



# Multimodal Deep Networks for Text and Image-Based Document Classification

Nicolas Audebert, Catherine Herold<sup>(✉)</sup>, Kuider Slimani, and Cédric Vidal

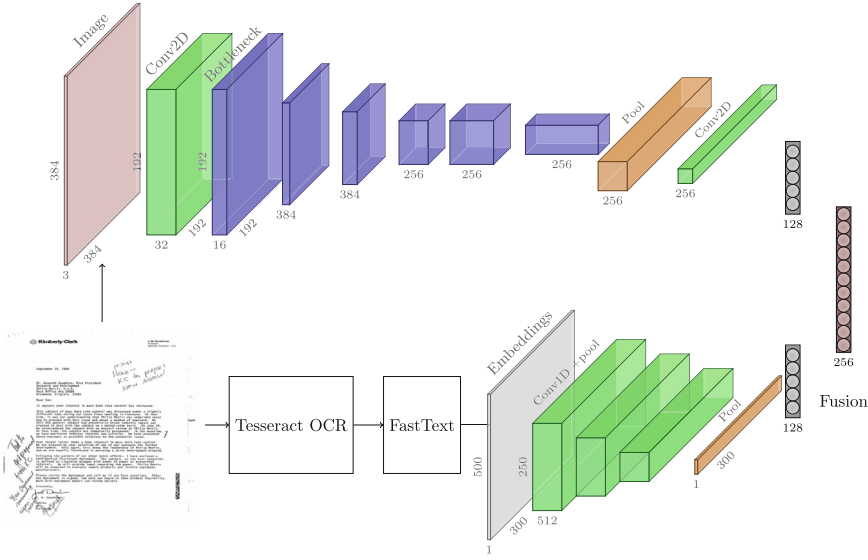
Quicksign, 38 rue du Sentier, 75002 Paris, France  
{nicolas.audebert,catherine.herold,kuider.slimani,  
cedric.vidal}@quicksign.com  
<https://www.quicksign.com/en/>

**Abstract.** Classification of document images is a critical step for accelerating archival of old manuscripts, online subscription and administrative procedures. Computer vision and deep learning have been suggested as a first solution to classify documents based on their visual appearance. However, achieving the fine-grained classification that is required in real-world setting cannot be achieved by visual analysis alone. Often, the relevant information is in the actual text content of the document, although this text is not available in digital form. In this work, we introduce a novel pipeline based on off-the-shelf architectures to deal with document classification by taking into account both text and visual information. We design a multimodal neural network that is able to learn both the image and from word embeddings, computed on noisy text extracted by OCR. We show that this approach allows us to improve single-modality classification accuracy by several points on the small Tobacco3482 and large RVL-CDIP datasets, even without clean text information. We release a post-OCR text classification (<https://github.com/Quicksign/ocrized-text-dataset>) that complements the Tobacco3482 and RVL-CDIP ones to encourage researchers to look into multi-modal text/image classification.

**Keywords:** Document classification · Text classification · Multimodal learning

## 1 Introduction

The ubiquity of computers and smartphones has incentivized governments and companies alike to digitize most of their processes. Onboarding new clients, paying taxes and proving one's identity is more and more done through a computer, as the rise of online banking has shown in the last few years. Industrial and public archives are also ongoing serious efforts to digitize their content in an effort for preservation, e.g. for old manuscripts, maps and documents with a historical value. This means that previously physical records, such as forms and identity documents, are now digitized and transferred electronically. In some cases, those records are produced and consumed by fully automated systems that rely

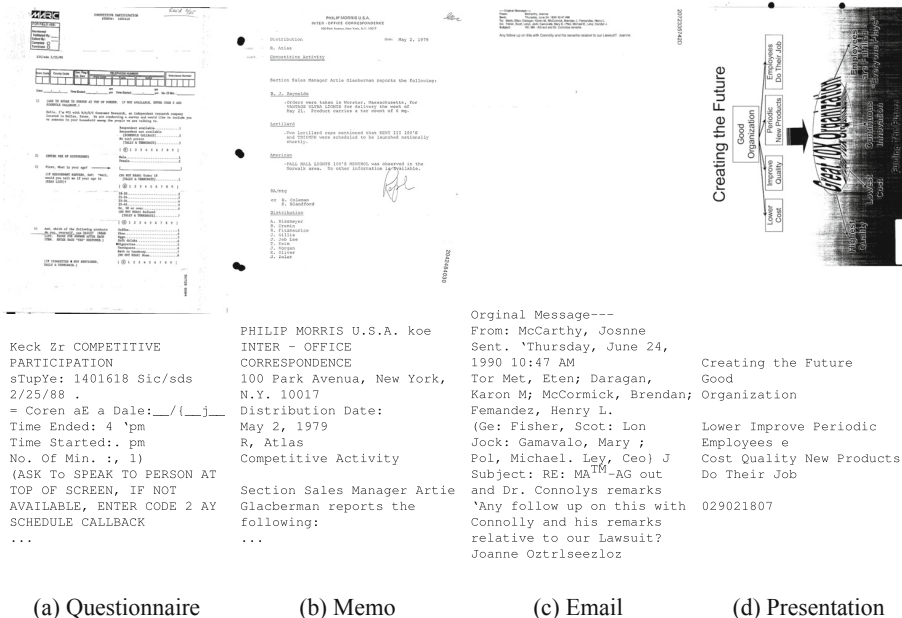


**Fig. 1.** Multimodal classifier for hybrid text/image classification. Training is performed end-to-end on both textual and visual features.

on machine-readable formats, such as XML or PDF with text layers. However, most of these digital copies are generated by end-users using whatever mean they have access to, i.e. scanners and cameras, especially from smartphones. For this reason, human operators have remained needed to proofread the documents, extract selected fields, check the records’ consistency and ensure that the appropriate files have been submitted. Automation through expert systems and machine learning can help accelerate this process to assist and alleviate the burden of this fastidious work for human workers.

A common task involved in data filing processes is document recognition, on which depends the class-specific rules that command each file. For example, a user might be asked to upload several documents such as a filled subscription form, an ID and a proof-of-residence. In this work, we tackle the document classification task to check that all required files have been sent so that they are filed accordingly.

Yet, if discriminating between broad classes of documents can be achieved based on their appearance only (e.g. separating passports from banking information), fine-grained recognition often depends on the textual content of the documents. For example, different tax forms might share their layout, logos and templates while the content in itself vastly differs. Computer vision has been interested for some time in optical character recognition (OCR) to extract text from images. However, dealing with both the textual and visual contents remains an open problem. In the past years, deep learning has been established as the new state-of-the-art for image classification and natural language processing. For



(a) Questionnaire

(b) Memo

(c) Email

(d) Presentation

**Fig. 2.** Document samples from the RVL-CDIP [11] dataset with corresponding text extracted by Tesseract OCR.

fine-grained document recognition, we expect the model to leverage both image and text information.

This work introduces a multimodal deep network that learns from both a document image and its textual content automatically extracted by OCR to perform its classification. We design a pragmatic pipeline for end-to-end heterogeneous feature extraction and fusion under time and cost constraints. We show that taking both the text and the document appearance into account improves both single modality baselines by several percents on two datasets from the document recognition literature. We detail some limitations of the current academic datasets and give leads for an application in an industrial setting with unclean data, such as photographed documents.

## 2 Related Work

Analyzing digitized documents is an old task in computer vision that was boosted by the dissemination of computers in offices and then of digital cameras and smartphones in everyday life. To allow for textual search and easy indexing, the critical part of digitization is extracting text content from documents that have been scanned or photographed. Indeed, either when scanning or taking a picture of the document, its actual text is lost, although it is implicitly embedded in the


pixel values of the image. Numerous optical character recognition (OCR) algorithms have been designed to transform images into strings of characters [17,37]. Despite those efforts perfectly reading any type of document remains challenging due to the wide variety of fonts and languages. Layout analysis is a way to preprocess the data to detect text areas and find the text orientation in order to enforce a better local and global consistency [15,20].

Document image analysis is also one of the first topic where modern deep learning has been applied. The first convolutional neural network (CNN) [21] was originally designed for classification of digits and letters. The computer vision community deployed consequent efforts to achieve image-based document classification without text, as shown by a 2007 survey [7] which focuses on document image classification without OCR results. As an example, [19] introduced SURF visual features with a bag-of-words scheme to perform document image classification and retrieval. In 2015, [11] introduced a large labeled image document dataset which sparked interest and generated several studies of deep CNN on this topic [1,9,36], inspired by the success of these networks on ImageNet and tuning data augmentation policies, transfer learning strategies and domain adaptation for document classification. In the same idea, [35] also investigated such deep architectures to classify identity documents. [2] goes even further by trying to segment the full layout of a document image into paragraphs, titles, ornaments, images etc. These models focus on extracting strong visual features from the images to classify the documents based on their layout, geometry, colors and shape.

On the other hand, text-based document classification has also long been investigated. In 1963, [6] introduced an algorithmic approach to classify scientific abstracts. More recently, [23] experimented with one-class SVM for document classification based on various text features, such as TF-IDF. [33] used Latent Dirichlet Allocation to perform topic modeling and used it as a generative approach to document classification. The recent appearance of learned word embeddings approaches such as word2vec [24] or ELMo [30] paved the way to a large body of works related to recurrent and attention mechanisms for text classification. For example, [39] proposed a bidirectional recurrent network with a hierarchical attention mechanism that learns both at the word and sentence levels to improve document classification.

Some works tried to reconcile the text-based and image-based approaches to exploit both information sources. [26] performs OCR to detect keywords in images which are then encoded as colored boxes before passing the image through a CNN. While a clever trick, this does not leverage the representation power of word embeddings. Closer to our approach, [38] goes further by generating text feature maps that are combined with visual feature maps in a fully convolutional network. However, the considered documents are synthetic and the network is trained using perfectly clean texts and images, which is unrealistic for practical uses. More similar to us, [4] learns to combine bag of words and bag of visual words features for industrial document images using a statistical model combining outputs of two single-modality classifiers. While using shallow features, they

show that using both information allows for a better accuracy when the OCR is unreliable, which is often the case in an industrial setting.



**Fig. 3.** MobileNetV2 uses inverted residual blocks to reduce the number of channels that are forwarded in subsequent layers. Figure from [34].

In this paper, we go further in this direction and propose a new baseline with a hybrid deep model. In order to classify OCRized document images, we present a pragmatic pipeline perform visual and textual feature extraction using off-the-shelf architectures. To leverage the complementary information present in both modalities, we design an efficient end-to-end network that jointly learn from text and image while keeping computation cost at its minimum. We build on existing deep models (MobileNet and FastText) and demonstrate significant improvements using our fusion strategy on two document images dataset.

### 3 Learning on Text and Image

#### 3.1 Visual Features


There is a large literature both in general image recognition and in image document classification. Recent works have established deep convolutional neural networks as the *de facto* state of the art on many competitions in object recognition, detection and segmentation, e.g. ImageNet. Deep features, extracted by pretrained or fine-tuned deep CNNs, constitute a strong baseline for visual recognition tasks [32]. Based on this, we choose to fine-tune a CNN pretrained on ImageNet in order to extract visual features on our images, as suggested in several recent document classification publications [1, 11, 36]. As we aim to perform inference on a large volume of data with time and cost constraints, we focus on a lightweight architecture with competitive classification performance, in our case the MobileNet v2 model [34].

MobileNetV2 [34] consists in a stack of bottleneck blocks. Based on the residual learning principle [13], each bottleneck block transforms a feature map first by expanding it by increasing its number of channels with a  $1 \times 1$  convolutional layer with identity activation. Then, a  $3 \times 3$  depthwise convolution is performed, followed by a ReLU and a final  $1 \times 1$  convolution with ReLU. For efficiency issues, this block inverts the traditional residual block since the expansion is performed inside the block, whereas residual blocks compress and then

reexpand the information, as illustrated in Fig. 3. The final MobileNetV2 contains 19 residual bottleneck layers. Compared to other state of the art CNNs, MobileNetV2’s accuracy is on-par with VGG-16 while being significantly faster.

### 3.2 Textual Features

Since our use case focuses on document images in which the text has not been transcribed, we need to perform an OCR step. To this end, we use the Tesseract OCR engine [17] in its 4.0 version which is based on an LSTM network. Tesseract is configured in English to use full page segmentation and the LSTM engine. In practice, this means that Tesseract will try to detect the text orientation in the image and perform the needed affine transformation and rotation if any. Tesseract also deals with the image binarization using Otsu’s thresholding to identify black text on white background [27]. This will suffice on the datasets described in Sect. 4.1, although we found Tesseract challenging to apply on real-world images, especially pictures which are not flat and grayscale scans.



(a) Document distribution w.r.t of non-dictionary words % in the Tobacco3482-Tesseract corpus.

Word pair	Similarities		
specifically	GloVe	ELMo	FastText
Specificallily	0.71	0.68	0.96
filter	GloVe	ELMo	FastText
fiilter	0.91	0.73	0.96
alcohol	GloVe	ELMo	FastText
Aleohol	0.40	0.69	0.88
Largely	GloVe	ELMo	FastText
Largly	0.25	0.81	0.98

(b) Word embeddings similarity for misspelled words.

**Fig. 4.** Tesseract OCR outputs noisy text that does not entirely overlap with the assumptions usually held when training word embeddings for NLP.

Recent literature in NLP suggests that pretrained word embeddings offer a strong baseline which surpasses traditional shallow learning approaches. Many word embeddings have been designed following the initial success of *word2vec* [24], such as GloVe [29] or more recently the contextualized word embeddings from ELMo [30].

However, those word embeddings assume a good tokenization of the words, i.e. most embeddings remove digits, ignore punctuation and do not deal with out-of-vocabulary (OOV) words. Since these embeddings are learned on clean corpus (e.g. Wikipedia or novels), tokenization is fairly straightforward. OOV words are either assigned a random embedding or mapped to the closest in-vocabulary word based on the Levenshtein distance.

Unfortunately, outputs of the Tesseract OCR are noisy and not as clean as the training data from these embeddings. Even in grayscale, well-oriented

documents, OCR might have trouble dealing with diacritics, exotic fonts or curved text, as illustrated by the extracts from Fig. 2. Moreover, specific user domains (e.g. banking or medieval manuscripts) might use rare words, codes, abbreviations or overall jargon that is absent from general-purpose word embeddings. Since we face many possible misspellings in the extracted text, we cannot use the previous workarounds for OOV embeddings since it would inject a lot of non-discriminant features in our text representation. In average, on the Tobacco3482 corpus, a document processed by Tesseract OCR contains 136 words with 4 characters or more. Of those, only 118 in average are in the GloVe embeddings [29]<sup>1</sup> and only 114 are in Enchant’s spellchecker US English dictionary. Overall, approximately 26% of the corpus is absent from the US English dictionary and 23% from the GloVe embeddings. The document distribution with respect to the proportion of out-of-vocabulary words is shown in Fig. 4a. Although most of the documents are concentrated around 10% of OOVs, there is a significant long tail including several dozens of documents that contain only words outside of the English language.

Therefore, we turn to character-based word embeddings that are able to deal with OOV words by assigning them plausible word vectors that preserve both a semantic and a spelling similarity. One possibility was to use the mimicking networks from [31] that learn to infer word embeddings such as GloVe, but based only on subword information. More complex embeddings such as FastText [5, 16] and ELMo [30], which produce vectors using respectively n-grams and subword information, can also address this problem. Finally, the Magnitude library [28] uses two alternative strategies to deal with OOV words:

- Assigning a *deterministic* random vector. These vectors do not capture semantic sense, however similar words based on the Levenshtein-Damerau distance will have similar vectors. Misspellings will therefore not be close to the original word, but similar lingo words will be close.
- Using character n-grams inspired by [5] and interpolation with in-vocabulary words, Magnitude can generate vectors for OOV words which are sensible based on existing learned embedding.

Preliminary data exploration shows that subword-aware embeddings perform better at preserving similarity despite misspellings, as illustrated in Fig. 4b. We therefore focus our interest on the FastText embedding, which is faster than ELMo since the latter requires passing the context through a bidirectional LSTM during inference. It is worth noting that this raises concern for characters that have not been seen by FastText. We found experimentally that Tesseract OCR generated no character that was OOV for FastText on the documents we considered.

Finally, it is necessary to convert those word embeddings into a document embedding. We consider two approaches:

- The simple baseline for sentence embedding suggested in [3], which consists in a weighted average of word embeddings altered by PCA.

---

<sup>1</sup> Based on the Wikipedia 2014 + Gigaword 5 datasets.

- Using variable-length document embeddings consisting in a sequence of word embeddings.

The first approach is suitable as generic feature while the second requires a statistical model able to deal with sequences, such as recurrent or convolutional neural networks. For both methods, we use the SpaCy small English model [14] to perform the tokenization and punctuation removal. Individual word embeddings are then inferred using FastText [5] pretrained on the Common Crawl dataset.

### 3.3 Multimodal Features

Once text and image features have been extracted, we feed them to a multi-layer perceptron following [10]. To do so, we need to combine both feature vectors into one. Two approaches can be envisioned:

- Adaptive averaging of both feature vectors. This aligns both feature spaces so that scalars at the same index become compatible by summation, i.e. that each dimension of the vectors have a similar semantic meaning.
- Concatenating both vectors. This does not imply that both feature spaces can be aligned and delegates to the fusion MLP the task of combining the two domains.

Both fusion strategies are differentiable, therefore the whole network can be trained in an end-to-end fashion. Moreover, the model is modular and each feature extractor can be swapped for another model, e.g. MobileNet can be exchanged with any other popular CNN and FastText could be replaced by subword-level NLP models, even differentiable ones that could allow fine-tuning the embeddings. In this work, we try to keep things simple and build on robust base networks in order to clearly understand how the data fusion impacts model performance. Preliminary experiments showed that the summation fusion significantly underperformed compared to pure image baseline. We suggest that this is provoked by the impossibility of aligning the text and image feature spaces without breaking their discriminating power, resulting in suboptimal space. Therefore, we move on with the concatenation strategy for the rest of this paper. The complete pipeline is illustrated in Fig. 1.

## 4 Experimental Setup

### 4.1 Datasets

**Tobacco3482.** The Tobacco3482 dataset [19] contains 3482 black and white documents, a subset from the Truth Tobacco Industry Documents<sup>2</sup> archives of legal proceedings against large American tobacco companies. There are annotations for 10 classes of documents (e.g. email, letter, memo. . .). Following common practices, we perform k-fold cross-validation using 800 documents for training and the rest for testing. Results are averaged over 3 runs.

<sup>2</sup> <https://www.industrydocuments.ucsf.edu/tobacco/>.



**Table 1.** Preliminary tuning of the single-modality baselines on Tobacco3482.

(a) Preliminary experiments on Tobacco3482 for the text baseline.			(b) Preliminary experiments on Tobacco3482 for the image baseline.		
Model	OA	$F_1$	Model	OA	$F_1$
MLP (document)	70.8%	0.69	MobileNetV2	84.5%	0.82
CNN 1D (word sequence)	73.9%	0.71	MobileNetV2 (w/ DA)	83.9%	0.82

OA = overall accuracy,  $F_1$  = class-balanced  $F_1$  score. DA = data augmentation.

**RVL-CDIP.** The RVL-CDIP dataset [11] is comprised of 400 000 grayscale digitized documents from the Truth Tobacco Industry Documents. There are annotations for 16 classes of documents (e.g. email, letter, invoice, scientific report...), each containing 25 000 samples. We use the standard train/val/test split from [11] with 320 000 documents for training, 40 000 for validation and 40 000 for testing.

**Text Generation.** The Tobacco3482 and RVL-CDIP are image-based datasets. In order to evaluate our multi-modal networks, we wish to learn from both visual and textual content. Therefore we use the Tesseract OCR library<sup>3</sup> to extract text from the grayscale images. We perform this operation on both datasets. We release the OCR text dataset openly<sup>4</sup> to encourage other researchers to replicate our work or test their own model for post-OCR text classification or multi-modal text/image classification.

## 4.2 Models

This subsection describes the implementation details of our deep networks. All models are implemented in TensorFlow 1.12 using the Keras API and trained using a NVIDIA Titan X GPU. Hyperparameters were manually selected on a subset of Tobacco3482 and fixed for all experiments.

**Text Baseline.** Seeing that our representation of textual data can be either a document embedding or a sequence of word embeddings, we compare two models for our text baseline.

The first model is an improved Multi-Layer Perceptron (MLP) with ReLU activations, Dropout and Batch Normalization (BN) after each layer. The network has a fixed width of 2048 neurons for all layers except the last one, which produces a 128 feature vector, classified by a softmax layer. Weights are randomly initialized using He’s initialization [12]. The averaged document embedding [3] is used as an input for this classifier.

<sup>3</sup> <https://github.com/tesseract-ocr/tesseract/>.

<sup>4</sup> The QS-OCR dataset is available at: <https://github.com/Quicksign/ocrized-text-dataset>.

The second model is a one-dimensional convolutional neural network designed inspired by previous work for sentence classification [18]. The CNN is 4-layers deep and interlaces 1D convolutions with a window of size 12 with maxpooling with a stride of 2. Each layer consists in 512 channels with ReLU activation. The final feature map is processed by a max-pooling-through-time layer that extracts maximal features on the sequence on top of which we apply Dropout for regularization. A fully connected layer then maps the features to the softmax classifier. The input word sequence is zero-padded up to 500 words for documents with less 500 words.

We experiment on the Tobacco3482 dataset in order to evaluate which text model to choose. Results are reported in Table 1a. Without surprise, the CNN 1D outperforms significantly the MLP classifier. The pattern recognition abilities of the convolutional network makes it possible to interpret the word sequences by leveraging contextual information. Since only some part of the text might be relevant, averaging over all word embeddings dilute the discriminating information. Moreover, noisy embeddings due to garbage output from Tesseract (e.g. incoherent strings where OCR has failed) are included in the final document embedding. However, when dealing with word sequences, convolutional layers and temporal max-pooling help extracting only the relevant information. Therefore, we choose to include the 1D CNN as the text component in our multimodal architecture.

This model is denoted **TEXT** in the rest of the paper. It is optimized using Stochastic Gradient Descent with momentum for 100 epochs, with a learning rate of 0.01, a momentum of 0.9 and a batch size of 40<sup>5</sup>.

**Image Baseline.** We investigate as our base CNN the lightweight MobileNetV2 [34] which focuses on computing efficiency, albeit at the cost of a slightly lower top-1 accuracy on ImageNet compared to other state of the art CNN. We train the CNN on grayscale document images resized at  $384 \times 384$ . Although this warps the aspect ratio, [36] reports better accuracy than when using padding at the same resolution. As the model is designed for RGB images, the grayscale channel is duplicated three times. This allows us to initialize the network by loading its pretrained weights on ImageNet, which accelerates convergence and slightly improves accuracy through transfer learning.

This model is denoted **IMAGE** in the rest of the paper. It is optimized using Stochastic Gradient Descent with momentum for 200 epochs, with a learning rate of 0.01, a momentum of 0.9 and a batch size of 40.

As reported in Table 1b, preliminary experiments on the Tobacco3482 with random JPEG artifacts, saturation and contrast alterations did not significantly alter the classifier’s accuracy compared to no augmentation. This is explained by the low variability between the grayscale document images. All images are grayscale with dark text on white background with horizontal text lines, therefore color and geometric augmentation are not necessary. However, [36] report some success using shear transform, which we did not consider in this work. It is

<sup>5</sup> Hyperparameters are manually tuned on a small validation set.

worth noting that compared with previous literature on the RVL-CDIP dataset, e.g. [1, 11, 36], we do not average predictions over multiple crops at inference time for speed concerns. This might explain why our visual baseline underperforms the current state of the art in this state (although this does not question the gains due to the multi-modal network).

**Table 2.** Overall accuracy on the RVL-CDIP dataset.

Model	<b>IMAGE</b>	<b>TEXT</b>	<b>FUSION</b>	CNNs [11]	VGG-16 [1]	AlexNet+SPP [36]
OA	89.1%	74.6%	90.6%	89.8%	90.97%	90.94%

OA = Overall Accuracy.

**Fusion.** For our multimodal network, we consider the same model as our baselines except that the final layers are cut-off. For the **TEXT** model, the last layer produces an output vector of dimension 128 instead of the number of classes. For the **IMAGE** model, we aggregate the last convolutional features using global average pooling on each channel, which produces a feature vector of dimension 1280. We then map this feature vector using a fully connected layer to a representation space of dimension 128.

**Table 3.** Overall accuracy and  $F_1$  scores on the Tobacco3482 datasets.

Model	OA	$F_1$	Adv.	Email	Form	Letter	Memo	News	Notes	Report	Res.	Sci.
CNNs [11]	79.9	–	–									
<b>TEXT</b>	73.8	0.71	0.60	0.96	0.76	0.71	0.79	0.67	0.62	0.43	<b>0.97</b>	0.57
<b>IMAGE</b>	84.5	0.82	<b>0.94</b>	0.96	0.85	0.83	0.90	0.89	0.83	0.61	0.80	0.62
<b>FUSION</b>	<b>87.8</b>	<b>0.86</b>	0.93	<b>0.98</b>	<b>0.88</b>	<b>0.86</b>	<b>0.90</b>	<b>0.90</b>	<b>0.85</b>	<b>0.71</b>	0.96	<b>0.68</b>
Oracle	92.1	0.91	0.94	0.99	0.94	0.92	0.93	0.93	0.89	0.81	0.97	0.79

Adv. = Advertisement, Res. = Resume, Sci. = Scientific.

This model is denoted **FUSION** in the rest of the paper. It is optimized using Stochastic Gradient Descent with momentum for 200 epochs, with a learning rate of 0.01, a momentum of 0.9 and a batch size of 40.

## 5 Discussion

### 5.1 Performances

Model performances scores on Tobacco3482 and RVL-CDIP are reported in Tables 2 and 3. Behaviour of all models is consistent both on the smaller dataset and on the very large one. In both cases, the **TEXT** baseline is significantly underperforming the **IMAGE** one. Indeed, as could be seen in Fig. 2, Tesseract OCR outputs noisy text. This includes words that have been misspelled – which

are correctly dealt with by the FastText embeddings – and new words that are hallucinated due to poor binarization or salt-and-pepper noise in the image. Moreover, layout and visual information tends to be more informative based on how the classes were defined: scientific papers, news and emails follow similar templates while advertisements present specific graphics. However, in both cases, this simple document embedding is enough to classify more than 70% of the documents, despite its roughness.

Using the **IMAGE** model only, we reach accuracies competitive with the state of the art. MobileNetV2 alone does on-par is with the holistic CNN ensemble from [11] and is competitive with fine-tuned GoogLeNet and ResNet-50 [1] (90.97%).

On both datasets, the fusion scheme is able to improve the overall accuracy by  $\simeq 1.5\%$  which demonstrates the relevance of our approach. While the document embedding we chose is simple, it appears to be at least partially robust to OCR noise and to preserve enough information about the document content to boosts CNN accuracy on document image classification even further. We also report the results from an oracle, which corresponds to the perfect fusion of the **TEXT** and **IMAGE** baselines, i.e. a model that would combine the predictions from both single-modality networks and always choose the right one. The oracle corresponds to the theoretical maximal accuracy boost that we could expect from the **FUSION** model. On Tobacco3482, the oracle corresponds to a 7.6% absolute improvement (9% relative). In our case, the **FUSION** model improves the best single-source baseline by an absolute 3.3% (4% relative), which is significant although still leaves the door open to further improvements. More importantly, the gains are consistent on all classes of interest, almost never underperforming one of the two base networks on any class. This confirm the proposed approach as the two sources, image and text, give complementary information to classify a document.

## 5.2 Processing Time

Although some applications of document image recognition can be performed offline, most of the time users upload a document and expect near real-time feedback. User experience engineering [25] indicates than less than 1 s is the maximum latency the user can suffer before the interface feels sluggish, and 10 s is the maximum delay before they start losing their attention. On the RVL-CDIP dataset, Tesseract processes a document image in  $\simeq 910$  ms in average on an Intel Core i7-8550U CPU using 4 threads, including loading the image from disk. This means that every additional latency induced by the network inference time is critical since it will negatively affect the user experience.

On the same CPU, the full inference using the **FUSION** model takes  $\simeq 360$  ms including loading, resizing and normalizing the image. The complete process including Tesseract OCR therefore takes less than  $\simeq 1300$  ms which is acceptable in a system requiring user input. Of those, 130 ms are spent in the 1D CNN (including reading the file and performing FastText inference) and 230 ms in MobileNetV2 (including image preprocessing). The overhead added by the final

fusion layer is negligible. We stress that this is using a standard TensorFlow without any CPU-specific compilation flags, which could speed up the inference further. On a NVIDIA Titan X GPU, the **FUSION** network runs in 110 ms (50 ms for **TEXT**, 60 ms for MobileNetV2), which brings the total just above the 1 s recommendation. In our case, using compute-efficient architectures allow us to avoid running on an expensive and power-hungry GPU.

As a comparison basis, other architecture choices that we dismissed earlier would have resulted in poorer performance and the network would not be usable in a near real-time user application. For example, the Xception network [8] takes 630 ms to run during inference with the same parameters and hardware. For the text model, an LSTM-based RNN with a similar depth takes many seconds to run.

Note that, although this does not reduced the perceived delay for one user, the global throughput of the system can be improved by batching the images. Two Tesseract processes can leverage the full eight cores from an Intel Core i7-8550U CPU. In this setting, processing an image takes  $\simeq 660$  ms in average. Thanks to the batch efficiency of neural networks, the average processing time becomes  $\leq 750$  ms on GPU and  $\leq 1000$  ms on CPU. This is particularly helpful when users have several documents to upload that can be processed concurrently.

### 5.3 Limitations

One of the main limitation of this work stems from the public document image datasets available. Indeed, in a real-world application, document images can be grayscale, RGB, scanned images and photographs with various rotations, brightness, contrast and hue values. The Tobacco documents are all oriented in the right way, which makes it easier for Tesseract to perform OCR. Moreover, documents have been scanned by professionals who tried to maximize their legibility while user-generated often presents poor quality.

While it was not required here, data augmentation is definitely required for practical applications to encompass the large variety of environmental conditions in which documents are digitized. This is especially true for rotations, since it is often not possible to ensure that users will capture the document with the right orientation and Tesseract does not always correctly detects it. For industrial-grade applications dealing with user-generated content, such a data augmentation is necessary to alleviate overfitting and reduce the gap between train and actual data. Preprocessing page segmentation and layout analysis tools, such as dhSegment [2] can also bring significant improvements by renormalizing image orientation and cropping the document before sending it to the classifier.

Moreover, as we have seen, the post-OCR word embeddings include lots of noisy or completely wrong words that generate OOV errors. In practical applications, we found beneficial to perform a semantic tokenization and named entity recognition using SpaCy. This allows us to perform a partial spellchecking, e.g. using `symspell`<sup>6</sup> to correct words that have been misread by Tesseract.

<sup>6</sup> <https://github.com/wolfgarbe/SymSpell>.

act, without affecting proper nouns or domain-specific abbreviations and codes. If this can deal frequent misspellings of words, it might also suppress out-of-vocabulary words such as alphanumeric codes. Therefore, learning domain specific, character-based or robust-to-OCR embeddings [22] is an interesting lead for future research, as the current interest in the ICDAR2019 competition on Post-OCR Text Correction shows<sup>7</sup>.

## 6 Conclusion

In this work, we tackled the problem of document classification using both image and text contents. Based only on an image of a digitized document, we try to perform a fine-grained classification using visual and textual features. To do so, we first used Tesseract OCR to extract the text from the image. We then compute character-based word embeddings using FastText on the noisy Tesseract output and generate a document embedding which represents our text features. Their counterpart visual features are learned using MobileNetv2, a standard CNN from the state of the art. Using those pragmatic approaches, we introduce an end-to-end learnable multimodal deep network that jointly learns text and image features and perform the final classification based on a fused heterogeneous representation of the document. We validated our approach on the Tobacco3482 and RVL-CDIP datasets showing consistent gains both on small and large datasets. This shows that there is a significant interest into hybrid image/text approach even when clean text is not available for document image classification and we aim to further investigate this topic in the future.

## References

1. Afzal, M.Z., Kölsch, A., Ahmed, S., Liwicki, M.: Cutting the error by half: investigation of very deep CNN and advanced training strategies for document image classification. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), vol. 01, pp. 883–888, November 2017. <https://doi.org/10.1109/ICDAR.2017.149>
2. Ares Oliveira, S., Seguin, B.L.A., Kaplan, F.: dhSegment: a generic deep-learning approach for document segmentation, August 2018
3. Arora, S., Liang, Y., Ma, T.: A simple but tough-to-beat baseline for sentence embeddings. In: Proceedings of the International Conference on Learning Representations (ICLR), November 2016
4. Augereau, O., Journet, N., Vialard, A., Domenger, J.: Improving classification of an industrial document image database by combining visual and textual features. In: 2014 11th IAPR International Workshop on Document Analysis Systems, pp. 314–318, April 2014. <https://doi.org/10.1109/DAS.2014.44>
5. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching word vectors with subword information. *Trans. Assoc. Comput. Linguist.* **5**, 135–146 (2017)
6. Borko, H., Bernick, M.: Automatic document classification. *J. ACM* **10**(2), 151–162 (1963). <https://doi.org/10.1145/321160.321165>

<sup>7</sup> <https://sites.google.com/view/icdar2019-postcorrectionocr>.

7. Chen, N., Blostein, D.: A survey of document image classification: problem statement, classifier architecture and performance evaluation. *Int. J. Doc. Anal. Recogn. (IJ DAR)* **10**(1), 1–16 (2007). <https://doi.org/10.1007/s10032-006-0020-2>
8. Chollet, F.: Xception: deep Learning with depthwise separable convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, United States, pp. 1800–1807, July 2017. <https://doi.org/10.1109/CVPR.2017.195>
9. Das, A., Roy, S., Bhattacharya, U., Parui, S.K.: Document image classification with intra-domain transfer learning and stacked generalization of deep convolutional neural networks. In: *2018 24th International Conference on Pattern Recognition (ICPR)*, pp. 3180–3185, August 2018. <https://doi.org/10.1109/ICPR.2018.8545630>
10. Eitel, A., Springenberg, J.T., Spinello, L., Riedmiller, M., Burgard, W.: Multimodal deep learning for robust RGB-D object recognition. In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 681–687, September 2015. <https://doi.org/10.1109/IROS.2015.7353446>
11. Harley, A.W., Ufkes, A., Derpanis, K.G.: Evaluation of deep convolutional nets for document image classification and retrieval. In: *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pp. 991–995, August 2015. <https://doi.org/10.1109/ICDAR.2015.7333910>
12. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1026–1034, December 2015. <https://doi.org/10.1109/ICCV.2015.123>
13. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, United States, pp. 770–778, June 2016. <https://doi.org/10.1109/CVPR.2016.90>
14. Honnibal, M., Montani, I.: spaCy 2: natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing (2017, to appear)
15. Imade, S., Tatsuta, S., Wada, T.: Segmentation and classification for mixed text/image documents using neural network. In: *Proceedings of 2nd International Conference on Document Analysis and Recognition (ICDAR 1993)*, pp. 930–934, October 1993. <https://doi.org/10.1109/ICDAR.1993.395584>
16. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431. Association for Computational Linguistics (2017)
17. Kay, A.: Tesseract: an open-source optical character recognition engine. *Linux J.* **2007**(159), 2 (2007)
18. Kim, Y.: Convolutional neural networks for sentence classification. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1746–1751, October 2014. <https://doi.org/10.3115/v1/D14-1181>
19. Kumar, J., Ye, P., Doermann, D.: Structural similarity for document image classification and retrieval. *Pattern Recogn. Lett.* **43**, 119–126 (2014). <https://doi.org/10.1016/j.patrec.2013.10.030>
20. Le, D.X., Thoma, G.R., Wechsler, H.: Classification of binary document images into textual or nontextual data blocks using neural network models. *Mach. Vis. Appl.* **8**(5), 289–304 (1995). <https://doi.org/10.1007/BF01211490>

21. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proc. IEEE* **86**(11), 2278–2324 (1998). <https://doi.org/10.1109/5.726791>
22. Malykh, V., Logacheva, V., Khakhulin, T.: Robust word vectors: context-informed embeddings for noisy texts. In: *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-Generated Text*, pp. 54–63, November 2018
23. Manevitz, L.M., Yousef, M.: One-class SVMs for document classification. *J. Mach. Learn. Res.* **2**(Dec), 139–154 (2001)
24. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. In: *ICLR 2013*, January 2013
25. Nielsen, J.: *Usability Engineering*. Morgan Kaufmann Publishers Inc., San Francisco (1993)
26. Noce, L., Gallo, I., Zamberletti, A., Calefati, A.: Embedded textual content for document image classification with convolutional neural networks. In: *Proceedings of the 2016 ACM Symposium on Document Engineering (DocEng 2016)*, pp. 165–173. ACM, New York (2016). <https://doi.org/10.1145/2960811.2960814>
27. Otsu, N.: A threshold selection method from gray-level histograms. *IEEE Trans. Syst. Man Cybern.* **9**(1), 62–66 (1979). <https://doi.org/10.1109/TSMC.1979.4310076>
28. Patel, A., Sands, A., Callison-Burch, C., Apidianaki, M.: Magnitude: a fast, efficient universal vector embedding utility package. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 120–126, November 2018
29. Pennington, J., Socher, R., Manning, C.: GloVe: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543. Association for Computational Linguistics (2014). <https://doi.org/10.3115/v1/D14-1162>
30. Peters, M., et al.: Deep contextualized word representations. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 2227–2237. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/N18-1202>
31. Pinter, Y., Guthrie, R., Eisenstein, J.: Mimicking word embeddings using Subword RNNs. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 102–112. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/D17-1010>
32. Razavian, A.S., Azizpour, H., Sullivan, J., Carlsson, S.: CNN features off-the-shelf: an astounding baseline for recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 512–519, June 2014. <https://doi.org/10.1109/CVPRW.2014.131>
33. Rubin, T.N., Chambers, A., Smyth, P., Steyvers, M.: Statistical topic models for multi-label document classification. *Mach. Learn.* **88**(1), 157–208 (2012). <https://doi.org/10.1007/s10994-011-5272-5>
34. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.: MobileNetV2: inverted residuals and linear bottlenecks. In: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, June 2018. <https://doi.org/10.1109/CVPR.2018.00474>



35. Sicre, R., Awal, A.M., Furon, T.: Identity documents classification as an image classification problem. In: Battiato, S., Gallo, G., Schettini, R., Stanco, F. (eds.) ICIAP 2017. LNCS, vol. 10485, pp. 602–613. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-68548-9\\_55](https://doi.org/10.1007/978-3-319-68548-9_55)
36. Tensmeyer, C., Martinez, T.: Analysis of convolutional neural networks for document image classification. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), vol. 01, pp. 388–393, November 2017. <https://doi.org/10.1109/ICDAR.2017.71>
37. Wong, K.Y., Casey, R.G., Wahl, F.M.: Document analysis system. IBM J. Res. Dev. **26**(6), 647–656 (1982). <https://doi.org/10.1147/rd.266.0647>
38. Yang, X., Yumer, E., Asente, P., Kraley, M., Kifer, D., Giles, C.L.: Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4342–4351, July 2017. <https://doi.org/10.1109/CVPR.2017.462>
39. Yang, Z., Yang, D., Dyer, C., He, X., Smola, A., Hovy, E.: Hierarchical attention networks for document classification. In: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 1480–1489. Association for Computational Linguistics (2016). <https://doi.org/10.18653/v1/N16-1174>