

Efficient Social Network Multilingual Classification using Character, POS n -grams and Dynamic Normalization

Carlos-Emiliano González-Gallardo¹, Juan-Manuel Torres-Moreno^{1,2}, Azucena Montes Rendón³ and Gerardo Sierra⁴

¹Laboratoire Informatique d'Avignon, Université d'Avignon et des Pays de Vaucluse, Avignon, France

²École Polytechnique de Montréal, Montréal, Canada

³Centro Nacional de Investigación y Desarrollo Tecnológico, Cuernavaca, Mexico

⁴GIL-Instituto de Ingeniería, Universidad Nacional Autónoma de México, Ciudad de México, Mexico

Keywords: Text Mining, Machine Learning, Classification, n -grams, POS, Blogs, Tweets, Social Network.

Abstract: In this paper we describe a dynamic normalization process applied to social network multilingual documents (Facebook and Twitter) to improve the performance of the Author profiling task for short texts. After the normalization process, n -grams of characters and n -grams of POS tags are obtained to extract all the possible stylistic information encoded in the documents (emoticons, character flooding, capital letters, references to other users, hyperlinks, hashtags, etc.). Experiments with SVM showed up to 90% of performance.

1 INTRODUCTION

Social networks are a resource that millions of people use to support their relationships with friends and family (Chang et al., 2015). In June 2016, Facebook¹ reported to have approximately 1.13 billion active users per day; while in June, Twitter² reported an average of 313 million active users per month.

Due to the large amount of information that is continuously produced in social networks, to identify certain individual characteristics of users has been a problem of growing importance for different areas like forensics, security and marketing (Rangel et al., 2015); problem that Author profiling aims from an automatic classification approach with the premise that depending of the individual characteristics of each person; such as age, gender or personality, the way of communicating will be different.

Author profiling is traditionally applied in texts like literature, documentaries or essays (Argamon et al., 2003; Argamon et al., 2009); these kind of texts have a relative long length (Peersman et al., 2011) and a standard language. By the other hand, documents produced by social networks users have some characteristics that differ from regular texts and prevents that

they can be analyzed in a similar way (Peersman et al., 2011). Some characteristics that social networks texts share are their length (significantly shorter that traditional texts) (Peersman et al., 2011), the large number of misspellings, the use of non-standard capitalization and punctuation.

Social networks like Twitter have their own rules and features that users use to express themselves and communicate with each other. It is possible to take advantage of these rules and extract a greater amount of stylistic information. (Gimpel et al., 2011) introduce this idea to create a Part of Speech (POS) tagger for Twitter. In our case, we chose to perform a dynamic context-dependent normalization. This normalization allows to group those elements that are capable of providing stylistic information regardless of its lexical variability; this phase helps to improve the performance of the classification process. In (González-Gallardo et al., 2016) we presented some preliminary results of this work, that we will extend with more experiments and further description.

The paper is organized as follows: in section 2 we provide an overview of Author profiling applied to age, gender and personality traits prediction. In section 3, a brief presentation of character and POS n -grams is presented. In section 4 we detail the methodology used in the dynamic context-dependent normal-

¹Facebook website: <http://www.facebook.com>

²Twitter website: <http://www.twitter.com>

ization. Section 5 presents the datasets used in the study. The learning model is detailed in section 6. The different experiments and results are presented in section 7. Finally, in section 8 we present the conclusions and some future job prospects.

2 AGE, GENDER AND PERSONALITY TRAITS PREDICTION

(Koppel et al., 2002) performed a research with a corpus made of 566 documents of the *British National Corpus (BNC)*³ in which they identified automatically the gender of the author using function words and POS *n*-grams, obtaining a precision of 80% with a linear separator.

(Doyle and Kešelj, 2005) constructed a simple classifier to predict gender in a collection of 495 essays of the *BAWE* corpus⁴. The classifier consisted in measuring the distance between the text to classify and two different profiles (man and woman), being the smallest one the class that was assign to the text. A precision of 81% was reported using character, word and POS *n*-grams as features.

In (Peersman et al., 2011), the authors collected 1.5 millions of samples of the Netlog social network to predict the age and gender of the authors. They used a model created with support vector machines (SVM) (Vapnik, 1998) taking into account *n*-grams of words and characters as features. The reported precision was of 64.2%.

(Carmona et al., 2015) predicted age, gender and personality traits of tweets authors using a high lever representation of features. This representation is composed of discriminatory features (Second Order Atribbutes) (Lopez-Monroy et al., 2013) and descriptive features (Latent Semantic Analysis) (Wiemer-Hastings et al., 2004). For age and gender prediction in Spanish tweets they obtained a precision of 77.27% and for personality traits a RMSE value of 0.1297.

In (Grivas et al., 2015), a classification model with SVM was used to predict age, gender and personality traits of Spanish tweets while a SVR model was used to predict personality traits of tweets authors. In this approach two types of feature groups were considered: structural and stylometric. Structural features were extracted form the unprocessed tweets while stylometric features were obtained after html tags, urls, hashtags and references to other users are removed.

A precision of 72.73% was reported for age and gender prediction, while a RMSE value of 0.1495 was obtained for personality traits prediction.

3 CHARACTER AND POS *n*-grams

n-grams are sequences of elements of a selected textual information unit (Manning and Schütze, 1999) that allow the extraction of content and stylistic features from text; features that can be used in tasks such as Automatic Summarization, Automatic Translation and Text Classification.

The information unit to use changes depending on the task and the type of features that will be extracted. For example, in Automatic Translation and Summarization is common to use word and phrase *n*-grams (Torres-Moreno, 2014; Giannakopoulos et al., 2008; Koehn, 2010). Within Text Classification, for Plagiarism Detection, Author Identification and Author Profiling; character, word and POS *n*-grams are used (Doyle and Kešelj, 2005; Stamatatos et al., 2015; Oberreuter and Velásquez, 2013).

The information units selected in this research are characters and POS tags. With the character *n*-grams the goal is to extract as many stylistic elements as possible: characters frequency, use of suffixes (gender, number, tense, diminutives, superlatives, etc.), use of punctuation, use of emoticons, etc (Stamatatos, 2006; Stamatatos, 2009).

POS *n*-grams provide information about the structure of the text: frequency of grammatical elements, diversity of grammatical structures and the interaction between grammatical elements. The POS tags were obtained using the Freeling⁵ POS tagger, which follows the POS tags proposed by EAGLES⁶. To fully control the standardization process and make it independent of a detector of names, we preferred to perform a specific normalization for both datasets instead of using the Freeling functionalities (Padró and Stanilovsky, 2012).

The POS tags provided by Freeling have several levels of detail that provide insight into the different attributes of a grammatical category depending of the language that is being analyzed. In our case we used only the first level of detail that refers to the category itself. The example in Table 1 shows the Spanish word "*especialización*" (specialization), all its attributes and the selected tag.

³BNC website <http://www.natcorp.ox.ac.uk/>

⁴BAWE corpus website: <http://www2.warwick.ac.uk/fac/soc/al/research/collect/bawe/>

⁵Freeling website: <http://nlp.lsi.upc.edu/freeling/node/1>

⁶EAGLES website: <http://www.ilc.cnr.it/EAGLES96/annotate/node9.html>

Table 1: Part-of-speech tagging of the Spanish word: *especialización*.

Word: "especialización"			
Attribute	Code	Value	Tag
Category	N	Noun	N
Type	C	Common	
Gender	F	Female	
Number	S	Singular	
Case	0	-	
Semantic	0	-	
Gender	0	-	
Grade	0	-	

4 DYNAMIC CONTEXT-DEPENDENT NORMALIZATION

The freedom that social networks users have to encode their messages derives in a varied lexicon. To minimize this variations it is necessary to normalize those elements that are able to provide stylistic information regardless of its lexical variability (references to other users, hyperlinks, hashtags, emoticons, etc). This is what we call Dynamic Context-dependent Normalization which is separated in two phases: Text normalization and POS relabeling.

- Text normalization

The goal of this phase is to avoid the lexical variability that is present when a user tends to do certain actions in their texts; for example tagging another user or creating a link to a website. In the Twitter case, references to other users are defined as

```
@{username}
```

The amount of values that can be assigned to the label *username* is potentially infinite (depending on the number of users available to the social network). To avoid such variability, the Dynamic Context-dependent Normalization standardize this element in order to highlight the intention: make a reference to a user. So, the tweet

```
I was just watching ``update 10.``
@MKBHD http://t.co/P9Dn7t8zS1
```

will be normalized to

```
I was just watching ``update 10.``
@us http://t.co/P9Dn7t8zS1
```

Hyperlinks have a similar behavior: the number of links to Internet sites is also potentially infinite.

The important fact is that a reference to an external site is being implemented; so all text strings that match the pattern:

```
http[s]://{external_site}
```

are normalized. Following with the previous example, the tweet

```
I was just watching ``update 10.``
@us http://t.co/P9Dn7t8zS1
```

will be normalized to

```
I was just watching ``update 10.``
@us htt
```

- POS relabeling

All these lexical variations provide important grammatical information that must be preserved, but conventional POS taggers are unable to maintain. Therefore, it is necessary to relabel certain elements to keep their interaction with the rest of the POS tags of the text.

Using Freeling to tag the previous example and taking into account only the categories of the POS tag, the sequence obtained is

```
P V R V N Z F N N
```

As it can be seen, the reference to the user and the link to the site is lost and it is no possible to know how those elements interact with the rest of the tags. To overcome this limitation, references to other users, hyperlinks and hashtags are relabeled so that they have their own special tag dealing to the next sequence:

```
P V R V N Z REF@USERNAME REF#LINK
```

Figure 1 shows a general architecture of the system.

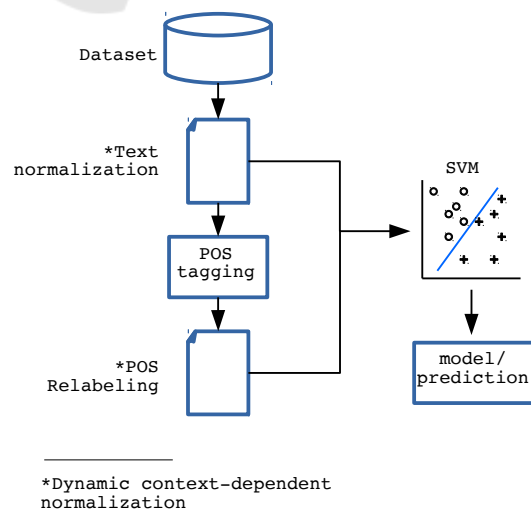


Figure 1: General architecture of classification system.

5 DATASETS

In order to test various contexts, we used corpora from two social networks: Twitter and Facebook.

The multilingual corpus PAN-CLEF2015 (Twitter) is labeled by gender, age and personality traits; whereas the multilingual corpus PAN-CLEF2016 (Twitter) is just labeled by gender and age. The Mexican Spanish corpus “Comments of Mexico City through time” (Facebook) is only labeled by gender.

5.1 PAN-CLEF2015

The PAN-CLEF2015 corpus⁷ (Rangel et al., 2005) is conformed by 324 samples distributed in four languages: Spanish, English, Italian and Dutch. A larger description of the corpus can be seen in (González-Gallardo et al., 2016).

5.2 PAN-CLEF2016_s

The multilingual PAN-CLEF2016_s corpus is a set of the training corpus for the Author Profiling task of PAN 2016.⁸ It is conformed by 935 samples labeled by age and gender which are distributed in three languages: Spanish, English and Dutch. For our experiments, only the gender label was considered.

Table 2 shows the distribution of gender samples per language.

Table 2: PAN-CLEF2016_s, distribution of samples by gender.

	Samples	
	Female	Male
Spanish	52%	48%
English	52%	48%
Dutch	50%	50%

5.3 Comments of Mexico City through time (CCDMX)

The CCDMX corpus consists of 5 979 Mexican Spanish comments from the Facebook page *Mexico City through time*⁹. In this Facebook page, pictures of Mexico City are posted so that people can share their memories and anecdotes. The average length of each comment is 110 characters.

⁷PAN challenge’s website: <http://pan.webis.de/>

⁸The corpus is downloadable at the Website: <http://www.uni-weimar.de/medien/webis/events/pan-16/pan16-code/pan16-author-profiling-twitter-downloader.zip>

⁹Website of the blog: <http://www.facebook.com/laciudaddemexicoeneltiempo>

The CCDMX corpus was manually annotated by bachelor’s linguistic students of the Group of Linguistic Engineering (GIL) of the UNAM in 2014¹⁰. It is only labeled by gender, being slightly higher the number of comments that belong to the “Men” class (see table 3).

Table 3: CCDMX, distribution of samples by gender.

	Comments	Percentage
Female	2573	43%
Male	3406	57%
Total of comments	5 979	100%

6 LEARNING MODEL

For the experiments we used support vector machines (SVM) (Vapnik, 1998), a classical model of supervised learning, which has proven to be robust and efficient in various NLP tasks.

In particular, to perform experiments we use the Python package *Scikit-learn*¹¹ (Pedregosa et al., 2011) using a linear kernel (LinearSVC), which produced empirically the best results.

6.1 Features

The character and POS *n*-gram windows used were generated with a unit length of 1 to 3.

Thus, for example, the word “*self-defense*” is represented by the following character *n*-grams:

$\{s, e, l, f, -, d, e, f, e, n, s, e, -, s, e, e, l, l, f, f, -, d, d, e, e, f, f, e, n, n, s, s, e, e, -, s, e, s, e, l, e, l, f, f, -, f, -, d, -, d, e, d, e, f, e, f, e, n, e, n, s, e, -\}$

And the normalize tweet “@us @us You owe me one, Cam!” which POS sequence is “REF@USERNAME REF@USERNAME N V P N F N F”, is represented by the POS *n*-grams:

$\{REF@USERNAME, REF@USERNAME, N, V, P, N, F, N, F, REF@USERNAME REF@USERNAME, REF@USERNAME N, N V, V P, P N, N F, F N, N F, REF@USERNAME REF@USERNAME N, REF@USERNAME N V, N V P, V P N, P N F, N F N, F N F\}$

The best results were obtained using a linear frequency scale in all cases except from the POS *n*-grams for Spanish texts, in which the logarithmic

¹⁰Website of the corpus: <http://corpus.unam.mx>

¹¹Scikit-learn is downloadable at the Website: <http://scikit-learn.org>

function

$$\log_2(1 + \text{number_of_occurrences})$$

is applied. In figure 2 is possible to see the linearized POS 2-grams values of the PAN-CLEF2015 training corpus.




Figure 2: Linearized frequencies of the Spanish PAN-CLEF2015 corpus.

6.2 Experimental Protocol

Four experiments were performed with the PAN-CLEF2015 corpus; one for each language. 70% of the samples was used for training the learning model and the remaining 30% was used during evaluation time.

Regarding to the PAN-CLEF2016.s corpus, three experiments were performed; one for Dutch, one for English and another one for Spanish. We used the training model created with the PAN-CLEF2015 corpus and applied it to this corpus. All the samples available in the PAN-CLEF2016.s corpus were used as test samples for gender classification.

Three experiments were performed with the CCDMX corpus.

- First, all the comments were used as test samples using the learning model generated with the Spanish training samples of the PAN-CLEF2015 corpus.
- For the second experiment, samples of 50 comments were created. Thus 121 samples were tested using the same learning model of the first experiment.
- Finally, the third experiment was a sum of various micro experiments. For each micro experiment a different sample size was tested: 1, 2, 4, 8, 16, 24, 32, 41, 50, 57 and 64 comments per sample. 70% of the samples were used to train the learning model and 30% to test it.

7 RESULTS

In order to evaluate the performance of the model in all the datasets, a group of classical measures was implemented. Accuracy (A), Precision (P), Recall (R) and F-score (F) (Manning and Schütze, 1999) were used to evaluate gender and age prediction. In relation to the personality traits prediction in the PAN-CLEF2015 corpus, the Root Mean Squared Error (RMSE) measure was used.

7.1 PAN-CLEF2015 Corpus

Tables 4 to 7 show the results obtained from the PAN-CLEF2015 corpus for Spanish and English.

Evaluation measures reports practically 100% in gender classification for Italian and Dutch (González-Gallardo et al., 2016). This is probably because the number of existing samples for both languages was very small. We think it would be worth trying with a larger amount of data to validate the results in these two languages.

- Spanish.

Table 4: PAN-CLEF2015; gender and age results (Spanish).

Gender	P	R	F	A
Male	0.929	0.867	0.897	0.900
Female	0.875	0.930	0.902	
Age	P	R	F	A
18-24	0.750	1	0.857	0.800
25-34	0.750	0.875	0.807	
35-49	1	0.667	0.800	
>50	1	0.500	0.667	

Table 5: PAN-CLEF2015; personality traits results (Spanish).

Trait	RMSE
E	0.106
N	0.128
A	0.158
C	0.164
O	0.138
Mean	0.139

- English.

Table 6: PAN-CLEF2015; gender and age results (English).

Gender	P	R	F	A
Male	0.826	0.826	0.826	0.826
Female	0.826	0.826	0.826	
Age	P	R	F	A
18-24	0.895	0.944	0.919	0.848
25-34	0.789	0.833	0.810	
35-49	0.800	0.667	0.727	
>50	1	0.750	0.857	

Table 7: PAN-CLEF2015; personality traits results (English).

Trait	RMSE
E	0.182
N	0.182
A	0.150
C	0.123
O	0.162
Mean	0.160

7.2 PAN-CLEF2016_s Corpus

As it can be see in tables 8 to 10, the performance of the experiments is low in all three languages. We think this is caused by the difference that exists in the size of the samples used to train the model (PAN-CLEF2015 corpus) and the size of the PAN-CLEF2016_s corpus samples. Another possible reason is that the quantity of noise in the PAN-CLEF2016_s corpus was too much to be handled by the trained models; thus creating mistaken predictions.

- Spanish.

Table 8: PAN-CLEF2016_s; gender results (Spanish).

Gender	P	R	F	A
Male	0.494	0.967	0.654	0.505
Female	0.700	0.071	0.130	

- English.

Table 9: PAN-CLEF2016_s; gender results (English).

Gender	P	R	F	A
Male	0.578	0.531	0.554	0.587
Female	0.594	0.638	0.615	

- Dutch.

Number of samples: 382.

Table 10: PAN-CLEF2016_s; gender results (Dutch).

Gender	P	R	F	A
Male	0.774	0.251	0.379	0.589
Female	0.553	0.927	0.693	

7.3 CCDMX Corpus

The first experiment (E1) performed with this corpus aimed to discover how much impact had the difference between the training and test samples. The training phase was done with 70% of the samples from the PAN -CLEF2015 corpus (Spanish). Remember that a sample of this corpus is a group of about 100 tweets.

Table 11 shows the results of the 5 979 samples that were tested.

Table 11: CCDMX, E1 results.

	P	R	F	A
Male	0.598	0.631	0.614	0.549
Female	0.474	0.439	0.456	

In the second experiment (E2) we chose to generate samples of 50 comments, size that represent a reasonable compromise between number of samples and number of characters per sample (about 5K characters).

A total of 121 samples were tested with the learning model of E1. The results are slightly better than E1 but the domain difference seems to affect greatly the system performance (Table 12).

Table 12: CCDMX; E2 results.

	P	R	F	A
Male	0.657	0.942	0.774	0.686
Female	0.818	0.346	0.486	

A third experiment (E3) was done with this corpus; eleven micro experiments with different number of comments per sample were performed: 1, 2, 4, 8, 16, 24, 32, 41, 50, 57, 64. This was done to measure the impact of the sample size variation in the performance of the algorithm.

Figures 3 to 6 show the performance of E3 (accuracy, precision, recall and F-score). As it can be seen, something curious happens when the sample is of 50 comments length. At that point the general performance drops about 8%. This is due to a statistical effect that the data shows with the residual sample. For both male and female, a residual sample of 22 comments is present. It is necessary to mention that this effect does not affect the general results of E3, which shows to be consistent with the rest of the experiments.




Figure 3: CCDMX; Accuracy.




Figure 4: CCDMX; Precision.

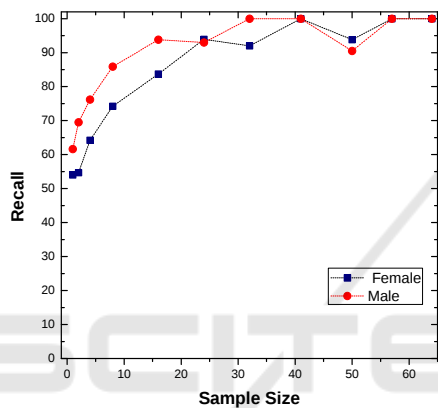


Figure 5: CCDMX; Recall.

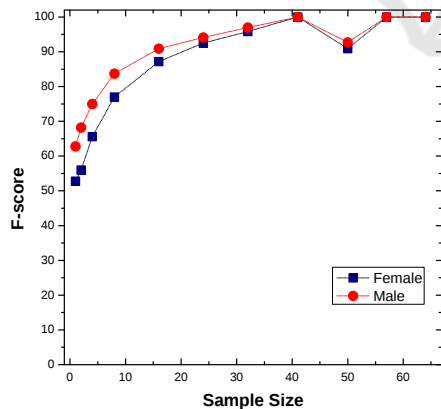


Figure 6: CCDMX; F-score.

8 CONCLUSIONS AND FUTURE WORK

Character and POS n -grams have shown to be a useful resource for the extraction of stylistic features in

short texts. With character n -grams it was possible to extract emoticons, punctuation exaggeration (character flooding), use of capital letters and different kinds of emotional information encoded in tweets and comments. The interaction between the special elements of the social networks, like references to users, hashtags or links, and the rest of the POS tags were captured with the POS n -grams after. Also, for Spanish and English it was possible to capture the most representative series of two and three grammatical elements; for the Italian and Dutch we were able to capture the most common grammatical elements. The Dynamic Context-dependent Normalization showed to be effective in the different domains (Facebook and Twitter), but also showed to have problems when a cross domain evaluation was performed. The poor results obtained in E1 reflect that it is important to keep the ratio size between train and test samples. E2 results show that the change of domain greatly affects the prediction capacity of the system; fact that E3 and the results obtained with the PAN-CLEF2016_s corpus reinforce. A proposal for future work is to repeat E3 applying cross validation. This will show a better distribution of the training and test samples eliminating the possibility of a biased result. Also, this probably balance the 50 comments samples and improves its performance. An interesting development for future work would be to apply the algorithm presented in this article to comments of multimedia sources to separate in an automatic way the different sectors of opinion makers.

ACKNOWLEDGEMENTS

This project was partially financed by the project: CONACyT-México No. 215179 *Caracterización de huellas textuales para el análisis forense*. We also appreciate the financing of the european project CHISTERA CALL - ANR: *Access Multilingual Information opinionS (AMIS)*, (France - Europe).

REFERENCES

- Argamon, S., Koppel, M., Fine, J., and Shimoni, A. R. (2003). Gender, genre, and writing style in formal written texts. *Text*, 23(3):321–346.
- Argamon, S., Koppel, M., Pennebaker, J. W., and Schler, J. (2009). Automatically profiling the author of an anonymous text. *Communications of the ACM*, 52(2):119–123.
- Carmona, M. Á. Á., López-Monroy, A. P., Montes-y-Gómez, M., Pineda, L. V., and Escalante, H. J. (2015).

- Inaoe's participation at pan'15: Author profiling task. In *Working Notes of CLEF 2015, Toulouse, France*.
- Chang, P. F., Choi, Y. H., Bazarova, N. N., and Löckenhoff, C. E. (2015). Age differences in online social networking: Extending socioemotional selectivity theory to social network sites. *Journal of Broadcasting & Electronic Media*, 59(2):221–239.
- Doyle, J. and Kešelj, V. (2005). Automatic categorization of author gender via n-gram analysis. In *6th Symposium on Natural Language Processing, SNLP*.
- Giannakopoulos, G., Karkaletsis, V., and Vouros, G. (2008). Testing the use of n-gram graphs in summarization sub-tasks. In *Text Analysis Conference (TAC)*.
- Gimpel, K., Schneider, N., O'Connor, B., Das, D., Mills, D., Eisenstein, J., Heilman, M., Yogatama, D., Flanagan, J., and Smith, N. A. (2011). Part-of-speech tagging for twitter: Annotation, features, and experiments. In *49th ACL HLT: Short Papers v2*, pages 42–47, Stroudsburg. ACL.
- González-Gallardo, C.-E., Torres-Moreno, J.-M., Montes Rendón, A., and Sierra, G. (2016). Perfilado de autor multilingüe en redes sociales a partir de n-gramas de caracteres y de etiquetas gramaticales. *Linguamática*, 8(1):21–29.
- Grivas, A., Krithara, A., and Giannakopoulos, G. (2015). Author profiling using stylometric and structural feature groupings. In *Working Notes of CLEF 2015 - Conference and Labs of the Evaluation forum, Toulouse, France, September 8-11, 2015*.
- Koehn, P. (2010). *Statistical Machine Translation*. Cambridge University Press, NY, USA, 1st edition.
- Koppel, M., Argamon, S., and Shimon, A. R. (2002). Automatically categorizing written texts by author gender. *Literary and Linguistic Computing*, 17(4):401–412.
- Lopez-Monroy, A. P., Gomez, M. M.-y., Escalante, H. J., Villasenor-Pineda, L., and Villatoro-Tello, E. (2013). Inaoe's participation at pan'13: Author profiling task. In *CLEF 2013 Evaluation Labs and Workshop*.
- Manning, C. D. and Schütze, H. (1999). *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge.
- Oberreuter, G. and Velásquez, J. D. (2013). Text mining applied to plagiarism detection: The use of words for detecting deviations in the writing style. *Expert Systems with Applications*, 40(9):3756–3763.
- Padró, L. and Stanilovsky, E. (2012). Freeling 3.0: Towards wider multilinguality. In *Language Resources and Evaluation Conference (LREC 2012)*, Istanbul, Turkey. ELRA.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, É. (2011). Scikit-learn: Machine Learning in Python. *Machine Learning Research*, 12:2825–2830.
- Peersman, C., Daelemans, W., and Van Vaerenbergh, L. (2011). Predicting age and gender in online social networks. In *3rd Int. Workshop on Search and mining user-generated contents*, pages 37–44. ACM.
- Rangel, F., Rosso, P., Potthast, M., Stein, B., and Daelemans, W. (2005). Overview of the 3rd Author Profiling Task at PAN 2015. In *CLEF 2015 Labs and Workshops, Notebook Papers*. Cappellato L. and Ferro N. and Gareth J. and San Juan E. (Eds.): CEUR-WS.org.
- Rangel, F., Rosso, P., Potthast, M., Stein, B., and Daelemans, W. (2015). Overview of the 3rd author profiling task at pan 2015. In *CLEF*.
- Stamatatos, E. (2006). Ensemble-based Author Identification Using Character N-grams. In *3rd Int. Workshop on Text-based Information Retrieval*, pages 41–46.
- Stamatatos, E. (2009). A Survey of Modern Authorship Attribution Methods. *American Society for information Science and Technology*, 60(3):538–556.
- Stamatatos, E., Potthast, M., Rangel, F., Rosso, P., and Stein, B. (2015). Overview of the pan/clef 2015 evaluation lab. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 518–538.
- Torres-Moreno, J.-M. (2014). *Automatic Text Summarization*. Wiley-Sons, London.
- Vapnik, V. N. (1998). *Statistical Learning Theory*. Wiley-Interscience, New York.
- Wiemer-Hastings, P., Wiemer-Hastings, K., and Graesser, A. (2004). Latent semantic analysis. In *16th Int. joint conference on Artificial intelligence*, pages 1–14.