or non-linguistic input, e.g., language-independent clues such as visual (Alowibdi et al., 2013) or network information (Jurgens, 2013; Plank and Hovy, 2015; Ljubešić et al., 2017). A related angle is cross-genre profiling. In both settings lexical models have limited portability due to their bias towards the language/genre they have been trained on (Rangel et al., 2016; Busger op Vollenbroek et al., 2016; Medvedeva et al., 2017).

Lexical bias has been shown to affect in-language human gender prediction, too. Flekova et al. (2016) found that people tend to rely too much on stereotypical lexical indicators, while Nguyen et al. (2014) show that more than 10% of the Twitter users do actually not employ words that the crowd associates with their biological sex. Our features abstract away from such lexical cues while retaining predictive signal.

5 Conclusions

Bleaching text into abstract features is surprisingly effective for predicting gender, though lexical infor-

⁵We experimented with training on 20 tweets rather than 200, and with different n-gram sizes (e.g., 1–4). Despite slightly better results, we decided to use the trained models as they were to employ the same settings across all experiments (200 tweets per users, $n = 5$), with no further tuning.

mation is still more useful within language (RQ1). However, models based on lexical clues fail when transferred to other languages, or require large amounts of unlabeled data from a similar domain as our experiments with the multilingual embedding model indicate. Instead, our bleached models clearly capture some signal beyond the lexicon, and perform well in a cross-lingual setting (RQ2). We are well aware that we are testing our cross-language bleached models in the context of closely related languages. While some features (such as PunctA, or Frequency) might carry over to genetically more distant languages, other features (such as Vowels and Shape) would probably be meaningless. Future work on this will require a sensible setting from a language typology perspective for choosing and testing adequate features.

In our novel study on human proficiency for cross-lingual gender prediction, we discovered that people are also abstracting away from the lexicon. Indeed, we observe that they are able to detect gender by looking at tweets in a language they do not know (RQ3) with an accuracy of 60% on average.

Acknowledgments

We would like to thank the three anonymous reviewers and our colleagues for their useful feedback on earlier versions of this paper. Furthermore, we are grateful to Chloé Braud for helping with the French human evaluation part. We would like to thank all of our human participants.

References

- Jalal S. Alowibdi, Ugo A. Buy, and Philip Yu. 2013. Language independent gender classification on twitter. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM '13*, pages 739–743, New York, NY, USA. ACM.
- David Bamman, Jacob Eisenstein, and Tyler Schnoebelen. 2014. Gender identity and lexical variation in social media. *Journal of Sociolinguistics*, 18(2):135–160.
- Angelo Basile, Gareth Dwyer, Maria Medvedeva, Josine Rawee, Hessel Haagsma, and Malvina Nissim. 2017. N-gram: New groningen author-profiling model. In *Proceedings of the CLEF 2017 Evaluation Labs and Workshop – Working Notes Papers, 11-14 September, Dublin, Ireland (Sept. 2017)*.
- John D. Burger, John Henderson, George Kim, and Guido Zarrella. 2011. *Discriminating Gender on Twitter*. In *Proceedings of the 2011 Conference on*

- Empirical Methods in Natural Language Processing*, pages 1301–1309, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Chris Emmery, Grzegorz Chrupała, and Walter Daelemans. 2017. Simple queries as distant labels for predicting gender on twitter. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 50–55, Copenhagen, Denmark. Association for Computational Linguistics.
- Song Feng, Ritwik Banerjee, and Yejin Choi. 2012. Characterizing stylistic elements in syntactic structure. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1522–1533. Association for Computational Linguistics.
- Lucie Flekova, Jordan Carpenter, Salvatore Giorgi, Lyle Ungar, and Daniel Preotiuc-Pietro. 2016. [Analyzing biases in human perception of user age and gender from text](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 843–854, Berlin, Germany. Association for Computational Linguistics.
- David Jurgens. 2013. That’s what friends are for: Inferring location in online social media platforms based on social relationships. *ICWSM*, 13(13):273–282.
- Moshe Koppel, Shlomo Argamon, and Anat Rachel Shimoni. 2002. Automatically categorizing written texts by author gender. *Literary and Linguistic Computing*, 17:401–412.
- Nut Limsopatham and Nigel Collier. 2016. [Bidirectional lstm for named entity recognition in twitter messages](#). In *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT)*, pages 145–152, Osaka, Japan. The COLING 2016 Organizing Committee.
- Nikola Ljubešić, Darja Fišer, and Tomaž Erjavec. 2017. Language-independent gender prediction on twitter. In *Proceedings of the Second Workshop on NLP and Computational Social Science*, pages 1–6, Vancouver, Canada. Association for Computational Linguistics.
- Maria Medvedeva, Hessel Haagsma, and Malvina Nissim. 2017. An analysis of cross-genre and in-genre performance for author profiling in social media. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction - 8th International Conference of the CLEF Association, CLEF 2017, Dublin, Ireland, September 11-14, 2017, Proceedings*, pages 211–223.
- Shachar Mirkin, Scott Nowson, Caroline Brun, and Julien Perez. 2015. Motivating personality-aware machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1102–1108.
- Frederick Mosteller and David Wallace. 1964. Inference and disputed authorship: The federalist.
- Dong Nguyen, Dolf Trieschnigg, A Seza Dođruöz, Rilana Gravel, Mariët Theune, Theo Meder, and Francisca De Jong. 2014. Why gender and age prediction from tweets is hard: Lessons from a crowdsourcing experiment. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1950–1961.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Slav Petrov and Dan Klein. 2007. [Improved inference for unlexicalized parsing](#). In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pages 404–411, Rochester, New York. Association for Computational Linguistics.
- Barbara Plank. 2017. All-in-1 at ijcnlp-2017 task 4: Short text classification with one model for all languages. *Proceedings of the IJCNLP 2017, Shared Tasks*, pages 143–148.
- Barbara Plank and Dirk Hovy. 2015. [Personality traits on twitter—or—how to get 1,500 personality tests in a week](#). In *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 92–98, Lisboa, Portugal. Association for Computational Linguistics.
- Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. [Multilingual Part-of-Speech Tagging with Bidirectional Long Short-Term Memory Models and Auxiliary Loss](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Daniel Preotiuc-Pietro, Vasileios Lampos, and Nikolaos Aletras. 2015. An analysis of the user occupational class through twitter content. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1754–1764.
- Francisco Rangel, Paolo Rosso, Martin Potthast, and Benno Stein. 2017. Overview of the 5th author profiling task at pan 2017: Gender and language variety identification in twitter. *Working Notes Papers of the CLEF*.
- Francisco Rangel, Paolo Rosso, Martin Potthast, Benno Stein, and Walter Daelemans. 2015. Overview of the 3rd author profiling task at pan 2015. In *CLEF*, page 2015. sn.

- Francisco Rangel, Paolo Rosso, Ben Verhoeven, Walter Daelemans, Martin Potthast, and Benno Stein. 2016. Overview of the 4th Author Profiling Task at PAN 2016: Cross-genre Evaluations. In *Working Notes Papers of the CLEF 2016 Evaluation Labs. CEUR Workshop Proceedings*. CLEF and CEUR-WS.org.
- Delip Rao, David Yarowsky, Abhishek Shreevats, and Manaswi Gupta. 2010. Classifying latent user attributes in twitter. In *Proceedings of the 2nd international workshop on Search and mining user-generated contents*, pages 37–44. ACM.
- Sebastian Ruder. 2017. A survey of cross-lingual embedding models. *arXiv preprint arXiv:1706.04902*.
- Ruchita Sarawgi, Kailash Gajulapalli, and Yejin Choi. 2011. Gender attribution: tracing stylometric evidence beyond topic and genre. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 78–86. Association for Computational Linguistics.
- Jonathan Schler, Moshe Koppel, Shlomo Argamon, and James W. Pennebaker. 2006. Effects of Age and Gender on Blogging. In *AAAI Spring Symposium: Computational Approaches to Analyzing Weblogs*, pages 199–205. AAAI.
- Tobias Schnabel and Hinrich Schütze. 2014. Flors: Fast and simple domain adaptation for part-of-speech tagging. *Transactions of the ACL*, 2:15–26.
- H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791.
- Samuel L Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv preprint arXiv:1702.03859*.
- Ben Verhoeven, Walter Daelemans, and Barbara Plank. 2016. Twisty: A multilingual twitter stylometry corpus for gender and personality profiling. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- M. Busger op Vollenbroek, T. Carlotto, T. Kreutz, M. Medvedeva, C. Pool, J. Bjerva, H. Haagsma, and M. Nissim. 2016. Gronup: Groningen user profiling notebook for PAN at CLEF. In *CLEF 2016 Evaluation Labs and Workshop – Working Notes Papers*.