

# MOROCO: The Moldavian and Romanian Dialectal Corpus

Andrei M. Butnaru and Radu Tudor Ionescu

Department of Computer Science, University of Bucharest

14 Academiei, Bucharest, Romania

butnaruandreimadalin@gmail.com

raducu.ionescu@gmail.com

## Abstract

In this work, we introduce the **Moldavian and Romanian Dialectal Corpus (MOROCO)**, which is freely available for download at <https://github.com/butnaruandrei/MOROCO>. The corpus contains 33564 samples of text (with over 10 million tokens) collected from the news domain. The samples belong to one of the following six topics: culture, finance, politics, science, sports and tech. The data set is divided into 21719 samples for training, 5921 samples for validation and another 5924 samples for testing. For each sample, we provide corresponding dialectal and category labels. This allows us to perform empirical studies on several classification tasks such as (i) binary discrimination of Moldavian versus Romanian text samples, (ii) intra-dialect multi-class categorization by topic and (iii) cross-dialect multi-class categorization by topic. We perform experiments using a shallow approach based on string kernels, as well as a novel deep approach based on character-level convolutional neural networks containing Squeeze-and-Excitation blocks. We also present and analyze the most discriminative features of our best performing model, before and after named entity removal.

## 1 Introduction

The high number of evaluation campaigns on spoken or written dialect identification conducted in recent years (Ali et al., 2017; Malmasi et al., 2016; Rangel et al., 2017; Zampieri et al., 2017, 2018) prove that dialect identification is an interesting and challenging natural language processing (NLP) task, actively studied by researchers in nowadays. Due to the recent interest in dialect identification, we introduce the **Moldavian and Romanian Dialectal Corpus (MOROCO)**, which is composed of 33564 samples of text collected from the news domain.



Figure 1: Map of Romania and the Republic of Moldova.

Romanian is part of the Balkan-Romance group that evolved from several dialects of Vulgar Latin, which separated from the Western Romance branch of languages from the fifth century (Coteanu et al., 1969). In order to distinguish Romanian within the Balkan-Romance group in comparative linguistics, it is referred to as *Daco-Romanian*. Along with Daco-Romanian, which is currently spoken in Romania, there are three other dialects in the Balkan-Romance branch, namely Aromanian, Istro-Romanian, and Megleno-Romanian. Moldavian is a subdialect of Daco-Romanian, that is spoken in the Republic of Moldova and in northeastern Romania. The delimitation of the Moldavian dialect, as with all other Romanian dialects, is made primarily by analyzing its phonetic features and only marginally by morphological, syntactical, and lexical characteristics. Although the spoken dialects in Romania and the Republic of Moldova are different, the two countries share the same literary standard (Minahan, 2013). Some linguists (Pavel, 2008) consider that the border between Romania and the Republic of Moldova (see Figure 1) does not correspond to any significant isoglosses to justify a dialectal division. One question that arises in this context

is whether we can train a machine to accurately distinguish literary text samples written by people in Romania from literary text samples written by people in the Republic of Moldova. If we can construct such a machine, then what are the discriminative features employed by this machine? Our corpus formed of text samples collected from Romanian and Moldavian news websites, enables us to answer these questions. Furthermore, MOROCO provides a benchmark for the evaluation of dialect identification methods. To this end, we consider two state-of-the-art methods, string kernels (Butnaru and Ionescu, 2018; Ionescu and Butnaru, 2017; Ionescu et al., 2014) and character-level convolutional neural networks (CNNs) (Ali, 2018; Belinkov and Glass, 2016; Zhang et al., 2015), which obtained the first two places (Ali, 2018; Butnaru and Ionescu, 2018) in the Arabic Dialect Identification Shared Task of the 2018 VarDial Evaluation Campaign (Zampieri et al., 2018). We also experiment with a novel CNN architecture inspired the recently introduced Squeeze-and-Excitation (SE) networks (Hu et al., 2018), which exhibit state-of-the-art performance in object recognition from images. To our knowledge, we are the first to introduce Squeeze-and-Excitation networks in the text domain.

As we provide category labels for the collected text samples, we can perform additional experiments on various text categorization by topic tasks. One type of task is intra-dialect multi-class categorization by topic, i.e. the task is to classify the samples written either in the Moldavian dialect or in the Romanian dialect into one of the following six topics: culture, finance, politics, science, sports and tech. Another type of task is cross-dialect multi-class categorization by topic, i.e. the task is to classify the samples written in one dialect, e.g. Romanian, into six topics, using a model trained on samples written in the other dialect, e.g. Moldavian. These experiments are aimed at showing if the considered text categorization methods are robust to the dialect shift between training and testing.

In summary, our contribution is threefold:

- We introduce a novel large corpus containing 33564 text samples written in the Moldavian and the Romanian dialects.
- We introduce Squeeze-and-Excitation networks to the text domain.
- We analyze the discriminative features that help the best performing method, string ker-

nels, in (i) distinguishing the Moldavian and the Romanian dialects and in (ii) categorizing the text samples by topic.

We organize the remainder of this paper as follows. We discuss related work in Section 2. We describe the MOROCO data set in Section 3. We present the chosen classification methods in Section 4. We show empirical results in Section 5, and we provide a discussion on the discriminative features in Section 6. Finally, we draw our conclusion in Section 7.

## 2 Related Work

There are several corpora available for dialect identification (Ali et al., 2016; Alsarsour et al., 2018; Bouamor et al., 2018; Francom et al., 2014; Johannessen et al., 2009; Kumar et al., 2018; Samardžić et al., 2016; Tan et al., 2014; Zaidan and Callison-Burch, 2011). Most of these corpora have been proposed for languages that are widely spread across the globe, e.g. Arabic (Ali et al., 2016; Alsarsour et al., 2018; Bouamor et al., 2018), Spanish (Francom et al., 2014), Indian (Kumar et al., 2018) or German (Samardžić et al., 2016). Among these, Arabic is the most popular, with a number of four data sets (Ali et al., 2016; Alsarsour et al., 2018; Bouamor et al., 2018; Zaidan and Callison-Burch, 2011), if not even more.

**Arabic.** The Arabic Online news Commentary (AOC) (Zaidan and Callison-Burch, 2011) is the first available dialectal Arabic data set. Although AOC contains 3.1 million comments gathered from Egyptian, Gulf and Levantine news websites, the authors labeled only around 0.05% of the data set through the Amazon Mechanical Turk crowdsourcing platform. Ali et al. (2016) constructed a data set of audio recordings, Automatic Speech Recognition transcripts and phonetic transcripts of Arabic speech collected from the Broadcast News domain. The data set was used in the 2016, 2017 and 2018 VarDial Evaluation Campaigns (Malmasi et al., 2016; Zampieri et al., 2017, 2018). Alsarsour et al. (2018) collected the Dialectal ARabic Tweets (DART) data set, which contains around 25K manually-annotated tweets. The data set is well-balanced over five main groups of Arabic dialects: Egyptian, Maghrebi, Levantine, Gulf and Iraqi. Bouamor et al. (2018) presented a large parallel corpus of 25 Arabic city dialects, which was created by translating selected sentences from the travel domain.

**Other languages.** The Nordic Dialect Corpus (Johannessen et al., 2009) contains about 466K spoken words from Denmark, Faroe Islands, Iceland, Norway and Sweden. The authors transcribed each dialect by the standard official orthography of the corresponding country. Francom et al. (2014) introduced the ACTIV-ES corpus, which represents a cross-dialectal record of the informal language use of Spanish speakers from Argentina, Mexico and Spain. The data set is composed of 430 TV or movie subtitle files. The DSL corpus collection (Tan et al., 2014) comprises news data from various corpora to emulate the diverse news content across different languages. The collection is comprised of six language variety groups. For each language, the collection contains 18K training sentences, 2K validation sentences and 1K test sentences. The ArchiMob corpus (Samardžić et al., 2016) contains manually-annotated transcripts of Swiss German speech collected from four different regions: Basel, Bern, Lucerne and Zurich. The data set was used in the 2017 and 2018 VarDial Evaluation Campaigns (Zampieri et al., 2017, 2018). Kumar et al. (2018) constructed a corpus of five Indian dialects consisting of 307K sentences. The samples were collected by scanning, passing through an OCR engine and proofreading printed stories, novels and essays from books, magazines or newspapers.

**Romanian.** To our knowledge, the only empirical study on Romanian dialect identification was conducted by Ciobanu and Dinu (2016). In their work, Ciobanu and Dinu (2016) used only a short list of 108 parallel words in a binary classification task in order to discriminate between Daco-Romanian words versus Aromanian, Istro-Romanian and Megleno-Romanian words. Different from Ciobanu and Dinu (2016), we conduct a large scale study on 33K documents that contain a total of about 10 million tokens.

### 3 MOROCO

In order to build MOROCO, we collected text samples from the top five most popular news websites in Romania and the Republic of Moldova, respectively. Since news websites in the two countries belong to different Internet domains, the text samples can be automatically labeled with the corresponding dialect. We selected news from six different topics, for which we found at least 2000 text samples in both dialects. For each dialect, we illustrate the distribution of text samples per topic




Figure 2: The distribution of text samples per topic for the Moldavian and the Romanian dialects, respectively. Best viewed in color.

in Figure 2. In both countries, we notice that the most popular topics are finance and politics, while the least popular topics are culture and science. The distribution of topics for the two dialects is mostly similar, but not very well-balanced. For instance, the number of Moldavian politics samples (5154) is about six times higher than the number of Moldavian science samples (877). However, MOROCO is well-balanced when it comes to the distribution of samples per dialect, since we were able to collect 15403 Moldavian text samples and 18161 Romanian text samples.

It is important to note that, in order to obtain the text samples, we removed all HTML tags and replaced consecutive space characters with a single space character. We further processed the samples in order to eliminate named entities. Previous research (Abu-Jbara et al., 2013; Nicolai and Kondrak, 2014) found that named entities such as country names or cities can provide clues about the native language of English learners. We decided to remove named entities in order to prevent classifiers from taking the decision based on features that are not truly indicative of the dialects or the topics. For example, named entities representing city names in Romania or Moldova can provide clues about the dialect, while named entities representing politicians or football players names can provide clues about the topic. The identified named entities are replaced with the token  $\$NE\$$ . In the experiments, we present results before and after named entity removal, in order to illustrate the effect of named entities.

In order to allow proper comparison in future research, we divided MOROCO into a training, a validation and a test set. We used stratified sampling in order to produce a split that preserves the distribution of dialects and topics across all subsets. Table 1 shows some statistics of the number

Set	#samples	#tokens
Training	21,719	6,705,334
Validation	5,921	1,826,818
Test	5,924	1,850,977
Total	33,564	10,383,129

Table 1: The number of samples (#samples) and the number of tokens (#tokens) contained in the training, validation and test sets included in our corpus.

of samples as well as the number of tokens in each subset. We note that the entire corpus contains 33564 samples with more than 10 million tokens in total. On average, there are about 309 tokens per sample.

Since we provide both dialectal and category labels for each sample, we can perform several tasks on MOROCO:

- Binary classification by dialect – the task is to discriminate between the Moldavian and the Romanian dialects.
- Moldavian (MD) intra-dialect multi-class categorization by topic – the task is to classify the samples written in the Moldavian dialect into six topics.
- Romanian (RO) intra-dialect multi-class categorization by topic – the task is to classify the samples written in the Romanian dialect into six topics.
- MD→RO cross-dialect multi-class categorization by topic – the task is to classify the samples written in the Romanian dialect into six topics, using a model trained on samples written in the Moldavian dialect.
- RO→MD cross-dialect multi-class categorization by topic – the task is to classify the samples written in the Moldavian dialect into six topics, using a model trained on samples written in the Romanian dialect.

## 4 Methods

**String kernels.** Kernel functions (Shawe-Taylor and Cristianini, 2004) capture the intuitive notion of similarity between objects in a specific domain. For example, in text mining, string kernels can be used to measure the pairwise similarity between text samples, simply based on character n-grams. Various string kernel functions have been proposed to date (Ionescu et al., 2014; Lodhi et al., 2002; Shawe-Taylor and Cristianini, 2004). Recently, the presence bits string kernel and the histogram intersection kernel obtained state-of-the-art results in a broad range of

text classification tasks such as dialect identification (Butnaru and Ionescu, 2018; Ionescu and Butnaru, 2017; Ionescu and Popescu, 2016), native language identification (Ionescu et al., 2016; Ionescu and Popescu, 2017), sentiment analysis (Giménez-Pérez et al., 2017; Ionescu and Butnaru, 2018; Popescu et al., 2017) and automatic essay scoring (Cozma et al., 2018). In this paper, we opt for the presence bits string kernel, which allows us to derive the primal weights and analyze the most discriminative features, as explained by Ionescu et al. (2016). For two strings over an alphabet  $\Sigma$ ,  $x, y \in \Sigma^*$ , the presence bits string kernel is formally defined as:

$$k_n^{0/1}(x, y) = \sum_{s \in \Sigma^n} \text{in}_s(x) \cdot \text{in}_s(y),$$

where  $\text{in}_s(x)$  is 1 if string  $s$  occurs as a substring in  $x$ , and 0 otherwise. In our empirical study, we experiment with character n-grams in a range, and employ the Kernel Ridge Regression (KRR) binary classifier. During training, KRR finds the vector of weights that has both small empirical error and small norm in the Reproducing Kernel Hilbert Space generated by the kernel function. The ratio between the empirical error and the norm of the weight vector is controlled through the regularization parameter  $\lambda$ .

**Character-level CNN.** Convolutional networks (LeCun et al., 1998; Krizhevsky et al., 2012) have been employed for solving many NLP tasks such as part-of-speech tagging (Santos and Zadrozny, 2014), text categorization (Johnson and Zhang, 2015; Kim, 2014; Zhang et al., 2015), dialect identification (Ali, 2018; Belinkov and Glass, 2016), machine translation (Gehring et al., 2017) and language modeling (Dauphin et al., 2017; Kim et al., 2016). Many CNN-based methods rely on words, the primary reason for this being the aid given by word embeddings (Mikolov et al., 2013; Pennington et al., 2014) and their ability to learn semantic and syntactic latent features. Trying to eliminate the pre-trained word embeddings from the pipeline, some researchers have decided to build end-to-end models using characters as input, in order to solve text classification (Zhang et al., 2015; Belinkov and Glass, 2016) or language modeling tasks (Kim et al., 2016). At the character-level, the model can learn unusual character sequences such as misspellings or take advantage of unseen words during test time. This appears to be particularly helpful in dialect identification, since some state-




Figure 3: **Left:** The architecture of the baseline character-level CNN composed of seven blocks. **Right:** The modified CNN architecture, which includes a Squeeze-and-Excitation block after each convolutional block.

of-the-art dialect identification methods (Butnaru and Ionescu, 2018; Ionescu and Butnaru, 2017) use character n-grams as features.

In this paper, we draw our inspiration from Zhang et al. (2015) in order to design a lightweight character-level CNN architecture for dialect identification. One way proposed by Zhang et al. (2015) to represent characters in a character-level CNN is to map every character from an alphabet of size  $t$  to a discrete value using a 1-of- $t$  encoding. For example, having the alphabet  $\Sigma = \{a, b, c\}$ , the encoding for the character  $a$  is 1, for  $b$  is 2, and for  $c$  is 3. Each character from the input text is encoded, and only a fixed size  $l$  of the input is kept. In our case, we keep the first  $l = 5000$  characters, zero-padding the documents that are under length. We compose an alphabet of 105 characters that includes uppercase and lowercase characters, Moldavian and Romanian diacritics (such as  $\tilde{a}$ ,  $\hat{a}$ ,  $\hat{i}$ ,  $\mathring{s}$  and  $\mathring{t}$ ), digits, and 33 other symbol characters. Characters that do not appear in the alphabet are encoded as a blank character.

As illustrated in the left-hand side of Figure

3, our architecture is seven blocks deep, containing one embedding layer, three convolutional and max-pooling blocks, and three fully-connected blocks. The first two convolutional layers are based on one-dimensional filters of size 7, the third one being based on one-dimensional filters of size 3. A thresholded Rectified Linear Units (ReLU) activation function (Nair and Hinton, 2010) follows each convolutional layer. The max-pooling layers are based on one-dimensional filters of size 3 with stride 3. After the third convolutional block, the activation maps pass through two fully-connected blocks having thresholded ReLU activations. Each of these two fully-connected blocks is followed by a dropout layer with the dropout rate of 0.5. The last fully-connected layer is followed by softmax, which provides the final output. All convolutional layers have 128 filters, and the threshold used for the thresholded ReLU is  $10^{-6}$ . The network is trained with the Adam optimizer (Kingma and Ba, 2015) using categorical cross-entropy as loss function.

**Squeeze-and-Excitation Networks.** Hu et al. (2018) argued that the convolutional filters close to the input layer are not aware of the global appearance of the objects in the input image, as they operate at the local level. To alleviate this problem, Hu et al. (2018) proposed to insert Squeeze-and-Excitation blocks after each convolutional block that is closer to the network’s input. The SE blocks are formed of two layers, *squeeze* and *excitation*. The activation maps of a given convolutional block are first passed through the squeeze layer, which aggregates the activation maps across the spatial dimension in order to produce a channel descriptor. This layer can be implemented through a global average pooling operation. In our case, the size of the output after the squeeze operation is  $1 \times 128$ , since our convolutional layers are one-dimensional and each layer contains  $d = 128$  filters. The resulting channel descriptor enables information from the global receptive field of the network to be leveraged by the layers near the network’s input. The squeeze layer is followed by an excitation layer based on a self-gating mechanism, which aims to capture channel-wise dependencies. The self-gating mechanism is implemented through two fully-connected layers, the first being followed by ReLU activations and the second being followed by sigmoid activations, respectively. The first fully-connected layer acts




Figure 4: Dialect classification results on the validation set provided by the KRR classifier based on the presence bits string kernel with n-grams in the range 5-8. Results are reported for various  $\lambda$  values from  $10^{-3}$  to  $10^{-6}$ . Best viewed in color.

as a bottleneck layer, reducing the input dimension (given by the number of filters  $d$ ) with a reduction ratio  $r$ . This is achieved by assigning  $d/r$  units to the bottleneck layer. The second fully-connected layer increases the size of the output back to  $1 \times 128$ . Finally, the activation maps of the preceding convolutional block are then reweighted (using the  $1 \times 128$  outputs provided by the excitation layer as weights) to generate the output of the SE block, which can then be fed directly into subsequent layers. Thus, SE blocks are just alternative pathways designed to recalibrate channel-wise feature responses by explicitly modeling interdependencies between channels. We insert SE blocks after each convolutional block, as illustrated in the right-hand side of Figure 3.

## 5 Experiments

**Parameter tuning.** In order to tune the parameters of each model, we used the MOROCO validation set. We first carried out a set of preliminary dialect classification experiments to determine the optimal choice of n-grams length for the presence bits string kernel and the regularization parameter  $\lambda$  of the KRR classifier. We present results for these preliminary experiments in Figure 4. We notice that both  $\lambda = 10^{-4}$  and  $\lambda = 10^{-5}$  are good regularization choices, with  $\lambda = 10^{-5}$  being slightly better for all n-grams lengths between 5 and 8. Although 6-grams, 7-grams and 8-grams attain almost equally good results, the best choice according to the validation results is to use 6-grams. Therefore, in the subsequent experiments, we employ the presence bits string kernel based on n-grams of length 6 and KRR with  $\lambda = 10^{-5}$ .

For the baseline CNN, we set the learning rate to  $5 \cdot 10^{-4}$  and use mini-batches of 128 samples

during training. We use the same parameters for the SE network. Both deep networks are trained for 50 epochs. For the SE blocks, we set the reduction ratio to  $r = 64$ , which results in a bottleneck layer with two neurons. We also tried lower reduction ratios, e.g. 32 and 16, but we obtained lower performance for these values.

**Results.** In Table 2 we present the accuracy, the weighted  $F_1$ -scores and the macro-averaged  $F_1$ -scores obtained by the three classification models (string kernels, CNN and SE networks) for all the classification tasks, on the validation set as well as the test set. Regarding the binary classification by dialect task, we notice that all models attain good results, above 90%. SE blocks bring only minor improvements over the baseline CNN. Our deep models, CNN and CNN+SE, attain results around 93%, while the string kernels obtain results above 94%. We thus conclude that written text samples from the Moldavian and the Romanian dialects can be accurately discriminated by both shallow and deep learning models. This answers our first question from Section 1.

Regarding the Moldavian intra-dialect 6-way categorization (by topic) task, we notice that string kernels perform quite well in comparison with the CNN and the CNN+SE models. In terms of the macro-averaged  $F_1$  scores, SE blocks bring improvements higher than 1% over the baseline CNN. In the MD $\rightarrow$ RO cross-dialect 6-way categorization task, our models attain the lowest performance on the Romanian test set. We would like to note that in both cross-dialect settings, we use the validation set from the same dialect as the training set, in order to prevent any use of information about the test dialect during training. In other words, the settings are intra-dialect with re-

Task	Method	Validation			Test		
		accuracy	weighted $F_1$	macro $F_1$	accuracy	weighted $F_1$	macro $F_1$
Binary classification by dialect	KRR + $k_6^{0/1}$	94.21	94.20	94.15	94.13	94.11	94.06
	CNN	93.00	93.00	92.95	92.75	92.76	92.71
	CNN + SE	93.02	93.01	92.95	92.99	92.98	92.93
MD categorization (by topic)	KRR + $k_6^{0/1}$	92.49	92.46	90.45	92.68	92.63	90.57
	CNN	85.42	85.25	79.28	83.92	83.73	76.82
	CNN + SE	86.23	85.97	80.51	84.39	84.01	77.85
MD→RO categorization (by topic)	KRR + $k_6^{0/1}$	92.49	92.46	90.45	68.21	67.59	67.59
	CNN	85.42	85.25	79.28	55.04	56.27	53.67
	CNN + SE	86.23	85.97	80.51	56.31	57.01	53.85
RO categorization (by topic)	KRR + $k_6^{0/1}$	74.38	74.49	77.68	74.98	75.11	78.76
	CNN	68.04	67.10	63.84	68.14	67.43	64.98
	CNN + SE	68.76	68.67	67.74	69.04	69.07	68.77
RO→MD categorization (by topic)	KRR + $k_6^{0/1}$	74.38	74.49	77.68	82.31	82.17	75.47
	CNN	68.04	67.10	63.84	72.49	71.60	62.70
	CNN + SE	68.76	68.67	67.74	74.84	74.87	67.42

Table 2: Accuracy rates, weighted  $F_1$  scores and macro-averaged  $F_1$ -scores (in %) for the five evaluation tasks: binary classification by dialect, Moldavian intra-dialect 6-way categorization (by topic), MD→RO cross-dialect 6-way categorization, Romanian (RO) intra-dialect 6-way categorization, and RO→MD cross-dialect 6-way categorization. Results are reported for three baseline models: KRR based on the presence bits string kernel (KRR+ $k_6^{0/1}$ ), convolutional neural networks (CNN), and Squeeze-and-Excitation convolutional neural networks (CNN+SE).

spect to the validation set and cross-dialect with respect to the test set. The Romanian intra-dialect 6-way categorization task seems to be much more difficult than the Moldavian intra-dialect categorization task, since all models obtain scores that are roughly 20% lower. In terms of the macro-averaged  $F_1$  scores, SE blocks bring improvements of around 4% over the baseline CNN. However, the results of CNN+SE are still much under those of the presence bits string kernel. Regarding the RO→MD cross-dialect 6-way categorization task, we find that the models learned on the Romanian training set obtain better results on the Moldavian (cross-dialect) test set than on the Romanian (intra-dialect) test set. Once again, this provides additional evidence that the 6-way categorization by topic task is more difficult for Romanian than for Moldavian. In all the intra-dialect or cross-dialect 6-way categorization tasks, we observe a high performance gap between deep and shallow models. These results are consistent with the recent reports of the VarDial evaluation campaigns (Malmasi et al., 2016; Zampieri et al., 2017, 2018), which point out that shallow approaches such as string kernels (Butnaru and Ionescu, 2018; Ionescu and Butnaru, 2017) surpass deep models in dialect and similar language discrimination tasks. Although deep models obtain generally lower results, our proposal of integrating Squeeze-and-Excitation blocks seems to be a steady step towards improving CNN models

Task	NER	Test		
		accuracy	weighted $F_1$	macro $F_1$
Classification by dialect	No	95.61	95.60	95.56
	Yes	94.13	94.11	94.06
MD categorization	No	93.23	93.19	91.36
	Yes	92.68	92.63	90.57
MD→RO categorization	No	68.80	68.23	68.49
	Yes	68.21	67.59	67.59
RO categorization	No	76.07	76.19	80.10
	Yes	74.98	75.11	78.76
RO→MD categorization	No	82.57	82.46	76.00
	Yes	82.31	82.17	75.47

Table 3: Accuracy rates, weighted  $F_1$  scores and macro-averaged  $F_1$ -scores (in %) of the KRR based on the presence bits string kernel for the five evaluation tasks, before and after named entity removal (NER).

for language identification, as SE blocks improve performance across all the experiments presented in Table 2, and, in some cases, the performance gains are considerable.

## 6 Discussion

**Named entity removal.** In Table 3, we presents comparative results before and after named entity removal (NER). We selected only the KRR based on the presence bits string kernel for this comparative study, since it provides the best performance among the considered baselines. The experiment reveals that named entities can artificially raise the performance by more than 1% in some cases, which is consistent with observations in previous works (Abu-Jbara et al., 2013; Nicolai and Kon-

NER	Top 6-grams for culture		Top 6-grams for finance		Top 6-grams for politics	
	original	translation	original	translation	original	translation
No	<b>[teatru]</b> <b>[scenă]</b> <b>[Eurovi]sion</b> <b>[scriit]or</b> <b>Euro[vision]</b>	theater scene Eurovision contest writer Eurovision contest	<b>[econom]ie</b> <b>[achita]t</b> <b>[tranza]cție</b> <b>di[n Mold]ova</b> <b>Un[iCredi]t</b>	economy payed transaction of Moldova UniCredit Bank	<b>[. PSD ]</b> <b>[parlam]ent</b> <b>Liviu D[ragnea]</b> <b>Igor[ Dodon]</b> <b>Dacian [Cioloș]</b>	Social-Democrat Party parliament ex-leader of PSD president of Moldova ex-prime minster of Romania
Yes	<b>[muzică]</b> <b>[piesă]</b> <b>[artist]</b> <b>[actoru]l</b> <b>s[pectac]ol</b>	music piece artist the actor show	<b>[ bănci]</b> <b>[monede]</b> <b>[afacer]i</b> <b>[export]uri</b> <b>p[roduse]</b>	banks currencies business exports products	<b>[politi]ca</b> <b>[președ]inte</b> <b>[primar]</b> <b>p[artidu]l</b> <b>[democr]ație</b>	the politics president mayor the party democracy
	Top 6-grams for science		Top 6-grams for sports		Top 6-grams for tech	
No	<b>[studiu]</b> <b>[știință]</b> <b>[ NASA ]</b> <b>Max [Planck]</b> <b>[Pământ]</b>	study science NASA Max Planck Earth	<b>[Simona] Halep</b> <b>[campio]n</b> <b>Simona[ Halep]</b> <b>o[limpic]</b> <b>[echipe]</b>	a tennis player champion a tennis player Olympic teams	<b>[Intern]et</b> <b>Fac[cebook]</b> <b>Mol[dtelec]om</b> <b>com[unicaț]ii</b> <b>[ telev]iziune</b>	Internet Facebook telecom operator in Moldova communications television
Yes	<b>[cercet]are</b> <b>[astron]omie</b> <b>[planet]a</b> <b>[univer]sitatea</b> <b>[teorie]</b>	research astronomy the planet the university theory	<b>[fotbal]</b> <b>[meciu]l</b> <b>[jucăto]r</b> <b>[antren]orul</b> <b>[clubul]</b>	football the match player the coach the club	<b>[mașini]</b> <b>[utiliz]ator</b> <b>t[elefon]</b> <b>[ compa]nie</b> <b>[tehnol]ogie</b>	cars user telephone company technology

Table 5: Examples of n-grams from the six different categories in MOROCO, that are weighted as more discriminative by the KRR based on the presence bits string kernel, before and after named entity removal (NER). The n-grams are placed between squared brackets and highlighted in bold. The n-grams are posed inside words and translated to English.

NER	Top 6-grams for MD		Top 6-grams for RO	
	original	translation	original	translation
No	<b>[Pământ]</b> <b>[Moldov]a</b> <b>[cîteva]</b> <b>M[oldova]</b> <b>cuv[întul ]</b>	Earth Moldova some Moldova the word	<b>[Români]a</b> <b>n[ews.ro]</b> <b>[Pământ]</b> <b>Nicu[lescu ]</b> <b>[Bucure]ști</b>	Romania a website Earth family name Bucharest
Yes	<b>[ sînt ]</b> <b>[ cînd ]</b> <b>[decît ]</b> <b>t[enisme]n</b> <b>[ pînă ]</b>	am / are when than tennis player until	<b>[ român]esc</b> <b>[ județ]</b> <b>[ cînd ]</b> <b>[ firme]</b> <b>[ vorbi]</b>	Romanian county when companies talk

Table 4: Examples of n-grams from the Moldavian and the Romanian dialects, that are weighted as more discriminative by the KRR based on the presence bits string kernel, before and after named entity removal (NER). The n-grams are placed between squared brackets and highlighted in bold. The n-grams are posed inside words and translated to English.

drak, 2014).

**Discriminative features.** In order to understand why the KRR based on the presence bits string kernel works so well in discriminating the Moldavian and the Romanian dialects, we conduct an analysis of some of the most discriminative features (n-grams), which are listed in Table 4. When named entities are left in place, the classifier chooses the country names (Moldova and Romania) or the capital city of Romania (Bucharest) as discriminative features. When named entities are removed, it

seems that Moldavian words that contain the letter ‘î’ inside, e.g. ‘cînd’, are discriminative, since in Romanian, the letter ‘î’ is only used at the beginning of a word (inside Romanian words, the same sound is denoted by ‘â’, e.g. ‘când’). While Moldavian writers prefer to use ‘tenismen’ to denote ‘tennis player’, Romanians prefer to use ‘jucător de tenis’ for the same concept. Although both terms, ‘tenismen’ and ‘jucător de tenis’, are understood in Romania and the Republic of Moldova, our analysis reveals that preference for one term or the other is not the same.

In a similar manner, we look at examples of features weighted as discriminative by the KRR based on the presence bits string kernel for categorization by topic. Table 5 lists discriminative n-grams for all the six categories inside MOROCO, before and after NER. When named entities are left in place, we notice that the KRR classifier selects some interesting named entities as discriminative. For example, news in the politics domain make a lot of references to politicians such as Liviu Dragnea (the ex-leader of the Social-Democrat Party in Romania), Igor Dodon (the current president of Moldova) or Dacian Cioloș (an ex-prime minster of Romania). News that mention NASA (the National Aeronautics and Space Administration) or the Max Planck institute are



Corpus	#dialects	#tokens per sample	Accuracy
Romanian (ours)	2	309.3	94.13
Arabic	5	22.6	76.27
German	4	7.9	66.36

Table 6: Accuracy rates (in %) of the KRR based on string kernels for Romanian dialect identification versus Arabic (Ali et al., 2016) and German (Samardžić et al., 2016) dialect identification, respectively. The results for the Arabic and German dialect identification tasks are taken from our previous work (Ionescu and Butnaru, 2017). For each corpus, we include the number of dialects (#dialects) and the average number of tokens in each sample (#tokens per sample).

likely to be classified in the science domain by  $KRR+k_6^{0/1}$ . After Simona Halep reached the first place in the Women’s Tennis Association (WTA) ranking, a lot of sports news that report on her performances started to appear, which determines the classifier to choose ‘Simona’ or ‘Halep’ as discriminative n-grams. References to the Internet or the Facebook social network indicate that the respective news are from the tech domain, according to our classifier. When named entities are removed, KRR seems to choose plausible words for each category. For instance, it relies on n-grams such as ‘muzică’ or ‘artist’ to classify a news sample into the culture domain, or on n-grams such as ‘champion’ or ‘fotbal’ to classify a news sample into the sports domain.

**Difficulty with respect to other dialects.** In our previous work (Ionescu and Butnaru, 2017), we have applied the KRR based on string kernels for Arabic dialect identification and German dialect identification. In the case of Arabic, we have reached performance levels of around 76% for discriminating between five dialects. In the same time, we have reached performance levels of around 66% for discriminating between four German dialects. As shown in Table 6, it seems to be much easier to discriminate between Romanian dialects, as the accuracy is near 94%. However, there are some important differences between these tasks. First of all, the random chance baseline is much high for our binary classification task, as we only have to choose between two dialects: Moldavian or Romanian. Second of all, the number of tokens per sample is much higher for the samples in our corpus compared to the samples provided in the Arabic (Ali et al., 2016) or the German (Samardžić et al., 2016) corpora. Before

drawing the conclusion that Romanian dialects are easier to discriminate than other dialects, we have to make sure that the experiments are conducted in similar conditions. We leave this discussion for future work.

## 7 Conclusion

In this paper, we presented a novel and large corpus of Moldavian and Romanian dialects. We also introduced Squeeze-and-Excitation networks to the NLP domain, performing comparative experiments using shallow and deep state-of-the-art baselines. We would like to stress out that the methods presented in this paper are only provided as baselines in order to enable comparisons in future work. Our intention was not that of providing top accuracy rates on the MOROCO corpus. In this context, we acknowledge that better accuracy rates can be obtained by combining string kernels using a range of n-grams, as we have already shown for other dialects and tasks in our previous works (Butnaru and Ionescu, 2018; Cozma et al., 2018; Ionescu and Butnaru, 2017, 2018). Another option for improving performance is to combine string kernels and neural networks into an ensemble model. We leave these ideas for future exploration.

Although Romanian and Moldavian are supposed to be hard to discriminate, since Romania and the Republic of Moldova share the same literary standard (Minahan, 2013), the empirical results seem to point in the other direction, to our surprise. However, we should note that the high accuracy rates attained by the proposed classifiers can be explained through a combination of two factors. First of all, the text samples are formed of 309 tokens on average, being at least an order of magnitude longer than samples in typical dialectal corpora (Ali et al., 2016; Samardžić et al., 2016). Second of all, the text samples can be discriminated in large part due to different word choices, as shown in the analysis of the most discriminative features provided in Section 5. Word preference seems to become easily distinguishable when news samples of around 309 tokens (multiple sentences) are used. In future work, we aim to determine if the same level of accuracy can be obtained when single sentences will be used as samples for training and testing.

## Acknowledgments

We thank reviewers for their useful comments.

## References

- Amjad Abu-Jbara, Rahul Jha, Eric Morley, and Dragomir Radev. 2013. Experimental Results on the Native Language Identification Shared Task. In *Proceedings of BEA-8*, pages 82–88.
- Ahmed Ali, Najim Dehak, Patrick Cardinal, Sameer Khurana, Sree Harsha Yella, James Glass, Peter Bell, and Steve Renals. 2016. Automatic Dialect Detection in Arabic Broadcast Speech. In *Proceedings of INTERSPEECH*, pages 2934–2938.
- Ahmed Ali, Stephan Vogel, and Steve Renals. 2017. Speech Recognition Challenge in the Wild: Arabic MGB-3. In *Proceedings of ASRU*, pages 316–322.
- Mohamed Ali. 2018. Character level convolutional neural network for arabic dialect identification. In *Proceedings of VarDial*, pages 122–127.
- Israa Alsarsour, Esraa Mohamed, Reem Suwaileh, and Tamer Elsayed. 2018. DART: A Large Dataset of Dialectal Arabic Tweets. In *Proceedings of LREC*, pages 3666–3670.
- Yonatan Belinkov and James Glass. 2016. A Character-level Convolutional Neural Network for Distinguishing Similar Languages and Dialects. In *Proceedings of VarDial*, pages 145–152.
- Houda Bouamor, Nizar Habash, Mohammad Salameh, Wajdi Zaghrouani, Owen Rambow, Dana Abdulrahim, Ossama Obeid, Salam Khalifa, Fadhl Eryani, Alexander Erdmann, and Kemal Oflazer. 2018. The MADAR Arabic Dialect Corpus and Lexicon. In *Proceedings of LREC*, pages 3387–3396.
- Andrei M. Butnaru and Radu Tudor Ionescu. 2018. UnibucKernel Reloaded: First Place in Arabic Dialect Identification for the Second Year in a Row. In *Proceedings of VarDial*, pages 77–87.
- Alina Maria Ciobanu and Liviu P. Dinu. 2016. A Computational Perspective on the Romanian Dialects. In *Proceedings of LREC*, pages 3281–3286.
- Ion Coteanu, Gheorghe Bolocan, and Matilda Caragiu Mariojeanu. 1969. *Istoria Limbii Române (History of the Romanian Language)*, volume II. Romanian Academy, Bucharest, Romania.
- Mădălina Cozma, Andrei M. Butnaru, and Radu Tudor Ionescu. 2018. Automated essay scoring with string kernels and word embeddings. In *Proceedings of ACL*, pages 503–509.
- Yann Dauphin, Angela Fan, Michael Auli, and David Grangier. 2017. Language Modeling with Gated Convolutional Networks. In *Proceedings of ICML*, pages 933–941.
- Jerid Francom, Mans Hulden, Adam Ussishkin, Julieta Fumagalli, Mikel Santesteban, and Julio Serrano. 2014. ACTIV-ES: A comparable, cross-dialect corpus of ‘everyday’ Spanish from Argentina, Mexico, and Spain. In *Proceedings of LREC*, pages 1733–1737.
- Jonas Gehring, Michael Auli, David Grangier, and Yann Dauphin. 2017. A Convolutional Encoder Model for Neural Machine Translation. In *Proceedings of ACL*, pages 123–135.
- Rosa M. Giménez-Pérez, Marc Franco-Salvador, and Paolo Rosso. 2017. Single and Cross-domain Polarity Classification using String Kernels. In *Proceedings of EACL*, pages 558–563.
- Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-Excitation Networks. In *Proceedings of CVPR*, pages 7132–7141.
- Radu Tudor Ionescu and Andrei M. Butnaru. 2017. Learning to Identify Arabic and German Dialects using Multiple Kernels. In *Proceedings of VarDial*, pages 200–209.
- Radu Tudor Ionescu and Andrei M. Butnaru. 2018. Improving the results of string kernels in sentiment analysis and Arabic dialect identification by adapting them to your test set. In *Proceedings of EMNLP*, pages 1084–1090.
- Radu Tudor Ionescu and Marius Popescu. 2016. UnibucKernel: An Approach for Arabic Dialect Identification based on Multiple String Kernels. In *Proceedings of VarDial*, pages 135–144.
- Radu Tudor Ionescu and Marius Popescu. 2017. Can string kernels pass the test of time in native language identification? In *Proceedings of BEA-12*, pages 224–234.
- Radu Tudor Ionescu, Marius Popescu, and Aoife Cahill. 2014. Can characters reveal your native language? A language-independent approach to native language identification. In *Proceedings of EMNLP*, pages 1363–1373.
- Radu Tudor Ionescu, Marius Popescu, and Aoife Cahill. 2016. String kernels for native language identification: Insights from behind the curtains. *Computational Linguistics*, 42(3):491–525.
- Janne Bondi Johannessen, Joel James Priestley, Kristin Hagen, Tor Anders Åfarli, and Øystein Alexander Vangnes. 2009. The Nordic Dialect Corpus—an advanced research tool. In *Proceedings of NODAL-IDA*, pages 73–80.
- Rie Johnson and Tong Zhang. 2015. Effective Use of Word Order for Text Categorization with Convolutional Neural Networks. In *Proceedings of NAACL*, pages 103–112.
- Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of EMNLP*, pages 1746–1751.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-Aware Neural Language Models. In *Proceedings of AAAI*, pages 2741–2749.

- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of NIPS*, pages 1106–1114.
- Ritesh Kumar, Bornini Lahiri, Deepak Alok, Atul Kr. Ojha, Mayank Jain, Abdul Basit, and Yogesh Dawar. 2018. Automatic Identification of Closely-related Indian Languages: Resources and Experiments. In *Proceedings of LREC*.
- Yann LeCun, Leon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324.
- Huma Lodhi, Craig Saunders, John Shawe-Taylor, Nello Cristianini, and Christopher J. C. H. Watkins. 2002. Text classification using string kernels. *Journal of Machine Learning Research*, 2:419–444.
- Shervin Malmasi, Marcos Zampieri, Nikola Ljubešić, Preslav Nakov, Ahmed Ali, and Jörg Tiedemann. 2016. Discriminating between Similar Languages and Arabic Dialect Identification: A Report on the Third DSL Shared Task. In *Proceedings of VarDial*, pages 1–14.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Proceedings of NIPS*, pages 3111–3119.
- James Minahan. 2013. *Miniature Empires: A Historical Dictionary of the Newly Independent States*. Routledge.
- Vinod Nair and Geoffrey E. Hinton. 2010. Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings of ICML*, pages 807–814.
- Garrett Nicolai and Grzegorz Kondrak. 2014. Does the Phonology of L1 Show Up in L2 Texts? In *Proceedings of ACL*, pages 854–859.
- Vasile Pavel. 2008. Limba română – unitate în diversitate (Romanian language – there is unity in diversity). *Romanian Language Journal*, XVIII(9–10).
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of EMNLP*, pages 1532–1543.
- Marius Popescu, Cristian Grozea, and Radu Tudor Ionescu. 2017. HASKER: An efficient algorithm for string kernels. Application to polarity classification in various languages. In *Proceedings of KES*, pages 1755–1763.
- Francisco Rangel, Paolo Rosso, Martin Potthast, and Benno Stein. 2017. Overview of the 5th author profiling task at PAN 2017: Gender and language variety identification in twitter. In *Working Notes Papers of the CLEF*.
- Tanja Samardžić, Yves Scherrer, and Elvira Glaser. 2016. ArchiMob—A corpus of spoken Swiss German. In *Proceedings of LREC*, pages 4061–4066.
- Cicero D. Santos and Bianca Zadrozny. 2014. Learning character-level representations for part-of-speech tagging. In *Proceedings of ICML*, pages 1818–1826.
- John Shawe-Taylor and Nello Cristianini. 2004. *Kernel Methods for Pattern Analysis*. Cambridge University Press.
- Liling Tan, Marcos Zampieri, Nikola Ljubešić, and Jörg Tiedemann. 2014. Merging Comparable Data Sources for the Discrimination of Similar Languages: The DSL Corpus Collection. In *Proceedings of BUCC*, pages 11–15.
- Omar F. Zaidan and Chris Callison-Burch. 2011. The Arabic Online Commentary Dataset: An Annotated Dataset of Informal Arabic with High Dialectal Content. In *Proceedings of ACL: HLT*, volume 2, pages 37–41.
- Marcos Zampieri, Shervin Malmasi, Nikola Ljubešić, Preslav Nakov, Ahmed Ali, Jörg Tiedemann, Yves Scherrer, and Noëmi Aepli. 2017. Findings of the VarDial Evaluation Campaign 2017. In *Proceedings of VarDial*, pages 1–15.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Ahmed Ali, Suwon Shon, James Glass, Yves Scherrer, Tanja Samardžić, Nikola Ljubešić, Jörg Tiedemann, Chris van der Lee, Stefan Grondelaers, Nelleke Oostdijk, Antal van den Bosch, Ritesh Kumar, Bornini Lahiri, and Mayank Jain. 2018. Language Identification and Morphosyntactic Tagging: The Second VarDial Evaluation Campaign. In *Proceedings of VarDial*, pages 1–17.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level Convolutional Networks for Text Classification. In *Proceedings of NIPS*, pages 649–657.