



# Post-ocr text correction for Bulgarian historical documents

Angel Beshirov<sup>1</sup> · Milena Dobрева<sup>1</sup> · Dimitar Dimitrov<sup>1</sup> · Momchil Hardalov<sup>2</sup> · Ivan Koychev<sup>1</sup> · Preslav Nakov<sup>3</sup>

Received: 15 May 2024 / Revised: 19 September 2024 / Accepted: 19 January 2025  
© The Author(s) 2025

## Abstract

The digitization of historical documents is crucial for preserving the cultural heritage of the society. An essential step in this process is converting scanned images to text using Optical Character Recognition (OCR), which can enable further search, information extraction, etc. Unfortunately, this is a challenging problem as standard OCR tools are not tailored to deal with historical orthography or challenging layouts. Thus, it is standard to apply an additional text correction step on the OCR output when dealing with such documents. In this work, we focus on Bulgarian, and we create the first benchmark dataset for evaluating the OCR text correction for historical Bulgarian documents written in the first standardized Bulgarian orthography: the Drinov orthography from the 19th century. We further develop a method for automatically generating synthetic data in this orthography, as well as in the subsequent Ivanchev orthography, by leveraging vast amounts of contemporary literature Bulgarian texts. We then use state-of-the-art LLMs and encoder-decoder framework which we augment with diagonal attention loss and copy and coverage mechanisms to improve the post-OCR text correction. The proposed method reduces the errors introduced during the recognition. It improves the quality of the documents by 25%, which is an increase of 16% compared to the state-of-the-art on the ICDAR 2019 Bulgarian dataset. We release our data and code at <https://github.com/angelbeshirov/post-ocr-text-correction>.

**Keywords** Post-OCR text correction · Synthetic data · Orthographic variety · LLMs · Character-level sequence-to-sequence model

## 1 Introduction

In recent years, many libraries and museums have stepped towards datafication, trying to make their documents and

papers fully accessible, searchable and processable in digital form. Having this information in digital form helps preserve the cultural heritage of the society. The documents are digitized using Optical Character Recognition (OCR), which is the process of identifying typed or handwritten printed characters within a document and converting them automatically into computer codes, which can then be processed as textual data<sup>1</sup>.

However, for documents written in old Bulgarian spelling, the process often introduces errors in recognition because of the challenging layouts and orthographic variety in the documents due to the nine language reforms applied to the Bulgarian language before 1945. As a result of these errors, other applications like Named Entity Recognition, Part-of-speech tagging, text summarization and others are negatively impacted. According to [1], if the word error rate of the digitized text increases from 0% to 2.7%, the F1-score of the NER decreases by 3%. Nowadays, most modern OCR engines

---

✉ Angel Beshirov  
beshirov@uni-sofia.bg

Milena Dobрева  
milena.dobрева@gmail.com

Dimitar Dimitrov  
mitko.bg.ss@gmail.com

Momchil Hardalov  
momchilh@amazon.com

Ivan Koychev  
koychev@fmi.uni-sofia.bg

Preslav Nakov  
preslav.nakov@gmail.com

<sup>1</sup> FMI, Sofia University “St. Kliment Ohridski”, Sofia, Bulgaria

<sup>2</sup> AWS AI Labs, Barcelona, Spain

<sup>3</sup> Mohamed bin Zayed University of Artificial Intelligence, Abu Dhabi, UAE

<sup>1</sup> Work done while Momchil was at Sofia University before joining Amazon.

recognize around 99% of the characters in high-quality documents. With an average word length of five characters, this results in one error for every twenty words. For Bulgarian historical documents, the percentage of correctly recognized characters is lower as some are no longer used in the current spelling.

The large amount of transcribed data required to train an OCR model from scratch is unavailable for the resource-poor Bulgarian historical spellings. Instead, our focus is on post-correcting the produced output, and we show that using this approach can significantly improve the quality of the raw OCR'd text. Additionally, the characteristics of the Bulgarian language are similar to those of other Slavic languages, so the methods presented in this paper can easily be adapted for them as well.

We have created a benchmark dataset for the first standardized Bulgarian orthography: the Drinov orthography from the 19th century. For the subsequent Ivanchev orthography, we use the dataset provided by the ICDAR 2019 competition [2]. Our proposed method reduces the errors introduced during the recognition. It improves the quality of the documents by 25%, which is an increase of 16% compared to the best contender from the ICDAR 2019 on the Bulgarian dataset. Finally, due to scarce training data for these historical Bulgarian spellings, we developed a method that can convert modern Bulgarian spelling into a historical one, thus allowing us to leverage the large amount of contemporary Bulgarian literature texts. The method generated synthetically noised samples using the confusion matrix of the isolated misspellings from each dataset. We have released our data and code on GitHub.<sup>2</sup>

In summary, our contributions are the following:

1. A new benchmark dataset for post-OCR text correction on the first standardized Bulgarian orthography: the Drinov orthography from the 19th century. This is the first dataset that uses this orthography.
2. A method for synthetic data generation, which automatically transforms modern Bulgarian spelling into a resource-poor historical spelling using concrete spelling and grammatical norms.
3. A novel post-OCR text correction method that utilizes pre-trained LLMs and encoder-decoder framework augmented with diagonal attention loss, copy and coverage mechanisms.

## 2 Related work

Post-OCR text correction has been an important problem ever since the inception of OCR technology in the 1960s.

<sup>2</sup> <https://github.com/angelbeshirov/post-ocr-text-correction>.

According to [3], there are three main approaches to improve the quality of the OCR output: changing the input images, optimizing the OCR system and post-processing the output. In the current research, we focus on the last option.

The post-OCR text correction process consists of two interdependent tasks: error detection and error correction. Additionally, there are two types of errors that the OCR model could make: real-word and non-word errors [4]. If a token is not a lexicon entry, it is considered a non-word error, and if a valid word occurs in a wrong context, then it is a real-word error. Real-word errors are harder to detect and correct since they can only be detected by leveraging the context around the word. For example, consider the phrases “glowing brightly” and “growing bnghtly”, a non-word error is “bnghtly” and “growing” is a real-word error.

In the beginning, the first approaches were based on dictionaries [5], [6], [7]. The error detection task consists of checking if a word is present or not in a predefined dictionary. For the error correction, the most similar candidates from the dictionary are computed using similarity measures like Levenshtein distance [8], longest common sub-sequence or confusion weight. The International Conference on Document Analysis and Recognition (ICDAR) organized two competitions in 2017 and 2019 for correcting mistakes in digitized documents [9] [2]. In the first competition, the best-performing model [10] was based on machine translation. They use an ensemble of statistical and neural machine translation models on the character level. In the second competition in 2019, the winner was the Clova AI team. They fine-tuned a pre-trained BERT model [11] with convolutional and fully connected layers to identify OCR errors. Afterwards, a character-level sequence-to-sequence model was used to correct the erroneous tokens.

Since then, the interest in applying pre-trained language models for post-OCR text correction has increased. Embeddings from BERT are used in [12] for detecting errors, after which a method based on machine translation corrects them. Another approach is in [13] where again BERT embeddings are used to detect errors after which a BERT model is fine-tuned on MLM to correct these mistakes. In [14] a sequence-to-sequence correction model was trained to correct a synthetically corrupted corpus, constructed by injecting uniformly distributed insertion, deletion and substitution errors into the dataset. The use of synthetic dataset led to an overall improvement. The above studies clearly demonstrate the potential of the pre-trained language models for post-OCR text correction.

## 3 Methodology

In this section, we define the problem and the pipeline of post-OCR text correction. Afterwards, we describe the error

detection and correction in details with the models which we have used to tackle both of these problems.

### 3.1 Problem formulation

The task of post-OCR text correction is defined formally as follows: given a sequence of characters  $C = c_1c_2c_3 \dots c_n$ , the goal is to find the correct sequence of characters  $W = w_1w_2w_3 \dots w_m$ , which matches with the original printed text. Note that it is possible  $n \neq m$  for an error as some symbols might have been missed or a word ending might not have been recognized correctly by the OCR engine. If  $C = W$ , then we will consider the word to be correct, otherwise it is incorrect.

There are two types of errors that the OCR engine can make - real-word and non-word errors. If  $C \notin D$  where  $D$  is some dictionary we will say that the error is of type non-word error. If  $C \in D$  then we will say that the error is a real-word error.

The pipeline for the post-OCR text correction process is shown in Fig. 1. The first step in the pipeline is preprocessing, where the data is tokenized, cleaned and prepared to be fed to the next error detection step. The error detection classifies each token as correct or incorrect. Afterwards, the classified as erroneous tokens are fed into the error corrector model which either directly corrects them or suggests a list of candidates if the text correction is executed in a semi-automatic setting. Splitting the task to error detection and correction adds the benefit of increasing the proportion of OCR errors which are fed into the corrector model and decreases the amount of erroneously changed correct tokens since the tokens classified as correct are not passed through the error corrector model.

### 3.2 Error detection

The goal of the error detection task is to determine whether a token was correctly recognized by the OCR engine. In other words, it is a binary classification. In the text correction pipeline shown in Fig. 1, only the erroneous tokens will be passed downstream to the error correction task. This is why we want to minimise the incorrectly classified tokens as erroneous, i.e. false positives. Therefore, for the error detection task precision is preferred over recall.

We have used two naive baselines for the error detection task. The first one is based on the CLaDA-BG dictionary [15] which contains over 1.1 million unique words in all of their correct forms. During inference, if the token is present in the dictionary, then it is considered to be correct, otherwise, it is classified as incorrect. The second baseline uses fastText [16], which is a library for the classification and learning of word representations. As the model is trained on subword n-grams, it can handle out-of-vocabulary words, unlike the

CLaDA-BG dictionary. The data is transformed into the format `__label__ < 0|1 > < word >`, where 0 is the label for correct words and 1 for incorrect.

Pre-trained language models are very effective and achieve state-of-the-art results in solving different NLP tasks, including token classification. We experimented with various language models for solving the task in a context-dependent manner. The best-performing model is the multilingual DeBERTaV3 [17], which improves the BERT [11] and RoBERTa [18] models using disentangled attention and enhanced mask decoder. The model was trained using the CC100 data, containing 100+ languages including Bulgarian.

For all pre-trained language models, the process of fine-tuning is very similar and consists of the following steps. Initially, we get the tokenized data from the preprocessing step aligned at the token level and we apply the specific model tokenization, which can be WordPiece [19], BPE [20] or SentencePiece [21] for different language models. When a token is split into multiple sub-tokens, each of the sub-tokens receives the label of the token, i.e. if a token is erroneous, then all of its sub-tokens will also be labelled as erroneous. The embeddings produced are then fed as inputs to the token classification language model, which consists of a language model and a fully-connected layer. The produced output contains classified sub-tokens, which have to be merged back into token-level predictions. This happens by considering every sub-token and if any of them are predicted to be erroneous, then the whole token is classified as erroneous. The experiments with all language models are shown in Sect. 5.

### 3.3 Error correction

The error correction task transforms each token classified as erroneous from the detector model to the sequence of characters. We use k-nearest neighbors as a naive baseline model. The method consists of finding the nearest neighbors from the CLaDA-BG dictionary [15] for each erroneous token using Levenshtein distance as a proximity measure. We only consider distances of one and two. If multiple tokens are found with the same distance, then word frequency is used to pick the best one and if no candidates are found, then the token remains unchanged.

As a second approach, we employ neural machine translation at character level to translate the classified as erroneous tokens into their corrected versions. Our model is a character-level sequence-to-sequence model with attention. The encoder is a bidirectional LSTM [22] and transforms the characters into a sequence of hidden states  $\{h_i\}$ . During the decoding process, the model uses an attention mechanism in order to decide what information to use from the encoder's hidden states. The embeddings between the encoder and the decoder are shared and cross-entropy loss between the pre-

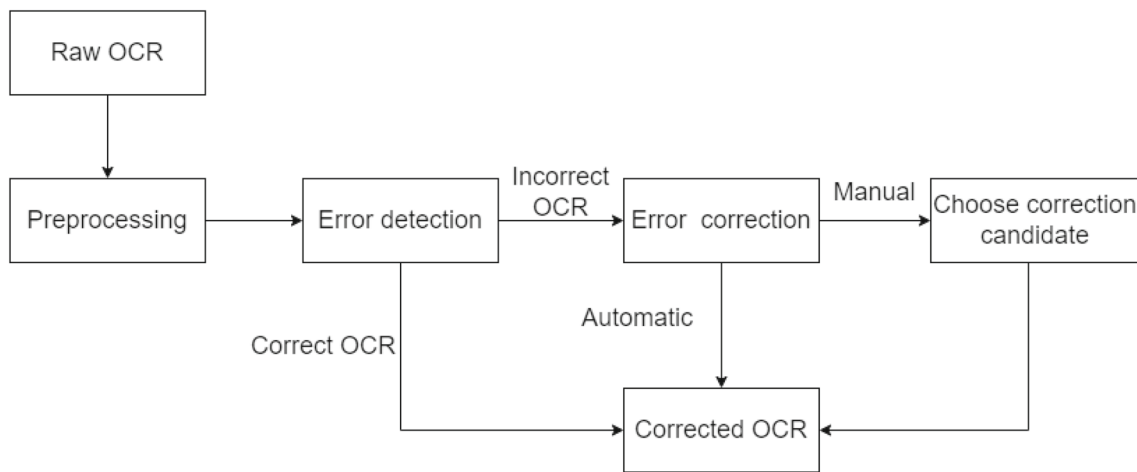


Fig. 1 Pipeline for post-OCR text correction

dicted output and the gold standard text is used as an initial loss function.

Additionally, some enhancements are implemented to the base sequence-to-sequence model, similar to the ones proposed in [23]. The first enhancement is diagonal attention loss. According to [4], the most frequent OCR errors are substitution errors, which means that the OCR engine confused a character with another one. This is expected, as the OCR is most likely to confuse a character rather than hallucinate or completely ignore pixels from the original document. The nature of the OCR is such that, errors are happening for characters that look similar. This is not the case for example in regular spell checking where common mistakes are for characters which are next to each other on the keyboard or the order of the characters is reversed due to fast typing. Therefore, during correction, reordering of characters almost never occurs, so we expect the attention weights of the model to have higher values close to the diagonal of the encoder. The diagonal loss  $L_{diag}$  is added to the final model loss so that the elements from the attention vector that are not within  $m$  steps from the current timestep  $t$  are added to the loss. In this way, we encourage elements away from the diagonal to have lower attention scores.

$$L_{diag} = \sum_t \left( \sum_{i=1}^{k-m} \alpha_{t,i} + \sum_{i=t+m}^N \alpha_{t,i} \right)$$

The second enhancement is a copy mechanism, which enables the model to copy characters from the source directly rather than generating a character on each step. This allows us to leverage the overlap of correctly recognized characters between the raw OCRed text and the gold standard output. We use the copy mechanism first proposed in [24]. It calculates a *generation probability*  $P_g$  for each time step  $t$  which is used as a soft switch to choose whether to generate a char-

acter or copy directly from the source. Copying a character from the input sequence happens with probability  $1 - P_g$  by sampling the encoder attention distribution. For time step  $t$ , the *generation probability* is calculated from the context vector  $c_t$ , the decoder state  $s_t$  and the decoder input  $x_t$ :

$$P_g = \sigma(w_c^T c_t + w_s^T s_t + w_x^T x_t + b)$$

The final enhancement is the coverage mechanism, which keeps track of the attention distribution over all past time steps in a coverage vector:

$$g_t = \sum_{i=0}^{t-1} \alpha_i$$

Intuitively, the coverage vector is a distribution over the erroneous token’s characters, which represents how much coverage each character has received. Additionally, we also add it as a separate parameter when calculating the attention distribution  $\alpha_t$ :

$$e_{t,i} = \text{vtanh}(W_1 S_{t-1} + W_2 h_i + W_g g_{t,i})$$

$$\alpha_t = \text{softmax}(e_{t,i})$$

The coverage vector is also used to encourage the model to attend to different locations by adding a coverage loss  $L_c$ :

$$L_c = \sum_t \sum_{i=0}^N \min(\alpha_{t,i}, g_{t,i}).$$

The final loss of the model is calculated as:

$$L = L_{ce} + L_{diag} + L_c.$$

In Sect. 5 we have evaluated the combinations of these enhancements and we see significant improvement when all of them are applied at the same time in the final sequence-to-sequence model.

## 4 Data

This section describes the datasets used for training and evaluating the models as well as the process of generating the synthetic dataset.

### 4.1 Datasets

There is a substantial orthographic variety in the Bulgarian historical documents due to the nine language reforms applied to the Bulgarian language. In 1945 the current Bulgarian orthography was accepted and prior to that, there were two major orthographies that were used - the Drinov orthography and the Ivanchev orthography.

We have used the Bulgarian dataset [25] provided by the ICDAR 2019 competition on post-OCR text correction. The dataset contains around 200 documents written in the Ivanchev orthography. The data comes aligned at the character level, but we needed it at the token level. To get that alignment, we tokenize the gold standard based on white space and for each token, we take the same number of characters from the aligned OCR'd version. By splitting the data this way we avoid a situation where the incorrectly split OCR'd word also splits the GS words. After that, we remove the special alignment symbols: “@” and “#”, and the resulting tokens are compared in order to set the labels as 0 for correct and 1 for incorrect. An example sentence from the ICDAR 2019 dataset is shown in Fig. 2.

Since the ICDAR 2019 dataset contains documents only in the Ivanchev orthography we release a small benchmark dataset, which contains documents in the Drinov orthography - the Drinov Orthography for Post-OCR Correction dataset [26]. The corpus was created by annotating a historical newspaper collection provided to us by the National Library “Ivan Vazov” (NLIV) in Plovdiv, Bulgaria.<sup>3</sup> We consider printed versions of these documents, which we manually annotate and align at the character level in the same format as the one from the ICDAR 2019 competition. The number of sentences, words and errors, as well as mean CER  $\mu$  and standard deviation CER  $\sigma$  of the two datasets are shown in Table 1.

However, some sentences are too noisy for learning, as they either contain too many typos, unrecognized or misaligned words. We discard those during training and evaluation by using the normalized Levenshtein distance between the OCR'd sentence and its corresponding gold standard.

We only consider sentences having normalized Levenshtein distance less than 0.5. The distribution of the Levenshtein distances from the two datasets are shown in Fig. 4.

### 4.2 Synthetic data

Since the available data is scarce for the Bulgarian historical spellings, we generate a large amount of synthetically constructed misspelled-correct sentence pairs to support training. We have used around 100k sentences from old Bulgarian books available online in the current orthography. The process of synthetic data generation consists of the following two steps: first, we transform the data from modern to the historical spelling of the respective dataset. In order to do that, we have developed a method which automatically converts from modern to one of the two major historical spellings. The method uses spelling transformation rules [27] and a morpho-syntactic tagging model for Bulgarian [28] to transform between the two orthographies. For example, the Bulgarian verb “пяя”, which translates to “sing” in English, after the conversion is applied results in “пѣя” and “пѣя” in the Drinov and Ivanchev orthographies respectively. After the word is converted, we insert noise into it by replacing a character with another one using the confusion matrix of each dataset (ICDAR 2019 and DOPOC). The confusion matrix contains the character error frequencies from the documents and is created automatically by tokenizing the OCR aligned and GS aligned sentences and comparing their characters. The amount of generated noise should be similar to the actual noise from real-world OCR'd documents, otherwise it will have detrimental impact on the accuracy of the model [29]. This is accomplished by the calculated probabilities from the confusion matrix. An example of synthetically generated sentence pair is shown in Fig. 3.

## 5 Experiments

This section describes the experimental setup and performance of the evaluated models for error detection and correction tasks on both datasets.

### 5.1 Setup

For evaluating the error detection models we use the standard metrics for classification tasks: precision, recall and F1-score. To evaluate error correction we use % of improvement, which is calculated by comparing the weighted sum of the Levenshtein distances between the noisy raw OCR'd text  $ocr$  and the corrected one  $c$  with the gold standard  $gs$ .

$$\% \text{ improvement} = \frac{\text{lev}(ocr, gs) - \text{lev}(c, gs)}{\text{len}(gs)}$$

<sup>3</sup> <https://digital.libplovdiv.com/en>.



[OCR\_toInput] Не, не, не бѣше братъ ми-мень Не бн оставилъ той така, Ме би  
 ме хвърлила, ивъв тжга, Та вторъ љт,тъ съ него раз- дѣлень,  
 [OCR\_aligned] Не, не, не бѣше братъ ми-мень Не бн оставилъ той така, Ме би  
 ме хвърлила, ивъв тжга, Та вторъ љт,тъ съ него раз- дѣлень  
 [GS\_aligned] Не, не, не бѣше братъ ми-мень Не би оставилъ той така, Ме би  
 ме хвърлилѣъ вѣвъ тжга, Та вторъ ѣѣпжтъ съ него разѣѣдѣлень

**Fig. 2** Example sentence from the ICDAR 2019 dataset

**Table 1** Sentence, word, and error quantities, mean CER and standard deviation CER

Dataset	# Sentences	# Words	# Errors	$\mu$ CER	$\sigma$ CER
ICDAR 2019	4900	68,511	25,703	16.65	16.13
DOPOC	227	5152	589	2.94	1.67

той казва, че ще издигне черкова, дето е била пещерата, за споменъ, а тя му  
 казва, че ще си продаде всичкитѣ ялмази, за да раздаде милостиня на сиромаситѣ  
 той казва, чс щс издигне черкова, дето е бмла пещсрага, за споменъ, а тя му  
 казба, че щс си продаде всичкитѣ ялмазм, за да раздаде милостѣня на сиромаситѣ

**Fig. 3** Example of a synthetically generated sentence pair in the Ivanchev orthography. The first sentence is correct, while the second is misspelled

The train dataset is split further into train and development datasets. The train dataset contains 90% of the examples and the development dataset contains the rest 10%. For the pre-trained language models, we have used the Huggingface's transformers library [30]. We are fine-tuning them for 10 epochs using AdamW as an optimizer and cross-entropy loss between the predicted output and the gold standard text. The learning rate is  $3e-5$  and the sequence length is 100. The error correction model was trained for 60 epochs and beam search was used at inference time. We trained all models on Tesla T4 GPU 16GB with early stopping on the development set. The models are trained on the ICDAR 2019 dataset to which we add synthetically generated data for the specific orthography as described in Sect. 4.2.

The evaluated language models for error detection are: BERT [11], DistilBERT [31], RoBERTa<sup>4</sup> [18], XLM-RoBERTa [32], XLM-MLM [33] and mDeBERTaV3 [17], as well as two naive baselines - CLaDA-BG and fastText.

## 5.2 Results


The results from the experiments are shown in Table 2. The best-performing model for both datasets is the multilingual DeBERTaV3 with synthetic data during training. On the ICDAR 2019 competition the best-performing model (CCC) achieves F1-score of 0.77 on the Bulgarian dataset, which we manage to improve by +0.13. DeBERTaV3 gets a 0.65 F1-score on the DOPOC dataset, which is lower than the results from the ICDAR 2019 dataset. We attribute these lower detection scores to two factors: better document quality of the

dataset which makes the errors harder to detect and lack of original annotated train dataset in the Drinov orthography. Without synthetic data during training, the model is not able to learn a reasonable distribution of the Drinov orthography. However, after adding the synthetically noised data, we see a significant improvement from 0.46 to 0.65 F1-score for DeBERTaV3. The CCC baseline was not run on the DOPOC dataset since the hyperparameters of the model were not shared in order to properly finetune it on the Drinov orthography. Thus, we only use it as a baseline comparison against the best-performing model from the ICDAR 2019 competition.

In Table 3 the results for the error correction models described in Sect. 3.3 are shown. The kNN baseline model is actually worsening the quality of the raw OCR because it does not consider the individual mistaken characters and is often picking up candidates which are correct according to the dictionary but in reality, have greater edit distance than the original text. The character-based Seq2Seq model achieves good results compared to the baseline. After adding copy mechanism to the original Seq2Seq model, the performance is notably better for both orthographies. This can be explained since the first pass OCR manages to predict most of the characters correctly. The best-performing model combines all enhancements described in Sect. 3.3 and is trained using both the original and synthetic datasets. The Seq2Seq Final achieves 25.4% of improvement on the ICDAR 2019 dataset, which improves the results of the best-performing model from the competition (CCC) by +16.4% of improvement. On the DOPOC dataset, we are able to properly correct most of the erroneous tokens and get a 26.9% of improvement on the original raw OCRed documents. The above results clearly demonstrate that by leveraging synthetic data we can

<sup>4</sup> <https://huggingface.co/iarfmoose/roberta-base-bulgarian>.

**Fig. 4** Distribution of the normalized Levenshtein distance for the ICDAR 2019 and DOPOC datasets



**Table 2** Error detection: evaluation results in terms of P, R, and F1

Model	Synthetic data	ICDAR 2019			DOPOC		
		P	R	F1	P	R	F1
CCC▼	No	—	—	0.77	—	—	—
CLaDA-BG∇	No	0.77	0.55	0.64	0.28	0.56	0.37
	Yes	0.79	0.57	0.66	0.30	0.53	0.37
FastText∇	No	0.66	0.81	0.73	0.33	0.61	0.43
	Yes	0.83	0.46	0.60	0.46	0.29	0.35
mBERT	No	0.79	0.80	0.81	0.68	0.26	0.38
	Yes	0.83	0.84	0.84	0.60	0.48	0.53
DistilmBERT	No	0.78	0.80	0.80	0.68	0.24	0.36
	Yes	0.83	0.83	0.82	0.61	0.42	0.49
RoBERTa	No	0.75	0.79	0.77	0.71	0.26	0.39
	Yes	0.75	0.81	0.78	0.59	0.38	0.46
XLM-RoBERTa	No	0.81	0.85	0.83	0.78	0.29	0.42
	Yes	0.84	0.88	0.85	0.70	0.52	0.60
XLM-MLM-100	No	0.81	0.84	0.83	0.75	0.26	0.39
	Yes	0.83	0.84	0.84	0.61	0.47	0.53
mDeBERTaV3	No	0.87	0.89	0.88	0.78	0.33	0.46
	Yes	0.89	0.91	<b>0.90</b>	0.74	0.59	<b>0.65</b>

We present the results with and without synthetic data and highlight the best model for each orthography. ▼ denotes ICDAR 2019 baseline and ∇ denotes naive baseline model

Bold highlights the best-performing results for both datasets

**Table 3** Error correction: evaluation results in terms of % improvement

Model	Synthetic Data	ICDAR 2019	DOPOC
CCC▼	No	9%	—
kNN∇	No	−2.4%	−3.2%
	Yes	−1.6%	−2.1%
Seq2Seq Base	No	17.1%	19.3%
	Yes	19.8%	22.1%
Seq2Seq Copy	No	19.2%	21.7%
	Yes	21%	25.3%
Seq2Seq Final	No	22.7%	25.3%
	Yes	<b>25.4%</b>	<b>26.9%</b>

kNN is k-nearest neighbors with Levenshtein distance, Seq2Seq Base is a base character-level sequence-to-sequence model with attention. Seq2Seq Copy includes a copy mechanism and Seq2Seq Final includes all enhancements—diagonal attention loss, copy and coverage mechanisms. ▼ denotes ICDAR 2019 baseline and ∇ denotes naive baseline model

Bold highlights the best-performing results for both datasets

further improve the quality of digitized historical documents written in any of the two main Bulgarian historical spellings.

## Error analysis

In Table 4 we have shown the types of errors that an OCR engine makes with examples in English. Our models are able to detect and correct most of the errors introduced during the OCR process. Examples of errors fixed by the post-OCR correction method are shown in Fig. 5. The model is also

**Table 4** OCR error types with examples

Error type	Example
Misrecognized character	more → moro
Missing character	there → here
Hallucination	where → wherever
Run-on	he ran → heran
Incorrect split	locomotive → loco motive



	Errors fixed by post-correction		Errors <b>not</b> fixed by post-correction	
Raw OCR	разбн <del>т</del> тъ	исповедни <del>ц</del> итб	сми <del>в</del> атъ	вид <del>н</del> яъ
GS OCR	разб <del>и</del> тъ	исповѣд <del>н</del> ици <del>т</del> ѣ	стѣ <del>ж</del> пватъ	в димъ
Post-corrected	разб <del>и</del> тъ	исповѣд <del>н</del> ици <del>т</del> ѣ	сми <del>в</del> атъ	вид <del>н</del> я

**Fig. 5** Our model manages to fix most of the errors, as shown in the examples from the first two columns. However, in some cases it fails to properly correct the error, as demonstrated in the two subsequent columns

**Table 5** Quantities of word segmentation errors in the ICDAR 2019 and DOPOC datasets

# errors	ICDAR 2019	DOPOC
Total	25703	589
Word segmentation	3503	44
Run-on	2423	11
Incorrect split	1080	33

able to successfully correct multiple incorrectly recognized characters.

However, we have noticed that the models do not perform well when it comes to fixing word segmentation errors, which are incorrectly recognized word endings. These types of errors are divided into two types [4]: run-on and incorrect split errors. In case of a run-on error, the model usually simply misspells the input word without adding the necessary space, whereas incorrect split errors often result in misspelled words rather than being properly combined. An example of an unsuccessful attempt to fix a run-on error is shown in the rightmost column in Fig. 5. The counts of word segmentation errors from our datasets are shown in Table 5.

There is some research done on specifically fixing word segmentation errors, in [34] a character-level Seq2Seq model is utilized for correcting these types of errors separately. The training and test data are scientific texts from the electronic submissions of the ACL collection. Word-level data is first transformed into character-level by removing all white spaces between words in the input data and replacing them with “###”. The model is then trained on this data to learn the word boundaries represented by the special symbols “###” and use it to correct word boundary errors in the test dataset with achieved precision and recall of above 0.95.

## 6 Conclusion and future work

In this paper, we have presented a new post-OCR text correction method for Bulgarian historical documents. Our proposed methods reduce the errors introduced during recognition and notably improve the quality of the documents. The

characteristics of the Bulgarian language are very similar to other Slavic languages, so the methods can easily be adapted for them as well with minimal changes required.

To solve the error detection task, we used various pre-trained LLMs. The best performing one is DeBERTaV3 which achieves 0.90 F1-score on the ICDAR 2019 dataset and 0.65 F1-score on the DOPOC dataset. Character-level sequence-to-sequence model augmented with diagonal attention loss, copy and coverage mechanisms was used for tackling the error correction task which improves the text quality of the ICDAR dataset by 25% and of the DOPOC dataset by 27%. Our results show that we get an improvement over the best performing model in the ICDAR 2019 competition [2] with +0.13 F1-score for error detection and +16% for the error correction task. Additionally, we develop a method for generating noised synthetic data, which automatically transforms modern Bulgarian spelling into a historical one by using concrete spelling and grammatical norms which have an outcome in the spelling. This allows us to utilize many Bulgarian literature text resources available online and adapt them for any of the two main historical spellings.

In future work, we plan to address the word segmentation errors mentioned in Sect. 5.2. The models do not perform very well when correcting this type of error.

**Acknowledgements** This work is partially funded by project UNITE BG05M2OP001–1.001-0004 by the OPSESG and by the EU through NextGenerationEU funds, through the National Recovery and Resilience Plan of the Republic of Bulgaria, project No BG-RRP–2.004-0008.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. David, Miller., Sean, Boisen., Richard, Schwartz., Rebecca, Stone., Ralph, Weischede.: Named entity extraction from noisy input: Speech and OCR. In *Proceedings of Sixth Applied Natural Language Processing Conference*, pages 316–324, (2000)
2. Christophe, Rigaud., Antoine, Doucet., Mickael, Coustaty., Jean-Philipp, Moreux.: ICDAR 2019 competition on post-OCR text correction. In *Proceedings of 2019 International Conference of Information Retrieval (ICDAR)*, pages 1588–1593, (2019)
3. Martin, Volk., Lenz, Furrer., Rico, Sennrich.: Strategies for reducing and correcting OCR errors. In *Sporleder, C., van den Bosch, A., Zervanou, K. (eds) Language Technology for Cultural Heritage. Theory and Applications of Natural Language Processing*, pages 3–22, (2011)
4. Thi-Tuyet-Hai, Nguyen., Adam, Jatowt., Mickael, Coustaty., Nhu-Van, Nguyen., Antoine, Doucet.: Deep statistical analysis of ocr errors for effective post-ocr processing. In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, 29–38, (2019)
5. Christian, Strohmaier., Ch Ringlstetter., Schulz, Klaus., Stoyan, Mihov.: Lexical postcorrection of OCR-results: The web as a dynamic secondary dictionary? In *Proceedings of the 7th International Conference on Document Analysis and Recognition*, 1133–1137, (2003)
6. Klaus, Schulz., Stoyan, Mihov., Petar, Mitankin.: Fast selection of small and precise candidate sets from dictionaries for text correction tasks. In *Proceedings of the 9th International Conference on Document Analysis and Recognition (ICDAR 2007)*, pages 471–475, (2007)
7. Paula, Estrella., Pablo, Paliza.: Ocr correction of documents generated during argentina’s national reorganization process. In *Proceedings of the 1st International Conference on Digital Access to Textual Cultural Heritage*, pages 119–123, (2014)
8. Vladimir I. Levenshtein.: Binary codes capable of correcting deletions, insertions and reversals. In *Soviet physics doklady*, (1966)
9. Guillaume, Chiron., Antoine, Doucet., Mickaël, Coustaty., Jean-Philipp, Moreux.: ICDAR 2017 competition on post-OCR text correction. In *Proceedings of the 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, 1423–1428, (2017)
10. Chantal, Amrhein., Simon, Clematide.: Supervised OCR error detection and correction using statistical and neural machine translation methods. *Language Technology and Computational Linguistic*, (2018)
11. Jacob, Devlin., Ming-Wei, Chang., Kenton, Lee., Kristina, Toutanova.: BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June (2019). Association for Computational Linguistics
12. Thi, Tuyet Hai Nguyen., Adam, Jatowt., Nhu-Van, Nguyen., Mickaël, Coustaty.: Neural machine translation with BERT for post-OCR error detection and correction. In *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries*, pages 333–336, (2020)
13. Shaohua, Zhang., Haoran, Huang., Jicong, Liu., Hang, Li.: Spelling error correction with soft-masked BERT. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 882–890, Online, July (2020). Association for Computational Linguistics
14. Rui, Dong., David, Smith.: Multi-input attention for unsupervised OCR correction. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2363–2372, Melbourne, Australia, July (2018). Association for Computational Linguistics
15. Kiril, Simov.: Integrated language and knowledge resources for a bulgarian-centric knowledge graph. In *Digital Presentation and Preservation of Cultural and Scientific Heritage*, pages 65–74, (2019)
16. Armand, Joulin., Edouard, Grave., Piotr, Bojanowski., Tomáš, Mikolov.: Bag of tricks for efficient text classification. *CoRR*, arXiv:1607.01759, (2016)
17. Pengcheng, He., Xiaodong, Liu., Jianfeng, Gao., Weizhu, Chen.: DeBERTa: Decoding-enhanced bert with disentangled attention. In *International Conference on Learning Representations*, (2021)
18. Yinhan, Liu., Myle, Ott., Naman, Goyal., Jingfei, Du., Mandar, Joshi., Danqi, Chen., Omer, Levy., Mike, Lewis., Luke, Zettlemoyer., Veselin, Stoyanov.: Roberta: A robustly optimized BERT pretraining approach. *CoRR*, arXiv:1907.11692, (2019)
19. Yonghui, Wu., Mike, Schuster., Zhifeng, Chen., Quoc V., Le., Mohammad, Norouzi., Wolfgang, Macherey., Maxim, Krikun., Yuan, Cao., Qin, Gao., Klaus, Macherey., Jeff, Klingner., Apurva, Shah., Melvin, Johnson., Xiaobing, Liu., Lukasz, Kaiser., Stephan, Gouws., Yoshikiyo, Kato., Taku, Kudo., Hideto, Kazawa., Keith, Stevens., George, Kurian., Nishant, Patil., Wei, Wang., Cliff, Young., Jason, Smith., Jason, Riesa., Alex, Rudnick., Oriol, Vinyals., Greg, Corrado., Macduff, Hughes., Jeffrey, Dean.: Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, arXiv:1609.08144, (2016)
20. Rico, Sennrich., Barry, Haddow., Alexandra, Birch.: Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany, August (2016). Association for Computational Linguistics
21. Taku, Kudo.: Subword regularization: Improving neural network translation models with multiple subword candidates. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia, July (2018). Association for Computational Linguistics
22. Sepp, Hochreiter., Jürgen, Schmidhuber.: Long short-term memory. In *Neural computation*, pages 1735–1780, (1997)
23. Shruti, Rijhwani., Antonios, Anastasopoulos., Graham, Neubig.: OCR Post Correction for Endangered Language Texts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5931–5942, Online, November (2020). Association for Computational Linguistics
24. Abigail See., Peter J Liu., Christopher D., Manning.: Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada, July (2017). Association for Computational Linguistics
25. Christophe, Rigaud., Antoine, Doucet., Mickael, Coustaty., Jean-Philipp, Moreux.: Dataset of ICDAR 2019 competition on post-OCR text correction, (2019)
26. Beshirov, Angel, Dobрева, Milena: Dimitar Dimitrov. Ivan Koychev, and Preslav Nakov. Drinov orthography for post-OCR correction dataset, Momchil Hardalov (2024)
27. Charalozova, Katya: Codification of the Norms of the Bulgarian Standard Language from the End of the 19th and the Beginning of the 20th Century (1879–1921). Prof. Marin Drinov Publishing House, Sofia (2022)
28. Nikola, Ljubešić., Petya, Osenova., Kiril, Simov.: The CLASSLA-StanfordNLP model for morphosyntactic annotation of standard bulgarian 1.0, 2020. Slovenian language resource repository CLARIN.SI

29. Eva, D'hondt., Cyril, Grouin., Brigitte, Grau.: Generating a training corpus for OCR post-correction using encoder-decoder model. In Greg Kondrak and Taro Watanabe, editors, *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1006–1014, Taipei, Taiwan, November (2017). Asian Federation of Natural Language Processing
30. Thomas, Wolf., Lysandre, Debut., Victor, Sanh., Julien, Chaumond., Clement, Delangue., Anthony, Moi., Pierric, Cistac., Tim, Rault., Rémi, Louf., Morgan, Funtowicz., Jamie, Brew.: Hugging-face's transformers: State-of-the-art natural language processing. *CoRR*, [arXiv:1910.03771](https://arxiv.org/abs/1910.03771), (2020)
31. Victor, Sanh., Lysandre, Debut., Julien, Chaumond., Thomas, Wolf.: Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, [arXiv:1910.01108](https://arxiv.org/abs/1910.01108), (2019)
32. Alexis, Conneau., Kartikay, Khandelwal., Naman, Goyal., Vishrav, Chaudhary., Guillaume, Wenzek., Francisco, Guzmán., Edouard, Grave., Myle, Ott., Luke, Zettlemoyer., Veselin, Stoyanov.: Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online, July (2020). Association for Computational Linguistics
33. Guillaume, Lample., Alexis, Conneau.: Cross-lingual language model pretraining. *CoRR*, [arXiv:1901.07291](https://arxiv.org/abs/1901.07291), (2019)
34. Vivi Nastase and Julian Hitschler. Correction of OCR word segmentation errors in articles from the ACL collection through neural machine translation methods. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.