International Journal of Computational Linguistics and Applications, Vol. 6, No. 1, 2015, pp. 45–62
Received 19/02/2015, Accepted 15/05/2015, Final 20/06/2015.
ISSN 0976-0962, http://ijcla.bahripublications.com

Computational Study of Stylistics: A Clustering-based Interestingness Measure for Extracting Relevant Syntactic Patterns

MOHAMED-AMINE BOUKHALED
FRANCESCA FRONTINI
GAUVAIN BOURGNE
JEAN-GABRIEL GANASCIA
Université Pierre et Marie Curie and CNRS (UMR7606), France

ABSTRACT

In this contribution, we present a computational stylistic study of the French classic literature texts based on a data-driven approach where discovering interesting linguistic patterns is done without any prior knowledge. We propose an objective interestingness measure to extract meaningful stylistic syntactic patterns from a given author's work. Our hypothesis is based on the fact that the most characterising linguistic patterns should significantly reflect the author's stylistic choice in that the positions of theirs occurrences are controlled by the author's purpose, while the irrelevant linguistic patterns are distributed randomly in the text. Since it does not rely on the counts of occurrences of the syntactic patterns in texts, this measure can work reasonably well with both large and small text samples. The analysed results show the effectiveness in extracting interesting syntactic patterns from a single text, and this seems particularly promising for the analyses of such texts that, for their characteristics or for historical reasons, cannot support a comparative study.

Keywords: Computational stylistics, interestingness measure, sequential pattern mining, syntactic style

1. INTRODUCTION

Computational stylistic is the subfield of computational linguistics that aims to extract patterns of style characterizing a particular type of texts using some computational and automatic methods [5]. Addressing questions of style, it shares many commonalities with computational authorship attribution [18] in which one assign a text of unknown authorship to one of some candidate authors based on the stylistic information extracted from documents written by them. Rather than concentrating on those subconscious traits that may constitute an author's fingerprint, computational stylistics seeks to study those features of an author's style that are not only distinctive but also intentionally used by the author. Computational stylistic interferes in many other related tasks such as stylistic text classification [8], stylistic-based text generation [7] automatic readability and complexity assessment [12]. However the field of research in which the notion of style find its strong arguments is the computer-assisted literary analysis and computer-based literary criticism. Stylometric techniques have been used for nearly sixty years to study questions relating style (see [17] for a discussion and overview). First works focused more on lexical traits such as word counts, later more complex grammatical traits have been taken into account.

From the methodological point of view, two different types of approach have emerged:

- Classification approaches, that can be simplified as such: an a priori classification is found in literature (such as Shakespeare's comedies vs tragedies); some relevant linguistic features are identified and counted (such as function words) and finally clustering techniques are used to see whether the a priori distinction holds or not [5].
- Hermeneutic approach, in which texts are analysed in order to automatically extract significant features that may later be used by domain experts to produce a better informed and data driven critical analysis of texts [9, 14].

What such approaches often have in common is the fact that the distinctiveness of a text or of some of its elements is measured by comparison to other texts. The present work follows the hermeneutic approach, in that it seeks to extract significant syntactic patterns. However, the proposed methodology is based only on intrinsic evaluation, so that it can be applied to one text at a time.

In our contribution, we present a computational stylistic study of the French classic literature texts based on a data-driven approach where discovering interesting linguistic patterns is done without any prior knowledge.

We propose an objective interestingness measure to assist linguists in studying the syntactic style and in extracting meaningful linguistic patterns from a given author's work. More specifically, this interestingness measure is based on the position in which a pattern appears in the text, rather than its frequency. It is meant to support stylistic textual analysis by:

1. Verifying the degree of importance of each linguistic pattern (syntagmatic segments with gaps) via a new clustering-based interestingness measure.
2. Automatically inducing a list of linguistic features that are significant, representative for an author's work.

Thus, the goal of this paper is to present a practical intrinsic measure for extracting syntactic patterns from texts for stylistic analysis. This measure is motivated by statistical and linguistic considerations. Since it does not rely on the occurrences' counts of the syntactic patterns in texts, this measure can work reasonably well with both large and small text samples and allows the extraction of significant syntactic patterns.

The rest of the paper is structured as follow. We first give a linguistic motivation to our contribution with a brief overview of the methodology in Section 2. Then, in Section 3 we quantitatively analyze the syntactic order to show that the positions of the syntactic forms in the text are more relevant than their frequencies. We present the statistical formulation of the proposed intrinsic interestingness measure in Section 4. Finally a

thorough analysis and discussion of the patterns extracted with this methodology will be given in Section 5 for a selected classic French novel in order to illustrate the kind of stylistic patterns that can be identified.

2. MOTIVATION

In our study, we consider a syntagmatic approach. The text is first segmented into a set of sentences, and then each sentence is mapped into a sequence of syntactic (part of speech or POS-tag) items. For example, the sentence "Le silence profond régnait nuit et jour dans la maison" is first mapped to a sequence of:

$$< DET, NOM, ADJ, VER, NOM, KON, NOM, PRP, DET, NOM, SENT >!$$

Then sequential patterns of a determined length are extracted, such as:

$$< DET >< NOM >< ADJ >$$
$$< NOM >< ADJ >< VER >< NOM >$$
$$< KON >< NOM >< * >< DET >< NOM >$$

However, as sequential pattern mining is known to produce a large quantity of patterns even from relatively small samples of texts, an interestingness measure should be applied on these patterns in order to identify the most important ones. To the best of our knowledge, all the interestingness measures proposed in the literature to deal with this issue are based on the support of the pattern, that is, the frequency of occurrence of those patterns in the texts. However, frequency based methods are argued not to be precise in studying linguistic phenomena unless a huge quantity of texts is used.

Our hypothesis is based on the fact that the most characterizing linguistic patterns should significantly reflect the author's stylistic choice which makes theirs positions of occurrences controlled by the author's purpose, while the irrelevant linguistic patterns are distributed randomly in the text. Following this idea, the assumption made in this approach is that

the higher the importance of a linguistic pattern, the more its occurrences cluster together detaching themselves from a random distribution. By this methodology we search for patterns whose frequency is much higher in single portions of texts than in others, thus making each of them the locally most prominent pattern. The clustering phenomenon can be visualized in Figure 1 where we have plotted the absolute positions in the text of two different syntactical patterns. One can clearly notice that despite the two patterns having the same support (counts of sentences where they appear); they significantly behave differently in terms of their distribution of occurrences' positions. This property gives them a different linguistic relevancy value.

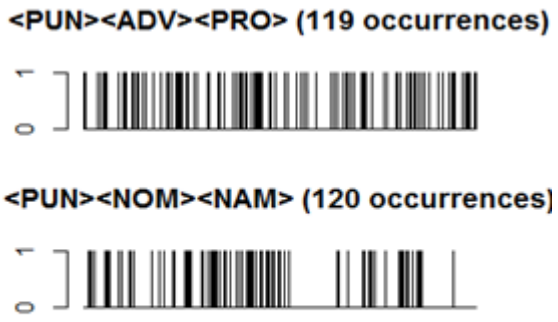

Figure 1. *Positions of occurrences in the text, counted by sentences from the first to the last one, of two different patterns with approximately same support, but with different distribution of positions. (A very thin vertical line is drawn at the position of each occurrence).*

3. QUANTITATIVE ANALYSIS OF THE SYNTACTIC ORDER

The written text is a very syntactically regulated phenomenon in the sense that not all the syntactical combinations are allowed to construct a well-formed syntactic sequence that can carry a semantic meaning.

There are two main factors acting at two different levels that regulate the syntactic order of a text. The first one is the grammar that acts on the phrase level by restricting the syntactic variations via a set of syntactic rules. These syntactic rules forbid certain syntactic sequences that are considered invalid, and allow other valid ones (syntactic sequence that respects these rules). The second element is the genre of the text which acts at the sentence level. In fact it is clear that a text written to be a poem will significantly differ from a text written to be a novel in terms of the syntactic forms that are incorporated on each ones of these two different types of genre. This is due to the linguistic constraints imposed by the rhetoric of the genre. These two elements will introduce a certain statistical order into the syntactic sequences.

Much of the works done so far to extract linguistic patterns - including syntactic patterns – from text is based on some statistical methods that relay mainly on the frequency of these linguistic patterns in that text. This means that the importance and significance of the patterns are evaluated according to its counts in the texts. In fact, this turns out to be not always true. For example, let's consider that we want to extract relevant syntactic patterns that are representative of the style of an author in a classic French novel. If we consider the part of speech 3-gram (simple syntactic pattern) as a minimum unit of the syntactic structure, we find that the statistical hierarchy of this unit can be represented by Zipf's law (see Figure 2). The Zipf low [20] states that the frequency of any unit is inversely proportional to its rank. Thus a frequent unit will occur approximately times as often as the direct following less frequent unit where is a constant that can be empirically deduced. The Zipf low, as known in statistics, has tendency to decrease the significance of the low frequency events, even if they are implicitly relevant to the studied subject. This applies also in the context of linguistics [10].

The property of the Zipf low, which makes the POS 3-gram behave in such a way on the one hand, and the statistical fluctuation on the other hand, increase the weaknesses of the methods based on the frequency as main element to evaluate the

relevancy. In fact, the frequency properties and their application have been discussed in the computational linguistic community. As a result, the frequency-based methods were claimed to be an unsophisticated way to discriminate the relevant linguistic forms in many cases, especially for the case of extracting relevant words from texts [15]. For instances, mutual information estimates based directly on counts are subject to overestimation when the counts involved are relatively small, and z-scores method that are based on the normality assumption substantially overestimate the significance of rare events [6].

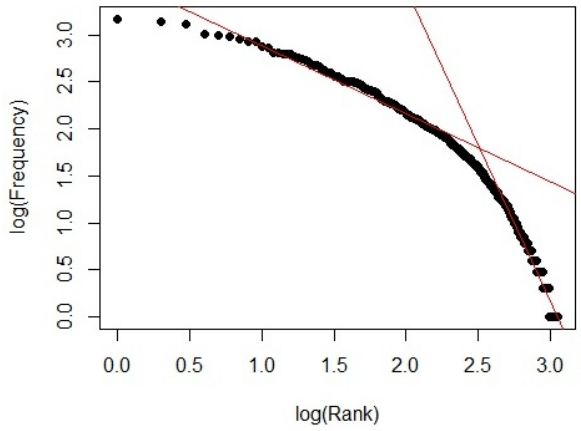

Figure 2. *A plot of POS 3-gram frequency in Hugo's novel "Notre-Dame de Paris". The plot is in log-log coordinates. x is rank of a word in the frequency table; y is the total number of the 3-gram's occurrences*

4. PROPOSED METHODOLOGY

Our method is based on two steps. First a sequential pattern mining algorithm is used to extract sequential syntactic patterns from the text. Secondly, the proposed interestingness measure is applied on the extracted set of syntactic patterns in order to identify the most relevant ones. Section 4.1 introduces some elements about the sequential syntactic pattern extraction

process. Then, the formulation and the statistical details of the proposed interestingness measure are presented in Section 4.2.

4.1. *Extracting sequential syntactic pattern*

Sequential data mining is a data mining subdomain introduced by [1] which is concerned with finding interesting characteristics and patterns in sequential databases, such as biological databases and customer databases in which data is mined to extract behavior sequential patterns. In what follow, for the sake of clarity, we will limit our definitions and annotations to those necessary to understand our experiment.

Table 1. *Sequence database SDB*

Sequence ID	Sequence
1	< a, b, d, e >
2	< a, b, c, e >
3	< b, d, e >

Let's consider a set of literals called items, denoted by $I = \{i_1, ..., i_n\}$. A sequence S (single-item sequence) is an ordered list of items, denoted by $S = <i_1...i_n>$ where $i_1...i_n$ are items. A sequence database *SDB* is a set of tuples *(id, S),* where *id* is the sequence identifier and S a sequence. Interesting characteristics can be extracted from such databases using sequential rules and pattern mining.

We define the sequential pattern as an ordered subset of I formed by the sequence combination of k recurring elements of the set *I.* The sequential pattern admits also the presence of one or many variable items denoted by "*" that can be substituted by any other items. Several algorithms have been developed to efficiently extract this type of patterns [19]. For example, if we run this algorithm on the *SDB* containing the three sequences presented in Table 1, we will get as a result sequential patterns, such us: "$< a >< b >$" with support equal to 2, which means that this pattern is respected by two sequences in the *SDB* (i.e., there exist two sequences of the *SDB* where we find the item *a*, we also find *b* just afterward in the same sequence), or:

"$< b > < * > < e >$" with support equal to 3, which means that this pattern is respected by tree sequences in the *SDB* (i.e., in all the sequences of the *SDB* one can find the item b followed by any item, followed by e).

Quiniou et al. [13] in their study have shown the interest of using sequential data mining methods for the stylistic analysis of large texts. They have shown that relevant and understandable patterns that are characteristic of a specific type of text can be extracted using sequential data mining techniques such as sequential pattern mining. They have considered the text as a set of sentences and each sentence as a sequence of ordered syntactic (POS tags1) or lexical (lemma) items using TreeTagger [16]. Each item in the sequence corresponds to one token (word) in the sentence respecting the order. Using this configuration as input for the sequential pattern mining algorithm, they pointed out that their method is better than machine learning methods such as Hidden Markov Models or Conditional Random Fields, in the sense that it produces outputs that are more understandable by humans. In our study, we will consider a similar configuration to that used by Quiniou et al. The text is first segmented into a set of sentences, then each sentence is mapped into a sequence of syntactic items (POS tags[1]). The whole text will produce a sequential database. Then, sequential patterns of a determined length are extracted from this syntactic sequential database using the sequential pattern algorithm. To avoid the effect of statistical fluctuation on the patterns with small support, we set the minimum support threshold constraint to 1%, that is, we focus only on patterns that are presents in at least 1% of the sentences of the analysed text. Example about syntactic patterns are presented and discussed in Section 5.

4.2. *Filtering the most relevant sequential syntactic pattern*

The positions of the pattern occurrences with support equal to in the ordered set of sentences of the text are denoted by $P = (p_1, p_2, p_3, ...p_n)$ where p_1 is the rank of the sentence in which

[1] POS tag list for French: <www.cis.uni-muenchen.de/~schmid/tools/ TreeTagger/data /french-tagset.html>.

the pattern appears for the -th time $i \in (1,..,n)$ with. The set of distances between two successive occurrences of a pattern can be denoted by $D = (d_1, d_2, d_3,...,d_{n-1})$ where $di = p_{i+1} - p_1$.

Given that configuration, in order to quantify the degree of importance of each linguistic pattern and thus evaluate its relevance based on its clustering behaviour, we use the parameter which is defined as the standard deviation of the distribution of the set of distances normalised by the expected value of such as:

$$\sigma = \frac{\sqrt{E(D^2) - E(D)^2}}{E(D)} \quad (1)$$

For the case where the distribution of the pattern position set P is completely random, one should expect the corresponding distances set D to follow a Geometric distribution and thus to have a parameter $\sigma = 1$, and the larger is the clustering the bigger is σ [11]. However, patterns with different random distributions of D (different random clustering settings in a text) would have a significant difference in their corresponding parameter σ. To avoid this, we normalise the pattern parameter σ with the Geometric parameter σ_{geo} where $\sigma_{geo} = \sqrt{1 - (1/E(d))}$ such as:

$$\sigma_{norm} = \frac{\sigma}{\sigma_{geo}} \quad (2)$$

Such method was already successfully used in physics to quantify energy level of disordered system and in information retrieval to extract key words and keys phrase from informative texts [4].

5. Experimental Validation and Discussion

In this section, some of the significant patterns extracted and ranked in terms of relevancy with the proposed method from Victor Hugo's novel "*Notre-Dame de Paris*" (NDdP) are described and discussed.

In order to appreciate the results, we shall compare them with the mere study of frequent patterns (frequency as interestingness measure). Frequency based approaches may be used in comparative works, as in such cases it is possible to filter out patterns that are frequent in all texts and thus uninteresting. In an intrinsic approach this is not possible, thus the most frequent patterns are not very informative per se:

For example, the Pattern (1): $< VER > < DET > < NOM >$ (support = 3101) is one of the most frequent ones in NDdP. By comparing NDdP with other novels by contemporary authors, we might find out that Hugo under/over uses this syntactic structure, and possibly draw some conclusions. But in isolation this doesn't tell us much, as it is clear that the verb-object structure is very common in French.

The same is true for the Pattern (2): $< NOM > < PRP > < NOM >$ (support = 3101) that presents a simple noun phrase modified by a prepositional phrase which is a frequent structure in French.

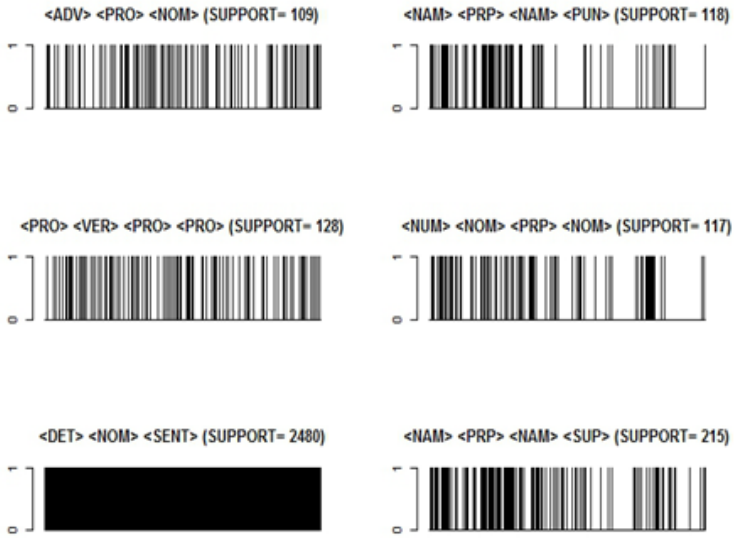

Figure 3. *Positions plot of three relevant syntactic patterns (right column) versus three non-relevant patterns (left column)*

On the contrary, by using the proposed method the extracted patterns seem to bear a strong relation to this particular text, its story line and the literary genre it instantiates, namely that of the historical novel.

Let us here take into account some examples:

- Pattern (3): *< NOM > < PRP > < NAM > < PUN >* support = 340 instances in the text:

  — …tour de Notre-Dame ,
  — …hôtel de Bourbon ,
  — …murailles de Paris ,
  — …prince de Conty :
  — …dauphin de Vienne ;
  — …échevin de Bruxelles ;

In Pattern (3) the proper name is often a location, especially at the beginning of the novel where descriptive parts are more frequent for the purpose of guiding the reader into the topography of medieval Paris. Later it serves the purpose of locating the plot. Other instances of this pattern are used to mention characters, especially historical ones, by their title and provenance. This also is very typical of a literary genre where historical elements are combined with fictional ones.

The skewedness of the distribution of this pattern is also due to the fact that some of the parts of the novel are more descriptive; they serve the purpose of introducing the historical setting and the characters, while other parts develop the action and thus do not introduce many new characters. Other extracted patterns have a similar function.

Pattern (4) is often used to introduce characters by first stating their name and title. It is worth noticing that NDdP presents a plethora of minor characters (see positions plot of this pattern in Figure 3):

- Pattern (4): *< NAM > < PRP > < NAM > < PUN >* support= 118, instances in the text:

- …Marguerite de Flandre,
- …Jehan de Troyes ,
- …Frollo du Moulin ,

The same is true for Pattern (5) which is also used to introduce character's full names (patronymics).

- Pattern (5): *< PUN > < NOM > < NAM >* support = 120

Pattern (6) instead is often instantiated in presentative structures such as (6.A and 6.B), with the topic (here the person) in focus position, used for changes of scenes, to introduce new characters.

In the final part of NDdP, this pattern instantiates other kinds of structures, such as (6.C), which are used to represents actions.

- Pattern (6) : *< PRP > < VER > < NAM >*, support = 113

  - (6.A) Il y avait pourtant une créature humaine que Quasimodo exceptait de sa malice et de sa haine pour les autres, et qu'il aimait autant, plus peut-être que sa cathédrale; ***c'était Claude*** Frollo. [it was Claude]
  - (6.B) ***C' était Quasimodo***, sanglé, cerclé, ficelé, garrotté et sous bonne garde. [it was Quasimodo]
  - (6.C) …acculés à Notre-Dame qu' ils assaillaient encore et ***que défendait Quasimodo***,… [that Quasimodo was defending]

As we have seen, many of the significant patterns extracted with this technique contain the NAM (proper name) tag. This happens not only with sequential pattern mining, but also with other statistical pattern mining methods, as in general proper names are less frequent than other tags and their skewed distribution causes them to emerge in significance measurements. In a study that focuses purely on syntax it may be worth merging this class with the one of common names.

Pattern (7) does not contain proper names, and seems very relevant for the text in question. Among the instances of this

pattern we find many vivid and precise descriptions, as is evident especially in (7.B) where Hugo lists all the different divisions that used to be in charge of the defence of the former stronghold of "Châtelet", in Paris.

- Pattern (7): *< NUM > < NOM > < PRP > < NOM >* support = 117 , instances in the text:

  - (7.A) …le fracas de tous les gros doubles pétards de la Saint-Jean, la décharge de ***vingt arquebuses à croc***, la détonation de cette fameuse serpentine de la Tour de Billy, [twenty arquebuses]
  - (7.B) …les cent ***vingt sergents à cheval***, les cent ***vingt sergents à verge***, le chevalier du guet avec son guet, son sous-guet, son contre-guet et son arrière-guet? [twenty seargents on horse]

The few analysed examples indicate that the presented technique is effective in extracting interesting syntactic patterns from a single text, and this seems particularly promising for the analyses of such texts that, for their characteristics or for historical reasons, cannot support a comparative study as they are, in some way, unique. This might be the case of great poems from the antiquity, such as the Iliad or the Odyssey or even contemporary works whose style is too peculiar for comparison, such as James Joyces's Ulysses.

On the other hand, this technique, as well as other similar ones, prompts the question of what is really captured by significant patterns. Some structures may be significant because they are typical of an author's style, its fingerprint – as we may say borrowing a metaphor often used in attribution studies, or they may be dictated by functional needs, due to the particular topic of the work, or to the conventions of the chosen genre. This is particularly true for syntactic analysis, where the functional constraints on the authorial freedom are more evident.

Using Biber & Conrad [3] definitions, it is worth asking how far it is possible to distinguish style from register and genre when

analysing syntactic structure, especially considering that stylometric studies have given their best results with features such as word and sentence length, or lexical richness. In particular this technique, based on the uneven distribution of patterns, will invariably capture local changes in register motivated by the different elements of the novel (introduction, descriptions, scenes of action, dialogue, ...), alongside with stylistic traits.

It is always hard in linguistics to separate the form from the function. For this reason it is important to study syntactic patterns in the light of the sentences from which they are drawn to avoid false conclusions. Nevertheless the technique seems efficient in demoting those frequent constructions that are typical of French syntax in general without the need of a reference corpus; at the same time the syntactic structure of the extracted patterns and their use in vivid descriptions, in the presentation of characters and in the reconstruction of scenes do seem to resonate with the particular use of language typical of Victor Hugo's prose in NDdP.

6. CONCLUSION

In this paper, we have presented an objective interestingness measure to extract meaningful stylistic syntactic patterns from a given author's work. Our hypothesis is based on the fact that the most characterizing linguistic patterns should significantly adopt a clustering behaviour which detaches them from a random positioning in the text. To evaluate the effectiveness of the proposed method, we conducted an experiment on a classic French Novel. The analysed results show the effectiveness in extracting interesting syntactic patterns from a single text.

Based on the current study, we have identified several future research directions. First, we will explore combining the support with the distribution to calculate the interestingness measure. Second, this study will be expanded to include morpho-syntactic patterns. Third, we intend to experiment with other languages and text sizes using standard corpora employed in the field of computational stylistics at large.

REFERENCES

1. Agrawal, R., Imieli'nski, T. & Swami, A. 1993. Mining association rules between sets of items in large databases. In *ACM SIGMOD Record* (pp. 207-216).
2. Biber, D. 2006. *University Language: A Corpus-based Study of Spoken and Written Registers*. John Benjamins Publishing.
3. Biber, D. & Conrad, S. 2009. *Register, Genre, and Style*. Cambridge University Press.
4. Carpena, P. et al. 2009. Level statistics of words: Finding keywords in literary texts and symbolic sequences. *Physical Review E*, 79/3, 35102.
5. Craig, H. 2004. Stylistic analysis and authorship studies. *A Companion to Digital Humanities*, 3, 233-334.
6. Dunning, T. 1993. Accurate methods for the statistics of surprise and coincidence. *Computational Linguistics*, 19/1, 61-74.
7. Hovy, E. H. 1990. Pragmatics and natural language generation. *Artificial Intelligence*, 43/2, 153-197.
8. Kessler, B., Numberg, G. & Schütze, H., 1997. Automatic detection of text genre. In proceedings of the *35th Annual Meeting of the Association for Computational Linguistics and Eighth Conference of the European Chapter of the Association for Computational Linguistics* (pp. 32-38).
9. Mahlberg, M. 2012. *Corpus Stylistics and Dickens's Fiction*. Routledge.
10. Montemurro, M. A. 2001. Beyond the Zipf-Mandelbrot law in quantitative linguistics. *Physica A: Statistical Mechanics and its Applications*, 300/3, 567-578.
11. Ortuno, M. et al. 2002. Keyword detection in natural languages and DNA. *EPL (Europhysics Letters)*, 57/5, 759.
12. Pitler, E. & Nenkova, A. 2008. Revisiting readability: A unified framework for predicting text quality. In proceedings of the *Conference on Empirical Methods in Natural Language Processing* (pp. 186-195).
13. Quiniou, S. et al. 2012. What about sequential data mining techniques to identify linguistic patterns for stylistics? In *Computational Linguistics and Intelligent Text Processing* (pp. 166-177). Springer.

14. Ramsay, S. 2011. *Reading Machines: Toward an Algorithmic Criticism*. University of Illinois Press.
15. Salton, G. & McGill, M. J. 1983. Introduction to modern information retrieval.
16. Schmid, H. 1994. Probabilistic part-of-speech tagging using decision trees. In proceedings of the *International Conference on New Methods in Language Processing* (pp. 44-49).
17. Siemens, R. & Schreibman, S. 2013. *A Companion to Digital Literary Studies*. John Wiley & Sons.
18. Stamatatos, E. 2009. A survey of modern authorship attribution methods. *Journal of the American Society for information Science and Technology*, 60/3, 538-556.
19. Viger, P. F. et al. 2014. SPMF: A java open-source pattern mining library. *Journal of Machine Learning Research*, 15, 3389-3393.
20. Zipf, G. K. 1949. Human behavior and the principle of least effort.

ACKNOWLEDGMENT

This work was supported by French state funds managed by the ANR within the Investissements d'Avenir programme under reference ANR-11-IDEX-0004-02.

MOHAMED-AMINE BOUKHALED
LIP6 (LABORATOIRE D'INFORMATIQUE DE PARIS 6),
UNIVERSITÉ PIERRE ET MARIE CURIE AND CNRS (UMR7606),
ACASA TEAM, 4, PLACE JUSSIEU,
75252-PARIS CEDEX 05, FRANCE.
E-MAIL: <MOHAMED.BOUKHALED@LIP6.FR>

FRANCESCA FRONTINI
LIP6 (LABORATOIRE D'INFORMATIQUE DE PARIS 6),
UNIVERSITÉ PIERRE ET MARIE CURIE AND CNRS (UMR7606),
ACASA TEAM, 4, PLACE JUSSIEU,
75252-PARIS CEDEX 05, FRANCE.
E-MAIL: <FRANCESCA.FRONTINI@LIP6.FR>

GAUVAIN BOURGNE
LIP6 (LABORATOIRE D'INFORMATIQUE DE PARIS 6),
UNIVERSITÉ PIERRE ET MARIE CURIE AND CNRS (UMR7606),
ACASA TEAM, 4, PLACE JUSSIEU,
75252-PARIS CEDEX 05, FRANCE.
E-MAIL: <GAUVAIN.BOURGNE@LIP6.FR>

JEAN-GABRIEL GANASCIA
LIP6 (LABORATOIRE D'INFORMATIQUE DE PARIS 6),
UNIVERSITÉ PIERRE ET MARIE CURIE AND CNRS (UMR7606),
ACASA TEAM, 4, PLACE JUSSIEU,
75252-PARIS CEDEX 05, FRANCE.
E-MAIL: <JEAN-GABRIEL.GANASCIA@LIP6.FR>