

When Tesseract Does It Alone

Optical Character Recognition of Medieval Texts

Vít Novotný 

Faculty of Informatics, Masaryk University
Botanická 68a, 602 00 Brno, Czech Republic
witiko@mail.muni.cz
<https://mir.fi.muni.cz/>

Abstract. Optical character recognition of scanned images for contemporary printed texts is widely considered a solved problem. However, the optical character recognition of early printed books and reprints of Medieval texts remains an open challenge. In our work, we present a dataset of 19th and 20th century letterpress reprints of documents from the Hussite era (1419–1436) and perform a quantitative and qualitative evaluation of speed and accuracy on six existing ocr algorithms. We conclude that the Tesseract family of ocr algorithms is the fastest and the most accurate on our dataset, and we suggest improvements to our dataset.

Keywords: Optical character recognition, ocr, Historical texts

1 Introduction

The aim of the ANISTO project is to make documents from the Hussite era (1419–1436) available to the general public through a web-hosted searchable database. Although scanned images of letterpress reprints from the 19th and 20th century are available at the Czech Medieval Sources online (CMS online) web site,¹ accurate optical character recognition (ocr) algorithms are required to extract searchable text from the scanned images.

In this paper, we compare the speed and accuracy of six ocr algorithms on the CMS online dataset both quantitatively and qualitatively. In Section 2, we describe the ocr algorithms. In Section 3, we describe our dataset, how it was pre-processed and used in our quantitative and qualitative evaluation. In Section 4, we discuss the results of our evaluation. In Section 5, we offer concluding remarks and ideas for future work in the ocr of Medieval texts.

2 Related Work

Optical character recognition makes it possible to convert scanned images to digital text. For the ANISTO project, an ocr algorithm should:

¹ <https://sources.cms.flu.cas.cz/>

1. support different European languages, mainly Czech, German, and Latin,
2. detect language of the text to enhance the searchability of our database, and
3. use a standard output format to prevent a vendor lock-in.

In this section, we will describe and discuss the ocr algorithms we considered.

2.1 Google Cloud Vision AI

Google Cloud Vision AI is a paid ocr service made available by Google in 2015.²

Google Vision AI supports Czech, German, and Latin, among other languages,³ and detects language at the level of individual letters. Regrettably, only a non-standard json output format is supported, leading to a vendor lock-in.

Google Cloud Vision AI provides two features: `DOCUMENT_TEXT_DETECTION` and `TEXT_DETECTION`. `DOCUMENT_TEXT_DETECTION` performs the ocr of printed documents, whereas `TEXT_DETECTION` solves the more general task of detecting text in arbitrary images, such as camera images of traffic signs. In our experiments, we used the `DOCUMENT_TEXT_DETECTION` feature.

2.2 Tesseract

Tesseract [6] has been developed by Hewlett-Packard in the 1980s and released under a free open-source license in 2005. Since 2006, the development of the project has been funded by Google, and it is presumed that parts of Tesseract are also used in Google Cloud Vision AI.

Initially, Tesseract only supported English. Since version 2, Tesseract has also supported Western languages, including German. Since version 3, Tesseract has also supported Czech, [7], detected language at the level of words, and added support for the standard HOcr XML output format. [1] Since version 3.04, Tesseract has also supported Latin. Since version 4, Tesseract has supported a new LSTM-based ocr engine,⁴ which is regrettably not GPU-accelerated.

In our experiments, we used Tesseract 4.00 in three configurations:

1. `--oem 0`, which uses the non-LSTM ocr engine (f.k.a. Tesseract 3),
2. `--oem 1`, which uses the LSTM ocr engine (f.k.a. Tesseract 4), and
3. `--oem 2`, which ensembles the two ocr engines (f.k.a. Tesseract 3 + 4).

For all configurations, we used the `--psm 3` page segmentation mode, the Czech, German, and Latin (`-l ces+deu+lat`) language models, and the medium-size pre-trained models,⁵ which, compared to the best models,⁶ support the non-LSTM ocr engine of Tesseract 3.

² <https://cloud.google.com/vision/docs/release-notes>

³ <https://cloud.google.com/vision/docs/languages#supported-langs>

⁴ <https://tesseract-ocr.github.io/tessdoc/NeuralNetsInTesseract4.00>

⁵ <https://github.com/tesseract-ocr/tessdata.git>

⁶ https://github.com/tesseract-ocr/tessdata_best.git

2.3 OCR-D

OCR-D [2] is a free open-source project that has been funded by the German Research Foundation and developed by the Berlin-Brandenburg Academy of Sciences, the Herzog August Bibliothek, the Berlin State Library, and the Karlsruhe Institute of Technology. The main goal of the project is to digitize the German cultural heritage 16th–19th century in connection with the vD16, vD17, and vD18 cataloging and archival projects. In February 2020, the project has entered phase 3 and OCR-D is now being deployed to intent organizations.

Unlike Google Cloud Vision AI and Tesseract, which provide fully-automated OCR workflows, OCR-D’s workflows are fully-configurable⁷ and include image enhancement, binarization, cropping, denoising, deskewing, dewarping, segmentation, clipping, line OCR, text alignment, and post-correction. Another distinguishing feature of OCR-D is the ocr4all web frontend,⁸ which enables semi-automatic OCR with human input.

As a part of the line OCR workflow step, OCR-D supports the LSTM-based and GPU-accelerated Calamari engine,⁹ which has achieved the state-of-the-art performance on historical texts typeset in Fraktur. [9,4] With Calamari, OCR-D regrettably does not detect the language of the text, but it supports the standard HOCR XML output format [1] like Tesseract.

In our experiments, we used the recommended workflow for OCR-D,¹⁰ which includes Calamari as the line OCR engine. Since Calamari is GPU-accelerated, we tested OCR-D both with and without a GPU to see the difference in speed.

3 Methods

In this section, we will describe the CMS online dataset, how it was pre-processed and then used to evaluate the OCR algorithms discussed in the previous section.

3.1 Data Preprocessing

In the CMS online dataset, we received 302,909 low-resolution and 168,113 high-resolution scanned images of letterpress reprints from the 19th and 20th centuries. For all low-resolution images, we received Google Cloud Vision AI OCR outputs, although it is unknown if these were produced by the TEXT_DETECTION feature, or the more appropriate DOCUMENT_TEXT_DETECTION feature, and if they were produced using the low-resolution or the high-resolution images.

Although the high-resolution images have an estimated average density of 414 DPI and are suitable for OCR, the low-resolution images have an estimated average density of only 145 DPI. Therefore, we had to design an algorithm that would link the low-resolution images to the matching high-resolution images in

⁷ <https://ocr-d.de/en/workflows>

⁸ <https://www.uni-wuerzburg.de/en/zpd/ocr4all/>

⁹ <https://github.com/Calamari-OCR/calamari>

¹⁰ <https://ocr-d.de/en/workflows#best-results-for-selected-pages>

order to produce a test dataset containing high-resolution images together with corresponding Google Cloud Vision AI ocr outputs as ground truth.

In designing the algorithm, we assumed that there was no other difference between matching low-resolution and high-resolution images other than downscaling. Additionally, we manually inspected the low-resolution and high-resolution images to find a subset of 187,267 (62%) low-resolution images guaranteed to cover all the books from which the high-resolution images originated.

A pseudocode of our linking algorithm is given in Algorithm 1. Using the algorithm, we linked 65,348 (39%) high-resolution scanned images with Google Vision AI ground truth ocr outputs, which served as our test dataset.

Algorithm 1: Linking low-resolution and high-resolution images.

Result: Linked low-resolution and high-resolution images.
 Preprocess all images by rescaling them to 512×512 px and binarizing;
 Index the preprocessed high-resolution images in a vector database;
foreach *preprocessed low-resolution image* **do**
 | Retrieve the ground truth Google Vision AI output;
 | Retrieve the 100 preprocessed high-resolution images nearest to the
 | preprocessed low-resolution image by the Hamming distance;
 | **foreach** *neighboring preprocessed high-resolution image* **do**
 | | Process the high-resolution image by Tesseract 4;
 | **end**
 | Rerank the 100 nearest high-resolution images by TF-IDF cosine similarity
 | between the Google Vision AI ground truth and the Tesseract 4 output;
 | **if** *the nearest high-resolution image is the same after reranking* **then**
 | | Prelink the low-resolution image with its nearest high-resolution image;
 | **end**
end
foreach *book* **do**
 | **if** *all low-resolution images in the book are prelinked* **then**
 | | Link all low-resolution images in the book with their prelinked
 | | high-resolution images;
 | **end**
end

3.2 Quantitative Evaluation

Speed To evaluate the speed of the ocr algorithms, we measured and report the wall clock time in days to process the test dataset on a single CPU/GPU. For Tesseract, we used the *apollo*, *asteria04*, *mir*, *hypnos1*, *turnus01*, and *nymfe*{23..74} nodes at the Faculty of Informatics, Masaryk University, totalling 492 CPU cores. For OCR-D, we used the *epimetheus*{1..4} nodes for evaluating the CPU speed, and *turnus03* for evaluating the GPU speed. For OCR-D, we also measured and report the speed of each workflow step separately.

Accuracy To evaluate the accuracy of the ocr algorithms, we measured and report the Character Error Rate (CER) and the Word Error Rate (WER): [8]

$$\text{ErrorRate}(A, B) = \frac{\text{EditDistance}(A, B)}{\text{Maximum}(|A|, |B|)} \times 100\%.$$

For WER, we lower-cased and deaccented the texts in addition to tokenization in order to better model our full-text search use case. Although CER and WER are correlated, we were mainly interested in WER, which is harder and better corresponds to our full-text search use case.

To discover the origin of the ground truth, we evaluated the accuracy of the Google Cloud Vision AI using both the low-resolution images (f.k.a. G-low) and the high-resolution images (f.k.a. G-high), not just the high-resolution images.

3.3 Qualitative Evaluation

To better understand the strengths and weaknesses of the individual ocr algorithms, we compared their outputs on three different scanned images:

1. A random image from the test dataset.
2. The image with the best WER using the least accurate ocr algorithm.
3. The image with the worst WER using the most accurate ocr algorithm.

To inspect the quality of the ground truth, we manually classified the differences between the ground truth and the ocr output on the three images as either improvements or errors, and we report the ratio of improvement to error.

4 Results

In this section, we will describe the results of the evaluation described in the previous section. We report both quantitative and qualitative evaluation results.

4.1 Quantitative Evaluation

Speed Table 1 shows that Tesseract 3 is the fastest ocr algorithm. The second fastest Tesseract 4 is more than twice as slow as Tesseract 3 because of the higher computational complexity of the non-gpu-accelerated LSTM-based ocr engine. OCR-D achieves speed that is comparable with Tesseract 4 despite the GPU acceleration of Calamari, which is likely because of the fully-configurable workflows and the high disk I/O caused by the locking and the common updates

Table 1. The wall clock time to process the test dataset with different ocr algorithms ordered from the fastest to the slowest. Since Google Cloud Vision AI is not self-hosted, it was not included in the speed evaluation.

	<i>Tesseract 3</i>	Tesseract 4	OCR-D (GPU)	Tesseract 3 + 4	OCR-D (CPU)
Time (days)	61.12	127.69	140.39	172.11	174.07




Fig. 1. The wall clock time to process the test dataset on the level of the individual workflow steps of OCR-D on a single CPU (left) and on a single GPU (right).

to a large METS XML document containing the history of all workflow steps executed for all scanned images in a book. Tesseract 3 + 4 is the second slowest with the time roughly equivalent to the sum of the times of Tesseract 3 and Tesseract 4. OCR-D without GPU acceleration is the slowest OCR algorithm.

Figure 1 shows that without GPU acceleration, more than one third of OCR-D’s time is taken up by the Calamari line OCR workflow step. GPU acceleration causes more than 2× speed increase of this workflow step, whereas the other steps do not benefit from GPU acceleration.

Accuracy Table 2 shows that Google Cloud Vision AI achieved perfect accuracy with neither low-resolution nor high-resolution images, which leads us to the conclusion that the ground truth had been produced using the less appropriate TEXT_DETECTION feature. The higher accuracy of Google Cloud Vision AI with low-resolution images leads us to believe that the ground truth had been produced using the low-resolution images rather than the high-resolution images. Both findings raise doubts about the quality of the ground truth.

Table 2 also shows that Tesseract 4 is the most accurate OCR algorithm with the second most accurate being Tesseract 3. Ensembling Tesseract 3 and Tesser-

Table 2. The accuracy of different OCR algorithms on the test dataset ordered from the most accurate to the least accurate. G-low and G-high correspond to the Google Cloud Vision AI with low-resolution images and high-resolution images, respectively.

	G-low	G-high	Tesseract 4	Tesseract 3	Tesseract 3 + 4	OCR-D
WER (%)	3.85	5.24	13.77	16.04	18.70	30.94
CER (%)	2.44	3.38	9.81	10.60	12.00	19.88

act 4 leads to the second worst accuracy, which suggests a poor confidence heuristic of Tesseract in selecting words from the suggestions of both OCR algorithms. OCR-D is the least accurate OCR algorithm, likely because we used the recommended Calamari models trained on German Fraktur,¹¹ whereas most of our images are scans of letterpress prints in Antiqua typefaces with Czech accents.

4.2 Qualitative Evaluation

Random image We randomly selected page 5 from *Archiv český 01*,¹² with the following ground truth text:

I. PSANJ ČESKÁ CJSARE SIGMUNDA
od roku 1414 do 1437.
Panu Čeňkowi z Wartenberka a z Weselj, neyw. purkrabj Pražskému: aby králowně Sofii na statejch gegjch wěnnjch nátisku činiti nedopauštěl.
Z Teplice u Ferrary, 1414, 24 Mart.
Sigmund z božie milosti Římský a Uherský oc. král.
Urozený wěrný milý! Slyšíme, že někteří páni w Čechách najoswiecenější kněžnu pant Sofii, králewnu Českú, sestru naši milú, mienie a chystajj sie, jie na jejím wěně mimo práwo a mimo panský nález tisknúti; jenžto nerádi slyšíme, anižbychom toho rádi dopustili, by sie jie to od koho mělo státi. Protož od tebe žádáme i prosime, byliliby žeby jmenowanú králewnu, sestru naši milú, mimo práwo kto tisknúti, a nebo na jejiem wěně překázeti chtěl, aby podle ní stál, a jie wěrně pro nás pomohl, aby od swého nebyla tištěna. Na tom nám zvláštní službu učiníš a ukážeš. Dán w Teplici u Ferrarii, weder matky božie Annuntia-tionis, léta králowstwíe našich Uherského uc. w XXVII., a Římského w čtwtřém létě,
Ad mandatum D. Regis: Michael de Priest.
Urozenému Čeňkowi z Wesele
wěrnému nám zvláště milému.

Tesseract 3 achieved 6.36% WER and 3.96% CER. Over 42% of changes were improvements, raising doubts about the quality of the ground truth:

J. PSANJ ČESKÁ CJSARĚ SIGMUNDA
od roku 1414 do 1437.
Panu Čeňkowi z Wartenberka a z Weselj, neyw. purkrabj Pražskému: aby králowně Sofii na. statcjch gegjch wěnnjch nátisku činiti nedopauštěl.
Z Teplice uFermry, 1414, 24 Mart.
Sigmund z božie milosti Římský a Uherský oc. král.
Urozený wěrný milý! Slyšíme, že někteří páni w Čechách najoswiecenější kněžnu paní Sofii, králewnu Českú, sestru nasi milú, mienie a chystajj sie, jie na jejím wěně mimo práwo a mimo panský nález tisknúti; jenžto nerádi slyšíme, anižbychom toho rádi dopustili, by sie jie to od koho mělo státi. Protož od tebe žádáme i prosime, byliliby Zeby jmenowanú králewnu, sestru naši milú, mimo práwo kto tisknúti, a nebo na jejiem wěně překázeti chtěl, aby podle ní stál, a jie wěrně pro nás pomohl, aby od swého nebyla tištěna. Na tom nám zvláštní službu učiníš & ukážeš. Dán W Teplici u Ferrarii, wečer matky božie Annuntia-tionis, léta králowstwíe našich Uherského oc. w XXVII., & Římského w čtwtřém létě.
Ad mandatum D. Regis: Michael de Priest.
Urozenému Čeňkowi z Wesele
Wěrnému nám zvláště milému.

Tesseract 4 achieved 8.57% WER, 4.95% CER. 28% changes were improvements:

I. PSANJ ČESKÁ CJSARĚ SIGMUNDA
od roku 1414 do 1437.
Panu Čeňkowi z Wartenberka a z Weselj, neyw. purkrabj Pražskému: aby králowně Sofii na statcjch gegjch wěnnjch nátisku činiti nedopauštěl.
Z Teplice u Ferrary, 1414, 24 Mart.
Sigmund z božie milosti Římský a Uherský oc. král.
Urozený wěrný milý! Slyšíme, že někteří páni w Čechách najoswiecenější kněžnu paní Sofii, králewnu Českú, sestru naši milú, mienie a chystajj sie, jie na jejím wěně mimo práwo a mimo panský nález tisknúti; jenžto nerádi slyšíme, anižbychom toho rádi dopustili, by sie jie to od koho mělo státi. Protož od tebe žádáme i prosime, byliliby Zeby jmenowanú králewnu, sestru naši milú, mimo práwo kto tisknúti, a nebo na jejiem wěně překázeti chtěl, aby podle ní stál, a jie wěrně pro nás pomohl, aby od swého nebyla tištěna. Na tom nám zvláštní sluZbu uciniš & ukážeš. Dán w Teplici u Ferrarii, wečer matky boZie Annuntia-tionis, léta králowstwíe našich Uherského sc. w XXVII., a Římského w čtwtřém létě,
Ad mandatum D. Regis: Michael de Priest.
Urozenému Ceňkowi z Wesele
wěrnému nám zvláště milému.

¹¹ <https://ocr-d-repo.scc.kit.edu/models/calamari/GT4HistOCR/model.tar.xz>

¹² <https://sources.cms.flu.cas.cz/src/index.php?cat=10&bookid=792&page=5>

Tess. 3 + 4 achieved 12.43% WER, 8.35% CER. 12% changes were improvements:

I. PSANJ ČESKÁ CJSARĚ SIGMUNDA
od roku 1414 do 1437.

Panu Čeňkowi z Wartenberka a z Weselj, neyw. purkrabj Pražskému: aby králowně Sofii na statcjch gegjch wěnnjch nătisku činiti nedopauštěl.
Z Teplice u Ferrary, 1414, 24 Mart.

Sigmund z božie milosti Římský a Uherský oc. král.

Urozeny wörnny mily! Sly&tme, ze někerí páni w Cechách najoswiecenější kněfnu pani Sofii, kráľewnu Českú, sestru nasi milú, mienie a chystaji sie, jie na jejim wéné mimo práwo a mimo panský nález tisknúti; jenZto nerádi slysime, aniZbychom toho rádi dopustili, by sie jie to od koho mělo státi. ProtoZ od tebe Zádáme i prosime, byloliby Zeby jmenowanú kráľewnu, sestru naši milú, mimo práwo kto tisknúti, a nebo na jejim wéné prekázeti chtěl, aby podlé ni stál, a jie wěrné pro nás pomohl, aby od swého nebyla tistěna. Na tom nám zwlástrni sluZbu uciniš a ukáášZes. Dáan w Teplici u Ferrari, wecer matky bozie Annuntia-tionis, léta kráľowstwie nasich Uherského oc. w XXVII, a Římského w ětwtřém létě,

Ad mandatum D. Regis: Michael de Priest.

Urozenému Ceňkowi z Weselě
wěrnému nám zwláste milému.

OCR-D achieved 27.17% WER, 13.95% CER. Only 5% of changes were improvements, which indicates that the ground truth can still differentiate good and bad OCR outputs. Due to the training on German Fraktur without Czech accents, most errors originate from accented letters:

J. PaAN SINA
od roku 1414 do 1437.

Panu Cekowi 2 artemberka a 2 Veselj, neyw. purkrabj Praiskému: aby kráľown Sofii na statejch gegjch wěnnjeh nătisku činiti nedopauštěl.

Teplice u Ferrar, 1414, 24 Mart.

Sigmund z boie milosti Římský a Uherský oc. král.

rozeny wörnny mily! Slysime, e někerí páni w Cechách najoswiecenější kněnu pani Sofii, kralewnu Cesk, sestru nasi milu, mienie a chystaji sie, jie na jejim wéné mimo práwo a mimo panský nález tisliti; jenáto nerádi slysime, anizbychom toho rádi dopustili, by sie jie to od koho mělo státi. Proto od tebe adame i prosime, byloliby eby jmenowanu kralewnu, sestru nasi milu, mimo práwo kto tisknuti, a nebo na jejim wéné prekázeti chtěl, aby podlé ni stál, a jie wvřně pro nás pomohl, aby od swého nebyla tistěna. Na tom nám zwlástrni sluibu ucinis a ukáes. Dán w Teplici u Ferrarii, wecer matky boie Annuntia-tionis, léta kráľowstwie nasich Uherského oe. w XXVI., a Římského w ětwtřém létě.

Ad mandatum D. Regis: Michael de Priest.

Urozenému Ceňkowi 2 Veselě
wěrnému nám zlastě milému.

Best image OCR-D achieved the best accuracy (0% WER, 0.29% CER) on page 59 of *CIM I*.¹³ None of the changes were improvements, confirming our thesis that the ground truth can reliably distinguish good and bad OCR outputs:

vllam exigere, petere aut recipere volumus pecuniam, aut exigí, peti

...

vngelta in dicta ciuitate uidelicet pannorum, mercium institarum et braxaturas ceruisie cum prouentibus ipsorum iuxta placita, per nos et ipsos ciues nostros hincinde habita, infra spacium dicatorum annorum cum adicione quinti anni pro se recipere debeant et habere, quousque omnia et singula debita per ipsos ciues nostros quocumque modo contracta in integrum fuerint persoluta, impedimento nostro et cuiuslibet non obstante. Debet quoque ipsum vngeltum sic recipi atque dari, uidelicet quod vendens pannos siue merces institarum cuius vel hospes de qualibet sexagena grossorum Sex paruos denarios vsuales et emens merces easdem totidem paruos in prima vendicione et empicione tantum et non vltorius soluere est astrictus, et quilibet braxans ceruisiam in ipsa ciuitate de vna braxatura ceruisie vnum grossum pro vngelto ipsis ciuibus dare debet. Si quis vero sub vna sexagena grossorum vendiderit vel emerit in mercibus, vt predicatur, quidquam, de huiusmodi vendicione et empicione pro vngelto nichil dabit. Addicimus etiam, quod omnes mercatores Pragam cum pannis quibuscunque uel mercibus institarum ibidem non emptis transire volentes,

...

Liber vetustissimus statutorum c. 993 str. 61 V archivum m. Prahy.

Worst image Tesseract 4 achieved the worst accuracy (100% WER, 81.56% CER) on page 1297 of *CIM II*.¹⁴ All changes were improvements, since the ground truth did not detect two-column page layout, unlike Tesseract 4. This confirms our thesis that although the ground truth can distinguish good and bad OCR outputs, it cannot distinguish OCR outputs that are better than the ground truth.

¹³ <https://sources.cms.flu.cas.cz/src/index.php?cat=12&bookid=117&page=230>

¹⁴ <https://sources.cms.flu.cas.cz/src/index.php?cat=12&bookid=118&page=1327>

5 Conclusion and Future Work

In our work, we have compared the speed and the accuracy of six ocr algorithms on the cms online dataset from the AHISTO project.

Based on our results, we conclude that the ground truth ocr outputs in our dataset are low-quality and should be replaced or supplemented either by human judgements, or by synthetic data produced from searchable PDF documents.

As far as the ground truth can be trusted, Tesseract 4 is the second fastest and the most accurate ocr algorithm, which also detects language at the level of words. Pre-detecting the language of paragraphs using n -gram frequency analysis and then using Tesseract 4 only with the language models corresponding to the detected languages can further improve its accuracy. [3, Section 4.4]

OCR-D with the Calamari line ocr is comparable to Tesseract 4 in speed and is likely to produce more accurate results than Tesseract 4. However, the performance of Calamari was undermined by its poor pre-trained models. Additionally, Calamari does not detect language. However, ocr-d can align ocr outputs from Calamari and Tesseract 4, which can bring the best of both worlds.

Applying super-resolution algorithms [5] to the low-resolution scanned images can make them suitable for ocr and the AHISTO project.

Acknowledgments. First author’s work was funded by the South Moravian Centre for International Mobility as a part of the Brno Ph.D. Talent project and also by TAČR Ěta, project number TL03000365. I also sincerely thank my colleagues from the NLP Centre and the MIR research group at the Masaryk University for their insight.

References

1. Breuel, T.M.: The hOCR microformat for OCR workflow and results. In: ICDAR 2007. vol. 2, pp. 1063–1067. IEEE (2007)
2. Neudecker, C., Baierer, K., Federbusch, M., Boenig, M., Würzner, K.M., Hartmann, V., Herrmann, E.: OCR-D: An end-to-end open source OCR framework for historical printed documents. In: Proceedings of the 3rd International Conference on Digital Access to Textual Cultural Heritage. pp. 53–58 (2019)
3. Panák, R.: Digitalizace matematických textů. Master’s thesis, Faculty of Informatics, Masaryk University (2006), <https://is.muni.cz/th/pspz5/>
4. Reul, C., Springmann, U., Wick, C., Puppe, F.: State of the art optical character recognition of 19th century fraktur scripts using open source engines (2018), <https://arxiv.org/pdf/1810.03436.pdf>
5. Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1874–1883 (2016)
6. Smith, R.: An overview of the Tesseract OCR engine. In: ICDAR 2007. vol. 2, pp. 629–633. IEEE (2007)

7. Smith, R., Antonova, D., Lee, D.S.: Adapting the Tesseract open source OCR engine for multilingual OCR. In: Proceedings of the International Workshop on Multilingual OCR. pp. 1–8 (2009)
8. Soukoreff, R.W., MacKenzie, I.S.: Measuring errors in text entry tasks: An application of the levenshtein string distance statistic. In: CHI'01 extended abstracts on Human factors in computing systems. pp. 319–320 (2001)
9. Wick, C., Reul, C., Puppe, F.: Calamari-a high-performance tensorflow-based deep learning package for optical character recognition (2018), <https://arxiv.org/pdf/1807.02004.pdf>