

Combining Deep Learning and String Kernels for the Localization of Swiss German Tweets

Mihaela Găman

Department of Computer Science
University of Bucharest
14 Academiei, Bucharest, Romania
mp.gaman@gmail.com

Radu Tudor Ionescu

Department of Computer Science
Romanian Young Academy
University of Bucharest
14 Academiei, Bucharest, Romania
raducu.ionescu@gmail.com

Abstract

In this work, we introduce the methods proposed by the UnibucKernel team in solving the Social Media Variety Geolocation task featured in the 2020 VarDial Evaluation Campaign. We address only the second subtask, which targets a data set composed of nearly 30 thousand Swiss German Jodels. The dialect identification task is about accurately predicting the latitude and longitude of test samples. We frame the task as a double regression problem, employing a variety of machine learning approaches to predict both latitude and longitude. From simple models for regression, such as Support Vector Regression, to deep neural networks, such as Long Short-Term Memory networks and character-level convolutional neural networks, and, finally, to ensemble models based on meta-learners, such as XGBoost, our interest is focused on approaching the problem from a few different perspectives, in an attempt to minimize the prediction error. With the same goal in mind, we also considered many types of features, from high-level features, such as BERT embeddings, to low-level features, such as characters n-grams, which are known to provide good results in dialect identification. Our empirical results indicate that the handcrafted model based on string kernels outperforms the deep learning approaches. Nevertheless, our best performance is given by the ensemble model that combines both handcrafted and deep learning models.

1 Introduction

The organizers of the 2020 VarDial Evaluation Campaign (Găman et al., 2020) proposed a shared task targeted towards the geolocation of short texts, e.g. tweets, namely the Social Media Variety Geolocation (SMG) task. Typically formulated as a double regression problem, the task is about predicting the location, expressed in latitude and longitude, from where the text received as input was posted on a certain social media platform. Twitter and Jodel are the platforms used for data collection, divided by the language area in three subtasks, namely:

- Standard German Jodels (DE-AT) - formed of conversations initiated in Germany and Austria in regional dialectal forms (Hovy and Purschke, 2018).
- Swiss German Jodels (CH) - based on a smaller number of Jodel conversations from Switzerland (Hovy and Purschke, 2018).
- BCMS Tweets - from the area of Bosnia and Herzegovina, Croatia, Montenegro and Serbia where the macro-language used is BCMS, with both similarities and a fair share of variation among the component languages (Ljubešić et al., 2016).

In this paper, we focus only on the second subtask, SMG-CH, proposing a variety of handcrafted and deep learning models, as well as an ensemble model that combines all our previous models through meta-learning. Our first model is a Support Vector Regression (SVR) classifier (Chang and Lin, 2002) based on string kernels, which are known to perform well in other dialect identification tasks (Butnaru

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

and Ionescu, 2018b; Ionescu and Popescu, 2016; Ionescu and Butnaru, 2017). Our second model is a character-level convolutional neural network (CNN) (Zhang et al., 2015), which is also known to provide good results in dialect identification (Butnaru and Ionescu, 2019; Tudoreanu, 2019). Due to the high popularity and the outstanding results of Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) in solving mainstream NLP tasks, we decided to try out a Long Short-Term Memory (LSTM) network (Hochreiter and Schmidhuber, 1997) based on German BERT embeddings as our third model. Lastly, we combine our three models into an ensemble that employs Extreme Gradient Boosting (XGBoost) (Chen and Guestrin, 2016) as meta-learner. We conducted experiments on the development set provided by the organizers, in order to decide which models to choose for our three submissions for the SMG-CH subtask. Our results indicate that the ensemble model attains the best results. Perhaps surprisingly, our shallow approach based on string kernels outperforms both deep learning models. Our observations are consistent across the development and the test sets provided by the organizers.

The rest of this paper is organized as follows. We present related work on dialect identification and geolocation of short texts in Section 2. Our approaches are described in more detail in Section 3. We present the experiments and empirical results in Section 4. Finally, our conclusions are drawn in Section 5.

2 Related Work

One of the initial works on text-based geotagging (Ding et al., 2000) aims at automatically finding the geographic scope of web pages, in a classification setup relying on named location entities such as cities and states. The authors used gazetteers as the source of the location mappings, proposing a rather heuristic approach. Gazetteers, constitute a tool used in one of the three general approaches taken so far in text-based geolocation, this tool being adopted in a number of works (Lieberman et al., 2010; Quercini et al., 2010; Cheng et al., 2010). In this line of research, some researchers employed rule-based methods (Bilhaut et al., 2003), while others plugged named entity recognition into various machine learning techniques (Gelernter and Mushegian, 2011; Qin et al., 2010). The main disadvantage of these methods is that they rely on the existence of specific mentions of locations in text, rather than inferring them in a not so straightforward manner. These direct mentions of places do not represent a safe assumption, especially when it comes to social media platforms such as Twitter, which is used as the data source in some of these studies (Cheng et al., 2010). The other two main categories of approaches for text-based geolocation rely on either unsupervised learning (Ahmed et al., 2013; Hong et al., 2012; Eisenstein et al., 2010) or supervised classification (Wing and Baldrige, 2011; Kinsella et al., 2011). The unsupervised methods can be described in large part as clustering techniques based on topic models.

There are some studies on user geolocation in social media, that look at this task from a supervised learning perspective (Rout et al., 2013) and can be included in the second set of approaches for geotagging. However, in such works, other details (e.g. social ties) in the users profile have been considered rather than their written content. Although these works cover geolocation prediction in social media, they do not use text as input. Our current interest in studying language variation for the geolocation of users in social media has been covered in the literature in a series of works (Rahimi et al., 2017; Han et al., 2014; Doyle, 2014; Roller et al., 2012; Eisenstein et al., 2010), employing various machine learning techniques, that range from probabilistic graphical models (Eisenstein et al., 2010) and adaptive grid search (Roller et al., 2012) to Bayesian methods (Doyle, 2014) and neural networks (Rahimi et al., 2017).

The related work to date covers a wide range of languages and dialects, including Dutch (Wieling et al., 2011), British (Szmrecsanyi, 2008), American (Huang et al., 2016; Eisenstein et al., 2010) and even African American Vernacular English (Jones, 2015). Most related to our work is the study of Hovy and Purschke (2018), which targets the German language and its variations and, in addition to the previously mentioned endeavours, performs a quantitative analysis against a dialect map. Moreover, Hovy and Purschke (2018) collected 16.8 million online posts from the German-speaking area with the aim of learning document representations of cities. Among these posts, some were from the German speaking side of Switzerland, being part of the SMG shared task, more specifically the SMG-CH subtask that we

are addressing. The authors aimed at capturing enough regional variations in the written language, serving as input in automatically distinguishing the geographical region of speakers. The focus was on larger regions covering a given dialect, the proposed approach being based on clustering. Given the shared task formulation, we take a different approach and use the provided data in a double regression setup, addressing the problem both from a shallow perspective and a deep learning perspective, respectively.

3 Methods

3.1 ν -Support Vector Regression based on String Kernels.

String Kernels. Lodhi et al. (2001) introduced string kernels as a means of comparing two documents, based on the inner product generated by all substrings of length n , typically known as character n -grams. Since then, string kernels have found many applications, from sentiment analysis (Giménez-Pérez et al., 2017; Ionescu and Butnaru, 2018; Popescu et al., 2017), automated essay scoring (Cozma et al., 2018) and sentence selection (Masala et al., 2017) to native language identification (Ionescu et al., 2014; Ionescu et al., 2016; Ionescu and Popescu, 2017; Popescu and Ionescu, 2013) and dialect identification (Butnaru and Ionescu, 2018b; Butnaru and Ionescu, 2019; Ionescu and Butnaru, 2017).

In this work, we employ string kernels as described in (Butnaru and Ionescu, 2019), specifically using the efficient algorithm for building string kernels of Popescu et al. (2017). We note that the number of character n -grams is usually much higher than the number of samples, so representing the text samples as feature vectors may require a lot of space. String kernels provide an efficient way to avoid storing and using the feature vectors (primal form), by representing the data through a kernel matrix (dual form). Each cell in the kernel matrix represents the similarity between some text samples x_i and x_j . In our experiments, we use the presence bits string kernel (Popescu and Ionescu, 2013) as the similarity function. For two strings x_i and x_j over a set of characters S , the presence bits string kernel is defined as follows:

$$k^{0/1}(x_i, x_j) = \sum_{g \in S^n} \#(x_i, g) \cdot \#(x_j, g), \quad (1)$$

where n is the length of n -grams and $\#(x, g)$ is a function that returns 1 when the number of occurrences of n -gram g in x is greater than 1, and 0 otherwise.

ν -Support Vector Regression. Support Vector Machines (SVM) (Cortes and Vapnik, 1995) represent a popular method initially designed for binary classification, which was subsequently repurposed for regression, under the SVR (Drucker et al., 1997) acronym (i.e. Support Vector Regression). Similar to SVM, SVR uses the notion of support vectors and margin in order to find an optimal estimator. In the original ϵ -SVR formulation (Drucker et al., 1997), there is an ϵ -insensitive region, i.e. ϵ tube, defined in the optimization function. The goal is to find the flattest tube containing most of the training samples, while also minimizing the prediction error and model complexity. Different from linear regression, ϵ -SVR fits the error within the maximum margin ϵ , instead of minimizing the error directly (Smola and Schölkopf, 2004). In our experiments, we employ an equivalent SVR formulation known as ν -SVR (Chang and Lin, 2002), where ν is the configurable proportion of support vectors to keep with respect to the number of samples in the data set. In ν -SVR, the margin ϵ is automatically estimated to its optimal value. Using ν -SVR, the optimal solution can converge to a small model, with only a few support vectors. This is especially useful in our use case, as the data set provided in the SMG-CH subtask does not contain too many samples. Another reason to employ ν -SVR in our regression task is that it was found to surpass other regression methods in complex word identification (Butnaru and Ionescu, 2018a).

3.2 Character-Level Convolutional Neural Network

Character Embeddings. From the pioneering works in language modelling at the character level (Gasthaus et al., 2010; Wood et al., 2009) to date (Georgescu et al., 2020), a broad range of neural architectures rely on characters as features. Among these, we can mention Recurrent Neural Networks (RNNs) (Sutskever et al., 2011), LSTM networks (Ballesteros et al., 2015), CNNs (Kim et al., 2016; Zhang et al., 2015) and transformer models (Al-Rfou et al., 2019). Characters are the base units in building words that exist in the vocabulary of most languages. Knowledge of words, semantic structure

or syntax is not required when working with characters. Robustness to spelling errors and words that are outside the vocabulary (Ballesteros et al., 2015) constitute other advantages explaining the growing interest in using characters as features. In our paper, we employ a convolutional neural network working at the character level (Zhang et al., 2015). The employed CNN is equipped with a character embedding layer, automatically learning a 2D representation of text formed of character embedding vectors, that is further processed by the convolutional layers.

Convolutional Neural Networks. Inspired by the visual cortex of mammals (Fukushima, 1980), CNNs have been extensively used in image classification (LeCun et al., 1989; LeCun et al., 2004; Krizhevsky et al., 2012), subsequently being adapted for various NLP tasks (Kim, 2014; Zhang et al., 2015). CNNs are composed of convolutional blocks, consisting of convolutions and pooling operations, usually followed by a sequence of dense layers and ending with an output layer, with the number of neurons equal to the number of values that we are interested in predicting. In the experiments, we employ a character-level CNN (Zhang et al., 2015) with squeeze-and-excitation (SE) blocks, introduced by Butnaru and Ionescu (2019). Since this method has been previously applied in Romanian dialect identification with good results (Butnaru and Ionescu, 2019; Tudoreanu, 2019), we consider it a good candidate for our text geolocation task. We therefore change the original architecture by replacing the Softmax classification layer with a regression layer formed of two units, one predicting the latitude and one predicting the longitude, respectively. We train our character-level CNN towards minimizing the mean squared error (MSE) loss function with respect to the ground-truth latitude and longitude.

3.3 Long Short-Term Memory Networks based on BERT Embeddings

BERT Embeddings. Transformers (Vaswani et al., 2017) represent a very important advance in Natural Language Processing, with many benefits over the traditional sequential neural architectures. Based on an encoder-decoder architecture with attention, transformers proved to be better at modelling long-term dependencies in sequences, while being effectively trained as the sequential dependency of previous tokens is removed. Unlike other contemporary attempts at using transformers in language modelling (Radford et al., 2018), BERT (Devlin et al., 2019) incorporates context from both directions in the process of building deep language representations, in a self-supervised fashion. The masked language modeling technique enables BERT to pre-train these deep bidirectional representations, that can be further fine-tuned and adapted for a variety of tasks, without significant architectural updates. We also make use of this property in the current work, employing a TensorFlow version of a German BERT model¹. The model has been trained on the latest German Wikipedia dump, the OpenLegalData dump and a collection of news articles, summing up to a total of 12 GB of text files. We add the pre-trained German BERT model to be fine-tuned in an end-to-end fashion along with our LSTM architecture for geolocation of Swiss German short texts.

Long Short-Term Memory Networks. RNNs (Werbos, 1988) operate at the sequence level, attaining state-of-the-art performance on various problems involving time series (Weiss et al., 2018). Major drawbacks in regular RNNs are the phenomenons of exploding and vanishing gradients, which can be caused by an increase in the length of the input sequence (Hochreiter et al., 2001). LSTM networks (Hochreiter and Schmidhuber, 1997) represent a flavour of RNN, designed to overcome the aforementioned challenges faced when working with RNNs. An LSTM unit has a more complex structure, including a memory cell to remember dependencies in the input and three gates acting as regulators: input, output and, in more recent versions, forget gates, which enable the cell to reset its state (Gers et al., 2000). The LSTM architecture used in this work is inspired by the one described in (Onose et al., 2019), which has been successfully employed in Romanian dialect identification. We train our LSTM model using the mean squared logarithmic error as loss function. We opted for the aforementioned loss in favor of the mean squared error, as the latter loss function did not produce optimal results for our LSTM.

¹<https://github.com/deepset-ai/FARM>

3.4 Ensemble Learning

XGBoost. Gradient tree boosting (Friedman, 2001) is based on training a tree ensemble model in an additive fashion. This technique has been successfully used in classification (Li, 2010) and ranking (Burges, 2010) problems, obtaining notable results in reputed competitions such as the Netflix Challenge (Bennett et al., 2007). Furthermore, gradient tree boosting is the ensemble method of choice in real-world pipelines running in production (He et al., 2014). XGBoost (Chen and Guestrin, 2016) is a tree boosting model targeted at solving large-scale tasks with limited computational resources. This approach aims at parallelizing tree learning while also trying to handle various sparsity patterns. Overfitting is addressed through shrinkage and column subsampling. Shrinkage acts as a learning rate, reducing the influence of each individual tree. Column subsampling is borrowed from Random Forests (Breiman, 2001), bearing the advantage of speeding up the computations. In the experiments, we employ XGBoost as a meta-learner over the individual predictions of each of the models described above. We opted for XGBoost in detriment of average voting and a ν -SVR meta-learner, both providing comparatively lower performance levels in a set of preliminary ensemble experiments.

4 Experiments

4.1 Data Set

The data set for the SMG-CH subtask contains a training set of 22,600 samples, with one sample per line, each formed of a piece of text and a pair of coordinates representing the position on Earth, i.e. latitude and longitude. The development set is composed of 3,086 samples, provided in the same format. The test set consists in 3,097 samples without coordinates. We note that the centroid computed on the training data has a latitude of 47.26 degrees and a longitude of 8.33 degrees, confirming that the average location is on the territory of Switzerland.

4.2 Parameter Tuning

ν -SVR based on string kernels. In the experiments, we use ν -SVR with a pre-computed string kernel, employing the efficient algorithm proposed in (Popescu et al., 2017). In a set of preliminary experiments, we employed various blended spectrum string kernels based on various n-gram ranges that include n-grams from 3 to 7 characters long. The best performance in terms of both mean absolute error (MAE) and mean squared error (MSE) were attained by a string kernel based on the blended spectrum of 3 to 5 character n-grams. These results are consistent with those reported by Ionescu and Butnaru (2017), suggesting that the 3-5 n-gram range is optimal for German dialect identification. The resulting kernel matrix is used as input in a double regression setup, with a ν -SVR model for predicting the latitude (in degrees), and another ν -SVR model for predicting the longitude (in degrees), respectively. We tried out values ranging from 10^{-4} to 10^4 for the regularization penalty C , during the hyperparameter tuning phase. Similarly, for the proportion of support vectors ν , we considered 10 values covering the interval $(0, 1]$ with a step of 0.1. For both regression models, the best value for the parameter C is 10. As for the parameter ν , the default value of 0.5 seems to yield the best results.

Character-level CNN. Except for the last layer, the architecture used in our experiments is identical to the architecture employed by Butnaru and Ionescu (2019) for Romanian dialect identification. An input of maximum 5000 characters (zero-padding is used as necessary) is expected into the network, with the characters initially encoded with their position in the vocabulary. For each character, a vectorial representation of 128 elements is learned in the embedding layer. Three convolutional blocks follow, each being composed of a convolutional layer with 128 one-dimensional filters of size 7 for the first two blocks, and of size 3 for the last block, respectively. Each convolutional block also performs downsampling through max-pooling operations with a filter of size 3. Squeeze-and-excitation (SE) attention modules are integrated after each pooling layer. The outputs are then flattened and given as input into the regression layer, containing two neurons, one for predicting the latitude and the other for predicting the longitude. Adam (Kingma and Ba, 2015) is used as the optimization algorithm, in an attempt to minimize the MSE loss. We trained our model on mini-batches of 128 samples for 100 epochs with early stopping, using a

learning rate of $5 \cdot 10^{-4}$. The network converged in 60 epochs, after observing no improvements for the last 7 epochs.

LSTM based on BERT embeddings. In conjunction with the LSTM, we fine-tuned a BERT model that is pre-trained on a German corpus, as detailed in Section 3.1. Thus, we input the data into a BERT layer, initialized with the corresponding pre-trained parameters. We set the maximum sequence length to 310, which is around the mean sequence length in the SMG-CH data set. The BERT layer is followed by two LSTM layers of size 128 each, both having *tanh* activation. We use dropout for regularization, randomly removing 20% of the neurons. We tried out various optimization algorithms such as Adam, RMSProp and stochastic gradient descent (SGD) with momentum. We obtained the best convergence using SGD with a momentum rate of 0.9 and a learning rate of $\alpha = 10^{-1}$. The training automatically ended after 18 epochs because of early stopping, as there were no more improvements registered for the loss value.

Extreme Gradient Boosting. We employed XGBoost as a meta-learner, training it over the predictions of all the other models. In our case, XGBoost provided the best results with the number of estimators set to 1000, a maximum depth of 10^{-4} and the learning rate $\alpha = 10^{-2}$.

4.3 Preliminary Results

In the development phase, as there was no metric specified in the description of the SMG task, we treated it as any other regression problem and evaluated the predictions in terms of both MAE and MSE. In Table 1, we present the results obtained by our four models on the development set. As the organizers released the ground-truth labels for the test set after the competition, we also include the MAE and the MSE on the test set, for reference.

Method	MAE		MSE	
	Development	Test	Development	Test
ν -SVR + string kernels	0.2306	0.2289	0.1066	0.1049
character-level CNN	0.2937	0.3123	0.1552	0.1633
LSTM + BERT embeddings	0.3594	0.3618	0.2226	0.2259
XGBoost ensemble	0.2234	0.2207	0.1043	0.1017

Table 1: Results in terms of Mean Absolute Error (MAE) and Mean Squared Error (MSE) obtained on the development set and on the test set by the proposed handcrafted, deep and ensemble algorithms. The reported MAE and MSE values represent the average value computed over the latitude and the longitude, both being expressed in degrees.

Considering the results presented in Table 1, it is clear that the algorithm that achieves the best MAE and MSE values is the ensemble based on the XGBoost meta-learner. This does not come as a surprise, as the ensemble combines the predictions of three individual models, each being based on a different type of features and a different learning model. While these aspects are complementary in theory, the results indicate that this is also the case in practice.

Additionally, we note that the performance achieved by ν -SVR comes close to the one attained by the ensemble, leaving behind the two neural models based on characters and fine-tuned BERT embeddings. This confirms the efficiency of string kernels over deep learning approaches observed in related works (Butnaru and Ionescu, 2019; Găman and Ionescu, 2020), which seems to be independent of the task to be solved.

Another comment regarding the results outlined in Table 1 is that the errors on the development set do not fall far from the ones obtained on the test set. However, the ensemble model as well as the ν -SVR based on string kernels obtain slightly lower errors at test time compared to the errors reported on the development set. The opposite seems to happen with the deep models, as both the LSTM based on BERT embeddings and the character-level CNN yield slightly higher errors on the test data than on the development data.

Every participant was allowed to make three submissions to compete against other shared task participants. Based on the results reported on the development set, we have decided to choose the character-

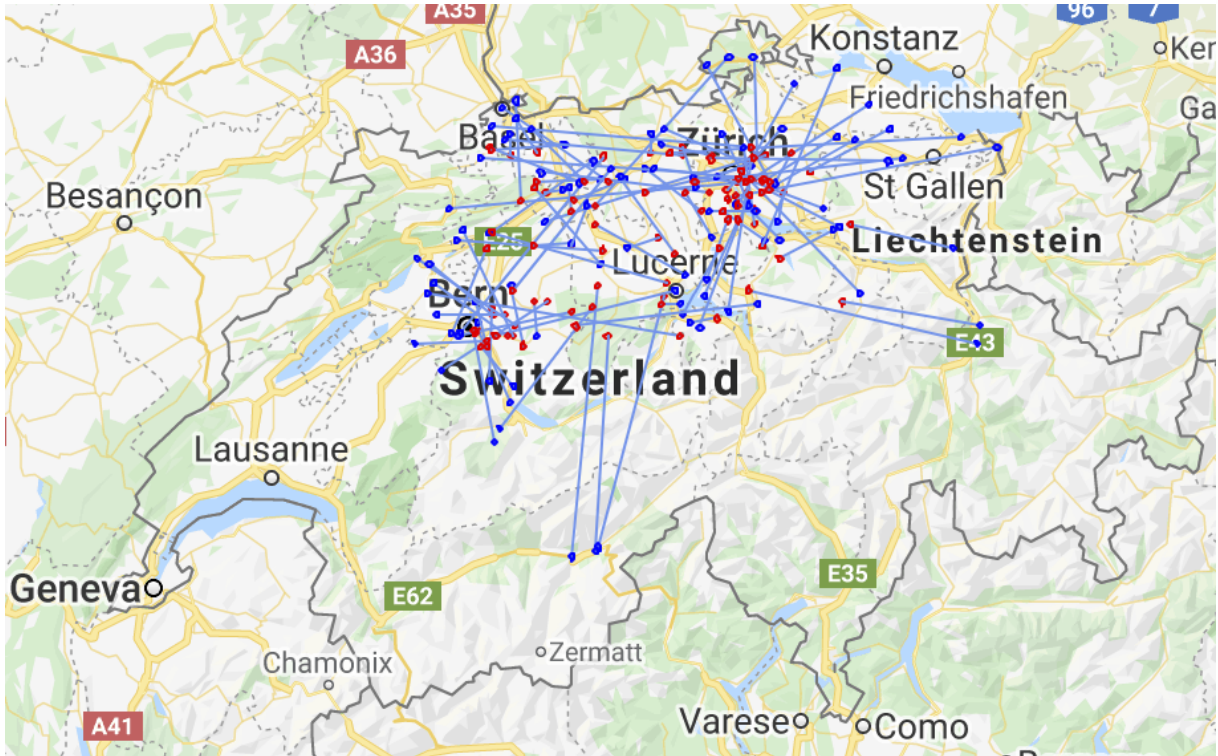


Figure 1: Distances between ground-truth locations (blue) and predicted locations (red) for a subset of 100 Swiss German Jodels randomly selected from the official test set. Best viewed in color.

level CNN, the ν -SVR based on string kernels and the XGBoost ensemble as candidates for the SMG-CH challenge. We excluded the LSTM based on BERT embeddings, since it attains the highest errors among the considered models.

4.4 Final Results

Method	Submission	Median distance	Mean distance	Clustering Accuracy
character-level CNN	#1	40.23	42.87	29.79%
ν -SVR + string kernels	#2	26.78	31.49	51.13%
XGBoost ensemble	#3	25.57	30.52	53.88%

Table 2: The final results of our team (UnibucKernel) obtained in the SMG-CH subtask, with the metrics picked by the organizers, oriented on clustering by city and on distances expressed in kilometers.

Table 2 shows our final results obtained on the test set, considering the metrics chosen by the organizers, which are oriented on distances (in kilometers) and on clustering accuracy. Considering the official metrics, our best submission placed us in the top six participants.

We observe that our best performing algorithm, namely the XGBoost ensemble combining the predictions of both deep and shallow methods based on various types of features, achieves a median distance of 25.57 km, a mean distance of 30.52 km and a clustering accuracy of 53.88%. Consistent with our findings on the development set, the ν -SVR based on string kernels does not fall far behind the XGBoost ensemble, obtaining a mean distance and median distance that is about one kilometer higher and a clustering accuracy that is nearly 2.75% lower. The deep character-level CNN seems significantly worse, with around 15 km and more than 10 km higher errors in terms of the median and the mean distances as compared to the other two submitted models, and around 21% lower clustering accuracy. In our opinion, these results stand proof that neural networks might not be the holy grail in every possible situation.

We believe that the proposed methods attain decent results, given the challenging nature of the problem

at hand. However, it is clear to us that they could benefit from some improvements, considering the values obtained in the final evaluation phase. One important step in this direction is to visualize the errors, at scale, on the map of Switzerland, for a better understanding of the patterns that our best performing algorithm drew with its predictions. Thus, we overlap the predicted and the ground-truth locations on the map of Switzerland, illustrating the result in Figure 1. The points depicted on the map are described by their 2D coordinates, latitude and longitude, as in the data set provided for the task. The ground-truth locations are colored in blue, while the predictions are illustrated in red. For each pair of ground-truth and predicted location, there is a line connecting the two points, giving us a better idea regarding the errors made by the XGBoost ensemble, in terms of distance. For the visualization presented in Figure 1, we have randomly selected a subset of 100 points from the test set, with non-overlapping ground-truth locations. We hereby notice that including all the data points from the test set would generate a visualization that is too cluttered and hard to understand. Hence, we opted for a smaller number of points for a better visualization experience. Considering the annotated map illustrated in Figure 1, we observe that our predictions tend to be clustered around the main cities in the German-speaking side of Switzerland, such as Zürich, Bern, Lucerne and Basel. This bias towards the mentioned cities might be induced by the data samples from the training set, likely not having a well-distributed variance in terms of the locations or the texts used in the learning process. We also observe that the ground-truth locations exhibit a higher variance than the predicted locations. One possible solution would be to manually adjust the variance of the predicted location to match the variance of the actual locations.

5 Conclusion

In the current work, we tackled the SMG-CH shared subtask of the 2020 VarDial Evaluation Campaign. We addressed this challenge from a shallow perspective, with handcrafted models such as a ν -SVR based on string kernels, as well as from a deep learning perspective, with neural models such as an LSTM based on BERT embeddings and a character-level CNN, respectively. Additionally, we combined the proposed models into an ensemble, employing the XGBoost meta-learner. We obtained our best results with the XGBoost ensemble, which benefits from complementary information from the handcrafted and deep models. We therefore brought one more proof regarding the effectiveness of ensemble learning in general, and of XGBoost, in particular. Another important conclusion is that our shallow model based on string kernels outperforms the two deep neural networks. We consider this as yet another indicator of the high discriminative power that string kernels can bring to a fairly standard learning model, i.e. the ν -SVR.

In future work, we aim to explore ways to improve our performance with respect to the metrics proposed by the shared task organizers. Currently, it seems that training the models to simply minimize the MSE or the MAE values is not effective, as our best model was significantly outperformed by the model proposed by the shared task organizers themselves.

Acknowledgements

This work was supported by a grant of the Romanian Ministry of Education and Research, CNCS - UEFISCDI, project number PN-III-P1-1.1-TE-2019-0235, within PNCDI III. This article has also benefited from the support of the Romanian Young Academy, which is funded by Stiftung Mercator and the Alexander von Humboldt Foundation for the period 2020-2022.

References

- Amr Ahmed, Liangjie Hong, and Alex J. Smola. 2013. Hierarchical geographical modeling of user locations from social media posts. In *Proceedings of WWW*, pages 25–36.
- Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. 2019. Character-Level Language Modeling with Deeper Self-Attention. In *Proceedings of AAAI*, pages 3159–3166.
- Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2015. Improved Transition-Based Parsing by Modeling Characters instead of Words with LSTMs. In *Proceedings of EMNLP 2015*, pages 349–59.

- James Bennett, Stan Lanning, et al. 2007. The Netflix Prize. In *Proceedings of KDD*, volume 2007, page 35.
- Frédéric Bilhaut, Thierry Charnois, Patrice Enjalbert, and Yann Mathet. 2003. Geographic reference analysis for geographic document querying. In *Proceedings of HLT-NAACL-GEOREF*, pages 55–62.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- Christopher J.C. Burges. 2010. From RankNet to LambdaRank to LambdaMART: An Overview. *Learning*, 11(23-581):81.
- Andrei Butnaru and Radu Tudor Ionescu. 2018a. UnibucKernel: A kernel-based learning method for complex word identification. In *Proceedings of BEA-13*, pages 175–183.
- Andrei M. Butnaru and Radu Tudor Ionescu. 2018b. UnibucKernel Reloaded: First Place in Arabic Dialect Identification for the Second Year in a Row. In *Proceedings of VarDial*, pages 77–87.
- Andrei M. Butnaru and Radu Tudor Ionescu. 2019. MOROCO: The Moldavian and Romanian Dialectal Corpus. In *Proceedings of ACL*, pages 688–698.
- Chih-Chung Chang and Chih-Jen Lin. 2002. Training ν -Support Vector Regression: Theory and Algorithms. *Neural Computation*, 14:1959–1977.
- Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A scalable tree boosting system. In *Proceedings of KDD*, pages 785–794.
- Zhiyuan Cheng, James Caverlee, and Kyumin Lee. 2010. You are where you tweet: a content-based approach to geo-locating twitter users. In *Proceedings of CIKM*, pages 759–768.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine Learning*, 20(3):273–297.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL*, pages 4171–4186.
- Junyan Ding, Luis Gravano, and Narayanan Shivakumar. 2000. Computing geographical scopes of web resources. In *Proceedings of VLDB*.
- Gabriel Doyle. 2014. Mapping dialectal variation by querying social media. In *Proceedings of EACL 2014*, pages 98–106.
- Harris Drucker, Christopher J.C. Burges, Linda Kaufman, Alex J. Smola, and Vladimir Vapnik. 1997. Support vector regression machines. In *Proceedings of NIPS*, pages 155–161.
- Jacob Eisenstein, Brendan O’Connor, Noah A Smith, and Eric Xing. 2010. A latent variable model for geographic lexical variation. In *Proceedings of EMNLP 2010*, pages 1277–1287.
- Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Kunihiko Fukushima. 1980. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202.
- Jan Gasthaus, Frank Wood, and Yee Whye Teh. 2010. Lossless Compression Based on the Sequence Memoizer. In *Proceedings of DCC*, page 337–345.
- Judith Gelernter and Nikolai Mushegian. 2011. Geo-parsing messages from microtext. *Transactions in GIS*, 15(6):753–773.
- Mariana-Iuliana Georgescu, Radu Tudor Ionescu, Nicolae-Catalin Ristea, and Nicu Sebe. 2020. Non-linear Neurons with Human-like Apical Dendrite Activations. *arXiv preprint arXiv:2003.03229*.
- Felix A. Gers, Jürgen Schmidhuber, and Fred Cummins. 2000. Learning to forget: Continual prediction with LSTM. *Neural Computation*, 12(10):2451–2471.
- 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.

- Mihaela Găman and Radu Tudor Ionescu. 2020. The Unreasonable Effectiveness of Machine Learning in Moldavian versus Romanian Dialect Identification. *journal=arXiv preprint arXiv:2007.15700*.
- Mihaela Găman, Dirk Hovy, Radu Tudor Ionescu, Heidi Jauhiainen, Tommi Jauhiainen, Krister Lindén, Nikola Ljubešić, Niko Partanen, Christoph Purschke, Yves Scherrer, and Marcos Zampieri. 2020. A Report on the VarDial Evaluation Campaign 2020. In *Proceedings of VarDial*.
- Bo Han, Paul Cook, and Timothy Baldwin. 2014. Text-based twitter user geolocation prediction. *Journal of Artificial Intelligence Research*, 49:451–500.
- Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, et al. 2014. Practical lessons from predicting clicks on ads at facebook. In *Proceedings of ADKDD*, pages 1–9.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780.
- Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, Jürgen Schmidhuber, et al. 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. *A Field Guide to Dynamical Recurrent Neural Networks*, pages 237–244.
- Liangjie Hong, Amr Ahmed, Siva Gurumurthy, Alex J. Smola, and Kostas Tsioutsoulouklis. 2012. Discovering geographical topics in the twitter stream. In *Proceedings of WWW*, pages 769–778.
- Dirk Hovy and Christoph Purschke. 2018. Capturing Regional Variation with Distributed Place Representations and Geographic Retrofitting. In *Proceedings of EMNLP*, pages 4383–4394.
- Yuan Huang, Diansheng Guo, Alice Kasakoff, and Jack Grieve. 2016. Understanding us regional linguistic variation with twitter data analysis. *Computers, Environment and Urban Systems*, 59:244–255.
- 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.
- Taylor Jones. 2015. Toward a description of African American Vernacular English dialect regions using “Black Twitter”. *American Speech*, 90(4):403–440.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-Aware Neural Language Models. In *Proceedings of AAAI*, pages 2741–2749.
- Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *Proceedings of EMNLP*, pages 1746–1751.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of ICLR*.
- Sheila Kinsella, Vanessa Murdock, and Neil O’Hare. 2011. “I’m eating a sandwich in Glasgow” modeling locations with tweets. In *Proceedings of SMUC*, pages 61–68.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of NIPS*, pages 1097–1105.
- Yann LeCun, Bernhard Boser, John S. Denker, Donnie Henderson, Richard E. Howard, Wayne Hubbard, and Lawrence D. Jackel. 1989. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4):541–551.

- Yann LeCun, Fu Jie Huang, and Leon Bottou. 2004. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of CVPR*, volume 2, pages II–104.
- Ping Li. 2010. Robust Logitboost and Adaptive Base Class (ABC) Logitboost. In *Proceedings of UAI*, pages 302–311.
- Michael D Lieberman, Hanan Samet, and Jagan Sankaranarayanan. 2010. Geotagging with local lexicons to build indexes for textually-specified spatial data. In *Proceedings of ICDE*, pages 201–212.
- Nikola Ljubešić, Tanja Samardžić, and Curdin Derungs. 2016. TweetGeo - A Tool for Collecting, Processing and Analysing Geo-encoded Linguistic Data. In *Proceedings of COLING*, pages 3412–3421.
- Huma Lodhi, John Shawe-Taylor, Nello Cristianini, and Christopher J.C.H. Watkins. 2001. Text Classification Using String Kernels. In *Proceedings of NIPS*, pages 563–569.
- Mihai Masala, Stefan Ruseti, and Traian Rebedea. 2017. Sentence selection with neural networks using string kernels. In *Proceedings of KES*, pages 1774–1782.
- Cristian Onose, Dumitru-Clementin Cercel, and Ștefan Trăușan-Matu. 2019. SC-UPB at the VarDial 2019 Evaluation Campaign: Moldavian vs. Romanian Cross-Dialect Topic Identification. In *Proceedings of VarDial*, pages 172–177.
- Marius Popescu and Radu Tudor Ionescu. 2013. The Story of the Characters, the DNA and the Native Language. In *Proceedings of BEA-8*, pages 270–278.
- 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.
- Teng Qin, Rong Xiao, Lei Fang, Xing Xie, and Lei Zhang. 2010. An efficient location extraction algorithm by leveraging web contextual information. In *Proceedings of GIS*, pages 53–60.
- Gianluca Quercini, Hanan Samet, Jagan Sankaranarayanan, and Michael D Lieberman. 2010. Determining the spatial reader scopes of news sources using local lexicons. In *Proceedings of GIS*, pages 43–52.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding with unsupervised learning. *Technical report, OpenAI*.
- Afshin Rahimi, Trevor Cohn, and Timothy Baldwin. 2017. A neural model for user geolocation and lexical dialectology. *arXiv preprint arXiv:1704.04008*.
- Stephen Roller, Michael Speriosu, Sarat Rallapalli, Benjamin Wing, and Jason Baldrige. 2012. Supervised text-based geolocation using language models on an adaptive grid. In *Proceedings of EMNLP*, pages 1500–1510.
- Dominic Rout, Kalina Bontcheva, Daniel Preoțiuc-Pietro, and Trevor Cohn. 2013. Where’s@ wally? A Classification Approach to Geolocating Users Based on their Social Ties. In *Proceedings of HT*, pages 11–20.
- Alex J. Smola and Bernhard Schölkopf. 2004. A tutorial on support vector regression. *Statistics and computing*, 14(3):199–222.
- Ilya Sutskever, James Martens, and Geoffrey Hinton. 2011. Generating Text with Recurrent Neural Networks. In *Proceedings of ICML*, pages 1017–1024.
- Benedikt Szmezcanyi. 2008. Corpus-based dialectometry: Aggregate morphosyntactic variability in british english dialects. *International Journal of Humanities and Arts Computing*, 2(1-2):279–296.
- Diana Tudoreanu. 2019. DTeam @ VarDial 2019: Ensemble based on skip-gram and triplet loss neural networks for Moldavian vs. Romanian cross-dialect topic identification. In *Proceedings of VarDial*, pages 202–208.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of NIPS*, pages 5998–6008.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. On the Practical Computational Power of Finite Precision RNNs for Language Recognition. In *Proceedings of ACL*, pages 740–745.
- Paul J. Werbos. 1988. Generalization of backpropagation with application to a recurrent gas market model. *Neural Networks*, 1(4):339–356.

- Martijn Wieling, John Nerbonne, and R. Harald Baayen. 2011. Quantitative social dialectology: Explaining linguistic variation geographically and socially. *PloS One*, 6(9):e23613.
- Benjamin Wing and Jason Baldridge. 2011. Simple supervised document geolocation with geodesic grids. In *Proceedings of ACL*, pages 955–964.
- Frank Wood, Cédric Archambeau, Jan Gasthaus, Lancelot James, and Yee Whye Teh. 2009. A Stochastic Memoizer for Sequence Data. In *Proceedings of ICML*, pages 1129–1136.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level Convolutional Networks for Text Classification. In *Proceedings of NIPS*, pages 649–657.