

Unsupervised Context-Sensitive Spelling Correction of English and Dutch Clinical Free-Text with Word and Character N-Gram Embeddings

Pieter Fivez
Simon Šuster
Walter Daelemans

PIETER.FIVEZ@UANTWERPEN.BE
SIMON.SUSTER@UANTWERPEN.BE
WALTER.DAELEMANS@UANTWERPEN.BE

CLiPS, University of Antwerp, Prinsstraat 13, 2000 Antwerp, Belgium

Abstract

We present an unsupervised context-sensitive spelling correction method for clinical free-text that uses word and character n-gram embeddings. Our method generates misspelling replacement candidates and ranks them according to their semantic fit, by calculating a weighted cosine similarity between the vectorized representation of a candidate and the misspelling context. To tune the parameters of this model, we generate self-induced spelling error corpora. We perform our experiments for two languages. For English, we greatly outperform off-the-shelf spelling correction tools on a manually annotated MIMIC-III test set, and counter the frequency bias of a noisy channel model, showing that neural embeddings can be successfully exploited to improve upon the state-of-the-art. For Dutch, we also outperform an off-the-shelf spelling correction tool on manually annotated clinical records from the Antwerp University Hospital, but can offer no empirical evidence that our method counters the frequency bias of a noisy channel model in this case as well. However, both our context-sensitive model and our implementation of the noisy channel model obtain high scores on the test set, establishing a state-of-the-art for Dutch clinical spelling correction with the noisy channel model.¹

1. Introduction

The problem of automated spelling correction has a long history, dating back to the late 1950s.² Traditionally, spelling errors are divided into two categories: non-word misspellings, the most prevalent type of misspellings, where the error leads to a nonexistent word, and real-word misspellings, where the error leads to an existing word, either caused by a typo (e.g. *I hote* → *hope so*), or as a result of grammatical (e.g. *their - there*) or lexical (e.g. *aisle - isle*) confusion. The spelling correction task can be divided into three subtasks: detection of misspellings, generation of replacement candidates, and ranking of these candidate replacements to correct the misspelling. The nature of the detection subtask is dependent on the type of error: non-word misspellings are typically defined as tokens absent from a reference lexicon, while for real-word misspellings, the detection task is postponed by considering all tokens as replaceable, using the confidence of the candidate ranking module to determine which tokens should be treated as misspellings. The generation of replacement candidates is typically performed by including all items from a lexicon which fall within a pre-defined edit distance of the misspelling (e.g. all items within a Levenshtein distance of 3). The ranking component is the most complex of the three subtasks, and is the main topic of this paper.

The genre of clinical free-text poses an interesting challenge to the spelling correction task, since it is notoriously noisy. English corpora contain observed spelling error rates which range from 0.1% (Liu et al. 2012) and 0.4% (Lai et al. 2015) to 4% and 7% (Tolentino et al. 2007), and even 10% (Ruch et al. 2003). Moreover, clinical text also has variable lexical characteristics, caused by a broad range of domain- and subdomain-specific terminology and language conventions. These properties of clinical text can render traditional spell checkers ineffective (Patrick et al. 2010). Recently, Lai et al. (2015) have achieved nearly

1. Source code, which includes a script to extract the annotated English test data from MIMIC-III (for those who have access to the corpus), can be found at <https://github.com/clips/clinspell>. Due to privacy concerns, we are not allowed to share the annotated Dutch test data.

2. A good overview is given by Mitton (2010) and Jurafsky and Martin (2016).

80% correction accuracy on a test set of clinical notes with their noisy channel model. However, their ranking model does not use any contextual information, while the context of a misspelling can provide important clues for the spelling correction process, for instance to counter the frequency bias of a context-insensitive system based on corpus frequency. As an example, consider the misspelling *goint* → *going* present in the MIMIC-III (Johnson et al. 2016) clinical corpus. While in many domains, *going* will be a relatively frequent word type and will consequently be picked by a corpus frequency-based system, it is actually outnumbered in MIMIC-III by the more prevalent word types *joint* and *point*, which are other replacement candidates for the same misspelling. In other words, corpus frequency is not a reliable metric in such cases. Flor (2012) also pointed out that ignoring contextual clues harms performance where a specific misspelling maps to different corrections in different contexts, e.g. *iron deficiency due to enemia* → *anemia* vs. *fluid injected with enemia* → *enema*. A noisy channel model like the one by Lai et al. (2015) will choose the same item for both corrections.

Our proposed unsupervised context-sensitive method exploits contextual clues by using neural embeddings to rank misspelling replacement candidates according to their semantic fit in the misspelling context. Neural embeddings have proven useful for a variety of related tasks, such as unsupervised normalization (Sridhar 2015) and reducing the candidate search space for spelling correction (Pande 2017). We hypothesize that, by using neural embeddings, our method can counter the frequency bias of a noisy channel model. We test our system on manually annotated misspellings from the MIMIC-III corpus. We also conduct experiments on Dutch data, since there is still a need for a Dutch spelling correction method for clinical free-text (Cornet et al. 2012). By replicating our English research setup for Dutch, we simultaneously examine the language adaptability of our context-sensitive model, and establish a state-of-the-art for Dutch clinical spelling correction. We test our Dutch model on manually annotated misspellings from clinical records collected at the Antwerp University Hospital (UZA). In our experiments for both English and Dutch, we focus on already detected non-word misspellings for developing and testing our spelling correction method, following Lai et al (2015). Note that our method could also be applied to real-word errors. However, since our strategy for collecting an empirical test set of misspellings, which we describe in section 3.4, can not be used for real-word errors, we do not address them in this article.

2. Approach

Since we focus on already detected non-word misspellings, our system only deals with two subtasks of the spelling correction task, namely, generating candidate replacements and ranking them.

2.1 Candidate Generation

We generate replacement candidates in 2 phases, using the reference lexicons described in section 3.1. First, we extract all items within a Damerau-Levenshtein edit distance of 2 from a reference lexicon. Secondly, to allow for candidates beyond that edit distance, we also apply the Double Metaphone matching popularized by the open source spell checker Aspell³. This algorithm converts lexical forms to an approximate phonetic consonant skeleton, and matches all Double Metaphone representations within a Damerau-Levenshtein edit distance of 1. The Double Metaphone representation is an intentionally approximate phonetic representation, which is created with an elaborate set of rules, and whose principles of design include mapping voiced/unvoiced consonant pairs to the same encoding, encoding any initial vowel with ‘A’, and disregarding all non-initial vowel sounds. For example, the Double Metaphone representation of *antibiotic* is *ANTPTK*.

2.2 Candidate Ranking

Our approach computes the cosine similarity between the vector representation of a candidate and the composed vector representations of the misspelling context, weights this score with other parameters, and uses it as the ranking criterium. This setup is similar to the contextual similarity score by Kilicoglu et al. (2015),

3. <http://aspell.net/metaphone/>

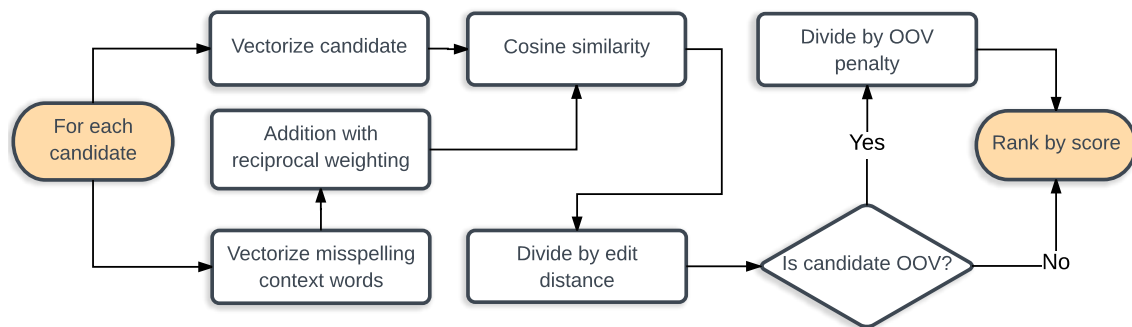


Figure 1: The final architecture of our model. Within a specified window size (9 for English, 10 for Dutch), it vectorizes every context word on each side if it is present in the vector vocabulary, applies reciprocal weighting, and sums the representations. It then calculates the cosine similarity with each candidate vector, and divides this score by the Damerau-Levenshtein edit distance between the candidate and misspelling. If the candidate is OOV, the score is divided by an OOV penalty.

which proved unsuccessful in their experiments. However, their experiments were preliminary. They used a limited context window of 2 tokens, could not account for candidates which are not observed in the training data, and did not investigate whether a bigger training corpus would lead to vector representations which scale better to the complexity of the task.

We undertake a more thorough examination of the applicability of neural embeddings to the spelling correction task. To tune the parameters of our context-sensitive spelling correction model in an unsupervised way, we automatically generate development corpora with artificial, randomly created spelling errors for three different scenarios following the procedures described in section 3.3. These three types of generated spelling error corpora, which we refer to as *setups*, are increasingly difficult for the spelling correction task. We apply the same setups to both English and Dutch. **Setup 1** is generated from the same corpus which is used to train the neural embeddings, and exclusively contains corrections which are present in the vocabulary of these neural embeddings. **Setup 2** is generated from a corpus in a different clinical subdomain, and also exclusively contains in-vector-vocabulary corrections. **Setup 3** presents the most difficult scenario, where we use the same corpus as for Setup 2, but only include corrections which are not present in the embedding vocabulary (OOV). In other words, here our model has to deal with both domain change and data sparsity.

Correcting OOV tokens in Setup 3 is made possible by using a combination of word and character n-gram embeddings. We train these embeddings with the fastText model (Bojanowski et al. 2017), an extension of the popular Word2Vec model (Mikolov et al. 2013), which creates vector representations for character n-grams and sums these with word unigram vectors to create the final word vectors. FastText allows for creating vector representations for misspelling replacement candidates absent from the trained embedding space, by only summing the vectors of the character n-grams.

We report our development experiments with the different setups in section 4.1. The final architecture of our model for both English and Dutch is described in Figure 1. We evaluate this model on our test data in section 4.2.

3. Materials

We tokenize all English data with the Pattern tokenizer (De Smedt and Daelemans 2012), and all Dutch data with Ucto⁴. All text is lowercased⁵, and we remove all tokens that include anything different from alphabetic

4. <https://languagemachines.github.io/ucto/>

5. While this has consequences for the nature of the task, it is a salient aspect of training good embeddings. Lowercasing reduces sparsity, therefore leading to more reliable representations, especially in the case of low frequency words.

Table 1: A comprehensive overview of our corpora described in section 3.3 and 3.4.

Language	Corpus type	Domain	Data used	Instances
ENGLISH	DEVELOPMENT: SETUP 1	critical care	MIMIC-III	5,000
	DEVELOPMENT: SETUP 2	breast/colon cancer	THYME	5,000
	DEVELOPMENT: SETUP 3	breast/colon cancer	THYME	1,500
	TEST	critical care	MIMIC-III	873
DUTCH	DEVELOPMENT: SETUP 1	critical care	UZA	5,000
	DEVELOPMENT: SETUP 2	breast/colon cancer	UZA	5,000
	DEVELOPMENT: SETUP 3	breast/colon cancer	UZA	350
	TEST	miscellaneous	UZA	490

Table 2: Examples of automatically generated spelling errors and some replacement candidates for the English development setups.

	Misspelling	Candidates
Setup 1	<i>unchanged</i> → <i>unchained</i>	unchanged, unchained, uncharged, unhinged
Setup 2	<i>chronic</i> → <i>chornic</i>	chronic, choreic, cornice, chloric
Setup 3	<i>accrued</i> → <i>accued</i>	accrued, accused, accuse, accede

characters or hyphens. Table 1 gives a comprehensive overview of the English and Dutch development and test corpora we describe in section 3.3 and 3.4.

3.1 Lexicons

To construct reference lexicons, we fuse general dictionaries with specialized resources. For our English lexicon, we use a union of the general dictionary from Jazzy⁶, a Java open source spell checker (47,160 items), and the UMLS[®] SPECIALIST lexicon⁷ (304,840 items), which contains a broad range of specialized clinical terms. This amounts to 319,579 unique lexical items. For our Dutch lexicon, we use as general dictionary the publicly available word list from Stichting OpenTaal⁸ (320,913 tokens), which has the official quality label of the Dutch Language Union. As specialized resource, we extract terminology from two clinical resources, namely, the Belgian Bilingual Biclassified Thesaurus (23,794 items) constructed by the universities of Ghent and Brussels, and the UMLS[®] Metathesaurus⁹ (77,646 items). This amounts to 371,559 unique lexical items.

3.2 Neural embeddings

We train a fastText skipgram model using the default parameters, except for the dimensionality, which we raise to 300, since we want to make sure that the embeddings are able to capture subtle semantic relationships in a training corpus of our size. For our English experiments, we train on 425M words from the MIMIC-III corpus, which contains medical records from critical care units. For our Dutch experiments, we train on 720M words from clinical records collected at the Antwerp University Hospital (UZA). These records span a decade in time, and cover various genres (notes, letters, protocols, reports) as well as a wide range of clinical subdomains, including gastroenterology, pulmonology, and critical care.

6. <http://jazzy.sourceforge.net>

7. <https://lexsrv3.nlm.nih.gov/LexSysGroup/Projects/lexicon/current/web/index.html>

8. <https://www.opentaal.org>

9. https://www.nlm.nih.gov/research/umls/knowledge_sources/metathesaurus/

Table 3: Examples of automatically generated spelling errors and some replacement candidates for the Dutch development setups.

	Misspelling	Candidates
Setup 1	<i>mediane</i> → <i>medciane</i>	mediane, mediale, medianen, Mediene
Setup 2	<i>beperkt</i> → <i>beprekt</i>	beperkt, betrekt, verrekt, gerekt, bevekt
Setup 3	<i>megacyste</i> → <i>megacyte</i>	megacyste, megabyte, megabytes

3.3 Development corpora

In order to tune our model parameters in an unsupervised way, we automatically create self-induced error corpora. We generate these development corpora by randomly sampling lines from a reference corpus, randomly sampling a single word per line if the word is present in our reference lexicon, transforming these words with either 1 (80%) or 2 (20%) random Damerau-Levenshtein operations to a non-word, and then extracting these misspelling instances with a context window of up to 10 tokens on each side. Table 1 gives an overview of all the development corpora and the data used to generate them. Table 2 and 3 give examples from all development corpora for both languages. For **Setup 1**, we perform our corpus creation procedure for critical care records, a domain which is present in the data used to train our neural embeddings. We exclusively sample words present in our vector vocabulary, resulting in 5,000 tokens for both English and Dutch. For **Setup 2**, we perform our procedure for records which exclusively cover the domain of brain and colon cancer, which is not represented in our neural embedding corpora. For English, we use the THYME (Styler IV et al. 2014) corpus. For Dutch, we use data which originally belonged to our neural embeddings training data, but which was located and held out before our experiments. We once again exclusively sample in-vector-vocabulary words, resulting in 5,000 tokens for both English and Dutch. For **Setup 3**, we again perform our procedure for the cancer corpora, but this time we exclusively sample OOV words, resulting in 1,500 tokens for English and 350 for Dutch. While this last setup can seem exaggerated or overly artificial, we want to explicitly isolate these cases from the other setups, since the distribution of OOVs is entirely dependent on the vocabulary overlap between the data being corrected and the data used to train the neural embeddings. In other words, it is relative with respect to the specific use case of our model in practice. On the one hand, we use this setup to estimate how well our trained model can generalize to other subdomains and corpora with only partially overlapping vocabulary; on the other hand, we use this setup to regulate the role of OOV correction candidates, as we discuss in section 4.1.

3.4 Test corpora

No benchmark test sets are publicly available for clinical spelling correction. A straightforward annotation task is costly and can lead to small corpora, such as the one by Lai et al., which contains just 78 misspelling instances. Therefore, we adopt a more cost-effective annotation approach. In a corpus, we spot misspellings by looking at items with a frequency of 5 or lower which are absent from our lexicon.¹⁰ We then extract and annotate instances of these misspellings along with their context. For English, we use the MIMIC-III data, resulting in 873 contextually different tokens of 357 unique error types.¹¹ For Dutch, we use a recent set of clinical records from the Antwerp University Hospital, which covers the same genres and domains as the neural embeddings training data. This results in 490 contextually different tokens of 359 unique error types. Tables 4 and 5 give examples from both test corpora.

10. While this excludes frequent error types, and is therefore far from an optimal strategy, it is hard to estimate the possible deceiving effect of this strategy without knowing the frequency distribution of spelling errors in the MIMIC-III corpus.

11. A script to extract this data can be found at <https://github.com/clips/clinspell>.

Table 4: Examples of empirically observed misspellings and some replacement candidates from our English test set, per Damerau-Levenshtein edit distance.

	Misspelling	Candidates
Edit distance 1	<i>sclerosing</i> → <i>sclerosin</i>	sclerosing, sclerosis, sclerotin, sclerostin
Edit distance 2	<i>symptoms</i> → <i>sympots</i>	symptoms, symptom, spots, symbols
Edit distance 3	<i>phlebitis</i> → <i>phebilitis</i>	phlebitis, cheilitis, pyelitis, phallitis

Table 5: Examples of empirically observed misspellings and some replacement candidates from our Dutch test set, per Damerau-Levenshtein edit distance.

	Misspelling	Candidates
Edit distance 1	<i>letsels</i> → <i>letels</i>	letsels, lepels, netels, zetels, zetsels
Edit distance 2	<i>weinig</i> → <i>wijnig</i>	weinig, pijnig, wijzig, tijdig, wijn
Edit distance 3	<i>verminderde</i> → <i>verminderderde</i>	verminderde, verminderende

4. Results

We first develop our model for each language by tuning the parameters with the development corpora. We then test this tuned model on the test data. We discuss the results and their implications in the next section. To evaluate the performance of our model, we use **first-best** accuracy as criterion, i.e., the percentage of misspellings which are properly corrected by the first-ranked replacement suggestion of our model. We use two variations of first-best accuracy, the terminology of which we borrow from Reynaert (2008): **true** first-best accuracy, which is the accuracy given the system’s dictionary; and **upper-bound** first-best accuracy, which removes the effect of dictionary shortcomings, by adding all correct word forms for the errors to be corrected to the system’s spelling dictionary. The latter criterion allows for measuring the upper bound on correction attainable by our system.

4.1 Development

To develop our model, we investigate a variety of parameters:

Vector composition functions

- (a) addition
- (b) multiplication
- (c) max embedding by Wu et al. (2015)

Edit distance penalty

- (a) Damerau-Levenshtein
- (b) Double Metaphone
- (c) Damerau-Levenshtein + Double Metaphone
- (d) Spell score by Lai et al.

Context parameters

- (a) Window size (1 to 10)
- (b) Reciprocal weighting
- (c) Removing stop words using the English stop word list from scikit-learn (Pedregosa et al. 2011) or the Dutch stop word list from Pattern (De Smedt and Daelemans 2012)
- (d) Including a vectorized representation of the misspelling

Table 6: True first-best correction accuracies for our 3 English development setups.

	Setup 1	Setup 2	Setup 3
Context	90.24	88.20	57.00
Noisy Channel	85.02	85.86	39.73

Table 7: True first-best correction accuracies for our 3 Dutch development setups.

	Setup 1	Setup 2	Setup 3
Context	87.94	89.10	82.00
Noisy Channel	86.90	85.80	66.57

We perform a grid search for Setup 1 and Setup 2 to discover which parameter combination leads to the highest accuracy averaged over both corpora. In this setting, we only allow for candidates which are present in the vector vocabulary. We then introduce OOV candidates for Setup 1, 2 and 3, and experiment with penalizing them, since their representations are less reliable. As these representations are only composed out of character n-gram vectors, with no word unigram vector, they are susceptible to noise caused by the particular nature of the n-grams; namely, sometimes the semantic similarity of OOV vectors to other vectors can be inflated in cases of strong orthographic overlap. OOV replacement candidates are more often redundant than necessary, as in most practical use cases of the correction model (where there is considerable vocabulary overlap between the embedding domain and the correction domain), the majority of correct misspelling replacements will be present in the trained vector space. Therefore, we try to penalize OOV representations to the extent that they do not cause noise in cases where they are redundant, but still rank first in cases where they are the correct replacement. We tune this OOV penalty by maximizing the accuracy for Setup 3 while minimizing the performance drop for Setup 1 and 2, using a weighted average of their correction accuracies.

The final architecture of our model for both English and Dutch is described in full in Figure 1, showing all used parameters. As the description shows, the models for both languages only differ in optimal window size (9 for English, 10 for Dutch). To compare our method against a reference noisy channel model in the most direct way, we implement the ranking component of Lai et al.’s model in our pipeline (**Noisy Channel**). This component requires corpus frequencies, which we extract from the same data that we use to train the embeddings. Our context-sensitive model (**Context**) outperforms the noisy channel for each corpus in our development phase, for both English and Dutch, as shown in Table 6 and 7. Moreover, as the results for Setup 3 show, our method generalizes considerably better to OOV misspellings, as we explicitly intended in the development of our model.

4.2 Test

Table 8 shows the English correction accuracies for **Context** and **Noisy Channel** as off-the-shelf tools, compared to two existing tools. The first tool is HunSpell, a popular open source spell checker used by Google Chrome and Firefox. The second tool is the original implementation of Lai et al.’s model, which they shared with us. Table 9 shows the Dutch correction accuracies for **Context** and **Noisy Channel** as off-the-shelf tools, as compared to HunSpell.

The performance of our models on the test sets is held back by the incomplete coverage of our reference lexicons. For English, missing corrections are mostly highly specialized medical terms, or inflections of more common terminology. For Dutch, this includes relatively infrequent compounds as well. Compounds in Dutch, as opposed to English, are mostly orthographically concatenated into one lexical item. Since Dutch language users tend to be very productive with compounding, this leads to a whole range of standard language that is hard to cover exhaustively in a lexicon. We use the upper-bound first-best correction accuracy to examine the performance of our ranking models with disregard to such circumstances. Tables 8 and 9 show that the performance according to this metric is comparable to the true first-best correction accuracy for the development corpora.

Table 8: The correction accuracies for our English test experiments, evaluated for two different scenarios. *True first-best accuracy*: gives the first-best accuracies of all off-the-shelf tools. *Upper-bound first-best accuracy*: gives the first-best accuracies of our implemented models for the scenario where correct replacements missing from the lexicon are included in the lexicon before the experiment.

Evaluation	HunSpell	Lai et al.	Context	Noisy Channel
TRUE FIRST-BEST ACCURACY	52.69	61.97	88.21	87.85
UPPER-BOUND FIRST-BEST ACCURACY			93.02	92.66

Table 9: The correction accuracies for our Dutch test experiments, evaluated for two different scenarios. *True first-best accuracy*: gives the accuracies of all off-the-shelf tools. *Upper-bound first-best accuracy*: gives the accuracies of our implemented models for the scenario where correct replacements missing from the lexicon are included in the lexicon before the experiment.

Evaluation	HunSpell	Context	Noisy Channel
TRUE FIRST-BEST ACCURACY	64.29	76.53	79.71
UPPER-BOUND FIRST-BEST ACCURACY		87.75	92.45

5. Discussion

In terms of correction accuracy, our context-sensitive model and our own implementation of Lai et al.’s ranking model outperform off-the-shelf tools for both English and Dutch, establishing a state-of-the-art for spelling correction of clinical free-text. The salient difference in performance between Lai et al.’s system and our specific implementation of their noisy channel model highlights the influence of (lack of) training resources and development decisions on the general applicability of spelling correction models. Moreover, it shows the strength of the noisy channel model in scenarios where the scale of the resources is sufficient (in this case, 425M words for English and 720M words for Dutch) to reliably estimate prior probabilities from corpus frequencies.

However, sufficient empirical resources to estimate a fine-grained likelihood (namely, a large corpus of empirically observed errors from which a reliable error model can be extracted) are still absent for the clinical domain. Therefore, the likelihood of Lai et al.’s ranking model is estimated with a rudimentary spell score, which is a weighted combination of Damerau-Levenshtein and Double Metaphone edit distance. While this error model leads to a noisy channel model which is robust in performance, as shown by our test results, it also leads to a pragmatic performance ceiling where more heavily distorted replacement candidates are downplayed to safeguard robustness of performance, regardless of their possible empirical association with the misspelling. As a result, our noisy channel model is still prone to cases of frequency bias, including the example of frequency bias which we have provided in the introduction of this paper: our noisy channel model does not succeed in correcting the MIMIC-III misspelling *goint* to the correct form *going* due to the higher corpus frequency of, and therefore higher prior probability assigned to, the word type *point*. While the difference in frequency is salient, it is not insurmountable for a likelihood reflecting a proper error model, which in this case would typically reflect that *goint* is more probable to be a typo of *going* than of *point*. However, the rudimentary spell score does not reflect that notion. This example illustrates that, regardless of the theoretical validity of the noisy channel, we are still very much bound to the practical reality of its implementation, including the state of resources.

Our method tries to improve on the clinical spelling correction process considering the availability of actual incomplete resources. As it stands, a noisy channel model like the one by Lai et al. still occasionally suffers from frequency bias; it is not able to correct a specific misspelling type to different corrections in different contexts, and is not sufficiently equipped to deal with word types that are not observed in training data. Our unsupervised context-sensitive model targets these weaknesses. Figures 2 and 3 show the correction




Figure 2: The English correction accuracies for **Context** and **Noisy Channel** for Setup 1, Setup 2, and the test set, grouped per relative frequency of the correct replacement compared to other replacement candidates. *rel freq = 1*: highest corpus frequency of all candidates. *rel freq = 2*: second highest corpus frequency of all candidates. *rel freq > 2*: corpus frequency lower than second highest of all candidates.

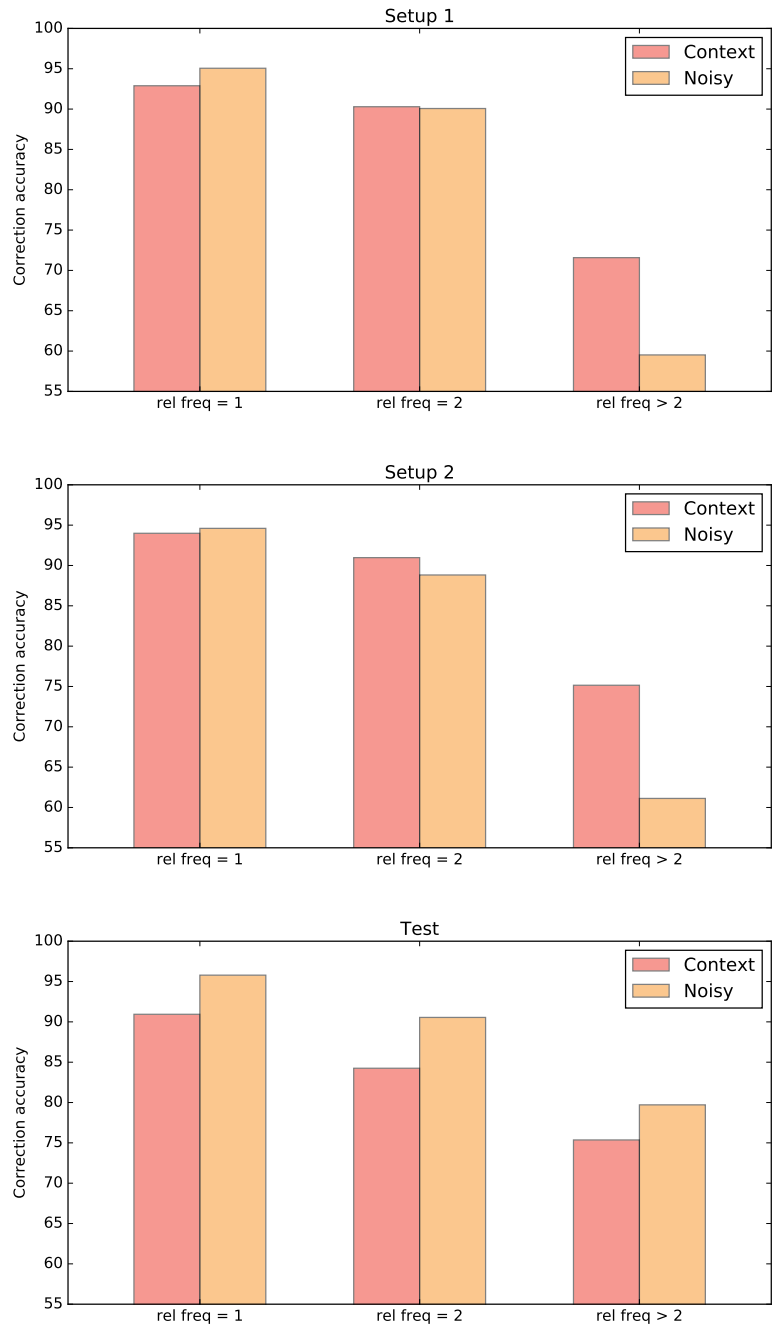


Figure 3: The Dutch correction accuracies for **Context** and **Noisy Channel** for Setup 1, Setup 2, and the test set, grouped per relative frequency of the correct replacement compared to other replacement candidates. *rel freq = 1*: highest corpus frequency of all candidates. *rel freq = 2*: second highest corpus frequency of all candidates. *rel freq > 2*: corpus frequency lower than second highest of all candidates.




Figure 4: 2-dimensional t-SNE projection of the vectorized context of the English test misspelling *goint* and 4 replacement candidates in the trained MIMIC-III vector space. Dot size denotes corpus frequency, numbers denote cosine similarity. The English misspelling context is *new central line lower extremity bypass with sob now [goint] to [be] intubated*. While the noisy channel chooses the more frequent *point*, our model correctly chooses the most semantically fitting *going*.

accuracies for three scenarios: one where the most frequent candidate is the correct one ($rel\ freq = 1$), one where the second most frequent candidate is the correct one ($rel\ freq = 2$), and one where the correct candidate has a lower relative frequency ($rel\ freq > 2$). Figure 2 confirms the hypothesis that our context-sensitive model counters the frequency bias of a noisy channel model for our English experiments. The results for our development corpora show that in cases where $rel\ freq = 1$, the noisy channel scores similar or slightly better, as expected. This trend is reflected in the test results. In cases where $rel\ freq = 2$, our model scores slightly better. This trend is not reflected in the test results. In fact, it is reversed. Lastly, in cases where $rel\ freq > 2$, our model scores much better. This trend is reflected in the test results, if to a smaller extent. However, the relatively small sample size (a difference of 6 correct instances on a total of 243) should be kept in mind. Figure 4 visualizes an example of frequency bias, where the *goint* misspelling which we discussed earlier is correctly handled by our model as opposed to the noisy channel model.

Figure 3 shows that the performance our context-sensitive model exhibits the same characteristics for the Dutch development corpora as for the English development corpora. However, this time none of the trends are reflected in the test results, which leads to our model being outperformed by the noisy channel model. This discrepancy raises the question to what extent the artificial nature of the development corpora leads to reliable models for empirical data. If the distributions of the several data types differ greatly, this undermines our unsupervised approach, which implicitly assumes that the distributions will not differ that greatly. To investigate this, we performed a grid search for both the English and Dutch test corpus, to examine which parameter combination leads to the best-performing model. For the English test data, this parameter combination is similar to our actual model derived from our development experiments. In other words, the underlying assumption of our unsupervised approach is confirmed.

For the Dutch test data, however, the optimal parameter combination differs dramatically from our developed model. It includes two parameters which are absent from our developed model described in Figure 1: the context representation also includes a vectorized representation of the misspelling itself, and the edit distance weighting adds Double Metaphone edit distance to the Damerau-Levenshtein edit distance. Moreover, the optimal context window size is 2, which is considerably smaller than for the originally developed model. With this parameter combination, the output of the model for the Dutch test data is exactly similar

to the output of the noisy channel model. These analyses suggest that the distribution of the Dutch test data differs greatly from that of the development data. This discrepancy can be caused by the sparsity of the Dutch test data, which covers the same amount of error types as the English test data, but much fewer contextually different instances. The only conclusion we can draw is that the nature of our test set is possibly skewed in a way that does not allow for a thorough comparative evaluation of our models. As it stands, however, we have no empirical evidence that our Dutch context-sensitive model actually counters the frequency bias of our noisy channel. While we want to avoid too much speculation as to the reason why, these results invite inquiry into how important context actually is for Dutch clinical spelling correction.

When we look at the output of our context-sensitive model for both English and Dutch, we can categorize the errors it makes in 3 different types. The first type of errors concerns, predictably, misspellings for which the contextual clues are too unspecific. This lack of useful contextual information is sometimes caused by occurrences of other misspellings in the context window, and poses a fundamental challenge to our method. The second type of errors concerns cases where the contextual clues are actually misleading. This happens for instance in cases where a word type has multiple senses which are not strongly related. Our Dutch test set contains the misspelling *posten* → *polsen*, where from the context it appears that *polsen* has the more infrequent sense of ‘polling someone about something’ instead of the prevalent sense ‘wrists’. Since this word type shares one vector representation for both senses, the contextual information does not turn out to be strong enough for correcting the misspelling to the correct word type. Lastly, while our development experiments have tried to minimize the noise spread by OOV candidates, it is still noticeable in some instances.

6. Conclusion and future research

In this paper, we have proposed an unsupervised context-sensitive model for clinical spelling correction which uses word and character n-gram embeddings. This simple ranking model, which can be tuned to a specific language and domain by generating self-induced error corpora, tries to counter the frequency bias of a noisy channel model by exploiting contextual clues.

As an implemented spelling correction tool for English clinical free-text, our method outperforms both a broadly used and a domain-specific off-the-shelf tool for empirically observed misspellings in MIMIC-III. Moreover, a detailed analysis of its performance shows that it does in fact counter the frequency bias of a noisy channel model. However, the relatively small sample size for this analysis should be kept in mind.

As an implemented spelling correction tool for Dutch clinical free-text, our method outperforms a broadly used off-the-shelf tool for empirically observed misspellings in collected data from the Antwerp University Hospital. However, our Dutch test set offers no empirical evidence that it counters the frequency bias of a noisy channel model. It is unclear whether this is caused by the sparsity of the test set.

Future research can investigate whether our method transfers well to other genres and domains. Secondly, it can address the three problem areas we have identified at the end of our discussion in section 5, namely, unspecific contextual clues, multiple word senses of a single word type, and noise spread by OOV candidates. Lastly, it is worthwhile to investigate how successfully our model can be applied to real-word errors.

7. Acknowledgements

This research was carried out in the framework of the Accumulate VLAIO SBO project, funded by the government agency Flanders Innovation & Entrepreneurship (VLAIO). We would also like to thank Kim Luyckx for providing access to the Dutch data; Elyne Scheurwegs for preparing and managing the Dutch data; Stéphan Tulkens for his logistic support with coding; and Kenneth Lai, Maxim Topaz, Foster R. Goss, and Li Zhou for sharing their system with us.

References

- Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov (2017), Enriching word vectors with subword information, *Transactions of the Association for Computational Linguistics* **5**, pp. 135–146.
- Cornet, Ronald, Armand van Eldik, and Nicolette De Keizer (2012), Inventory of tools for Dutch clinical language processing, *Proceedings of the 24th European Medical Informatics Conference*.
- De Smedt, Tom and Walter Daelemans (2012), Pattern for Python, *Journal of Machine Learning Research* **13**, pp. 2031–2035.
- Flor, Michael (2012), Four types of context for automatic spelling correction, *TAL* **53** (3), pp. 61–99.
- Johnson, Alistair E.W., Tom J. Pollard, Lu Shen, Li Wei, H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark (2016), MIMIC-III, a freely accessible critical care database, *Scientific Data*.
- Jurafsky, Daniel and James H. Martin (2016), Spelling correction and the noisy channel, Draft of November 7, 2016.
- Kilicoglu, Halil, Marcelo Fiszman, Kirk Roberts, and Dina Demner-Fushman (2015), An ensemble method for spelling correction in consumer health questions, *AMIA Annual Symposium Proceedings* pp. 727–73.
- Lai, Kenneth H., Maxim Topaz, Foster R. Goss, and Li Zhou (2015), Automated misspelling detection and correction in clinical free-text records, *Journal of Biomedical Informatics* **55**, pp. 188–195.
- Liu, Hongfang, Stephen T. Wu, Dingcheng Li, Siddharta Jonnalagadda, Sunghwan Sohn, Kavishwar Waghlikar, Peter J. Haug, Stanley M. Huff, and Christopher G. Chute (2012), Towards a semantic lexicon for clinical natural language processing, *AMIA Annual Symposium Proceedings*.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean (2013), Efficient estimation of word representations in vector space, *Proceedings of Workshop at International Conference on Learning Representations*.
- Mitton, Roger (2010), Fifty years of spellchecking, *Writing Systems Research* **2** (1), pp. 1–7.
- Pande, Harshit (2017), Effective search space reduction for spell correction using character neural embeddings, *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* pp. 170–174.
- Patrick, J., M. Sabbagh, S. Jain, and H. Zheng (2010), Spelling correction in clinical notes with emphasis on first suggestion accuracy, *2nd Workshop on Building and Evaluating Resources for Biomedical Text Mining* pp. 2–8.
- Pedregosa, Fabian, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, and et al. (2011), Scikit-learn: Machine learning in Python, *The Journal of Machine Learning Research* **12**, pp. 2825–2830.
- Reynaert, Martin (2008), All, and only, the errors: more complete and consistent spelling and OCR-error correction evaluation, *Proceedings of the International Conference on Language Resources and Evaluation (LREC)* pp. 1867–1872.
- Ruch, Patrick, Robert Baud, and Antoine Geissbühler (2003), Using lexical disambiguation and named-entity recognition to improve spelling correction in the electronic patient record, *Artificial Intelligence in Medicine* **29**, pp. 169–184.
- Sridhar, Vivek Kumar Rangarajan (2015), Unsupervised text normalization using distributed representations of words and phrases, *Proceedings of NAACL-HLT 2015* pp. 8–16.

- Styler IV, William F., Steven Bethard, Sean Finan, Martha Palmer, Sameer Pradhan, Piet C. de Groen, Brad Erickson, Timothy Miller, Chen Lin, Guergana Savova, and James Pustejovsky (2014), Temporal annotation in the clinical domain, *Transactions of the Association for Computational Linguistics* **2**, pp. 143–154.
- Tolentino, Herman D., Michael D. Matters, Wikke Walop, Barbara Law, Wesley Tong, Fang Liu, Paul Fontelo, Katrin Kohl, and Daniel C. Payne (2007), A UMLS-based spell checker for natural language processing in vaccine safety, *BMC Medical Informatics and Decision Making*.
- Wu, Yonghui, Jun Xu, Yaoyun Zhang, and Hua Xu (2015), Clinical abbreviation disambiguation using neural word embeddings, *Proceedings of the 2015 Workshop on Biomedical Natural Language Processing (BioNLP)* pp. 171–176.