Preparing Big Manuscript Data for auto Collation with minimal HTR training

Elpida Perdiki

Journal of Data Mining and Digital Humanities

EPIsciences overlay journals

Preparing Big Manuscript Data for (auto) Collation with Minimal HTR Training

Elpida Perdiki[1 *]

1 Department of Greek Philology, Democritus University of Thrace, Greece

*Corresponding author: eperdiki@helit.duth.gr

Abstract

HTR (Handwritten Text Recognition) technologies have progressed enough to offer high-accuracy results in recognising handwritten documents, even on a synchronous level. Despite the state-of-the-art algorithms and software, historical documents (especially those written in Greek) remain a real-world challenge for researchers. A large number of unedited or under-edited works of Greek Literature (ancient or byzantine, especially the latter) exist to this day due to the complexity of producing critical editions. To critically edit a literary text, scholars need to pinpoint text variations on several manuscripts, which requires fully (or at least partially) transcribed manuscripts. For a large manuscript tradition (i.e., a large number of manuscripts transmitting the same work), such a process can be a painstaking and time-consuming project. To that end, HTR algorithms that train AI models can significantly assist, even when not resulting in entirely accurate transcriptions. Deep learning models, though, require a quantum of data to be effective. This, in turn, intensifies the same problem: big (transcribed) data require heavy loads of manual transcriptions as training sets. In the absence of such transcriptions, this study experiments with training sets of various sizes to determine the minimum amount of manual transcription needed to produce usable results. HTR models are trained through the Transkribus platform (transkribus.eu) on manuscripts from multiple works of a single Byzantine author, John Chrysostom. By gradually reducing the number of manually transcribed texts and by training mixed models from multiple manuscripts, economic transcriptions of large bodies of manuscripts (in the hundreds) can be achieved. Results of these experiments show that if the right combination of manuscripts is selected, and with the transfer-learning tools provided by Transkribus, the required training sets can be reduced by up to 80%. Certain peculiarities of Greek manuscripts, which lead to easy automated cleaning of resulting transcriptions, could further improve these results. The ultimate goal of these experiments is to produce a transcription with the minimum required accuracy (and therefore the minimum manual input) for text classification. Less could be more if we can correctly evaluate HTR learning and results. This case study proposes a solution for researching/editing authors and works that were popular enough to survive in hundreds (if not thousands) of manuscripts and are, therefore, unfeasible to be evaluated by humans.

keywords

big data; byzantine manuscripts; deep learning; HTR models; Transkribus;

INTRODUCTION

The humanitarian spirit of Antiquity and Byzantium has passed down to younger generations a multitude of manuscripts that preserve ancient and byzantine Greek literary texts. Many of these manuscripts remain unedited or under-edited due to the complexity of producing critical editions. This process requires heavy loads of manuscript research until all disparate text variations (instances of manuscripts that contain the same opus transmitting different text) are collected and collated to the last detail. Especially in cases of rich manuscript traditions (i.e., a significant number of manuscripts transmitting the same work), this process is not only

tedious, but it could also take years to complete. Therefore, some otherwise well-known authors remain archived in libraries or under poorly edited publications. Such is the case of the ~880 manuscripts of Homer, ~3,991 of the New Testament, or ~21,482 of John Chrysostom's opera (as currently listed in [*Pinakes*, n.d.], yet those numbers might be even higher).

HTR technology could greatly assist the monumental task of massively and accurately collating hundreds, if not thousands, of manuscripts. After all, no collation can be done without a diplomatic transcription of the sources. Although HTR systems have evolved significantly in the last decades, and despite the several state-of-the-art systems readily available to the scholars' community, the peculiarities of handwritten historical documents remain a real challenge. This phenomenon is especially true for documents written in Ancient Greek, in which special characters, such as accents (at least five unique characters for accents and six combinations of those) and ligatures or abbreviations of letters, perplex character recognition even more. Apart from these factors, HTR algorithms –assisted by AI neural networks and can therefore train highly accurate models– require a quantum of training data to be effective. In order to produce these input data, one should return to the same old process: manual transcription of a bulk of manuscripts.

Due to the scarcity of these transcriptions and the human efforts millstone of producing them ex nihilo, this paper examines the limits of HTR technology by defining the optimum amount of data needed to train an AI model successfully. In the second phase of this research (currently ongoing), the HTR-produced transcriptions are tested further as data input in experiments of manuscripts auto-collation (yet not for stemmatical analysis). The aim is to produce a classification system by which all instances of a text can be traced back to their ancestors through a series of branching points –much like the phylogenetics method in Biology, but with DNA sequences replaced by manuscripts [Macé and Baret, 2004; Spencer *et al.*, 2004].

For all the following methodological experiments, a set of 11 manuscripts with homilies of John Chrysostom served as a case study on HTR and manuscript collation testing. This author was chosen for two main reasons: a) his opera are numbered at ~21,482 manuscripts, which is equal to almost half a million words –unfeasible, thus, for humans to transcribe–, and b) almost 3,000 of these manuscripts are known for the *double recension* phenomenon, simply meaning that there are at least two prominent manuscript families, known as recensions, from which one is the revision of the other [Konstantinidou, 2021; Perdiki and Konstantinidou, 2021]. Thus, to classify these thousands of manuscripts into the relevant recension, one should first extract the raw text data from the manuscripts. For both tasks, exploitation of pertinent technology seems necessary to rapidly and massively handle the bulk of data.

HTR experiments were conducted on the Transkribus platform.

I LITERATURE REVIEW

Text recognition systems are well-researched and continuously developing. Currently, there are two main systems for image text extraction: OCR (Optical Character Recognition), targeting printed text, and HTR for handwritten documents, where character recognition is not straightforward [Firmani *et al.*, 2018; Ströbel *et al.*, 2022]. HTR is furtherly divided into offline (meaning recognition from a scanned document image) or online (text recognition

while the text is being written) [Ingle *et al.*, 2019]. Furthermore, the ever-increasing need for transcription of historical documents, currently archived in libraries and collections worldwide, has led to the development of HTR systems, mainly focused on ancient or medieval handwriting.

The most recent bibliography suggests applications such as Tesseract [Patel *et al.*, 2012; White, 2012; *Tesseract Open Source OCR Engine*, n.d.],[1] TensorFlow [Pang, Nijkamp and Wu, 2020; TensorFlow Developers, 2021], Kraken [Schoen and Saretto, 2022; *kraken — OCR system*, n.d.], eScriptorium [Kiessling *et al.*, 2019], Transkribus [Kahle *et al.*, 2017; Muehlberger *et al.*, 2019], or μDoc [Tsochatzidis *et al.*, 2021]. Despite being listed in line here, it should be noted that the above methods cannot be measured reliably against each other measured against each other, as they are diverse in architecture and function (i.e., Kraken is an OCR engine, Tensorflow a library, while eScriptorium and Transkribus are interface platforms). Regardless, already conducted experiments [Ströbel and Clematide, 2019; Ströbel, Clematide and Volk, 2020; Clérice, 2022] have demonstrated that from the HTR mentioned above tools, Transkribus, Kraken, and e-Scriptorium (which implements Kraken) are the most successful in producing low CER (Character Error Rate) text recognitions. Part of this success is due to the different and more efficient layout analysis performed by both Transkribus and eScriptorium, an analysis that does not restrict segmentation in rectangular regions, as handwritten text can expand in many forms and directions [Stokes *et al.*, 2021]. However, since only some Humanities scholars are tech savvy, it was decided to exploit the Transkribus system for experimenting with Greek manuscripts HTR. The reason behind this decision was twofold. First of all, while all other systems are executed via CLI (Command Line Interface), assuming coding fluency, Transkribus is offered as a GUI (Graphical User Interface) and Web-based application,[2] making it accessible to most researchers [Transkribus Expert Client, 2021; Transkribus Lite, n.d.]. Currently, eScriptorium also offers a Web-based platform upon registration and further contact with their team [Kiessling *et al.*, 2019; Stokes *et al.*, 2021]. However, their interface was partially developed when this research began experiments. That being said, we intend to expand our training methods with the exploitation of eScriptorium in the near future. Secondly, eScriptorium/Kraken needs high computational power in order to train models [Stokes *et al.*, 2021]. Unfortunately, not many scholars have access to high-performing hardware. On the contrary, since Transkribus is connected to the Innsbruck server, all computations and training are performed there [Kahle *et al.*, 2017; Muehlberger *et al.*, 2019, pp. 959, 962]. So, each user can train models even from a low-cost laptop.

II METHODS OF MANUSCRIPTS AUTO-TRANSCRIPTION

1.1 Data Availability

As described previously, the 11 case study manuscripts were used as training data. The manuscripts are dated to the 10th-14th century and transmit John Chrysostom's *Homilies on St. Paul's Epistles to Titus*. Homilies 1 and 5 were used as data sets in all experiments conducted based on the availability of digital images.

[1] Nevertheless, there are also some voices of opposition regarding the results of Tesseract compared to similar systems, as stated in [Smith, 2007].

[2] Currently, eScriptorium offers a Web-based platform upon registration and further contact with the eScriptorium team.

Most digital reproductions of manuscripts are under some degree of copyright protection. As a result, data availability is not always a straightforward process. Quite the opposite, the author produced ground truth data from scratch for the following experiments, which are fully provided to the research community. Instead of the digital files of the manuscripts, links to the libraries' digital archives are given, if applicable. The detailed dataset can be found at [Perdiki, 2023]. Finally, from all the trained models, the general model, the one with the better score, will be released in Transkribus soon (after some needed fine-tuning).

1.2 Methodology

Transkribus documentation denotes that for a successful HTR model training, at least 15,000 words of diplomatic transcription input are required. However, since such transcriptions are unavailable, producing them from scratch would demand heavy economic and human resources. Early experiments on Transkribus indicated that most erroneous outputs involved misrecognition of accents, punctuation or word tokens splitting (due to the *scripta continua* form of the writing style), see Figure 1. Most of the time, the last character of a word token (usually pronouns or conjunctions with an average of three characters in length) was erroneously connected with the following or the previous word token. In addition, probably for the same reason,[3] when tokenisation fails, so does accent recognition. See, i.e., line 5 of Figure 1, where instead of *μὴ νεκροὺς* HTR recognises *μὴν ἐκροὺς* and so adds the smooth breathing diacritical mark above the letter *ε*. This diacritics addition is an interesting mistake; according to ancient Greek grammar, when a word begins with the letter *ε*, it commonly has the smooth breathing mark. Transkribus is learning fast. Other times, accent recognition fails because accents are already misplaced in the manuscripts. For instance, cf. Figure 2 and line 3 of Figure 1, where the grave accent of the adjective *πολὺς* is not recognised since it is misplaced above the last character, the consonant *ς*. It should also be mentioned that accent misrecognition might be negatively affected by the non-expansion of the line region to the upper margin (see Figure 2; the line region is depicted with the blue rectangle and baseline with the purple underline). Following Transkribus guidelines regarding the higher importance of baseline region in HTR results [Muehlberger *et al.*, 2019, p. 959; *Line region*, n.d.], adjustments were made only on the relevant baselines in all mentioned experiments. This inaction on line-region adjustment proves to be a limitation and should be corrected in future experiments.

Such errors, as mentioned above, can be easily cleaned to lower CER significantly. Moreover, despite successful OCR demands of 90% accuracy, the complex peculiarities of HTR allow for a lower level of accuracy, especially if we consider that keyword spotting techniques return accurate results even with a 30% CER [Muehlberger *et al.*, 2019, p. 963; Tomoiaga *et al.*, 2019; Stokes *et al.*, 2021; Ströbel *et al.*, 2022]. Consequently, aiming for better HTR results, a maximum 20% CER threshold was decided for a model to be deemed adequately accurate (i.e., to serve as data for text classification).[4] That being said, it should be pointed out that this threshold must decided upon firm grounds, so current benchmarking experiments with HTR results to be published in the near future are adapting a different methodology to test the limits of such a metric accurately. Although important in evaluation [Sánchez *et al.*, 2019, p. 124; Kang *et al.*, 2022], using WER instead of CER was ruled out due to the errors

[3] Although, extensive experiments should be conducted before a conclusion is made on the matter.

[4] Of course, this percentage of error rate is not an ideal aim for HTR performance. Yet, it is a cost-efficient choice in order to balance low input data and high output accuracy.

described in the previous section (i.e., a wrong accent in a letter despite being only one erroneous character will return a WER, a result which is inaccurate if such errors are normalised). Lastly, previous research concluded that, although AI machine learning algorithms require a quantum of data to be effective, there is certainly a limit to the data set volume or the training epochs number in order to avoid overfitting [Rabus, 2019; Ströbel, Clematide and Volk, 2020; Perdiki and Konstantinidou, 2021].

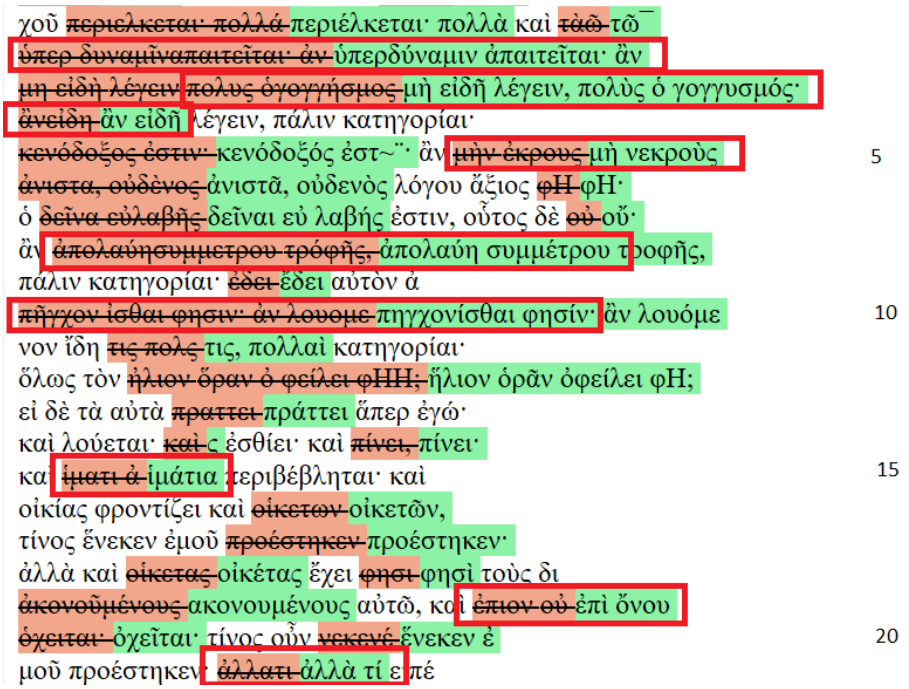

Figure 1. Highlighted in red rectangles are instances of HTR errors due to scripta continua. I.e., line 21 depicts how tokens *ἀλλὰ τί* (transl.: but what) are falsely recognised as the single (meaningless) token *ἀλλατι*.

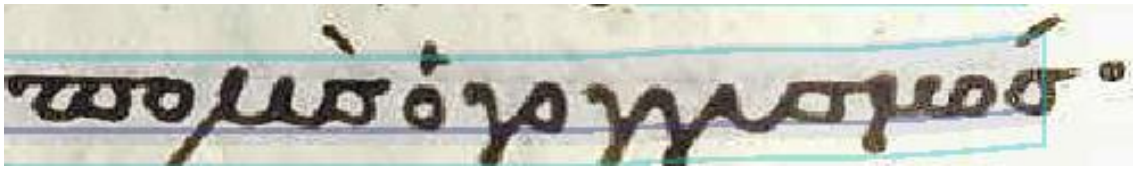

Figure 2. An instance of erroneous accent recognition. The depicted text should be transcribed as *πολὺς ὁ γογγυσμός*.

III RESULTS

Upon these three criteria, experiments were conducted under four main methods and mainly under the same system configurations (reported in Table 1). The first method was HTR model training with gradual data set reduction to define the minimum amount of data needed to produce usable results. As depicted in Figure 3, 24 models were trained (via CITlab HTR+ method) from 8 different manuscripts (three models per manuscript),[5] with a decreasing number of words: transcription input of ~3,000, ~2,000 and ~1,000 words from John Chrysostom's *1st Homily*, with a minimum of 50 epochs of each training set. A 10% portion of the data input was reserved in each training set as validation data. Most models performed below the 20% CER threshold, even under the low 1,000-word input test. The few exceptions of poor recognition results overlapped with some low-quality manuscript digitisations. Usually, the breaking point of the model training was around 5-10 epochs.

	Training Set (words)	**Validation Set**	**Epochs[6]**	**Early Stopping[7]**	**Base model**	**Training system**
Single models I	3,000	10%	50	N/A	No	HTR+
	2,000	10%	50	N/A	No	HTR+
	1,000	10%	50	N/A	No	HTR+
Mixed models	2,000	10%	50	N/A	No	HTR+
General model I	9,000	10%	50	N/A	Yes	HTR+
Single models II	3,000	10%	50	20	No	PyLaia
	3,000	10%	250	20	No	PyLaia
General model II	25,621	10%	50	20	Yes	PyLaia
	25,621	10%	250	20	Yes	PyLaia

Table 1. System configurations for all experiments

[5] Two manuscripts of the dataset, E and W, were excluded from this testing due to their insufficient number of words.

[6] In most cases, the number of epochs was set up to 50. This fixed choice was an attempt to standardise configurations so as to clarify which factors affect the HTR performance, i.e., the number of training data, the number of epochs, the quality of the image data or otherwise.

[7] The early stopping technique is not applicable in the HTR+ method (the system of choice for the majority of the experiments). PyLaia, on the contrary, does provide that option, so early stopping was set to the predefined number of 20. Since Transkribus had announced that HTR+ would not be supported any further than November 2022 [HTR+, no date], the exploitation of both systems was decided as a form of systems performance comparison.

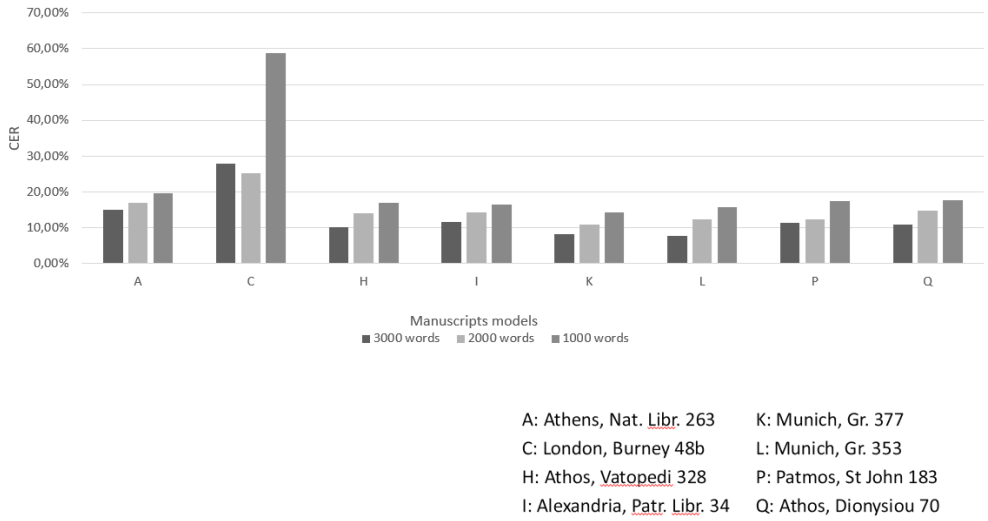

Figure 3. CER results of decreasing training data (number of words of manually generated training data), as in [Perdiki and Konstantinidou, 2021].

With the aim of testing script similarity out of text recognition and limiting the manual production of data input even further, a cyclical application of each trained model to 10 manuscripts was performed. The experiment hypothesis of this method was whether an already trained model could accurately recognise the text of a different but similar writing style manuscript. This process would also serve as a manuscript clustering method if proven successful. However, as seen in Figure 4, the resulting 90 text recognitions were mainly inaccurate. Only 9 out of 90 combinations recovered text with a lower than 20% CER, despite exploiting the 3,000-word input training sets. As a result, since these nine successful applications were not apparent in advance as similar writing styles, clustering algorithms that would predict script similarity seem necessary.

A: Athens, Nat. Libr. 263
C: London, Burney 48b
E: Paris, Bibl. Nat., Gr. 745
H: Athos, Vatopedi 328
I: Alexandria, Patr. Libr. 34
K: Munich, Gr. 377
L: Munich, Gr. 353
P: Patmos, St John 183
Q: Athos, Dionysiou 70
W: Munich, Gr 211

Manuscripts/ HTR Models	Q	H	E	L	K	W	C	A	P	I
Q	10.4%	34.69%	72.69%	17.27%	14.33%	27.43%	25,.3%	12.16%	25.93%	15.51%
H	52.13%	10.03%	78.01%	32.89%	47.67%	46.67%	67.17%	38.46%	61.06%	39.41%
E	73.28%	94.12%	74.66%	67.23%	83.29%	73.68%	72.96%	77.18%	83.66%	68.99%
L	28.75%	33.79%	74.21%	7.72%	31.74%	41.59%	37.33%	36.84%	41.60%	32.65%
K	27.73%	27.73%	75.58%	22.90%	8.23%	39.84%	57.14%	30.14%	55.05%	32.74%
W	25.82%	38.82%	73.94%	24.13%	26.02%	31.31%	47.10%	24.62%	44.37%	23.21%
C	48.69%	55.14%	79.66%	36.41%	48.04%	55.54%	27.83%	47.96%	47.94%	47.45%
A	16.12%	38.67%	72.43%	26.59%	17.88%	30.61%	36.49%	14.20%	30.27%	14.68%
P	26.14%	46.35%	73.78%	24.37%	27.89%	34.35%	30.05%	19.20%	11.29%	17.78%
I	25.29%	44.51%	73.22%	29.01%	28.31%	36.10%	44.62%	25.42%	31.53%	11.62%

Figure 4. Cyclic application of each model to all manuscripts, as in [Perdiki and Konstantinidou, 2021].

The third method of experimenting with HTR extended the second method's hypothesis. If one mixes training data from more than one manuscript, as a data augmentation process, the data set will be enlarged without demanding extra manual input. So, a hypothesis was made to test whether this would result in mixed models of high accuracy. Furthermore, combining all transcriptions in a single training set would test the possibility of building an optimum model capable of accurately transcribing any of our Greek manuscripts. The nine best matches, produced under the second method experiments, served as the data for the manuscripts' combinations. The data set was formed out of randomly selected pages from each manuscript. These combined data consisted of ~2,000-word input transcription per combination. Each model was trained with 50 epochs via the CITlab HTR+ method. The validation set was formed out of a random 10% of the training data. Afterwards, the trained mixed models were applied to each of the training set's manuscripts, as in Figure 5 (i.e., the Q&L model was trained from Q and L manuscripts' combined data and then applied to each for text recognition). These models performed at 80% accuracy, within threshold limits, with a breaking point around 5-10 epochs.[8] Lastly, a 9,000-word input from all 10 joined manuscripts, trained on 50 epochs and CITlab HTR+ method, validated from a random 10% of the data, and applied to every single manuscript, resulted in top-end CER performance (down to 4,48%, see Figure 5).

A: Athens, Nat. Libr. 263
C: London, Burney 48b
E: Paris, Bibl. Nat., Gr. 745
H: Athos, Vatopedi 328
I: Alexandria, Patr. Libr. 34
K: Munich, Gr. 377
L: Munich, Gr. 353
P: Patmos, St John 183
Q: Athos, Dionysiou 70
W: Munich, Gr 211

Manuscripts/ HTR models	Q & L	Q & K	Q & W	Q & A	I & A	I & P	I & Q	A & K	SINGLE MODEL
Q	13.18%	10.89%	14.19%	11.04%			13.31%		4.48%
H									7.51%
E									25.91%
L	59.89%								5.91%
K		10.82%						11.21%	9.5%
W			16.4%						17.04%
C									16.23%
A				11.87%	16.46%			15.41%	12.33%
P						11.62%			11.4%
I					14.63%	13.17%	12.26%		11.4%

Figure 5. CER of models with mixed training sets, as in [Perdiki and Konstantinidou, 2021].

[8] The poor performance of the Q+L model, when applied to the L manuscript, has yet to be fully explained. Apart from the writing style, the only difference between the two manuscripts was that Q's data set was coloured digitisation, whereas L's was grey-scaled microfilm digitisation. However, that was also the case with the K manuscript, yet the relevant CERs returned under 20%. On the other hand, E's bad score is probably due to the insufficient amount of data and the poor image quality; this particular manuscript had a missing folio which drastically affected the HTR results. Its inclusion in the data set was deliberate to test the degree of effect such manuscript damages might have on our methodology.

The fourth and last methodology on HTR experiments via the Transkribus platform came as validation. With the same data set of the 11 manuscripts mentioned above and under the same methodology, the most successful of the above experiments (the first and third method) was performed on a different training set (John Chrysostom's *5th Homily* transcription) in two testing phases. Firstly, the training set consisted of ~3,000-word input and the HTR method was altered to PyLaia with 250 epochs. The validation set was a random 10% portion of the training data. The resulting CER was lower than 10% (breaking point on 20-30 epochs), yet higher compared to previous experiments with 50 epochs of training. In addition to the CITlab HTR+ method's better performance, it appears that, with insufficient data, a higher epoch number returns the worst results, as already shown by Rabus [Rabus, 2019]. Secondly, another attempt was made to build a general model, as all manuscripts are characterised by a certain (perceived) script uniformity and clarity (none of the manuscripts was heavy in ligatures, abbreviations, or damaged areas), however unique in writing style. By joining up the *5th Homily's* transcriptions (from 9 out of the 11 manuscripts), a 25,621-word input trained the general model. In an attempt to further augment the data of the experiments, fine-tune the model, and improve, thus, the recognition results, the model with the best performance – during phase 1 of the fourth method experiments– was added to the training process as a base model. The final CER on that last experiment was 0.60% on the training set, which translates to 3.90% on the validation set (at 3-10 epochs range breaking point), the minimum CER of all conducted experiments (see Table 2).

Training Data	**Model Name**	**50 epochs CER (HTR+)**	**250 epochs CER (PyLaia)**
Q: Athos, Dionysiou 70	Q-3000[9]	11.62%	14.41%
H: Athos, Vatopedi 328	H-3000	10.03%	13.70%
A: Athens, Nat. Libr. 263	A-3000	10.03%	12.20%
I: Alexandria, Patr. Libr. 34	I-3000	13.00%	14.60%
D: Venice, ONB theol. gr.14	D-3000	(n/a)[10]	14.20%
E: Paris, Bibl. Nat., Gr. 745	E-3000	(n/a)	14.90%
K: Munich, Gr. 377	K-3000	8.93%	12.30%
L: Munich, Gr. 353	L-3000	8.12%	13.00%
General Model	GM	17.18%	3.90%

Table 2. CER of 50 and 250 epochs models.

Limitations

The end-all of the conducted experiments was to determine the minimum amount of data needed to produce an accurate transcription and to deliver data usable in text classification. The classification ought not to be a hierarchical clustering as stemmatological analysis is out of the scope of this research. Current experiments with data mining techniques prove to be promising for the task. However, since data quality evaluation for text classification needs to be further fine-tuned, results on the matter will be published in a separate future paper.

[9] The "3000" tag indicates the amount of word input on the training data set.

[10] Manuscripts D and E were excluded from the 50 epochs training due to insufficient number of training data.

ISSN 2416-5999, an open-access journal

As pointed out previously in this paper, the metrics used in our experiments should be further adjusted. Although the 20% CER threshold has been proven adequate so far for successful classification of automated transcriptions, benchmarking experiments may further lower the required accuracy. Furthermore, the uneven distribution of errors on each page probably affects the erroneous HTR results, which should also be accounted for in future testing. More than that, the low number of input data equals insufficient data on the validation set. These experiments produce seemingly unexpected results. In other words, since our case study was a limit test, traditional metrics cannot always be applied. We seek to find more suitable evaluation techniques as we continue our experiments.

The aforementioned dual approach to early stopping (none for HTR+ and 20 for the PyLaia method) might prove to be a limitation on our results. As already explained, both systems were used to compare performance fairly. Nevertheless, since early stopping is not available for the HTR+ training, perhaps the reasonable choice would be to nullify PyLaia early stopping. The systems configurations need further improvement on the matter.

Conclusion

Computational processes can highly assist philological research when dealing with a bulk of data. Time-consuming and painstaking tasks, often leading to errors due to their complexity, produce fruitful results when conducted via special algorithms. This paper presented methodologies for exploiting a specific HTR tool to enhance manuscript tradition research.

The Transkribus platform proved highly efficient in training HTR models and recognising text from digitised manuscripts. Even with minimal training data input, the accuracy of the produced models was high. With further testing and fine-tuning, developing general models that could transcribe a good portion of Greek manuscripts is more than possible. Mass transcription from historical documents can fuel the research with much-needed data. Ongoing experiments, planned to be published soon, are testing whether it is possible to perform algorithmic classification even with some inaccuracy in HTR results.

Indeed, machine learning benefits from data plethora, but sometimes data augmentation and manipulation can produce functioning results. According to the research questions, humans can evaluate the process and fine-tune algorithms to high performance. By outsourcing tedious and prone to errors tasks to computing power and accuracy, researchers can concentrate on more analytical quests and lead the way forward.

References

Clérice, T. You Actually Look Twice At it (YALTAi): using an object detection approach instead of region segmentation within the Kraken engine. arXiv, 2022. Available at: https://doi.org/10.48550/arXiv.2207.11230.

Transkribus Expert Client, 2021. Available at: http://web.archive.org/web/20211113063459/https:/readcoop.eu/transkribus/download/ (Accessed: 19 July 2022).

Firmani, D. *et al.* Towards Knowledge Discovery from the Vatican Secret Archives. In Codice Ratio -- Episode 1: Machine Transcription of the Manuscripts. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018:263–272. Available at: https://doi.org/10.1145/3219819.3219879.

HTR+ READ-COOP, n.d.. Available at: https://readcoop.eu/glossary/htr-plus/ (Accessed: 24 February 2023).

Ingle, R.R. *et al.* A Scalable Handwritten Text Recognition System. *2019 International Conference on Document Analysis and Recognition (ICDAR)*, 2019:17–24. Available at: https://doi.org/10.1109/ICDAR.2019.00013.

Kahle, P. *et al.* Transkribus - A Service Platform for Transcription, Recognition and Retrieval of Historical Documents. *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, Kyoto: IEEE, 2017:19–24. Available at: https://doi.org/10.1109/ICDAR.2017.307.

Kang, L. *et al.* Pay attention to what you read: Non-recurrent handwritten text-Line recognition. *Pattern Recognition*, 2022;129:108766. Available at: https://doi.org/10.1016/j.patcog.2022.108766.

Kiessling, B. *et al.* eScriptorium: An Open Source Platform for Historical Document Analysis. *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, 2019:19–19. Available at: https://doi.org/10.1109/ICDARW.2019.10032.

Konstantinidou, M. The Double Tradition of John Chrysostom's Exegetical Works: Revisions Revisited. M. Vinzent, G. Bady, and C. Broc-Schmezer (eds) *Studia Patristica. Vol. CXIV - Papers presented at the Eighteenth International Conference on Patristic Studies held in Oxford 2019*. Peeters Publishers (Volume 11: John Chrysostom through Manuscripts, Editions and History), 2021:5–26. Available at: https://doi.org/10.2307/j.ctv27vt5gb.4.

kraken — OCR system, n.d. Available at: http://kraken.re/ (Accessed: 18 September 2020).

Line region READ-COOP, n.d. Available at: https://readcoop.eu/glossary/line-region/ (Accessed: 13 February 2023).

Macé, C. and Baret, P.V. Why Phylogenetic Methods Work : The Theory of Evolution and Textual Criticism. *Linguistica computazionale* [Preprint], 2004:(24/25). Available at: https://doi.org/10.1400/54380.

Muehlberger, G. *et al.* Transforming scholarship in the archives through handwritten text recognition: Transkribus as a case study. *Journal of Documentation*, 2019;75(5):954–976. Available at: https://doi.org/10.1108/JD-07-2018-0114.

Pang, B., Nijkamp, E. and Wu, Y.N. Deep Learning With TensorFlow: A Review. *Journal of Educational and Behavioral Statistics*, 2020;45(2):227–248. Available at: https://doi.org/10.3102/1076998619872761.

Patel, C., Patel, A. and Patel, D. Optical Character Recognition by Open source OCR Tool Tesseract: A Case Study. *International Journal of Computer Applications*, 2012;55(10):50–56. Available at: https://doi.org/10.5120/8794-2784.

Perdiki, E. and Konstantinidou, M. Handling Big Manuscript Data. *Classics@*. Edited by C. Clivaz and V.A. Garrick, 2021;18(Ancient Manuscripts and Virtual Research Environments, special issue). Available at: https://classics-at.chs.harvard.edu/classics18-perdiki-and-konstantinidou/ (Accessed: 21 January 2022).

Perdiki, E. List of manuscripts containing John Chrysostom's Homilies and the relevant manual transcriptions (Version 1) [Data set]. Zenodo, 2023. https://doi.org/10.5281/zenodo.7681133

Pinakes, Institut de Recherche et d'histoire des textes, n.d. Available at: https://pinakes.irht.cnrs.fr/ (Accessed: 8 October 2020).

Rabus, A. Recognizing handwritten text in Slavic manuscripts: A neural-network approach using Transkribus. *Scripta & e-Scripta*, 2019;19:9–32.

Sánchez, J.A. *et al.* A set of benchmarks for Handwritten Text Recognition on historical documents. *Pattern Recognition*, 2019;94:122–134. Available at: https://doi.org/10.1016/j.patcog.2019.05.025.

Schoen, J. and Saretto, G.E. Optical Character Recognition (OCR) and Medieval Manuscripts: Reconsidering Transcriptions in the Digital Age. *Digital Philology: A Journal of Medieval Cultures*, 2022;11(1):174–206. Available at: https://doi.org/10.1353/dph.2022.0010.

Smith, R. An Overview of the Tesseract OCR Engine. *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, 2007:629–633. Available at: https://doi.org/10.1109/ICDAR.2007.4376991.

Spencer, M. *et al.* Phylogenetics of artificial manuscripts. *Journal of Theoretical Biology*, 2004;227(4):503–511. Available at: https://doi.org/10.1016/j.jtbi.2003.11.022.

Stokes, P.A. *et al.* The eScriptorium VRE for Manuscript Cultures. *Classics@*. Edited by C. Clivaz and V.A. Garrick, 2021;18(Ancient Manuscripts and Virtual Research Environments, special issue). Available at: https://classics-at.chs.harvard.edu/classics18-stokes-kiessling-stokl-ben-ezra-tissot-gargem/ (Accessed: 16 February 2023).

Ströbel, P. and Clematide, S. Improving OCR of Black Letter in Historical Newspapers: The Unreasonable Effectiveness of HTR Models on Low-Resolution Images. *Digital Humanities 2019* [Preprint], 2019. Available at: https://doi.org/10.5167/UZH-177164.

Ströbel, P.B. *et al.* Evaluation of HTR models without Ground Truth Material. *Proceedings of the Thirteenth Language Resources and Evaluation Conference. LREC 2022*, Marseille, France: European Language Resources Association, 2022:4395–4404. Available at: https://aclanthology.org/2022.lrec-1.467 (Accessed: 17 January 2023).

Ströbel, P.B., Clematide, S. and Volk, M. How Much Data Do You Need? About the Creation of a Ground Truth for Black Letter and the Effectiveness of Neural OCR. *Proceedings of the 12th Language Resources and Evaluation Conference.*

LREC 2020, Marseille, France: European Language Resources Association, 2020:3551–3559. Available at: https://www.aclweb.org/anthology/2020.lrec-1.436 (Accessed: 14 October 2020).
TensorFlow Developers. 'TensorFlow'. Zenodo, 2021. Available at: https://doi.org/10.5281/zenodo.4758419.
Tesseract Open Source OCR Engine, n.d. Available at: https://web.archive.org/web/20220125061256/https:/github.com/tesseract-ocr/tesseract (Accessed: 20 October 2022).
Tomoiaga, C. *et al.* Field typing for improved recognition on heterogeneous handwritten forms. arXiv, 2019. Available at: http://arxiv.org/abs/1909.10120 (Accessed: 27 October 2022).
Transkribus Lite, n.d. Available at: http://web.archive.org/web/20220119164148/https:/transkribus.eu/lite/ (Accessed: 19 July 2022).
Tsochatzidis, L. *et al.* HTR for Greek Historical Handwritten Documents. *Journal of Imaging*, 2021;7(12):260. Available at: https://doi.org/10.3390/jimaging7120260.
White, N. Training Tesseract for Ancient Greek OCR. *Εὔτυπον*, 2012;28–29:1–11.