

Evaluation of HTR models without Ground Truth Material

Phillip Benjamin Ströbel¹, Simon Clematide¹, Martin Volk¹, Raphael Schwitter²,
Tobias Hodel³, David Schoch³

¹Department of Computational Linguistics, ²Department of Classics, ³Walter Benjamin Kolleg

^{1,2}University of Zurich, ³University of Bern

{pstroebe, siclemat, volk}@cl.uzh.ch

raphael.schwitter@sglp.uzh.ch

{tobias.hodel, david.schoch}@unibe.ch

Abstract

The evaluation of Handwritten Text Recognition (HTR) models during their development is straightforward: because HTR is a supervised problem, the usual data split into training, validation, and test data sets allows the evaluation of models in terms of accuracy or error rates. However, the evaluation process becomes tricky as soon as we switch from development to application. A compilation of a new (and forcibly smaller) ground truth (GT) from a sample of the data that we want to apply the model on and the subsequent evaluation of models thereon only provides hints about the quality of the recognised text, as do confidence scores (if available) the models return. Moreover, if we have several models at hand, we face a model selection problem since we want to obtain the best possible result during the application phase. This calls for GT-free metrics to select the best model, which is why we (re-)introduce and compare different metrics, from simple, lexicon-based to more elaborate ones using standard language models and masked language models (MLM). We show that MLM-based evaluation can compete with lexicon-based methods, with the advantage that large and multilingual transformers are readily available, thus making compiling lexical resources for other metrics superfluous.

Keywords: handwritten text recognition, digital humanities, evaluation, ground truth data, resources, model selection

1. Introduction

Optical Character Recognition (OCR) has become a well-established technique for digitising historical printed collections in libraries and archives. At the same time, *Handwritten Text Recognition* (HTR) is also increasingly finding its way into these institutions (Terras, in press). The digitisation efforts of libraries and archives simplify access to sought-after documents for researchers from various disciplines and allow them to pursue research questions they could either not or only very clumsily have answered without digital copies of the sources.

Still, we can only consider OCR and HTR to be “research facilitators” as long as they perform within reasonable accuracy ranges. Several studies have shown that inaccuracies in OCRred documents harm information retrieval and text mining techniques like named entity recognition and linking, topic modelling, and language modelling (Alex and Burns, 2014; Chiron et al., 2017; van Strien et al., 2020; Hamdi et al., 2020; Pontes et al., 2019; Hill and Hengchen, 2019).

But what do we mean by “reasonable accuracy ranges”? Holley (2009) shared the experiences of the *National Library of Australia Newspaper Digitisation Program*, which found that OCR contractors and libraries usually agree that “good” OCR means > 98%, “average” OCR between 90% and 98%, and “poor” OCR any score < 90% character accuracy, respectively. Springmann et al. (2016) defined ranges of > 95% and between 90% and 95% character accuracy, but one time referring to the former as “good” and later

as “excellent”, and to the latter as “reasonable” and in another place as “good”, respectively. Although opinions of what good OCR is differs, there is consensus that any score < 90% means poor quality¹.

While Holley (2009) and Springmann et al. (2016) were evaluating OCR for printed texts, these accuracy ranges naturally extend to text recognised from handwritten documents. However, HTR for personal handwriting styles is particularly challenging since a personal hand has much more variance than printed fonts or regular scripts like blackletter. Nonetheless, HTR has seen considerable improvements over the last decade, mainly thanks to the advent of neural networks (Graves and Schmidhuber, 2008; Graves et al., 2009). These developments fostered the creation of platforms like *Transkribus*², which facilitate the production of ground truth (GT), i.e., humanly transcribed material, and large-scale training of HTR models. Given suitable training data, the *Transkribus*-internal HTR+ model (Michael et al., 2018) can achieve character error rates (CERs) of about 5% (Mühlberger et al., 2019), therefore coming close to what Holley (2009) considered as “good” and Springmann et al. (2016) already as “excellent”.

Although HTR has made considerable progress, it has not yet found its way to a broader application in li-

¹In this paper, we mainly look at character error rates, which are simply the inverse of character accuracy. E.g., a 98% character accuracy corresponds to a 2% character error rate.

²<https://readcoop.eu/de/transkribus/>

braries and archives. A look at the digital platform for manuscript material for Swiss libraries and archives *e-manuscripta*³ shows that in the category “correspondence” containing ~45k titles, only 313, or 0.1%, contain transcriptions. With the output quality of modern HTR, libraries and archives can further process the scans that they accumulated over the years of mass digitisation and thus provide researchers and the public access to the content of manuscripts of various sorts. However, the deployment of HTR models to produce this content comes with uncertainties. While we can evaluate the performance of HTR models during the development phase via accuracy or error scores on validation and test sets, we cannot do so during the application phase. Hence, there are only a few options to estimate the quality of the resulting texts:

1. We compile a new GT from a representative sample of the new data.
2. We rely on confidence scores.
3. We find GT-free evaluation metrics.

As concerns the first option, compiling a GT is costly, time-consuming, and requires human experts. Moreover, we suspect that we would only obtain a rough estimate of the overall quality (Hodel et al., 2021). As for the second option, confidence scores are not always available, and if they are, it is unclear whether and to what extent they correlate with actual accuracy scores across different documents and distinct models (Springmann et al., 2016). This leaves the third option as the only remaining one, which we will address in this paper. We propose different metrics in Section 3.1 and investigate whether there is a significant correlation between the introduced metrics and the actual CER.

The scrutiny of metrics for performance evaluation has practical implications for libraries and archives: on the one hand, they can provide an estimate of the text quality, which is relevant for researchers working with this data. On the other hand, GT-free metrics allow for model selection when several models have produced hypotheses. Factors like the number of authors in a particular collection, the size of the GT, the skew of the author distribution, scan quality, and the like justify the training and testing of multiple models with different parameters and training data. In order to produce the best possible text during the application phase, it is necessary to apply all models and decide which model produced the best result for each document. In such a case, GT-free evaluation metrics can assist libraries and archives in making this choice. Hence, we also test whether GT-free evaluation metrics are suitable for model selection.

To summarise, in this paper, we investigate a collection of GT-free evaluation metrics to provide a quality

estimation of HTR model results. Moreover, we examine whether we can use the metrics to select the best performing models.

2. Related Work

Studies examining the evaluation of HTR output are few, which might be because baselines on popular data sets like the IAM (Marti and Bunke, 2002) database for handwritten English sentences are indicated in standard measures (CERs or word error rates) (Sánchez et al., 2019). Neudecker et al. (2021) pointed out that although the evaluation of OCR output is standardised by measures like CERs, different implementations in tools that assist the evaluation make it difficult to compare model performances. Additionally, Schambach (2010) stated that such standard measures to rank models “may be of limited significance for the decision to choose a recognition system for a real-life application.” His work showed that for the recognition of postal addresses, the final performance metric and, e.g., error limitation, i.e., the rejection of uncertain results, for different model architectures do not necessarily correlate.

Whenever confidence scores are available, we can exploit them to estimate the text quality. For example, Sarkar et al. (2001) used confidence score outputs by an OCR model to build a triage system for OCR documents. They classified documents to bypass manual inspection, which they successfully achieved for 41% of all documents in their validation set. Overall, they could show that their triage method significantly sped up the document verification process.

Clausner et al. (2016) suggested an OCR quality prediction system that assesses the accuracy of OCR results by any given system. Their classifier relies on 28 features, including metadata, image and layout information, textual features, and confidence scores. Their model was trained on newspaper data and provided reliable OCR quality estimates with an average error rate of 6.1% on a bag-of-words basis, indicating suitability for OCR quality estimation. In addition, a more detailed feature analysis of their classifier showed that the confidence scores are the third most important feature, only topped by lexical features.

Work by Springmann et al. (2016) on the automatic OCR quality estimation of historical printings also showed that the relationship between accuracy and confidence scores is significant (in contrast to Schambach (2010)).⁴ Their approach investigated two measures: a mean token lexicality score based on the text-channel-model by Reffle and Ringlstetter (2013) and confidence values that the OCR model assigns to its output characters. Moreover, they further improved their “standard” models that they had already trained by continuous training with different amounts of “pseudo ground

³<https://www.e-manuscripta.ch/>

⁴We believe that the extent of how much we can trust a model’s confidence scores strongly depends on the model itself.

truths” consisting of text output where standard models return high confidence values. Lastly, Springmann et al. (2016) used the confidence values in order to rank the models during training, which enabled them to choose the best one according to the confidence values after all training runs had been completed.

Alex and Burns (2014) illustrated the usefulness of lexicality in OCR outputs and advocated for a much simpler evaluation procedure. They examined the influence of OCR errors on named entity recognition in historical documents and found that especially recall is affected. They suggested measuring the quality of a text by simply computing the ratio of all words in a document and words occurring in an English dictionary. The test of their metric against human evaluation showed a correlation between their score and human judgements. Salah et al. (2015) developed an approach that operates without external language resources. Their method relied on two OCR outputs: an OCR result and a second OCR output by another system called the reference (i.e., an uncertain GT). They used word-based character agreement rates and mutual information scores as features for a Support Vector Regressor that predicts the word recognition rate. They found their approach reliable, with an R^2 score of 0.95.

We can also use error classifiers in OCRred texts as quality estimators. Nguyen et al. (2019) proposed an error detector that uses several word- and character- n -gram measures as features to train an OCR error detector. They also introduced a so-called *peculiarity index* which they defined as the root-mean-square of the indices of its trigrams. The higher the peculiarity index of a word, the more unusual its trigram composition is and, consequently, the more unlikely it is to be correct. Their error detector works best with all features included, reaching an F-score of 79% on monographs and 70% on periodicals.

3. Method

Our goal is to evaluate different GT-free metrics and to check for each of them if it is suitable to rank HTR models. In the first step, we train many models on data sets we specify in more detail in Section 4. Next, we apply all models to different test sets and compute all metrics introduced in this Section. This procedure results in a ranking for each metric. At the same time, we can rank all the models we have trained according to the CER they achieve on the test set, which we use as our reference ranking. Ideally, a metric’s ranking would be the same as the ranking according to the CER. We use Spearman’s rank correlation measure to learn whether the reference ranking correlates with the metrics’ rankings. Its coefficients tell us at what significance level each two ranking pairs correlate. Fig. 1 summarises the pipeline, which we will describe in more detail in this and the forthcoming sections.

3.1. Metrics

The code and language resources to run the metrics will be made available with the publication of the conference proceedings.⁵

3.1.1. Lexicon-based Metrics

Clausner et al. (2016) showed that lexical features are most suitable to predict OCR quality, while Alex and Burns (2014) indicated that a simple ratio of recognised words against a lexicon correlates well with human judgement about OCR quality. We thus include lexicon-based methods in our metrics.

Token Ratio The token ratio determines the percentage of tokens recognised by the HTR model that also occur in a reference. Let T be the tokens in the HTR result and \mathcal{V} the corresponding vocabulary. The token ratio is given by

$$\mathcal{R}_{\text{token}} = \frac{c(T \in \mathcal{V})}{c(T)}. \quad (1)$$

This measure performs a lexicon lookup, where the lexicon is based on a collection of texts in the same language.

Character N -gram Ratio Nguyen et al. (2019) included an n -gram analysis in the OCR error detection process. We also include character n -gram ratios against a lexicon, but without the peculiarity index that Nguyen et al. (2019) introduced.

We consider the ratios of 2- to 7-grams. For each n -gram variant, we build the corresponding n -gram vocabulary from the reference text and the HTR results. Be \mathcal{G} the set of n -grams in the reference text and N the n -grams in the HTR result. The n -gram ratio is given by

$$\mathcal{R}_{n\text{-gram}} = \frac{c(N \in \mathcal{G})}{c(N)}. \quad (2)$$

3.1.2. Language Modeling Perplexity

RNN-based HTR architectures learn an implicit character language model (LM) (Sabir et al., 2017). Other approaches pre-train a LM and incorporate it in an HTR model to boost performance (Kang et al., 2021). LMs thus form an integral part of the HTR process, so we suggest using external ones to evaluate HTR results. This is the first approach that uses LMs for HTR quality estimation to the best of our knowledge. We differentiate between statistical LMs and transformer-based LMs.

Statistical Language Modeling Perplexity The task of LMs is to predict a word given a sequence of preceding words. Probabilistic LMs learn from massive text corpora which words must follow each other. A statistical LM is given by

$$P_{MLE}(w_n | w_1 \dots w_{n-1}) = \frac{c(w_1 \dots w_n)}{c(w_1 \dots w_{n-1})} \quad (3)$$

⁵<https://github.com/pstroe/atr-eval>

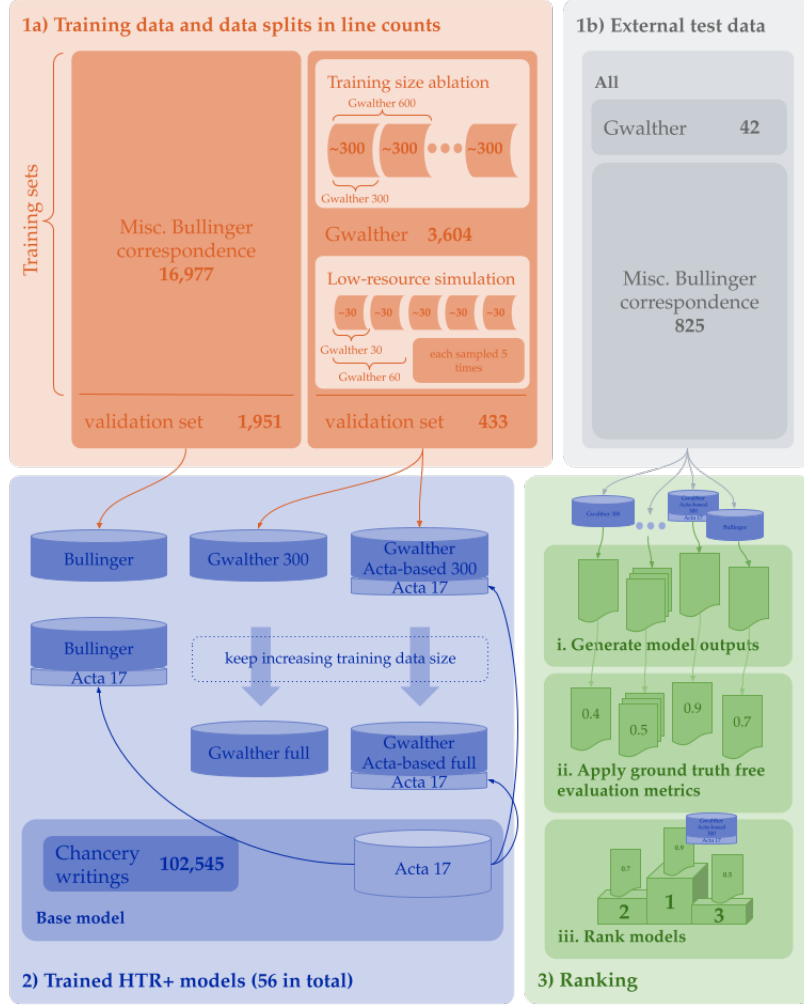


Figure 1: Training and evaluation procedure.

where *MLE* refers to *Maximum Likelihood Estimation* and n to the order of n -grams. For example, a bi-gram LM would predict the next word based on the preceding word.

We can evaluate LMs with the perplexity (PPL) score, a measure from information theory. For example, given a LM P and test set W consisting of words in their sentential contexts, the word-level PPL measures the surprisal of the model when computing the likelihood of the sequence of words:

$$PPL(W) = 2^{H(W)} \quad (4)$$

where H is the cross-entropy of the sequence W :

$$H(W) = -\frac{1}{N} \sum_{i=1}^N \log_2 P(w_i) \quad (5)$$

The lower the PPL score, the less surprised the model is, and hence, the better the text quality should be for HTR results.

Pseudo-Perplexity With the advent of deep learning, language modelling has considerably increased accuracy, not least thanks to transformer-based (Vaswani

et al., 2017) methods like, e.g., BERT (Devlin et al., 2018). BERT is an approach that uses the transformer blocks and attention heads and learns two tasks: (1) the prediction of tokens that are masked in the input and (2) the prediction of whether the next sentence is a natural continuation of the input sentence. As such, the authors designed BERT as a language representation model suitable to be fine-tuned on many tasks like natural language understanding, question answering, or grounded common sense inference. There is a large amount of research striving to improve and modify *masked language models* (MLMs) like BERT. One modification we will use is RoBERTa (Liu et al., 2019), which used longer training over more and larger batches and more data, the omission of the next sentence prediction, and other tweaks to make BERT more robust.

Whilst PPL is a suitable metric for traditional LMs, the MLM objective does not permit a direct calculation of PPL. Instead, Salazar et al. (2019) proposed the use of *pseudo-perplexities* (PPPLs). The PPPL is the exponential of the word-normalised sums of each token's conditional log probability in all output sentences to evaluate. More concretely, given an input sentence, we

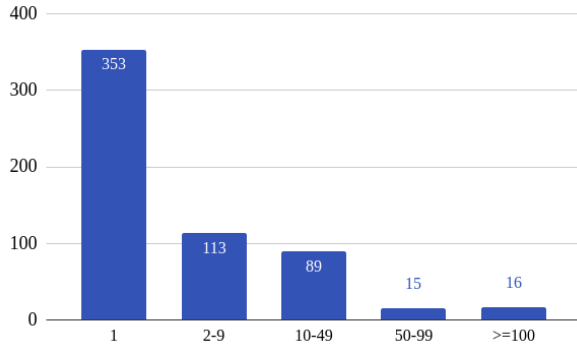


Figure 2: Author distribution of the Bullinger correspondence. Number of letters on the x-axis.

mask each token once, add up the conditional log probabilities of all tokens and normalise by the number of words in the sentence. We use the MLM scorer⁶ framework proposed by the authors for our experiments.

4. Data and Models

4.1. HTR Data

In contrast to the evaluation methods in Section 2, we work with historical handwritten data, more concretely, with samples from Heinrich Bullinger’s (1504-1575) extensive correspondence. The Zurich State Archives and the Zurich Central Library have preserved some 2,000 letters that Bullinger wrote and 10,000 letters that he received. 80% of the letters are Latin, most of the others in Early New High German. About 3,000 letters had already been manually transcribed, edited, and published in printed form in efforts that spanned three decades.⁷ Another 5,700 letters have been transcribed by various scholars and are available as electronic texts, albeit in uncertain quality. In our ongoing project *Bullinger digital*⁸, we make the scans and transcriptions accessible for researchers and the public. Moreover, we will produce automatic transcriptions for the roughly 3,000 letters for which no transcriptions exist and thus will face model selection and HTR quality estimation issues.

Bullinger received many more letters than he wrote, so the author distribution is heavily skewed. Fig. 2 exemplifies this fact; over 350 contemporaries only ever wrote once to Bullinger. This skew presents a considerable hurdle for HTR models since the presence of many different hands makes it more difficult to generalise. Fig. 3 shows examples of how the handwritings differ among authors. Especially in such a setting, it can be necessary to train multiple models with different author distributions and choose the model that performs best. Our training data for these models comes from two sources. Firstly, we found that a correspondent of Bullinger’s, Rudolf Gwalther, wrote a volume called

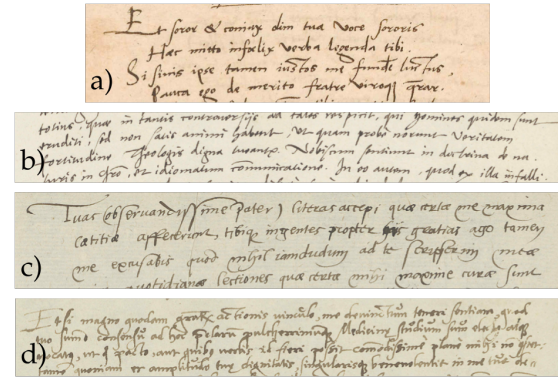


Figure 3: Different handwriting styles. a) poetry by Rudolf Gwalther, letters by b) Rudolf Gwalther, c) Matthieu Coignet, and d) Kaspar Wolf.

Lateinische Gedichte, whose scan images and partial transcriptions are available on *e-manuscripta*⁹. We loaded the images into the Transkribus interface, applied layout recognition to identify lines, and aligned the transcriptions with the lines. We made minor corrections during this process, mainly concerning inconsistent capitalisation and punctuation. The resulting “Gwalther” data set consists of $\sim 4,000$ lines and is one of our training sets (see box 1a on the right in Fig. 1).¹⁰ Secondly, we aligned the scan images that our suppliers have already delivered with the already published transcriptions by applying Transkribus’ Text2Image module, which resulted in an automatically aligned GT of roughly 20,000 lines that we call “Misc. Bullinger correspondence” (see box 1a on the left in Fig. 1). An analysis of the author distribution shows a heavily skewed picture with 69 different authors, where about 4,600 lines stem from Oswald Myconius and only two, e.g., from Georg Cassander.

We also compiled two separate test sets. On the one hand, we sampled 825 lines from a portion of scans without transcriptions (box 1b in Fig. 1) and manually transcribed the lines. Again, the author distribution is heavily skewed. Moreover, 35% of the 825 lines are by authors, which we also find in the “Misc. Bullinger correspondence”. On the other hand, we include a separate 57 lines letter from Gwalther that is not part of any training or test data. We observed that Gwalther’s writing style differs in the poetry volume and the letters, as Fig. 3 shows. Models trained on “Gwalther” data should nevertheless perform better on this test set than other models.

4.2. Language Resources

Since we worked with Latin correspondence, we collected Latin reference material (i.e., we did not use

⁶<https://github.com/awslabs/mlm-scoring>

⁷They are available on <http://teoigsed.uzh.ch/>.

⁸<https://www.bullinger-digital.ch/>

⁹<https://www.e-manuscripta.ch/zuz/content/titleinfo/1111284>

¹⁰This set is publicly available on <https://doi.org/10.5281/zenodo.4780947>.

texts in Early New High German). For the experiments in Section 5, we downloaded the Latin part of the CC-100 corpus¹¹ (Conneau et al., 2019; Wenzek et al., 2019). The raw file contains 2.5G of text, i.e., 390 million tokens. Next, we removed generic "Lorem ipsum" text, deduplicated the corpus and retained only lines consisting of characters in the Latin alphabet, numbers, and a selection of punctuation marks. Finally, we tokenised and normalised the corpus with the help of the *Classical Language Toolkit* CLTK (Johnson et al., 2021) and split it into sentences. After preprocessing ~206 million tokens remain.

4.3. Models

4.3.1. HTR Models

We trained all our HTR models with Transkribus and the HTR+ model architecture for 50 epochs. For both training data sets in box 1a in Fig. 1, we set aside 10% of the corpus for validation and 90% for training. Transkribus offers the possibility to include base models in the training process. We selected the publicly available model "Acta_17 HTR+."¹² A base model initialises the model's weights and allows for fine-tuning the model on novel data. This way, the model knows something about handwriting before seeing the new training data, leading to faster convergence and better performance. We trained a variant with and without a base model for every data set.

For the "Gwalther" data set, we additionally conducted a size ablation, starting at 300 lines (corresponding to roughly ten pages) up to the usage of the entire training set in increments of 300 lines, following a similar approach as detailed in Ströbel et al. (2020). Moreover, we simulated an extreme low-resource setting, in which we sampled 30, 90, and 150 lines five times and used those as training data.

We sketched the training procedure in box 2 in Fig. 1. This training regime results in a total of 56 models. We only considered the best performing model per training set size for the low-resource simulation setting, thus discarding 24 models. We also evaluated the base model Acta_17 HTR+ and the *Spruchakten M.2-11* model on the "Gwalther" validation set and the test sets.¹³ The reason for including this model is to check whether the effects of applying a model on paleographically more related data (since it's closer to Bullinger's time) are more pronounced than just corpus size.

We present the results of all the models in Table 1. We note a considerable improvement of performance the more data we add, while models trained with a base model consistently perform better than their counter-

parts without a base model. The Spruchakten model, although trained on data closer to Bullinger's time of living, performs worse than the Acta_17 model, which indicates that more data helps.

4.3.2. Language Models

We used the Latin corpus presented in Section 4.2 to train all LMs. We applied a data split of 90% for training and 10% for testing and kept it constant for all models.

Statistical LM We used the *KenLM* package (Heafield, 2011; Heafield et al., 2013) to compute a 5-gram Kneser-Ney interpolated LM on the Latin training data. It achieves a PPL score of 23.13 on the Latin corpus test set.

Masked Language Models For the evaluation with the help of MLMs, we used the pre-trained multilingual BERT model based on Devlin et al. (2018). Its training data consists of Wikipedia articles in 104 different languages, among which we also find Latin. It is available via the Hugging Face model hub (Wolf et al., 2019).¹⁴

Since the multilingual BERT model uses WordPiece tokenisation and a shared vocabulary of 110k subtokens, we hypothesise that a single-language MLM should be superior to multilingual BERT. Thus, we trained a RoBERTa model with standard settings on the training set of our Latin corpus. We make the model¹⁵, as well as the training and test data¹⁶ available on the Hugging Face hub.

5. Experiments

5.1. HTR Quality Estimation Suitability

After training the HTR and language models and preparing our language resource, we want to investigate whether the metrics introduced in Section 3.1 are suitable for HTR quality evaluation. We aim to establish a correlation between the CERs and the different metrics. Thus, we applied each model to our test set and computed the scores of our metrics. For evaluating PPPL, PPL, and token ratio, we preprocessed the HTR result in the same way as our Latin reference corpus to guarantee a fair comparison.

Our null hypothesis H_0 is that the metrics' scores on the test data and the CERs that the models achieve on the same data do not correlate. We follow Springmann et al. (2016)'s approach in fitting linear models with the CERs as dependent variables and the scores of the metrics as predictors using the R package (R Core Team, 2021). The adjusted R^2 informs us about the goodness of fit, where a higher value indicates the model can fit the data well. This procedure helps us determine whether a metric reliably estimates HTR quality.

¹¹<http://data.statmt.org/cc-100/>, accessed: 29.09.2021.

¹²The data (chancery writings) and model were compiled by Alvermann et al. from the University of Greifswald. Period: 1580-1705, size: ~600k words in Latin and Low German from 1k different authors, epochs: 1k.

¹³The model was trained on a subset of the Acta_17 training set: period: 1583-1653, size: 246k words, epochs: 1k.

¹⁴<https://huggingface.co/bert-base-multilingual-cased>

¹⁵<https://huggingface.co/pstroe/roberta-base-latin-cased>

¹⁶<https://huggingface.co/datasets/pstroe/cc100-latin>

	Number of lines for training														
	30	90	150	300	600	900	1200	1500	1800	2100	2400	2700	3000	3300	3600
No base	39.33	14.37	11.39	7.28	5.55	4.96	4.36	4.14	3.9	3.94	3.64	3.59	3.36	3.24	3.29
Acta-based	13.67	8.38	7.08	5.5	4.42	4.03	3.66	3.43	3.27	3.13	3.09	3.2	2.82	2.8	2.74
Other models (evaluated on the “Gwalther” validation set)															
Bullinger	6.99														
Bullinger+Acta	6.56														
Acta_17	14.66														
Spruchakten	15.95														

Table 1: CERs for all models trained on different training data and number of lines on their respective test sets.

5.2. Ranking Ability

The ranking tells us whether our metrics capture differences in HTR models. Our null hypothesis H_0 states that there is no correlation between the rankings of different models based on CERs and the rankings based on the metrics. Rejecting H_0 would prove the existence of such a correlation and validate the proposed metrics. Furthermore, the strength of the correlation between the CER and each of the metrics allows us to determine which metric the ranking should be based on.

We compute Spearman’s rank correlation coefficient ρ between the reference and metric rankings to check for correlations. It is given by

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)}, \quad (6)$$

where d_i denotes the difference between two ranks and n is the number of ranks. A comparison against critical values indicates whether the correlation is significant, thus rejecting H_0 .

Applied to our data, this means that we ran all our models on the Gwalther₅₇ and the Bullinger₈₂₅ test sets in block 1b in Fig. 1, computed the corresponding CERs, and ranked the models according to the CERs. These are our reference rankings. We added the reference ranking for the Gwalther₄₃₃ validation set in block 1a. Table 1 shows that many models perform in similar ranges, especially for the Acta-based models. Our metrics should reflect this by returning higher values when we apply the models to the validation set. We then took each model output for the Gwalther₄₃₃, Gwalther₅₇, and the Bullinger₈₂₅ test set and applied our metrics to it. This produces a score for each model and metric, which we used to rank the models. Lastly, we computed Spearman’s ρ by comparing each metric ranking to its corresponding reference ranking. For example, we compared the model ranking of the PPL metric on the Gwalther₅₇ test set against the model ranking according to the CERs on the same set.

5.3. Hit Top-N

The last investigation we conducted is a check for each metric on how well it identifies a model in the Top-1, Top-3, or Top-5. As we see in Table 1, the Top-5 Acta-based models have a difference of at most 0.35 per-

centage points, which lead us to assume that for practical applications, a method that would identify either of those five models as the best performing model was suitable for model selection. A single factor ANOVA test on the Gwalther₄₃₃ with a p-value of 0.71 hardens our suspicions that the means between the Top-5 performing models are not significantly different.

6. Results and Discussion

Table 2 shows the adjusted R^2 values for the models we fitted for each CER vs. metric pair. The superscript behind each score indicates the polynomial degree of the model. The higher the R^2 score, the better the model fits the data, i.e., the higher the correlation between metric and actual CER. We exemplify this with two metrics for the Gwalther₅₇ test set from the Table 2 in Fig. 4. Each black dot represents a model. We notice the linear relationship (blue line) between metric and CER. The big picture from HTR quality estimation suitability confirms that our metrics generally find it easier to score the Gwalther₄₃₃ set. This is only logical since we trained most of our models on the “Gwalther” set.

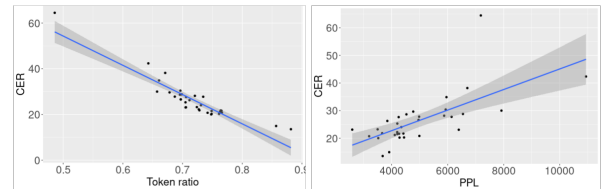


Figure 4: Model fit for two measures (Token ratio left and PPL right) for the Gwalther₅₇ test set. Regression line in blue, 95% confidence interval in grey.

To put our results in relation to previous work, we would like to point out that some of our R^2 scores are considerably higher than what Springmann et al. (2016) reported for their confidence vs. CER and lexicality vs. CER correlations. However, our results confirm the importance of the lexicality of the output and that its measuring is suitable for HTR quality estimation. As concerns character n -grams, the general tendency is, the higher the order of the n -gram, the better it correlates with the CER.

As for the (P)PPL metrics, especially the scoring with the multilingual BERT model results in reliable esti-

Metrics		Adj. R ²		
		Gwalther ₄₃₃	Bullinger ₈₂₅	Gwalther ₅₇
PPPL	BERT	0.99 ²	0.80 ²	0.92 ²
	RoBERTa	0.98 ¹	0.69 ²	0.88 ²
PPL	Statistical LM	0.65 ²	0.52 ¹	0.42 ¹
Token ratio		0.99 ¹	0.95 ⁴	0.92 ³
Character <i>n</i> -grams	2-gram	0.42 ²	-0.01 ¹	0.14 ²
	3-gram	-0.03 ¹	0.09 ¹	-0.03 ¹
	4-gram	0.80 ²	0.35 ¹	0.06 ¹
	5-gram	0.96 ²	0.79 ¹	0.81 ¹
	6-gram	0.99 ²	0.96 ³	0.96 ²
	7-gram	0.99 ²	0.99 ²	0.97 ²

Table 2: R² values for all metrics compared to the performance of the models on the respective validation or test sets. The superscript indicates the polynomial degree of the model we fitted to the data.

Metrics		Ranking Reference (CER)		
		Gwalther ₄₃₃	Bullinger ₈₂₅	Gwalther ₅₇
PPPL	BERT	<u>0.98</u>	<u>0.90</u>	<u>0.90</u>
	RoBERTa	<u>0.96</u>	<u>0.82</u>	<u>0.85</u>
PPL	Statistical LM	<u>0.78</u>	<u>0.65</u>	<u>0.71</u>
Token ratio		<u>0.98</u>	<u>0.95</u>	<u>0.91</u>
Character <i>n</i> -grams	2-gram	-0.28	0.28	0.13
	3-gram	0.01	<u>0.48</u>	<u>0.36</u>
	4-gram	<u>0.58</u>	<u>0.62</u>	0.15
	5-gram	<u>0.90</u>	<u>0.88</u>	<u>0.70</u>
	6-gram	<u>0.97</u>	<u>0.94</u>	<u>0.92</u>
	7-gram	<u>0.99</u>	<u>0.97</u>	<u>0.94</u>

Table 3: The Spearman correlation values for all measures. Significance levels: *0.05*, *0.025*, ***0.01***, *0.005*, ***0.001***.

mates for HTR quality. We found it surprising that a dedicated Latin model does not perform better than multilingual BERT. However, the upside of this finding is that we can use a pre-trained model out of the box for HTR quality estimation without having to train additional models. Moreover, the solid performance of PPPL makes the collection of language-specific data superfluous.

Moving on to the examination of the ranking capabilities of our metrics, we find that most metrics’ rankings exhibit strong correlations to the reference rankings, as Table 3 shows. Overall, the ranking correlation shows a similar picture as the R² values do. Again, the token ratio metric shows strong correlations. The PPL of the statistical LM cannot keep up with the other measures, while our newly proposed application of PPPL for ranking purposes perform equally well as the token ratio metric, at least on the Gwalther data.

Lastly, we take a look at the Top-N evaluation in Table 4. We see that metrics select the best model in 70% of the cases for the Bullinger₈₂₅ test set and in 60% of the cases for the Gwalther₅₇ test set, respectively. This is in stark contrast to the Gwalther₄₃₃ validation set, with only 30% Top-1 hits. This is mainly because the CERs of the models are very close to each other on the validation set, thus making it harder for alternative metrics to pick the best model. All in all, our metrics select one of the five best performing models 90% of the time

for the Bullinger₈₂₅ set and 80% for the Gwalther₅₇ set, thus confirming that our metrics are capable of model selection.

Metrics		Ranking reference		
		Gwalther ₄₃₃	Bullinger ₈₂₅	Gwalther ₅₇
PPPL	BERT	1	1	1
	RoBERTa	1	1	1
PPL	Statistical LM		3	
Token ratio		3	1	1
Character <i>n</i> -grams	2-gram			
	3-gram		3	5
	4-gram		1	3
	5-gram	5	1	1
	6-gram	1	1	1
	7-gram	3	1	1

Table 4: Top-N values for each metric.

7. Conclusion

This paper proposes metrics based on lexical features and (pseudo-)perplexity scores for HTR quality estimation and model selection. We confirmed other research that the lexicality of an HTR result is a strong indicator of quality. However, lexical resources are not always readily available, especially for low-resourced languages. While we dealt with Latin in this paper, we would like to recall that the Bullinger correspondence contains Early New High German texts (ENHG). We would need to assemble a big reference corpus for our or other lexicon-based methods to work. Our experiments confirmed that HTR quality estimation metrics based on transformers like multilingual BERT are suitable replacements. We leave it for future work whether pseudo-perplexity aligns with CER for historical German. After all, ENHG is not part of multilingual BERT’s training data, and it is far more irregular and includes many more spelling variants (especially compared to Latin).

The pseudo-perplexity score also proved capable of ranking HTR outputs. Applying different models to new data and evaluating the output with ground truth-free metrics also opens up new possibilities to further improve HTR quality. For example, when we apply the metrics to lines of different model outputs, we can puzzle together the lines with the best scores, thus combining the best possible outputs from different models.

Our findings are valuable for digitisation initiatives by libraries and archives, which are concerned with the automatic transcription of handwritten documents. Especially our suggested employment of transformers works out of the box and provides quality estimates, based on which rankings of different models are possible. Hence, libraries and archives could quickly produce estimates of the HTR (or OCR) and identify quality issues. It needs further research to investigate whether such metrics are suitable for mass application.

8. Acknowledgements

We express our gratitude to the UZH Foundation and the various sponsors of this project¹⁷. We thank the Zurich State Archive and the Zurich Central Library for providing the scans.

9. Bibliographical References

- Alex, B. and Burns, J. (2014). Estimating and rating the quality of optically characterised text. In *Proceedings of the First International Conference on Digital Access to Textual Cultural Heritage*, pages 97–102.
- Chiron, G., Doucet, A., Coustaty, M., Visani, M., and Moreux, J.-P. (2017). Impact of OCR errors on the use of digital libraries: Towards a better access to information. In *2017 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE.
- Clausner, C., Pletschacher, S., and Antonacopoulos, A. (2016). Quality Prediction System for Large-Scale Digitisation Workflows. In *2016 12th IAPR Workshop on Document Analysis Systems (DAS)*, pages 138–143, Santorini, Greece, April. IEEE.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. (2019). Unsupervised cross-lingual representation learning at scale. *arXiv preprint arXiv:1911.02116*.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Graves, A. and Schmidhuber, J. (2008). Offline handwriting recognition with multidimensional recurrent neural networks. *Advances in neural information processing systems*, 21:545–552.
- Graves, A., Liwicki, M., Fernández, S., Bertolami, R., Bunke, H., and Schmidhuber, J. (2009). A novel connectionist system for unconstrained handwriting recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(5):855–868.
- Hamdi, A., Jean-Caurant, A., Sidère, N., Coustaty, M., and Doucet, A. (2020). Assessing and minimizing the impact of ocr quality on named entity recognition. In *International Conference on Theory and Practice of Digital Libraries*, pages 87–101. Springer.
- Heafield, K., Pouzyrevsky, I., Clark, J. H., and Koehn, P. (2013). Scalable modified kneser-ney language model estimation. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 690–696.
- Heafield, K. (2011). Kenlm: Faster and smaller language model queries. In *Proceedings of the sixth workshop on statistical machine translation*, pages 187–197.
- Hill, M. J. and Hengchen, S. (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.
- Hodel, T. M., Schoch, D. S., Schneider, C., and Purcell, J. (2021). General models for handwritten text recognition: Feasibility and state-of-the-art. german kurrent as an example. *Journal of open humanities data*, 7(13):1–10.
- Holley, R. (2009). How good can it get? Analysing and improving OCR accuracy in large scale historic newspaper digitisation programs. *D-Lib Magazine*, 15(3/4).
- Johnson, K. P., Burns, P. J., Stewart, J., Cook, T., Besnier, C., and Mattingly, W. J. B. (2021). The Classical Language Toolkit: An NLP Framework for Pre-Modern Languages. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 20–29, Online, aug. Association for Computational Linguistics.
- Kang, L., Riba, P., Villegas, M., Fornés, A., and Rusiñol, M. (2021). Candidate fusion: Integrating language modelling into a sequence-to-sequence handwritten word recognition architecture. *Pattern Recognition*, 112:107790.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Marti, U.-V. and Bunke, H. (2002). The iam-database: an english sentence database for offline handwriting recognition. *International Journal on Document Analysis and Recognition*, 5(1):39–46.
- Michael, J., Weidemann, M., and Labahn, R. (2018). Htr engine based on nns p3. Technical report.
- Mühlberger, G., Seaward, L., Terras, M., and 51 more authors. (2019). Transforming scholarship in the archives through handwritten text recognition: Transkribus as a case study. *Journal of Documentation*, 75(5):954–976.
- Neudecker, C., Baierer, K., Gerber, M., Clausner, C., Antonacopoulos, A., and Pletschacher, S. (2021). A Survey of OCR Evaluation Tools and Metrics. In *Proceedings of the 6th International Workshop on Historical Document Imaging and Processing (HIP’21), Lausanne, Switzerland, September 6, 2021*. ACM.
- Nguyen, T.-T.-H., Jatowt, A., Coustaty, M., Nguyen, N.-V., and Doucet, A. (2019). Post-ocr error detection by generating plausible candidates. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 876–881.
- Pontes, E. L., Hamdi, A., Sidère, N., and Doucet, A. (2019). Impact of OCR quality on named entity link-

¹⁷see <https://www.bullinger-digital.ch/about>

- ing. In *International Conference on Asian Digital Libraries*, pages 102–115. Springer.
- R Core Team, (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Reffle, U. and Ringlstetter, C. (2013). Unsupervised profiling of OCRed historical documents. *Pattern Recognition*, 46:1346–1357, 05.
- Sabir, E., Rawls, S., and Natarajan, P. (2017). Implicit language model in lstm for ocr. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, volume 7, pages 27–31. IEEE.
- Salah, A. B., Moreux, J.-P., Ragot, N., and Paquet, T. (2015). Ocr performance prediction using cross-ocr alignment. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 556–560. IEEE.
- Salazar, J., Liang, D., Nguyen, T. Q., and Kirchhoff, K. (2019). Masked language model scoring. *arXiv preprint arXiv:1910.14659*.
- Sánchez, J. A., Romero, V., Toselli, A. H., Villegas, M., and Vidal, E. (2019). A set of benchmarks for handwritten text recognition on historical documents. *Pattern Recognition*, 94:122–134.
- Sarkar, P., Baird, H. S., and Henderson, J. (2001). Triage of ocr results using confidence scores. In *Document Recognition and Retrieval IX*, volume 4670, pages 216–222. International Society for Optics and Photonics.
- Schambach, M.-P. (2010). Reviewing performance metrics for handwriting recognition: Must-rejects and recognition graph scores. In *2010 12th International Conference on Frontiers in Handwriting Recognition*, pages 457–462. IEEE.
- Springmann, U., Fink, F., and Schulz, K. U. (2016). Automatic quality evaluation and (semi-) automatic improvement of ocr models for historical printings. <http://arxiv.org/abs/1606.05157>.
- Ströbel, P. B., Clematide, S., and Volk, M. (2020). How much data do you need? about the creation of a ground truth for black letter and the effectiveness of neural OCR. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 3551–3559, Marseille, France, May. European Language Resources Association.
- Terras, M. (in press). Inviting ai into the archives: The reception of handwritten recognition technology into historical manuscript transcription. In Lise Jaillant, editor, *Archives, Access and AI*, Digital Humanities Research. Transcript Verlag.
- van Strien, D., Beelen, K., Ardanuy, M., Hosseini, K., McGillivray, B., and Colavizza, G. (2020). Assessing the Impact of OCR Quality on Downstream NLP Tasks. In *Proceedings of the 12th International Conference on Agents and Artificial Intelligence*. SCITEPRESS - Science and Technology Publications.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Wenzek, G., Lachaux, M.-A., Conneau, A., Chaudhary, V., Guzmán, F., Joulin, A., and Grave, E. (2019). Ccnet: Extracting high quality monolingual datasets from web crawl data. *arXiv preprint arXiv:1911.00359*.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.