

Two Approaches to Diachronic Normalization of Polish Texts

Kacper Dudzic, Filip Graliński, Krzysztof Jassem, Marek Kubis, Piotr Wierzchoń
Adam Mickiewicz University, Poznań, Poland
{firstname.lastname}@amu.edu.pl

Abstract

This paper discusses two approaches to the diachronic normalization of Polish texts: a rule-based solution that relies on a set of handcrafted patterns, and a neural normalization model based on the text-to-text transfer transformer architecture. The training and evaluation data prepared for the task are discussed in detail, along with experiments conducted to compare the proposed normalization solutions. A quantitative and qualitative analysis is made. It is shown that at the current stage of inquiry into the problem, the rule-based solution outperforms the neural one on 3 out of 4 variants of the prepared dataset, although in practice both approaches have distinct advantages and disadvantages.

1 Introduction

This paper discusses two solutions to the problem of diachronic normalization, that is, the task of determining contemporary spelling for a given historical text. Diachronic normalization may concern the writing of individual words, punctuation, hyphenation, or separation of tokens. We believe that the methods described in this paper may be useful for linguistic research on historical texts. A practical use case for our work is to facilitate full-text search in historical texts – a query written in contemporary spelling may trigger a search for historical variants through the use of reversed-order diachronic normalization.

Similar experiments, for text normalization in a speech synthesis system from text, were described in (Sproat and Jaitly, 2016). Those authors claim that text normalization remains one of the few tasks in the field of natural language processing where handcrafted rules may yield better results than machine learning. This is due to the following reasons:

- Lack of training data; there is no economic motivation for creating training data for text normalization – unlike machine translation,

for example, for which training data are created "naturally";

- Low data density of interesting cases, i.e. words that should be somehow changed – unlike for example, phonemic transcription, where all words are converted to a new representation;
- Standard methods of evaluation which do not reward trivial cases (copying of input words), thus favoring human labor.

In our experiments, we compare the results of a rule-based approach with one based on machine learning. The rule-based approach relies on a set of handcrafted rules to normalize text. In the ML approach, we train a supervised normalization model on the basis of a corpus of Polish books for which both historical and current spellings are available.

2 Related work

The first attempts at rule-based diachronic normalization used for historical text in English were described by Rayson et al. (2007) and Baron et al. (2009). Similar studies were conducted for German (Archer et al., 2006). There, context rules operated at the level of letters instead of words. The normalization rules may be derived from corpora, as Bollmann et al. (2011) showed for German. Diachronic normalization may be also performed using a noisy channel model, as described by Oravecz et al. (2010) using the example of Old Hungarian texts. Research on diachronic normalization has also been conducted for Portuguese (Reynaert et al., 2012), Swedish (Pettersson et al., 2012), Slovene (Scherrer and Erjavec, 2013), Spanish (Porta et al., 2013), and Basque (Etxeberria et al., 2016).

Bollmann (2019) surveys historical spelling normalization methods for eight languages. He reports word-level accuracy for the evaluated systems. He claims that using CER is not justified, because it

strongly correlates with WER for systems showing reasonable accuracy. [Bollmann and Sjøgaard \(2016\)](#) use bi-directional LSTMs and multi-task learning to normalize texts in Early New High German. Their dataset consists of 44 texts from the Aselm corpus. The model presented is evaluated with respect to word-level accuracy. [Robertson and Goldwater \(2018\)](#) discuss the problem of evaluating historical text normalization systems. They emphasize the necessity of reporting accuracy for unobserved tokens and recommend confronting the normalization systems with a simple baseline that memorizes training samples.

[Jassem et al. \(2017, 2018\)](#) present an automatic method for diachronic normalization of Polish texts. The proposed method uses a formal language to model diachronic changes. [Graliński and Jassem \(2020\)](#) introduce a method for finding spelling variants in a diachronic corpus using word2vec.

3 Data

Training and evaluation of a diachronic normalizer requires a corpus of texts that preserve historical spelling along with their contemporized counterparts. As our aim is the normalization of Polish prose, we decided to collect texts for our corpus from two sources. Texts that preserve historical spelling were drawn from the Polish edition of the Wikisource project ([Wikimedia Foundation, 2023](#)), which provides proof-read transcriptions of printed books that have fallen into the public domain, encoded in the MediaWiki format. For contemporized texts, we used Wolne Lektury ([Modern Poland Foundation, 2023](#)), a digital library that aims to deliver new editions of school readers, free of charge. Although both sources encompass a wide variety of texts, ranging from poems and works of philosophy to dictionaries and historical documents, we narrowed our attention to novels, to facilitate the process of matching the original texts from Wikisource to their contemporized versions in Wolne Lektury with the use of metadata information available for novels in both sources. We initially sourced 308 novels from Wikisource and 279 from Wolne Lektury.

3.1 Preprocessing

All of the texts then underwent preprocessing. First, we split the texts into paragraphs, with the use of markup information preserved in XML files sourced from Wolne Lektury, and MediaWiki con-

tent collected from Wikisource. Next, regular expressions were used to remove leftover markup information, such as in-text metadata, formatting, or HTML tags, and to normalize some atypical characters. Accordingly, diacritical marks were removed from letter characters not belonging to the Polish alphabet, and non-ASCII variants of standard letter characters of the Latin alphabet were replaced by their ASCII counterparts. Finally, the same method was used to remove dialogue-specific text formatting and punctuation in paragraphs consisting of dialogue utterances, such as quotation dashes or character cues.

3.2 Alignment

To create aligned paragraph data, we first automatically matched all editions of novels existing across both data sources using fuzzy information similarity for author and title metadata. We then narrowed the matches to those that contained at least one edition in each of the sources.

Next, for each match of all editions of a novel, the oldest edition from Wikisource and the most recent edition from WolneLektury were identified using metadata information. Subsequently, the text paragraphs of both editions were extracted and aligned using the Hunalign tool, version 1.1 ([Varga et al., 2005](#)). Specifically, it was used to automatically create paragraph pairs consisting of a given text fragment with historical spelling from the oldest edition of a novel and the same text fragment but with contemporized spelling found in its newer edition, optionally automatically joining or splitting paragraphs where it was applicable. The paragraph alignment quality metric returned by Hunalign was consulted to provide additional filtering. The average alignment quality score across the entire text contents for each edition pair was used to identify and discard very low-scoring edition pairs, which turned out to be Polish translations of foreign novels made by different translators. In turn, per-paragraph alignment quality scores below 1.0 were used as an indicator to discard singular misaligned paragraphs.

3.3 Dataset creation

After completing all of the above steps and performing deduplication at the very end, we obtained a final corpus of 248,645 paragraph pairs originating from 87 eligible pairs of matched novel editions. Four dataset variants were created with this as the

basis. All variants involve a training and test split, but they differ in the following two respects:

Pruning was either applied or not. *Pruned* versions of the dataset are reduced in size by removing samples in which the paragraphs of the pair are identical. Applying pruning leads to a 64.83% decrease in the number of samples, a 47.34% decrease in the number of words, and a 47.23% decrease in the number of characters.

Separation of novels prior to the train/test split was either performed or not. In *separated* variants of the dataset, train and test sets are created from separate pools of novels with no overlap, so that all paragraphs from a given novel are contained in only one of the sets. Four novels were sampled from each of the quartiles determined with respect to the number of paragraphs contained in the corpus, to guarantee that each data subset contained a balanced volume of text. In the case of *non-separated* variants, the paragraphs are randomly sampled from the entire set of novels following the standard 80%/20% sampling ratio for train/test splits.

4 Experiments

4.1 Rule-based model

Our first solution to the problem of diachronic normalization relies on a set of deterministic rules. Henceforth, we will refer to this solution as *Transducers*. The rules were handcrafted initially and then adjusted semi-automatically. They were created mostly based on the expert literature describing changes in the Polish spelling system and by looking at a list of similar words having close embeddings. For most of the work on the rules, datasets for supervised learning were not consulted. Originally, the rules were written using the Thrax language (Tai et al., 2011) for defining transducer grammars, but more recently have been rewritten into a Java code base with normalization rules encoded using regular expressions. For instance, the rule:

```
Rule(  
  "([cs]|(?:\\A|(?<![cdsr]))z)  
  y([aąęiου])",  
  "$1j$2")
```

handles normalization of y into i in some circumstances (e.g. *decyzya* into *decyzja*). The decision to switch to Java was motivated by the fact that such a module can be easily incorporated, as a plugin, into Java-based open-source search engines (Lucene and Solr). When writing the rules, a conservative approach was taken: a rule was added only when the probability of unwanted changes to texts was very low. Apart from regular expressions, the rule-based solution uses a dictionary of transformations for specific words and dictionaries of exceptions, based on the ideas outlined in (Graliński and Jassem, 2020). The *Transducers* module also handles some OCR errors, but the coverage is rather low (only high-precision rules were applied).

Some further examples of the rules used are included in appendix A.

4.2 Neural normalization models

Diachronic normalization is an example of a language processing task that accepts text at the input and returns text at the output. Therefore, we decided to use the text-to-text transfer transformer architecture (T5, Raffel et al., 2020) as a basis for our supervised normalization models. Initial weights were taken from the pre-trained pT5 model (Chrabrowa et al., 2022), an encoder-decoder model that follows the T5 architecture. The pT5 model was initialized from its multilingual counterpart (mT5, Xue et al., 2021) and further trained on Polish language corpora. It achieves better performance than mT5 on Polish language benchmark tasks with a smaller number of parameters. For our experiments, we used the largest variant of this model available at Hugging Face.¹

We finetuned four neural diachronic normalization models, with one model for each variant of our dataset. The models were trained for three epochs, using Adam as the optimizer and a learning rate of 5e-05 with a linear scheduler. The batch size was kept at 1 due to the memory limitations of the GPU used for the experiments. Maximum input and output sequence token lengths followed the T5 model family’s default of 512. Longer input sequences were split into chunks of maximum length, processed separately, and then joined.

Table 2 reports the results of the evaluation of the neural normalization models, and compares them with the rule-based model. One may observe that the rule-based model is a strong baseline for

¹<https://huggingface.co/allegro/plt5-large>

Pruning	Separation	Split samples		Characters	Words
		Train	Test		
No	No	198,916	49,729	92,306,901	14,438,223
Yes	No	69,952	17,488	48,710,393	7,603,573
No	Yes	199,004	49,641	92,306,901	14,438,223
Yes	Yes	63,921	23,519	48,710,393	7,603,573

Table 1: Dataset statistics

Method	Pruning	Separation	CER	WER
Transducers	No	No	0.0164	0.0466
Neural	No	No	0.0488	0.0654
Transducers	Yes	No	0.0319	0.0827
Neural	Yes	No	0.0728	0.1011
Transducers	No	Yes	0.0182	0.0560
Neural	No	Yes	0.0632	0.0932
Transducers	Yes	Yes	0.0281	0.0844
Neural	Yes	Yes	0.0398	0.0737

Table 2: Evaluation results

the task, outperforming the neural models with respect to character error rate (CER) and word error rate (WER). However, the supervised model surpasses the rule-based solution in the case of a test set that consists of a separate set of novels (*Separation=Yes*) and excludes samples that should remain unmodified in the normalization process (*Pruning=Yes*).

5 Discussion

After performing a qualitative analysis of the results obtained using rule-based and neural normalization models, we observed for the neural networks: (1) flexibility in context interpretation, i.e., the ability to adapt to various contexts and understand linguistic nuances; (2) recognition of irregular patterns, i.e., the ability to identify and process non-standard and complex language forms; (3) context-based changes, i.e., considering a broad context, which can lead to changes that go beyond simple spelling rules. On the other hand, for rule-based normalization, it was noted that: (1) relying on specific, defined rules, i.e., focusing on the strict application of established spelling rules, *Transducers* are less flexible in interpretation, meaning that they have limited abilities to cope with irregularities and linguistic nuances; (2) *Transducers* fol-

low a literal interpretation of rules, which may not take into account the full context. The neural approach effectively normalizes examples of former single-word spelling, especially for conjunctions: *przyczem* → *przy czym* (Eng. *at the same time*), *poczem* → *po czym* (Eng. *thereafter*), *napewno* → *na pewno* (Eng. *certainly*), *niema* → *nie ma* (Eng. *there is no*). The rule-based approach, in turn, aptly converts regular orthographic phenomena: *egzystencya* → *egzystencja* (Eng. *existence*), *jenerał* → *generał* (Eng. *general*), *teorya* → *teoria* (Eng. *theory*). It also accurately transforms proper nouns: *Anglja* → *Anglia*, *Marjetka* → *Marietka*. The spelling changes – from *egzystencya* to *egzystencja*, *Anglja* to *Anglia*, etc. – were part of the Polish orthographic reform of 1936. This reform was aimed at simplifying and standardizing the Polish language’s spelling. It introduced several changes, including the replacement of the letter ‘y’ with ‘j’ or ‘i’ in certain contexts, and the introduction of the letter ‘j’ in place of ‘i’ in some cases to better reflect pronunciation. This reform significantly influenced the modern Polish language, aligning it more closely with its phonetics.

6 Conclusion

This paper has discussed two approaches to the diachronic normalization of Polish texts. We presented *Transducers*, a rule-based solution that relies on a set of deterministic, handcrafted rules, and a family of neural normalization models based on a text-to-text transfer transformer architecture. The experiments that we conducted showed that the rule-based approach is effective in the diachronic normalization task. However, the neural model surpassed the rule-based solution in the case of a test set that consists of a separate set of novels and excludes samples that should remain unchanged in the normalization process.

As the presented research is preliminary in nature, there are several promising directions to explore, which we are committed to doing in the near future. Among other ideas, we want to test the performance of hybrid solutions combining both approaches in distinct ways. We are also considering testing different model architectures and conducting further work on improving the quality of the training data used for the neural approach, as we believe it has the potential to eventually surpass the rule-based solution in most typical scenarios.

Limitations

We restrict our attention to the diachronic normalization of Polish texts. Generalizing the proposed methods to new languages will require, firstly, a new, handcrafted set of normalization rules being developed for the rule-based model presented in section 4.1; and secondly, a parallel corpus of texts that encompass both historical and current spelling, for the neural normalization models discussed in section 4.2.

Acknowledgments

The diachronic normalization methods presented in this work have been developed as part of the „Tools for text normalization and diachronic analysis” module of the Dariah.lab² infrastructure hosted at the Faculty of Mathematics and Computer Science of Adam Mickiewicz University in Poznań, Poland. Construction of Dariah.lab was the aim of the project: *Digital Research Infrastructure for the Arts and Humanities DARIAH-PL*, funded by the Intelligent Development Operational Programme,

²<https://lab.dariah.pl/en/project/about-project/>

Polish National Centre for Research and Development, ID: POIR.04.02.00-00-D006/20.

References

- Dawn Archer, Andrea Ernst-Gerlach, Sebastian Kempken, Thomas Pilz, and Paul Rayson. 2006. The identification of spelling variants in English and German historical texts: manual or automatic? In *Digital Humanities 2006*. Paris, France: CATI, Université Paris-Sorbonne.
- Alistair Baron, Paul Rayson, and Dawn Archer. 2009. Automatic standardization of spelling for historical text mining. In *Digital Humanities 2009*.
- Marcel Bollmann. 2019. [A large-scale comparison of historical text normalization systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3885–3898, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marcel Bollmann, Florian Petran, and Stephanie Dipper. 2011. Rule-based normalization of historical texts. In *Proceedings of the International Workshop on Language Technologies for Digital Humanities*.
- Marcel Bollmann and Anders Søgaard. 2016. [Improving historical spelling normalization with bi-directional LSTMs and multi-task learning](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 131–139, Osaka, Japan. The COLING 2016 Organizing Committee.
- Aleksandra Chrabrowa, Łukasz Dragan, Karol Grzegorzczak, Dariusz Kajtoch, Mikołaj Koszowski, Robert Mroczkowski, and Piotr Rybak. 2022. [Evaluation of transfer learning for Polish with a text-to-text model](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4374–4394, Marseille, France. European Language Resources Association.
- Izaskun Etxeberria, Inaki Alegria, Larraitz Uria, and Mans Hulden. 2016. Evaluating the noisy channel model for the normalization of historical texts: Basque, Spanish and Slovene. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC-2016)*.
- Filip Galiński and Krzysztof Jassem. 2020. [Mining historical texts for diachronic spelling variants](#). *Poznan Studies in Contemporary Linguistics*, 56(4):629–650.
- Krzysztof Jassem, Filip Galiński, and Tomasz Obrębski. 2017. Pros and Cons of Normalizing Text with Thrax. In *Proceedings of the 8th Language and Technology Conference. Human Language Technologies as a Challenge for Computer Science and Linguistics*.

- Krzysztof Jassem, Filip Graliński, Tomasz Obrebski, and Piotr Wierzchoń. 2018. [Automatic diachronic normalization of polish texts](#). *Investigationes Linguisticae*, 37:17–33.
- Modern Poland Foundation. 2023. About the Project. <https://wolnelektury.pl/info/o-projekcie/>. Accessed: 2023-12-16.
- Csaba Oravecz, Balint Sass, and Eszter Simon. 2010. Semi-automatic normalization of Old Hungarian codices. In *Proceedings of the ECAI 2010 Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH 2010)*.
- Eva Pettersson, Beata Megyesi, and Joakim Nivre. 2012. Rule-based normalisation of historical text – a diachronic study. In *Empirical Methods in Natural Language Processing: Proceedings of the 11th Conference on Natural Language Processing (KONVENS 2012)*.
- Jordi Porta, Jose-Luis Sancho, and Javier Gomez. 2013. Edit transducers for spelling variation in Old Spanish. In *Proceedings of the workshop on computational historical linguistics at NODALIDA*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Paul Rayson, Dawn Archer, Alistair Baron, and Nicholas Smith. 2007. Tagging historical corpora – the problem of spelling variation. In *Dagstuhl Seminar Proceedings. Schloss Dagstuhl-Leibniz-Zentrum für Informatik*.
- Martin Reynaert, Iris Hendrickx, and Rita Marquilha. 2012. Historical spelling normalization. a comparison of two statistical methods: Ticcl and vard2. In *Proceedings of the Second Workshop on Annotation of Corpora for Research in the Humanities (ACRH-2), Lisbon*, pages 87–98. Edicoes Colibri.
- Alexander Robertson and Sharon Goldwater. 2018. [Evaluating historical text normalization systems: How well do they generalize?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 720–725, New Orleans, Louisiana. Association for Computational Linguistics.
- Yves Scherrer and Tomasz Erjavec. 2013. Modernizing historical Slovene words with character-based SMT. In *BSNLP 2013-4th Biennial Workshop on Balto-Slavic Natural Language Processing*.
- Richard Sproat and Navdeep Jaitly. 2016. [RNN approaches to text normalization: A challenge](#). *CoRR*, abs/1611.00068.
- Terry Tai, Wojciech Skut, and Richard Sproat. 2011. Thrax: An open source grammar compiler built on OpenFst. In *IEEE Automatic Speech Recognition and Understanding Workshop*.
- Dániel Varga, László Németh, Péter Halácsy, András Kornai, Viktor Trón, and Viktor Nagy. 2005. Parallel corpora for medium density languages. In *Proceedings of the RANLP 2005*, pages 590–596.
- Wikimedia Foundation. 2023. About Wikisource. https://wikisource.org/wiki/Wikisource>About_Wikisource. Accessed: 2023-12-16.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

A Appendix

Here we present further examples of rules used in the rule-based diachronic normalization solution.

```
Rule("izk", "isk")
Rule("yja\\b", "ja")
Rule("(le|ó)dz\\Z", "$1c")
Rule("\\Aanti-?", "anty")
Rule("iemi\\Z", "imi")
Rule("emi\\Z", "ymi")
Rule(
    "(ac|owan|yjn|owat|jsz|tyczn|logiczn)
    em\\Z",
    "$1ym")
Rule(
    "([dfglnprtt])[jy]([aæęiou])",
    "$1i$2")
```