








Data Centric Domain Adaptation for Historical Text with OCR Errors

Luisa März^{1,2,4} , Stefan Schweter³ , Nina Poerner¹ , Benjamin Roth² ,
and Hinrich Schütze¹ 

¹ Center for Information and Language Processing, Ludwig Maximilian University,
Munich, Germany

maerz@cis.lmu.de, inquiries@cislmu.org

² Digital Philology, Research Group Data Mining and Machine Learning, University
of Vienna, Vienna, Austria

³ Bayerische Staatsbibliothek München, Digital Library/Munich Digitization Center,
Munich, Germany

⁴ NLP Expert Center, Data:Lab, Volkswagen AG, Munich, Germany
<https://www.cis.uni-muenchen.de/>

Abstract. We propose new methods for in-domain and cross-domain Named Entity Recognition (NER) on historical data for Dutch and French. For the cross-domain case, we address domain shift by integrating unsupervised in-domain data via contextualized string embeddings; and OCR errors by injecting synthetic OCR errors into the source domain and address data centric domain adaptation. We propose a general approach to imitate OCR errors in arbitrary input data. Our cross-domain as well as our in-domain results outperform several strong baselines and establish state-of-the-art results. We publish preprocessed versions of the French and Dutch Europeana NER corpora.

Keywords: Named Entity Recognition · Historical data · FLAIR

1 Introduction

Neural networks achieve good NER accuracy on high-resource domains such as modern news text or Twitter [2, 4]. But on historical text, NER often performs poorly. This is due to several challenges: i) Domain shift: Entities in historical texts can be different from contemporary entities, this makes it difficult for modern taggers to work with historical data. ii) OCR errors: historical texts – usually digitized by OCR – contain systematic errors not found in non-OCR text [14]. In addition these errors can change the surface form of entities. iii) Lack of annotation: Some historical text is now available in digitized form, but without labels, and methods are required for beneficial use of such data [16].

In this paper, we address *data centric domain adaptation* for NER tagging on historical French and Dutch data. Following Ramponi and Plank [20], data centric approaches do not adapt the model but the training data in order to

improve generalization across domains. We address both in-domain and cross-domain NER. In the cross-domain setup, we use supervised contemporary data and integrate unsupervised historical data via contextualized embeddings. We introduce artificial OCR errors into supervised modern data and find a way to perturb corpora in a general and robust way – independent of language or linguistic properties.

In the cross-domain setup as well as in-domain, our system outperforms neural and statistical state-of-the-art methods, achieving 69.3% F_1 for French and 63.4% for Dutch. With the in-domain setup, we achieve 77.9% for French and 84.2% for Dutch. If we only consider named entities that contain OCR errors, our domain-adapted cross-domain tagger even performs better (83.5% French/ 46.2% Dutch) than in-domain training (77.1% French/ 43.8% Dutch). Our main contributions are:

- Release of the preprocessed French and Dutch NER corpora¹;
- Developing synOCR to mimic historical data while exploiting the annotation of modern data;
- Training historical embeddings on a large amount of unlabeled historical data;
- Ensembling a NER system that establishes SOTA results for both languages and scenarios.

2 Methods

2.1 Architecture

We use the FLAIR NLP framework [1]. FLAIR taggers achieve SOTA results on various benchmarks and are well suited for NER. Secondly, there are powerful FLAIR embeddings. They are trained without explicit notion of words and model words as character sequences depending on their context. These two properties contribute to making atypical entities - even those with distorted surface - easier to recognize. In all our experiments, word embeddings are generated by a character-level RNN and passed to a word-level bidirectional LSTM with a CRF as the final layer. Depending on the experiment, we concatenate (\odot) additional embeddings and refer to that as *ensembling process*.

2.2 Noise Methods

Since digitizing by OCR introduces a lot of noise into the data, we recreate some of those phenomena in the modern corpora that we use for training. Our goal is to increase the similarity of historical (OCR'd) and modern (clean) data. An example drawn from the dutch training corpora can be found in Fig. 1. Words that are different from the original text are indicated in bold font.

Generation of Synthetic OCR (synOCR) Errors. This method processes every sentence by assigning a randomly selected font and a font size between 6

¹ <https://github.com/stefan-it/historic-domain-adaptation-icdar>

and 11 pt. Batches of 150 sentences are printed to PDF documents and then converted to PNG images. The images are perturbed using `imgaug`² with the following steps: (i) rotation, (ii) Gaussian blur and (iii) white or black pixel dropout. The resulting image is recognized using tesseract version 0.2.6.³. We re-align the recognized sentences with the clean annotated corpus to transfer the NER tags. For the alignment between original and degraded text we select a window of the bitext and calculate a character-based alignment cost. We then use the Wagner-Fisher algorithm [27] to obtain the best alignment path through the window and the lowest possible cost. If the cost is below a threshold, we shift the window to the mid-point of the discovered path. Otherwise, we iteratively increase the window size and re-align, until the threshold criterion is met. This procedure allows us to find an alignment with reasonable time and space resources, without risking to lose the optimal path in low-quality areas. Finally this results in an OCR-error enhanced annotated corpus with a range of recognition quality, from perfectly recognized to fully illegible. We refer to OCR-corrupted data as *synOCR'd data*.

Generation of Synthetic Corruptions. This method is applied to our modern corpora, again to introduce noise as we find it in historical data. Similar to [21], we randomly corrupt 20% of all words by (i) inserting a character or (ii) removing a character or (iii) transposing two characters. Therefore, we use the standard alphabet of French/Dutch. We re-align the corrupted tokens with the clean annotated tokens while maintaining the sentence boundaries to transfer the NER tags. Since the corruption method does not break the word boundaries we can simply map each corrupted word to the original one and retrieve the corresponding NER tag. We refer to synthetically corrupted data as *corrupted data*.

De tekst van het arrest is nog niet schriftelijk beschikbaar maar het bericht werd alvast bekendgemaakt door een communicatiebureau dat Floralux inhuurde .	<i>Original</i>
De :eks: van het attest is nog hie: schiizeujk beschikbaar maar bet helicht weéd alvas: bekendgemaakt door sen' communicacisbuxeau dac Floralux inhuurde , ' * — — —	<i>synOCR</i>
De tekst van het arrest is nog niet schriftelijk beschikbaar maar ht bCbericht werd alvast bekendgemaakt door eeRn communicatiebureau dat Floralux inhuurde .	<i>corrupted</i>

Fig. 1. Example from the dutch train set. Text in its original, the synOCR'd and the corrupted form.

² <https://github.com/aleju/imgaug>

³ <https://github.com/tesseract-ocr/>

2.3 Embeddings

We experiment with various common embeddings and integrate them in our neural system. Some of them are available in the community and some others we did train on data described in Sect. 3.1.

Flair Embeddings. [3] present contextual string embeddings which can be extracted from a neural language model. FLAIR embeddings use the internal states of a trained character language model at token boundaries. They are contextualized because a word can have different embeddings depending on its context. These embeddings are also less sensitive to misspellings and rare words and can be learned on unlabeled corpora. We also use *multilingual FLAIR embeddings*. They were trained on a mix of corpora from different domains (Web, Wikipedia, Subtitles, News) and languages.

Historical Embeddings. We train FLAIR embeddings on large unlabeled historical corpora from a comparable time period (see Sect. 3.1) and refer to them as *historical embeddings*.

BERT Embeddings. Since *BERT embeddings* [9] produce state-of-the-art results for a wide range of NLP tasks, we also experiment with multilingual BERT embeddings⁴. BERT embeddings are subword embeddings based on a bidirectional transformer architecture and can model the context of a word. For NER on CoNLL-03 [25], BERT embeddings do not perform as well as on other tasks [9] and we want to examine if this observation holds for a cross-domain scenario with different data.

FastText Embeddings. We do also use *FastText embeddings* [6] which are widely used in NLP. They can be efficiently trained and address character-level phenomena. Subwords are used to represent the target word (as a sum of all its subword embeddings). We use pre-trained FastText embeddings for French/Dutch⁵.

Character-Level Embeddings. Due to the OCR errors out-of-vocabulary problems occur. Lample et al. [15] create *character embeddings*, passing all characters in a sentence to a bidirectional LSTM. To obtain word representations, the forward and backward representations of all the characters of the word from this LSTM are concatenated. Having the character embedding, every single words vector can be formed even if it is out-of-vocabulary. Therefore, we do also compute these embeddings for our experiments.

3 Experiments

In the cross-domain setup, we train on modern data (clean or synOCR'd) and test on historical data (OCR'd). In the in-domain setup, we train and test on a set of historical data (OCR'd). We do use different combinations of embeddings and also use our noise methods in the experiments.

⁴ We use the cased variant from <https://huggingface.co/bert-base-multilingual-cased>

⁵ <https://fasttext.cc/docs/en/crawl-vectors.html>

3.1 Data

We use different data sources for our experiments from which some are openly available and some historical data come from an in-house project. For an overview of different properties (domain, labeling, size, language) see Table 1.

Annotated Historical Data. Our annotated historical data comes from the Europeana Newspapers collection⁶, which contains historical news articles in 12 languages published between 1618 and 1990. Parts of the German, Dutch and French data were manually annotated with NER tags in IO/IOB format for PER (person), LOC (location), ORG (organization) by Neudecker [17]. Each NER corpus contains 100 scanned pages (with OCR accuracy over 80%), amounting to 207K tokens for French and 182K tokens for Dutch.

We preprocess the data as follows. We perform sentence splitting, filter out metadata, re-tokenize punctuation and convert all annotations to IOB1 format. We split the data 80/10/10 into train/dev/test. We will make this preprocessed version available in CoNLL format.

Annotated Modern Data. For the French cross-domain experiments, we use the French WikiNER corpus [18]. WikiNER is tagged in IOB format with an additional MISC (miscellaneous) category; we convert the tags to our Europeana format. For better comparability we downsample (sentence-wise) the corpus from 3.5M to 525K tokens. Therefore, entire sentences were sampled uniformly at random without replacement. For Dutch, we use the CoNLL-02 corpus [24], which consists of four editions of the Belgian Dutch newspaper “De Morgen” from the year 2000. The data comprises 309K tokens and is annotated for PER, ORG, LOC and MISC. We convert the tags to our Europeana format.

Unlabeled Historical Data. For historical French, we use “Le Temps”, a journal published between 1861 and 1942 (initially under a different name), a similar time period as the Europeana Newspapers. The corpus contains 977M tokens and is available from the National Library of France.⁷ For historical Dutch, we use data from an in-house OCR project. The data is from the 19th century and it consists of 444M tokens. We use the unlabeled historical data to pre-train historical embeddings (see Sect. 2.3).

3.2 Baselines

We experiment with three baselines. (i) The Java implementation⁸ of the Stanford NER tagger [12]. (ii) A version of Stanford NER published by Neudecker [17]⁹ that was trained on Europeana. In contrast to our system they trained

⁶ <http://www.europeana-newspapers.eu/>

⁷ <https://www.bnf.fr/fr>

⁸ <https://nlp.stanford.edu/software/CRF-NER.html>

⁹ <https://lab.kb.nl/dataset/europeana-newspapers-ner>

Table 1. Number of tokens per dataset in our experiments.

	Domain	Data	Labeled	Size
French	Historical	Europeana NER	+	207K
	Modern	WikiNER	+	525K
	Historical	“Le Temps”	–	977M
Dutch	Historical	Europeana NER	+	182K
	Modern	CoNLL-02	+	309K
	historical	in-house OCR	–	444M

theirs on the entire amount of the labeled Europeana corpora with 4-fold cross validation. (iii) NN base. The neural network (see Sect. 2.1) with FastText, character and multilingual FLAIR embeddings, as recommended in Akbik et al. [1]. For French, we also list the result reported by Çavdar [8]. Since we do not have access to their implementation and could not confirm that their data splits conform to ours, we could not compute the combined F_1 score or test for significance.

Table 2. Results (F_1 scores on French/Dutch Europeana test set) of training on Europeana French/Dutch training set. *Hist. Embs.* are historical embeddings. Scores marked with * are significantly lower than *NN base* \odot *hist. Es*.

French models	Overall	PER	ORG	LOC
Çavdar		0.68	0.37	0.68
Stanford NER tagger	0.662*	0.569*	0.335*	0.753*
Stanford Neudecker	0.750*	0.750*	0.505	0.826*
NN base	0.741*	0.703*	0.320*	0.813*
NN base \odot hist. Embs	0.779	0.759	0.498	0.832
Dutch models	Overall	PER	ORG	LOC
Stanford NER tagger	0.696*	0.640*	0.333*	0.794*
Stanford Neudecker	0.623*	0.700*	0.253*	0.702*
NN base	0.818*	0.809*	0.442*	0.871*
NN base \odot hist. Embs	0.842	0.833	0.480	0.891

4 Results and Discussion

We evaluate our systems using the CoNLL-2000 evaluation script¹⁰, with F_1 score. To check statistical significance we use randomized testing [28] and results are considered significant if $p < 0.05$.

¹⁰ <https://www.clips.uantwerpen.be/conll2000/chunking/conllevel.txt>

4.1 In-domain Setup

For both languages we achieve the best results with NN base \odot historical embeddings. With this setup we can produce F_1 scores of around 80% for both languages, which outperforms all three baselines in the overall performance significantly. The results are presented in Table 2. For French, the overall F_1 score as well as the F_1 for LOC and ORG is best with NN base \odot historical embeddings. For ORG the pre-trained tagger of Neudecker [17] works best, which could be due to the gazetteer information they included and of course due to the fact that they train with the entire Europeana data. We hypothesize that the category with the most structural changes over time is ORG. In the military or ecclesiastical context in particular, there are a number of names that no longer exist (in this form). For Dutch we observe the best overall performance with NN base \odot historical embeddings except for all entity types.

4.2 Cross-Domain Setup

As shown in Table 3, NN base performs better than the statistical Stanford NER baseline, which is in line with the observations for the in-domain training. We experimented with concatenating BERT embeddings to NN base. For both languages this increases the performance (Table 3, NN base \odot BERT). The usage of the historical embeddings is also very beneficial for both languages. We can achieve our best results by using BERT for Dutch and by using historical embeddings for French. We conclude that the usage of modern pre-trained language models is crucial for the performance of NER taggers.

We generated synthetic corruptions for the WikiNER/CoNLL corpus. This could not outperform NN base for both languages. The training on synOCR'd WikiNER/CoNLL gives slightly worse results than NN base too. The corruption of the training data without the usage of any embeddings seems to harm performance drastically, what is in line with the observation of Hamdi et al. [13]. It is striking that the training on corrupted/synOCR'd Dutch gives especially bad results for PER compared to French. A look at the Dutch test set shows that many entities are abbreviated first names (e.g. in **A J van Roozendaal**) and are often misrecognized what leads to a performance decrease. For French the combination of NN base and historical embeddings, trained on corrupted data or on synOCR'd (*NN ensemble corrupted/ NN ensemble synOCR'd*) gives the best results and outperforms all other systems. For Dutch *NN ensemble corrupted* and *NN ensemble synOCR* give slightly worse results than NN base \odot BERT and NN base \odot historical embeddings, but performs better than the tagger trained on synOCR'd or corrupted data only (Table 3, NN ensemble).

Ablation Study. We analyze our results and examine the composition of NN ensemble synOCR more closely (since the results for NN ensemble corrupted are very similar we perform the analysis for NN ensemble synOCR as a representative for both NN ensemble).

The ablation study (see Table 4) shows that NN ensemble benefits from different information in combination. For French NN ensemble gives the best results only for PER. The overall performance increases if we do not use character level embeddings. There is a big performance loss if we omit the historical embeddings. If we do not train on synOCR'd data the performance decreases. For Dutch we can observe these facts even more clearly. If we do not train on synOCR'd data the F_1 score even increases. If omitting the historical embeddings we loose performance as well.

Table 3. Results of training on WikiNER/CoNLL corpus. Scores marked with * are significantly lower than *NN ensemble*.

French models	Overall	PER	ORG	LOC
Çavdar		0.48	0.11	0.56
Stanford NER tagger	0.536*	0.451*	0.059*	0.618*
NN base	0.646*	0.636*	0.096*	0.721*
NN base \odot BERT	0.660	0.639*	0.163	0.725*
NN base \odot hist. Embs.	0.672	0.661*	0.015*	0.748
Corrupted WikiNER	0.627*	0.635*	0.085*	0.710*
synOCR'd WikiNER	0.619*	0.590*	0.078	0.710
NN ensemble corrupted	0.693	0.624	0.063	0.783
NN ensemble synOCR	0.684	0.710	0.111	0.744
Dutch models	Overall	PER	ORG	LOC
Stanford NER tagger	0.371*	0.217*	0.083*	0.564*
NN base	0.567*	0.493*	0.085*	0.700*
NN base \odot BERT	0.634	0.572	0.250	0.771
NN base \odot hist. Embs.	0.632	0.568	0.084	0.738*
Corrupted CoNLL	0.535*	0.376*	0.155*	0.717*
synOCR'd CoNLL	0.521*	0.327*	0.061*	0.721*
NN ensemble corrupted	0.606*	0.439*	0.158	0.799
NN ensemble synOCR	0.614	0.481	0.157	0.775

To find out why our implementation of the assumption – synOCR increases the similarity of the data and improves results – does not have the expected effect, we analyze the test sets. It shows, that only 10% of the French and 6% of the Dutch entities contain OCR errors. Therefore the wrong predictions are mostly not due to the OCR errors, but due to the inherent difficulty of recognizing entities cross-domain. This also explains why synthetic noisification does not consistently improve the system. In addition there are some illegible lines in the synOCR'd corpora consisting of dashes and metasymbols, what is not similar to real OCR errors.

Table 4. Ablation study. Results of training on the clean and the synOCR'd WikiNER/CoNLL corpus.

French models	Overall	PER	ORG	LOC
NN ensemble synOCR	0.684	0.710	0.011	0.744
- char	0.693	0.681	0.078	0.758
- word	0.686	0.664	0.080	0.756
- hist. Embs.	0.619	0.590	0.078	0.710
- synOCR'd data	0.672	0.661	0.015	0.748
Dutch models	Overall	PER	ORG	LOC
NN ensemble synOCR	0.614	0.382	0.157	0.775
- char	0.600	0.404	0.119	0.780
- word	0.584	0.430	0.102	0.745
- hist. Embs.	0.521	0.327	0.061	0.721
- synOCR'd data	0.632	0.568	0.084	0.738

D' Arras <LOC> d' où nous avons débouché , nous sommes arrivés aux premières maisons de **Siiint-lûrenl-Blangy** <LOC>. *NN_ensemble*

D' Arras <LOC> d' où nous avons débouché , nous sommes arrivés aux premières maisons de **Siiint-lûrenl-Blangy** <O>. *Stanford NER tagger*

Fig. 2. Example sentence from the French test set.

To verify our assumption we also compare the different systems only on the entities with OCR errors. Here NN ensemble outperforms both of the cross-domain baselines (Table 5, Stanford NER tagger, NN base cross-domain). Compared to the French results Dutch is a lot worse. A look at the entities shows that in the Dutch test set there are many hyphenated words where both word parts are labeled. However, if looking at the parts of the word individually, a clear assignment to an entity type cannot be made, which leads to difficulties with tagging. Though it is plausible that NN ensemble can capture specific phenomena in the historical data better, since the difference between the domains is reduced by the synthetic noisyfication and the historical embeddings. The example in Fig. 2 drawn from the test set shows, that NN ensemble can handle noisy entities well in contrast to e.g. the Stanford NER tagger. Thus in a scenario with many OCR errors the NN ensemble performs well.

5 Related Work

There is some research on using natural language processing for improving OCR for historical documents [5, 26] and also on NER for historical documents [11]. In the latter - a shared task for Named Entity Processing in historical documents - Ehrmann et al. find that OCR noise drastically harms systems performance. Like

Table 5. Results on entities with OCR errors in the French/Dutch test set. Scores marked with * are significantly lower than *NN ensemble*.

Models	French	Dutch
Stanford NER tagger	0.661*	0.207*
NN base in-domain	0.771	0.438
NN base cross-domain	0.783	0.200*
NN ensemble synOCR	0.835	0.462

us several participants (e.g. [7, 23]) also use language models that were trained on historical data to boost the performance of NER taggers. Schweter and Baiter [22] explore NER for historical German data in a cross-domain setting. Like us, they train a language model on unannotated in-domain data and integrate it into a NER tagger. In addition to the above mentioned work, we employ “OCR noisification” (Sect. 2.2) and examine the influence of different pretrained embeddings systematically. Çavdar [8] addresses NER and relation extraction on the French Europeana Newspaper corpus. Ehrmann et al. [10] investigate the performance of NER systems on Swiss historical Newspapers and show that historical texts are a great challenge compared to contemporary texts. They find that the LOC class entities causes the most difficulties in the recognition of named entities. The recent work of Hamdi et al. [13] investigates the impact of OCR errors on NER. To do so, they also perturb modern corpora synthetically with different degrees of error rates. They experiment with Spanish, Dutch and English. Like us they perturb the Dutch CoNLL corpus and train NER taggers on that data. Unlike us they do also train on a subset of the perturbed corpus. We test on a subset of the Dutch Europeana corpus. Hamdi et al. [13] show that neural taggers perform better compared to other taggers like the Stanford NER tagger and they also prove that performance decreases drastically if the OCR error rate increases. Piktus et al. [19] learn misspelling-oblivious FastText embeddings from synthetic misspellings generated by an error model for part-of-speech tagging. We use a similar corruption method, but we also use synOCR and historical embeddings for NER.

6 Conclusion

We proposed new methods for in-domain and cross-domain Named Entity Recognition (NER) on historical data and addressed data centric domain adaptation. For the cross-domain case, we handle domain shift by integrating non-annotated historical data via contextualized string embeddings; and OCR errors by injecting synthetic OCR errors into the modern data. This allowed us to get good results when labeled historical data is not available and the historical data is noisy. For training on contemporary corpora and testing on historical corpora we achieve new state-of-the-art results of 69.3% on French and 63.4% on Dutch. For the in-domain case we obtain state-of-the-art results of 77.9% for French and

84.2% for Dutch. There is an increasing demand for advancing the digitization of the world’s cultural heritage. High quality digitized historical data, with reliable meta information, will facilitate convenient access and search capabilities, and allow for extensive analysis, for example of historical linguistic or social phenomena. Since named entity recognition is one of the most fundamental labeling tasks, it would be desirable that advances in this area translate to other labeling tasks in processing of historical data as well.

Acknowledgement. This work was funded by the European Research Council (ERC #740516).

A Appendix

Detailed Information About Experiments and Data

The computing infrastructure we use for all our experiments is one GeForce GTX 1080Ti GPU with an average runtime of 12 h per experiment. For the French and Dutch baseline model *NN base* we count 15,895,683 parameters each. For the French *NN ensemble* model there are 88,264,777 parameters and 96,895,161 parameters for the Dutch *NN ensemble*.

The Europeana Newspaper Corpus is split 80/10/10 into train/dev/test (Table 6). The downsampled French WikiNER corpus is split 70/15/15 into train/dev/test and the Dutch CoNLL-02 corpus is already split in its original version. The downloadable version of the data can be found here: <https://github.com/stefan-it/historic-domain-adaptation-icdar>.

Table 6. Number of tokens for each datasplit.

Dataset	Train	dev	Test
French Europeana	167,723	18,841	20,346
Dutch Europeana	147,822	16,391	18,218
French WikiNER	411,687	88,410	88,509
Dutch ConNLL-02	202,930	68,994	37,761

References

1. Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., Vollgraf, R.: FLAIR: an easy-to-use framework for state-of-the-art NLP. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations), pp. 54–59. Association for Computational Linguistics, Minneapolis (June 2019). <https://www.aclweb.org/anthology/N19-4010>

2. Akbik, A., Bergmann, T., Vollgraf, R.: Pooled contextualized embeddings for named entity recognition. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long and Short Papers), vol. 1, pp. 724–728. Association for Computational Linguistics, Minneapolis (June 2019)
3. Akbik, A., Blythe, D., Vollgraf, R.: Contextual string embeddings for sequence labeling. In: 27th International Conference on Computational Linguistics, COLING 2018, pp. 1638–1649 (2018)
4. Baeveski, A., Edunov, S., Liu, Y., Zettlemoyer, L., Auli, M.: Cloze-driven pretraining of self-attention networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 5359–5368. Association for Computational Linguistics, Hong Kong (November 2019). <https://www.aclweb.org/anthology/D19-1539>
5. Berg-Kirkpatrick, T., Durrett, G., Klein, D.: Unsupervised transcription of historical documents. In: Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 207–217. Association for Computational Linguistics, Sofia (August 2013). <https://www.aclweb.org/anthology/P13-1021>
6. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguist.* **5**, 135–146 (2017). <https://www.aclweb.org/anthology/Q17-1010>
7. Boros, E., et al.: Robust named entity recognition and linking on historical multilingual documents. In: Conference and Labs of the Evaluation Forum (CLEF 2020). Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum, vol. 2696, pp. 1–17. CEUR-WS Working Notes, Thessaloniki (September 2020). <https://hal.archives-ouvertes.fr/hal-03026969>
8. Çavdar, M.: Distant supervision for French relation extraction (2017)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long and Short Papers), vol. 1, pp. 4171–4186. Association for Computational Linguistics, Minneapolis (June 2019). <https://www.aclweb.org/anthology/N19-1423>
10. Ehrmann, M., Colavizza, G., Rochat, Y., Kaplan, F.: Diachronic evaluation of NER systems on old newspapers. In: KONVENS (2016)
11. Ehrmann, M., Romanello, M., Flückiger, A., Clematide, S.: Extended overview of CLEF HIPE 2020: named entity processing on historical newspapers. Zenodo (October 2020)
12. Finkel, J.R., Grenager, T., Manning, C.: Incorporating non-local information into information extraction systems by Gibbs sampling. In: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL 2005), pp. 363–370. Association for Computational Linguistics, Ann Arbor (June 2005). <https://www.aclweb.org/anthology/P05-1045>
13. Hamdi, A., Jean-Caurant, A., Sidère, N., Coustaty, M., Doucet, A.: Assessing and minimizing the impact of OCR quality on named entity recognition. In: Hall, M., Merčun, T., Risse, T., Duchateau, F. (eds.) TPD 2020. LNCS, vol. 12246, pp. 87–101. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-54956-5_7
14. Jean-Caurant, A., Tamani, N., Courboulay, V., Burie, J.: Lexicographical-based order for post-OCR correction of named entities. In: 14th IAPR International

- Conference on Document Analysis and Recognition, ICDAR 2017, Kyoto, Japan, November 9–15, 2017, pp. 1192–1197. IEEE (2017). <https://doi.org/10.1109/ICDAR.2017.197>
15. Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., Dyer, C.: Neural architectures for named entity recognition. In: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 260–270. Association for Computational Linguistics, San Diego (June 2016). <https://www.aclweb.org/anthology/N16-1030>
 16. Martinek, J., Lenc, L., Král, P., Nicolaou, A., Christlein, V.: Hybrid training data for historical text OCR, pp. 565–570 (September 2019). <https://doi.org/10.1109/ICDAR.2019.00096>
 17. Neudecker, C.: An open corpus for named entity recognition in historic newspapers. In: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016), pp. 4348–4352. European Language Resources Association (ELRA), Portorož (May 2016). <https://www.aclweb.org/anthology/L16-1689>
 18. Nothman, J., Ringland, N., Radford, W., Murphy, T., Curran, J.R.: Learning multilingual named entity recognition from Wikipedia. *Artif. Intell.* **194**, 151–175 (2013)
 19. Piktus, A., Edizel, N.B., Bojanowski, P., Grave, E., Ferreira, R., Silvestri, F.: Misspelling oblivious word embeddings. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Long and Short Papers), vol. 1, pp. 3226–3234. Association for Computational Linguistics, Minneapolis (June 2019). <https://www.aclweb.org/anthology/N19-1326>
 20. Ramponi, A., Plank, B.: Neural unsupervised domain adaptation in NLP—a survey (2020)
 21. Schick, T., Schütze, H.: Rare words: a major problem for contextualized embeddings and how to fix it by attentive mimicking. *CoRR abs/1904.06707* (2019). <http://arxiv.org/abs/1904.06707>
 22. Schweter, S., Baiter, J.: Towards robust named entity recognition for historic German. In: Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019), pp. 96–103. Association for Computational Linguistics, Florence (August 2019). <https://www.aclweb.org/anthology/W19-4312>
 23. Schweter, S., März, L.: Triple E - effective ensembling of embeddings and language models for NER of historical German. In: Cappellato, L., Eickhoff, C., Ferro, N., Névéal, A. (eds.) Working Notes of CLEF 2020 - Conference and Labs of the Evaluation Forum, Thessaloniki, Greece, September 22–25, 2020. CEUR Workshop Proceedings, vol. 2696. CEUR-WS.org (2020). http://ceur-ws.org/Vol-2696/paper_173.pdf
 24. Tjong Kim Sang, E.F.: Introduction to the CoNLL-2002 shared task: language-independent named entity recognition. In: COLING-02: The 6th Conference on Natural Language Learning 2002 (CoNLL-2002) (2002). <https://www.aclweb.org/anthology/W02-2024>
 25. Tjong Kim Sang, E.F., De Meulder, F.: Introduction to the CoNLL-2003 shared task: language-independent named entity recognition. In: Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003, pp. 142–147 (2003). <https://www.aclweb.org/anthology/W03-0419>
 26. Vobl, T., Gotscharek, A., Reffle, U., Ringlstetter, C., Schulz, K.U.: Pocoto - an open source system for efficient interactive postcorrection of ocred historical texts. In: Proceedings of the First International Conference on Digital Access to Textual

- Cultural Heritage, DATECH 2014, pp. 57–61. ACM, New York (2014). <http://doi.acm.org/10.1145/2595188.2595197>
27. Wagner, R.A., Fischer, M.J.: The string-to-string correction problem. *J. ACM* **21**(1), 168–173 (1974). <https://doi.org/10.1145/321796.321811>
 28. Yeh, A.: More accurate tests for the statistical significance of result differences. In: The 18th International Conference on Computational Linguistics, COLING 2000, vol. 2 (2000). <https://www.aclweb.org/anthology/C00-2137>