






Key-Value Information Extraction from Full Handwritten Pages

Solène Tarride¹(✉) , Mélodie Boillet^{1,2} , and Christopher Kermorvant^{1,2} 

¹ TEKLIA, Paris, France
starride@teklia.com

² LITIS, Normandy University, Rouen, France

Abstract. We propose a Transformer-based approach for information extraction from digitized handwritten documents. Our approach combines, in a single model, the different steps that were so far performed by separate models: feature extraction, handwriting recognition and named entity recognition. We compare this integrated approach with traditional two-stage methods that perform handwriting recognition before named entity recognition, and present results at different levels: line, paragraph, and page. Our experiments show that attention-based models are especially interesting when applied on full pages, as they do not require any prior segmentation step. Finally, we show that they are able to learn from key-value annotations: a list of important words with their corresponding named entities. We compare our models to state-of-the-art methods on three public databases (IAM, ESPOSALLES, and POPP) and outperform previous performances on all three datasets.

Keywords: Key-value extraction · Named-Entity Recognition · Handwritten Document · Segmentation-free Approach

1 Introduction

Although machine learning and deep learning techniques are nowadays commonly used in the field of automatic processing of historical documents [12], scientific work still often focuses on some specific processing steps in isolation. It is common to develop models either for page analysis or line detection, for handwriting recognition or for information extraction. Processing chains are still often developed as a sequence of these steps independently. However, these processing chains suffer from several drawbacks. Firstly, errors accumulate along the chain: if the line detection step is bad, write recognition will be highly impacted and information extraction impossible. On the other hand, the implementation of these chains and their maintenance is complex: each step requires specific skills and annotated data for each model and any update of a part of the chain has an impact on all downstream processes. Finally, the different modules are developed independently and there is no global optimization of the processing chain. For all these reasons, the development of models allowing the extraction of information directly from the image, by an end-to-end approach, with a single model, would be very beneficial.

As far as automatic recognition is concerned, three main types of projects are currently being carried out on collections of historical documents, depending on the intended use. The first type of project aims to carry out a complete transcription of the documents to allow full-text searches [16, 27]. The processing chain then focuses on the page analysis stage to extract a maximum number of lines of text and the handwriting recognition stage to best recognize the text. The result of the processing is then exploited thanks to a search engine that allows queries to be made and documents to be identified according to their content. The second type of processing aims to produce electronic editions of documents [11]. In this case, the emphasis is obviously on the quality of the recognition, but also the fidelity to the text of the document and the reading order. The result of the automatic processing is in this case always submitted to the correction of an expert before publication. The last type of project aims at extracting information from documents in order to populate a database with the information they contain [24]. In addition to the document analysis and handwriting recognition stages, these projects also incorporate an information extraction stage, often in the form of named entity extraction. It is this third type of project, the most complex in its implementation, that we are interested in this work.

Information extraction chains for historical handwritten documents are usually composed of the following steps: line detection or document layout analysis (DLA), handwriting recognition (HTR) and named entity extraction (NER). In this paper, we first reconsider the possibility of combining the HTR and NER models into a single model. Then we study whether it is possible to extend this model to the processing of a complete page without going through a line detection step. Finally, we show that it is possible to go even further and train a single model for the extraction of target information, of the key-value type, without going through an explicit transcription.

The rest of this paper is organized as follows. In Sect. 2, we review the state-of-the-art for information extraction in handwritten documents. We describe our methodology and experiments in Sect. 3. The experimental results are presented and analyzed in Sect. 4. Finally, in Sect. 5 we discuss the conclusions and outline future works.

2 Related Work

Recent advances in computer vision and natural language processing have led to major breakthroughs in the field of automatic document understanding. Deep learning-based systems are now capable of automatically extracting relevant information from historical documents. Interest in this field has been encouraged by the emergence of competitions, such as the Information Extraction on Historical Handwritten Records competition [8] on the ESPOSALLES database [20], as well as the publication of named entity recognition annotations for other databases, such as IAM-NER [26] and POPP [4].

Two main approaches exist to address automatic information extraction from handwritten documents:

- **Sequential approaches** consist in dividing the problem into two successive tasks: handwritten text recognition, and named entity recognition;
- **Integrated approaches** consist in combining text and named entity recognition in a single-step.

Each of these approaches can work at several levels: either on words, lines, paragraphs, or directly on full pages. Segmentation-based systems work on pre-segmented text zones (words, lines, or paragraphs), while segmentation-free systems work directly on full pages. Performing handwriting recognition on smaller zones is usually easier to achieve, but requires a prior segmentation step. As opposed, handwriting recognition on full pages is more challenging (memory management, reading order), but does not require any prior segmentation.

2.1 Sequential Approaches

In sequential approaches, HTR is performed first, then, NER is applied on recognized text. Note that HTR and NER can be applied at different levels: HTR is usually performed at line-level, and NER at paragraph or page-level.

Segmentation-Based Systems. Five systems were introduced during the ICDAR2017 Competition on Information Extraction in Historical Handwritten Records [8] on ESPOSALLES. Most participants used CRNN trained with CTC to recognize handwritten text. Named entity recognition was then performed using logical rules based on regular expressions or CRF tagging. Other methods were proposed after the competition.

Prasad et al. [17] propose a two-stage system combining a CRNN-CTC neural network for HTR on text line images, followed by a BLSTM layer over the feature layer for NER.

Tuselmann et al. [26] also introduce a two-stage system for information extraction that combines a Transformer model [10] for HTR on word images, and a LSTM-CRF model with word embeddings obtained using a pre-trained RoBERTa for NER. They highlight the advantages of two-stage methods for information extraction, as these methods yield state-of-the art results and are easy to improve using post-processing techniques, close dictionary, or pre-trained embeddings.

Monroc et al. [15] compare different off-the-shelf NER libraries on handwritten historical documents: SpaCy [9], FLAIR [1], and Stanza [19]. They perform experiments on three datasets in an end-to-end setting, and study the impact of text line detection and text line recognition on NER performances. Their results highlight that line detection errors have a greater impact than handwriting recognition errors. This conclusion suggests that working directly on pages could prevent segmentation errors from impacting the final entity recognition.

Segmentation-Free Systems. To the best of our knowledge, no system performing sequential HTR and NER at page-level has been proposed so far. However, many segmentation-free HTR models working directly at page-level [5, 6, 28] have been introduced recently. Any of these models could easily be combined with off-the-shelf NER libraries for segmentation-free information extraction.

2.2 Integrated Approaches

Integrated approaches combine HTR and NER in a single step by modeling named entities with special tokens. This can be achieved with or without prior segmentation.

Segmentation-Based Systems. Toledo et al. [25] and Rowtula et al. [22] introduce models that work at word-level. Their systems are able to recognize and classify word images into semantic categories.

Both Carbonell et al. [3] and Tarride et al. [23] propose neural networks that predict characters and semantic tags from line images, respectively, using a CRNN model trained with CTC and an attention-based network. Both of these studies suggest that working on records would allow the model to capture more contextual information.

In [4], the authors use the same approach on French census images from the POPP dataset, as they predict text characters and special tokens for empty cells and column separators. Although, this dataset does not directly include named entities, each word is linked to a specific column and can be seen as a named entity (name, surname, date of birth, place of birth...).

Finally, Rouhou et al. [21] are the first to introduce a Transformer model for combined HTR and NER at record-level on the ESPOSALLES database. They highlight the interest of performing this task on records to benefit from more contextual information. As each page contains several records, this model still requires record segmentation. Moreover, they use a special token for line breaks, as they observe this improves performance.

Segmentation-Free Systems. Carbonell et al. [2] are the first to propose a model that works directly at page-level on ESPOSALLES. Their system is able to jointly learn word bounding boxes, word transcription and word semantic category on ESPOSALLES. However, a major limitation of this method is that it requires word bounding boxes during training.

The Transformer proposed by Rouhou et al. [21] could be applied to full pages in its current stage, although this task has not been tackled by the authors.

Finally, the Document Attention Network (DAN) [5] is able to recognize text on full pages with reading order. It is based on the Transformer architecture and jointly learns characters and special tokens that represent layout information. It is likely that this method is also able to recognize named entities, or in other words, tokens that are not spatially localized but have a semantic meaning. However, the authors did not perform any experiments on named entity recognition.

2.3 Discussion

The literature review opens up three main questions that are discussed in the following.

What Is the Best Approach for Information Extraction? Although this question has been well studied in the past, no consensus has been reached. On the one hand, researchers have shown the interest of sequential methods which can be optimized at every stage (with a language model, a dictionary, or pre-trained embeddings) [15,26]. On the other hand, the advantages of integrated methods have also been demonstrated [23], notably because they benefit from shared contextual features and avoid cascade errors.

Can We Extract Relevant Information from Full Pages? Different methods were designed to work at different levels, some of them requiring prior segmentation of text lines or paragraphs. However, in real-world scenarios, text areas are not known and must therefore be detected automatically, which can introduce segmentation errors. It has been established that segmentation errors have a greater impact on information extraction than handwriting recognition errors [15]. Recently, Transformers have proved their ability to learn from paragraphs and pages [5,21], enabling segmentation-free information extraction. Learning directly from pages increases the task difficulty, but avoids the need for prior segmentation. Moreover, working directly on pages makes it possible to benefit from a larger context [21].

Are Integrated Models Able to Learn from Key-Value Annotations? As sequential approaches rely on HTR, they require the entire transcription before retrieving named entities. However, integrated methods could potentially learn from key-value annotations, which corresponds to a list of words with their corresponding named entities. In this scenario, ground-truth is also easier and faster to produce, as annotators would only have to annotate important words as well as their semantic category. This approach could also be applied in a lot of practical applications where full transcriptions are not available, such as genealogical crowdsourced information (civil status, or personal records). This question has not been studied yet in the context of information extraction.

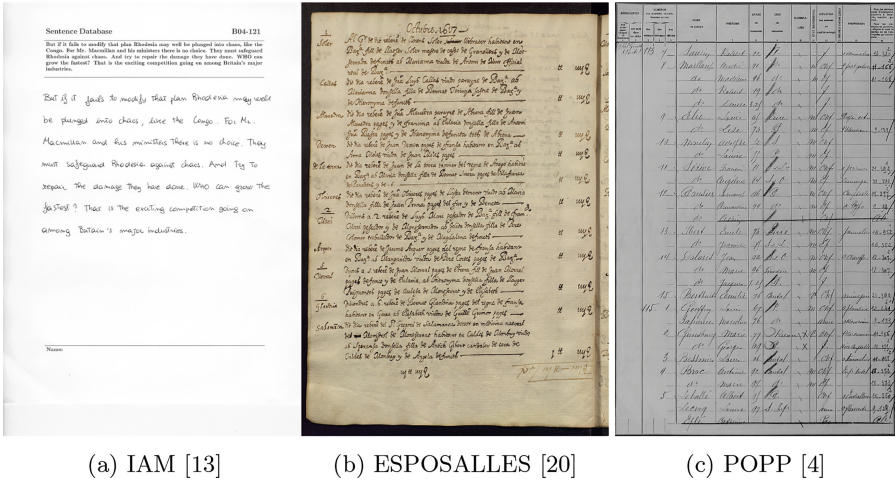
In the next section, we describe the experiments designed to address these three questions.

3 Methodology and Experiments

In this section, we introduce the datasets used during our experiments, present our methodology, and describe the different experiments conducted in this study.

3.1 Datasets

During our experiments, we worked on three public datasets of different kind.



(a) IAM [13] (b) ESPOSALLES [20] (c) POPP [4]

Fig. 1. Examples of pages from the three datasets used in this work

IAM. The IAM dataset [13] is composed of modern documents written in English by 500 writers. It includes 747 training pages with corresponding transcriptions. NER annotations have been made available by Tüselmann et al. [26]. A page from IAM is presented in Fig. 1a.

For our experiments, we use the RWTH split with 18 entities: Cardinal, Date, Event, FAC, GPE, Language, Law, Location, Money, NORP, Ordinal, Organization, Person, Percent, Product, Quantity, Time and Work of art. The details of this split are provided in the appendix. Less than 10% of words are associated to an entity. Due to the large number of classes, some entities have very few examples in the training set. We perform experiments at two levels: text line and page. When working on pages, we remove the header so that the model does not see the printed transcription instruction.

ESPOSALLES. The ESPOSALLES dataset [8] is a collection of historical marriage records from the archives of the Cathedral of Barcelona. The corpus is composed of 125 pages. Each document is written in old Catalan by a single writer. It includes 125 pages with word, line and record segmentations. The details of this split are provided in the appendix. A page from ESPOSALLES is presented in Fig. 1b.

Each word is transcribed and labeled with a semantic category (name, surname, occupation, location, state, other) and a person (husband, wife, husband’s father, husband’s mother, wife’s father, wife’s mother, other person, none). More than 50% of words are associated to an entity. As there is no validation set, we keep 25% of training pages for validation. We perform experiments at three levels: text line, record, and page.

POPP. The POPP dataset contains tabular documents from the 1926 Paris census whose statistics are detailed in Table 1. It contains 160 pages written in French, each page contains 30 lines. A page from POPP is presented in Fig. 1c.

Each row is divided in 10 columns: surname, name, birthdate, birthplace, nationality, civil status, link, education level, occupation, employer. In our experiments, we use the column name as a named entity. As a consequence, 100% words are associated to an entity. We perform experiments at two levels: text line and page.

Table 1. Statistics of the POPP dataset

				(b) Entities by split			
				Train	Validation	Test	
(a) Pages, lines, words, and entities by split				Surname	3,100	392	375
	Train	Validation	Test	First name	3,853	476	478
Pages	128	16	16	Birthdate	3,824	469	466
Lines	3,837	480	479	Location	4,789	600	584
Words	29,581	3,681	3,569	Nationality	283	17	30
Entities	29,581	3,681	3,569	Civil status	2,277	292	225
				Link	3,667	449	412
				Education level	25	4	12
				Occupation	4,488	529	535
				Employer	3,275	453	452

3.2 Methods

Three methods are introduced and compare in this work.

Two-Stage Workflow. The first method is a traditional two-stage workflow for information extraction that combines two steps. First, an HTR system is applied for text recognition on line-level images, then, SpaCy¹ [9] is used for named entity recognition. We compare two systems for the HTR task: PyLaia [18] and DAN [5].

- PyLaia² is an open source model for handwritten text recognition. It combines 4 convolutional layers and 3 recurrent layers, and is trained with the CTC loss function. The last layer is a linear layer with a softmax activation function that computes probabilities associated with each character of

¹ <https://spacy.io>.

² <https://github.com/jpuigcerver/PyLaia>.

the vocabulary. We use early stopping to avoid overfitting: the training is stopped after 50 epochs without improvement. PyLaia is trained on text line images.

- DAN³ is an open source attention-based Transformer model for handwritten text recognition that can work directly on images of paragraph or page. It is trained with the cross-entropy loss function. The last layer is a linear layer with a softmax activation function that computes probabilities associated with each character of the vocabulary. For each dataset, we train DAN on zones with the strongest semantic consistency: on records for ESPOSALLES, on pages for IAM, and lines for POPP.

For NER, we use SpaCy, a production-oriented NLP library that includes transformer-based pipelines with support of English (for IAM), Catalan (for ESPOSALLES), and French (for POPP). Like DAN, SpaCy is trained on records for ESPOSALLES, pages for IAM, and lines for POPP. For ESPOSALLES, we train two SpaCy models: one for the *category* label and one for the *person* label. Comparing two HTR systems with the same SpaCy model allows us to study the impact of handwriting recognition errors on the overall performance.

Integrated Workflow. The second method consists in training a model to recognize directly characters and NER tokens.


We train DAN models for this task, later referred to as *HTR+NER*. The model is trained at different levels to evaluate the impact of context: on lines and pages for IAM, on lines, records and pages for ESPOSALLES, on lines and pages for POPP. NER tokens are considered like characters by the network and are localized before relevant words, as illustrated in Table 2. For ESPOSALLES, we use a unique tag combining the *category* and *person* information (ex: <name_wife>Maria), as we found out that using two separate tags led to poorer performance. This observation is consistent with the findings of Carbonell et al. [3] and Rouhou et al. [21]. Finally, we also trained DAN with curriculum learning, e.g. trained for *HTR* and fine-tuned for *HTR+NER* and found out that the network reach similar performance. For clarity, we only provide results without curriculum learning.

Integrated Workflow with Key-Value Annotations. Our last experiment consists in training DAN on key-value annotations, so as to only predict relevant information with the relevant text and the corresponding named-entity. This task is referred to as *Key-value HTR+NER* in the rest of the article. To achieve this, words that are not linked to any entities are removed from transcriptions, as illustrated in Table 2. As a result, the model must learn to directly extract important words with their named entities, and ignore any other word. In this

³ <https://github.com/FactoDeepLearning/DAN>.

scenario, the two-stage approach cannot be used, as the full transcription is not available. This task is very challenging on IAM, as 90% of words are not linked to any entities, and more than 5% of pages do not have any entities. As a result, the training data is very sparse. The task is easier for ESPOSALLES, as 50% of words are linked to an entity. Finally, in POPP, every word is related to a named entity, so the *HTR+NER* and *Key-value HTR+NER* tasks are the same.

Table 2. Example of different transcriptions of the same record from the Esposalles database. Each transcription is used for a different task. *HTR*: the model predicts characters, *HTR+NER*: the model predicts characters and NER tokens, *Key-value HTR+NER*: the model predicts characters and NER tokens only for relevant words, ignoring words that are not associated with NER tokens.



Task	Transcription
<i>HTR</i>	dit dia rebere de Jua Oliveres pages de Lissa demunt viudo ab Maria donzella filla de Juan Pruna pages del far y de Beneta
<i>HTR+NER</i>	dit dia rebere de <N-H>Jua <SN-H>Oliveres <O-H>pages de <L-H>Lissa demunt <S-H>viudo ab <N-W>Maria <S-W>donsella filla de <N-WF>Juan <SN-WF>Pruna <O-WF>pages del <L-WF>far y de <N-WM>Beneta
<i>Key-value HTR+NER</i>	<N-H>Jua <SN-H>Oliveres <O-H>pages <L-H>Lissa <S-H>viudo <N-W>Maria <S-W>donsella <N-WF>Juan <SN-WF>Pruna <O-WF>pages <L-WF>far <N-WM>Beneta

4 Experimental Results

In this section, we introduce the evaluation metrics and present the results obtained on each dataset. We also compare our work with state-of-the-art methods and discuss the results.

4.1 Metrics

For all three datasets, performances are evaluated by the same standard character recognition and entity recognition metrics, as detailed in the following paragraphs. An additional metric is used to evaluate the experiments on ESPOSALLES.

4.2 HTR Metrics

The quality of handwriting recognition is evaluated using the character error rate (CER) and word error rate (WER). The full text is evaluated, and named entity tokens are ignored in integrated methods at this step of the evaluation.

4.3 NER Metrics

We use the Nerval⁴ evaluation toolkit to evaluate named entity recognition results. In Nerval [14], the automatic transcription is aligned with the ground truth at character level. Predicted and ground truth words are considered a match if their edit distance is less than 30%. From this alignment, precision, recall and F1-score are computed.

4.4 IEHHR Metrics

Finally, for the ESPOSALLES dataset, we also compute the IEHHR metric that was introduced in the ICDAR 2017 Competition on Information Extraction in Historical Handwritten [8]. This metric jointly evaluates HTR and NER. Only words associated with named entities are taken into account in this evaluation. The “basic” score is equal to 100-CER if the *category* tag is correct, 0 otherwise. The “complete” score is equal to 100-CER if both the *category* and *person* tags are correct, 0 otherwise.

4.5 Evaluation Results

We present handwritten text recognition results in Table 3 and named entity recognition results in Table 4. For ESPOSALLES, we also provide the results for information extraction in Table 5 and obtain state-of-the-art results on the public IEHHR benchmark⁵.

What is the Best Model for HTR? Results in Table 3 show that DAN is better than PyLaia for HTR on all three datasets. The DAN model trained only for *HTR* is generally better than the model directly trained for *HTR+NER*. The results show that DAN is always better than PyLaia for handwriting recognition: CER and WER are always lower with DAN. The WER reaches 1.37% on ESPOSALLES, 13.66% on IAM, and 18.09% on POPP. Finally, we note that DAN can be more performant on larger text zones. Indeed, DAN performs better on pages on IAM, and on records on ESPOSALLES. On the other hand, on POPP, the best performances are obtained on text lines. This observation can be explained by the fact that POPP documents are tables in which the lines are independent.

⁴ <https://gitlab.com/tekliia/ner/nerval>.

⁵ <https://rrc.cvc.uab.es/?ch=10&com=evaluation&task=1>.

Table 3. Evaluation results for handwritten text recognition on IAM, ESPOSALLES, and POPP. Results are given for test sets. NER tokens are not taken into account for this evaluation.

(a) IAM (RWTH split)				
Model	Task	CER (%)	WER (%)	Input
VAN [6]	<i>HTR</i>	4.45	14.55	Line
PyLaia	<i>HTR</i>	7.79	24.73	Line
DAN	<i>HTR</i>	4.30	13.66	Page
DAN	<i>HTR+NER</i>	5.12	16.17	Line
DAN	<i>HTR+NER</i>	4.82	14.61	Page
(b) ESPOSALLES				
Method	Task	CER (%)	WER (%)	Input
Seq2seq [23]	<i>HTR</i>	2.82	8.33	Line
Seq2seq [23]	<i>HTR+NER</i>	1.81	6.10	Line
PyLaia	<i>HTR</i>	0.76	2.62	Line
DAN	<i>HTR</i>	0.46	1.37	Record
DAN	<i>HTR+NER</i>	0.48	1.75	Line
DAN	<i>HTR+NER</i>	0.39	1.51	Record
DAN	<i>HTR+NER</i>	3.61	4.23	Page
(c) POPP				
Model	Task	CER (%)	WER (%)	Input
VAN [4]	<i>HTR</i>	7.08	19.05	Line
PyLaia	<i>HTR</i>	17.19	37.43	Line
DAN	<i>HTR</i>	8.18	18.09	Line
DAN	<i>HTR+NER</i>	7.83	24.57	Line
DAN	<i>HTR+NER</i>	11.74	30.78	Page

Table 4. Evaluation results for named entity recognition on IAM, ESPOSALLES, and POPP. Results are given for test sets. Evaluation results are computed using Nerval, which computes an alignment between ground truth and predicted entities.

(a) IAM (RWTH split)						
Method	P (%)	R (%)	F1 (%)	Input Type		
Tülselmann et al.* [26]	60.4	50.9	54.2	Word/Record		
Rowtula et al.* [22]	33.8	30.9	32.3	Word/Record		
Todolo et al.* [25]	26.4	10.8	14.9	Word/Record		
Dessurt [7]	-	-	40.4	Page		
Ground-truth + SpaCy	74.9	76.2	75.5	-/Page		
PyLaia + SpaCy	56.5	49.0	52.5	Line/Page		
DAN + SpaCy	61.8	57.9	59.8	Page/Page		
DAN	37.1	30.8	33.7	Line		
DAN	37.2	27.0	31.3	Page		
DAN	0	0	0	Page (key-value)		

* Different computation method due to pre-existing word alignment.

(b) ESPOSALLES							
Method	Person			Category			Input Type
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)	
Tülselmann et al.* [26]	99.3	99.2	99.3	98.5	98.2	98.3	Word/Record
Rowtula et al.* [22]	97.0	96.2	96.6	97.1	97.0	97.0	Word/Record
Todolo et al.* [25]	98.5	97.8	98.1	98.5	97.8	98.1	Word/Record
Ground-truth + SpaCy	98.6	98.4	98.5	98.3	98.7	98.5	-/Record
PyLaia + SpaCy	95.9	94.0	94.9	95.6	94.3	95.0	Line/Record
DAN + SpaCy	97.9	97.9	97.9	97.6	98.1	97.8	Record/Record
DAN	96.0	96.1	96.1	96.9	97.0	96.9	Line
DAN	97.9	98.2	98.1	97.4	97.8	97.6	Record
DAN	95.0	98.4	96.6	94.2	97.6	95.9	Page
DAN	97.0	97.4	97.2	96.7	97.1	96.9	Record (key-value)

* Different computation method due to pre-existing word alignment.

(c) POPP				
Method	P (%)	R (%)	F1 (%)	Input type
Ground-truth + SpaCy	95.6	97.3	96.4	-/Line
PyLaia + SpaCy	75.6	77.0	76.3	Line/Line
DAN + SpaCy	82.8	85.3	84.0	Line/Line
DAN	85.6	86.2	85.9	Line
DAN	83.8	86.9	85.3	Page

Table 5. IEHHR scores given for the test set of ESPOSALLES dataset.

Method	Basic (%)	Complete (%)	Input Type
Baseline HMM [8]	80.28	63.11	Line/Line
CITlab ARGUS-1 [8]	89.54	89.17	Line/Line
CITlab ARGUS-2 [8]	91.63	91.19	Line/Line
CITlab ARGUS-3 [8]	91.94	91.58	Line/Line
CVC [25]	90.59	89.40	Line/Line
Naver Labs [17]	95.46	95.03	Line/Line
IRISA [23]	94.7	94.0	Line
IRISA multi-task [23]	95.2	94.4	Line
InstaDeep GNN/Transformer ⁸	96.22	96.24	Record
InstaDeep Transformer [21]	96.25	95.54	Record
TEKLI Kaldi + Flair [15]	96.96	-	Line/Record
Ground-truth + SpaCy	97.51	97.57	-/Record
PyLaia + SpaCy	96.58	96.58	Line/Record
DAN + SpaCy	97.13	97.11	Record/Record
DAN	96.26	94.47	Line
DAN	97.03	96.93	Record
DAN	95.45	95.04	Page
DAN (key-value)	96.48	96.31	Record (key-value)

What is the Impact of HTR Errors on NER? Results in Table 4 help us understand the impact of handwriting recognition errors on NER performance. The second block of each subtable compares the results using ground transcription or predicted transcriptions (PyLaia or DAN). On ESPOSALLES, both HTR systems are very performant with CER below 1%. As a result, NER performance remains very good. However, on IAM and POPP, PyLaia and DAN yield a higher CER. As a consequence, the F1 score drops by 15 points for a 5% CER on IAM, and by 10 points for a 10% CER on POPP.

What is the Best Approach for Information Extraction? The best performance on IAM is achieved with a two-stage method, combining DAN (HTR) and SpaCy (NER). These results support the observations of Tüselmann et al. [26], and can be explained because there are few entities in the dataset. As a result, DAN struggles to learn semantic information, while SpaCy benefits from pre-trained embeddings for the English language. However, on POPP, DAN trained for *HTR+NER* outperforms the two-stage approach combining DAN and SpaCy, although SpaCy does benefit from pre-trained French embeddings. There are two possible explanations for this result. First, POPP documents contain mostly names and surnames, which may not be included in the embeddings. Second, since these are tabular documents, word localization determines the semantic

category, as each column corresponds to a specific named entity. Unlike DAN, SpaCy does not have any information regarding the word localization. Finally, on ESPOSALLES, both approaches yield similar results: SpaCy recognizes the *category* labels better while DAN recognizes the *person* labels better.

What is the Performance of Segmentation-Free Methods? It is interesting to note that DAN often performs better on pages (IAM) or records (ESPOSALLES) than on text lines. And yet, the text recognition task is traditionally done on text lines, which requires prior automatic or manual segmentation. But manual segmentation is time-consuming, and automatic segmentation can introduce many errors that affect the performance of handwriting or named entity recognition [15]. Therefore, results presented on pages cannot be directly compared to the results on text lines or records, as the task is much harder. In order to compare these results fairly, segmentation-based workflows should be evaluated on automatically segmented text lines or records. It is likely that segmentation-free workflows will outperform segmentation-based workflows in an end-to-end evaluation setting.

Is DAN Able to Learn from Key-Value Annotations? Finally, we evaluate the ability of DAN to learn from key-value annotations. On ESPOSALLES, where 50% of words are linked to an entity, DAN manages to learn from key-value annotations. It learns to recognize relevant words and to ignore the others. Although its performances are slightly lower than when trained with full transcripts, they remain very competitive. In contrast, DAN fails to learn on IAM, in which only 10% of words are linked to an entity. The model can be trained for a few epochs before overfitting. As a result, it does not predict anything on the test set. Finally, on POPP, all words are linked to an entity, so this experiment is similar to the one with full annotations, as there are no words to ignore during training.

5 Conclusion

In this paper, we focus on information extraction in digitized handwritten documents. We compare an integrated approach trained for joint HTR and NER with a traditional two-stage approach that performs HTR before NER. We present results at different levels: pages, paragraphs and lines and reach state-of-the-art performance on three datasets.

Our experiments show that integrated approaches trained jointly for HTR and NER can outperform two-stage approaches when word localization has an impact on the NER label (POPP). As opposed, two-stage approaches are better when applied on datasets with few entities (IAM) as they can benefit from pre-trained embeddings. In other cases (ESPOSALLES), two-stage and integrated approaches reach similar performance 97.11% and 96.93% respectively, for the complete IEHHR score on records on ESPOSALLES. We also demonstrate that

applying these models directly on pages leads to very acceptable performances, either better than when applied on lines (ESPOSALLES, IAM), or with a minor performance loss (POPP). The interest of this method is enhanced by the lack of need for prior automatic segmentation, which is known to impact handwriting recognition performances [15]. Finally, we show that, under certain conditions, integrated methods are able to learn from key-value annotations, e.g. from a list of relevant words with their corresponding named entities. On ESPOSALLES, the model trained on key-value annotations reaches a complete recognition score of 96.31%. This observation is encouraging as it would allow training models from incomplete information manually, which considerably reduces the effort needed for manual transcription.

In future works, we are interested in measuring the impact of segmentation errors when evaluating end-to-end systems for information extraction. We also would like to identify the conditions needed to train a model on key-value annotations. Finally, we want to improve DAN for the task of information extraction. For example, the training loss could also be adapted to differentiate NER tokens from characters. Performance could be improved by using pre-trained embeddings like in SpaCy. Since DAN and SpaCy rely on character- and word-embeddings respectively, it would be interesting to find a common representation at sub-word level.

Appendix

Detailed splits for IAM and Esposalles

We provide the detailed splits used for IAM in Table 6 and ESPOSALLES in Table 7. For IAM, we use the RWTH split. For ESPOSALLES, we use the official split, with 25% of training data used for validation.

Impact of curriculum learning

We evaluate the impact of curriculum learning for the task of *HTR+NER* in Table 8 and 9. The DAN model trained with curriculum learning is pre-trained on the *HTR* task, then fine-tuned on the *HTR+NER* task. The results show that curriculum learning does not always have a positive impact on final performances.

Table 6. Statistics of the IAM dataset (RWTH split)

				(b) Entities by split		
				Train	Validation	Test
				Person	1,399	252 603
				GPE	731	38 129
				Organization	825	39 100
				NORG	282	19 79
				Date	1,000	57 178
(a) Pages, lines, words, and entities by split				Cardinal	409	75 130
	Train	Validation	Test	Work of Art	294	41 110
Pages	747	116	336	Time	167	24 114
Lines	6,482	976	2,915	FAC	126	37 71
Words	55,111	8,900	25,931	Quantité	107	17 66
Entities	5,868	654	1,713	Location	124	16 41
				Ordinal	104	19 38
				Product	78	6 24
				Percent	91	6 4
				Event	61	2 15
				Law	43	6 0
				Language	15	0 5
				Money	12	0 6

Table 7. Statistics of the ESPOSALLES dataset

				(b) Entities by split		
				Train	Validation	Test
(a) Pages, records, lines, words, and entities by split				Name	3,774	1,223 1,312
	Train	Validation	Test	Surname	2,033	634 694
Pages	75	25	25	Location	3,440	1,069 1,087
Records	731	267	253	Occupation	2,273	737 797
Lines	2,328	742	757	State	868	274 319
Words	23,893	7,608	8,026	Wife	2,093	678 768
Entities	12,388	3,937	4,238	Wife's father	2,745	847 908
				Wife's mother	566	188 189
				Husband	4,334	1,493 1,563
				Husband's father	1,838	476 518
				Husband's mother	462	1401 156
				Other person	350	115 136

For POPP, we also trained the model for key-value *HTR+NER* in a random order, e.g. with named entities in a random order. Results show that DAN is also able to learn with a random reading order, although the error rates are a bit higher than when the model is trained with the correct reading order.

Table 8. Impact of curriculum learning on handwritten text recognition on IAM, ESPOSALLES, and POPP. Results are given for test sets. NER tokens are not taken into account for this evaluation.

(a) IAM (RWTH split)				
Model	Task	CER (%)	WER (%)	Input
DAN	<i>HTR</i>	4.86	15.78	Line
DAN	<i>HTR</i>	4.30	13.66	Page
DAN	<i>HTR+NER</i>	5.12	16.17	Line
DAN curriculum	<i>HTR+NER</i>	5.01	16.32	Line
DAN	<i>HTR+NER</i>	4.82	14.61	Page
DAN curriculum	<i>HTR+NER</i>	4.30	13.65	Page
(b) ESPOSALLES				
Method	Task	CER (%)	WER (%)	Input
DAN	<i>HTR</i>	0.54	2.13	Line
DAN	<i>HTR</i>	0.46	1.37	Record
DAN	<i>HTR</i>	2.77	3.58	Page
DAN	<i>HTR+NER</i>	0.48	1.75	Line
DAN curriculum	<i>HTR+NER</i>	0.64	2.02	Line
DAN	<i>HTR+NER</i>	0.39	1.51	Record
DAN curriculum	<i>HTR+NER</i>	0.89	1.97	Record
DAN	<i>HTR+NER</i>	3.61	4.23	Page
DAN curriculum	<i>HTR+NER</i>	2.23	3.15	Page
(c) POPP				
Model	Task	CER (%)	WER (%)	Input
DAN	<i>HTR</i>	8.18	18.09	Line
DAN	<i>HTR+NER</i>	7.83	24.57	Line
DAN curriculum	<i>HTR+NER</i>	8.06	24.85	Line
DAN curriculum + random order	<i>HTR+NER</i>	9.53	27.01	Line

Table 9. Impact of curriculum learning for named entity recognition on IAM, ESPOSALLES, and POPP. Results are given for test sets. Evaluation results are computed using Nerval, which computes an alignment between ground truth and predicted entities.

(a) IAM (RWTH split)						
Method	P (%)	R (%)	F1 (%)	Input Type		
DAN	37.1	30.8	33.7	Line		
DAN curriculum	33.0	23.3	27.3	Line		
DAN	37.2	27.0	31.3	Page		
DAN curriculum	38.2	29.1	33.1	Page		

(b) ESPOSALLES							
Method	Person			Category			Input Type
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)	
DAN	96.0	96.1	96.1	96.9	97.0	96.9	Line
DAN curriculum	95.6	94.0	94.8	96.3	95.5	95.9	Line
DAN	97.9	98.2	98.1	97.4	97.8	97.6	Record
DAN curriculum	97.3	97.5	97.4	96.5	97.3	96.9	Record
DAN	95.0	98.4	96.6	94.2	97.6	95.9	Page
DAN curriculum	96.4	97.3	96.9	95.4	97.2	96.3	Page
DAN	97.0	97.4	97.2	96.7	97.1	96.9	Record (key-value)
DAN curriculum	96.7	96.3	96.5	96.0	96.1	96.0	Record (key-value)

(c) POPP					
Method	P (%)	R (%)	F1 (%)	Input type	
DAN	85.6	86.2	85.9	Line	
DAN curriculum	85.4	86.2	85.8	Line	
DAN curriculum + random order	84.6	84.8	84.7	Line	

References

1. Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., Vollgraf, R.: FLAIR: an easy-to-use framework for state-of-the-art NLP. In: 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)(Demonstrations), pp. 54–59 (2019)
2. Carbonell, M., Fornés, A., Villegas, M., Lladós, J.: A neural model for text localization, transcription and named entity recognition in full pages. *Pattern Recogn. Lett.* **136**, 219–227 (2020). <https://doi.org/10.1016/j.patrec.2020.05.001>
3. Carbonell, M., Villegas, M., Fornés, A., Lladós, J.: Joint recognition of handwritten text and named entities with a neural end-to-end model. In: 2018 13th IAPR International Workshop on Document Analysis Systems (DAS), pp. 399–404. IEEE Computer Society, Los Alamitos, CA, USA, April 2018. <https://doi.org/10.1109/DAS.2018.52>
4. Constum, T., et al.: Recognition and information extraction in historical handwritten tables: toward understanding early 20th century Paris census. In: 15th International Workshop on Document Analysis Systems (DAS), pp. 143–157, May 2022. https://doi.org/10.1007/978-3-031-06555-2_10

5. Coquenet, D., Chatelain, C., Paquet, T.: DAN: a segmentation-free document attention network for handwritten document recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 1–17 (2023). <https://doi.org/10.1109/TPAMI.2023.3235826>
6. Coquenet, D., Chatelain, C., Paquet, T.: End-to-end handwritten paragraph text recognition using a vertical attention network. *IEEE Trans. Pattern Anal. Mach. Intell.* **45**, 508–524 (2023). <https://doi.org/10.1109/TPAMI.2022.3144899>
7. Davis, B., Morse, B., Price, B., Tensmeyer, C., Wigington, C., Morariu, V.: End-to-end Document Recognition and Understanding with Dessurt (2022). <https://doi.org/10.48550/ARXIV.2203.16618>
8. Fornés, A., Romero, V., Baro, A., Toledo, J., Sánchez, J.A., Vidal, E., Lladós, J.: ICDAR2017 competition on information extraction in historical handwritten records. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), pp. 1389–1394, November 2017. <https://doi.org/10.1109/ICDAR.2017.227>
9. Honnibal, M., Montani, I., Van Landeghem, S., Boyd, A.: spaCy: Industrial-strength Natural Language Processing in Python (2020). <https://doi.org/10.5281/zenodo.1212303>
10. Kang, L., Toledo, J.I., Riba, P., Villegas, M., Fornés, A., Rusiñol, M.: Convolve, attend and spell: an attention-based sequence-to-sequence model for handwritten word recognition. In: German Conference on Pattern Recognition, pp. 459–472 (2019)
11. Kiessling, B., Tissot, R., Stokes, P., Stökl Ben Ezra, D.: eScriptorium: an open source platform for historical document analysis. In: 2019 International Conference on Document Analysis and Recognition Workshops (ICDARW) (2019). <https://doi.org/10.1109/ICDARW.2019.10032>
12. Lombardi, F., Marinai, S.: Deep learning for historical document analysis and recognition—a survey. *J. Imaging* **6**, 110 (2020)
13. Marti, U.V., Bunke, H.: The IAM-database: an English sentence database for offline handwriting recognition. *Int. J. Doc. Anal. Recogn.* **5**, 39–46 (2002). <https://doi.org/10.1007/s100320200071>
14. Miret, B., Kermorvant, C.: Nerval: a python library for named-entity recognition evaluation on noisy texts (2021). <http://gitlab.com/tekli/ner/nerval>
15. Monroc, C.B., Miret, B., Bonhomme, M.L., Kermorvant, C.: A comprehensive study of open-source libraries for named entity recognition on handwritten historical documents. In: Document Analysis Systems, pp. 429–444 (2022). https://doi.org/10.1007/978-3-031-06555-2_29
16. Muehlberger, G., et al.: Transforming scholarship in the archives through handwritten text recognition: Transkribus as a case study. *J. Doc.* **75**, 954–976 (2019). <https://doi.org/10.1108/JD-07-2018-0114>
17. Prasad, A., Déjean, H., Meunier, J., Weidemann, M., Michael, J., Leifert, G.: Bench-marking information extraction in semi-structured historical handwritten records. In: CoRR (2018). <http://arxiv.org/abs/1807.06270>
18. Puigcerver, J.: Are multidimensional recurrent layers really necessary for handwritten text recognition? In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), vol. 01, pp. 67–72 (2017). <https://doi.org/10.1109/ICDAR.2017.20>
19. Qi, P., Zhang, Y., Zhang, Y., Bolton, J., Manning, C.D.: Stanza: a python natural language processing toolkit for many human languages. In: 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 101–108, January 2020. <https://doi.org/10.18653/v1/2020.acl-demos.14>

20. Romero, V., et al.: The ESPOSALLES database: an ancient marriage license corpus for off-line handwriting recognition. *Pattern Recogn.* **46**, 1658–1669 (2013). <https://doi.org/10.1016/j.patcog.2012.11.024>
21. Rouhou, A.C., Dhiaf, M., Kessentini, Y., Salem, S.B.: Transformer-based approach for joint handwriting and named entity recognition in historical document. *Pattern Recogn. Lett.* **155**, 128–134 (2022). <https://doi.org/10.1016/j.patrec.2021.11.010>
22. Rowtula, V., Krishnan, P., Jawahar, C.V.: POS tagging and named entity recognition on handwritten documents. In: *Proceedings of the 15th International Conference on Natural Language Processing* (2018)
23. Tarride, S., Lemaitre, A., Couasnon, B., Tardivel, S.: A comparative study of information extraction strategies using an attention-based neural network. In: *Document Analysis Systems*, pp. 644–658 (2022). https://doi.org/10.1007/978-3-031-06555-2_43
24. Tarridea, S., et al.: Large-scale genealogical information extraction from handwritten Quebec parish records. *Int. J. Document Anal. Recogn.* (2023)
25. Toledo, J.I., Carbonell, M., Fornés, A., Lladós, J.: Information extraction from historical handwritten document images with a context-aware neural model. *Pattern Recogn.* **86**, 27–36 (2019). <https://doi.org/10.1016/j.patcog.2018.08.020>
26. Tüselmann, O., Wolf, F., Fink, G.A.: Are end-to-end systems really necessary for ner on handwritten document images? In: *Document Analysis and Recognition - ICDAR 2021*, pp. 808–822 (2021). https://doi.org/10.1007/978-3-030-86331-9_52
27. Vidal, E., et al.: The Carabela project and manuscript collection: large-scale probabilistic indexing and content-based classification. In: *In proceedings of the 17th International Conference on Frontiers in Handwriting Recognition (ICFHR 2020)* (2020)
28. Yousef, M., Bishop, T.: OrigamiNet: weakly-supervised, segmentation-free, one-step, full page text recognition by learning to unfold. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14698–14707, June 2020