

Received May 5, 2021, accepted May 12, 2021, date of publication May 21, 2021, date of current version June 2, 2021.

Digital Object Identifier 10.1109/ACCESS.2021.3082689

# Boosting Offline Handwritten Text Recognition in Historical Documents With Few Labeled Lines

JOSÉ CARLOS ARADILLAS<sup>1</sup>, JUAN JOSÉ MURILLO-FUENTES<sup>1</sup>, (Senior Member, IEEE),  
AND PABLO M. OLMOS<sup>2,3</sup>, (Member, IEEE)

<sup>1</sup>Departamento de Teoría de la Señal y Comunicaciones, Escuela Técnica Superior de Ingeniería, Universidad de Sevilla, 41004 Sevilla, Spain

<sup>2</sup>Departamento de Teoría de la Señal y Comunicaciones, Universidad Carlos III de Madrid, 28903 Madrid, Spain

<sup>3</sup>Gregorio Marañón Health Research Institute, 28007 Madrid, Spain

Corresponding author: Juan José Murillo-Fuentes (murillo@us.es)

This work was supported in part by the Spanish Government Ministerio de Educación (MEC) under Grant FPU16/04190 and Project MINECOTEC2016-78434-C3-2-3-R; in part by the Comunidad de Madrid under Grant IND2017/TIC-7618, Grant IND2018/TIC-9649, Grant IND2020/TIC-17372, and Grant Y2018/TCS-4705; in part by the Banco Bilbao Vizcaya Argentaria (BBVA) Foundation through the Domain Alignment and Data Wrangling with Deep Generative Models (Deep-DARWiN) Project; and in part by the European Union by the Fondo Europeo de Desarrollo Regional (FEDER) and the European Research Council (ERC) through the European Union Horizon 2020 Research and Innovation Program under Grant 714161.

**ABSTRACT** In this paper we address the problem of offline handwritten text recognition (HTR) in historical documents when few labeled samples are available and some of them contain errors in the train set. Our three main contributions are: first, we analyze how to perform transfer learning (TL) from a massive database to a smaller historical database, analyzing which layers of the model need fine-tuning. Second, we analyze methods to efficiently combine TL and data augmentation (DA). Finally, we propose an algorithm to mitigate the effects of incorrect labeling in the training set. The methods are analyzed over the ICFHR 2018 competition database, Washington and Parzival. Combining all these techniques, we demonstrate a remarkable reduction of CER (up to 6 percentage points in some cases) in the test set with little complexity overhead.

**INDEX TERMS** Connectionist temporal classification (CTC), convolutional neural networks (CNN), data augmentation (DA), deep neural networks (DNN), historical documents, long-short-term-memory (LSTM), offline handwriting text recognition (HTR), outlier detection, transfer learning.

## I. INTRODUCTION

The transcription of historical manuscripts is paramount for a better understanding of our history, as it allows for direct access to the contents, quite facilitating searches and studies. Also, the classification and indexing of transcript text can be easily automated. Handwritten text recognition (HTR) tasks in historical datasets have been studied by many authors in the last few years [1]–[10]. In HTR, transcribing each author can be considered a different task, since the distribution of both model input and output varies from writer to writer. At the input, we have variations not only in the calligraphy but also, depending on the digitization process, in the image resolution, contrast, color, or background. On the other hand, at the output, the labels usually correspond to different languages

and historical periods, with differences in the character set, the semantics and the lexicon.

Usually, the process for automatic transcription of a document comprises 4 phases: 1) digitization of the document to obtain an image of every page in the document in electronic format; 2) segmentation of each page into corresponding regions with lines of text; 3) transcription of each line of text and finally 4) application of a dictionary or language model to correct errors in the transcription of texts as well as in the composition of the complete texts from the lines obtained in step 3).

While segmentation is an important issue in HTR [2], [11]–[13], this paper focuses on the transcription phase. In recent years, there has been a trend towards models based on deep neural networks (DNNs) [8]. Once the DNN model to be used has been designed, an enormous number of training samples are required to minimize the number of transcription errors, measured in character error rate

The associate editor coordinating the review of this manuscript and approving it for publication was Gang Mei<sup>1</sup>.

(CER) or word error rate (WER), given by the Levenshtein distance [14] between the ground truth (GT) and the output of the model.

However, in reality most of the times we might only have a limited number of lines for a given author and document. Besides, transcription of part of the documents to get labeled samples is expensive either in time or money. Take [1] as an example, where the manual transcription process of a document by an expert in paleography took an average of 35 minutes per page. In this scenario, allowing for a reduction in the transcript needed would greatly improve the viability and cost of the process.

The goal of the methods proposed in this paper is to significantly reduce the number of annotated lines, thus reducing the monetary cost of the process. Contributions are as follows:

- A DNN model with transfer learning (TL) by studying in detail which layers should be retained and which ones should be retrained. We found this model to provide the best performance within the models using long-short-term-memory (LSTM).
- A thoughtful analysis of HTR when little labeled data is available, studying the evolution with the number of transcript lines, highlighting how sensitive the models are to the number of lines in the train set when this number is low.
- A model combining TL and data augmentation (DA). We show that for a reduced number of lines, using DA can be counterproductive.
- An algorithm to detect and avoid mislabeled lines is proposed and tested, with good results in small training datasets. Furthermore, in some scenarios, wrong labels are fixed.

The remaining of the paper is organized as follows: in Section II previous works reporting solutions to the problem of HTR over small datasets are summarized; in Section III the architecture and models used in this paper are presented; in Section IV the application of TL and DA for HTR is analyzed; in Section V an algorithm is proposed to detect and prone mislabeled lines in the training set; the paper ends with Section VI, drawing main conclusions.

## II. RELATED WORK AND CONTRIBUTIONS

### A. DNN MODEL

State-of-the-art architectures combine a convolutional neural network (CNN) [15] with a recurrent neural network (RNN) with LSTM cells [16]. This type of network models the conditioned probability,  $p(l|x)$ , of a character sequence of arbitrary length,  $l$ , given an image,  $x$ , of fixed height and arbitrary width. Note that varying the number of characters yields a new design of the last layer of the DNN model. These models are configured to minimize the connectionist temporal classification (CTC) cost function proposed by Graves in [17]. In some works 2D-LSTM [16] networks are used [18]–[21]. This RNN has two main drawbacks. On the one hand, it has an extremely large number of parameters that make learning

difficult. On the other hand it is not parallelizable [22]. For these reasons, it has been discarded here.


Regarding the state-of-the-art DNN models for HTR, some recent works are in the line of avoiding recurrence in the models, developing models based in fully-convolutional networks such as the (Gated) Fully Convolutional Networks (G)FCN [23]–[26]. This kind of model reduces the number of parameters in the architecture.

### B. TRANSFER LEARNING

In the HTR problem with a reduced training set, TL was applied by Soullard *et al.* in [7]. The main idea behind TL is initializing the parameters of a model by those learned from a huge dataset beforehand, denoted as *source*. Then, the available labeled set of samples of the dataset of interest, the *target*, is used to refine the parameters of the model, usually just a subset of them [27]–[30]. In [30], they analyze how to reduce the dataset shift and enhance the feature transferability in task-specific layers of deep networks. Hence, with TL we start learning a different task to avoid learning the whole set of parameters from scratch, preventing overfitting and favoring convergence. In [7], they proposed a method that applies TL in both the optical and the language model. In this and other similar previous proposals on TL, the authors applied DA in both training and test steps.

### C. DATA AUGMENTATION

DA consists in augmenting the training set with synthetically generated samples. Similar to TL, it reduces the tendency to overfit when training models with a large number of parameters and limited labeled data. In DA for image classification problems, the training set is increased by modifying the original images through transformations such as scaling, rotation, or flipping images, among others [31]. Several authors have proposed specific DA techniques for HTR: in [32] the authors apply methods for augmentation and normalization to improve HTR by allowing the network to be more tolerant of variations in handwriting by profile normalization. In [33] they show some affine transformation methods for data augmentation in HTR. In [34] and [35] they synthesize new lines images by concatenating characters from different datasets. [34] does it from cursive characters while in [35] they do it from a database of handwritten Chinese characters. Similar to [32], in [36] they also apply some elastic distortions to the original images. In [4] the authors improve the performance by augmenting the training set with specially crafted multi-scale data. They also propose a model-based normalization scheme that considers the variability in the writing scale at the recognition phase. In these works, they apply DA in relatively large well-known datasets, but here we show that the regularization effect of any DA technique has no impact when doing the fine-tuning adaptation to a singular writer in small databases. Accordingly, we conclude that the combination of TL and DA applied to small datasets has to be done carefully, to reduce the final error.



**FIGURE 1.** The adapted CRNN architecture from [39] used as baseline. The number of channels of each CNN layer is shown in this scheme. Pooling layers after the first, second and third CNN layer are also depicted.

#### D. MISLABELED SAMPLES

Mislabeled detection in HTR has been seldom faced. In [37] they face a specific problem in the IAM database: crossed out words that are labeled with the symbol “#”. The authors propose a method to avoid how this specific label affects the performance. That method is focused on the specific problem of crossed-out text and how it is annotated in the GT. The algorithm we propose in Section V is more general, addressing this and other possible problems. In further related work in [38], the authors apply a method to align the output of a segmentation process with the available GT.

#### III. ARCHITECTURE AND DATABASES

In the HTR pipeline, there are several ways to improve performance of a DNN model: preprocessing steps, the architecture used, regularization techniques, optimization, language model and dictionary, among others. The methods proposed in this paper are developed for an exemplary state-of-the-art DNN architecture but can be easily used in the pipeline of any other HTR system to reduce transcript errors. For fair comparison, in this paper, we use the same DNN model for all the experiments. Extra correction steps such as adding a language model (LM) are not included but could be applied to further improve the performance. In this section, we focus on the selection of the model of the DNN and the databases used.

#### A. ARCHITECTURE

In this work, we implement a network architecture based on the convolutional recurrent neural network (CRNN) presented in [39]. This approach avoids the use of two-dimensional LSTM (2D-LSTM) layers, applying convolutional layers as feature extractors and a stack of 1D bidirectional LSTM (BLSTM) layers to perform classification. Previous DNN architectures for HTR consisted of a

combination of 2D-LSTM layers and convolutional layers, with a collapsing stage before the output layer in order to reshape the features tensors from 2D to 1D [18], [19]. The use of 2D-LSTM layers at the first stages has several drawbacks such as the need of more memory in the allocation of activations and buffers during back-propagation and a longer runtime is required to train the networks since parallel computation cannot be implemented in contrast to a CNN [22]. Recently, it has been shown that CNN in the lower layers of an HTR system obtains similar features than an RNN containing 2D-LSTM units [22].

The CRNN architecture proposed in [39] is comprised of seven CNN with a max-pooling step at the output of four of them, followed by a stack of two BLSTM layers at the top of the network. In [6] we have shown that the CRNN in Fig. 1, the one used in this work,<sup>1</sup> achieves better performance than the original one proposed in [39]. It uses a CNN with 5 layers at the bottom, with a  $3 \times 3$  and  $1 \times 1$  stride kernel, the number of filters are 16, 32, 48, 64 and 80, respectively. We use LeakyReLU as the activation function. A  $2 \times 2$  max-pooling is also applied at the output of the first 3 layers, to reduce the size of the input sequence. At the output of the CNN, a column-wise concatenation is carried out with the purpose of transforming the 3D tensors of size  $w \times h \times d$  (width  $\times$  height  $\times$  depth) into 2D tensors of size  $w \times (h \times d)$  where  $w$  and  $h$  are the width and height of the input image divided by 8, i.e., after 3 stages of  $2 \times 2$  max-pooling. The depth,  $d = 80$ , is the number of features of the last CNN layer. Therefore, at the output of the CNN, we have sequences of length  $w$  and depth  $h \times 80$  features.

After the CNN stage, 5 1D BLSTM recurrent layers of 256 units with hyperbolic tangent functions and without peephole connections.. Since at the output of each BLSTM layer we

<sup>1</sup>Implementation is publicly available in <https://github.com/josarajar/HTRTF>

GT: those in authority to find some simple

**FIGURE 2.** IAM handwritten text sample: image of a line and its transcript.

GT: As there are several Contracts made by me to

**FIGURE 3.** Washington handwritten text sample: image of a line and its transcript.

have 256 features in each direction, we perform a depth-wise concatenation to adapt the input of the next layer, to the overall size of 512. Dropout regularization [18], [40] is applied at the output of every layer, except for the first convolutional one, with rates 0.2 for the CNN layers and 0.5 for the BLSTM layers.

Finally, each column of features after the 5th BLSTM layer, with depth 512, is mapped into the  $L + 1$  output labels with a fully connected layer, where  $L$  is the number of characters in the alphabet of each database, e.g., 79, 83, 96 or 102 in the IAM, Washington, Parzival or International Conference on Frontiers in Handwriting Recognition (ICFHR) 2018 Competition databases, respectively. The additional dimension is needed for the blank symbol of the CTC [17], which concludes this architecture. Overall, this CNN-BLSTM-CTC model has  $9.58 \times 10^6$  parameters roughly, depending on the number of characters in each database.

The architecture is implemented in the open-source framework TensorFlow in Python, using the GPU-enabled version. We use the Adam algorithm, a learning rate of 0.003,  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . The parameters are updated using the gradients of the CTC loss on each batch of 16 text lines. We apply an early stopping criterion of 10 epochs without average improvement.

The selected model was the one with the best performance out of the 7 + 3, 8 + 0, 4 + 4, 5 + 5, and 6 + 6 configurations, where  $A + B$  corresponds to  $A$  convolutional followed by  $B$  BLSTM layers. On the other hand, for the CTC we tried best path decoding and beam search decoding, with no significant improvement of the later, despite its computational complexity.

## B. DATABASES

In this paper we focus on HTR over eight databases: IAM [41], Washington [42], Parzival [42], and the five ones provided at the ICFHR 2018 Competition [5].

### 1) THE IAM DATABASE

The IAM database [41] contains 13353 labeled text lines of modern English handwritten by 657 different writers. The images were scanned at a resolution of 300 dpi and saved as PNG images with 256 gray levels. An image of this database is included in Fig. 2 alongside the GT transcript. The database

is partitioned into training, validation, and test sets of 6161, 900, and 2801 lines, respectively.<sup>2</sup> Here, the validation and test sets provided are merged in a unique test set. There are 79 different characters in this database, including capital and small letters, numbers, some punctuation symbols, and white space.

### 2) THE RIMES DATABASE

The RIMES database is a collection of French letters handwritten by 1,300 volunteers who have participated in the RIMES database creation by writing up to 5 emails. The RIMES database thus contains 12,723 pages corresponding to 5605 mails of two to three pages. In our experiments, we take a set of 12111 lines extracted from the International Conference on Document Analysis and Recognition (ICDAR) 2011 line-level competition. There are 100 different characters in this database.

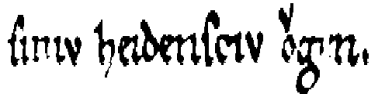
### 3) THE WASHINGTON DATABASE

The Washington database contains 565 text lines of the George Washington letters, handwritten by two writers in the 18th century. Although the language is also English, the text is written in longhand script and the images are binarized as illustrated in Fig. 3, see [3] for a description of the differences between binarized and binarization-free images when applying HTR tasks. In this database, four possible partitions are provided to train and validate. In this work, we have randomly chosen one of them. The train, validation, and test set contain 325, 168 and 163 handwritten lines, respectively. There are 83 different characters in the database.

### 4) THE PARZIVAL DATABASE

The Parzival database contains 4477 text lines handwritten by three writers in the 13th century. In this case, the lines are binarized like in the Washington database, but the text is written in gothic script. A sample is included in Fig. 4. There are 96 different characters in this database. Note that the Parzival database has a large number of text lines in comparison to the Washington one. We have randomly chosen a training set of the approximately same size as in the Washington training to emulate learning with a small dataset, the main goal of this work.

<sup>2</sup> The names of the images of each set are provided in the *Large Writer Independent Text Line Recognition Task*.



GT: *finiv heidenfciv ögen.*

**FIGURE 4.** Parzival handwritten text sample: image of a line and its transcript.

#### 5) THE ICFHR 2018 COMPETITION OVER READ DATASET

The set of documents of the ICFHR2018 Competition on Automated Text Recognition on a READ Dataset (<https://readcoop.eu/>) was proposed to compare the performance of approaches learning with few labeled pages. The dataset provided for the competition consists of 22 documents segmented at line level [5]. They are written in Italian and modern and medieval German. Each of them was written by only one writer but in different periods and various languages. The training data is divided into a general set (of 17 documents) and a document-specific set (of 5 documents) called Konzilsprotokolle\_C, Schiller, Ricordi, Patzig, and Schwerin of an equal script as in the test set. Hereafter, *general* is used to denote available source labeled databases different from the one of interest, while *document-specific* denotes particular target documents. Also, the Konzilsprotokolle\_C dataset, of the University of Greifswald, will be abbreviated as Konzil. The general database comprises roughly 25 pages per document (the precise number of pages varies such that the number of contained characters is almost equal per document). It will be denoted hereafter by ICFHR18-G. For the 5 document-specific databases the authors provide 16 labeled pages plus 15 unlabeled pages. One can check for the error in the transcription of these databases by sending the authors the transcription of these 15 pages. The results of the transcription are then published on the web of the contest. In Fig. 5, samples from five specific target documents are displayed. The standard Unicode Normalization Form Compatibility Decomposition - NFKD is applied to the GT to provide a common character set over such different documents, with 102 characters. The goal of the competition is to fit a model to transcript each of the 5 specific target documents with the lowest CER possible, using the 17 source documents available for training. For each document-specific target dataset, four experiments are conducted, simulating that you have 0, 1, 4, or 16 annotated pages available for training.

In Table 1 we include the number of training and test lines available for every database. While all lines in the test sets will be used, the number of lines of the training set used vary through the experiments and will be indicated in every case.

## IV. ON THE DATA AUGMENTATION AND TRANSFER LEARNING TRADEOFF

As our first contribution, in this section, we analyze the joint performance of TL and DA methods when applied to HTR.

### A. TRANSFER LEARNING

To cope with a reduced set of labeled inputs, we could first train the DNN model using as source available labeled large

**TABLE 1.** Number of lines available for training and test in each dataset.

|            | Full train set size | Test set size |
|------------|---------------------|---------------|
| RIMES      | 10163               | 778           |
| IAM        | 6152                | 2912          |
| Washington | 325                 | 163           |
| Parzival   | 350                 | 1328          |
| ICFHR18-G  | 11424               | 2878          |
| Konzil     | 351                 | 118           |
| Schiller   | 238                 | 90            |
| Ricordi    | 273                 | 110           |
| Patzig     | 473                 | 168           |
| Schwerin   | 782                 | 275           |

datasets. Then, we could apply TL or domain adaptation strategies [43] to tune the learned model to transcript a target document. As discussed in Section I, we usually deal with different tasks, where TL has shown useful to share the results of the learning between tasks.

Formally, in HTR, deep learning algorithms have been usually used to solve problems over a domain  $\mathcal{D} = \{\mathcal{X}, P(\mathbf{x})\}$ , where  $P(\mathbf{x})$  is the marginal probability. Typically  $\mathbf{x}$  is the image for a segmented line in the text. The task consists of two components: a label space  $\mathcal{Y}$  and an objective predictive function  $f(\cdot)$  (denoted by  $\mathcal{T} = \{\mathcal{Y}, f(\cdot)\}$ ), which can be learned from the training data. The data consists of pairs  $\{\mathbf{x}_i, y_i\}$ , where  $\mathbf{x}_i \in \mathcal{X}$  and  $y_i \in \mathcal{Y}$  [44] and  $f(\mathbf{x}) = Q(y|\mathbf{x})$  can be interpreted as the conditional probability distribution.

Given a source domain  $\mathcal{D}_S$  and a learning task  $\mathcal{T}_S$ , transfer learning aims to help improve the learning of another target predictive function  $f_T(\cdot)$  in  $\mathcal{D}_T$  using the knowledge in  $\mathcal{D}_S$  and  $\mathcal{T}_S$ . In this work, we are interested in *inductive transfer learning* in which the target task is different from the source task, as the domains are different ( $\mathcal{D}_S \neq \mathcal{D}_T$ ). Here we perform TL by retraining a DNN model where 1) all weights are initialized to the ones of the DNN learned for  $\mathcal{D}_S$  and  $\mathcal{T}_S$  and 2) the parameters of lower layers can be frozen to the values of the ones obtained after training with other available source datasets, used as off-the-shelf feature extractors [45].

In [6] we analyzed preliminary TL results over the Washington and Parzival databases, by using the IAM database as the source, and we investigated which layers should be kept fixed to then apply a fine-tuning process to the others. We concluded that the best choice is to unfreeze all the layers, where the first one can be eventually fixed. In most cases, fixing only the first CNN layer leads to the best performance.

In Table 2 we extend the analysis in [6] to the five specific documents in the ICFHR 2018 Competition dataset, where the 17 documents of the general set of the database, in the ICFHR18-G, are used as the source. Results are included when fixing layers 1 to 3 of the CNN, as fixing other layers provided larger errors in all cases. The lowest achieved errors are highlighted in boldface. Training set size is given in the number of lines. It can be observed that, among all databases, the best performance is achieved when unfreezing all layers or, at most, only the first layer is kept frozen. Hereafter the TL is applied by freezing just the first layer of the DNN model. The results shown in all tables hereafter indicate mean values



FIGURE 5. From top to bottom: Konzil, Schiller, Ricordi, Patzig and Schwerin handwritten text samples with their transcripts.

TABLE 2. TL performance: Mean CER (%) and bootstrapped confidence interval at 95%, in brackets, of the model in Fig. 1 using TL for the Washington, Parzival, Konzil, Schiller, Ricordi, Patzig and Schwerin datasets (see Section III) as target domains.

|            | All free                | Fixed layers               |                     |                     |
|------------|-------------------------|----------------------------|---------------------|---------------------|
|            |                         | CNN 1                      | CNN 1,2             | CNN 1,2,3           |
| Washington | <b>5.32</b> [5.22-5.41] | 5.41 [5.24-5.43]           | 5.53 [5.41-5.64]    | 6.30 [5.73-7.1]     |
| Parzival   | <b>3.30</b> [3.24-3.36] | <b>3.30</b> [3.21-3.34]    | 3.52 [3.36-3.62]    | 3.63 [3.47-3.69]    |
| Konzil     | 4.51 [4.33-4.61]        | <b>4.37</b> [4.24-4.54]    | 4.42 [4.36-5.49]    | 4.53 [4.43-4.61]    |
| Schiller   | <b>9.40</b> [9.31-9.46] | 9.42 [9.34-9.46]           | 9.48 [9.40-9.54]    | 10.11 [9.21-10.32]  |
| Ricordi    | 11.21 [11.14-11.25]     | <b>11.20</b> [10.11-11.23] | 11.28 [11.19-11.34] | 11.60 [11.41-11.72] |
| Patzig     | 10.63 [10.51-10.70]     | <b>10.60</b> [10.52-10.65] | 10.68 [10.57-10.74] | 12.4 [12.33-12.46]  |
| Schwerin   | <b>3.50</b> [3.46-3.53] | <b>3.50</b> [3.47 - 3.51]  | 3.91 [3.81 - 3.94]  | 4.22 [4.15-4.26]    |

of CER or WER. To get the statistics, the model in Fig. 1 is trained 10 times, where the parameters to initialize are independently and randomly set. In Table 2, a non-parametric bootstrapped confidence interval at 95% [46] is also included. For the remaining tables, the confidence intervals can be found in the supplementary material.

In Table 2 we analyze different strategies of applying TL, with no DA, for the Washington and Parzival target datasets with the IAM database as the source and Konzil, Schiller, Ricordi, Patzig, and Schwerin datasets (see Section III) as target domains with ICFHR18-G as the source. In each of ICFHR 2018 document-specific datasets, 12 pages are used for training, see the corresponding number of lines in Table 1. We conclude that the good choice is to freeze the first convolutional layer of the model (column “CNN1”). This solution will be used later in Subsection IV-C.

B. DATA AUGMENTATION

In [32] the authors compare various DA approaches using both RIMES [42] and IAM [41] databases as benchmarks,

where transcription is made on the word level. Note that these databases have a considerably large number of labeled lines.

When not applying any augmentation technique, they get a CER of 5.35 % (IAM) and 3.69 % (RIMES). The best CER values reported in [32] by using various DA techniques are 3.93 % and 1.36 %, respectively. Which is equivalent to an improvement of approximately 2 percentage points in both databases.

Let us now extend the same analysis to scenarios with small training datasets: Washington, Parzival, Konzil, Schiller, Ricordi, Patzig, and Schwerin databases. As throughout the paper, the transcriptions are made on line level. Results for the IAM, RIMES, and the ICFHR18-G, i.e., the 17 documents of the general dataset in the ICFHR 2018 database, are also analyzed as references. In Table 3 we include the CER of our DNN model with no DA and two different DA techniques, affine transformation [33] and random warp grid distortion (RWGD) [32], for all databases in Subsection III-B.

We augment the training set by generating ten copies of every line in the training set. One of these copies is the original line without distortions.

In Table 3, for the largest databases, the DA improvement is around 2 percentage points (2 percentage points in IAM, 1.9 percentage points in RIMES, and 2.5 percentage points in ICFHR18-G). However, in the small databases, the CER reduction is remarkable, in the range of 5 percentage points to 23.6 percentage points, see CERs highlighted in boldface. Note that the results in [32] are different to the ones in Table 3 because while in [32] transcription is done at the word level here whole lines are processed. This explains that in IAM without DA we get CER 7.2% while in [32] 5.35%

**TABLE 3. DA performance: Mean CER and WER (%) with affine transformations [33] and RWGD [32] DA approaches evaluated for all datasets in Subsection III-B. The DNN is trained from scratch using the number of lines indicated by ‘Train size’. Largest DA CER reductions are highlighted in boldface.**

|            | Train size | No method |      | Affine Transf. |      | RWGD [32]   |      |
|------------|------------|-----------|------|----------------|------|-------------|------|
|            |            | CER       | WER  | CER            | WER  | CER         | WER  |
| RIMES      | 10163      | 4.4       | 10.8 | 2.7            | 10.7 | <b>2.5</b>  | 10.4 |
| IAM        | 6152       | 7.2       | 22.2 | 5.9            | 20.3 | <b>5.3</b>  | 19.7 |
| Washington | 325        | 41.1      | 85.3 | 18.7           | 69.2 | <b>17.5</b> | 65.2 |
| Parzival   | 350        | 18.2      | 63.0 | 14.1           | 56.4 | <b>12.9</b> | 53.6 |
| ICFHR18-G  | 11424      | 12.2      | 43.7 | 10.6           | 40.1 | <b>9.7</b>  | 38.6 |
| Konzil     | 351        | 37.1      | 95.4 | 26.2           | 93.4 | <b>21.5</b> | 90.1 |
| Schiller   | 238        | 45.4      | 88.2 | 32.4           | 87.6 | <b>30.1</b> | 85.5 |
| Ricordi    | 273        | 37.2      | 93.1 | 36.2           | 91.1 | <b>35.2</b> | 90.3 |
| Patzig     | 473        | 24.5      | 86.3 | 18.5           | 81.3 | <b>17.1</b> | 80.5 |
| Schwerin   | 782        | 21.1      | 76.6 | 17.4           | 72.9 | <b>16.5</b> | 71.2 |

is reported. In any case, it can be concluded that, since the DA acts as a regularization technique to avoid overfitting, the CER reduction is greater as the size of the training set is reduced. At this point, it is most interesting to compare the results of TL and DA techniques when applied independently. It can be observed that TL exhibits, by far, a much larger CER reduction. Next, we face the design and analysis of both techniques combined, where RWGD will be used as the DA approach.

**C. COMBINING DATA AUGMENTATION AND TRANSFER LEARNING**

When comparing DA with TL, the large databases are excluded from the comparison. They play the role of source databases in the TL approach, specifically, the IAM is the source dataset when Washington and Parzival are targets and ICFHR18-G in the Konzil, Schiller, Ricordi, Patzig, and Schwerin case. The RIMES database is only used to enhance the comparisons in this section.

In the combination of TL and DA techniques, there are several possible designs. Here we propose the following two schemes. In a first approach we perform DA at both the learning from the source dataset and the retraining of the model with the target one:

- 1) Train the model from scratch with a source dataset, applying DA.
- 2) Retrain the model with the target dataset, applying DA.

We name this proposal DA-TL-DA. In a second proposal, denoted by DA-TL, no DA is applied to the target:

- 1) Train the model from scratch with a source dataset, applying DA.
- 2) Retrain the model with the target dataset, **without** applying DA.

We perform the same experiments as in Subsection IV-A, obtaining the results included in Table 4. For the sake of completeness we also report the results in the two previous subsections for TL with the first layer frozen and DA with RWGD. In the first step of the DA-TL and DA-TL-DA methods, the model has been trained from scratch with the IAM database. After that, we fine-tune the model using data from the Parzival and Washington databases. In Table 5 the results are shown when training the model in Fig. 1 from scratch with

**TABLE 4. TL and DA combined performance: Mean CER (%) evaluated for Washington and Parzival datasets using TL and DA with IAM database as the source. The number of annotated lines used in training is included as ‘Train size’.**

|          | Train size #lines | No method            | TL                | DA                   | DA-TL-DA          | DA-TL                                  |
|----------|-------------------|----------------------|-------------------|----------------------|-------------------|--|
|          |                   | Washington           | 150<br>250<br>325 | 51.6<br>46.4<br>41.1 | 9.4<br>7.1<br>5.4 | 22.8<br>20.4<br>17.5                   |
| Parzival | 150<br>250<br>350 | 21.9<br>20.7<br>18.2 | 5.8<br>4.0<br>3.3 | 15.7<br>14.2<br>12.9 | 6.0<br>4.2<br>3.4 | <b>5.6</b><br><b>3.8</b><br><b>3.3</b> |

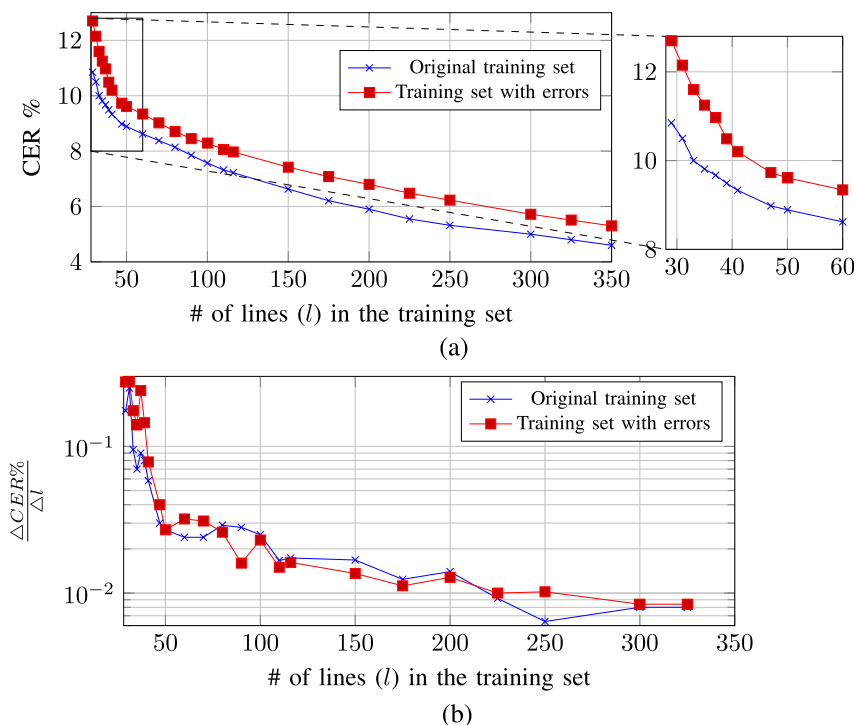
**TABLE 5. TL and DA combined performance: Mean CER (%) evaluated in ICFHR 2018 Competition Specific datasets as targets using TL and DA with ICFHR18-G as source. The number of annotated pages used in the training is included as ‘Train size’.**

|          | Training set size. # pages                           | None                         | TL   | DA                           | DA-TL-DA                       | DA-TL  |
|----------|--|------------------------------|--|------------------------------|--------------------------------|--|
|          |  | Konzil                       | 0<br>1 (29 lines)<br>4 (116 lines)<br>12 (351 lines) | -<br>48.10<br>45.30<br>37.10 | 15.50<br>10.85<br>6.54<br>4.37 | -<br>37.30<br>28.71<br>21.50                                 |
| Schiller | 0<br>1 (21 lines)<br>4 (84 lines)<br>12 (238 lines)  | -<br>53.70<br>48.40<br>45.40 | 24.60<br>17.36<br>12.25<br>9.42                      | -<br>39.50<br>33.20<br>30.10 | -<br>21.40<br>14.00<br>10.00   | <b>24.60</b><br><b>17.31</b><br><b>12.22</b><br><b>9.38</b>  |
| Ricordi  | 0<br>1 (19 lines)<br>4 (88 lines)<br>12 (273 lines)  | -<br>56.20<br>43.52<br>37.21 | 39.19<br>23.66<br>21.17<br>11.20                     | -<br>51.0<br>40.81<br>35.20  | -<br>24.10<br>21.10<br>10.91   | <b>39.19</b><br><b>22.71</b><br><b>21.02</b><br><b>11.14</b> |
| Patzig   | 0<br>1 (38 lines)<br>4 (156 lines)<br>12 (473 lines) | -<br>42.54<br>37.63<br>24.50 | 41.50<br>27.91<br>16.40<br>10.60                     | -<br>35.31<br>30.50<br>17.13 | -<br>31.42<br>18.32<br>11.21   | <b>41.50</b><br><b>26.7</b><br><b>16.11</b><br><b>10.00</b>  |
| Schwerin | 0<br>1 (68 lines)<br>4 (264 lines)<br>12 (782 lines) | -<br>38.40<br>29.31<br>21.10 | 34.53<br>12.15<br>5.73<br>3.51                       | -<br>30.20<br>24.30<br>16.50 | -<br>10.62<br>5.30<br>3.30     | <b>34.53</b><br><b>10.80</b><br><b>5.52</b><br><b>3.41</b>   |

the ICFHR18-G, and being fine-tuned on the 5 specific target data sets provided. For the sake of completeness we include in Table 5 the results for 0 pages in the target dataset, i.e., when no labeled sampled from the target is used. Note that in this case DA-TL-DA cannot be applied.

In the light of Table 4 and Table 5, it can be concluded that applying DA over the target training set once TL is applied, i.e., DA-TL-DA, either does not reduce the CER or even it slightly increases it, compared to the result of the TL approach alone or the DA-TL method. Except for Schwerin, in which DA-TL-DA slightly improves DA-TL. Put in other words, in general, it is harmful to apply DA to the target dataset if TL has been applied, when just a reduced number of labeled lines are available in the target. On the other hand, DA+TL achieves improvements up to 5 % in the ICFHR 2018 target documents, usually increasing with the reduction of the training set.

From the discussion above, and bearing Table 4 and Table 5 in mind, it can be concluded that DA-TL is a robust approach. When fine-tuning a DNN that has been previously trained with a similar task (a huge database of HTR samples), the starting point is reasonably good as we can observe in Table 5 for the training set sizes of 0 pages in all datasets. We show a good generalization ability of the model for the TL and DA+TL without further training with the target dataset. Afterwards, the DNN model is trained with the target database. Only a few samples are available in the target set, which represent just a limited part of the support



**FIGURE 6.** (a) CER (%) divided by the number of annotated lines,  $l$ , used and (b) decrement of CER (%) divided by the number of new labeled lines added to obtain it,  $\Delta l$ , in the training of the DNN model with DA-TL approach using the ICFHR18-G dataset as the source and the Konzil dataset as the target with no artificial errors (x) corrupted with artificial errors (■).

of its marginal distribution,  $P_T(\mathbf{x})$ . After TL, the parameters of the DNN encode information from both the source and the target training sets. At this point, we conjecture that by using DA in the target dataset and further refining the parameters, the DNN model overfits to the augmented versions of the target samples, forgetting the knowledge learned from the source one, that very much helps to transcript inputs out of the support generated by augmenting the target set.

This leads to an increase of the final CER.

**V. THE CORRUPTED LABEL PURGING (CLP) ALGORITHM**

In this section, we focus on the impact of the number of lines and their quality in the target dataset on the learning process of the DNN model. We first analyze the impact of the performance on the number of *healthy* lines, i.e., lines with no transcription errors in the training dataset. Then we study how this performance degrades with label errors. Finally, we propose an algorithm to detect and remove potential label errors in the dataset.

**A. PERFORMANCE VARIATION WITH THE NUMBER OF LABELED SAMPLES**

When a small number of lines is available in the target training set, deep learning models are quite sensitive to small variations in the number of labeled lines. In this subsection, this sensitivity is evaluated on a specific dataset from the ICFHR 2018 Competition [5]. The chosen training dataset consists of 16 pages from the Konzil, which is segmented


at line level. Similar results were obtained for the other datasets.

The ICFHR 2018 target datasets have 16 labeled pages each. Unless otherwise indicated, 4 of them will be used for testing purposes while up to 12 pages will be used for training. Usually, 10 % out of the used training set is devoted to validation. The ICFHR18-G dataset is used as source database in the TL-DA approach.


In Fig. 6.(a) the blue curve in x represents the TL-DA CER versus the available number of lines,  $l$ , of the target training set in the range 29-350 lines, corresponding to 1 and 12 pages, respectively. In the left part of the figure, the CER decreases at a rate of 1 percentage point every 4 new lines added to the training set. After approximately 50 lines, the decreasing rate of the CER changes to approximately 1 percentage point every 100 lines. This is evidenced in Fig. 6.(b) where the absolute value of the variation of the CER in percentage points is depicted  $\Delta CER$ , with the increment of the number of annotated lines used in the target to achieve it,  $\Delta l$ . It can be concluded that the sensitivity to the number of samples in the training set is significantly larger for small training sets.

In Fig. 6 we also include the “Training set with errors” curve (■) which corresponds to the analysis above but where labeling errors have been artificially introduced, as follows. The annotation of a line is modified with probability  $L$ . Then, with modified labeling, a character is changed with probability  $R$ . In both cases following a Bernoulli distribution. Every changed character is replaced by an independently and





**FIGURE 7.** Sample of a completely mislabeled text at Ricordi dataset.



**FIGURE 8.** Sample of special annotations in the GT at the Ricordi dataset.

randomly selected character, following a discrete uniform probability. In Fig. 6, where  $L = 0.2$  and  $R = 0.3$ , it is interesting to note that the impact of labeling errors in the CER value is more dramatic for small training sets while the rate at which the CER decreases with the number of lines added remains roughly unaltered.

## B. TYPES OF TRANSCRIPTION ERRORS

Before proposing approaches to detect mislabels in the training set, we discuss three typical types and causes of errors in the datasets, as follows.

- 1) *Mislabeled characters.* When labeling a training set, the most common mistake is to confuse a character with another, usually a similar one. This can be seen in the well-known IAM database [41], where it is indicated that some lines could have some annotation errors in the labels. This type of error is the one simulated in Fig. 6.
- 2) *Label Misalignment.* The second kind of detected error happens due to a misalignment in the labels. This could be caused by, e.g., a mistake in the name given to some images in the database. This error is encountered several times in the Ricordi dataset from the ICFHR 2018 Competition [5] as illustrated in Fig. 7. It can be observed in this example that the transcript does not correspond to the handwritten text in the image above. On the contrary, it is quite close to the model output, after being trained with several lines of the dataset.
- 3) *Special annotations in the ground truth.* Perhaps, the most common source of error is due to special annotations that some transcribers or database managers introduce in some datasets to include some notes inline. In [37] they found this problem in the IAM database: crossed out words that are labeled with the symbol “#” followed by the word behind the blot. Training the model with this labeling might lead to unpredictable behavior, since the model could replace the text using “#” at different parts of the text. The model will either be able to recognize the text behind the blot or replace the word with the symbol “#”, or both. Another special annotation is included in Fig. 8, where they write in brackets extra characters that are not in the handwritten

text. The output of a model trained with samples of the same dataset is shown below the GT. Despite in this line, the CER is about 35%, it can be observed that the model output is quite similar to the handwritten text.

Manually annotating historical documents remains a challenging task that is prone to errors, even for experts in the field. As discussed in the previous section, when a huge set of annotated samples is available, deep learning models do not suffer from a few mislabeled samples, as they better generalize. However, when a limited set of annotated lines of a specific writer is available to train, mislabeled lines induce overfitting to transcripts with errors, which are quite hard to tackle via regularization. In the example shown in Fig. 6 we illustrate this problem when just a few mislabeled lines are introduced.

---

### Algorithm 1 Corrupted Labels Purging (CLP)

---

**Given inputs:** source set  $x_S \in \mathcal{X}_S$  and  $y_S \in \mathcal{Y}_S$ , target set  $x_T \in \mathcal{X}_T$  and  $y_T \in \mathcal{Y}_T$  and threshold,  $\epsilon$ .

- 1) Fit the prediction function  $f_S(y_S|x_S)$  with the source training set  $\{x_S, y_S\}$ .
- 2) Split the target training set into  $N$  subsets  $\{x_{T_1}, x_{T_2}, \dots, x_{T_N}\}, \{y_{T_1}, y_{T_2}, \dots, y_{T_N}\}$ .

**for**  $n = 1, \dots, N$  **do**

- 3) Initialize the prediction function  $f_{T_n}(\cdot) = f_S(\cdot)$ .
- 4) Fine tune the prediction function with the whole target set except for the  $n$ th,  $\{x_{T_i \neq T_n}\}, \{y_{T_i \neq T_n}\}$ .
- 5) Include in the new target set,  $\{x_{T'}, y_{T'}\}$ , all pairs  $\{x_{T_n}^{(i)}, y_{T_n}^{(i)}\}$  whose predictions  $f_{T_n}(y_{T_n}^{(i)}|x_{T_n}^{(i)})$  have errors below a CER threshold,  $\epsilon$ .

**end for**

- 6) Initialize the prediction function  $f_T(\cdot) = f_S(\cdot)$ .
- 7) Fine tune the prediction function,  $f_{T'}(y_{T'}|x_{T'})$ , to the modified target set  $\{x_{T'}, y_{T'}\}$ .

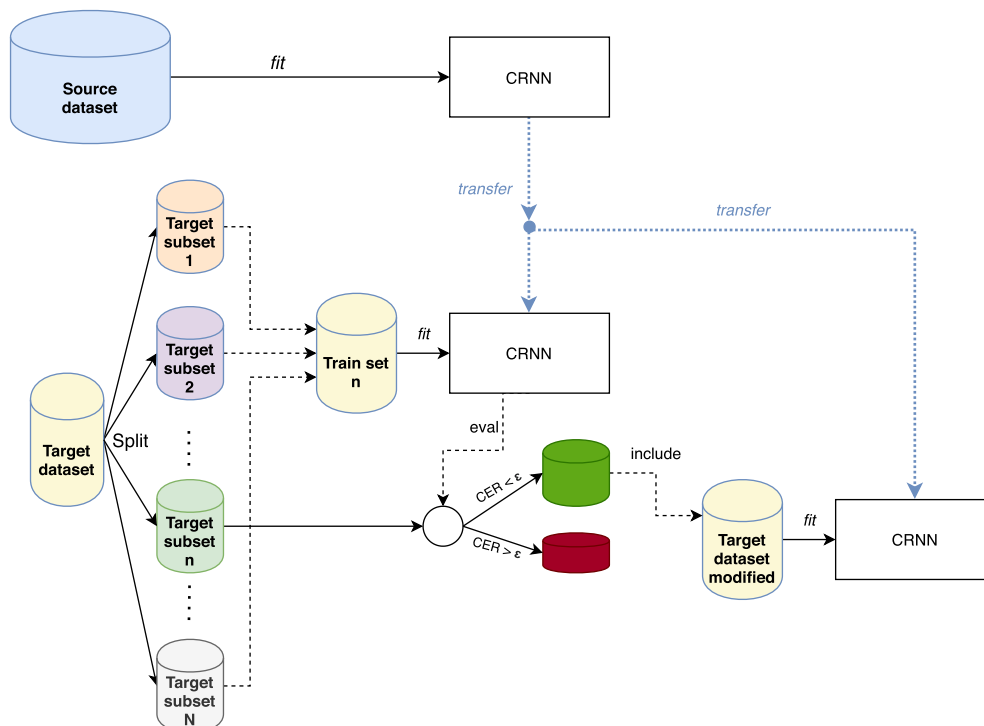
**Output:**

Function  $f_T(y_T|x_T)$  over the target domain  $\mathcal{D}_T$ .

---

## C. MISLABEL DETECTION ALGORITHM

As one of our main contributions, we propose an algorithm to detect and remove mislabeled lines from the training set,



**FIGURE 9.** Corrupted labels purging algorithm. The algorithm applied over target subset  $n$  is depicted. The same procedure should be applied to all the subsets to build the *Target dataset modified*.

detailed in Algorithm 1. A block diagram of the algorithm is also depicted in Fig. 9. It divides the target training dataset into  $N$  subsets. For every subset,  $n$ , the method performs DA-TL using the rest of subsets,  $k = 1, \dots, N, k \neq n$ , as training sets and it evaluates the CER metric over the subset  $n$ . Lines with CER above a threshold,  $\epsilon$ , in the  $n$ th subset are detected as wrongly transcribed and discarded. Hence, we are implementing some sort of  $k$ -fold validation, in which the size of each validation partition is reduced after removing problematic lines. Finally, the DA-TL is applied to the resulting target database. Note the algorithm performs  $N + 1$  different training steps. However, the  $N$  steps concerning the target subsets could be run in parallel, since they are independent of each other. Hence, the run time of applying this algorithm is approximately double the run time of regular training.

In Fig. 10 we include the histogram of the CER per line for the 5 ICFHR 2018 document-specific datasets using the CLP algorithm with  $N = 2$ . The ICFHR18-G was used as the source. The histograms were estimated with the CER of the outputs of the  $n = 1, 2$  stages computed with the lines not used during training, see the output of “Target subset  $n$ ” blocks in Fig. 9. In the left column, models have been trained with 4 pages while in the right column they have been trained with 12 pages. Lines are corrupted with artificial errors with probability  $L = 0.1$ , while every character in the label of a line is changed with probability  $R = 0.3$  to a random value.

Conservatively, we believe that a 10% average number of corrupted lines represent a label error rate similar to the one we encounter in real databases. It is interesting to observe that

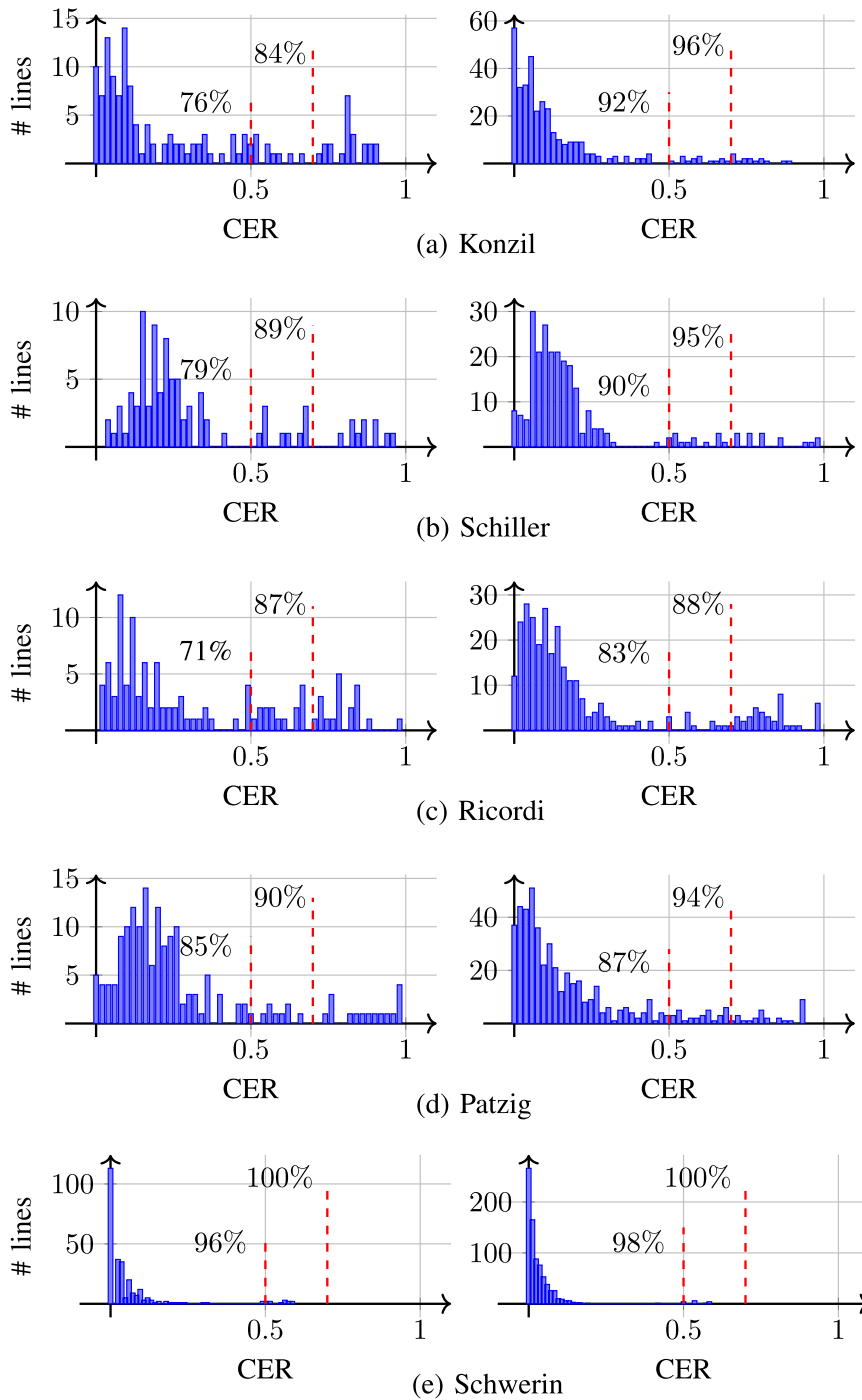
the results for the Schwerin dataset are remarkably better than for the others because it has a significantly larger number of lines per page. Besides, in the Ricordi dataset, the histogram for 12 pages exhibits large values around 0.8. This dataset is known to have label misalignments.

#### D. CLP THRESHOLD ANALYSIS

The selection of the threshold is central to the algorithm performance. In Fig. 10 the CER of the healthy lines is mainly distributed around a mode value, to the left of each histogram, while outliers exhibit larger values. As representative values to be studied, after extensive simulations, we restrict our analysis to the thresholds  $\epsilon = 0.5$  and  $\epsilon = 0.7$ , for an average rate  $L = 0.1$  of artificially modified lines, and  $R = 0.3$ . In Fig. 10 we indicate the percentage of lines with CER equal or lower than 0.5 and 0.7, left and right red dashed lines in the subfigures, respectively. We conclude that almost 10% of lines have a CER above  $\epsilon = 0.7$  when 4 pages for training are available and the same occurs in the case of 12 pages when  $\epsilon = 0.5$ .

The selection for  $\epsilon$  should not lead to the deletion of healthy lines, otherwise, the overall CER would raise. On the other hand, the threshold must ensure a sensitivity when corrupted lines are encountered.

In the following, we study the CLP algorithm in two different scenarios. The first experiment we perform consists in applying the CLP algorithm to the ICFHR 2018 target datasets, with 4 and 12 pages as target training set size. Then we evaluate the CLP for the Washington and



**FIGURE 10.** Histogram of CER with DA-TL and ICFHR18-G as source dataset for the 5 document-specific datasets using 4 pages (left) and 12 pages (right) of the target dataset. Lines and characters were corrupted with probabilities  $L = 0.1$  and  $R = 0.3$  respectively. The histograms were evaluated with the outputs of the  $N = 2$  target subsets. Red dashed lines indicate the percentage of lines with  $CER \leq \epsilon$  with  $\epsilon = 50\%$  and  $\epsilon = 70\%$ , left and right lines, respectively.

Parzival databases, with 150 and 325 lines as target training set size. The same procedure is followed through all the scenarios:

- 1) Fit the model to the source set.
- 2) Run DA-TL plus CLP with  $N = 2$ .

### 1) ICFHR 2018 COMPETITION RESULTS

We test the CLP algorithm over real databases where we do not have any prior knowledge about the pattern of labeling errors. We do also include artificial errors to evaluate the CLP robustness.

**TABLE 6.** Mean CER (%) evaluated in Konzil, Schiller, Ricordi, Patzig and Schwerin target documents in the ICFHR2018 Competition datasets for DA-TL, DA-TL+CLP with  $\epsilon = 50\%$  and  $\epsilon = 70\%$ . DA-TL was applied with both a training set of 4 pages and 12 pages. The annotation for a line is corrupted with probability  $L = 0.1$  and a character within it randomly replaced with probability  $R$ .  $R = 0$  indicates no error introduced in the labelings. The number of removed lines by the CLP algorithm is included in parentheses in the last two columns. The best-achieved value in every row is in boldface.

| Dataset  | Train set size          | $R$ | Baseline     | $\epsilon = 50\%$  | $\epsilon = 70\%$  |
|----------|-------------------------|-----|--------------|--------------------|--------------------|
| Konzil   | 4 Pages<br>(116 lines)  | 0%  | <b>7.6</b>   | 8.5(-31)           | 7.9(-7)            |
|          |                         | 30% | 8.7          | 8.3 (-41)          | <b>7.82</b> (-14)  |
|          |                         | 50% | 9.1          | 8.2 (-39)          | <b>7.9</b> (-16)   |
|          | 12 Pages<br>(351 lines) | 0%  | <b>4.6</b>   | 5.3 (-1)           | <b>4.6</b> (-0)    |
|          |                         | 30% | 5.3          | <b>4.6</b> (-29)   | 5.0 (-25)          |
|          |                         | 50% | 5.5          | <b>4.8</b> (-35)   | 5.0(-28)           |
| Schiller | 4 Pages<br>(84 lines)   | 0%  | <b>13.27</b> | 14.72 (-12)        | 13.61 (-5)         |
|          |                         | 30% | 15.19        | 14.81 (-17)        | <b>14.43</b> (-10) |
|          |                         | 50% | 15.64        | 14.96 (-22)        | <b>13.87</b> (-12) |
|          | 12 Pages<br>(244 lines) | 0%  | <b>9.42</b>  | 9.76 (-2)          | 9.42 (-0)          |
|          |                         | 30% | 11.31        | <b>10.41</b> (-22) | 10.62 (-22)        |
|          |                         | 50% | 12.75        | 10.61 (-24)        | <b>10.51</b> (-25) |
| Ricordi  | 4 Pages<br>(88 lines)   | 0%  | 21.1         | 18.2 (-16)         | <b>18.2</b> (-16)  |
|          |                         | 30% | 23.2         | 20.8 (-32)         | <b>20.5</b> (-27)  |
|          |                         | 50% | 24.31        | 21.94 (-44)        | <b>20.81</b> (-27) |
|          | 12 Pages<br>(295 lines) | 0%  | 9.7          | <b>9.4</b> (-38)   | 9.4 (-38)          |
|          |                         | 30% | 10.47        | <b>9.23</b> (-41)  | 9.49 (-38)         |
|          |                         | 50% | 10.8         | <b>9.53</b> (-52)  | 9.75 (-44)         |
| Patzig   | 4 Pages<br>(156 lines)  | 0%  | <b>18.32</b> | 18.93 (-7)         | <b>18.32</b> (-0)  |
|          |                         | 30% | 21.41        | 21.6 (-27)         | <b>21.1</b> (-18)  |
|          |                         | 50% | 21.84        | 22.12 (-27)        | <b>21.31</b> (-18) |
|          | 12 Pages<br>(473 lines) | 0%  | <b>11.5</b>  | 11.96 (-15)        | 11.54 (-4)         |
|          |                         | 30% | 12.28        | 12.23 (-61)        | <b>11.98</b> (-52) |
|          |                         | 50% | 12.8         | 12.67 (-63)        | <b>12.35</b> (-54) |
| Schwerin | 4 Pages<br>(264 lines)  | 0%  | 5.3          | 5.3 (-0)           | <b>5.3</b> (-0)    |
|          |                         | 30% | 5.36         | <b>5.31</b> (-14)  | 5.36 (-0)          |
|          |                         | 50% | 5.39         | <b>5.32</b> (-26)  | 5.33 (-12)         |
|          | 12 Pages<br>(782 lines) | 0%  | 3.3          | 3.3 (0)            | <b>3.3</b> (0)     |
|          |                         | 30% | 3.36         | <b>3.31</b> (-14)  | 3.36 (-0)          |
|          |                         | 50% | 3.53         | <b>3.34</b> (-75)  | 3.39 (-22)         |

The results of these analyses are reported in Table 6 and Table 7. Their three last columns include the results for the DA-TL with no CLP as ‘Baseline’, for the DA-TL+CLP with  $\epsilon = 50\%$ , and then for the DA-TL+CLP with  $\epsilon = 70\%$ . For every target dataset and training set size three rows are used to report the CER (%) when no artificial errors are introduced,  $R = 0$ , for  $R = 30\%$  and  $R = 50\%$ .

In this first case, the ICFHR18-G dataset is used as the source. The 17 documents of this corpus have a total number of 11424 lines. The DA-TL plus CLP was applied to the five target documents in the competition: Konzil, Schiller, Ricordi, Patzig, and Schwerin. The results are included in Table 6, where it is included the average value for the CER and the number of removed lines by the CLP algorithm.

In view of the results, we highlight the following aspects. First note that, when errors are induced, the threshold  $\epsilon = 70\%$  performs better in most of the cases when the training set is of 4 pages while the threshold  $\epsilon = 50\%$  is the best choice for 12 pages. Exceptions can be observed in Patzig and Schwerin corpora. For the Patzig dataset, we conclude that  $\epsilon = 70\%$  is the best choice for every case. This is due to the distribution of the errors in this dataset, which has a larger variance, and therefore more lines are above the  $\epsilon = 50\%$  CER, it can be seen in Fig. 10. In the Schwerin corpus, the threshold 50% has the best CER in all cases, the opposite that in the Patzig dataset. This is due to the distribution of the errors in this dataset that, due to the larger number of lines

used, has lower variance and most lines are below 10% CER (see Fig. 10).

It is also interesting to remark that in the Ricordi case, the algorithm improves the CER in the original dataset, i.e., without synthetic errors. This is explained by the fact that in this dataset, as already discussed, there are some mislabeled lines like in the case illustrated in Fig. 7. Additionally, note that for  $R = 0$  and  $\epsilon = 70\%$  a large number of removed lines is quite an indicator of the dataset containing errors in the annotated lines.

For the sake of completeness, we include in Fig. 11 the evolution of the CER versus the number of lines,  $l$ , in Fig. 6 including the CER for the proposed algorithm (CLP) (o). The introduction of the CLP improves the TL-DA approach when the dataset has corrupted lines. In the range,  $l = [40, 50]$  the TL-DA with CLP with  $l = 40$  achieves the same CER as the DA-TL with  $l = 50$  lines in the training set.

## 2) WASHINGTON AND PARZIVAL RESULTS

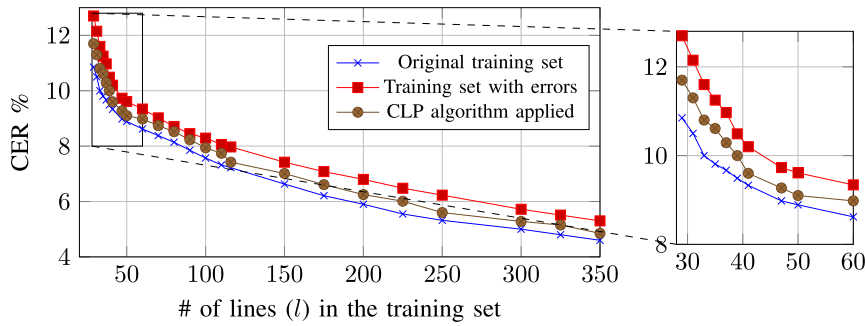
In this second analysis, the model is pre-trained with the IAM database as the source dataset to train the model with DA-TL for the Washington and Parzival targets. There are two main differences to the previous study of the ICFHR 2018 datasets: 1) the number and set of characters are different from the source and targets datasets and 2) we compare the CER of both targets in terms of the number of lines instead of the number of pages, where we consider two cases, 150 lines and 325 lines, similar to the number of lines used in the previous scenario.

The first rows in Table 7 include the results obtained after fine-tuning the model to the Washington dataset. In this study, the threshold  $\epsilon = 70\%$  is the best option when the number of lines is 150 while  $\epsilon = 50\%$  exhibits the lowest CER when the number of lines is 325. This is equivalent to the 4 and 12 pages in the Konzil, Schiller, and Ricordi cases in which the number of lines is similar. For these thresholds: we get an improvement of 0.8 and 0.63 in the case of 150 lines and no deterioration over the original dataset. In the case of 325 lines, we get a boost of 0.4 and 0.5 and no deterioration over the original dataset.

Results obtained after fine-tuning the model to the Parzival dataset are also included in Table 7, see the lower rows. Similar conclusions can be drawn except for  $R = 30\%$  and 150 lines, where the 50% exhibits the best CER. If we choose the threshold as in the previous cases, 70% and 50%, we still get a slight improvement or at least, no deterioration.

## E. CORRECTING LABEL MISALIGNMENT

In Section V-B we summarized the different types of transcription errors. One of these errors is due to the misalignment of the annotations with the images. When a high number of lines are classified as mislabels, this type of error can be addressed by searching within the outputs of the DNN model for the whole target dataset, the transcript best fitting every annotation in the GT, hence aligning annotations and



**FIGURE 11.** CER (%) divided by the number of annotated lines,  $l$ , with the DA-TL approach using the ICFHR18-G dataset as source and the Konzil dataset as target with no artificial errors (x), with artificial errors (■) and with artificial errors and CLP used (●).

**TABLE 7.** Mean CER (%) evaluated in Washington and Parzival documents for DA-TL, CLP with threshold  $\epsilon = 50\%$  and CLP with threshold  $\epsilon = 70\%$ . DA-TL was applied with the IAM dataset as the source and using 150 and 325 lines from the target. The annotation for a line is corrupted with probability  $L = 0.1$  and a character within it randomly replaced with probability  $R$ .  $R = 0$  indicates no error introduced in the labelings. The number of removed lines by the CLP algorithm is included in parentheses in the last two columns.

| Dataset    | Train set size | R    | Baseline | $\epsilon = 50\%$ | $\epsilon = 70\%$ |
|------------|----------------|------|----------|-------------------|-------------------|
| Washington | 150 lines      | 0 %  | 9.4      | 9.5 (-6)          | <b>9.4 (-2)</b>   |
|            |                | 30 % | 11.3     | 10.6 (-20)        | <b>10.5(-14)</b>  |
|            |                | 50 % | 11.5     | 11.1 (-31)        | <b>10.87(-19)</b> |
|            | 325 lines      | 0 %  | 5.3      | <b>5.3 (-2)</b>   | 5.3 (-0)          |
|            |                | 30 % | 6.1      | <b>5.7 (-26)</b>  | 6.1 (-0)          |
|            |                | 50 % | 6.3      | <b>5.8 (-34)</b>  | 6.3 (-0)          |
| Parzival   | 150 lines      | 0 %  | 5.8      | 5.8 (-0)          | <b>5.8 (-0)</b>   |
|            |                | 30 % | 6.4      | <b>6.0 (-15)</b>  | 6.2 (-2)          |
|            |                | 50 % | 6.6      | 6.2 (-20)         | <b>6.1(-14)</b>   |
|            | 325 lines      | 0 %  | 3.3      | <b>3.3 (-0)</b>   | 3.3 (-0)          |
|            |                | 30 % | 3.5      | <b>3.5 (-0)</b>   | 3.5 (-0)          |
|            |                | 50 % | 3.5      | <b>3.4 (-35)</b>  | 3.5 (-0)          |

images in the dataset. This approach is quite similar to the one proposed in [38].

In the case of the Ricordi dataset in the ICFHR 2018 competition, we realized that the CLP detected a high number of mislabeled lines in the dataset. Note the large numbers of removed lines in Table 6 for  $\epsilon = 70\%$  and this dataset with  $R = 0$ . By simply visual inspection we confirmed that the error was of the type of misalignment of images and annotations. Here, we apply the CLP plus the simple automatic alignment approach described above.

The comparison between simply removing the mislabeled lines and correcting the alignment of the database is shown in Table 8. In this table, one can observe a significant dropping in the CER when correcting these misalignments of the lines. In the training with 4 pages, the overall decrease is 3.7 percentage points. In the 12 pages analysis, the CER drops 0.3 percentage points when removing the lines while it further decreases 0.8 percentage points when correcting them. Note also that the gain is higher when a lower number of annotated lines are used.

**F. COMPARISON TO THE STATE-OF-THE-ART**

By using the proposed 5+5 DNN model with CNN and BLSTM layers followed by a CTC, we conclude by analyzing the results of the novel DA-TL approach over the ICFHR

**TABLE 8.** Comparison between the CLP algorithm with line removal and the CLP plus alignment of the GT before detection. The mean CER (%) is evaluated for the Ricordi document with a training set of size 4 pages (88 lines) and 12 pages (295 lines).

| Train set size       | Method          | Baseline | $\epsilon = 50\%$ | $\epsilon = 70\%$ |
|----------------------|-----------------|----------|-------------------|-------------------|
| 4 Pages (88 lines)   | CLP             | 21.1     | 18.2              | 18.2              |
|                      | CLP + alignment | 21.1     | 17.4              | 17.4              |
| 12 Pages (351 lines) | CLP             | 9.7      | 9.4               | 9.4               |
|                      | CLP + alignment | 9.7      | 8.9               | 8.9               |

**TABLE 9.** CER ICFHR 2018 Competition results for LSTM based models: upper part, other previous approaches and, in the lower part, the results for the approaches in this work. Lowest mean values in both parts are highlighted in boldface.

|               | CER (%) per training size |       |       |       | CER (%) per document |          |         |        |          | Mean         |
|---------------|---------------------------|-------|-------|-------|----------------------|----------|---------|--------|----------|--------------|
|               | 0                         | 1     | 4     | 16    | Konzil               | Schiller | Ricordi | Patzig | Schwerin |              |
| OSU [32]      | 31.40                     | 17.74 | 13.27 | 9.02  | 9.39                 | 21.10    | 23.27   | 23.17  | 12.98    | <b>17.86</b> |
| ParisTech [4] | 32.25                     | 19.80 | 16.98 | 14.72 | 10.49                | 19.05    | 35.60   | 23.83  | 17.02    | 20.94        |
| LITIS [47]    | 35.30                     | 22.51 | 16.89 | 11.34 | 9.14                 | 25.69    | 30.50   | 25.18  | 18.04    | 21.51        |
| PRHLT         | 32.79                     | 22.15 | 17.89 | 13.33 | 8.65                 | 18.39    | 35.07   | 26.26  | 18.65    | 21.54        |
| RPPDI [21]    | 30.80                     | 28.40 | 27.25 | 22.85 | 11.90                | 21.88    | 37.29   | 32.75  | 28.55    | 27.32        |
| TL            | 32.77                     | 19.51 | 15.12 | 8.26  | 9.16                 | 21.00    | 29.39   | 23.25  | 13.54    | 18.93        |
| DA-TL         | 31.55                     | 19.21 | 14.91 | 8.16  | 8.58                 | 21.68    | 27.84   | 22.35  | 12.50    | 17.83        |
| CLP           | 30.13                     | 19.10 | 12.40 | 7.93  | 8.59                 | 21.69    | 22.81   | 22.35  | 12.51    | <b>17.39</b> |

2018 Competition.<sup>3</sup> The results included in Table 9 were reported by the organizers of the competition. The contestants provide the transcript of the 15 test pages for every document in the target set: Konzil, Schiller, Ricordi, Patzig, and Schwerin. Then, the organizers evaluate the CER, publicly publishing the results. In this table our results are compared against the 5 original contestants in the competition: OSU [32], ParisTech [4], LITIS [47], PRHLT and RPPDI. These approaches use DNN models based on CNN, LSTM, and CTC, where some variant of the LSTM is used. Some of them use DA in the target and LM. The recent work published by Yousef *et al.* [9] using a DNN model based on a fully gate convolutional network (GCN), outperformed the LSTM based approaches, with a mean value of 13.02 % providing a 23.35 % CER for a 0-page training size.

The results of the proposal in this work are included in the lowest rows of Table 9 where, following the conclusions in Subsection V-D, we used  $\epsilon = 70\%$  for the 1 and 4 pages training and  $\epsilon = 50\%$  for the 16 pages. Also, the CLP includes an alignment stage. Results are presented in three groups of columns. First, the average CER (%) for the 5 target dataset is included when 0, 1, 4 and 16 pages of the target datasets are used. The second group of 5 columns reports the

<sup>3</sup> The results are publicly available in the ICFHR competition website: <https://scriptnet.iit.demokritos.gr/competitions/10/viewresults/>

average CER (%) for the learning with 0, 1, 4, and 16 pages in the dataset for every document. The mean value per row is included in the last column.

## VI. CONCLUSION

In this paper, we analyze, for small training sets and in the framework of historical HTR, two well-known techniques in almost every deep learning application: TL and DA. We show that TL improves the CER between 10-40 percentage points when applied to small training sets, of the order of 300 text lines. DA also drops the CER in a range of 2-20 percentage points when a network is trained from scratch. Before performing TL, applying DA in the source dataset does reduce the CER. However, applying DA to the target datasets jointly with TL exhibits worse results than using TL alone. Hence, we propose the DA-TL approach where the DA is applied to the source dataset in the TL process.

Besides, we highlight that the DNN models are very sensitive to the number of lines in the train set when this number is low. Therefore, errors in annotated lines of small target datasets have a greater impact than in large datasets, for the same proportions of mislabels. To avoid that, we propose a method that can detect the mislabeled lines and remove it from the training set. Furthermore, we fix errors of the misalignment type, by searching for the true labels in the datasets.

Comparing to the state-of-the art in the ICFHR 2018 Competition, it can be observed that the DA-TL and CLP outperform all approaches within the CNN+LSTM+CTC class hence underlining the importance of the issues discussed: DA is important but in the source dataset, TL is to be considered, and mislabeling detection and correction is important if the dataset exhibits errors. Besides, the CLP introduces a residual 0.01 percentage points of loss if the datasets have no errors in the labels while the reduction is important if they have, see the results for the Ricordi corpus where a reduction of 6.58 percentage points is achieved. The presence of errors in this database was detected by checking the number of removed lines by the CLP.

At this point, it is interesting to mention that other variations of the algorithm have been tried to further improve the performance. In this sense, we tried to evaluate the CTC loss [17] to select a threshold  $\epsilon$ . We found it complex to deal with because it depends on several factors like the number of epochs in the training or if batch normalization has been applied. In future work, we expect to improve the algorithm in this way. Another promising research line could be introducing TL-DA and CLP in other DNN models, such as the based on GCN [9], that has a quite low value for 0 pages, to further improve the CER. Besides, introducing LM in the proposed DA-TL and CLP approaches could be also investigated.

## REFERENCES

- [1] N. Serrano, F. Castro, and A. Juan-Cfscar, "The RODRIGO database," in *Proc. LREC*, 2010, pp. 19–21.
- [2] R. Saabni, A. Asi, and J. El-Sana, "Text line extraction for historical document images," *Pattern Recognit. Lett.*, vol. 35, no. 1, pp. 23–33, Jan. 2014.
- [3] M. R. Yousefi, M. R. Soheili, T. M. Breuel, E. Kabir, and D. Stricker, "Binarization-free OCR for historical documents using LSTM networks," in *Proc. 13th Int. Conf. Document Anal. Recognit. (ICDAR)*, Aug. 2015, pp. 1121–1125.
- [4] E. Chammas, C. Mokbel, and L. Likforman-Sulem, "Handwriting recognition of historical documents with few labeled data," in *Proc. 13th IAPR Int. Workshop Document Anal. Syst. (DAS)*, Apr. 2018, pp. 43–48.
- [5] T. Strauß, G. Leifert, R. Labahn, T. Hodel, and G. Mühlberger, "ICFHR2018 competition on automated text recognition on a READ dataset," in *Proc. 16th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Aug. 2018, pp. 477–482.
- [6] J. C. A. Jaramillo, J. J. Murillo-Fuentes, and P. M. Olmos, "Boosting handwriting text recognition in small databases with transfer learning," in *Proc. 16th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Aug. 2018, pp. 429–434.
- [7] Y. Soullard, W. Swaileh, P. Tranouez, T. Paquet, and C. Chatelain, "Improving text recognition using optical and language model writer adaptation," in *Proc. Int. Conf. Document Anal. Recognit. (ICDAR)*, Sep. 2019, pp. 1175–1180.
- [8] J. A. Sánchez, V. Romero, A. H. Toselli, M. Villegas, and E. Vidal, "A set of benchmarks for handwritten text recognition on historical documents," *Pattern Recognit.*, vol. 94, pp. 122–134, Oct. 2019.
- [9] M. Yousef, K. F. Hussain, and U. S. Mohammed, "Accurate, data-efficient, unconstrained text recognition with convolutional neural networks," *Pattern Recognit.*, vol. 108, Dec. 2020, Art. no. 107482.
- [10] J. C. Aradillas, J. J. Murillo-Fuentes, and P. M. Olmos, "Improving offline HTR in small datasets by purging unreliable labels," in *Proc. 17th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Dortmund, Germany, Sep. 2020, pp. 25–30.
- [11] G. M. Binmakhashen and S. A. Mahmoud, "Document layout analysis: A comprehensive survey," *ACM Comput. Surveys*, vol. 52, no. 6, pp. 1–36, Jan. 2020.
- [12] S. He and L. Schomaker, "DeepOtsu: Document enhancement and binarization using iterative deep learning," *Pattern Recognit.*, vol. 91, pp. 379–390, Jul. 2019.
- [13] S. Ares Oliveira, B. Seguin, and F. Kaplan, "DhSegment: A generic deep-learning approach for document segmentation," in *Proc. 16th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Aug. 2018, pp. 7–12.
- [14] J. Oncina and M. Sebban, "Learning stochastic edit distance: Application in handwritten character recognition," *Pattern Recognit.*, vol. 39, no. 9, pp. 1575–1587, Sep. 2006.
- [15] Y.-C. Wu, F. Yin, and C.-L. Liu, "Improving handwritten chinese text recognition using neural network language models and convolutional neural network shape models," *Pattern Recognit.*, vol. 65, pp. 251–264, May 2017.
- [16] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [17] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *Proc. 23rd Int. Conf. Mach. Learn. (ICML)*, Jun. 2006, pp. 369–376.
- [18] V. Pham, T. Bluche, C. Kermorvant, and J. Louradour, "Dropout improves recurrent neural networks for handwriting recognition," in *Proc. 14th Int. Conf. Frontiers Handwriting Recognit.*, Sep. 2014, pp. 285–290.
- [19] P. Voigtlaender, P. Doetsch, and H. Ney, "Handwriting recognition with large multidimensional long short-term memory recurrent neural networks," in *Proc. 15th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Oct. 2016, pp. 228–233.
- [20] D. Castro, B. Bezerra, and M. Valenca, "Boosting the deep multidimensional long-short-term memory network for handwritten recognition systems," in *Proc. 16th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Aug. 2018, pp. 127–132.
- [21] B. Moysset and R. Messina, "Are 2D-LSTM really dead for offline text recognition?" *Int. J. Document Anal. Recognit.*, vol. 22, no. 3, pp. 193–208, Sep. 2019.
- [22] J. Puigcerver, "Are multidimensional recurrent layers really necessary for handwritten text recognition?" in *Proc. 14th IAPR Int. Conf. Document Anal. Recognit. (ICDAR)*, Nov. 2017, pp. 67–72.
- [23] D. Coquenec, C. Chatelain, and T. Paquet, "Recurrence-free unconstrained handwriting text recognition using gated fully convolutional network," in *Proc. 17th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Sep. 2020, pp. 19–24.

- [24] A. F. de Sousa Neto, B. L. D. Bezerra, A. H. Toselli, and E. B. Lima, "HTR-Flor++: A handwritten text recognition system based on a pipeline of optical and language models," in *Proc. ACM Symp. Document Eng.*, New York, NY, USA, Sep. 2020, pp. 1–4.
- [25] F. Naiemi, V. Ghods, and H. Khalesi, "A novel pipeline framework for multi oriented scene text image detection and recognition," *Expert Syst. Appl.*, vol. 170, May 2021, Art. no. 114549.
- [26] N. Daniyar, B. Kairat, K. Maksat, and A. Anel, "Classification of handwritten names of cities using various deep learning models," in *Proc. 15th Int. Conf. Electron., Comput. Comput. (ICECCO)*, Dec. 2019, pp. 1–4.
- [27] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 3320–3328.
- [28] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "DeCAF: A deep convolutional activation feature for generic visual recognition," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Jun. 2014, vol. 32, no. 1, pp. 647–655.
- [29] L. Zhu, Z. Huang, Z. Li, L. Xie, and H. T. Shen, "Exploring auxiliary context: Discrete semantic transfer hashing for scalable image retrieval," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 11, pp. 5264–5276, Nov. 2018.
- [30] M. Long, Y. Cao, Z. Cao, J. Wang, and M. I. Jordan, "Transferable representation learning with deep adaptation networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 41, no. 12, pp. 3071–3085, Dec. 2019.
- [31] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *Proc. Brit. Mach. Vis. Conf.*, 2014. [Online]. Available: <http://dx.doi.org/10.5244/C.28.6>
- [32] C. Wigginton, S. Stewart, B. Davis, B. Barrett, B. Price, and S. Cohen, "Data augmentation for recognition of handwritten words and lines using a CNN-LSTM network," in *Proc. 14th IAPR Int. Conf. Document Anal. Recognit. (ICDAR)*, vol. 1, Nov. 2017, pp. 639–645.
- [33] A. Poznanski and L. Wolf, "CNN-N-Gram for Handwriting Word recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2305–2314.
- [34] P. Krishnan and C. V. Jawahar, "Matching handwritten document images," in *Computer Vision—ECCV*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham, Switzerland: Springer, 2016, pp. 766–782.
- [35] X. Shen and R. Messina, "A method of synthesizing handwritten chinese images for data augmentation," in *Proc. 15th Int. Conf. Frontiers Handwriting Recognit. (ICFHR)*, Oct. 2016, pp. 114–119.
- [36] P. Y. Simard, D. Steinkraus, and J. C. Platt, "Best practices for convolutional neural networks applied to visual document analysis," in *Proc. 7th Int. Conf. Document Anal. Recognit.*, Aug. 2003, pp. 958–963.
- [37] H. Nisa, J. A. Thom, V. Ciesielski, and R. Tennakoon, "A deep learning approach to handwritten text recognition in the presence of struck-out text," in *Proc. Int. Conf. Image Vis. Comput. New Zealand (IVCNZ)*, Dec. 2019, pp. 1–6.
- [38] A. B. Salah, J. P. Moreux, N. Ragot, and T. Paquet, "OCR performance prediction using cross-OCR alignment," in *Proc. 13th Int. Conf. Document Anal. Recognit. (ICDAR)*, Aug. 2015, pp. 556–560.
- [39] B. Shi, X. Bai, and C. Yao, "An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 11, pp. 2298–2304, Nov. 2017.
- [40] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [41] U.-V. Marti and H. Bunke, "The IAM-database: An English sentence database for offline handwriting recognition," *Int. J. Document Anal. Recognit.*, vol. 5, no. 1, pp. 39–46, Nov. 2002.
- [42] A. Fischer, A. Keller, V. Frinken, and H. Bunke, "Lexicon-free handwritten word spotting using character HMMs," *Pattern Recognit. Lett.*, vol. 33, no. 7, pp. 934–942, May 2012.
- [43] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [44] S. Jialin Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010.
- [45] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features Off-the-shelf: An astounding baseline for recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2014, pp. 512–519.
- [46] B. Efron, "Better bootstrap confidence intervals," *J. Amer. Stat. Assoc.*, vol. 82, no. 397, pp. 171–185, Mar. 1987.
- [47] W. Swaileh, T. Paquet, Y. Soullard, and P. Tranouez, "Handwriting recognition with multigrams," in *Proc. 14th IAPR Int. Conf. Document Anal. Recognit. (ICDAR)*, vol. 1, Nov. 2017, pp. 137–142.



**JOSÉ CARLOS ARADILLAS** received the B.E. and M.Sc. degrees in telecommunication engineering from the Universidad de Sevilla, Spain, in 2015 and 2017, respectively, where he is currently pursuing the Ph.D. degree with the Department of Signal Theory and Communications. He is also an Assistant Professor with the Universidad de Sevilla. He has been a Visiting Researcher with the Max Planck Institute, Tübingen, Germany, and the Universidad Carlos III de Madrid, Spain. His current research interests include machine learning for communications and deep learning applied to cultural heritage.



**JUAN JOSÉ MURILLO-FUENTES** (Senior Member, IEEE) received the degree in telecommunication engineering from the Universidad de Sevilla, in 1996, and the Ph.D. degree in telecommunication engineering from the Universidad Carlos III de Madrid (UC3M), Spain, in 2001. He has been a Visiting Researcher with the University of Cambridge and UCL. Since 2016, he has been a Full Professor with the Universidad de Sevilla, teaching several bachelor's and master's courses (M.Sc. and Ph.D.) related to digital communications and machine learning. He is also a member of the Signal Processing and Learning Group, UC3M. He has published extensively, particularly in interdisciplinary fields, with more than 80 journal articles and conference papers in his active record. His research interests include algorithm development for signal processing and machine learning, and their applications to digital communication systems and cultural heritage.



**PABLO M. OLMOS** (Member, IEEE) received the M.Sc. degree and the Ph.D. degree in telecommunication engineering from the Universidad de Sevilla, in 2008 and 2011, respectively. He is currently an Assistant Professor with the Universidad Carlos III de Madrid. He has held appointments as a Visiting Researcher with Princeton University, EPFL, Notre Dame University, ENSEA, and Bell Labs. His research interests include approximate inference methods for Bayesian machine learning to information theory and digital communications.

...