



This is a self-archived – parallel-published version of an original article. This version may differ from the original in pagination and typographic details. When using please cite the original.

Jenna Kanerva, Cassandra Ledins, Siiri Käpyaho, Filip Ginter

OCR Error Post-Correction with LLMs in Historical Documents: No Free Lunches
2025

Publisher's PDF

Kanerva, Jenna, et al. "OCR Error Post-Correction with LLMs in Historical Documents: No Free Lunches." Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL-2025), edited by Špela Arhar Holdt et al., University of Tartu Library, Estonia, 2025, pp. 38–47. ACLWeb, <https://aclanthology.org/2025.resourceful-1.8/>.

CC-BY

OCR Error Post-Correction with LLMs in Historical Documents: No Free Lunches

Jenna Kanerva, Cassandra Ledins, Siiri Käpyaho, and Filip Ginter

TurkuNLP, Department of Computing
University of Turku, Finland
{jmnybl, caledi, siakap, figint}@utu.fi

Abstract

Optical Character Recognition (OCR) systems often introduce errors when transcribing historical documents, leaving room for post-correction to improve text quality. This study evaluates the use of open-weight LLMs for OCR error correction in historical English and Finnish datasets. We explore various strategies, including parameter optimization, quantization, segment length effects, and text continuation methods. Our results demonstrate that while modern LLMs show promise in reducing character error rates (CER) in English, a practically useful performance for Finnish was not reached. Our findings highlight the potential and limitations of LLMs in scaling OCR post-correction for large historical corpora.

1 Introduction

Digitizing and transcribing historical documents and literature is vital for preserving our cultural heritage and making it accessible for modern digital research methods. The transcription process relies on OCR, which naturally introduces noise into the output. The noise varies in severity depending on the quality of the source material and the OCR technology used, impacting the research usage of the data (Chiron et al., 2017b). The OCR output at two noise levels is illustrated in Figure 1. Although modern OCR systems are becoming increasingly accurate (Li et al., 2023), reprocessing large collections of historical literature remains a significant challenge, as the resources available to the institutions maintaining these collections are often insufficient for such an undertaking. Consequently, OCR error post-correction has been suggested as means of improving the historical collections without the need to repeat the whole transcription process (Nguyen et al., 2021).

Mild noise (0.04 CER):

A work of art, (be it what it may, house,
picture, book, or garden,) however
beautiful in it's underparts, loses half
it's value, if the general scope
of it be not obvious to conception.

Severe noise (0.19 CER):

bke up at Sx in the Mo.r aig. ll the
eaving Withr he went from Cbaud to Cbhh
every Soday, «d from Play. bote~PI0oaB
cu evi Niuht m the Week, but vd

Figure 1: Example extracts of texts at two different OCR noise levels from the ECCO dataset of 18th century literature.

Recent studies (Boros et al., 2024; Bourne, 2024) have proposed the application of LLMs to the task, with varying degrees of success. Currently, there is no clear consensus as to how LLMs can be applied to the task and how to deal with the various methodological challenges it poses. Our objective is to address some of these challenges as well as to assess several open LLMs to correct OCR-generated text when prompted to. We study hyperparameter optimization, quantization levels, input lengths, output post-processing and several novel correction methods so as to benchmark and improve the LLM performance on this task.

As our long-term goal is to post-correct two large historical datasets, one in English and the other in Finnish, we focus on these two languages as well as open-weight LLMs, since post-correction of large datasets with commercial models is infeasible cost-wise.

2 Related work

Despite decades of active research, post-correction of historical documents remains a challenge. The ICDAR 2017 and 2019 shared tasks (Chiron et al.,

2017a; Rigaud et al., 2019) addressed the lack of adequate benchmarks for evaluating OCR performance across several languages, introducing two tracks: detecting OCR errors, and correcting previously detected errors. This setting has naturally guided the development towards two-stage systems, and the best performing models in the ICDAR 2019 edition were based on the BERT model fine-tuned separately for each task (Rigaud et al., 2019). Such two-step approaches are still actively pursued, with e.g. Beshirov et al. (2024) applying a BERT classifier for error detection, and an LSTM-based seq2seq model for error correction in Bulgarian.

Recently, LLMs have been effectively applied to text correction problems, for example, Penteadó and Perez (2023) and Östling et al. (2024) demonstrated that LLMs perform well in grammatical error correction. Naturally, LLMs have been proposed also to the OCR post-correction task, in line with the two broad paradigms of LLM use: fine-tuning for the post-correction task and purely prompt-based zero-shot generation. Fine-tuning was applied e.g. by Soper et al. (2021) who fine-tune the BART model on the English subset of the ICDAR 2017 data and apply it to English Newspaper text. Veninga (2024) fine-tunes the character-based ByT5 model on the ICDAR 2019 data, with a prompt-based Llama model as a baseline. Similarly, Madarász et al. (2024) apply the mT5 model to historical Hungarian scientific literature, and Dereza et al. (2024) applies the BART model to historical Irish–English bilingual data.

In the zero-shot, prompt-based line of work, Boros et al. (2024) evaluated a variety of models and prompts on several multilingual historical datasets. Interestingly, the results of the study were mostly negative, concluding that LLMs (including the commercial GPT-4 model) are not effective at correcting transcriptions of historical documents, in many cases the LLM actually decreasing the quality instead of improving it. Bourne (2024) conducted a similar study on three historical English datasets, arriving at the opposite conclusion. They achieved over 60% reduction of character error rate at best, with most of the evaluated models improving the data quality.

Finally, several studies also pursue approaches that include the original image as an input, together with the OCR output to be post-corrected. Here, e.g. Chen and Ströbel (2024) combine a state-of-the-art transformer-based OCR system with the character-

based CharBERT model for handwritten text recognition, and Fahandari et al. (2024) propose a model iterating between OCR and post-correction steps for Farsi. Such image-text approaches are, nevertheless, beyond the scope of the present study.

3 Data

In our study, we utilize manually corrected samples of two large historical datasets, one for English and the other for Finnish.

3.1 English ECCO-TCP

Eighteenth Century Collections Online (ECCO) (Gale) is a dataset of over 180,000 digitized publications (books and pamphlets) originally printed in the 18th century Britain and its overseas colonies, Ireland, as well as the United States. While mainly in English, some texts appear in other languages. The collection was created by the software and education company Gale by scanning and OCRing the publications. ECCO has significantly impacted 18th-century historical research despite its known limitations (Gregg, 2021; Tolonen et al., 2021).

While ECCO contains only OCR engine output, the ECCO-TCP initiative¹ provides highly accurate, human-made text versions for 2,473 publications from the original collection (Gregg, 2022). In this study, we use a dataset from the Helsinki Computational History Group², where clean ECCO-TCP texts are paired with their corresponding ECCO OCR publications, creating an OCR ground truth dataset (Hill and Hengchen, 2019). The data is paired on page level, resulting in a dataset of 338K pages.

To prepare data for post-correction evaluation, we applied several filtering steps. First, we excluded 1,436 pages (0.4%) marked as blank in ECCO-TCP, ECCO OCR, or both. We also removed 5,782 pages (1.7%) containing fewer than 150 non-whitespace characters in either collection. Further filtering was necessary in cases of substantial mismatch between OCR and GT pages, typically with large chunks of text missing in either OCR or GT, or otherwise an obvious lack of correspondence. A brief manual analysis identified as typical causes (1) very noisy OCR output with a large amount of non-alphanumeric characters, likely from OCR engine transcribing an image; (2)

¹<https://textcreationpartnership.org/tcp-texts/ecco-tcp-eighteenth-century-collections-online/>

²<https://github.com/COMHIS>

Dataset	Language	Pages	OCR words	GT words	OCR w./pg.	CER	WER
ECCO-TCP	English	301,937	67,549,822	64,519,266	223.72	0.07	0.22
NLF GT	Finnish	449	449,088	461,305	1000.20	0.09	0.28

Table 1: Dataset statistics after preprocessing in terms of whitespace delimited words. *OCR w./pg.* denotes for mean OCR words per page, and CER and WER are average page-level character and word error rates in the data, weighted by the page length. For details about the metrics, see Section 4.1.

OCR and GT containing approximately the same text, but in different order due to misidentified reading order or column layout; and (3) significant length differences between pages, possibly from errors in automated page alignments, unrecognized regions left out in the OCR process, or omissions in the ECCO-TCP data.

To identify such instances, we align each OCR and GT page pair on their non-whitespace characters³ and slide a 100-character window across the alignment. If in any window less than 10% of characters were aligned as a match, the page was discarded from the dataset. In total, 28,907 (8.6%) pages were removed by this process.

Our filtering produced a dataset of 301,937 well-aligned pages (89.3% of the initial ECCO-TCP pages). While we do not filter by language, nearly all the data is in English, with other languages appearing only very rarely.

3.2 Finnish NLF Ground Truth Data

For Finnish, we use the National Library of Finland (NLF) OCR ground truth dataset⁴ of Kettunen et al. (2018, 2020), specifically intended for OCR quality evaluation. The data draws from the National Library’s collection of digitized newspapers, and consists of 479 pages randomly chosen from 188 different Finnish newspapers and journals published between 1836–1918, all printed in the Fraktur font.

The ground truth was created by manually correcting the OCR system output with reference to the original scans. The dataset contains texts produced by three different OCR software (ABBYY FineReader 7/8, ABBYY FineReader 11, and Tesseract) along with the ground truth. In this work, we use output produced by ABBYY FineReader 7/8, which is the OCR engine that has been used to digitize the majority of the NLF collection and therefore gives most useful information for a possi-

³Using global string alignment as implemented in the PairwiseAligner module in biopython.

⁴<http://digi.nationallibrary.fi>

ble future post-correction effort targeting it.

We applied the same filtering procedure as for the ECCO-TCP data, resulting in the removal of 29 pages (6%), and we further removed one page written in Swedish. The final dataset consists of 449 pages, with 449,088 OCR words, and 461,305 GT words of text. The key statistics for both datasets are shown in Table 1.

4 Experiments

First, we set out to evaluate the basic performance of different LLMs on the OCR error correction task and establish how the generation and model hyperparameters (e.g. sampling parameters and quantization) affect the results.

The page lengths in our data vary, with the ECCO-TCP pages on average at 200 words, and the Finnish newspaper pages at about 1,000 words. To improve comparability of the results, we split the pages to segments of 200 OCR words for English and 100 OCR words for Finnish, both corresponding to roughly 300 sub-words in OpenAI’s GPT-4 tokenization for the language in question. The length of roughly 300 sub-words was established as suitable in our initial experiments, however, we will carry out a more detailed evaluation of segment lengths as a separate experiment in Section 6.

Since the Finnish data is originally word-aligned, obtaining these shorter-than-page segments is trivial. For the English data, which is only page-aligned, we utilize the character-level OCR-GT alignments produced during data filtering (described in Section 3.1), allowing us to find corresponding points. In cases where the segment boundary falls within a region of poor alignment, we shift the boundary to the next reliable word (the word starting the next aligned region). Therefore, the exact segment length may vary depending on how well the OCR and GT strings could be aligned.

Given the substantial volume of our data, and the number of LLM runs necessary in our experiments, we randomly sample for each language a development and a test set, each containing 200

examples (i.e. segments of about 300 sub-words in length). These constitute 244K+243K GT characters for English, and 162K+165K GT characters for Finnish. The development set is used to set the generation parameters and the test set is used to report all results, unless otherwise stated.

4.1 Evaluation Metric

As a primary evaluation metric, we use Character Error Rate (CER) defined as the sum of character substitutions, deletions and insertions, divided by the length of the ground truth string. In line with the common practice in OCR post-correction literature, we mainly report relative CER reduction defined for one examples as $CER\% = (CER_{orig} - CER_{post})/CER_{orig} \times 100$ where *orig* and *post* refer to before and after correction, respectively. The overall CER% is calculated as an average of example-wise CER% weighted by example lengths in terms of OCR character count. The example-wise CER% values are further clipped not to go below -100% to prevent extremely large negative scores in cases where most of the text is omitted. The CER% therefore works on a range between -100% and 100%.

Many downstream applications utilizing historical corpora, such as various literature search interfaces, operate at the level of words rather than characters. Therefore, the main results are reported also in terms of Word Error Rate (WER) and its relative improvement (WER%). This metric is much like CER, but on the level of words.⁵

Finally, we apply few normalization steps before the evaluation. First, all unicode whitespace characters (`\s+`) are normalized into a single whitespace. Secondly, in line with similar works (Duong et al., 2021; Kettunen and Pääkkönen, 2016), we apply two normalization steps to address systematic differences between historical and modern spellings. In the English ECCO-TCP ground truth data, the long-s character `ſ` appears in places where modern English would use `s`. Similarly, in older Finnish historical texts `w` is often used where modern Finnish uses `v`. These spelling variations do not alter meaning, and we choose to disregard them by applying Unicode NFKC normalization, which handles both canonically equivalent and compatible transformations (including converting `ſ` to `s`) for both languages. Addition-

⁵We use the HuggingFace *evaluate* package implementation of both CER and WER.

ally, for Finnish, we replace all occurrences of `w` with `v` before evaluation, as modern Finnish does not use `w` except in loanwords or proper names, which occur only very rarely, making the difference negligible.

4.2 Models and Generation Parameters

We evaluate several top-tier open-weight models as well as OpenAI’s GPT-4o (v. 2024-08-06). The latter is included mostly for comparison, since it would not be cost-wise feasible to apply it to post-correction at a large scale, unlike open-weight models which can be applied on academic super-computing infrastructure. The evaluated open models are: Llama-3-8B-Instruct, Llama-3.1-8B-Instruct, and Llama-3.1-70B-Instruct from Meta (AI@Meta, 2024a,b), Mixtral-8x7B-Instruct-v0.1 from MistralAI (Jiang et al., 2024), and Gemma-2-9B-it and Gemma-2-27B-it from Google (Mesnard et al., 2024). It is noteworthy that while several of the open models are multilingual, none officially report supporting Finnish.

All models are run on the Ollama framework⁶ (v. 0.3.8) for fast inference, using the default 4-bit quantization unless otherwise stated. Other quantization levels are experimented separately, and reported in Section 5.2. All parameters of the framework and models are set to default except for the ones explored in Section 4.4. Note that we will not repeat the "Instruct" in model names in tables and figures for space considerations.

4.3 LLM Overgeneration Removal

LLM outputs often include undesired content in addition to the requested output. In most cases, the undesired text appears either before the corrected text (e.g. "*Here is your corrected text.*"), or after the corrected text has been provided (e.g. hallucinated continuation, or an additional commentary). This was noted also by Boros et al. (2024), who applied simple heuristics for removing any unwanted text, such as removal of whitespace, parts of the prompt, and specific phrases commonly appearing in the model’s output, together with trimming the predicted text if it exceeded 1.5 times the original.

Therefore, we base our overgeneration removal on automatically aligning the generated output against the original input on character level, and filtering out leading and trailing texts which do not align well to the input. For this purpose, we utilize

⁶<https://ollama.com/>

Model	English		Finnish	
	CER	WER	CER	WER
	%	%	%	%
Llama-3-8B	7.3	31.4	-68.8	-28.2
Llama-3.1-8B	19.5	37.7	-65.7	-30.1
Llama-3.1-70B	38.7	46.3	-47.0	-8.9
Mixtral-8x7B	-14.9	19.1	-76.5	-40.5
Gemma-2-9B	28.2	38.4	-24.0	-4.1
Gemma-2-27B	35.6	37.8	-19.1	0.0
GPT-4o	58.1	59.1	11.9	33.5

Table 2: Overall CER and WER relative improvement.

character-level local sequence alignment⁷ of the model’s output and the OCR text, and recover the region between the first and the last aligned characters. The alignment is configured to ignore whitespace and the ‘-’ character, to avoid text formatting discrepancies having an impact on the outcome of the alignment.

4.4 Parameter Optimization

The model generation parameters naturally affect the quality of the output and we therefore optimize the most critical parameters of the open-weight model generation on a held-out development set. As discussed earlier, this set is not used in any subsequent experiments.

Using the Optuna hyperparameter optimization library (Akiba et al., 2019), we set the temperature, top_k and top_p parameters. For each model and each language, we test 100 runs with different parameter combinations. Subsequently, the 10 best runs in terms of CER were selected, creating a range of possible best parameters. These ranges generally overlap across models but not across languages, therefore we pick a set of parameters for each language. The final parameters are chosen as the median value of the 10 best runs of every model. For English, the parameters are temperature 0.26, top_k 65, and top_p 0.66. For Finnish, the final parameters are temperature 0.14, top_k 30, and top_p 0.60.

5 Results

The main results are shown in Table 2. For English, six out of seven models achieve positive improve-

⁷Unlike global sequence alignment, its local counterpart does not penalize leading and trailing misalignments. We use the implementation in the *biopython* package, with open-gap-score -1 and extend-gap-score -0.5

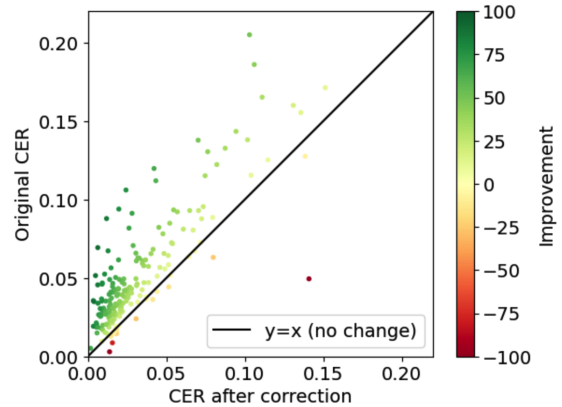


Figure 2: CER before and after correction on English test data (Llama-3.1-70B).

ment, ranging from 7.3% (Llama-3-8B) to 58.1% (GPT-4o) in terms of CER%. GPT-4o outperforms all open models by a large margin, the next best model (Llama-3.1-70B) being almost 20pp worse. However, the Llama-3.1-70B still shows a notable improvement of 38.7%. In Figure 2 we illustrate the CER values for English test examples before and after the Llama-3.1-70B correction. Most examples demonstrate improved CER scores, regardless of whether they initially had mild or severe noise levels.

In terms of WER%, all models show positive improvement on English, the two best models achieving an improvement of 59.1% (GPT-4o) and 46.3% (Llama-3.1-70B). The relative order of the models seems to mostly follow the number of model parameters, bigger models generally performing better, except for Mixtral which is clearly worse than the others, and the two Gemma models performing almost equally in terms of WER%, although the Gemma-2-27B version clearly outperforms the 8B model in terms of CER%.

For Finnish, on the other hand, GPT-4o is the only model achieving a positive improvement in either metric, albeit considerably smaller in absolute terms than for English with 11.9 CER% and 33.5 WER%. Seeing these entirely negative results for Finnish, we are forced to conclude that prompt-based OCR post-correction is presently infeasible for this language using any of the tested open-weight models. This is disappointing, but not surprising since none of the models officially support Finnish.⁸

⁸We made also preliminary experiments with the well-known Finnish Poro model of Luukkonen et al. (2024), but the results were considerably worse than the models in our

OCR: ... mirkavalta on mahmgonnen , mutta ei ole musasta sekään että mallion mallaa ja armoa mähennetään.
 GPT: ... virkavalta on vahingollinen, mutta ei ole viisasta sekään että vallan valtaa ja armoa vähennetään.
 GT : ... wirkawalta on wahingollinen, mutta ei ole wiisasta sekään että waltion waltaa ja arwoa wähennetään.

OCR: Mr. Powell's mind was R.aid upon God in to st.aa dil: a m:;:'.cir, that after wv. had sung an hymn, ...
 GPT: Mr. Powell's mind was stayed upon God in so steadfast a manner, that after we had sung a hymn, ...
 GT : Mr. Powell's mind was ftaid upon God in fo fteadfauft a manner, that after we had fung an hymn, ...

OCR error GPT modernization GPT error

Figure 3: An example in both languages illustrating historical language artifacts alongside the corresponding GPT-4o generated output.

Model	Overg. removal	
	w/o	with
Llama-3-8B	-74.1	7.3
Llama-3.1-8B	-57.4	19.5
Llama-3.1-70B	-53.6	38.7
Gemma-2-9B	28.1	28.2
Gemma-2-27B	35.3	35.6
GPT-4o	53.7	58.1

Table 3: The CER improvement on English test data with and w/o the overgeneration removal step.

The striking effect of a common but meaning-preserving difference between historical and modern language becomes apparent when measuring the effect of modern spelling produced by the LLMs such as the `f` vs. `s` and `w` vs. `v` variation discussed in Section 4.1. A typical example for both languages is illustrated in Figure 3. Without the applied normalization, the results of GPT-4o would have been 34.9 CER% and 35.5 WER% (compared to 58.1% and 59.1% with normalization) for English, and -10.1 CER% and -4.8 WER% (compared to 11.9% and 33.5%) for Finnish. This demonstrates a substantial impact on the reported scores, and while the relative model ranking is unlikely to change, we can see that the conclusion w.r.t. this model’s performance on Finnish would have been the opposite, and the improvements seen in English would have been lot smaller.

Given the entirely negative results for Finnish with the open-weight models, we carry out all further analyses on English only. Furthermore, we also remove the Mixtral-8x7B from follow-up experiments as it performs notably worse than the other models.

5.1 LLM Overgeneration Removal

Next we measure the effect of the alignment-based overgeneration removal method described in Section 4.3, and we did not pursue it any further.

i.e. we evaluate the raw model-generated output against the post-processed version of the generated output. The results are shown in Table 3. For the Llama family models, the results without this post-processing step are highly negative, whereas all Llama models achieve positive improvements when this step is applied. This highlights the necessity of post-processing for the Llama models, which very systematically generate an additional explanation together with the requested output. An example of a typical Llama generation is:

```
Here is the corrected text: {{answer}}
I corrected the following errors:
* "pi\&ure" -> "picture"
* "it's" -> "its" (multiple instances)
* "gneralfcope" -> "general scope"
...
```

On the other hand, Gemma models seem to be largely unaffected, as they generally tend to not produce any additional text. For the GPT-4o model, we also notice a small gain when applying the post-processing, as it occasionally generates explanatory phrases like *"Here is the corrected text:"*.

5.2 Quantization and Performance

Since the historical text collections to which post-correction would potentially be applied comprise millions of pages of text, it is necessary to strike balance between accuracy and computational resources. Among the most important factors here is model quantization, i.e. real number representation with fewer bits. High degrees of quantization substantially reduce model memory footprint and increase inference speed, but can be assumed to potentially degrade model performance. We therefore evaluate the models at the 4 bit Q4_0 quantization (default setting in Ollama), and at the standard 16 bit fp16 floating point representation.

The results are reported in Table 4. As expected, the fp16 quantization performs better for all the evaluated models, with a gain of 2.5-4.7pp, except for Llama-3.1-8B where we do not experience a

Model	CER%		Memory (Gb)	
	q4	fp16	q4	fp16
Llama-3-8B	7.3	12.0	6.3	16.1
Llama-3.1-8B	19.5	19.4	6.3	16.1
Llama-3.1-70B	38.7	42.6	43.6	132.1*
Gemma-2-9B	28.2	30.7	8.9	20.9
Gemma-2-27B	35.6	38.1	19.2	58.9

Table 4: CER improvement on English test data using 4bit quantized (q4) and fp16 models, alongside peak memory usage. * in the memory consumption indicates the number was obtained using the HuggingFace library, as we were not able to run the Llama-3.1-70B model with fp16 through Ollama.

significant difference between 4bit and fp16 models. The relative ranking of the selected models is preserved regardless the quantization level, and using fp16 does not help less performing models to outrank any of the originally best performing 4bit quantized models. The improvement comes at a high cost in terms of memory footprint. As seen in the table, the best improvement is unsurprisingly achieved by the largest model, where the memory requirement increases from 43Gb to 132Gb. It is of consideration that even with 4bit quantization, using the largest Llama-3.1-70B model would necessitate 2 GPUs (assuming 32GB GPU memory), instantly doubling the GPU hours required to complete the task compared to other models which can fit on one GPU.

6 Segment Length and Continuation

The results in the previous sections were reported for text segments about 300 sub-words in length. The actual texts in the historical collections are naturally substantially longer, necessitating splitting the input into segments of appropriate length. This raises two related questions: (1) how long the input segments should optimally be for best post-correction accuracy, and (2) how should the outputs be combined to minimize degradation on segment boundaries.

Our English data is on the level of pages, which we cannot simply naively concatenate, we need other means to obtain sufficiently long documents. For this experiment, we sample long pages of at least 600 whitespace delimited OCR words in length from the development data, taking at maximum two pages from any one book. This resulted in a sample of 53 development set pages.

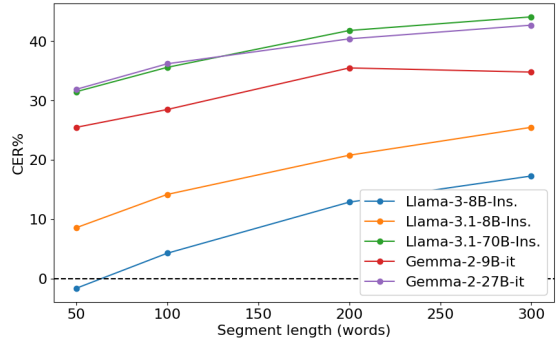


Figure 4: CER% improvement for English when using different segment lengths.

These sampled pages are then divided into non-overlapping segments of 50, 100, 200, and 300 words, using the same alignment-based splitting strategy as described in Section 4. The segments are corrected individually and the CER% improvement over the segments is calculated. The results are shown in Figure 4. Shorter segments (50–100 words) get notably worse CER% score for all models, with the gains diminishing past about 200–300 words, but our page-level data does not have long-enough examples to allow us to reach the point where the performance would start consistently decreasing as the segments become too long. In the future, we plan to develop a book-level version of the data, and study the correction performance on even longer segments.

6.1 Post-correction on Segment Boundaries

Presently, post-correction studies either do not address segment-wise correction of longer texts as it is not necessary for the datasets they study, or split the input into non-overlapping segments, whose corrections are simply concatenated. This may potentially disrupt text continuity since neither word nor sentence boundaries can be reliably adhered to in the noisy OCR input. Furthermore, it also means that no left context is available for the correction of the beginning of each segment. This may potentially have a negative effect on the correction in the region around segment boundary. Here we quantify this effect and explore several straightforward methods for its mitigation.

For the prompt optimization, we use the same sample of 53 pages as in the previous section, and the final results are reported on a similar sample of 50 pages from the test data. In order to maximize the number of examples of segment boundaries for evaluation, we generate—from each page—pairs of

segments 200+200 words long, with stride of 100 words. With this method, each page of at least 600 words produces a minimum of 3 examples of neighboring segment pairs. The final development and test samples include 194 and 208 such examples, respectively.

On these examples, we evaluate the following methods of post-correction on segment boundaries: (1) *Baseline*: Each segment is corrected independently with the same prompt and the outputs are concatenated; (2) *Left-corrected-concatenate (LCC)*: The left segment is corrected first and given as prior context for the correction of the right segment, the model is instructed to only correct the right segment; and (3) *Left-uncorrected-concatenate (LUC)*: The uncorrected left segment is provided for context in the correction of the right segment. The primary advantage of this method is that correcting the right segment does not need the left segment to be corrected first, making parallelization of the process much simpler technically.

The overall results across these strategies are shown in Table 5 and suggest that at present only the two largest models are able to follow the more complex prompts necessitated when merging neighboring segment corrections. The smaller models occasionally suffer from omitting part of the text to correct, which did not seem to occur if only one text was given at a time. For the two models improving on performance, we inspect the boundary effect more closely in Table 6 where we report CER% calculated on ± 10 words around the boundary of the two segments.⁹ Here we see that both methods effectively incorporate the provided additional context, substantially improving the post-correction of the right segment at the boundary, whereas the baseline system’s performance on the right side is notably worse compared to its performance on the left side.

7 Conclusion

We set out to establish the ability of recent open-weight models to post-correct OCR errors in a zero-shot, prompt-based setting. In the first set of experiments, we established that for historical English these models achieved notable improvements (Llama-3.1-70B-Instruct reaching a CER improvement of 38.7%), even though still far behind the

⁹The ± 10 word boundaries were human-verified to ensure that the evaluation occurred at the same boundary, even in more complex examples where words were omitted and/or added.

Model	Method		
	Bas.	LCC	LUC
Llama-3-8B	-0.2	-2.9	-2.8
Llama-3.1-8B	13.7	8.4	10.6
Llama-3.1-70B	33.2	36.0	34.6
Gemma-2-9B	29.2	27.9	28.3
Gemma-2-27B	38.1	39.9	39.7

Table 5: CER improvement on English test data with different correction methods.

Model	Method			
	Baseline		LCC	LUC
	L	R	R	R
Llama-3.1-70B	29.1	9.8	21.4	21.7
gemma-2-27b	34.5	18.7	29.7	33.7

Table 6: CER% around segment boundaries with different correction methods. L and R stand for left and right of the boundary.

commercial GPT-4o model (58.1 CER%). We also demonstrate the necessity of post-processing to remove any additional, model generated text, and present an effective string alignment technique to address this. We also highlight the effect of segment length, which may have a substantial negative impact on the outcome if set too short.

Unlike for English, for Finnish we find poor performance across the board and need to conclude that zero-shot post-correction with open-weight models remains currently out of reach for historical Finnish.

In a separate set of experiments we examine how segment-wise correction of long documents should be approached. We devise and evaluate a number of methods to incorporate additional context for the correction of individual segments. We find that some of these methods have a strong positive effect in the immediate proximity of segment boundaries, however, for smaller models the more complicated prompt may cause unexpected degradation in performance when the whole text is considered. Further work will be necessary to resolve these issues.

As future work, we will pursue a large correction run of the ECCO collection as well as a fine-tuned model for Finnish post-correction. All datasets and evaluation scripts used in this study are available at <https://github.com/TurkuNLP/ocr-postcorrection-lm> to support result replication and comparability.

Limitations

Our work includes certain limitations, which we will discuss next. First, during data preprocessing, we discarded a proportion of documents (~10% for English, ~6% for Finnish) that our correction methods may not be able to address. These documents include cases with severe alignment issues between OCR output and ground truth. We acknowledge that our post-correction method, which relies entirely on the OCR system’s output, cannot recover text where significant portions are missing, therefore setting an upper-boundary for the method. Further analysis is needed to investigate the causes of these gaps and to determine how much, if any, of this missing information could potentially be addressed through post-correction.

We also find that OCR post-correction evaluation suffers from various dataset and metric issues, some of which we have already discussed (e.g. normalization). In related work (including our own study conducted directly on our long-term target corpora), results are reported on varying datasets and evaluation metrics. These challenges make it difficult to achieve comparable results across studies and languages, potentially contributing to some of the contradictory conclusions reported in prior work. Clearly, more work will be needed to establish a set of standard benchmarks that resolve most of the data and evaluation issues.

Finally, reporting pure numeric improvements does not address all aspects of downstream data usability. While an improved word error rate has a direct, positive effect on certain applications (e.g. lexical search), its impact on others (e.g. close reading) may be less straightforward or proportional.

Acknowledgments

This work was carried out in the *Human Diversity* University profilation programme (PROFI-7) of the Research Council of Finland, as well as in the context of several other research projects supported by the Research Council of Finland. Computational resources were provided by CSC — the Finnish IT Center for Science.

References

- AI@Meta. 2024a. https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md Llama 3 model card.
- AI@Meta. 2024b. https://github.com/meta-llama/llama-models/blob/main/models/llama3_1/MODEL_CARD.md Llama 3.1 model card.
- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Angel Beshirov, Milena Dobрева, Dimitar Dimitrov, Momchil Hardalov, Ivan Koychev, and Preslav Nakov. 2024. Post-OCR text correction for Bulgarian historical documents. *ArXiv preprint arXiv:2409.00527*.
- Emanuela Boros, Maud Ehrmann, Matteo Romanello, Sven Najem-Meyer, and Frédéric Kaplan. 2024. Post-correction of historical text transcripts with large language models: An exploratory study. In *Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2024)*, pages 133–159. Association for Computational Linguistics.
- Jonathan Bourne. 2024. CLOCR-C: Context leveraging OCR correction with pre-trained language models. *ArXiv preprint arXiv:2408.17428*.
- Yung-Hsin Chen and Phillip B. Ströbel. 2024. TrOCR meets language models: An end-to-end post-correction approach. In *Proceedings of the Document Analysis and Recognition – ICDAR 2024 Workshops*, pages 12–26. Springer Nature Switzerland.
- Guillaume Chiron, Antoine Doucet, Mickaël Coustaty, and Jean-Philippe Moreux. 2017a. ICDAR 2017 competition on post-OCR text correction. In *Proceedings of the 14th IAPR International Conference on Document Analysis and Recognition (ICDAR 2017)*, volume 1, pages 1423–1428. IEEE.
- Guillaume Chiron, Antoine Doucet, Mickael Coustaty, Muriel Visani, and Jean-Philippe Moreux. 2017b. Impact of OCR errors on the use of digital libraries: Towards a better access to information. In *Proceedings of the 2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 1–4.
- Oksana Dereza, Deirdre Ní Chonghaile, and Nicholas Wolf. 2024. “To have the ‘million’ readers yet”: Building a digitally enhanced edition of the bilingual Irish-English newspaper *An Gaodhal* (1881-1898). In *Proceedings of the Third Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA) @ LREC-COLING-2024*, pages 65–78. ELRA and ICCL.
- Quan Duong, Mika Hämäläinen, and Simon Hengchen. 2021. An unsupervised method for OCR post-correction and spelling normalisation for Finnish. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 240–248. Linköping University Electronic Press.

- Ali Fahandari, Fatemeh Asadi Zeydabadi, Elham Shabaninia, and Hossein Nezamabadi-pour. 2024. Enhancing Farsi text recognition via iteratively using a language model. In *Proceedings of the 20th CSI International Symposium on Artificial Intelligence and Signal Processing (AISP)*.
- Gale. <https://www.gale.com/intl/primary-sources/eighteenth-century-collections-online> Eighteenth Century Collections Online.
- Stephen H. Gregg. 2021. *Old Books and Digital Publishing: Eighteenth-Century Collections Online*. Elements in Publishing and Book Culture. Cambridge University Press.
- Stephen H. Gregg. 2022. The nature of ECCO-TCP. *Digital Defoe: Studies in Defoe & His Contemporaries*, 14(1).
- Mark J. Hill and Simon Hengchen. 2019. Quantifying the impact of dirty OCR on historical text analysis: Eighteenth Century Collections Online as a case study. *Digital Scholarship in the Humanities*, 34(4):825–843.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Kimmo Kettunen, Jukka Kervinen, and Mika Koistinen. 2018. Creating and using ground truth OCR sample data for Finnish historical newspapers and journals. In *Proceedings of the Digital Humanities in the Nordic Countries Conference*.
- Kimmo Kettunen, Mika Koistinen, and Jukka Kervinen. 2020. Ground truth OCR sample data of Finnish historical newspapers and journals in data improvement validation of a re-OCRing process. *LIBER Quarterly: The Journal of the Association of European Research Libraries*, 30(1):1–20.
- Kimmo Kettunen and Tuula Pääkkönen. 2016. Measuring lexical quality of a historical Finnish newspaper collection — analysis of garbled OCR data with basic language technology tools and means. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 956–961. European Language Resources Association (ELRA).
- Minghao Li, Tengchao Lv, Jingye Chen, Lei Cui, Yijuan Lu, Dinei Florencio, Cha Zhang, Zhoujun Li, and Furu Wei. 2023. TrOCR: Transformer-based optical character recognition with pre-trained models. In *Proceedings of The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23)*.
- Risto Luukkonen, Jonathan Burdge, Elaine Zosa, Arne Talman, Ville Komulainen, Väinö Hatanpää, Peter Sarlin, and Sampo Pyysalo. 2024. Poro 34b and the blessing of multilinguality. *ArXiv preprint arXiv:2404.01856*.
- Gábor Madarász, Noémi Ligeti-Nagy, András Holl, and Tamás Váradi. 2024. OCR cleaning of scientific texts with LLMs. In *Natural Scientific Language Processing and Research Knowledge Graphs*, pages 49–58. Springer Nature Switzerland.
- Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, et al. 2024. <https://www.kaggle.com/m/3301> Gemma.
- Thi Tuyet Hai Nguyen, Adam Jatowt, Mickael Coustaty, and Antoine Doucet. 2021. Survey of post-OCR processing approaches. *ACM Comput. Surv.*, 54(6).
- Robert Östling, Katarina Gillholm, Murathan Kurfalı, Marie Mattson, and Mats Wirén. 2024. Evaluation of really good grammatical error correction. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 6582–6593, Torino, Italia. ELRA and ICCL.
- Maria Carolina Penteado and Fábio Perez. 2023. Evaluating GPT-3.5 and GPT-4 on grammatical error correction for Brazilian Portuguese. In *Proceedings of the LatinX in AI Workshop at ICML 2023*.
- Christophe Rigaud, Antoine Doucet, Mickaël Coustaty, and Jean-Philippe Moreux. 2019. ICDAR 2019 competition on post-OCR text correction. In *Proceedings of the 15th International Conference on Document Analysis and Recognition (ICDAR 2019)*, pages 1588–1593. IEEE.
- Elizabeth Soper, Stanley Fujimoto, and Yen-Yun Yu. 2021. BART for post-correction of OCR newspaper text. In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 284–290. Association for Computational Linguistics.
- Mikko Tolonen, Eetu Mäkelä, Ali Ijaz, and Leo Lahti. 2021. Corpus linguistics and Eighteenth Century Collections Online (ECCO). *Research in Corpus Linguistics*, 9:19–34.
- Martijn Veninga. 2024. LLMs for OCR post-correction. Master’s thesis, University of Twente.