

# Improving the results of string kernels in sentiment analysis and Arabic dialect identification by adapting them to your test set

Radu Tudor Ionescu<sup>1,2</sup> and Andrei M. Butnaru<sup>1</sup>

<sup>1</sup>University of Bucharest,  
Department of Computer Science  
14 Academiei, Bucharest, Romania  
butnaruandreimadalin@gmail.com

<sup>2</sup>Inception Institute  
of Artificial Intelligence (IIAI)  
Al Maryah Island, Abu Dhabi, UAE  
raducu.ionescu@gmail.com

## Abstract

Recently, string kernels have obtained state-of-the-art results in various text classification tasks such as Arabic dialect identification or native language identification. In this paper, we apply two simple yet effective transductive learning approaches to further improve the results of string kernels. The first approach is based on interpreting the pairwise string kernel similarities between samples in the training set and samples in the test set as features. Our second approach is a simple self-training method based on two learning iterations. In the first iteration, a classifier is trained on the training set and tested on the test set, as usual. In the second iteration, a number of test samples (to which the classifier associated higher confidence scores) are added to the training set for another round of training. However, the ground-truth labels of the added test samples are not necessary. Instead, we use the labels predicted by the classifier in the first training iteration. By adapting string kernels to the test set, we report significantly better accuracy rates in English polarity classification and Arabic dialect identification.

## 1 Introduction

In recent years, methods based on string kernels have demonstrated remarkable performance in various text classification tasks ranging from authorship identification (Popescu and Grozea, 2012) and sentiment analysis (Giménez-Pérez et al., 2017; Popescu et al., 2017) to native language identification (Popescu and Ionescu, 2013; Ionescu et al., 2014, 2016; Ionescu and Popescu, 2017), dialect identification (Ionescu and Popescu, 2016b; Ionescu and Butnaru, 2017; Butnaru and Ionescu, 2018) and automatic essay scoring (Cozma et al., 2018). As long as a labeled training set is available, string kernels can reach state-of-the-art results in various languages in-

cluding English (Ionescu et al., 2014; Giménez-Pérez et al., 2017; Cozma et al., 2018), Arabic (Ionescu, 2015; Ionescu et al., 2016; Ionescu and Butnaru, 2017; Butnaru and Ionescu, 2018), Chinese (Popescu et al., 2017) and Norwegian (Ionescu et al., 2016). Different from all these recent approaches, we use unlabeled data from the test set to significantly increase the performance of string kernels. More precisely, we propose two transductive learning approaches combined into a unified framework. We show that the proposed framework improves the results of string kernels in two different tasks (cross-domain sentiment classification and Arabic dialect identification) and two different languages (English and Arabic). To the best of our knowledge, transductive learning frameworks based on string kernels have not been studied in previous works.

## 2 Transductive String Kernels

**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 (Lodhi et al., 2002; Shawe-Taylor and Cristianini, 2004; Ionescu et al., 2014). Perhaps one of the most recently introduced string kernels is the histogram intersection string kernel (Ionescu et al., 2014). For two strings over an alphabet  $\Sigma$ ,  $x, y \in \Sigma^*$ , the intersection string kernel is formally defined as follows:

$$k^\cap(x, y) = \sum_{v \in \Sigma^p} \min\{\text{num}_v(x), \text{num}_v(y)\}, \quad (1)$$

where  $\text{num}_v(x)$  is the number of occurrences of n-gram  $v$  as a substring in  $x$ , and  $p$  is the length of  $v$ . The spectrum string kernel or the presence bits string kernel can be defined in a similar fashion

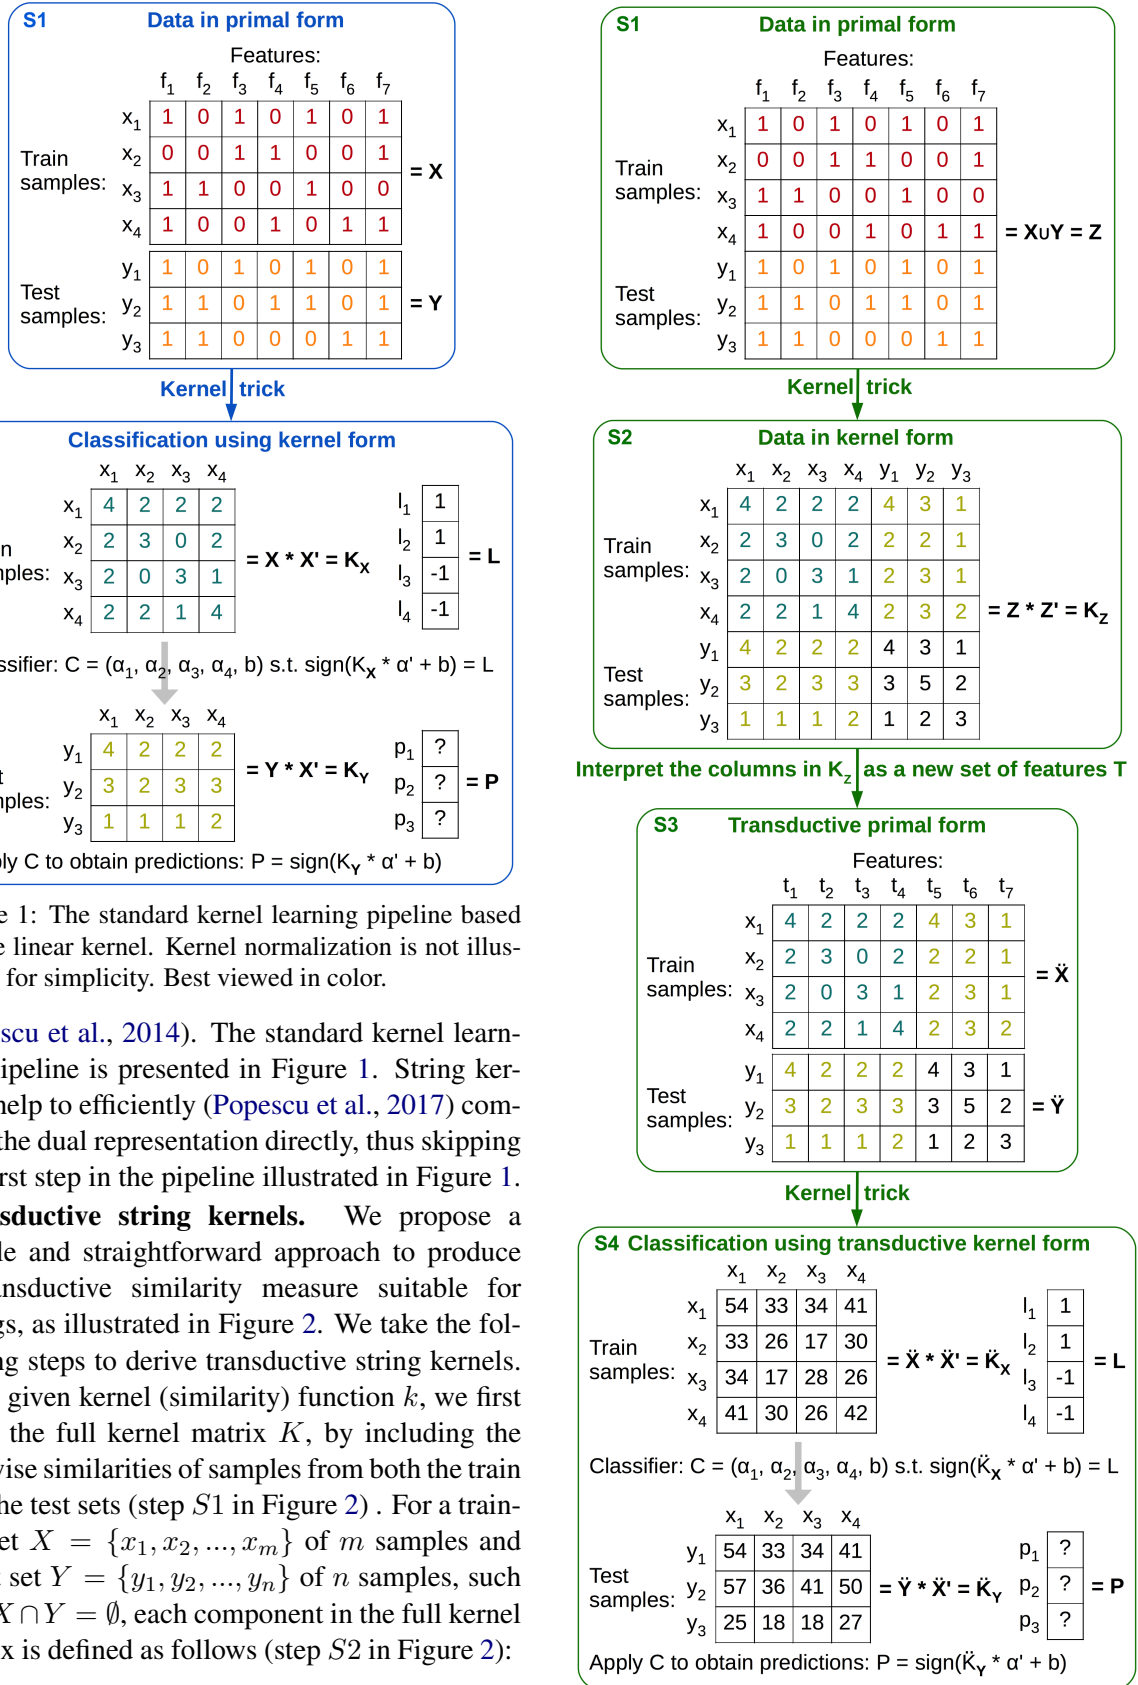


Figure 1: The standard kernel learning pipeline based on the linear kernel. Kernel normalization is not illustrated for simplicity. Best viewed in color.

(Ionescu et al., 2014). The standard kernel learning pipeline is presented in Figure 1. String kernels help to efficiently (Popescu et al., 2017) compute the dual representation directly, thus skipping the first step in the pipeline illustrated in Figure 1.

**Transductive string kernels.** We propose a simple and straightforward approach to produce a transductive similarity measure suitable for strings, as illustrated in Figure 2. We take the following steps to derive transductive string kernels. For a given kernel (similarity) function  $k$ , we first build the full kernel matrix  $K$ , by including the pairwise similarities of samples from both the train and the test sets (step S1 in Figure 2). For a training set  $X = \{x_1, x_2, \dots, x_m\}$  of  $m$  samples and a test set  $Y = \{y_1, y_2, \dots, y_n\}$  of  $n$  samples, such that  $X \cap Y = \emptyset$ , each component in the full kernel matrix is defined as follows (step S2 in Figure 2):

$$K_{ij} = k(z_i, z_j), \quad (2)$$

where  $z_i$  and  $z_j$  are samples from the set  $Z = X \cup Y = \{x_1, x_2, \dots, x_m, y_1, y_2, \dots, y_n\}$ , for all  $1 \leq i, j \leq m + n$ . We then normalize the kernel matrix by dividing each component by the square

Figure 2: The transductive kernel learning pipeline based on the linear kernel. Kernel normalization and RBF kernel transformation are not illustrated for simplicity. Best viewed in color.

root of the product of the two corresponding diagonal components:

$$\hat{K}_{ij} = \frac{K_{ij}}{\sqrt{K_{ii} \cdot K_{jj}}}. \quad (3)$$

We transform the normalized kernel matrix into a radial basis function (RBF) kernel matrix as follows:

$$\tilde{K}_{ij} = \exp\left(-\frac{1 - \hat{K}_{ij}}{2\sigma^2}\right). \quad (4)$$

As the kernel matrix is already normalized, we can choose  $\sigma^2 = 0.5$  for simplicity. Therefore, Equation (4) becomes:

$$\tilde{K}_{ij} = \exp\left(-1 + \hat{K}_{ij}\right). \quad (5)$$

Each row in the RBF kernel matrix  $\tilde{K}$  is now interpreted as a feature vector, going from step *S2* to step *S3* in Figure 2. In other words, each sample  $z_i$  is represented by a feature vector that contains the similarity between the respective sample  $z_i$  and all the samples in  $Z$  (step *S3* in Figure 2). Since  $Z$  includes the test samples as well, the feature vector is inherently adapted to the test set. Indeed, it is easy to see that the features will be different if we choose to apply the string kernel approach on a set of test samples  $Y'$ , such that  $Y' \neq Y$ . It is important to note that through the features, the subsequent classifier will have some information about the test samples at training time. More specifically, the feature vector conveys information about how similar is every test sample to every training sample. We next consider the linear kernel, which is given by the scalar product between the new feature vectors. To obtain the final linear kernel matrix, we simply need to compute the product between the RBF kernel matrix and its transpose (step *S4* in Figure 2):

$$\ddot{K} = \tilde{K} \cdot \tilde{K}'. \quad (6)$$

In this way, the samples from the test set, which are included in  $Z$ , are used to obtain new (transductive) string kernels that are adapted to the test set at hand.

**Transductive kernel classifier.** After obtaining the transductive string kernels, we use a simple transductive learning approach that falls in the category of self-training methods (McClosky et al., 2006; Chen et al., 2011). The transductive approach is divided into two learning iterations. In the first iteration, a kernel classifier is trained on the training data and applied on the test data, just as usual. Next, the test samples are sorted by the

classifier’s confidence score to maximize the probability of correctly predicted labels in the top of the sorted list. In the second iteration, a fixed number of samples (1000 in the experiments) from the top of the list are added to the training set for another round of training. Even though a small percent (less than 8% in all experiments) of the predicted labels corresponding to the newly included samples are wrong, the classifier has the chance to learn some useful patterns (from the correctly predicted labels) only visible in the test data. The transductive kernel classifier (TKC) is based on the intuition that the added test samples bring more useful information than noise, since the majority of added test samples have correct labels. Finally, we would like to stress out that *the ground-truth test labels are never used in our transductive algorithm*.

The proposed transductive learning approaches are used together in a unified framework. As any other transductive learning method, the main disadvantage of the proposed framework is that the *unlabeled* test samples from the target domain need to be used in the training stage. Nevertheless, we present empirical results indicating that our approach can obtain *significantly better* accuracy rates in cross-domain polarity classification and Arabic dialect identification compared to state-of-the-art methods based on string kernels (Giménez-Pérez et al., 2017; Ionescu and Butnaru, 2017). We also report better results than other domain adaptation methods (Pan et al., 2010; Bollegala et al., 2013; Franco-Salvador et al., 2015; Sun et al., 2016; Huang et al., 2017).

### 3 Polarity Classification

**Data set.** For the cross-domain polarity classification experiments, we use the second version of Multi-Domain Sentiment Dataset (Blitzer et al., 2007). The data set contains Amazon product reviews of four different domains: Books (B), DVDs (D), Electronics (E) and Kitchen appliances (K). Reviews contain star ratings (from 1 to 5) which are converted into binary labels as follows: reviews rated with more than 3 stars are labeled as positive, and those with less than 3 stars as negative. In each domain, there are 1000 positive and 1000 negative reviews.

**Baselines.** We compare our approach with several methods (Pan et al., 2010; Bollegala et al., 2013; Franco-Salvador et al., 2015; Sun et al., 2016; Giménez-Pérez et al., 2017; Huang et al.,

Method	DEK→B	BEK→D	BDK→E	BDE→K
SST	76.3	78.3	83.9	85.2
KE-Meta	77.9	80.4	78.9	82.5
$K_{0/1}$ (sota)	82.0	81.9	83.6	85.1
$K_{\cap}$ (sota)	80.7	80.7	83.0	85.2
$\ddot{K}_{0/1}$	82.9	83.2*	84.8*	86.0*
$\ddot{K}_{\cap}$	82.5	82.9*	84.5*	86.1*
$\ddot{K}_{0/1}$ + TKC	<b>84.1*</b>	<b>84.0*</b>	<b>85.4*</b>	86.9*
$\ddot{K}_{\cap}$ + TKC	83.8*	83.5*	85.0*	<b>87.1*</b>

Table 1: Multi-source cross-domain polarity classification accuracy rates (in %) of our transductive approaches versus a state-of-the-art (sota) baseline based on string kernels (Giménez-Pérez et al., 2017), as well as SST (Bollegala et al., 2013) and KE-Meta (Franco-Salvador et al., 2015). The best accuracy rates are highlighted in bold. The marker \* indicates that the performance is significantly better than the best baseline string kernel according to a paired McNemar’s test performed at a significance level of 0.01.

2017) in two cross-domain settings. Using string kernels, Giménez-Pérez et al. (2017) reported better performance than SST (Bollegala et al., 2013) and KE-Meta (Franco-Salvador et al., 2015) in the multi-source domain setting. In addition, we compare our approach with SFA (Pan et al., 2010), KMM (Huang et al., 2007), CORAL (Sun et al., 2016) and TR-TrAdaBoost (Huang et al., 2017) in the single-source setting.

**Evaluation procedure and parameters.** We follow the same evaluation methodology of Giménez-Pérez et al. (2017), to ensure a fair comparison. Furthermore, we use the same kernels, namely the presence bits string kernel ( $K_{0/1}$ ) and the intersection string kernel ( $K_{\cap}$ ), and the same range of character n-grams (5-8). To compute the string kernels, we used the open-source code provided by Ionescu and Popescu (2016a). For the transductive kernel classifier, we select  $r = 1000$  unlabeled test samples to be included in the training set for the second round of training. We choose Kernel Ridge Regression (Shawe-Taylor and Cristianini, 2004) as classifier and set its regularization parameter to  $10^{-5}$  in all our experiments. Although Giménez-Pérez et al. (2017) used a different classifier, namely Kernel Discriminant Analysis, we observed that Kernel Ridge Regression produces similar results ( $\pm 0.1\%$ ) when we employ the same string kernels. As Giménez-Pérez et al. (2017), we evaluate our approach in two cross-domain settings. In the multi-source setting, we train the models on all domains, except the one used for testing. In the single-source setting, we train the models on one of the four domains and we independently test the models on the remain-

ing three domains.

**Results in multi-source setting.** The results for the multi-source cross-domain polarity classification setting are presented in Table 1. Both the transductive presence bits string kernel ( $\ddot{K}_{0/1}$ ) and the transductive intersection kernel ( $\ddot{K}_{\cap}$ ) obtain better results than their original counterparts. Moreover, according to the McNemar’s test (Dietterich, 1998), the results on the DVDs, the Electronics and the Kitchen target domains are significantly better than the best baseline string kernel, with a confidence level of 0.01. When we employ the transductive kernel classifier (TKC), we obtain even better results. On all domains, the accuracy rates yielded by the transductive classifier are more than 1.5% better than the best baseline. For example, on the Books domain the accuracy of the transductive classifier based on the presence bits kernel (84.1%) is 2.1% above the best baseline (82.0%) represented by the intersection string kernel. Remarkably, the improvements brought by our transductive string kernel approach are statistically significant in all domains.

**Results in single-source setting.** The results for the single-source cross-domain polarity classification setting are presented in Table 2. We considered all possible combinations of source and target domains in this experiment, and we improve the results in each and every case. Without exception, the accuracy rates reached by the transductive string kernels are significantly better than the best baseline string kernel (Giménez-Pérez et al., 2017), according to the McNemar’s test performed at a confidence level of 0.01. The highest improvements (above 2.7%) are obtained when the source domain contains Books reviews and the target domain contains Kitchen reviews. As in the multi-source setting, we obtain much better results when the transductive classifier is employed for the learning task. In all cases, the accuracy rates of the transductive classifier are more than 2% better than the best baseline string kernel. Remarkably, in four cases (E→B, E→D, B→K and D→K) our improvements are greater than 4%. The improvements brought by our transductive classifier based on string kernels are statistically significant in each and every case. In comparison with SFA (Pan et al., 2010), we obtain better results in all but one case (K→D). With respect to KMM (Huang et al., 2007), we also obtain better results in all but one case (B→E). Remarkably, we surpass the

Method	D→B	E→B	K→B	B→D	E→D	K→D	B→E	D→E	K→E	B→K	D→K	E→K
SFA	79.8	78.3	75.2	81.4	77.2	<b>78.5</b>	73.5	76.7	85.1	79.1	80.8	86.8
KMM	78.6	-	-	-	-	72.2	<b>76.9</b>	-	-	-	-	83.6
CORAL	78.3	-	-	-	-	73.9	76.3	-	-	-	-	83.6
TR-TrAdaBoost	74.7	69.1	70.6	79.6	71.8	74.4	74.9	75.9	83.1	77.8	75.7	83.7
$K_{0/1}$ (sota)	82.0	72.4	72.7	81.4	74.9	73.6	71.3	74.4	83.9	74.6	75.4	84.9
$K_{\cap}$ (sota)	82.1	72.4	72.8	81.3	75.1	72.9	71.8	74.5	84.4	74.9	75.1	84.9
$\ddot{K}_{0/1}$	83.3*	74.5*	74.3*	83.0*	76.9*	74.9*	74.0*	76.0*	85.4*	77.6*	77.3*	86.0*
$\ddot{K}_{\cap}$	83.2*	74.2*	74.0*	82.8*	76.4*	75.1*	74.2*	75.9*	85.2*	77.6*	77.3*	85.9*
$\ddot{K}_{0/1} + \text{TKC}$	<b>84.9*</b>	<b>78.5*</b>	<b>76.6*</b>	84.0*	<b>79.6*</b>	76.4*	76.6*	<b>77.1*</b>	<b>86.4*</b>	<b>79.6*</b>	<b>80.9*</b>	<b>87.0*</b>
$\ddot{K}_{\cap} + \text{TKC}$	84.5*	<b>78.5*</b>	75.8*	<b>84.2*</b>	79.1*	76.5*	76.7*	76.8*	<b>86.4*</b>	79.4*	80.5*	<b>87.0*</b>

Table 2: Single-source cross-domain polarity classification accuracy rates (in %) of our transductive approaches versus a state-of-the-art (sota) baseline based on string kernels (Giménez-Pérez et al., 2017), as well as SFA (Pan et al., 2010), KMM (Huang et al., 2007), CORAL (Sun et al., 2016) and TR-TrAdaBoost (Huang et al., 2017). The best accuracy rates are highlighted in bold. The marker \* indicates that the performance is significantly better than the best baseline string kernel according to a paired McNemar’s test performed at a significance level of 0.01.

other state-of-the-art approaches (Sun et al., 2016; Huang et al., 2017) in all cases.

#### 4 Arabic Dialect Identification

**Data set.** The Arabic Dialect Identification (ADI) data set (Ali et al., 2016) contains audio recordings and Automatic Speech Recognition (ASR) transcripts of Arabic speech collected from the Broadcast News domain. The classification task is to discriminate between Modern Standard Arabic and four Arabic dialects, namely Egyptian, Gulf, Levantine, and Maghrebi. The training set contains 14000 samples, the development set contains 1524 samples, and the test contains another 1492 samples. The data set was used in the ADI Shared Task of the 2017 VarDial Evaluation Campaign (Zampieri et al., 2017).

**Baseline.** We choose as baseline the approach of Ionescu and Butnaru (2017), which is based on string kernels and multiple kernel learning. The approach that we consider as baseline is the winner of the 2017 ADI Shared Task (Zampieri et al., 2017). In addition, we also compare with the second-best approach (Meta-classifier) (Malmasi and Zampieri, 2017).

**Evaluation procedure and parameters.** Ionescu and Butnaru (2017) combined four kernels into a sum, and used Kernel Ridge Regression for training. Three of the kernels are based on character n-grams extracted from ASR transcripts. These are the presence bits string kernel ( $K_{0/1}$ ), the intersection string kernel ( $K_{\cap}$ ), and a kernel based on Local Rank Distance ( $K_{LRD}$ ) (Ionescu, 2013). The fourth kernel is an RBF kernel ( $K_{ivec}$ ) based on the i-vectors provided with the ADI data set (Ali et al., 2016). In our experiments, we employ the exact same kernels as Ionescu and Butnaru (2017) to ensure an unbiased comparison with their ap-

Method	Development	Test
Meta-classifier	-	71.65
$K_{0/1} + K_{\cap} + K_{LRD} + K_{ivec}$ (sota)	64.17	76.27
$\ddot{K}_{0/1} + \ddot{K}_{\cap} + \ddot{K}_{LRD} + \ddot{K}_{ivec}$	65.42*	77.08
$\ddot{K}_{0/1} + \ddot{K}_{\cap} + \ddot{K}_{LRD} + \ddot{K}_{ivec} + \text{TKC}$	<b>66.73*</b>	<b>78.35*</b>

Table 3: Arabic dialect identification accuracy rates (in %) of our adapted string kernels versus the 2017 ADI Shared Task winner (sota) (Ionescu and Butnaru, 2017) and the first runner up (Malmasi and Zampieri, 2017). The best accuracy rates are highlighted in bold. The marker \* indicates that the performance is significantly better than (Ionescu and Butnaru, 2017) according to a paired McNemar’s test performed at a significance level of 0.01.

proach. As in the polarity classification experiments, we select  $r = 1000$  unlabeled test samples to be included in the training set for the second round of training the transductive classifier, and we use Kernel Ridge Regression with a regularization of  $10^{-5}$  in all our ADI experiments.

**Results.** The results for the cross-domain Arabic dialect identification experiments on both the development and the test sets are presented in Table 3. The domain-adapted sum of kernels obtains improvements above 0.8% over the state-of-the-art sum of kernels (Ionescu and Butnaru, 2017). The improvement on the development set (from 64.17% to 65.42%) is statistically significant. Nevertheless, we obtain higher and significant improvements when we employ the transductive classifier. Our best accuracy is 66.73% (2.56% above the baseline) on the development set and 78.35% (2.08% above the baseline) on the test set. The results show that our domain adaptation framework based on string kernels attains the best performance on the ADI Shared Task data set, and the improvements over the state-of-the-art are statistically significant, according to the McNemar’s test.

## References

- 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.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boomboxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of ACL*, pages 187–205.
- D. Bollegala, D. Weir, and J. Carroll. 2013. Cross-Domain Sentiment Classification Using a Sentiment Sensitive Thesaurus. *IEEE Transactions on Knowledge and Data Engineering*, 25(8):1719–1731.
- 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 Workshop of COLING*, pages 77–87.
- Minmin Chen, Kilian Weinberger, and John Blitzer. 2011. Co-Training for Domain Adaptation. In *Proceedings of NIPS*, pages 2456–2464.
- Mădălina Cozma, Andrei Butnaru, and Radu Tudor Ionescu. 2018. Automated essay scoring with string kernels and word embeddings. In *Proceedings of ACL*, pages 503–509.
- Thomas G. Dietterich. 1998. Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Computation*, 10(7):1895–1923.
- Marc Franco-Salvador, Fermin L. Cruz, Jose A. Troyano, and Paolo Rosso. 2015. Cross-domain polarity classification using a knowledge-enhanced meta-classifier. *Knowledge-Based Systems*, 86:46–56.
- 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.
- Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex Smola. 2007. Correcting sample selection bias by unlabeled data. In *Proceedings of NIPS*, pages 601–608.
- Xingchang Huang, Yanghui Rao, Haoran Xie, Tak-Lam Wong, and Fu Lee Wang. 2017. Cross-Domain Sentiment Classification via Topic-Related TrAdaBoost. In *Proceedings of AAAI*, pages 4939–4940.
- Radu Tudor Ionescu. 2013. Local Rank Distance. In *Proceedings of SYNASC*, pages 221–228.
- Radu Tudor Ionescu. 2015. A Fast Algorithm for Local Rank Distance: Application to Arabic Native Language Identification. In *Proceedings of ICONIP*, volume 9490, pages 390–400.
- Radu Tudor Ionescu and Andrei Butnaru. 2017. Learning to Identify Arabic and German Dialects using Multiple Kernels. In *Proceedings of VarDial Workshop of EACL*, pages 200–209.
- Radu Tudor Ionescu and Marius Popescu. 2016a. Native Language Identification with String Kernels. In *Knowledge Transfer between Computer Vision and Text Mining*, Advances in Computer Vision and Pattern Recognition, chapter 8, pages 193–227. Springer International Publishing.
- Radu Tudor Ionescu and Marius Popescu. 2016b. UnibucKernel: An Approach for Arabic Dialect Identification based on Multiple String Kernels. In *Proceedings of VarDial Workshop of COLING*, 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 the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, 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.
- 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 and Marcos Zampieri. 2017. Arabic Dialect Identification Using iVectors and ASR Transcripts. In *Proceedings of the VarDial Workshop of EACL*, pages 178–183.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Effective Self-training for Parsing. In *Proceedings of NAACL*, pages 152–159.
- Sinno Jialin Pan, Xiaochuan Ni, Jian-Tao Sun, Qiang Yang, and Zheng Chen. 2010. Cross-domain Sentiment Classification via Spectral Feature Alignment. In *Proceedings of WWW*, pages 751–760.
- Marius Popescu and Cristian Grozea. 2012. Kernel methods and string kernels for authorship analysis. In *Proceedings of CLEF (Online Working Notes/Labs/Workshop)*.
- 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.

- Marius Popescu and Radu Tudor Ionescu. 2013. The Story of the Characters, the DNA and the Native Language. In *Proceedings of the Eighth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 270–278.
- John Shawe-Taylor and Nello Cristianini. 2004. *Kernel Methods for Pattern Analysis*. Cambridge University Press.
- Baochen Sun, Jiashi Feng, and Kate Saenko. 2016. Return of Frustratingly Easy Domain Adaptation. In *Proceedings of AAAI*, pages 2058–2065.
- 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 Workshop of EACL*, pages 1–15.