

Improving Lemmatization of Non-Standard Languages with Joint Learning

Enrique Manjavacas¹, Ákos Kádár², and Mike Kestemont¹

¹Computational Linguistics and Psycholinguistics Research Center, University of Antwerp, Belgium

²Tilburg Center for Cognition and Communication, Tilburg University, The Netherlands

{firstname, lastname}@uantwerpen.be, a.kadar@uvt.nl

Abstract

Lemmatization of standard languages is concerned with (i) abstracting over morphological differences and (ii) resolving token-lemma ambiguities of inflected words in order to map them to a dictionary headword. In the present paper we aim to improve lemmatization performance on a set of non-standard historical languages in which the difficulty is increased by an additional aspect (iii): spelling variation due to lacking orthographic standards. We approach lemmatization as a string-transduction task with an encoder-decoder architecture which we enrich with sentence context information using a hierarchical sentence encoder. We show significant improvements over the state-of-the-art when training the sentence encoder jointly for lemmatization and language modeling. Crucially, our architecture does not require POS or morphological annotations, which are not always available for historical corpora. Additionally, we also test the proposed model on a set of typologically diverse standard languages showing results on par or better than a model without enhanced sentence representations and previous state-of-the-art systems. Finally, to encourage future work on processing of non-standard varieties, we release the dataset of non-standard languages underlying the present study, based on openly accessible sources.

1 Introduction

Lemmatization is the task of mapping a token to its corresponding dictionary head-form to allow downstream applications to abstract away from orthographic and inflectional variation (Knowles and Mohd Don, 2004). While lemmatization is considered to be solved for analytic and resource-rich languages such as English, it remains an open challenge for morphologically complex (e.g. Estonian, Latvian) and low-resource languages with unstable orthography (e.g. historical languages).

Especially for languages with higher surface variation, lemmatization plays a crucial role as a pre-processing step for downstream tasks such as topic modeling, stylometry and information retrieval.

In the case of standard languages, lemmatization complexity arises primarily from two sources: (i) morphological complexity affecting the number of inflectional patterns a lemmatizer has to model and (ii) token-lemma ambiguities (e.g. “living” can refer to lemmas “living” or “live”) which require modeling sentence context information. In the case of historical languages, however, the aforementioned spelling variation introduces further complications. For instance, the regularity of the morphological system is drastically reduced since the evidence supporting token-lemma mappings becomes more sparse. As an example, while the modern Dutch lemma “jaar” (en. “year”) can be inflected in 2 different ways (“jaar”, “jaren”), in a Middle Dutch corpus used in this study it is found in combination with 70 different forms (“iare”, “ior”, “jaer”, etc.). Moreover, spelling variation increases token-lemma ambiguities by conflating surface realizations of otherwise unambiguous tokens—e.g. Middle Low German “bath” can refer to lemmas “bat” (en. “bad”) and “bidden” (en. “bet”) due to different spellings of the dental occlusive in final position.

Spelling variation is not exclusive of historical languages and it can be found in contemporary forms of communication, such as microblogs, with loose orthographic conventions (Crystal, 2001). An important difference, however, is that while for modern languages normalization is feasible (Schulz et al., 2016), for many historic languages such is not possible, because one is dealing with an amalgam of regional dialects that lacked any sort of supra-regional variant functioning as target domain (Kestemont et al., 2016).

In the present paper, we apply representation

learning to lemmatization of historical languages. Our method shows improvements over a plain encoder-decoder framework, which reportedly achieves state-of-the-art performance on lemmatization and morphological analysis (Bergmanis and Goldwater, 2018; Peters et al., 2018). In particular, we make the following contributions:

1. We introduce a simple joint learning approach based on a bidirectional Language Model (LM) loss and achieve relative improvements in overall accuracy of 7.9% over an encoder-decoder trained without joint loss and 30.72% over edit-tree based approaches.
2. We provide a detailed analysis of the linguistic and corpus characteristics that explain the amount of improvement we can expect from LM joint training.
3. We probe the hidden representations learned with the joint loss and find them significantly better predictors of POS-tags and other morphological categories than the representations of the simple model, confirming the efficiency of the joint loss for feature extraction.

Additionally, we test our approach on a typologically varied set of modern standard languages and find that the joint LM loss significantly improves lemmatization accuracy of ambiguous tokens over the encoder-decoder baseline (with a relative increase of 15.1%), but that, in contrast to previous literature (Chakrabarty et al., 2017; Bergmanis and Goldwater, 2018), the overall performance of encoder-decoder models is not significantly higher than that of edit-tree based approaches. Taking into account the type of inflectional morphology dominating in a particular language, we show that the benefit of encoder-decoder approaches is highly dependent on typological morphology. Finally, to assure reproducibility, all corpus preprocessing pipelines and train-dev-test splits are released. With this release, we hope to encourage future work on processing of lesser studied non-standard varieties.¹

¹ Datasets and training splits are available at <https://www.github.com/emanjavacas/pie-data>. Experiments are conducted with our framework `pie` available at: <https://www.github.com/emanjavacas/pie>. All our experiments are implemented using PyTorch (Paszke et al., 2017).

2 Related Work

Modern data-driven approaches typically treat lemmatization as a classification task where classes are represented by binary edit-trees induced from the training data. Given a token-lemma pair, its binary edit-tree is induced by computing the prefix and suffix around the longest common subsequence, and recursively building a tree until no common character can be found. Such edit-trees manage to capture a large proportion of the morphological regularity, especially for languages that rely on suffixation for morphological inflection (e.g. Western European languages), for which such methods were primarily designed.

Based on edit-tree induction, different lemmatizers have been proposed. For example, Chrupala et al. (2008) use a log-linear model and a set of hand-crafted features to decode a sequence of edit-trees together with the sequence of POS-tags using a beam-search strategy. A related approach is presented by Gesmundo and Samardži (2012), where edit-trees are extracted using a non-recursive version of the binary edit-tree induction approach. More recently, Cotterell et al. (2015) have used an extended set of features and a second-order CRF to jointly predict POS-tags and edit-trees with state-of-the-art performance. Finally, Chakrabarty et al. (2017) employed a softmax classifier to predict edit-trees based on sentence-level features implicitly learned with a neural encoder over the input sentence.

With the advent of current encoder-decoder architectures, lemmatization as a string-transduction task has gained interest partly inspired by the success of such architectures in Neural Machine Translation (NMT). For instance, Bergmanis and Goldwater (2018) apply a state-of-the-art NMT system with the lemma as target and as source the focus token with a fixed window over neighboring tokens. Most similar to our work is the approach by Kondratyuk et al. (2018), which conditions the decoder on sentence-level distributional features extracted from a sentence-level bidirectional RNN and morphological tags.

Recently, work on non-standard historical varieties has focused on spelling normalization using rule-based, statistical and neural string-transduction models (Pettersson et al., 2014; Bollmann and Søgaard, 2016; Tang et al., 2018). Previous studies on lemmatization of historical variants focused on evaluating off-the-shelf systems.

For instance, Eger et al. (2016) evaluates different pre-existing models on a dataset of German and Medieval Latin, and Dereza (2018) focuses on Early Irish. The most similar to the present paper in this area is work by Kestemont et al. (2016), which tackled lemmatization of Middle Dutch with a neural encoder that extracts character and word-level features from a fixed-length token window and predicts the target lemma from a closed-set of true lemmas.

Using Language Modeling as a task to extract features in a Transfer Learning setup has gained momentum only in the last year, partly thanks to overall improvements over previous state-of-the-art across multiple tasks (NER, POS, QA, etc.). Different models have been proposed around the same idea varying in implementation, optimization and task definition. For instance, Howard and Ruder (2018) present a method to fine-tune a pre-trained LM for text classification. Peters et al. (2018) learn task-specific weighting schemes over different layer features extracted by a pretrained bidirectional LM. Recently, Akbik et al. (2018) used context-sensitive word-embeddings extracted from a bidirectional character-level LM to improve NER, POS-tagging and chunking.

3 Proposed Model

Here we describe our encoder-decoder architecture for lemmatization. In Section 3.1 we start by describing the basic formulation known from the machine translation literature. Section 3.2 shows how sentential context is integrated into the decoding process as an extra source of information. Finally, Section 3.3 describes how we learn richer representations for the encoder through the addition of an extra language modeling task.

3.1 Encoder-Decoder

We employ a character-level Encoder-Decoder architecture that takes an input token x_t character-by-character and has as goal the character-level decoding of the target lemma l_t conditioned on an intermediate representation of x_t . For token x_t , a sequence of token character embeddings c_1^x, \dots, c_n^x is extracted from embedding matrix $W_{enc} \in \mathbb{R}^{|C| \times d}$ (where $|C|$ and d represent, respectively, the size of the character vocabulary and the embedding dimensionality). These are then passed to a bidirectional RNN encoder, that computes a forward and a backward sequence of hid-

den states: $\overrightarrow{h_1^{enc}}, \dots, \overrightarrow{h_n^{enc}}$ and $\overleftarrow{h_1^{enc}}, \dots, \overleftarrow{h_n^{enc}}$. The final representation of each character i is the concatenation of the forward and the backward states: $h_i^{enc} = [\overrightarrow{h_i^{enc}}; \overleftarrow{h_i^{enc}}]$.

At each decoding step j , a RNN decoder generates the hidden state h_j^{dec} , given the lemma character embedding c_j^l from embedding matrix $W_{dec} \in \mathbb{R}^{|L| \times d}$, the previous hidden state h_{j-1}^{dec} and additional context. This additional context consists of a summary vector r_j obtained via an attentional mechanism (Bahdanau et al., 2014) that takes as input the previous decoder state h_{j-1}^{dec} and the sequence of encoder activations $h_1^{enc}, \dots, h_n^{enc}$.² Finally, the output logits for character j are computed by a linear projection of the current decoder state h_j^{dec} with parameters $O \in \mathbb{R}^{H \times |L|}$, which are normalized to probabilities with the *softmax* function. The model is trained to maximize the probability of the target character sequence expressed in Equation 1 using teacher-forcing.

$$P(l_t|x_t) = \prod_{j=1}^m P(c_j^l | c_{<j}^l, r_j; \theta_{enc}, \theta_{dec}) \quad (1)$$

3.2 Adding sentential context

Lemmatization of ambiguous tokens can be improved by incorporating sentence-level information. Our architecture is similar to Kondratyuk et al. (2018) in that it incorporates global sentence information by extracting distributional features with a hierarchical bidirectional RNN over the input sequence of tokens x_1, \dots, x_m . For each token x_t , we first extract word-level features re-using the last hidden state of character-level bidirectional RNN Encoder from Section 3.1 $w_t = [\overrightarrow{h_t^{enc}}; \overleftarrow{h_t^{enc}}]$. Optionally, word-level features can be enriched with extra lookup parameters from an embedding matrix $W_{word} \in \mathbb{R}^{|V| \times e}$ – where V and e denote respectively the vocabulary size in words and the word embedding dimensionality.³ Given these word-level features w_t , the sentence-level features s_t are computed as the concatenation of

² We refer to Bahdanau et al. (2014) for the description of the attentional mechanism.

³ During development word embeddings did not contribute significant improvements on historical languages, and we therefore exclude them from the rest of the experiments. It must be noted, however, that word embeddings might still be helpful for lemmatization of standard languages where the type-token ratio is smaller as well as when pretrained embeddings are available.

forward and backward activations of an additional sentence-level bidirectional RNN $s_t = [\vec{s}_t; \overleftarrow{s}_t]$.

In order to perform sentence-aware lemmatization for token x_t , we condition the decoder on the sentence-level encoding s_t and optimize the probability given by Equation 2.

$$P(l_t|x_t) = \prod_{j=1}^m P(c_j^l | c_{<j}^l, r_j, s_t; \theta_{enc}, \theta_{dec}) \quad (2)$$

Our architecture ensures that both word-level and character-level features of each input token in a sentence can contribute to the sentence-level features at any given step and therefore to the lemmatization of any other token in the sentence. From this perspective, our architecture is more general than those presented in Kestemont et al. (2016); Bergmanis and Goldwater (2018), where sentence information is included by running the encoder over a predetermined fixed-length window of neighboring characters. Moreover, we let the character-level embedding extractor and the lemmatizer encoder share parameters in order to amplify the training signal coming into the latter. Figure 1 visualizes the proposed architecture.

3.3 Improved sentence-level features

We hypothesize that the training signal from lemmatization alone might not be enough to extract sufficiently high quality sentence-level features. As such we include an additional bidirectional word-level language-model loss over the input sentence. Given the forward and backward subvectors of the sentence encoding $s^t = [\vec{s}^t; \overleftarrow{s}^t]$, we train two additional softmax classifiers to predict token x^{t+1} given \vec{s}_t and x^{t-1} given \overleftarrow{s}_t with parameters O_{LMfwd} and $O_{LMbwd} \in \mathbb{R}^{S \times |V|}$.⁴

We train our model to jointly minimize the negative log-likelihood of the probability defined by Equation 2 and the LM probability defined by Equation 3.

$$P_{LM}(\mathbf{x}) = 1/2 \prod_{t=2}^n P(x_t|x_1, \dots, x_{t-1}) + 1/2 \prod_{t=1}^{n-1} P(x_t|x_{t+1}, \dots, x_n) \quad (3)$$

⁴ We have found the joint loss most effective when both forward and backward classifiers shared parameters.

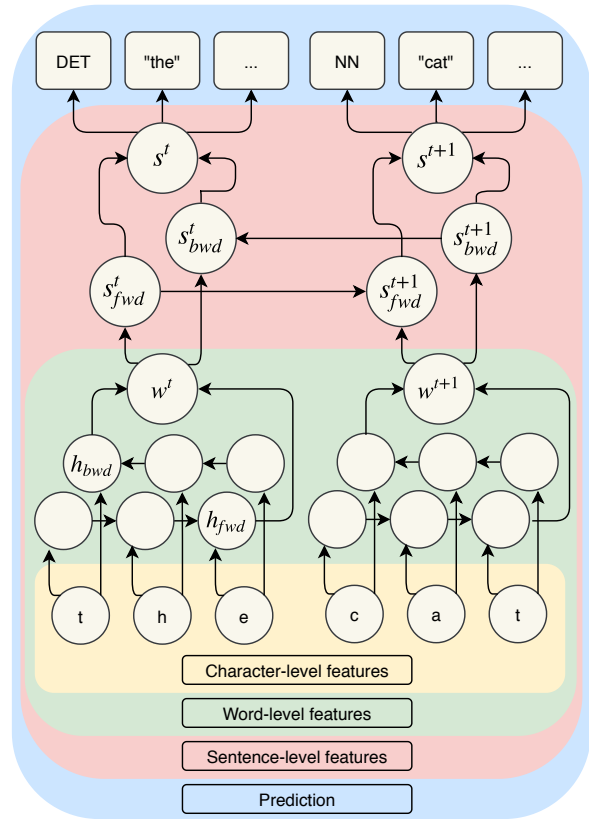


Figure 1: Hierarchical sentence encoding architecture with feature extraction at different levels.

Following a Multi-Task Learning (Caruana, 1997), we set a weight on the LM negative log-likelihood which we decrease over training based on lemmatization accuracy on development data to reduce its influence on training after convergence.

4 Experiments

Section 4.1 first introduces the datasets, both the newly introduced dataset of historical languages, and the dataset of modern standard languages sampled from Universal Dependencies (v2.2) corpus (Nivre et al., 2016). Finally, Section 4.2 describes model training and settings in detail.

4.1 Datasets

Historical Languages In recent years, a number of historical corpora have appeared thanks to an increasing number of digitization initiatives (Piotrowski, 2012). For the present study, we chose a representative collection of medieval and early modern datasets, favoring publicly available data, corpora with previously published results and datasets covering multiple genres and historic periods. We include a total of 8 corpora covering Middle Dutch, Middle Low German, Medieval

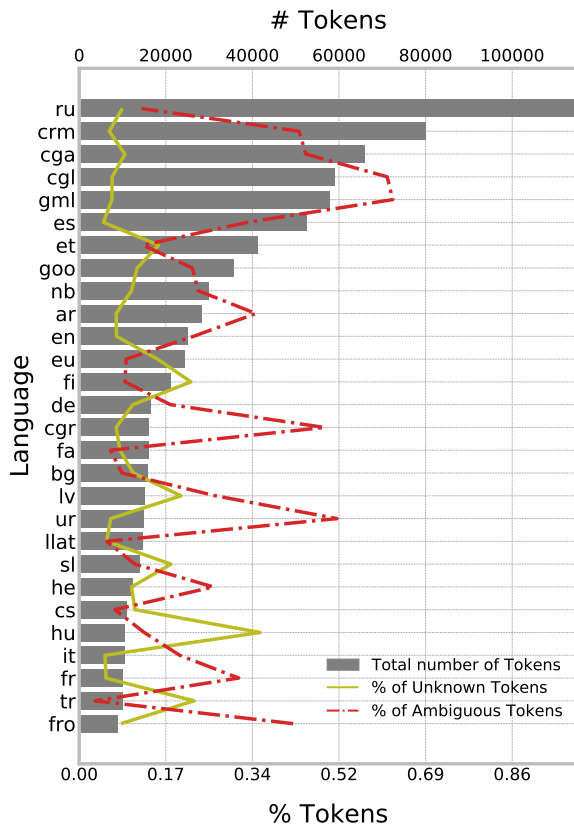


Figure 2: Statistics of total number of tokens, ambiguous and unknown tokens in the test sets. The full list of languages for both historical and standard corpora as well as the corresponding ISO 639-1 codes used in the present study can be found in the Appendix. Statistics for unknown and ambiguous tokens are shown as percentages.

French, Historical Slovene and Medieval Latin, which we take from the following sources.

Both *cga* and *cgl* contain medieval Dutch material from the Gysseling corpus curated by the Institute for Dutch Lexicology⁵ *cga* is a charter collection (administrative documents), whereas *cgl* concerns a variety of literary texts that greatly vary in length. *crm* is another Middle Dutch charter collection from the 14th century with wide geographic coverage (Van Reenen and Mulder, 1993; van Halteren and Rem, 2013). *cgr*, finally, is a smaller collection of samples from Middle Dutch religious writings that include later medieval texts (Kestemont et al., 2016). *fro* offers a corpus of Old French heroic epics, known as *chansons de geste* (Camps, 2016). *llat* dataset is taken from the Late Latin Charter Treebank, consisting of early medieval Latin documentary texts (Korkiakangas and Lassila, 2013). *goo* comes from the

⁵ <https://ivdnt.org/taalmaterialen>.

reference corpus of historical Slovene, sampled from 89 texts from the period 1584-1899 (Erjavec, 2015). *gml* refers to the reference corpus of Middle Low German and Low Rhenish texts, found in manuscripts, prints and inscriptions (Barteld et al., 2017). Finally, *cap* is a corpus of early medieval Latin ordinances decreed by Carolingian rulers (Eger et al., 2016).

Standard Languages For a more thorough comparison between systems across domains and a better examination of the effect of the LM loss, we evaluate our systems on a set of 20 standard languages sampled from the UD corpus, trying to guarantee typological diversity while selecting datasets with at least 20k words. We use the pre-defined splits from the original UD corpus (v2.2).⁶ Figure 2 visualizes the test set sizes in terms of total, ambiguous and unknown tokens for both historical and standard languages.

4.2 Models

We refer to the full model trained with joint LM loss by *Sent-LM*. In order to test the effectiveness of sentence information and the importance of enhancing the quality of the sentence-level feature extraction, we compare against a simple encoder-decoder model without sentence-level information (*Plain*) and a model trained without joint LM loss (*Sent*). Moreover, we compare to previous state-of-the-art lemmatizers based on binary edit-tree induction: *Morfette* (Chrupala et al., 2008) and *Lemming* (Cotterell et al., 2015), which we run with default hyperparameters.

For all our models, we use the same hyperparameter values as follows. All recurrent layers have 150 cells per layer and use GRUs (Cho et al., 2014). Encoder and Decoder have 2 layers but the sentence encoder has only 1. We apply 0.25 dropout (Srivastava et al., 2014) after the embedding layer and before the output layer and 0.25 variational dropout (Gal and Ghahramani, 2016) in between recurrent layers. Models are optimized with Adam (Kingma and Ba, 2015) using an initial learning rate of 1e-3 which is reduced by 25% after each epoch without improvement on development accuracy. Models are trained until fail-

⁶ The full list of languages for both historical and standard corpora as well as the corresponding ISO 639-1 codes used in the present study can be found in the Appendix. In the cases where train-dev-test splits were not pre-defined, we randomly split sentences using 10% and 5% for test and dev respectively.

	Full	Ambiguous	Unknown
Edit-tree	91.11	91.79	35.48
Plain	91.61	87.39	65.69
Sent	93.4	91.14	66.98
Sent-LM	94.0	92.81	65.39

Table 1: Average accuracy across historical languages. Lemming and Morfette are shown aggregated by taking the best performing model per dataset.

ing to achieve any improvement for 3 consecutive epochs. Initial LM loss weight is set to 0.2 and it is halved each epoch after two consecutive epochs without achieving any improvements on development perplexity.

We use sentence boundaries when given and otherwise use POS tags corresponding to full stops as clues. In any case, sentences are split into chunks of maximum 35 words to accommodate to limited memory. Target lemmas during both training and testing are lowercased in agreement with the implementation of Lemming and Morfette, which also do so. For models with joint loss, we truncate the output vocabulary to the top 50k most frequent words for similar reasons. We run a maximum of 100 optimization epochs in randomized batches containing 25 sentences each. The learning rate is decreased by a factor of 0.75, after every 2 epochs without accuracy increase on held-out data and learning stops after failing to improve for 5 epochs. Decoding is done with beam search with a beam size of 10.⁷

5 Results

As is customary, we report exact-match accuracy on target lemmas. Besides overall accuracy, we also compute accuracy of ambiguous tokens (i.e. tokens that map to more than 1 lemma in the training data) and unknown tokens (i.e. tokens that do not appear in the training data).

5.1 Historical languages

Table 1 shows the aggregated results over all datasets in our historical language corpus.⁸ In 4 cases (cga, cgl, crm and gml), Lemming failed to converge due to memory requirements exceeding 250G RAM due to the large amount of

⁷ For all languages, we observed relatively small gains ranging from 0.1% to 0.5% in overall accuracy.

⁸ We aggregate both edit-tree based approaches by selecting the best performing model for each corpus. When Lemming converge, the results were better than Morfette.

	Full	Ambiguous	Unknown
K-2016	91.88		51.64
Edit-tree	89.01	91.18	20.46
Plain	90.21	87.4	61.93
Sent	92.55	91.6	64.61
Sent-LM	93.25	93.31	62.1

Table 2: Accuracy for the Gysseling subcorpora.

edit-trees. Following Søgaard et al. (2014), we compute p-values with a Wilcoxon’s signed rank test. Sent-LM is the best performing model with a relative improvement of 7.9% ($p < .01$) over Sent and 30.72% ($p < .01$) over the edit-tree approach on full datasets and 10.27% ($p < .1$) and 18.66% ($p < .01$) on ambiguous tokens. Moreover, the edit-tree approach outperforms encoder-decoder models Plain and Sent on ambiguous tokens, and it is only due to the joint loss that the encoder-decoder paradigm gains an advantage. Finally, for tokens unseen during training, the best performing model is Sent with a relative error reduction of 47% ($p < .01$) over the edit-tree approach and 4.77% ($p < .1$) over Sent-LM.

Table 2 compares scores for a subset from the corpora coming from the Gysseling corpus, which have been used in previous work on lemmatization of historical languages. The model described by Kestemont et al. (2016) is included as K-2016 for comparison.⁹ It is apparent that both Sent and Sent-LM outperform K-2016 on full and unknown tokens. It is worth noting that K-2016, a model that uses distributed contextual features but no edit-tree induction, performs better than Plain – which highlights the importance of context for the lemmatization of historical languages –, and also better than the edit-tree approaches – which highlights the difficulty of tree induction on this dataset. We find Sent-LM to have a significant advantage over Sent on full and ambiguous tokens, but a disadvantage vs Sent and Plain on unknown tokens.

5.2 Standard languages

Table 4 shows overall accuracy scores aggregated across all languages.¹⁰ We observe that on average Sent-LM is the best model on full datasets.

⁹ Unfortunately, scores on ambiguous tokens were not reported and therefore cannot be compared.

¹⁰ Similarly to results on historical languages, we aggregate Morfette and Lemming due to the later failing to converge on et.

	Full			Ambiguous			Unknown		
	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
Edit-tree	96.34	93.02	98.37	92.47	92.47	97.5	84.99	74.66	91.39
Plain	96.21	94.6	97.5	88.87	90.51	95.34	85.43	83.78	86.37
Sent	96.42	94.41	97.84	91.12	92.01	96.65	85.44	84.02	86.67
Sent-LM	96.52	94.62	97.86	93.01	92.07	97.48	85.15	83.32	85.36

Table 3: Average accuracy across morphologically related standard languages. Type 1 encloses bg, cs, lv, ru and sl. Type 2 comprises et, fi, hu, tr. Finally, Type 3 encompasses de, en, es, fr, it and nb.

	Full	Ambiguous	Unknown
Edit-tree	96.1	94.35	83.26
Plain	95.93	91.44	83.02
Sent	96.19	93.25	82.61
Sent-LM	96.28	94.08	82.58

Table 4: Average accuracy across all 20 standard languages. Lemming and Morfette are shown aggregated by taking the best performing model per dataset.

However, in contrast to previous results, the edit-tree approach has an advantage over all encoder-decoder models for both ambiguous and unknown tokens.

Since the differences in performance are not statistically significant ($p > 0.05$), we seek to shed light on the advantages and disadvantages of the encoder-decoder and edit-tree paradigms by conducting a more fine-grained analysis with respect to the morphological typology of the considered languages. To this end, we group languages into morphological types depending on the dominant morphological processes of each language and aggregate scores over languages in each type:

Type 1. Balto-Slavic languages which are known for their strongly suffixing morphology and complex case system.

Type 2. Uralic and Altaic languages, which are characterized by agglutinative morphology and a tendency towards monoexponential case and vowel harmony.

Type 3. Western European languages with a tendency towards synthetic morphology and partially lacking nominal case.

Table 3 shows accuracy scores per morphological group for each model type. It is apparent that the Edit-tree approach is very effective for Type 3 languages both in ambiguous and unknown

tokens. In both Type 1 and Type 2 languages, the best overall performing model is Sent-LM. In the case of ambiguous tokens, Sent-LM achieves highest accuracy for Type 1 languages, but it is surpassed by the Edit-tree approach on Type 2 languages. Finally, in the case of unknown tokens, we observe a similar pattern to the historical languages where Plain and Sent have an advantage over Sent-LM.

6 Discussion

For clarity, we group the discussion of the main findings according to four major discussion points.

How does the joint LM loss help? As Section 5 shows, Sent-LM is the overall best model, and its advantage is biggest on ambiguous datasets, always outperforming the second-best encoder-decoder model on ambiguous tokens. For a more detailed comparison of the two models we tested the following two hypotheses: (i) the joint LM loss helps by providing sentence representations with stronger disambiguation capacities (ii) The joint LM loss helps in cases when the evidence of a token-lemma relationship is sparse —e.g in languages with highly synthetic morphological systems and in the presence of spelling variation.

As Figure 3 shows, improvement over Sent is correlated with percentage of token-lemma ambiguity in the corpus, providing evidence for hypothesis (i). Finally, as Figure 4 shows, improvement over Sent is correlated with higher token-lemma ratio, suggesting that the improvement is likely to be due to learned representations that better identify the input token. These two aspects help explain the efficiency of the joint learning approach on non-standard languages where high levels of spelling variation provide increased ambiguity by conflating unrelated forms and also lower evidence for token-lemma mappings.

Another factor certainly related to the efficiency

of the proposed joint LM-loss is the size of the training dataset. However, dataset size should be considered a necessary but not a sufficient condition for the feasibility of the joint LM-loss and has therefore weak explanation power for the performance of the proposed approach.

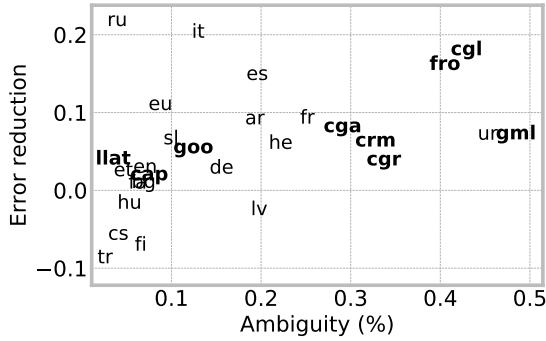


Figure 3: Error reduction of `Sent-LM` vs `Sent` by percentage of ambiguous tokens (Spearman’s $R = 0.53$; $p < .01$). Historical languages are shown in bold.

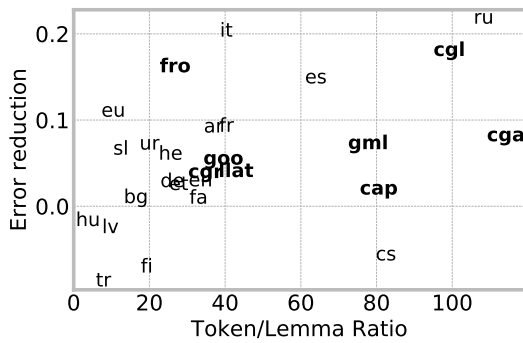


Figure 4: Error reduction of `Sent-LM` vs `Sent` by Token-Lemma ratio on 50k tokens of the training set (Spearman’s $R = 0.47$; $p < .05$). Historical languages are shown in bold.

LM loss leads to better representations In order to analyze the representations learned with the joint loss, we turn to “representation probing” experiments following current approaches on interpretability (Linzen et al., 2016; Adi et al., 2017). Using the same train-dev-test splits from the current study, we exploit additional POS, Number, Gender, Case and syntactic function (Dep) annotations provided in the UD corpora and compare the ability of the representations extracted by `Sent` and `Sent-LM` to predict these labels.¹¹ Model parameters are frozen and a linear softmax layer

¹¹ Note that not all tasks are available for all languages, due to some corpora not providing all annotations and some categories not being relevant for particular languages.

	Pos	Dep	Gender	Case	Num
Majority	82.61	62.93	85.83	80.76	83.92
<code>Sent</code>	79.27	64.14	83.39	83.01	81.01
<code>Sent-LM</code>	83.62	68.36	86.55	84.44	84.37
Support	29	20	15	19	20

Table 5: Overall accuracy of `Sent` and `Sent-LM` models and a `Majority` baseline on 5 probing tasks and actual number of languages per morphological category. All differences over `Sent` except for `Case` are significant at $p < 0.05$. Support is shown in terms of number of languages exhibiting such grammatical distinctions.

$Q \in \mathbb{R}^{H \times V}$ per task is learned using a cross-entropy loss function.¹² The results of this experiment are reported in Table 5. The classifier trained with `Sent-LM` outperforms the one with `Sent` on all considered labeling tasks, confirming the efficiency of the LM loss at extracting better representations.

Edit-tree vs. Encoder-Decoder Our fine-grained analysis suggests that the performance of the edit-tree and encoder-decoder approaches depends on the underlying morphological typology of the studied languages. Neural approaches seem to be stronger for languages with complex case systems and agglutinative morphology. In contrast, edit-tree approaches excel on more synthetic languages (e.g. Type 3) and languages with lower ambiguity (e.g. Type 2).

Figure 5 illustrates that as the number of edit-trees increase the encoder-decoder models start to excel. This is most likely due to the fact that, from an edit-tree approach perspective, a large number of trees creates a large number of classes, which leads to higher class imbalance and more sparsity. However, edit-tree based approaches do outperform representation learning methods for languages with lower number of trees, which leads to the intuition that the edit-tree formalism does provide a useful inductive bias to the task of lemmatization and it should not be discarded in future work. Our results, in fact, point to a future direction which applies the edit-tree formalism, but alleviates the edit-tree explosion by exploiting the relationships between the edit-tree classes potentially using representation learning methods.

¹² Models trained for 50 epochs using the Adam optimizer with default learning rate and training stops after 2 epochs without accuracy increase on dev set.

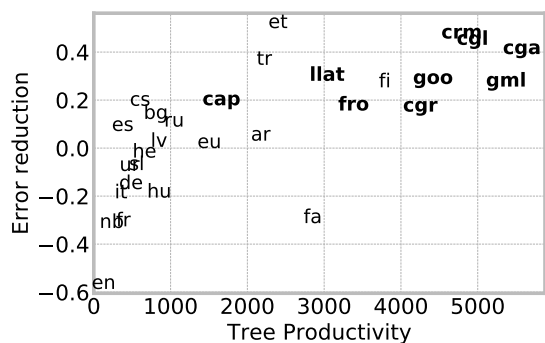


Figure 5: Error reduction of best encoder-decoder model vs best tree-edit model over tree productivity computed as number of unique binary edit-trees in the first 50k tokens of the training corpora (Spearman’s $R = 0.79$; $p < .001$). Historical languages are shown in bold.

Accuracy on unknown tokens We observe that while overall the joint loss outperforms the simpler encoder-decoder, it seems, however, detrimental to the accuracy on unknown tokens. This discrepancy is probably due to the fact that (i) unknown tokens are likely unambiguous and therefore less likely to profit from improved context representations and to (ii) our design choice of word-level language modeling, where the model is forced to predict UNK for unknown words. As *Sent-LM* is the overall best model, in future work we will explore character-level language modeling in order to harness the full potential of the joint-training approach even on unknown tokens.

7 Conclusion

We have presented a method to improve lemmatization with encoder-decoder models by improving context representations with a joint bidirectional language modeling loss. Our method sets a new state-of-the-art for lemmatization of historical languages and is competitive on standard languages. Our examination of the learned representations indicates that the LM loss helps enriching sentence representations with features that capture morphological information. In view of a typologically informed comparison of encoder-decoder and edit-tree based approaches, we have shown that the latter can be very effective for highly synthetic languages. Such result might have been overlooked in previous studies due to only considering a reduced number of languages (Chakrabarty et al., 2017) or pooling results across typology (Bergmanis and Goldwater, 2018). With respect to lan-

guages with higher ambiguity and token-lemma ratio, the encoder-decoder approach is preferable and the joint loss generally provides a substantial improvement. Finally, while other models use morphological information to improve the representation of context (e.g. edit-tree approaches), our joint language modeling loss does not rely on any additional annotation, which can be crucial in low resource and non-standard situations where annotation is costly and often not trivial.

Acknowledgments

We thank NVIDIA for donating 1 GPU that was used for the experiments in the present paper. We would also like to thank the anonymous reviewers for their valuable comments.

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2017. *Fine-grained analysis of sentence embeddings using auxiliary prediction tasks*. *ICLR '17*.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. *Contextual String Embeddings for Sequence Labeling*. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1638–1649. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. *Neural machine translation by jointly learning to align and translate*. *arXiv preprint arXiv:1409.0473*.
- Fabian Barteld, Katharina Dreessen, Sarah Ihden, and Ingrid Schröder. 2017. *Das Referenzkorpus Mittelniederdeutsch/Niederrheinisch (1200–1650)–Korpusdesign, Korpuserstellung und Korpusnutzung*. *Mitteilungen des Deutschen Germanistenverbandes*, 64(3):226–241.
- Toms Bergmanis and Sharon Goldwater. 2018. *Context Sensitive Neural Lemmatization with Lematus*. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies Volume 1*.
- Marcel Bollmann and Anders Søgaard. 2016. *Improving historical spelling normalization with bidirectional LSTMs and multi-task learning*. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 131–139. The COLING 2016 Organizing Committee.
- Jean-Baptiste Camps. 2016. *Geste: un corpus de chansons de geste*. Avec la collab. d’Elena Albarran, Alice Cochet & Lucence Ing.

- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Abhisek Chakrabarty, Onkar Arun Pandit, and Utpal Garain. 2017. [Context Sensitive Lemmatization Using Two Successive Bidirectional Gated Recurrent Networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1481–1491. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. [On the Properties of Neural Machine Translation: Encoder-Decoder Approaches](#). In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pages 103–111. Association for Computational Linguistics.
- Grzegorz Chrupala, Georgiana Dinu, and Josef van Genabith. 2008. [Learning Morphology with Morfette](#). In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Ryan Cotterell, Alexander Fraser, and Hinrich Schütze. 2015. [Joint Lemmatization and Morphological Tagging with LEMMING](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- D. Crystal. 2001. *Language and the Internet*. Cambridge University Press.
- Oksana Dereza. 2018. Lemmatization for Ancient Languages: Rules or Neural Networks? In *Conference on Artificial Intelligence and Natural Language*, pages 35–47. Springer.
- Steffen Eger, Rüdiger Gleim, and Alexander Mehler. 2016. [Lemmatization and Morphological Tagging in German and Latin: A Comparison and a Survey of the State-of-the-art](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France. European Language Resources Association (ELRA).
- Tomaž Erjavec. 2015. [Reference corpus of historical slovene goo300k 1.2](#). Slovenian language resource repository CLARIN.SI.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pages 1019–1027.
- Andrea Gesmundo and Tanja Samardžić. 2012. [Lemmatization as a Tagging Task](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 368–372, Jeju Island, Korea. Association for Computational Linguistics.
- Hans van Halteren and Margit Rem. 2013. [Dealing with orthographic variation in a tagger-lemmatizer for fourteenth century dutch charters](#). *Language Resources and Evaluation*, 47(4):1233–1259.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal Language Model Fine-tuning for Text Classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339. Association for Computational Linguistics.
- Mike Kestemont, Guy De Pauw, Renske van Nie, and Walter Daelemans. 2016. Lemmatization for variation-rich languages using deep learning. *Digital Scholarship in the Humanities*, 32(4):797–815.
- Diederik P. Kingma and Jimmy Lei Ba. 2015. [Adam: a Method for Stochastic Optimization](#). *International Conference on Learning Representations 2015*, pages 1–15.
- G. Knowles and Z. Mohd Don. 2004. The notion of a lemma. Headwords, roots and lexical sets. *International Journal of Corpus Linguistics*, 9(1):69–81.
- Daniel Kondratyuk, Tom Gavenčič, and Milan Straka. 2018. [LemmaTag: Jointly Tagging and Lemmatizing for Morphologically-Rich Languages with BRNNs](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4921–4928, Brussels, Belgium.
- Timo Korhonen and Matti Lassila. 2013. Abbreviations, fragmentary words, formulaic language: treebanking mediaeval charter material. In *Proceedings of The Third Workshop on Annotation of Corpora for Research in the Humanities*, pages 61–72.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the Ability of LSTMs to Learn Syntax-Sensitive Dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan T McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. [Universal Dependencies v1: A Multilingual Treebank Collection](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Portorož, Slovenia. European Language Resources Association (ELRA).
- Adam Paszke, Sam Gross, Soumith Chintala, and Gregory Chanan. 2017. Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration. *PyTorch: Tensors and dynamic neural networks in python with strong GPU acceleration*.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep Contextualized Word Representations](#). In *Proceedings of the 2018 Conference*

of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237. Association for Computational Linguistics.

Eva Pettersson, Beta Megyesi, and Joakim Nivre. 2014. [A Multilingual Evaluation of Three Spelling Normalisation Methods for Historical Text](#). In *Proceedings of the 8th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH)*, pages 32–41. Association for Computational Linguistics.

Michael Piotrowski. 2012. Natural language processing for historical texts. *Synthesis Lectures on Human Language Technologies*, 5(2):1–157.

Sarah Schulz, Guy De Pauw, Orphe De Clercq, Bart Desmet, Veronique Hoste, Walter Daelemans, and Lieve Macken. 2016. [Multimodular text normalization of dutch user-generated content](#). *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(4):61.

Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Hector Martínez Alonso. 2014. [What’s in a p-value in NLP?](#) In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 1–10. Association for Computational Linguistics.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A Simple Way to Prevent Neural Networks from Overfitting](#). *Journal of Machine Learning Research*, 15:1929–1958.

Gongbo Tang, Fabienne Cap, Eva Pettersson, and Joakim Nivre. 2018. [An Evaluation of Neural Machine Translation Models on Historical Spelling Normalization](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1320–1331. Association for Computational Linguistics.

P. Van Reenen and M. Mulder. 1993. Een gegevensbank van 14de-eeuwse middelnederlandse dialecten op computer. *Lexikos*, 3:259–279.

A Appendix

A.1 Dataset Statistics

Table 6 shows the dataset sources and codes from our Historical Languages.

Language	Dataset	Code
Medieval Latin	Capitularia	cap
	LLCT1	llat
Middle Dutch	Gys (Admin)	cga
	Gys (Literary)	cgl
	Religious	cgr
	Adelheid	crm
Medieval French	Geste	fro
Middle Low German	ReN	gml
Slovenian	goo300k	goo

Table 6: Corpus identifier and description in the historical languages dataset. “Gys” refers to the Gysseling corpus, which consists of several subsets.

Table 7 shows the languages from the UD corpus that were sampled for the study. We have used ISO 639-1 codes (instead of the more general ISO 639-2) in order to avoid clutter in the plots.

Language	Dataset	Code
Arabic	Arabic-PDAT	ar
Bulgarian	Bulgarian-BTB	bg
Czech	Czech-CAC	cs
German	German-GSD	de
English	English-EWT	en
Spanish	Spanish-AnCora	es
Estonian	Estonian-EDT	et
Basque	Basque-BDT	eu
Persian	Persian-Seraji	fa
Finnish	Finnish-TDT	fi
French	French-GSD	fr
Hebrew	Hebrew-HTB	he
Hungarian	Hungarian-Szeged	hu
Italian	Italian-ISDT	it
Latvian	Latvian-LVTB	lv
Norwegian (Bokmaal)	Norwegian-Bokmaal	nb
Russian	Russian-SynTagRus	ru
Slovenian	Slovenian-SSJ	sl
Turkish	Turkish-IMST	tr
Urdu	Urdu-UDTB	ur

Table 7: Standard language datasets from the Universal Dependencies (v2.2) corpus.