

Multi-label classification of promotions in digital leaflets using textual and visual information

Roberto Arroyo, David Jiménez-Cabello and Javier Martínez-Cebrián

Nielsen Connect R&D AI

Calle Salvador de Madariaga, 1, 28027, Madrid, Spain

<https://www.nielsen.com/>

{roberto.arroyo, david.jimenez, javier.martinezcebrian}@nielsen.com

Abstract

Product descriptions in e-commerce platforms contain detailed and valuable information about retailers assortment. In particular, coding promotions within digital leaflets are of great interest in e-commerce as they capture the attention of consumers by showing regular promotions for different products. However, this information is embedded into images, making it difficult to extract and process for downstream tasks. In this paper, we present an end-to-end approach that classifies promotions within digital leaflets into their corresponding product categories using both visual and textual information. Our approach can be divided into three key components: 1) region detection, 2) text recognition and 3) text classification. In many cases, a single promotion refers to multiple product categories, so we introduce a multi-label objective in the classification head. We demonstrate the effectiveness of our approach for two separated tasks: 1) image-based detection of the descriptions for each individual promotion and 2) multi-label classification of the product categories using the text from the product descriptions. We train and evaluate our models using a private dataset composed of images from digital leaflets obtained by Nielsen. Results show that we consistently outperform the proposed baseline by a large margin in all the experiments.

1 Introduction

The latest advances in Artificial Intelligence (AI) have provided new tools to enhance the automation of different recognition problems. We are witnessing a clear trend to merge different domains within AI to obtain better representations for the most complex problems. Many recent approaches merge textual and visual information by applying Natural Language Processing (NLP) and Computer Vision (CV), with the aim of solving problems that involve both text and images (Bai et al., 2018). Within this context, structured knowledge extraction from unstructured text is an open problem in the e-commerce literature (Arroyo et al., 2019). Regardless the source of the information (*e.g.* product websites, product images captured from stores or digital leaflets), it refers to a unified concept that can be denoted as “automated product coding”, *i.e.* the extraction of attribute values of e-commerce products (see Fig. 1).

Our present work focuses on the case of knowledge extraction from digital leaflets. Most retailers are replacing physical leaflets, that are directly collected from the stores, with digital leaflets that are uploaded in the cloud on the retailers websites. Comparing to real-world e-commerce platforms that contain billions of products with detailed descriptions (including ratings and opinions), digital leaflets include concise textual and visual information of promotions that applies to some of the products of the store assortment for a short period of time and thus it has to be updated regularly. The knowledge extraction of digital leaflets is of great interest for the e-commerce business as it impacts not only on several aspects of the consumers behaviour or seasonality, but also modifies relevant attributes of the products periodically, *e.g.* price or volume.

In this paper, we bring for the first time the problem of automated product coding in digital leaflets for e-commerce. In particular, we present an approach to predict the product categories for each of the promotions within the leaflet that serves as a strong baseline for future works. Technically, this is

This work is licensed under a Creative Commons Attribution 4.0 International License. Details:
<http://creativecommons.org/licenses/by/4.0/>.

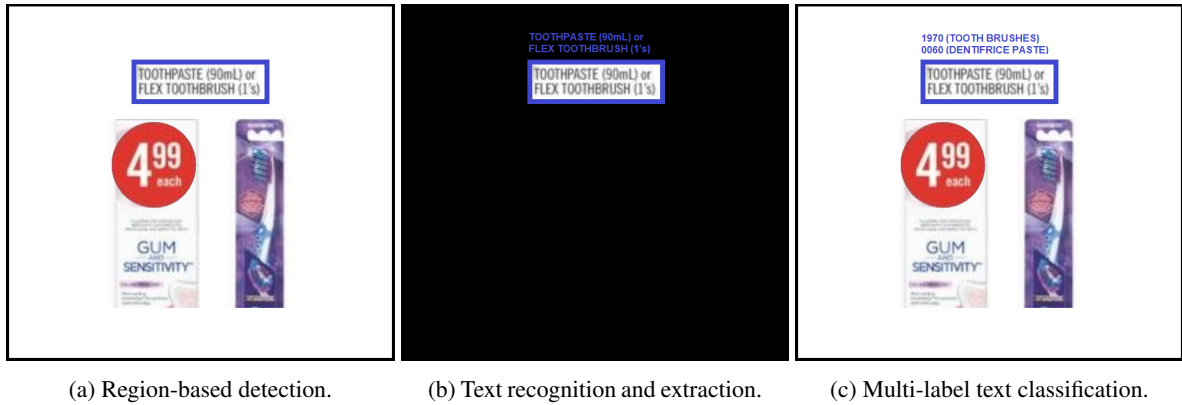


Figure 1: Key components of our approach for a single promotion within a digital leaflet.

a multi-label text classification problem as some promotions can potentially apply to several product categories. For that purpose, we hypothesize that most of the information of the promotions is self-contained in the products descriptions, so we first detect all these regions within the image that contain textual descriptions and extract the text using Optical Character Recognition (OCR) techniques. In Fig. 1, we show a visual representation of the three key components in the proposed approach: 1) region-based detection of the promotions description, 2) text recognition and extraction, and 3) multi-label text classification.

The building blocks depicted in Fig. 1 also relate to the different domains covered in the approach and they can be divided into the following three categories: 1) **CV**: a region detection architecture based on deep learning and image processing to detect the descriptions of each individual promotion within the digital leaflet, 2) **CV+NLP**: a text recognition method for extracting the textual information contained into the detected descriptions based on OCR and 3) **NLP**: a multi-label text classification model based on sub-word text embeddings and a shallow neural network.

The main contributions of the paper are three fold:

1. We bring for the first time the problem of automated item coding in digital leaflets for e-commerce platforms.
2. We formulate the process for predicting the categorization of each individual promotion within digital leaflets as a multi-label classification problem, which uses both CV and NLP techniques for properly fusing image and text information.
3. We conduct several experiments to assess the performance of the model for several aspects: a) detection of the product description region in promotions, b) multi-label classification of the product categories and c) multi-lingual capabilities.

The contents of the paper are structured as follows: related works are described in Section 2. The technical proposal presented in this paper is described in Section 3. The data used in the evaluations of our proposal, the experiments carried out to validate it and several comparative results are reviewed in Section 4. The final conclusions derived from this paper are discussed in Section 5.

2 Related Work

Prior works related to the fusion of NLP and CV have experienced an important growth in the last years due to the advances in deep learning and its influence in both domains. It covers several fields such as text retrieval (Gomez et al., 2018), image detection and classification (Bai et al., 2018) or automated item coding (Arroyo et al., 2019).

Regarding the first step of the system proposed in this paper, approaches based on region detection over images have received an incredible attention in the last few years. The popularization of deep learning

jointly with Convolutional Neural Networks (CNN) (Krizhevsky et al., 2012) has completely changed the traditional paradigm in CV. Standard CNNs are commonly applied only for image classification. However, R-CNNs (R stands for Region-based) are focused on object detection, which combines both detection and classification. Nowadays, techniques such as Faster R-CNN (Ren et al., 2015) are broadly extended to localize and classify objects over images. In this method, the regions in the R-CNN are detected by a selective search algorithm based on a Region Proposal Network (RPN). YOLO (Redmon et al., 2016) is also a technique very popularized for object detection which is focused on Single Shot Detection (SSD) (Liu et al., 2016). Similar proposals based on R-CNN architectures can be used in our approach to initially detect the regions where the text of the leaflets descriptions is located over the images.

Following with the second stage for the recognition of the texts contained in descriptions regions, Optical Character recognition (OCR) is a broad topic covered in the AI community and it aims at extracting text from images, thus working in the intersection between CV and NLP. The most recent approaches use deep learning techniques (Lee and Osindero, 2016) to examine images pixel by pixel, looking for shapes that match the character traits. Available OCR engines comprise implemented solutions that are open-source and proprietary. Calamari (Wick et al., 2020) or Tesseract (Zacharias et al., 2020) are some of the most effective open-source approaches, with lots of users around the world. However, proprietary solutions such as Google OCR¹ are currently obtaining better results in text recognition, including support for a larger number of languages. Our goal is to apply OCR-based algorithms over the regions previously detected using a R-CNN architecture in order to obtain the product descriptions over the images of leaflets.

In the final stage of our described approach, the textual information extracted from the detected regions is classified into their corresponding product categories. The state of the art in short text classification is recently moving to approaches based on DNNs (Deep Neural Networks). On the one hand, in (Joulin et al., 2017) the authors proposed to incorporate sub-word level information to train textual embeddings very efficiently for text classification. On the other hand, the BERT architecture described in (Devlin et al., 2018) also supposed a great milestone in natural language modeling introducing a self-supervised learning strategy that is able to incorporate an architecture based on Transformers (Vaswani et al., 2017) leveraging a large corpus for training. The embeddings obtained using BERT approaches highly correlate with the linguistic context within a sentence. Thus, proposals based on sub-word level information are very competitive compared to BERT models on those cases where we have unstructured textual information and probably OCR errors, such as the descriptions processed in most of the leaflets. It must be also noted that although BERT is focused on standard text processing, there are derived approaches that are also diving into text processing associated with images, such as ViLBERT (Lu et al., 2019) or VL-BERT (Su et al., 2020). The difference is that these recent approaches are not directly applied to classification over text contained in images, they are used for tasks that involve images and related external text, such as VQA (Visual Question Answering) (Antol et al., 2015).

3 Our Proposal for Digital Leaflets Categorization

The leaflets categorization proposal presented in this paper is focused on the prediction of multiple product categories from images such as the depicted in Fig. 1, which is showing a part of a catalog representing specific products. The solution can be divided into the following three main parts:

1. Detection of the regions related to the textual descriptions of each product in a promotion.
2. Recognition of the associated text inside the regions of the detected descriptions.
3. Classification of the recognized text into the different product categories of interest.

In this section, we introduce these three main components of the proposed approach, jointly with the whole leaflets categorization pipeline that combines them to obtain the final output.

¹<https://cloud.google.com/vision/docs/ocr>

3.1 Region-based Detection

The method designed for detecting the regions that contain the texts associated with product descriptions is based on an R-CNN schema, as presented in Fig 2. We decided to use this CV approach because the texts of product descriptions over the images have a specific appearance format that can be effectively visually differentiated, even when several templates and styles are used for varied retailers. Then, the texts from descriptions can be effectively separated from the rest of the texts in the image.




Figure 2: Visual representation of the text region detection over images in our R-CNN architecture.

Initially, we tried to differentiate the product descriptions from the rest of the texts in the digital leaflets by only considering predictions with high confidence in the multi-label text classification proposed for the third part of the system (described in detail in Section 3.3). Unfortunately, a great amount of texts out of scope were classified with high confidences by the model, so we decided to design the approach based on the R-CNN architecture to obtain more accurate results in the overall process.

As can be seen in Fig. 2, our architecture applies an internal Region Proposal Network (RPN) (Ren et al., 2015) with the aim of detecting the positions of the regions of interest. First of all, the image is resized before feeding it into the backbone CNN. This resizing is important to have similar detection schemas independently of the different sizes that could have the input images. For every point in the output, the network has to learn whether a text description region is present in the image at its corresponding position and estimate its size. Several anchors over the input image are used for each location from the backbone network. These anchors indicate possible objects in various sizes and aspect ratios at this location. As the RPN walks through each pixel in the feature map, it has to validate whether these corresponding anchors spanning the input image contain regions of interest. Besides, it has to refine the coordinates of anchors to provide bounding boxes as proposed regions associated with the different text of products descriptions. In order to help with this process, Non-Maximum Suppression (NMS) (Rothe et al., 2014) is applied as follows:

1. Choose the bounding box that has the highest confidence score.
2. Compute its overlap with the rest of bounding boxes and remove the bounding boxes that overlap more than an Intersection over Union (IoU) (Rezatofighi et al., 2019) threshold.
3. Return to the first step and iterate until there are no more boxes with a lower confidence score than the chosen box.

In order to train the detection model, our architecture requires Ground-Truth (GT) information about bounding boxes from sample images, with the aim of training the network to localize the regions of interest. In standard R-CNN architectures, a part of the network is in charge of classifying the bounding boxes into several classes. However, in our schema this is not required, because we do not need to differentiate

the class of the bounding boxes detected, what we need is to classify the internal textual information in the stage of text classification that is detailed in Section 3.3. Then, the standard network part for classifying the obtained visual embeddings is not applied in our model, only the part for localization based on RPN previously explained.

3.2 Text Recognition and Extraction

A method based on OCR is used in this second stage in order to recognize the text associated with the previously detected product descriptions. We use Google OCR as basis of our text recognition pipeline. Besides, the goal of our work is not focused on contributing a new complete OCR engine, which is a research out of the scope of this paper that considers a full detection, recognition and classification schema for leaflets categorization.

In our case, OCR converts leaflets images into machine-readable text data. The human visual system reads text by recognizing the patterns of light and dark, translating those patterns into characters and words, and then attaching meaning to it. Similarly, OCR attempts to mimic our visual system by using neural networks.

The approach applied in this stage to compute OCR returns the characters, words and paragraphs obtained from images and their locations. Initially, we implemented the idea of directly clustering the words recognized inside a bounding box detected for a product description, with the aim of providing the whole text string associated with that specific product description. However, we observed that the resulting text string sometimes contained errors due to other out-of-scope texts around the text of interest that interfere with it. To minimize the impact of this recognition issue, we decided to apply a mask to blacken all the parts of the image that are not contained inside the bounding boxes detected in the previous stage by the region detection model. Then, the blackened regions do not interfere with the regions of interest related to product descriptions during the OCR computation. The described blackened process is exemplified in Fig. 1 (b).

Finally, the text extracted by the OCR is post-processed to reduce typical errors, such as the ones associated with strange symbols incorrectly detected, problems derived from lower and upper case letters or dictionary-based corrections.

3.3 Multi-label Text Classification

After recognizing the text corresponding to product descriptions in digital leaflets images, a text classification model is applied to predict the different product categories of interest. Each product can be associated with more than one category, so this use case can be considered as an instance multi-label classification problem.

The proposed text classification model is based on FastText (Joulin et al., 2017), as it efficiently scales in the number of categories to predict. The defined architecture is a simple neural network that contains only one layer. The architecture generates a bag-of-words representation of the text, where the embeddings are fetched for every single word. After that, the embeddings are averaged to obtain a single embedding for the whole text in the hidden layer. Once the averaged embeddings are computed, the single vector is fed to independent binary classifiers for each label (one-vs-all loss). Character n-grams are used, which are really beneficial for text classification problems based on product descriptions (not natural language as it is usually known) and that may also include typos from the OCR. In order to visually understand the architecture and n-grams computation, Fig. 3 is presented.

For training the text classification model, a dataset with manually labeled annotations about the categories associated with each text description is required, as explained in detail in Section 4.1. The trained model is used to perform the inference of the categories related to each promotion description. The inference output gives a vector with probabilities for each available category. A threshold is used to filter the categories corresponding to an instance based on the obtained probabilities, with the aim of providing the multi-label classification. The categories with probabilities above this threshold are considered as positive.




Figure 3: FastText architecture used for our multi-label text classification with n-gram features $x_1, x_2, x_3, \dots, x_{n-2}, x_{n-1}, x_n$.

3.4 Whole Leaflets Categorization Pipeline

The diagram presented in Fig. 4 illustrates the overall setup of the solution for the whole leaflets categorization pipeline. An image corresponding to a digital leaflet is received as the input of the system. In the first step, the descriptions related to products are initially identified over the leaflet image using the previously trained region detection model. The detected regions are used for generating the masked image that is used in the text recognition stage to extract the texts of interest for each promotion. Finally, the model for text classification is applied for each description in order to compute the final output, which contains the categories associated for each promotion.

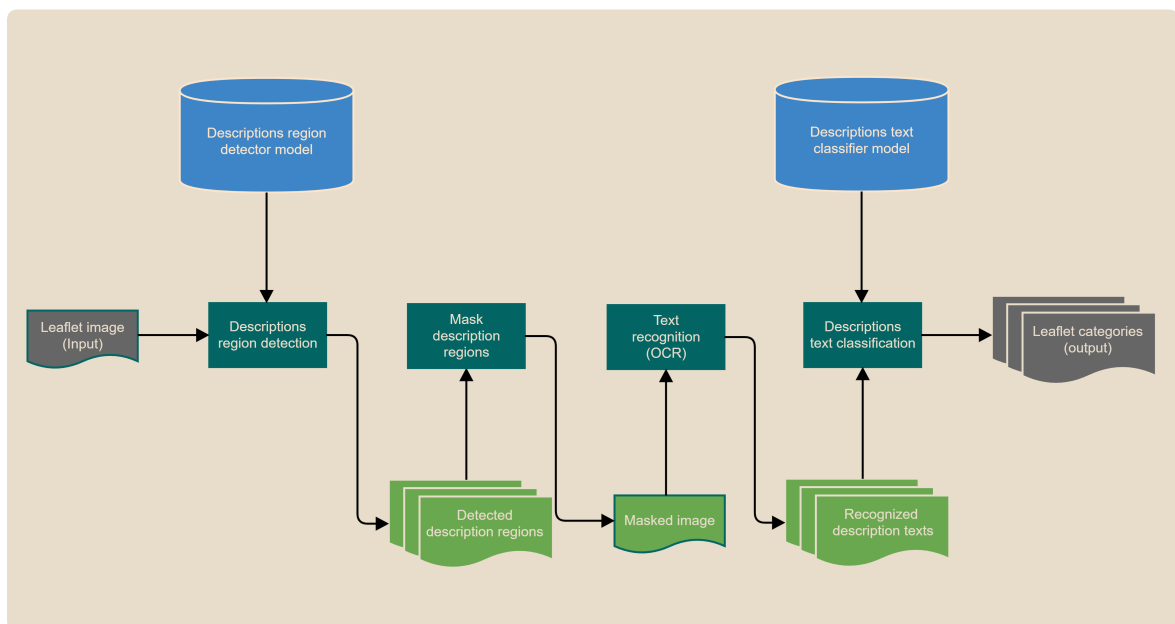


Figure 4: Diagram of the whole leaflets categorization pipeline.

4 Experiments and Results

With the aim of validating the proposed approach in a specific leaflets categorization use case, we prepared a set of experiments with leaflets data captured by Nielsen. In this section, we describe the experimental setup: the datasets, the hyperparameters used for training and the comparative results with different data and approaches.

4.1 Leaflets Dataset

To the best of our knowledge, there are not public datasets with GT information for leaflets categorization over images of catalogs, so we used our own labeled datasets from Nielsen internal data. We applied two different leaflets datasets for training and evaluation. Firstly, a "base" dataset with leaflets from only one retailer with textual descriptions in English. Secondly, an extended dataset with leaflets from four retailers with varied image formats to test generalization, and texts from two languages (English and French) to evaluate the multi-lingual capabilities of our approach. As the datasets are composed of proprietary images, we can not publicly share them. However, the main properties and statistics from both datasets are summarized in Table 1. It must be noted that the data distribution is long-tailed and thus unbalanced for both training and validation/test splits, as shown in Fig. 5. Then, this is an extra challenge for our models in order to be robust against the typical problems derived from long-tail datasets.

Dataset	#Languages	#Retailers	#Images	#Samples	#Categories	Avg. samples per cat.	Std. samples per cat.
Base	1	1	449	10,333	382	27.05	100.75
Extended	2	4	1,079	20,646	504	40.96	189.79

Table 1: Main statistics about the leaflets categorization datasets used for training and evaluation.

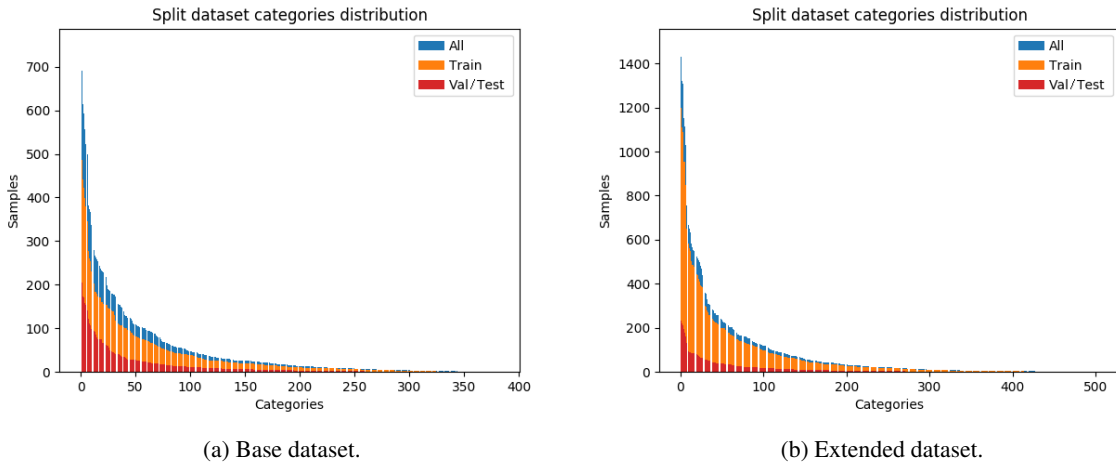


Figure 5: Distribution of category samples for the used leaflets datasets.

4.2 Hyperparameters Tuning

Region detection and text classification models require some hyperparameters tuning to obtain the best possible results. Some standard hyperparameters typically used in these models are configured.

In region detection, we decided to train our models using pre-trained weights on ImageNet (Deng et al., 2009) and based on a ResNet-101 backbone (He et al., 2016). Anchor scales and ratios for RPN are important hyperparameters, which are configured as $[2, 4, 8]$ and $[0.5, 1, 2]$, respectively. Learning rate was set to $1 \cdot 10^{-6}$ and regularization is applied by means of dropout (keep probabilities mode), which is set to 0.7. Besides, an Adam optimizer (Kingma and Ba, 2015) is used. A confidence threshold of 0.4 is applied to discard bounding boxes with low confidences. Trainings are iterated during 100 epochs.

The multi-label text classification hyperparameters configuration has a great dependence on the number of n-grams, which is a value finally set to 3 based on previous cross-validation experiments. Besides, the learning rate is set to 0.1 with a learning update rate of 100. A confidence threshold of 0.25 is applied to identify the categories of interest for a specific product description. We train for 30 epochs.

4.3 Results in Leaflets Categorization

In order to evaluate the performance of our approach, we apply metrics based on precision, recall and accuracy. Besides, we also use these metrics to obtain comparative results with respect to a standard baseline text classification, which is based on directly extracting OCR paragraphs on the wild from the image, without previously using a RPN detector for filtering texts related to product descriptions. In this baseline case, the texts recognized by the OCR with all the class probabilities below the text classifier confidence threshold are not considered as descriptions. Overall test results comparison is presented in Table 2 for the base dataset. As can be seen in these results, our approach yields an accuracy improvement of 24 points with respect to the standard baseline. These results confirm the enhancement given by our system with respect to the proposed baseline.

Method	Precision	Recall	Accuracy
Baseline (OCR on the wild + text classification)	0.64	0.66	0.48
Ours (RPN + OCR masked + text classification)	0.86	0.81	0.72

Table 2: Overall test results comparing a standard baseline approach vs ours in the base leaflets dataset.

It must be remarked that a confidence threshold of 0.25 is used for the multi-label text classification model of our proposal. This confidence represents the probabilities of having a correct prediction for a class, so the confidence threshold is used to filter predictions with low probabilities. The specific confidence threshold value is obtained by maximizing the accuracy values over threshold iterations from 0.00 to 1.00, as can be seen in the graph presented in Fig. 6 (b). To make fair comparisons, we also set the confidence threshold for the maximized accuracy value of the standard baseline approach, which is 0.40. The threshold iteration graph for the baseline method is shown in Fig. 6 (a).

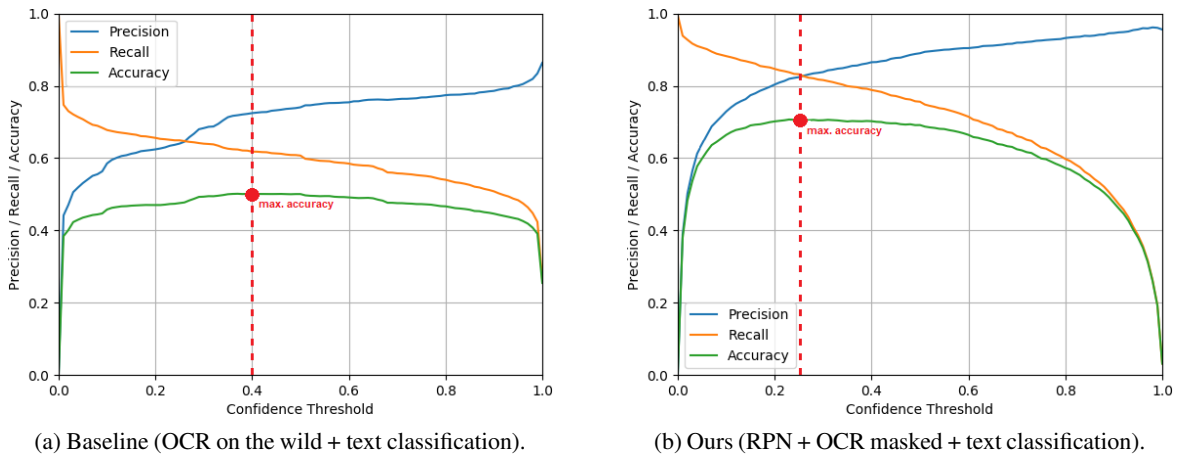


Figure 6: Graphs about sliding confidence threshold in text classification for the base leaflets dataset.

As a final insight about results, we trained our models using the extended dataset to check out how it generalizes to more retailers and languages. The obtained results can be seen in Table 3, where the models trained in the extended dataset are having a better performance in test. Moreover, in Fig. 7 we depict some qualitative results about some leaflets and their corresponding predictions. According to these results, it seems that the embeddings for the text classifier are able to generalize categorization to new languages. The reported accuracies must be understood taking into account the long-tail problems of the dataset exposed in Fig. 5, so the classes with less training samples are more difficult to predict.

Dataset	Precision	Recall	Accuracy
Base	0.86	0.81	0.72
Extended	0.87	0.86	0.76

Table 3: Comparison of results for our method in the base and extended leaflets datasets.



Figure 7: Qualitative results about leaflets examples and their corresponding predictions.

5 Conclusions

Along this paper, we have presented for the first time in the e-commerce research community (to the best of our knowledge) the problem of automated product coding for digital leaflets. In particular, we have addressed the problem of product classification for each promotion using image detection and multi-label text classification techniques. This schema provides a final proposal in the intersection between CV and NLP domains. Experimental results show that the described approach consistently outperforms a standard baseline in all the evaluated scenarios.

Future research includes expanding the multi-label classification of each promotion to knowledge extraction of different attributes, such as brand and product names, quantities, volumes, price or discounts. The final goal of this research line is to extract all the possible information contained in digital leaflets in order to fully understand their whole context.

We believe that the automated product coding in digital leaflets is at an early research stage but yet it is a very interesting approach in the future of e-commerce. Then, this paper has contributed the initial milestones for the dissemination and enhancement of this research topic across the e-commerce research community.

References

- S. Antol, A. Agrawal, J. Lu, M. Mitchell, D. Batra, C. L. Zitnick, and D. Parikh. 2015. VQA: Visual Question Answering. In *International Conference on Computer Vision (ICCV)*, pages 2425–2433.
- R. Arroyo, J. Tovar, F. J. Delgado, E. J. Almazan, D. G. Serrador, and A. Hurtado. 2019. Integration of Text-Maps in CNNs for Region Detection among Different Textual Categories. In *Workshops of Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–4.
- X. Bai, M. Yang, P. Lyu, Y. Xu, and Jiebo Luo. 2018. Integrating Scene Text and Visual Appearance for Fine-Grained Image Classification. *IEEE Access*, 6:66322–66335.
- J. Deng, W. Dong, R. Socher, L. Li, K. Li, and F. Li. 2009. ImageNet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255.
- J. Devlin, M. Chang, K. Lee, and K. Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 4171–4186.
- L. Gomez, A. Mafla, M. Rusinol, and D. Karatzas. 2018. Single Shot Scene Text Retrieval. In *European Conference on Computer Vision (ECCV)*, pages 728–744.
- K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov. 2017. Bag of Tricks for Efficient Text Classification. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 427–431.
- D. P. Kingma and J. Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, pages 1–15.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 1106–1114.
- C. Lee and S. Osindero. 2016. Recursive Recurrent Nets with Attention Modeling for OCR in the Wild. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2231–2239.
- W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. E. Reed, C. Fu, and A. C. Berg. 2016. SSD: Single Shot MultiBox Detector. In *European Conference on Computer Vision (ECCV)*, pages 21–37.
- J. Lu, D. Batra, D. Parikh, and S. Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 13–23.
- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. 2016. You Only Look Once: Unified, Real-Time Object Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788.
- S. Ren, K. He, R. Girshick, and J. Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 91–99.
- H. Rezatofighi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, and S. Savarese. 2019. Generalized Intersection over Union: A Metric and A Loss for Bounding Box Regression. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 658–666.
- R. Rothe, M. Guillaumin, and L. van Gool. 2014. Non-maximum Suppression for Object Detection by Passing Messages Between Windows. In *Asian Conference on Computer Vision (ACCV)*, pages 290–306.
- W. Su, X. Zhu, Y. Cao, Li B, L. Lu, F. Wei, and J. Dai. 2020. VL-BERT: Pre-training of Generic Visual Linguistic Representations. In *International Conference on Learning Representations (ICLR)*, pages 1–16.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is All you Need. In *International Conference on Neural Information Processing Systems (NIPS)*, pages 5998–6008.
- C. Wick, C. Reul, and F. Puppe. 2020. Calamari - A High-Performance Tensorflow-based Deep Learning Package for Optical Character Recognition. *Digital Humanities Quarterly*.
- E. Zacharias, M. Teuchler, and B. Bernier. 2020. Image Processing Based Scene-Text Detection and Recognition with Tesseract. *arXiv (CoRR)*.