

Experiments to Investigate the Utility of Nearest Neighbour Metrics Based on Linguistically Informed Features for Detecting Textual Plagiarism

Per Almqvist and Jussi Karlgren

Swedish Institute of Computer Science (SICS), Stockholm*
peralmq@kth.se jussi@sics.se

Abstract

Plagiarism detection is a challenge for linguistic models — most current implemented models use simple occurrence statistics for linguistic items. In this paper we report two experiments related to plagiarism detection where we use a model for distributional semantics and of sentence stylistics to compare sentence by sentence the likelihood of a text being partly plagiarised. The result of the comparison are displayed for visual inspection by a plagiarism assessor.

1 Plagiarism detection

Plagiarism is the act of copying or including another author’s ideas, language, or writing, without proper acknowledgment of the original source. Plagiarism analysis is a collective term for computer-based methods to identify plagiarism. (Stein et al., 2007a) Plagiarism analysis can be performed *intrinsically* — a text is examined for internal consistency, to detect suspicious passages that appear to diverge from the surrounding text, or *externally* — a text is inspected with respect to some known corpus to find passages with suspiciously similar content to other text.

In external plagiarism detection, it is assumed that the source document d_{src} for a given plagiarized document d_{plg} can be found in a target document collection D . Typically, plagiarism detection then proceeds in three stages:

1. candidate selection through retrieval of a set of candidate source documents D_{src} is retrieved from D_{plg} ;

This work is performed at SICS, supported by the Swedish Research Council (Vetenskapsrådet) through the project “Distributionally derived grammatical analysis models” as part of Per Almqvist’s M Sc program at the Royal Institute of Technology (KTH), Stockholm

2. candidates d_{src} from D_{src} is compared passage by passage with the suspicious document d_{plg} and every case where a passage from d_{plg} appears to be similar to some passage in some d_{src} is noted;
3. followed by some post-processing to remove false hits.(Stein et al., 2007b; Potthast et al., 2010)

2 PAN workshop series

A series of workshops on Plagiarism Analysis, Authorship Identification, and Near-Duplicate Detection, organised since 2007, have provided the field with a shared task and test materials in the form of gold standard text collections with manually and automatically constructed plagiarised sections marked for experimental purposes. Some of the plagiarised sections are *obfuscated* with word replacement, edits, and permutations. The research results from the workshops are comparable, since they are to a large extent performed on the same materials using the same starting points and same target measures.

Example results relevant to this study (and on the whole none too surprising) are that unobfuscated plagiarism can be detected with a reasonable accuracy by the top plagiarism detectors. The recall decreases slightly with increasing obfuscation and that longer stretches of plagiarised material are easier to detect than shorter segments.(Potthast et al., 2010)

3 Our experimental set-up

The base of the experiment described here is to test a finer-grained analysis of plagiarised texts than other previous work. We use a sentence-by-sentence comparison of the suspicious text (d_{plg}) with all sentences of each target text (d_{src}) in D_{src} using two different similarity measures: one based

Table 1: Stylometric sentence features

Name	Description
arg	argumentative (<i>merely, for sure, ...</i>)
cog	cognitive process (<i>remember, think, ...</i>)
com	complex (average word length > 6 characters or sentence length > 25 words)
date	one or more date references
fin	money symbols or percentage signs
fpp	first person pronouns
le	named entities (person, organization)
loc	location expression
neg	grammatical negations
num	numbers
pa	place adverbials (<i>inside, outdoors ...</i>)
pun	punctuation inside sentence
se	split infinitives or stranded prepositions
spp	second person pronouns
sub	subordinate clauses
ta	time adverbials (<i>early, presently, soon ...</i>)
tim	one or more time expressions
tpp	third person pronouns
uni	symbols representing units of measurement

on overall semantic similarity, the other on specific stylometric measures.

The experiment is not a full scale evaluation of our method but is intended to test the practicability of our approach. Given that we have a suspicious text and some reasonable number of candidate source texts (through some retrieval procedure) — can we detect the likelihood of plagiarism in a text by inspecting the sentence sequence of the suspicious text one by one? This paper reports a selected plot dry run of the methodology performed over a number of sample texts.

3.1 Data

The experiments are performed on the PAN-PC-09(Potthast et al., 2009)¹ corpus since it can be used free of charge for research and contains plagiarized passages which has previously been marked and labeled as plagiarism, so that we know beforehand which passages are plagiarism.

The corpus is divided in two sets, one for training and one for test. The training set is further divided into three parts (D_{plg} , D_{src} , and L). D_{plg} contain the documents that are suspicious and might plagiarize documents in D_{src} , where D_{src} contain only original documents that make out the sources of any plagiarism in D_{plg} , and L is the solution, the labeling that tells us which sentences in D_{plg} that plagiarize what sentence in D_{src} .

¹<http://www.uni-weimar.de/cms/medien/webis/research/corpora/pan-pc-09.html>

3.2 Nearest neighbour metrics

We use cosine similarity (as defined in equation 1) to represent how similar two vectors are.

$$sim_{COS}(\vec{x}, \vec{y}) = \frac{\vec{x} \cdot \vec{y}}{|\vec{x}| |\vec{y}|} = \frac{\sum_{i=1}^n x_i y_i}{\sqrt{\sum_{i=1}^n x_i^2} \sqrt{\sum_{i=1}^n y_i^2}} \quad (1)$$

For every sentence $s \in d_{plg}$ its nearest neighbour score (as defined in equation 2) is calculated.

$$\max(sim_{COS}(s, x)) \text{ for all sentences } x \in D_{src} \quad (2)$$

The nearest neighbour metric has the fortunate feature that a value of 1 describes identical or duplicate vectors. So if we were to find nearest neighbour values of 1 those two sentences would be very alike and therefore we would be able to assume that the newer sentence plagiarizes the older sentence.

In this experiment two settings for the experiment were used. In experiment one below we evaluate how well the nearest neighbour metric of two vectors in a semantic *word-space model* manage to detect plagiarism. In experiment two below we evaluate how well the nearest neighbour metric of two binary vectors based on 19 different stylometric features manage to detect plagiarism.

3.3 Target plots

As an example plagiarism inspection mechanism we plot the nearest neighbour metric with the sentences of a text along the x -axis against the score of the sentence. The objective is to find a stretch of material where several sentences have high nearest neighbour scores. As a comparison we will plot the gold standard plagiarism labeling of respective sentence and let the label for a sentence being plagiarism have value 1 and 0 otherwise. Now we can just plot our nearest neighbour scores and our modified labels against the sentences in the corpus.

3.4 Experiment 1: semantic similarity

The sentences of d_{plg} were compared by semantic similarity using a word-space model (Schütze, 1993) as a base for computing similarity between sentences. Each sentence was represented by the centroid of its constituent words in a word-space trained on the entire test corpus. The implementation was based on previous work on effective word-space models.

“The word-space model is a computational model of word meaning that utilizes the distributional patterns of words collected over large text data to represent semantic similarity between words in terms of spatial proximity.” (Sahlgren, 2006).

Experiments to investigate the utility of nearest neighbour metrics based on linguistically informed features for detecting textual plagiarism

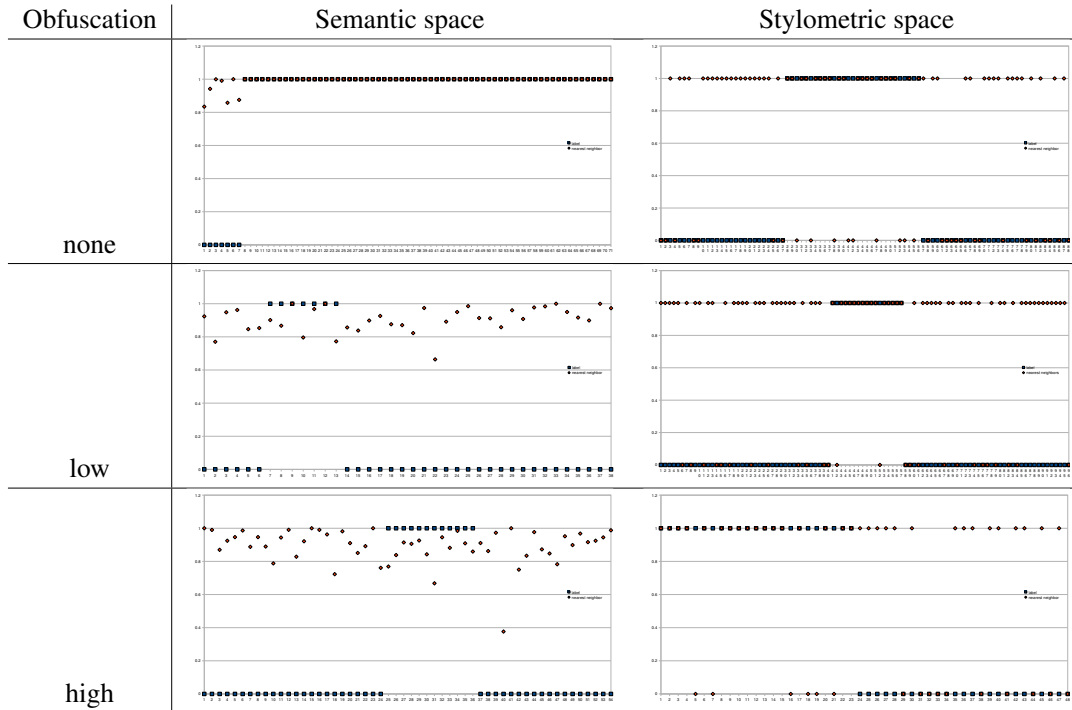


Table 2: Plots of semantic and stylometric similarity for texts with known plagiarized sections

The word-space model, from the work in (Kanerva, 1988) and (Sahlgren, 2006), models the meaning of words according to their distribution, creating a representation of their semantics based on where and how in the text the words appear. The word-space is a high dimensional vector space where every word is represented by a vector. Two words are semantically similar if their respective vectors are similar. For example the words "yellow" and "green" could be argued to have similar semantic meaning. So the vectors for yellow and green should be expected to be similar as seen in figure 1.

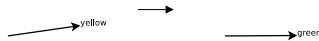


Figure 1: The vectors for the words "yellow" and "green" in a semantic space.

The word-space model is, as its name implies, mainly used to model words. It can however be used to model other linguistic entities such as sentences and documents using workarounds. A sentence can be represented by taking the centroid of the sentence's individual word's vectors. Therefore if the sentence "A yellow car." was changed to "A green car." the centroid ought not to change

too much since the only change to the centroid would be one vector that in the first case represented the word "yellow" and in the second case the word "green" and these vectors should be fairly similar, as seen in figure 2. In our model we use a semantic word-space to model sentences under the hypothesis that if a sentence were to be obfuscated its semantic similarity would be kept. We build a semantic space for the corpus under consideration and assign each sentence a representative centroid vector of 3000 real dimensions for every sentence in the corpus. We then perform, for every sentence vector \vec{s} from D_{plg} , the nearest neighbour search $nn(\vec{s}, s_{src})$ against all the vectors s_{src} in D_{src} .

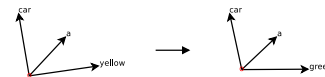


Figure 2: A centroid of a changed sentence.

3.5 Experiment 2: stylometric similarity

The 19 stylometric features that were chosen can be seen in table 1, and were chosen based on the work in (Biber, 1988) and (Karlgrén, 2000). Our intention was to capture the authors' writing styles. We tried to find features that would not

change if another author were to copy the text and even obfuscate it. Therefore we chose features that

- binds the texts to its topic, such as numbers or units of measurement.
- anchors the text to its context, i.e. named entities, location or time.
- captures peculiarities in the author's writing style: split infinitives or stranded prepositions.
- indicates how complex the language is, such as long sentences or subordinate clauses.

For every sentence in D_{plg} we extracted the stylometric features into a 19th dimensional binary vector \vec{f} . We then extracted 470 unique 19th dimensional binary vectors \mathbf{F}_{src} , based on the same stylometric features, from D_{src} . Then we performed the nearest neighbour search $nn(\vec{f}, \mathbf{F}_{src})$ against all the vectors \vec{f}_{src} in \mathbf{F}_{src} .

4 Results and Conclusions

Table 2 shows the results for the nearest neighbour scores for both experiments, run on a test text with a known plagiarized section with the corresponding source text. We have three plots representing different levels of obfuscation of plagiarism, namely; a high level of obfuscation, a low level of obfuscation, and no obfuscation. To determine the effectiveness of each nearness measure, the results (red rhomboids) are displayed together with an indication of which section of the text is plagiarized (blue squares) noted with a score of 1 and a score of 0 for the non-plagiarized sections.

4.1 Experiment 1: semantic similarity

We find that the semantic space model (1) is a good detector for no obfuscation; (2) does not hold up for obfuscated materials, neither for low or high obfuscation since it is based on the presence of each word in the text; and consequentially (3) needs tuning so that specifically topical terms are weighted up compared to less topical terms. This should be done specifically for the topic in the candidate document being examined, since presumably the topic under consideration is the most likely topic to be plagiarized.

4.2 Experiment 2: stylometric similarity

We find that the stylometric similarity score (1) which is a dramatic dimensionality reduction unsurprisingly gives a large number of false positives for all levels of obfuscation; and (2) gives a comparatively high precision even for a high level of obfuscation.

4.3 Directions

Coming experiments will establish whether the combination of the two knowledge sources and the preservation of sequence information in the candidate source texts might provide effective results for a plagiarism detection task. Previous experiments on sequence encoding of stylistic information seem to indicate that sequential information can contain the right type of information to distinguish writing style. (Karlgren and Eriksson, 2007)

References

- Douglas Biber. 1988. *Variation across Speech and Writing*. Cambridge University Press.
- Pentti Kanerva. 1988. *Sparse Distributed Memory*. MIT Press, Cambridge, MA, USA.
- Jussi Karlgren and Gunnar Eriksson. 2007. Authors, genre, and linguistic convention. In *SIGIR Workshop on Plagiarism Analysis, Authorship Identification, and Near-Duplicate Detection*.
- Jussi Karlgren. 2000. *Stylistic Experiments In Information Retrieval*. Ph.D. thesis, Stockholm University.
- Martin Potthast, Andreas Eiselt, Benno Stein, Alberto Barrón-Cedeño, and Paolo Rosso. 2009. PAN Plagiarism Corpus PAN-PC-09.
- Martin Potthast, Alberto Barrón-Cedeño, Andreas Eiselt, Benno Stein, and Paolo Rosso. 2010. Overview of the 2nd international competition on plagiarism detection. In Martin Braschler and Donna Harman, editors, *Notebook Papers of CLEF 2010 LABs and Workshops*. CLEF, Padua, Italy.
- Magnus Sahlgren. 2006. *The Word-Space Model*. Ph.D. thesis, Stockholm University.
- Hinrich Schütze. 1993. Word space. In *Proceedings of NIPS'93*, San Francisco. Morgan Kaufmann.
- Benno Stein, Moshe Koppel, and Efstathios Stamatatos. 2007a. Plagiarism analysis, authorship identification, and near-duplicate detection (PAN 07). *SIGIR Forum*, 42(2):68–71.
- Benno Stein, Sven Meyer zu Eissen, and Martin Potthast. 2007b. Strategies for retrieving plagiarized documents. In *Proceedings of SIGIR 07*. ACM.