THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Reinvestigating the Classification Approach to the Article and Preposition Error Correction

Citation for published version:

Grundkiewicz, R & Junczys-Dowmunt, M 2018, Reinvestigating the Classification Approach to the Article and Preposition Error Correction. in Z Vetulani, J Mariani & M Kubis (eds), *Human Language Technology. Challenges for Computer Science and Linguistics*. Lecture Notes in Computer Science (LNCS), vol. 10930, Springer, Pozna, Poland, pp. 112-122, 7th Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics, Poznan, Poland, 27/11/15.
https://doi.org/10.1007/978-3-319-93782-3_9

Digital Object Identifier (DOI):

[10.1007/978-3-319-93782-3_9](https://doi.org/10.1007/978-3-319-93782-3_9)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Human Language Technology. Challenges for Computer Science and Linguistics

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Reinvestigating the Classification Approach to the Article and Preposition Error Correction

Roman Grundkiewicz and Marcin Junczys-Dowmunt

Adam Mickiewicz University
ul. Wieniawskiego 1, 61-712 Poznań, Poland
{romang,junczys}@amu.edu.pl

Abstract. In this work, we reinvestigate the classifier-based approach to article and preposition error correction going beyond linguistically motivated factors. We show that state-of-the-art results can be achieved without relying on a plethora of heuristic rules, complex feature engineering and advanced NLP tools. A proposed method for detecting spaces for article insertion is even more efficient than methods that use a parser. We examine automatically trained word classes acquired by unsupervised learning as a substitution for commonly used part-of-speech tags. Our best models significantly outperform the top systems from CoNLL-2014 Shared Task in terms of article and preposition error correction.

Keywords: grammatical error correction, article errors, preposition errors, CoNLL-2014 shared task, detecting omitted words

1 Introduction

In the field of grammatical error correction (GEC), a large effort is made to design models and algorithms that incorporate linguistic knowledge. Heuristic rules, advanced tools or resources for natural language processing that were not created specifically with grammatical error correction in mind are commonly used. This results in a high degree of complexity with modest gains in overall performance. Results are difficult to reproduce and the integration of different systems is complicated. We believe, in accordance with Occam’s Razor, that between two models that solve the same problem on similar levels of quality the simpler one is to be preferred.

In this work, we reinvestigate the classifier-based grammatical error correction paradigm by reducing its dependence on heuristic rules and advanced natural language processing tools. We focus on two of the most frequent error types among English as a second language (ESL) learners: article and preposition errors.

The only features we allow ourselves to use are simple n -gram features of: surface level tokens, part-of-speech (POS) tags, and automatically trained word classes (AWC). Where possible we try to replace POS tags with AWC tags. The latter are language-independent tags produced by clustering vector space representations of words which in turn are learnt on large unannotated text [14].

These have been applied to a variety of NLP problems, such as document classification [21], statistical machine translation [12] or named entity recognition, and chunking [24], usually with beneficial effects.

Our main contributions are the following: Firstly, a new contextual method for detecting omitted articles is introduced that in practice outperforms previous methods. Secondly, we apply unsupervised word classes to classifier-based GEC. Finally, we show that it is possible to achieve state-of-the-art results for article and preposition error correction with almost no linguistic knowledge.

The remainder of the paper is organized as follows: Section 2 reviews recent research. Data sets, classification algorithms, feature sets, and evaluation schemes are described in Section 4. Section 3 deals with detection of spaces for potential article insertions. In Section 4, we present our results and compare them with top systems from the CoNLL 2014 shared task [16]. Conclusions and future work are presented in Section 5.

2 Related work

In this section we briefly discuss related work with predominantly classifier-based approaches which focused on correcting mistakes in article and preposition usage. For more comprehensive description of the field we refer the reader to the work of Leacock et al. [13] and the recent CoNLL shared tasks [17, 16].

The majority of researchers use lexical word forms, POS tags and structural information from shallow parser when designing features for article error correction classifiers. Features that encode linguistic knowledge are extracted, for example combinations of words preceding the article and a head word of the identified noun phrase [10, 8].

For instance, Rozovskaya et al. [18] design high-level features that encode POS tags and shallow parse properties. The authors show that adding rich features to the baseline system that uses only word n -grams is helpful. However, they do not compare these rich features with simple POS n -grams.

Features used for preposition error correction are usually less complex and base on lexical forms of surrounding words [11, 20]. Linguistically more complex knowledge is encoded in features that make use of various aspects of preposition complements [23] or additional features derived from a constituency and a dependency parse trees [22].

The results of Rozovskaya et al. [18, 19] are most similar to our work as all features are lexical, but the only type of n -grams tested in this work are pure word n -grams.

2.1 Confusion sets

Typically, a confusion set for article and determiner error correction consists of three units: $\{a, the, \emptyset\}$ ¹. This covers article insertion, deletion, and substitution

¹ $\emptyset$ stands for the English zero article.

errors. The distinction between *a* and *an* is usually made with heuristic rules during postprocessing. Since we made a point of not using any heuristic rules, our confusion set comprise both indefinite article variants, taking the final form: $\{a, an, the, \emptyset\}$.

In preposition error correction it is common to include in confusion set the top n most frequent English prepositions [8, 20, 2]. We restrict ourselves to the top twelve prepositions: $\{in, at, on, for, since, with, to, by, about, from, of, as\}$ that cover 88.6% of all preposition errors in NUCLE (see Section 2.2). In contrast to previous studies that consider only incorrectly selected prepositions, we handle also extraneous and (for final models) missing ones.

2.2 Data sets

For training and testing our models we use various versions of two learner data resources: the NUS Corpus of Learner English and the Lang-8 corpus. A brief summary of used corpora is presented in Tab. 1.

Table 1. Basic statistics of data sets used in experiments: size in sentences and error rates (in %) for article and preposition errors. For TS-2014 values are reported for two annotators separately.

Corpus	Size	ER _{art}	ER _{prep}
NUCLE	57,151	6.68	2.34
TS-2013	1,381	18.27	5.39
TS-2014 A0	1,312	10.23	5.08
TS-2014 A1	1,312	13.72	6.72
L8-NAIST	2,215,373	15.86	7.61
L8-WEB	3,386,887	18.55	9.22

The NUS Corpus of Learner English (NUCLE) [4] consists of 1,414 essays (57,151 sentences) which cover a wide range of topics, such as environmental pollution and health care. It was used as training data in two editions of the CoNLL shared task on Grammatical Error Correction [17, 16].

We also make use of the official test sets from the shared tasks (TS-2013 and TS-2014). This data covers similar topics as NUCLE, but is smaller (1,381 and 1,312 sentences) and has higher frequencies of both error types.

By “Lang-8” we refer to a collection of posts scrapped from a language exchange social networking website named Lang-8². We use the English part of the publicly available “Lang-8 Learner Corpora v1.0” [15] (L8-NAIST).

Furthermore, we have scrapped recent data from the Lang-8 website which resulted in a resource (L8-WEB) that is about one and a half times larger than L8-NAIST.

² <http://lang-8.com/>

2.3 Classification algorithm

Our largest models are trained on over 4 million training examples represented as binary feature vectors of a length that exceeds 1.5 million features. Therefore, we decided to use the L2-regularized logistic regression from LIBLINEAR [5] which supports large-scale multi-class classification. Logistic linear regression has been used before for correction of both, article and preposition errors [23, 11, 2].

2.4 Feature sets

During experiments we use various combinations of the following features:

- *source* — a source confused word encountered in the input text, i.e. the original article or preposition.
- *tokens* — n -grams of lowercased tokens around the confused word. All n -grams have lengths between one and four, and include or are adjacent to the position of confused words.
- *POS* — n -grams of part-of-speech tags obtained by the Stanford Part-Of-Speech Tagger. The tagset consists of 43 tags.
- *AWC* — n -grams of automatic word classes created with the *word2vec* toolkit³ [14]. The number of clusters and vector length were set to 200. Other than that, default options were used. We learnt word vectors from 75 millions of English sentences extracted from Common Crawl data⁴.
- *mixtags* — n -grams that consist of mixed tokens and tags, e.g. for tokens w_1, w_2, w_3 , and corresponding tags t_1, t_2, t_3 , the mixed n -grams are: t_1w_2 , w_1t_2 , $t_1w_2w_3$, $w_1t_2w_3$, $w_1w_2t_3$, $t_1t_2w_3$, $w_1t_2t_3$, etc.

For each word included in a confusion set encountered in the to-be-corrected text, we extract specific features which are later converted to binary feature vectors.

It should be emphasized that experimenting with various vector space representation models, size of space dimensions and number of clusters are not within the scope of this work.

2.5 Evaluation

We use the evaluation scheme and official test sets from the CoNLL-2014 shared task [16]. The system outputs submitted by participants are publicly available⁵, so that we can easily compare our models with top systems from this competition. The participants were free too use all resources that were publicly available, in particular the NUCLE corpus and test set from 2013.

³ <https://code.google.com/p/word2vec/>

⁴ <https://commoncrawl.org/>

⁵ <http://www.comp.nus.edu.sg/~nlp/conll14st.html>

System performance is measured by the MaxMatch (M^2) metric [3] which computes an $F_{0.5}$ score for the proposed corrections against a gold standard that has been similarly annotated as NUCLE.

The original test sets contain annotations of errors from 28 error categories. Evaluation focused on specific errors only results in very low recall in a 28 error type context, which disturbs the tuning process as well as final results due to harmonic properties of F-score. Therefore, we have modified gold standards for each test set by preserving only annotations for which the erroneous or corrected texts concern words from our confusion set (keeping deletions and insertions as well).

This method works better for us than rely on the original error categories, since many annotations that involve articles or prepositions are categorized differently (e.g. many article deletions are categorized as “local redundancy”).

3 Detection of article omissions

Article omissions represent a majority of article and determiner errors, for example, in NUCLE they constitute about 61.38% of all article errors. The most common solution for detecting positions where an article might have been incorrectly omitted is to use a shallow parser to identify spaces occurring before noun phrases [10, 18]. All noun phrases headed by a personal or demonstrative pronoun are excluded. Some research extends this by taking into account additional spaces following a preposition or a verb even when these are not identified by the parser.

On the other hand, a naive method which includes every space as a potential position for article insertion is considered to produce a lot of noise.

3.1 Detection by context comparison

We tested a new method of detection of spaces for potential article insertions based on the comparison of surrounding context. The proposed method consists of a training and a detecting stage.

During training, we extract n -grams from a text corpus consisting of l tokens to the left and r tokens to the right of each occurrence of words from the confusion set. Next, in the to-be-corrected text we flag each space for which a matching n -gram from the set of n -grams extracted during the training stage is found. Changing the minimum count c required for n -grams to be used for detection allows for control of the number of detected spaces. This procedure can be used with token n -grams and POS or AWC tags.

We estimated experimentally the values $l = 1$ and $r = 3$ for article errors, and $l = 3$, $r = 1$ for preposition errors (final models only). The n -grams were trained on a part of English Common Crawl Corpus consisting of ca. 75 million sentences.

3.2 Comparison of detection methods

We compared several methods for finding spaces for potential article insertions in the task of zero article detection. We used the entire NUCLE corpus as test set. The only positive class during evaluation (true positive) was the proper detection of a space where, according to the annotation, an article is missing. A good method should achieve a high recall and a low false positive rate (FPR). Results on NUCLE are presented in Tab. 2.

Table 2. Comparison of various methods for detecting spaces for potential article insertions. Results for NUCLE corpus: number of true positives (TP) and false negatives (FP), false-positive rate (FPR) and recall (R).

Method	TP	FP	FPR	R
naive	3,346	984,307	91.88	100.00
NP	2,871	186,484	68.20	85.80
NP _{verb,prep}	3,041	324,531	78.87	90.88
tokens ₅	1,059	35,543	29.02	31.56
AWC ₅	2,797	348,746	80.04	83.34
AWC ₅₀	2,157	178,290	67.22	64.27
AWC ₅₀₀	1,159	60,402	40.99	34.54
POS ₅₀	3,167	527,534	85.85	94.37
POS ₅₀₀	2,901	315,537	78.40	86.44
POS ₅₀₀₀	2,520	170,078	66.17	75.09

A naive method (*naive*) detects all of 984,307 spaces between words excluding spaces before and after *a*, *an* or *the*. A method that uses a shallow parser (*NP*) results in recall of 85.80, similarly to methods based on AWC n -grams with $c = 5$ (AWC₅) and POS n -grams (POS₅₀₀). But the latter almost double the number of false positives. Enforcing similar FPR requires to set $c = 50$ for AWC and $c = 5000$ for POS tags. For the context coverage method, recall can be adjusted by setting c , which should be determined based on the size of the training data. In the experiment, all n -grams were extracted from a part of English Common Crawl Corpus [1] consisting of ca. 75 million sentences.

We further evaluate these methods in the article error correction task in Section 4.1

4 Experimental results

We use 4-fold cross validation on NUCLE to adjust threshold values of the minimum classifier confidence required to accept its prediction. During each of the steps, additional data in the form of Lang-8 corpora is added as training data. Then, we train the classifier again on the entire data and for final evaluation we use an averaged confidence threshold.

To prevent the classifier from keeping the input text unchanged [2], the error rate of the training data was increased by randomly removing correct sentences. We experimentally set the error rate to 30% for article errors and to 20% for preposition errors. This procedure keeps 873,917 and 1,660,896 sentences in L8-NAIST and L8-WEB respectively for training article models, and 1,219,127 sentences in L8-WEB for training preposition models.

The TS-2013 is used to determine an error rate for tuning data in cross validation as there is a significant disproportion in error rates. We report results on both test sets, for article and for preposition models. The higher results on the CoNLL-2014 test set are due to it has been annotated by two annotators.

4.1 Methods for detecting article omissions

In order to compare the various methods of detecting spaces for possible article insertion (Section 3), we used the tuned L8-NAIST corpus as training data with feature set consisting of token n -grams. Table 3 presents the results.

Table 3. The comparison of different methods for article omission detection. All models are trained on L8-NAIST and use *source* and *tokens* features.

Method	TS-2013			TS-2014		
	P	R	$M_{0.5}^2$	P	R	$M_{0.5}^2$
naive	57.58	3.66	14.59	68.97	12.20	35.71
NP	54.44	9.44	27.87	58.33	20.83	42.89
$NP_{verb,prep}$	51.06	9.25	26.82	53.03	20.71	40.42
POS_{5000}	54.17	7.51	24.16	65.08	23.70	48.24
AWC_{50}	54.55	8.09	25.39	62.26	19.53	43.31

A naive method gives the lowest $F_{0.5}$ scores due to the high precision but low recall. Using a lower error rate in the training data shows a similar effect. Methods based on a shallow parser (NP) are more effective without augmenting them with spaces after each verb and preposition ($NP_{verb,prep}$) on both test sets. The proposed methods that compare surrounding context are significantly better on TS-2014 and reach slightly lower results on TS-2013. It is unclear why AWC n -grams are more effective than POS n -grams for TS-2013 and vice versa for TS-2014.

We also experiment with applying the proposed methods to handle missed preposition errors in our final models (Tab. 5). This increases the recall, since it enables making corrections that can not be detected otherwise, but may reduce precision.

4.2 Different feature sets and final models

Next, we compare different feature sets in Tab. 4. For article models we chose a method of detecting omissions that uses AWC n -grams due to its speed and

Table 4. Final results on the CoNLL-2013 and CoNLL-2014 test sets for article (top) and preposition (bottom) error correction. Article models are trained on L8-NAIST and use AWC_{50} method for omission detection. Preposition models are trained on L8-WEB. All models use *source* features.

Feature set	TS-2013			TS-2014		
	P	R	$M_{0.5}^2$	P	R	$M_{0.5}^2$
tokens	53.01	8.48	25.85	62.30	21.84	45.45
tokens+POS	45.22	10.02	26.56	60.20	32.42	51.39
tokens+POS+mix _{POS}	40.22	13.87	29.15	54.78	33.69	48.69
tokens+AWC	44.92	10.21	26.74	63.64	28.16	50.83
tokens+AWC+mix _{AWC}	42.31	12.72	28.87	56.38	29.44	47.66
tokens+POS+AWC	30.10	17.34	26.24	41.74	49.51	43.09
tokens	42.42	7.37	21.74	70.00	17.36	43.57
tokens+POS	40.00	8.42	22.86	59.46	18.03	40.74
tokens+POS+mix _{POS}	34.09	7.89	20.49	48.72	14.96	33.57
tokens+AWC	36.36	8.42	21.86	67.65	19.66	45.45
tokens+AWC+mix _{AWC}	24.49	6.32	15.54	50.94	21.95	40.30
tokens+POS+AWC	37.21	8.42	22.10	68.75	19.13	45.27

simplicity. For preposition models we used L8-WEB corpus as training data to get a sufficient number of training examples.

Models trained only on lexical features result in $F_{0.5}$ values that are slightly lower than results achieved by models that use more complex features. Using POS or AWC n -grams shows improvement in performance for both, article and preposition models. Although adding mixed n -grams is shown to improve performance in contextual spell checking [7], in our experiments it has a positive effect only for articles on TS-2013.

The results for final models trained on L8-WEB corpus are presented in Tab. /reftab:final. Training article models on larger corpus shows further improvement since more training examples are used. It also shows that POS tags (51.56) are more effective in article error correction than AWC tags (49.10).

For preposition errors, the highest result on TS-2014 (52.63) is achieved by a model using tokens and AWC tags and handling preposition omissions. Further investigation of automatic word classes and various numbers of classes is required.

4.3 Top systems from the CoNLL-2014 shared task

Finally, we compare our best models with top three systems participating in the CoNLL-2014 shared task.

The best system [6] (CAMB) participating in the task uses a hybrid approach, which includes both a rule-based and an SMT system augmented by a large web-based language model. The system of Rozovskaya et al. [19] (CUUI) for article error correction makes use of the averaged perceptron algorithm and POS-tagger

Table 5. Final results on the CoNLL-2013 and CoNLL-2014 test sets for article (top) and preposition (bottom) error correction. All models are trained on L8-WEB and use *source* features.

System	TS-2013			TS-2014		
	P	R	$M_{0.5}^2$	P	R	$M_{0.5}^2$
POS ₅₀₀₀ ; tokens+POS	47.26	13.29	31.28	57.89	35.87	51.56
AWC ₅₀ ; tokens+AWC	42.13	14.45	30.46	55.56	33.52	49.10
POS ₅₀₀₀ ; tokens+POS	34.78	8.42	21.39	58.14	20.66	42.66
AWC ₅₀ ; tokens+AWC	36.00	9.47	23.08	75.68	23.73	52.63

and chunker outputs to generate some of its features and correction candidates. For preposition errors a naive Bayes classifier is trained on n -grams counts from the Google n -gram corpus. AMU [9] is a phrase-based SMT system combining large training resources, task-specific parameter tuning and features.

The systems participating in the shared task were free to use test data from 2013. In addition to the NUCLE and test set from CoNLL-2013, all systems make use of other resources that are significant in size. The CAMB system uses Cambridge Learner Corpus, which is comparable in size with Lang-8 data. A module for preposition error correction in the CUUI system is trained on the Google 1T 5-gram Corpus. The AMU system is trained on data scraped from Lang-8 of similar size to our L8-WEB corpus.

Table 6. Top systems from the CoNLL-2014 shared task.

System	ArtOrDet			Prep		
	P	R	$M_{0.5}^2$	P	R	$M_{0.5}^2$
CAMB	39.00	65.00	42.39	41.15	51.63	42.89
CUUI	28.41	72.06	32.32	32.04	26.61	30.78
AMU	40.54	25.28	36.17	46.05	28.00	40.79
this work	57.89	35.87	51.56	75.68	23.73	52.63

System outputs submitted by participants contain corrections of errors of various types. Thus, we removed corrections that do not concern words from confusion sets (i.e. from system outputs we extracted corrections that concern article or preposition errors only), similarly as reported for the official test sets. Results are presented in Tab. 6.

Our best model for article error correction trained on token and POS features significantly beats the CAMB system by nearly 10% F-score (42.39 vs. 51.56). The top preposition model that uses AWC features and handles preposition omissions outperforms the top system from CoNLL-2014 in similar amount (42.89 vs. 52.63). We generally achieve a higher precision and lower recall than other systems.

5 Conclusions and future work

In this paper we reinvestigated the classifier-based approach in grammatical error correction by reducing the linguistic knowledge hidden in many aspects of system development. We have shown that state-of-the-art results can be achieved without applying a multitude of heuristic rules, complex feature engineering, and advanced NLP tools.

Although, for article error correction the best performance is achieved by models trained on POS n -grams, AWC n -grams also outperform lexical features and top systems participating in the CoNLL-2014 shared task. For preposition error correction, models that use AWC features and allow preposition insertions outperform other systems. Our results have shown that the proposed simple contextual method for detecting omitted articles is competitive with methods relying on chunker outputs.

This work allows to believe that automatic word classes trained with unsupervised methods are promising substitution for part-of-speech tags at least in some applications.

In the future, we plan a deeper examination of the application of automatic word classes to GEC. Other models for unsupervised learning of word representations should be tested, as well as different numbers of word clusters.

Acknowledgements. This work has been funded by the National Science Centre, Poland (Grant No. 2014/15/N/ST6/02330).

References

1. Buck, C., Heafield, K., Van Ooyen, B.: N-gram counts and language models from the common crawl. In: LREC. vol. 2, p. 4 (2014)
2. Cahill, A., Madnani, N., Tetreault, J.R., Napolitano, D.: Robust systems for preposition error correction using Wikipedia revisions. In: NAACL-HLT. pp. 507–517 (2013)
3. Dahlmeier, D., Ng, H.T.: Better evaluation for grammatical error correction. In: NAACL-HLT. pp. 568–572 (2012)
4. Dahlmeier, D., Ng, H.T., Wu, S.M.: Building a large annotated corpus of learner english: The NUS Corpus of Learner English. In: BEA8 Workshop. pp. 22–31 (2013)
5. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: A library for large linear classification. JMLR 9, 1871–1874 (2008)
6. Felice, M., Yuan, Z., Andersen, Ø.E., Yannakoudakis, H., Kochmar, E.: Grammatical error correction using hybrid systems and type filtering. CoNLL pp. 15–24 (2014)
7. Fossati, D., Di Eugenio, B.: A mixed trigrams approach for context sensitive spell checking. In: Computational Linguistics and Intelligent Text Processing, pp. 623–633 (2007)
8. Gamon, M., Gao, J., Brockett, C., Klementiev, A., Dolan, W.B., Belenko, D., Vanderwende, L.: Using contextual speller techniques and language modeling for ESL error correction. In: IJCNLP. vol. 8, pp. 449–456 (2008)

9. Grundkiewicz, R., Junczys-Dowmunt, M.: The AMU system in the CoNLL-2014 shared task: Grammatical error correction by data-intensive and feature-rich statistical machine translation. *CoNLL* pp. 25–33 (2014)
10. Han, N.R., Chodorow, M., Leacock, C.: Detecting errors in english article usage by non-native speakers. *JNLE* 12(02), 115–129 (2006)
11. Han, N.R., Tetreault, J.R., Lee, S.H., Ha, J.Y.: Using an error-annotated learner corpus to develop an ESL/EFL error correction system. In: *LREC* (2010)
12. Koehn, P., Hoang, H.: Factored translation models. In: *EMNLP-CoNLL*. pp. 868–876 (2007)
13. Leacock, C., Chodorow, M., Gamon, M., Tetreault, J.: Automated grammatical error detection for language learners. *Synthesis lectures on human language technologies* 3(1), 1–134 (2010)
14. Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: *Advances in Neural Information Processing Systems*. pp. 3111–3119 (2013)
15. Mizumoto, T., Hayashibe, Y., Komachi, M., Nagata, M., Matsumoto, Y.: The effect of learner corpus size in grammatical error correction of ESL writings. In: *COLING*. pp. 863–872 (2012)
16. Ng, H.T., Wu, S.M., Briscoe, T., Hadiwinoto, C., Susanto, R.H., Bryant, C.: The CoNLL-2014 shared task on grammatical error correction. In: *CoNLL*. pp. 1–14 (2014)
17. Ng, H.T., Wu, S.M., Wu, Y., Hadiwinoto, C., Tetreault, J.: The CoNLL-2013 shared task on grammatical error correction. In: *CoNLL* (2013)
18. Rozovskaya, A., Chang, K.W., Sammons, M., Roth, D.: The University of Illinois system in the CoNLL-2013 shared task. In: *CoNLL*. pp. 13–19 (2013)
19. Rozovskaya, A., Chang, K.W., Sammons, M., Roth, D., Habash, N.: The Illinois-Columbia system in the CoNLL-2014 shared task pp. 34–42 (2014)
20. Rozovskaya, A., Roth, D.: Generating confusion sets for context-sensitive error correction. In: *EMNLP*. pp. 961–970 (2010)
21. Sebastiani, F.: Machine learning in automated text categorization. *ACM computing surveys (CSUR)* 34(1), 1–47 (2002)
22. Tetreault, J., Foster, J., Chodorow, M.: Using parse features for preposition selection and error detection. In: *ACL*. pp. 353–358 (2010)
23. Tetreault, J.R., Chodorow, M.: The ups and downs of preposition error detection in ESL writing. In: *COLING*. pp. 865–872 (2008)
24. Turian, J., Ratinov, L., Bengio, Y.: Word representations: a simple and general method for semi-supervised learning. In: *ACL*. pp. 384–394 (2010)