Generating Errors: OCR post-processing for Icelandic

Anonymous Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Anonymouser Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Anonymousest Author
Affiliation / Address line 1
Affiliation / Address line 2
Affiliation / Address line 3
email@domain

Abstract

We describe work on enhancing the performance of transformer-based encoder-decoder models for OCR post-correction on modern and historical Icelandic texts, where OCRed data is scarce. We trained six models, four from scratch and two fine-tuned versions of Google's ByT5, on a combination of real data and texts populated with artificially-generated errors. Our results show that the models trained from scratch, as opposed to the fine-tuned versions, benefited the most from the addition of artificially generated errors.

1 Introduction

Optical Character Recognition (OCR) is used to digitize texts by converting scanned documents into machine-readable text. Unfortunately, OCR errors are prevalent, particularly when it comes to old texts, where data tends to be scarce, and post-correction is often required to improve the extracted text's accuracy (e.g. Nguyen et al. 2021).

Transformer-based encoder-decoder models have been shown to be effective in various natural language processing tasks, including machine translation (Vaswani et al., 2017; Chen et al., 2018) and text summarization (Garg et al., 2021). In this study, we investigate the use of such models for OCR post-correction under scarce data condition, as a sequence-to-sequence problem, similar to how neural machine translation (NMT) systems approach the problem of translation. To address the lack of resources available for training models when dealing with OCRed texts, we propose the use of artificially generated errors to improve the performance of the models, which has been shown to be an effective way of generating data for text correction (Kasewa et al., 2018). The main contribution of this study is an examination of the effectiveness of using artificially generated errors to improve the performance of transformer-based encoder-decoder models for OCR post-correction when data scarcity is a limiting factor. Furthermore, we publish our best performing model under an open license.[1]

Section 2 discusses related work while 3 describes the dataset and error generation methods used in this study. Section 4 presents the proposed models and training methods. Section 5 presents the experimental results and analysis, 6 discusses limitations and future work and 7 concludes.

2 Related Work

Previously, Daðason et al. (2014) developed a tool for post-processing Icelandic 19th century texts based on an error model containing statistical information on word and character errors and an n-gram language model. Their tool correctly identifies and corrects 52.9% of errors in their evaluation set, but they do not report on the number of new errors introduced in their output.

Poncelas et al. (2020) report a 63% error correction rate with their OCR post-processing tool on an English text from the 18th century text. They used a scoring system based on string-similarity to find possible substitutions for perceived errors and a language model to evaluate the edited sentences.

Richter et al. (2018) use a hidden Markov model alongside a modified version of the Viterbi algorithm and a dictionary to decode OCRed texts in Faroese into a hypothetical corrected version. They reduced the word error rate of 7.8% from the OCR base process to 5.4%.

3 Data

We used a combination of real OCRed texts, processed by ABBYY FineReader, and regular texts, populated with artificially generated errors, for training and evaluating our OCR post-correction

[1] Models available at https://anonymous.submission/models.

Original	Corrected	Frequency
p	*þ*	2,779
rn	*m*	166
m	*rn*	77

Table 1: Examples of the extracted errors.

models. The evaluation data solely comprised real OCRed texts, alongside their manually corrected counterparts, which were used to ensure that the models' performance is reflective of real-world OCR output. The ground truth (GT)[2], i.e. the data from which the errors were extracted, consist of around $375K$ tokens from 80 texts published between 1874 and 1913, which were manually corrected, while the data used in training and validation, which includes the OCRed texts as well as the texts populated with the artificial errors, amount to roughly $9.2M$ tokens.

The data into which the artificial errors were inserted were taken from the Icelandic Gigaword Corpus (IGC; Steingrímsson et al. 2018) and the Icelandic Text Archive (ITA).[3]

Overall, the training data consists of $7.8M$ tokens, the validation set, which is approximately 15% of the whole dataset, consists of around $1.4M$ tokens. The evaluation set, which all comprises manually corrected texts, only consists of $44K$ tokens. All datasets contain texts from different eras, including texts from the 19th century and the early 20th century, as well as texts from the last two decades of the 20th century. It should be noted that none of the data are based on texts printed in Gothic letters, which have been reported to be harder to recognize than other fonts (Furrer and Volk, 2011; Drobac et al., 2017). The evaluation set is divided into two parts, with $26K$ tokens being from modern texts and $18K$ tokens from texts from the 19th century and the early 20th century. This allows for an evaluation of the model's performance on different types of texts and OCR errors, which is crucial to ensure that the model is robust and generalizable.

It is important to note that the dataset used in this study is relatively small in size. One of the reasons is the scarcity of available corrected OCRed texts. Additionally, we observed that too large a proportion of modern texts resulted in the models over-generalizing and changing historical spellings to modern spellings. However, we aim for diplomatic transcription, preserving the original spelling. Therefore, we ensured that the dataset included texts from different eras while also avoiding over-generalization and alteration of historical spellings by limiting the amount of modern texts into which we inserted artificial errors.

[2]The GT is a product of the project *Language Change and Linguistic Variation in 19th-Century Icelandic and the Emergence of a National Standard*, led by Ásta Svavarsdóttir at the Árni Magnússon Institute (e.g. Svavarsdóttir et al. 2014).

[3]https://clarin.is/en/resources/textarchive/

	Model 1	Model 2
embeddings size	512	512
ffn embeddings	2,048	2,048
attention heads	4	4
encoder layers	5	5
decoder layer	5	5
tokenizer	WordPiece	SentencePiece
vocab. size	3,000	3,000

Table 2: The architecture of the two models trained from scratch.

3.1 Extracting the errors

The extraction of errors from the manually corrected OCRed texts was performed by analysing the $375K$ token dataset. The data were manually aligned, and then a line-by-line comparison was conducted between the OCRed texts and their corresponding manually corrected texts using Python's SequenceMatcher. In the process of extracting errors, tokens were considered to be the same if they shared the same index in a given line and had a similarity score [4] greater than 0.66. This two-fold restriction, taking into account both index and similarity, acted as a confidence threshold to ensure that the identified tokens were different versions of the same intended token.

The differences between the tokens, specifically focusing on character or character n-gram substitutions, such as *rn*→*m* and *þ*→*p*, were extracted as OCR errors. In total, 2,644 such errors were extracted, which were then filtered down to the 600 errors that occurred more than three times in the dataset. In addition, the frequency of each error was recorded, which allowed for the implementation of a weighting system during the artificial error generation process, ensuring that the errors were distributed in a way that somewhat reflected

[4]Calculated by finding "[...]the longest continuous matching subsequence that contains no "junk" elements", see: https://docs.python.org/3/library/difflib.html

	Older texts				Modern texts			
	base	Model 1	Model 2	ByT5	base	Model 1	Model 2	ByT5
chrF	94.8	94.80	96.00	96.22	95.2	95.52	95.75	96.09
chrF ERR	Ø	0.31%	25.65%	27.59%	Ø	5.31%	10.04%	17.37%
BLEU	97.2	97.19	98.24	98.54	97.7	97.63	98.06	98.24
BLEU ERR	Ø	-0.26%	37.34%	48.01%	Ø	-4.60%	14.30%	22.52%
WERR	Ø	-16.37%	35.04%	49.96%	Ø	-3.80%	17.02%	17.37%
CERR	Ø	-28.53%	18.34%	33.83%	Ø	-38.58%	-21.34%	-20.34%

Table 3: Our models trained on the GT compared to the base output from the OCR process.

their real-world frequency. Examples of extracted errors are shown in Table 1

3.2 Inserting the errors

To create the training dataset, we gathered texts from IGC and the ITA. Texts ranging from the 19th to the 21st century were collected to provide a diverse set of texts for training our models.

The extracted errors (see Section 3.1) were then inserted by randomly replacing characters or character n-grams via a lookup table. The randomness of the insertions of artificial errors had weights, based on the errors' frequency in the GT, in order to mimic the distribution of errors that occur in OCR output, so more frequent errors were made to appear more frequently in the dataset. However, to prevent the same errors from appearing excessively in the dataset, we used the $\log_{10} frequency$ of the errors to adjust the weights.

4 Models

In total, six models were trained. Two of them follow the architecture of model 1, laid out in table 2, two of them follow the architecture of model 2 in the same table, and two are a finetuned version of ByT5-base (Xue et al., 2022), a token-free model that operates directly on raw text. For every pair of the models, one was trained on the $375K$ tokens in the GT, and the other one on the whole dataset, around $7.8M$ tokens. This was done to study the artificially generated errors' impact on the models' output. We experimented with various hyperparameter configurations on the models we trained from scratch, but these specific configurations resulted in the highest performance.

It is well established that transformer models require large amounts of data to be trained effectively. In this study, our GT had a limited number of examples, which likely contributed to the poorer performance of the models trained from scratch on the smaller dataset.

5 Results

The six models we trained for post-processing of OCRed texts were applied to modern and historical texts to measure the impact and viability of using artificial errors to improve such models when available data is scarce. The results of these models were compared to the base output from the OCR process using four metrics: chrF (Popović, 2015) and BLEU (Papineni et al., 2002), character error rate and word error rate. The BLEU score is calculated by comparing texts on a word-level, while chrF score is calculated on a character-level and can be more accurate for inflected languages (Dowling et al., 2020).

Table 3 shows the results of our models trained on the GT compared to the base OCR output. Model 2 and the ByT5 model show moderate improvements for older texts, while model 1 performs similarly or worse than the base OCR output. Generally, the models do worse on the modern texts, as opposed to the historical ones, when only trained on the GT, which is to be expected as the GT solely consists of historical texts.

Table 4 shows the results of our models trained on the whole dataset compared to the base OCR output. The models all show substantial improvements compared to the models only trained on the GT, which suggests that the artificial errors have something to offer. Furthermore, the difference between WERR for the different text types was less than for the models only trained on the GT.

Note that while the models generally perform better on the historical texts, the addition of artificial errors improve their performance proportionally more on the modern ones. This could stem from the fact that the artificially-erroneous dataset

	Older texts				Modern texts			
	base	Model 1	Model 2	ByT5	base	Model 1	Model 2	ByT5
chrF	94.8	96.84	96.84	96.73	95.3	96.83	96.86	96.7
chrF ERR	Ø	39.35%	41.39%	37.38%	Ø	33.03%	33.49%	30.90%
BLEU	97.2	98.45	98.79	98.65	97.7	98.45	98.64	98.57
BLEU ERR	Ø	44.56%	56.92%	51.70%	Ø	31.70%	39.92%	36.91%
WERR	Ø	23.79%	52.60%	55.07%	Ø	18.00%	34.67%	42.97%
CERR	Ø	26.29%	47.55%	35.12%	Ø	10.01%	15.20%	1.93%

Table 4: Our models trained on the whole dataset compared to the base output from the OCR process.

includes modern texts, while the GT does not.

When evaluating the models on modern texts, we found that they were less capable in reducing errors in modern texts than in historical texts. This could be due to the fact that the GT only comprised historical data, suggesting that using solely historical OCRed texts is not a viable approach when designing an OCR post-processing tool for modern texts. The lower error rate reduction (ERR) on the modern texts assumably also stems from the higher base score on the modern texts, as opposed to the base score of the historical ones, leaving less room for improvement.

6 Limitations and future work

The cost of manually correcting OCR output is high, making it difficult to obtain a larger dataset for training. This has a direct impact on the ability of the models to perform well on a wider range of texts and OCR errors.

The models have the unwanted tendency to adapt to modern spellings when using a large amount of modern texts populated with artificial errors. This could lead to the alteration of historical spellings, which is not in line with our objectives, to produce diplomatic transcriptions. To mitigate this risk, more corrected texts are needed for the period of texts being OCRed.

Moreover, the use of artificially generated errors to enhance the performance of the models may not fully capture the complexity and diversity of real-world OCR errors. Future studies may benefit from incorporating a more diverse range of error types and more realistic error generation methods.

We are interested in investigating optimal methods for generating realistic errors to use in training the models. As previously mentioned, the artificial errors used in this study were generated by randomly inserting errors that were extracted from the GT into other texts. However, there may be more effective methods for generating errors that better simulate real-world OCR errors. By finding and implementing these methods, the performance of OCR error correction models could be further improved. Furthermore, it could be beneficial to explore different architectures or different data augmentation techniques, such as including multiple versions of the same texts.

7 Conclusion

Our findings demonstrate that while fine-tuning pre-trained models on smaller datasets is an effective approach to improving the performance of OCR error correction models, it is possible to achieve comparable results by training an encoder-decoder transformer model from scratch. Model 2, which was trained from scratch, emerged as the best performer in our study, achieving a 41.39% ERR in chrF and 56.92% in BLEU ERR on the historical texts, and a 34.32% ERR in chrF and 40.18% ERR in BLEU on the modern texts. These results indicate that with proper architectural design, it is possible to train effective OCR error correction models without relying on pre-trained models or large datasets.

The evaluation dataset used in this study was relatively small and further testing on larger datasets would provide a more robust. However, the use of artificially generated errors in the training process was found to be effective in countering the challenges posed by data scarcity.

The fact that the models' performance improved proportionally more on the modern texts after the introduction of the artificial errors, which were by and large inserted into modern texts, indicates that in order to train a designated OCR post-processing tool for modern texts, a dataset consisting of modern texts in needed.

References

Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Niki Parmar, Mike Schuster, Zhifeng Chen, et al. 2018. The best of both worlds: Combining recent advances in neural machine translation. arXiv:1804.09849.

Jón Friðrik Daðason, Kristín Bjarnadóttir, and Kristján Rúnarsson. 2014. The journal *Fjölnir* for everyone: The post-processing of historical OCR texts. In *Proceedings of Language Resources and Technologies for Processing and Linking Historical Documents and Archives – Deploying Linked Open Data in Cultural Heritage*, pages 65–62, Reykjavík.

Meghan Dowling, Sheila Castilho, Joss Moorkens, Teresa Lynn, and Andy Way. 2020. A human evaluation of English-Irish statistical and neural machine translation. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 431–440, Lisboa, Portugal. European Association for Machine Translation.

Senka Drobac, Pekka Sakari Kauppinen, and Bo Krister Johan Linden. 2017. OCR and post-correction of historical Finnish texts. In *Proceedings of the 21st Nordic Conference on Computational Linguistics, NoDaLiDa, 22–24 May 2017, Gothenburg, Sweden*. Linköping University Electronic Press.

Lenz Furrer and Martin Volk. 2011. Reducing OCR errors in Gothic-script documents. In *Proceedings of the Workshop on Language Technologies for Digital Humanities and Cultural Heritage*, pages 97–103, Hissar, Bulgaria. Association for Computational Linguistics.

Apar Garg, Saiteja Adusumilli, Shanmukha Yenneti, Tapas Badal, Deepak Garg, Vivek Pandey, Abhishek Nigam, Yashu Kant Gupta, Gyan Mittal, and Rahul Agarwal. 2021. News article summarization with pretrained transformer. In *Advanced Computing: 10th International Conference, IACC 2020, Panaji, Goa, India, December 5–6, 2020, Revised Selected Papers, Part I 10*, pages 203–211. Springer.

Sudhanshu Kasewa, Pontus Stenetorp, and Sebastian Riedel. 2018. Wronging a right: Generating better errors to improve grammatical error detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4977–4983, Brussels, Belgium. Association for Computational Linguistics.

Thi Tuyet Hai Nguyen, Adam Jatowt, Mickael Coustaty, and Antoine Doucet. 2021. Survey of post-ocr processing approaches. *ACM Comput. Surv.*, 54(6).

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Alberto Poncelas, Mohammad Aboomar, Jan Buts, James Hadley, and Andy Way. 2020. A tool for facilitating ocr postediting in historical documents. *arXiv preprint arXiv:2004.11471*.

Maja Popović. 2015. chrF: character n-gram F-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.

Caitlin Richter, Matthew Wickes, Deniz Beser, and Mitch Marcus. 2018. Low-resource post processing of noisy ocr output for historical corpus digitisation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.

Steinþór Steingrímsson, Sigrún Helgadóttir, Eiríkur Rögnvaldsson, Starkaður Barkarson, and Jón Guðnason. 2018. Risamálheild: A very large Icelandic text corpus. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

Ásta Svavarsdóttir, Sigrún Helgadóttir, and Guðrún Kvaran. 2014. Language resources for early Modern Icelandic. In *Proceedings of Language Resources and Technologies for Processing and Linking Historical Documents and Archives – Deploying Linked Open Data in Cultural Heritage*, pages 19–25, Reykjavík.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30*, Long Beach, California.

Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. ByT5: Towards a Token-Free Future with Pre-trained Byte-to-Byte Models. *Transactions of the Association for Computational Linguistics*, 10:291–306.