

# Sequence-to-Sequence Lexical Normalization with Multilingual Transformers

Ana-Maria Bucur\*

University of Bucharest  
ana-maria.bucur@drd.unibuc.ro

Adrian Cosma\*

University Politehnica of Bucharest  
cosma.i.adrian@gmail.com

Liviu P. Dinu

University of Bucharest  
ldinu@fmi.unibuc.ro

## Abstract

Current benchmark tasks for natural language processing contain text that is qualitatively different from the text used in informal day to day digital communication. This discrepancy has led to severe performance degradation of state-of-the-art NLP models when fine-tuned on real-world data. One way to resolve this issue is through lexical normalization, which is the process of transforming non-standard text, usually from social media, into a more standardized form. In this work, we propose a sentence-level sequence-to-sequence model based on mBART, which frames the problem as a machine translation problem. As the noisy text is a pervasive problem across languages, not just English, we leverage the multilingual pre-training of mBART to fine-tune it to our data. While current approaches mainly operate at the word or subword level, we argue that this approach is straightforward from a technical standpoint and builds upon existing pre-trained transformer networks. Our results show that while word-level, intrinsic, performance evaluation is behind other methods, our model improves performance on extrinsic, downstream tasks through normalization compared to models operating on raw, unprocessed, social media text.

## 1 Introduction

Social media is a pervasive part of our modern lives and provides us with a rich source of information and insight into human behaviour. User-generated content has been a valuable resource for the research community, especially in the form of text, but it is notoriously noisy and non-standard. Models that operate on social media posts go beyond marketing and advertisement applications, and have the potential to impact real human lives through, for instance, detecting loneliness (Guntuku et al., 2019), stress (Winata et al., 2018), life

satisfaction (Yang and Srinivasan, 2016), suicidal ideation (Matero et al., 2019; Cao et al., 2019), and mental health problems such as depression (Yates et al., 2017; Bucur et al., 2021a; Tadesse et al., 2019) and PTSD (Coppersmith et al., 2014; Amir et al., 2019).

Outside of a formal setting, users communicate freely in text form, resorting to abbreviations, slang or plain spelling mistakes or typos. Eisenstein (2013) further explored *bad language* on social media, in the sense of language that defies our expectation of *good* spelling, vocabulary and syntax. He identified several underlying factors for the cause of non-standard text: user illiteracy, length limits imposed by social media sites (i.e. Twitter), text input affordances (i.e. standard mobile keyboards or predictive entry), pragmatics (emoticons/emoji, abbreviations and expressive lengthening), and a social component. Nguyen et al. (2021) further explored the latter, concluding that some types of non-standard text have strong social meaning, and normalization could induce a loss of meaning.

However, it is well known that for most benchmark tasks, noisy/non-standard text has proven to be a real problem to NLP models, such as BERT (Kumar et al., 2020), trained on clean or curated data, but fine-tuned on tasks with noisy and inconsistent format.

To overcome this predicament, Eisenstein (2013) proposes two possible approaches: either *domain adaptation* or *normalization*. While domain adaptation is not specific to natural language processing, text normalization and cleaning have always been a central part of any modern text processing pipeline. Text normalization is the process of adapting an input text to a more standard form. It has proven to be effective in increasing performance on tasks such as POS tagging (van der Goot and Çetinoğlu, 2021), dependency parsing (van der Goot, 2019a) and sentiment analysis (Mandal and Nanmaran, 2018). Naturally, most text normalization pipelines

\* Equal contribution

are based on supervised models, which require carefully annotated data. However, annotating a large corpus of text in multiple languages is often cumbersome and expensive, and some approaches rely on synthetically generating corrupted text (Dekker and van der Goot, 2020; Ma, 2019).

Commonly, approaches are based on word-level normalization. One of the most prominent methods is MoNoise (van der Goot and van Noord, 2017), in which the text correction pipeline is similar to a classic ranked retrieval. However, MoNoise operates at the individual word and uses a spelling correction module and a word embedding module. While word embeddings can be made to account for a specific sentence context, it is mostly discarded.

Different from current methods, we aim to perform text normalization at a sentence level. This approach has several advantages, compared to word or subword methods: i) it can be naturally framed as a sequence-to-sequence type problem, ii) it is more straightforward, as it requires only one module, as opposed to a multi-stage pipeline (i.e. complex candidate generation and ranking), and iii) the same model can be trained on multiple languages at the same time, without increasing in size and computational processing.

In this edition of The Workshop on Noisy User-generated Text (W-NUT), organizers propose the shared task of multilingual lexical normalization<sup>1</sup>, in which participants are required to perform lexical normalization on 12 different languages (van der Goot et al., 2021a).

As such, we use the state-of-the-art multilingual sequence-to-sequence transformer model mBART (Tang et al., 2020) and fine-tune it for our task. mBART is one of the first models that can be fine-tuned simultaneously on multiple languages without performance loss. We show that framing text normalization as a neural machine translation problem is a viable method for text normalization, improving performance on extrinsic, downstream tasks compared to models that operate on raw, unprocessed social media text. We made the code publicly available on github.<sup>2</sup>

## 2 Related Work

The W-NUT workshop hosted a shared task on lexical normalization of user-generated content from

<sup>1</sup><http://noisy-text.github.io/2021/multi-lexnorm.html>

<sup>2</sup><https://github.com/bucuram/seq2seq-multilingual-normalization>

English tweets in its first edition (Baldwin et al., 2015a). The task received from the competing teams two categories of submissions, from constrained (using only the training data provided by the organizers) and unconstrained systems (using other publicly available data or tools).

The best model, from Jin (2015), generated candidates from the most similar canonical forms from the training data evaluated with the Jaccard Index. A random forest classifier was used to predict the suitable canonical form from all the candidates using features such as support and confidence, string similarity, and part of speech tags. The model was a constrained system, suggesting that the quality of the proposed model is more important than using additional data and tools. Other approaches were based on conditional random fields (CRF) (Akhtar et al., 2015; Supranovich and Patsepnia, 2015; Akhtar et al., 2015) and recurrent neural networks (RNN) (Min and Mott, 2015; Wagner and Foster, 2015) among others.

Notably, MoNoise (van der Goot and van Noord, 2017) has long been considered state-of-the-art in lexical normalization. MoNoise is a normalization model using spelling correction and word embeddings for candidate generation and a feature-based random forest classifier for candidate ranking. It is a modular normalization system easily reusable and adaptable (van der Goot and van Noord, 2017). The model was at the beginning developed only for English text. Still, then it was later expanded for multi-lingual lexical normalization covering languages such as Dutch, Spanish, Turkish, Slovenian, Croatian and Serbian (van der Goot, 2019b).

The lexical normalization task can also be formulated as a machine translation (MT) task. The noisy user-generated content is the source language, and the canonical form is the target language. Veliz et al. (2019) compare the MT approaches for lexical normalization, focusing on statistical neural translation (SMT) and neural machine translation (NMT) and obtaining better results using the SMT method. Furthermore, the authors show that the SMT approach works better in a low-resource setting than an NMT approach which requires a lot of data.

With the rise in popularity of pre-trained language models for natural language understanding and natural language generation, their ability to perform lexical normalization was also studied. By transforming the task into a token prediction

one, Muller et al. (2019) demonstrate that a BERT model can be used as a lexical normalization model in low resource settings.

Current methods for lexical normalization attempt to normalize at the character-level (Pennell and Liu, 2011; Ljubešić et al., 2014), syllable-level (Xu et al., 2015), word-level (van der Goot, 2019b; Jin, 2015) or sentence-level (Muller et al., 2019; Lourentzou et al., 2019). Luseti et al. (2018) propose an encoder-decoder approach for text normalization.

We propose to make use of the latest transformer models that are capable of multilingual translation in a sequence to sequence manner, namely mBART (Tang et al., 2020). However, we do not perform translation between languages, but instead, we use mBART as a denoising autoencoder, i.e. translating from *bad English* to *good English*. This way, we take the whole sentence into consideration when correcting the text. Moreover, this method is more straightforward and can scale to multiple languages without increasing computational demands.

### 3 Data and Evaluation

We further describe the dataset for this task and evaluation procedures.

**MultiLexNorm Dataset** The data provided by the organizers includes texts from 12 languages: Croatian, Danish, Dutch, English, German, Italian, Serbian, Slovenian, Spanish, Turkish and code-switched data for Indonesian-English and Turkish-German, as seen in Table 1. Some examples from the training data are shown in Table 2. For some languages in the dataset, the capitalization (Caps column) is also corrected, and words are split or merged (1-N/N-1 column). The dataset comprises Twitter posts from all languages, but some languages also have texts from additional sources. For example, Danish also has texts from Arto, Denmark’s first large-scale social media (Plank et al., 2020) and Dutch texts were also gathered from public Internet forums, and SMS messages (Schuur, 2020).

**W-NUT Evaluation Methodology** The organizers of the W-NUT workshop propose two types of evaluation procedures: *intrinsic*, word-level and *extrinsic*, downstream task performance (i.e. dependency parsing).

As *intrinsic* evaluation, the Error Reduction Rate

Language	Words	1-N/N-1	Caps	% normed
Croatian	75,276	-	+	8.98
Danish	11,816	+	+	8.66
Dutch	23,053	+	+	26.49
English	73,806	+	-	6.90
German	25,157	+	+	8.90
Indonesian-English	23,124	+	-	12.16
Italian	14,641	+	+	7.36
Serbian	91,738	-	+	7.73
Slovenian	75,276	-	+	15.66
Spanish	13,827	-	-	7.69
Turkish	7,949	-	+	36.60
Turkish-German	16,546	+	+	24.25

Table 1: Available languages in the training set. Each language has its own annotation guidelines, in which capitalization can be taken into account (Caps), or words can be split or merged (1-N/N-1). Moreover, some languages are code-switched, two different languages are used in a tweet.

introduced by van der Goot (2019b) is proposed:

$$ERR = \frac{TP - FP}{TP + FN} \quad (1)$$

Because accuracy is hard to compare across datasets with different numbers of raw words which have to be normalized, the ERR proposes an evaluation metric that can be used to compare the performance of systems across multiple datasets. It is computed as accuracy normalized for the number of raw words normalized in the gold dataset.

A system that always keeps the raw words has an ERR score of 0.0, while a perfect system will have ERR precisely 1.0. The ERR has a negative value when the system normalizes more words with a wrong form than the correct canonical form.

However, one downside of the ERR is that it fails to distinguish between FP and FN. Thus, in the case of FP, the system may provide a correct normalization, even if the annotators did not normalize the raw word.

Further, two baselines are provided: **Leave-As-Is (LAI)** - the output is the same as the raw input, the normalization is not performed - and **Most-frequent-Replacement (MFR)** - the output is the most frequent replacement from the training data. If the raw word is not found in the training set, no normalization is performed.

As a secondary evaluation, the organizers propose an *extrinsic* evaluation of the effect of normalization on the task of dependency parsing, previous research showing that lexical normalization improves the performance for this task (van der Goot, 2019a). A dependency parser is trained on

Language	Example raw	Example gold
Croatian (Ljubešić et al., 2017a)	dok je <b>bandić</b> bio <b>ćlan sdpa</b> tvrdilo se da je idealan	dok je <b>bandić</b> bio <b>ćlan sdpa</b> tvrdilo se da je idealan
Danish (Plank et al., 2020)	<b>Maerkeligt, tænker</b> jeg, og <b>går</b> ind igen.	<b>Mærkeligt, tænker</b> jeg, og <b>går</b> ind igen.
Dutch (Schoor, 2020)	<b>ja effe</b> slaapverhaal <b>vertelle</b> vo sophieke eh lol	<b>Ja even</b> slaapverhaal <b>vertellen</b> voor sophieke eh lol
English (Baldwin et al., 2015b)	he <b>obvi doesnt</b> understand that	he <b>obviously doesnt</b> 't understand that
German (Sidarenka et al., 2013)	Ich <b>werd dran</b> denken!	Ich <b>werde daran</b> denken!
Indonesian-English (Barik et al., 2019)	<b>msh bs</b> disebut sukses?	<b>masih bisa</b> disebut sukses?
Italian (van der Goot et al., 2020)	<b>zrate</b> prentento in <b>ciro kvelli</b> quelli kol <b>raffrettores</b>	<b>state</b> prendendo in <b>giro quelli</b> col <b>raffreddore</b>
Serbian (Ljubešić et al., 2017b)	ja sam ozbiljan <b>čovek</b>	ja sam ozbiljan <b>čovek</b>
Slovenian (Erjavec et al., 2017)	da se <b>naujo</b> zdaj še na planico <b>spravli!</b> ?	da se <b>ne</b> zdaj še na planico <b>spravili!</b> ?
Spanish (Alegria et al., 2013)	quiero <b>tranquileo</b> del bueno hoy..!!!	quiero <b>tranquilidad</b> del bueno hoy..!!!
Turkish (Çolakoğlu et al., 2019)	Avrupa ve <b>amerikada VALENTİNA DAY</b> diye geçer.	Avrupa ve <b>Amerika'da Valentina Day</b> diye geçer.
Turkish-German (van der Goot and Çetinoğlu, 2021)	<b>artık</b> ablamdan <b>bise yuruturum napim</b> :D	<b>Artık</b> ablamdan <b>bir şey yürütürüm ne yapayım</b> :D

Table 2: Noisy examples from each language and the corresponding canonical forms.

both raw and canonical data to evaluate the performance improvement of using the normalized versus the original data.

Moreover, we also evaluate the extrinsic performance of our model on two additional tasks: sentiment analysis on the SMILE dataset (Wang et al., 2016) and hate speech detection on OLID dataset (Zampieri et al., 2019a). Both datasets contain data collected from Twitter, making them good candidates for evaluating the semantic processing of noisy text.

**SMILE dataset** It consists of posts with mentions of several British museums gathered from Twitter to classify the emotions expressed by users towards art and cultural experiences from the museums. It contains 3,085 posts annotated with five emotions: anger, disgust, happiness, surprise and sadness; fear was not found in any Twitter posts.

**OLID dataset** It was the official dataset of the SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval 2019) (Zampieri et al., 2019b) and SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media (OffensEval 2020) (Zampieri et al., 2020). The dataset was also used in misogyny (Pamungkas et al., 2020), cyberbullying (Aind et al., 2020) and depression (Bucur et al., 2021b) research. It contains 14,100 tweets with a hierarchical annotation taxonomy with three levels: Level A - Offensive language identification (offensive vs non-offensive), Level B - categorization of Offensive language (targeted insults or threats vs untargeted profanity) and Level C - Offensive language target identification (individual vs group vs other). However, for our evaluation, we focus only on level A.

For evaluating on sentiment analysis (SMILE) and offensive language identification (OLID), we trained a simple word-level TF-IDF model together with a linear SVM with balanced weights. For

SMILE, we report average macro  $F_1$  score across 5 folds, and for OLID, we report macro  $F_1$  score on the test set.

## 4 Method

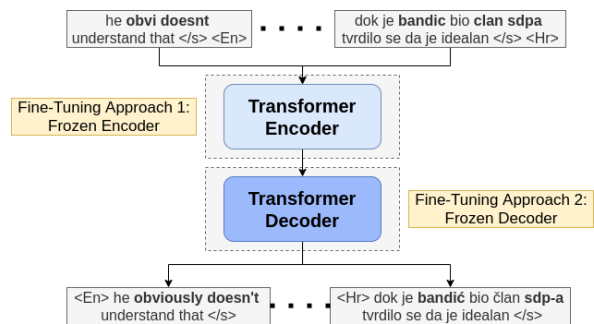


Figure 1: Fine-tuning a mBART model for lexical normalization on all available languages. We use the same model for all languages simultaneously.

Lewis et al. (2019) proposed BART in 2019, as a way to pre-train large-scale transformers for sequence-to-sequence tasks. Initially, the authors pre-trained an encoder-decoder transformer only for English, obtaining good results on multiple downstream NLP tasks. Further, mBART (Tang et al., 2020), follows the same procedure, but for multiple languages. The pre-training stage for both BART and mBART is akin to a denoising autoencoder, in which the model receives a noisy (in this case masked) sentence, and it learns to reconstruct it.

While mBART is fine-tuned on multiple language pairs, it is pretrained monolingually, and is capable of acting as an autoencoder for the same language. In our case, we make use of a pretrained mBART on 50 languages<sup>3</sup> from the transformers library (Wolf et al., 2020), and employ a procedure similar to the pre-training stage: a noisy sentence is

<sup>3</sup><https://huggingface.co/facebook/mbart-large-50>

Team Name	Avg.	da	de	en	es	hr	iden	it	nl	sl	sr	tr	trde
UFAL-2 (Samuel and Straka, 2021)	<b>67.30</b>	68.67	<b>66.22</b>	<b>75.60</b>	59.25	<b>67.74</b>	<b>67.18</b>	<b>47.52</b>	<b>63.58</b>	<b>80.07</b>	<b>74.59</b>	<b>68.58</b>	<b>68.62</b>
HEL-LJU-2 (Scherrer and Ljubešić, 2021)	53.58	56.65	59.80	62.05	35.55	56.24	55.33	35.64	45.88	66.97	66.44	51.18	51.18
MoNoise (van der Goot, 2019b)	49.02	51.27	46.96	74.35	45.53	52.63	59.79	21.78	49.53	61.91	59.58	28.21	36.72
TrinkaAI-2* (Kubal and Nagvenkar, 2021)	43.75	45.89	47.30	65.96	<b>61.33</b>	41.28	56.36	15.84	45.74	59.51	44.52	15.54	25.77
thunderml-1*	43.44	46.52	46.62	64.07	60.29	40.09	59.11	11.88	44.05	59.33	44.46	15.88	29.01
team-2	40.70	48.10	46.06	63.73	21.00	40.39	59.28	13.86	43.72	60.55	46.11	15.88	29.71
learnML-2	40.30	40.51	43.69	61.57	56.55	38.11	56.19	5.94	42.77	58.25	39.99	14.36	25.68
maet-1	40.05	48.10	46.06	63.90	21.00	40.39	59.28	5.94	43.72	60.55	46.11	15.88	29.71
MFR	38.37	49.68	32.09	64.93	25.57	36.52	61.17	16.83	37.70	56.71	42.62	14.53	22.09
CL-MoNoise (van der Goot, 2021)	12.05	7.28	16.55	4.13	4.99	26.41	2.41	0.00	16.22	8.77	20.09	17.57	20.16
(ours) Fixed Encoder (separate) + post proc.	10.65	<i>49.68</i>	-2.59	29.13	-7.90	26.41	-1.72	-8.91	-1.49	1.27	<i>42.62</i>	0.68	0.70
(ours) Fixed Encoder (separate)	6.73	<i>49.68</i>	-1.91	26.81	-9.36	-10.06	-7.22	-8.91	-2.09	-1.04	<i>42.62</i>	9.97	14.99
(ours) Fixed Encoder (separate) - stripped unicode	5.22	<i>49.68</i>	-1.91	26.81	-10.19	-9.86	-7.22	-31.68	-2.09	-1.13	<i>42.62</i>	1.01	6.57
(ours) ML - Fixed Decoder + post proc.	-6.54	<i>49.68</i>	12.50	27.41	-13.10	-111.84	-7.73	-8.91	16.82	-110.57	<i>42.62</i>	11.49	13.06
(ours) ML - Fixed Decoder	-11.79	<i>49.68</i>	20.05	22.12	-18.92	-127.60	-14.60	-25.74	16.69	-133.71	<i>42.62</i>	13.18	14.64
(ours) ML - Fixed Encoder + post proc.	-21.51	<i>49.68</i>	10.47	12.09	-28.69	-191.33	-9.97	-27.72	9.19	-141.80	<i>42.62</i>	9.63	7.62
(ours) ML - Fixed Encoder	-32.90	<i>49.68</i>	19.48	20.78	-40.12	-242.57	-24.23	-70.30	8.72	-180.75	<i>42.62</i>	11.15	10.69
LAI	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
MaChAmp (van der Goot et al., 2021b)	-21.25	-88.92	-93.36	50.99	25.36	42.62	39.52	-312.87	1.49	56.80	39.44	-12.67	-3.42

Table 3: Team standings, based on Error Reduction Rate (ERR). We kept the best result from each team, from clarity. (\* denotes late submissions).

fed to the model, and the output ought to be the normalized sentence. We only trained on the provided languages that are contained in the pre-training set of languages. As mBART is not pre-trained on code-switched languages, for IN-EN and TR-DE, we use the mBART model pre-trained only on the main language of each pair (e.g. IN for IN-EN and TR for TR-DE). For Danish and Serbian, we fall back to the MFR baseline.

Figure 1 showcases our fine-tuning procedure. We tried two different approaches for fine-tuning: Frozen Encoder and Frozen Decoder because, with a fixed encoder, the model suffers from the same OOV-type problems as a typical transformer. However, training with a fixed decoder allows the model to better adapt its representations to each language’s noisy version while maintaining its generative properties. For both approaches, we train a single model for all languages. Moreover, we also trained a separate mBART for each language, monolingually.

**Training details** For all runs, we fine-tune mBART for 50 epochs, using a batch size of 256 and with a cyclical learning rate scheduler (Smith, 2015) that linearly increases the learning rate from 0.00001 to 0.0001 and back across 5 epochs. The workshop organizers provided both the training data and the validation data on most languages. We omit validation on languages where the validation data is missing. The training was performed on an NVIDIA RTX 2070 graphics card. Since the memory requirements of an mBART model are quite high, we employed gradient accumulation to increase the batch size. In addition, we employed early stopping when the validation loss increased

for more than 3 epochs.

**Post-processing** Since our model outputs a whole sentence directly, the word-level evaluation requires the noisy input words to be aligned to their normalized counterpart. This phase is essential for sequence-to-sequence text normalization, as bad alignments will reduce the overall word-level performance score, especially in the 1-N/N-1 languages. As such, for the post-processing phase, we aligned input words with their normalized counterparts based on the Levenshtein distance between them. We used a linear sum assignment on the distance matrix to perform matching. Additionally, we matched the capitalization between corrections and left links, hashtags, and user mentions as they are.

## 5 Results

We further showcase the results of the pretrained mBART models fine-tuned on the available data: firstly, we kept the transformer encoder fixed and trained only the decoder, and secondly, we kept the decoder fixed and trained the encoder. During this fine-tuning process, we trained a single model for all languages. Further, for the CodaLab submission, we fine-tuned multiple models, one for each language, in the "fixed encoder" regime.

**Intrinsic Evaluation** Table 3 showcases intrinsic, word-level evaluation across languages. Our best model obtained an average ERR across languages of 10.65, corresponding to a separate mBART trained for each language, with the additional post-processing described in Section 4. In our case, training multilingually did improve performance on some languages (i.e. *DE*, *EN*, *NL*, *TR*), but over-

Team Name	Avg.	de-tweede	en-aae	en-monoise	en-tweebank2	it-postwita	it-twittiro	tr-iwt151
ÚFAL-2	<b>64.17</b>	<b>73.58</b>	<b>62.73</b>	<b>58.57</b>	59.08	<b>68.28</b>	72.22	54.74
HEL-LJU-2	63.73	73.49	60.64	56.27	60.30	68.11	72.32	54.95
MoNoise	63.44	73.20	62.27	56.83	58.90	67.55	70.69	54.61
MFR	63.31	72.86	60.32	56.74	<b>60.31</b>	67.34	70.72	54.89
TrinkaAI-2*	63.12	72.86	60.16	56.64	59.87	66.98	71.14	54.20
maet-1	63.09	72.80	59.44	56.64	59.80	67.41	71.07	54.45
team-2	63.03	72.80	59.44	56.64	59.80	67.19	70.86	54.45
thunderml-2*	63.02	72.67	59.57	56.74	59.25	67.34	71.35	54.24
thunderml-1*	62.95	72.52	59.31	56.74	59.86	67.09	71.00	54.09
CL-MoNoise	62.71	72.65	60.90	55.26	58.53	66.53	70.10	54.98
(ours) Fixed Encoder (separate)	62.53	72.57	59.57	54.20	59.81	66.74	69.99	54.84
LAI	62.45	72.71	59.21	53.65	59.99	66.49	70.06	<b>55.00</b>
MaChAmp	61.89	71.28	60.77	54.61	57.97	64.65	69.82	54.08

Table 4: Extrinsic evaluation results on dependency parsing task.

	SMILE	OLID Task A
Raw Text (Leave-As-Is)	22.65% $\pm$ 0.02	57.15%
mBART Fixed Encoder	<b>23.43% <math>\pm</math> 0.02</b>	<b>58.08%</b>

Table 5: Extrinsic evaluation on sentiment analysis (SMILE) and offensive language identification (OLID). Lexical normalization through fine-tuning mBART slightly improves performance.

Raw	Gold	Our	Correct?
i see, u can comeee	i see, you can come	i see, you can come	✓
ich geb heut einen aus	ich gebe heute einen aus	ich gebe heute einen aus	✓
Juhuuuuu	Juhu	Juhu	✓
fakt ap gaan. eig nu al mr kanniet	fakt op gaan. Eigenlijk nu al maar kan niet	echt ap gaan. eig nu al <b>mijn kanniet</b>	✗
"Why Germany says "nein"	"Why Germany says "nein"	" <b>Warum Deutschland sagt "nein"</b>	✗
i coulda swore .... lol nvm	i could have swore .... lol never mind	i <b>could</b> swore .... lol <b>never</b>	✗
todos lo sabemos jajajaja-jaja	todos lo sabemos jajajaja-jaja	todos lo sabemos ja	?
discussing w/ friend	discussing w/ friend	discussing <b>with</b> friend	?
n ni volver a dormirme	n ni volver a dormirme	<b>no</b> ni volver a dormirme	?

Table 6: Qualitative results on different languages with mBART Fixed Encoder. We present examples of correct normalization (✓), mistakes (✗), and questionable normalizations (?), in which the model correctly normalizes, but annotators do not.

all achieved lower scores, especially in the case for *HR* and *SL*. For those languages, the model severely diverged, and its output consisted only of repeating the first word in the sentence. As per our intuition, fixing the decoder results in better perfor-

mance when compared to fixing the encoder: the model learns to adapt its representations to account for the noisy text.

However, since our method is sentence-based, perfect alignment between words is cumbersome, with many cases of mismatch between punctuation. Moreover, merging or splitting words for normalization is also not taken into account in the post-processing phase.

**Extrinsic Evaluation** For extrinsic evaluation, we showcase the results for our best model in Table 4 on the dependency parsing downstream task from the workshop challenge. Even though our model is not in the top-performing models, the absolute difference in performance is minimal.

Moreover, we also evaluated the effect of lexical normalization on two other tasks - sentiment analysis on the SMILE dataset and offensive language identification on OLID (Table 5). We trained a word-level TF-IDF and a linear SVM with balanced weights for both datasets and reported a macro  $F_1$  score. Our lexical normalization improves results on both these tasks, compared to modelling the raw, unprocessed social media posts. This is because lexical normalization results in a smaller vocabulary for the documents, allowing the SVM model to operate on smaller dimensional data. Moreover, this evaluation procedure is arguably more realistic, as it does not require accurate post-processing to precisely align noisy words with their corrected version and match punctuation.

**Discussion** In Table 6 we showcase some examples for correct, incorrect and questionable text normalizations. The model is able to easily grasp contractions such as  $u \rightarrow you$  and expressive lengthening such as  $Juhuuuuuu \rightarrow Juhu$ . However, more complex word abbreviations such as *nvm* are quite challenging to generate, as the model only outputs

the first part (i.e. *never*). Moreover, code-switched languages are an inherent problem to our approach, as mBART is only trained to receive input from a single language and not code-switched. Interestingly, even though we specified the language code for German in the phrase "*Why Germany says 'nein'*", the model actually translates the English part into German: *Warum Deutschland sagt 'nein'*.

However, as the organizers have pointed out, there are inconsistencies in the training and testing data annotations. In some cases, some words are not normalized (i.e. *jajajajaj / w / n*) even though they were clearly lengthenings or contractions. Despite this, in some of these cases, our model was able to provide correct normalizations.

There also appears to be no correlation between training dataset size and final normalization performance. For example, in the case of Croatian, even though the dataset is the second largest, the performance is lower than for other languages. Thus, the lower performance in some languages may be a cause of the complexity of the language; for English, our model obtained the best results.

## 6 Conclusions

In this work, we presented a method to perform lexical normalization by fine-tuning a multilingual machine translation model on pairs of noisy and normalized sentences from various languages. We employed mBART, as it is currently the state-of-the-art in transformer-based multilingual machine translation, allowing us to fine-tune on all available languages simultaneously. Furthermore, we used mBART as a denoising autoencoder and tuned it in a supervised fashion.

As opposed to current two-stage methods for word candidate generation and ranking, our approach is more straightforward. Moreover, it scales to multiple languages without increasing computational demand (i.e. not increasing vocabulary size, increasing search space and others). Evaluation results show that our method, even though it lacks behind current methods on intrinsic, word-level evaluation, improves performance on downstream tasks.

For future work, we aim to develop our method for better post-processing of the output and increasing augmentation levels - i.e. injecting more noise in the form of spelling mistakes, backwards translations etc. Moreover, since our method is supervised, the quality and quantity of training data play

an essential role in the final performance. In this regard, we aim to explore ways to take into account inconsistent annotations.

## Acknowledgments

We would like to thank the reviewers for the insightful feedback provided. This research was partially supported by Blog Alchemy Limited.

## References

- Alwin T Aind, Akashdeep Ramnaney, and Divyashikha Sethia. 2020. Q-bully: a reinforcement learning based cyberbullying detection framework. In *2020 International Conference for Emerging Technology (INCET)*, pages 1–6. IEEE.
- Md Shad Akhtar, Utpal Kumar Sikdar, and Asif Ekbal. 2015. Iitp: Hybrid approach for text normalization in twitter. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 106–110.
- Inaki Alegria, Nora Aranberri, Víctor Fresno, Pablo Gamallo, Lluís Padró, Inaki San Vicente, Jordi Turmo, and Arkaitz Zubiaga. 2013. Introducción a la tarea compartida Tweet-Norm 2013: Normalización léxica de tuits en Español. In *Tweet-Norm@SEPLN*, pages 1–9.
- Silvio Amir, Mark Dredze, and John W Ayers. 2019. Mental health surveillance over social media with digital cohorts. In *Proceedings of the Sixth Workshop on Computational Linguistics and Clinical Psychology*, pages 114–120.
- Timothy Baldwin, Marie-Catherine de Marneffe, Bo Han, Young-Bum Kim, Alan Ritter, and Wei Xu. 2015a. Shared tasks of the 2015 workshop on noisy user-generated text: Twitter lexical normalization and named entity recognition. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 126–135.
- Timothy Baldwin, Marie Catherine de Marneffe, Bo Han, Young-Bum Kim, Alan Ritter, and Wei Xu. 2015b. [Shared tasks of the 2015 workshop on noisy user-generated text: Twitter lexical normalization and named entity recognition](#). In *Proceedings of the Workshop on Noisy User-generated Text*, pages 126–135, Beijing, China. Association for Computational Linguistics.
- Anab Maulana Barik, Rahmad Mahendra, and Mirna Adriani. 2019. [Normalization of Indonesian-English code-mixed Twitter data](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 417–424, Hong Kong, China. Association for Computational Linguistics.
- Ana-Maria Bucur, Adrian Cosma, and Liviu P Dinu. 2021a. Early risk detection of pathological gambling, self-harm and depression using bert. *CLEF (Working Notes)*.

- Ana-Maria Bucur, Marcos Zampieri, and Liviu P. Dinu. 2021b. [An exploratory analysis of the relation between offensive language and mental health](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3600–3606, Online. Association for Computational Linguistics.
- Lei Cao, Huijun Zhang, Ling Feng, Zihan Wei, Xin Wang, Ningyun Li, and Xiaohao He. 2019. Latent suicide risk detection on microblog via suicide-oriented word embeddings and layered attention. 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)*, pages 1718–1728.
- Talha Çolakoğlu, Umut Sulubacak, and Ahmet Cüneyd Tantuş. 2019. [Normalizing non-canonical Turkish texts using machine translation approaches](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 267–272, Florence, Italy. Association for Computational Linguistics.
- Glen Coppersmith, Craig Harman, and Mark Dredze. 2014. Measuring post traumatic stress disorder in twitter. In *Eighth international AAAI conference on weblogs and social media*.
- Kelly Dekker and Rob van der Goot. 2020. Synthetic data for english lexical normalization: How close can we get to manually annotated data? In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6300–6309.
- Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proceedings of the 2013 conference of the North American Chapter of the association for computational linguistics: Human language technologies*, pages 359–369.
- Tomaž Erjavec, Darja Fišer, Jaka Čibej, Špela Arhar Holdt, Nikola Ljubešić, and Katja Zupan. 2017. [CMC training corpus Janes-Tag 2.0](#). Slovenian language resource repository CLARIN.SI.
- Sharath Chandra Guntuku, Rachele Schneider, Arthur Pelullo, Jami Young, Vivien Wong, Lyle Ungar, Daniel Polsky, Kevin G Volpp, and Raina Merchant. 2019. Studying expressions of loneliness in individuals using twitter: an observational study. *BMJ open*, 9(11):e030355.
- Ning Jin. 2015. Ncsu-sas-ning: Candidate generation and feature engineering for supervised lexical normalization. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 87–92.
- Divesh Kubal and Apurva Nagvenkar. 2021. Multilingual sequence labeling approach to solve lexical normalization. In *Proceedings of the 7th Workshop on Noisy User-generated Text (W-NUT 2021)*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ankit Kumar, Piyush Makhija, and Anuj Gupta. 2020. [Noisy text data: Achilles’ heel of BERT](#). In *Proceedings of the Sixth Workshop on Noisy User-generated Text (W-NUT 2020)*, pages 16–21, Online. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Nikola Ljubešić, Tomaž Erjavec, and Darja Fišer. 2014. Standardizing tweets with character-level machine translation. In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 164–175. Springer.
- Nikola Ljubešić, Tomaž Erjavec, Maja Miličević, and Tanja Samardžić. 2017a. [Croatian Twitter training corpus ReLDI-NormTagNER-hr 2.0](#). Slovenian language resource repository CLARIN.SI.
- Nikola Ljubešić, Tomaž Erjavec, Maja Miličević, and Tanja Samardžić. 2017b. [Serbian Twitter training corpus ReLDI-NormTagNER-sr 2.0](#). Slovenian language resource repository CLARIN.SI.
- Ismeni Lourentzou, Kabir Manghnani, and Chengxiang Zhai. 2019. Adapting sequence to sequence models for text normalization in social media. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 13, pages 335–345.
- Massimo Lusetti, Tatyana Ruzsics, Anne Göhring, Tanja Samardzic, and Elisabeth Stark. 2018. Encoder-decoder methods for text normalization. In *Proceedings of the Fifth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2018)*, pages 18–28.
- Edward Ma. 2019. Nlp augmentation. <https://github.com/makcedward/nlpaug>.
- Soumil Mandal and Karthick Nanmaran. 2018. Normalization of transliterated words in code-mixed data using seq2seq model & levenshtein distance. In *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 49–53.
- Matthew Matero, Akash Idnani, Youngseo Son, Salvatore Giorgi, Huy Vu, Mohammad Zamani, Parth Limbachiya, Sharath Chandra Guntuku, and H Andrew Schwartz. 2019. Suicide risk assessment with multi-level dual-context language and bert. In *Proceedings of the sixth workshop on computational linguistics and clinical psychology*, pages 39–44.
- Wookhee Min and Bradford Mott. 2015. Ncsu\_sas\_wookhee: A deep contextual long-short term memory model for text normalization. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 111–119.



- Benjamin Muller, Benoît Sagot, and Djamé Seddah. 2019. Enhancing bert for lexical normalization. In *The 5th Workshop on Noisy User-generated Text (W-NUT)*.
- Dong Nguyen, Laura Rosseel, and Jack Grieve. 2021. On learning and representing social meaning in nlp: a sociolinguistic perspective. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 603–612.
- Endang Wahyu Pamungkas, Valerio Basile, and Viviana Patti. 2020. Misogyny detection in twitter: a multilingual and cross-domain study. *Information Processing & Management*, 57(6):102360.
- Deana Pennell and Yang Liu. 2011. A character-level machine translation approach for normalization of sms abbreviations. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 974–982.
- Barbara Plank, Kristian Nørgaard Jensen, and Rob van der Goot. 2020. [DaN+: Danish nested named entities and lexical normalization](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6649–6662, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- David Samuel and Milan Straka. 2021. ÚFAL at Multi-LexNorm 2021: Improving multilingual lexical normalization by fine-tuning ByT5. In *Proceedings of the 7th Workshop on Noisy User-generated Text (W-NUT 2021)*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yves Scherrer and Nikola Ljubešić. 2021. Sesame Street to Mount Sinai: BERT-constrained character-level Moses models for multilingual lexical normalization. In *Proceedings of the 7th Workshop on Noisy User-generated Text (W-NUT 2021)*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Youri Schuur. 2020. [Normalization for Dutch for improved pos tagging](#). Master’s thesis, University of Groningen.
- Uladimir Sidarenka, Tatjana Scheffler, and Manfred Stede. 2013. Rule-based normalization of German Twitter messages. In *Proc. of the GSCL Workshop Verarbeitung und Annotation von Sprachdaten aus Genres internetbasierter Kommunikation*.
- Leslie N. Smith. 2015. [No more pesky learning rate guessing games](#). *CoRR*, abs/1506.01186.
- Dmitry Supranovich and Viachaslau Patsepnia. 2015. Ihs\_rd: Lexical normalization for english tweets. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 78–81.
- Michael M Tadesse, Hongfei Lin, Bo Xu, and Liang Yang. 2019. Detection of depression-related posts in reddit social media forum. *IEEE Access*, 7:44883–44893.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Rob van der Goot. 2019a. An in-depth analysis of the effect of lexical normalization on the dependency parsing of social media. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 115–120.
- Rob van der Goot. 2019b. Monoise: A multi-lingual and easy-to-use lexical normalization tool. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 201–206.
- Rob van der Goot. 2021. CL-MoNoise: Cross-lingual lexical normalization. In *Proceedings of the 7th Workshop on Noisy User-generated Text (W-NUT 2021)*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rob van der Goot and Özlem Çetinoğlu. 2021. [Lexical normalization for code-switched data and its effect on POS tagging](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2352–2365, Online. Association for Computational Linguistics.
- Rob van der Goot, Alan Ramponi, Tommaso Caselli, Michele Cafagna, and Lorenzo De Mattei. 2020. [Norm it! lexical normalization for Italian and its downstream effects for dependency parsing](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6272–6278, Marseille, France. European Language Resources Association.
- Rob van der Goot, Alan Ramponi, Arkaitz Zubiaga, Barbara Plank, Benjamin Muller, Iñaki San Vicente Roncal, Nikola Ljubešić, Özlem Çetinoğlu, Rahmad Mahendra, Talha Çolakoğlu, Timothy Baldwin, Tommaso Caselli, and Wladimir Sidorenko. 2021a. MultiLexNorm: A shared task on multilingual lexical normalization. In *Proceedings of the 7th Workshop on Noisy User-generated Text (W-NUT 2021)*, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rob van der Goot, Ahmet Üstün, Alan Ramponi, Ibrahim Sharaf, and Barbara Plank. 2021b. [Massive choice, ample tasks \(MaChAmp\): A toolkit for multi-task learning in NLP](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 176–197, Online. Association for Computational Linguistics.

- Rob van der Goot and Gertjan van Noord. 2017. Monoise: Modeling noise using a modular normalization system. *Computational Linguistics in the Netherlands Journal*, 7:129–144.
- Claudia Matos Veliz, Orphée De Clercq, and Véronique Hoste. 2019. Comparing mt approaches for text normalization. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pages 740–749.
- Joachim Wagner and Jennifer Foster. 2015. Dcu-adapt: learning edit operations for microblog normalisation with the generalised perceptron. In *Proceedings of the Workshop on Noisy User-generated Text*, pages 93–98.
- Bo Wang, Maria Liakata, Arkaitz Zubiaga, Rob Procter, and Eric Jensen. 2016. Smile: Twitter emotion classification using domain. In *Proceedings of the 4th Workshop on Sentiment Analysis where AI meets Psychology co-located with 25th International Joint Conference on Artificial Intelligence*.
- Genta Indra Winata, Onno Pepijn Kampman, and Pascale Fung. 2018. Attention-based lstm for psychological stress detection from spoken language using distant supervision. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6204–6208. IEEE.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Ke Xu, Yunqing Xia, and Chin-Hui Lee. 2015. Tweet normalization with syllables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 920–928.
- Chao Yang and Padmini Srinivasan. 2016. Life satisfaction and the pursuit of happiness on twitter. *PloS one*, 11(3):e0150881.
- Andrew Yates, Arman Cohan, and Nazli Goharian. 2017. Depression and self-harm risk assessment in online forums. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2968–2978.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019a. [Predicting the type and target of offensive posts in social media](#). 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 1415–1420, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019b. SemEval-2019 Task 6: Identifying and Categorizing Offensive Language in Social Media (OffensEval). In *Proceedings of The 13th International Workshop on Semantic Evaluation (SemEval)*.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. [SemEval-2020 task 12: Multilingual offensive language identification in social media \(OffensEval 2020\)](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1425–1447, Barcelona (online). International Committee for Computational Linguistics.