

Discriminating Similar Languages: Evaluations and Explorations

Cyril Goutte¹, Serge Léger¹, Shervin Malmasi², Marcos Zampieri^{3,4}

National Research Council (NRC), Canada¹, Macquarie University, Australia²
Saarland University, Germany³, German Research Center for Artificial Intelligence (DFKI), Germany⁴
firstname.lastname@nrc.ca, shervin.malmasi@mq.edu.au, marcos.zampieri@uni-saarland.de

Abstract

We present an analysis of the performance of machine learning classifiers on discriminating between similar languages and language varieties. We carried out a number of experiments using the results of the two editions of the Discriminating between Similar Languages (DSL) shared task. We investigate the progress made between the two tasks, estimate an upper bound on possible performance using ensemble and oracle combination, and provide learning curves to help us understand which languages are more challenging. A number of difficult sentences are identified and investigated further with human annotation.

Keywords: language identification, language varieties, evaluation

1. Introduction

Discriminating between similar languages and language varieties is one of the main challenges of state-of-the-art language identification systems (Tiedemann and Ljubešić, 2012). Closely-related languages such as Indonesian and Malay or Croatian and Serbian are very similar both at their spoken and at their written forms making it difficult for systems to discriminate between them. Varieties of the same language, e.g. Spanish from South America or Spain, are even more difficult to detect than similar languages. Nevertheless, in both cases, recent work has shown that it is possible to train algorithms to discriminate between similar languages and language varieties with high accuracy (Goutte et al., 2014; Malmasi and Dras, 2015b).

This study looks in more detail into the features that help algorithms discriminating between similar languages, taking into account the results of two recent editions of the Discriminating between Similar Languages (DSL) shared task (Zampieri et al., 2014; Zampieri et al., 2015b). The analysis and results of this paper complement the information presented in the shared task reports and provide novel and important information for researchers and developers interested in language identification and particularly in the problem of discriminating between similar languages.

2. Related Work

Language identification in written texts is a well-established research topic in computational linguistics. Interest in the task is evidenced by early n-gram-based approaches (Dunning, 1994; Grefenstette, 1995) to more recent studies (Brown, 2013; Lui et al., 2014a; Brown, 2014; Simões et al., 2014). The interest in the discrimination of similar languages, language varieties, and dialects is more recent but it has been growing in the past few years. Examples of studies include the cases of Malay and Indonesian (Ranaivo-Malançon, 2006), Chinese varieties (Huang and Lee, 2008), South Slavic languages (Ljubešić et al., 2007; Ljubešić and Kranjčić, 2015), Portuguese varieties (Zampieri and Gebre, 2012), Spanish varieties (Zampieri et al., 2013; Maier and Gómez-Rodríguez, 2014), English varieties (Lui and Cook, 2013), Persian and Dari (Malmasi

and Dras, 2015a), Romanian dialects (Ciobanu and Dinu, 2016), and a number of studies on Arabic dialects (Elfardy and Diab, 2014; Zaidan and Callison-Burch, 2014; Tillmann et al., 2014; Malmasi et al., 2015a)

A number of shared tasks on language identification have been organized in the recent years ranging from general-purpose language identification (Baldwin and Lui, 2010) to more specific challenges such as the TweetLID shared task which focused on Twitter data (Zubiaga et al., 2014; Zubiaga et al., 2015), the shared task on Language Identification in Code-Switched Data (Solorio et al., 2014), and the two editions of the discriminating between similar languages (DSL) shared task. To our knowledge, however, no comprehensive analysis of the kind we are proposing in this paper has been carried out on the results obtained in a language identification shared task and our work fills this gap. The most similar analysis was applied to Native Language Identification (NLI)¹ using the 2013 NLI shared task dataset (Malmasi et al., 2015b).

In the next sections we present the systems that participated in the two editions of the DSL shared task.

2.1. DSL Shared Task 2014

The first edition of the DSL task was organized in 2014 within the scope of the workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects (VarDial) co-located with COLING. The organizers compiled and released a new dataset for this purpose, which they claim to be the first resource of its kind (Tan et al., 2014). The dataset is entitled *DSL Corpus Collection*, or DSLCC, and it includes short excerpts from journalistic texts from previously released corpora and repository.² Texts in the DSLCC v. 1.0 were written in thirteen languages or language varieties and divided into the following six groups: Group A (Bosnian, Croatian, Serbian), Group B (Indonesian, Malay), Group C (Czech, Slovak), Group D (Brazilian Portuguese, European Portuguese), Group E (Peninsu-

¹This task focuses on identifying the mother tongue of a learner writer based on stylistic cues; all the texts are in the same language (Malmasi and Dras, 2014; Malmasi and Dras, 2015c).

²See Tan et al. (2014) for a complete list of sources.

lar Spanish, Argentine Spanish), and Group F³ (American English, British English).

In the 2014 edition, eight teams participated and submitted results to the DSL shared task (eight teams in the closed and two teams in the open submission). Five of these teams wrote system description papers. The complete shared task report is available in Zampieri et al. (2014). We summarize the results in Table 1 in terms of accuracy (best performing entries displayed in bold).

Team	Closed	Open	System Description
NRC-CNRC	95.7	-	(Goutte et al., 2014)
RAE	94.7	-	(Porta and Sancho, 2014)
UMich	93.2	85.9	(King et al., 2014)
UniMelb-NLP	91.8	88.0	(Lui et al., 2014b)
QMUL	90.6	-	(Purver, 2014)
LIRA	76.6	-	-
UDE	68.1	-	-
CLCG	45.3	-	-

Table 1: DSL Shared Task 2014 - Accuracy results. Teams ranked by their results in the closed submission.

In the closed submission track the best performance was obtained by the NRC-CNRC (Goutte et al., 2014) team, which used a two-step classification approach to predict first the language group of the text, and subsequently the language. Both NRC-CNRC (Goutte et al., 2014) and QMUL (Purver, 2014), ranked 5th used linear support vector machines (SVM) classifiers with words and characters as features.

Two teams used information gain to estimate the best features for classification, UMich (King et al., 2014) and UniMelb-NLP (Lui et al., 2014b). These two teams were also the only ones teams which compiled and used additional training material to compete in the open submission track. As can be seen in Table 1, the performance of open submissions were worse than the closed submissions. Accuracy dropped from 93.2% to 85.9% for UMich, and from 91.8% to 88.0% for UniMelb-NLP. This is probably because along with diatopic variation, systems learn properties of datasets that are often topic or genre specific. Therefore, part of what was learned from the additional training corpora was not helpful for predictions on the test set.

The RAE team (Porta and Sancho, 2014) proposed an approach based on ‘white lists’ of words used exclusively in a given language or language variety and their closed submission ranked 2nd.

2.2. DSL Shared Task 2015

The 2015 edition of the DSL shared task was organized within the scope of the Joint Workshop on Language Technology for Closely Related Languages, Varieties and Dialects (LT4VarDial) co-located with RANLP.

For the DSL 2015, organizers released version 2.0 of the DSLCC which contained the same set of languages and language varieties as version 1.0 in groups A to E. The two

³There were many cases of republication (e.g. British texts republished by an American newspaper and tagged as American by the original sources) that made the task for this language group unfeasible. (Zampieri et al., 2014)

main modifications between the two versions are the exclusion of group F (British and American English) and the inclusion of group G (Bulgarian and Macedonian).⁴ A new addition in the DSL 2015 is the use of two test sets (A and B). In test set A instances are presented exactly as they appear in newspaper texts whereas in test set B named entities were substituted by placeholders. According to the organizers, the release of test set B aimed to evaluate the extent to which named entities influence classification performance. Ten teams submitted their results and eight of them published system description papers. Results of the DSL 2015 are described in detail in Zampieri et al. (2015b) In Table 2 we summarize the results obtained by all teams using test sets A and B in both open and close submissions. The best results for each submission type are displayed in bold.

Team	Closed	Open	System Description
Test Set A			
MAC	95.54	-	(Malmasi and Dras, 2015b)
MMS	95.24	-	(Zampieri et al., 2015a)
NRC	95.24	95.65	(Goutte and Léger, 2015)
SUKI	94.67	-	(Jauhiainen et al., 2015)
BOBICEV	94.14	-	(Bobicev, 2015)
BRUNIBP	93.66	-	(Ács et al., 2015)
PRHLT	92.74	-	(Franco-Salvador et al., 2015)
INRIA	83.91	-	-
NLEL	64.04	91.84	(Fabra-Boluda et al., 2015)
OSEVAL	-	76.17	-
Test Set B			
MAC	94.01	-	(Malmasi and Dras, 2015b)
SUKI	93.02	-	(Jauhiainen et al., 2015)
NRC	93.01	93.41	(Goutte and Léger, 2015)
MMS	92.78	-	(Zampieri et al., 2015a)
BOBICEV	92.22	-	(Bobicev, 2015)
PRHLT	90.80	-	(Franco-Salvador et al., 2015)
NLEL	62.78	89.56	(Fabra-Boluda et al., 2015)
OSEVAL	-	75.30	-

Table 2: DSL Shared Task 2015 - Accuracy results for open and closed submissions using test sets A and B. Teams ranked by their results in the closed submission.

We observe that for all systems, performance dropped from test set A (complete texts) to test set B (name entities removed). This was the expected outcome. However, performance is, in most cases, only 1 or 2 percentage points lower which is not as great as one could expect. This means that even if all place and person names, which are to a large extent country specific, are removed from Brazilian texts, systems are still able to distinguish them from, for example, Portuguese texts with very high accuracy. Another interesting aspect to observe is that, unlike in the DSL 2014, the use of additional training material helped system performance as can be observed in the NRC and NLEL submissions. This is mainly due to corpus comparability, as in 2015 team were allowed to use the DSLCC v. 1.0 and in 2014 they did not have such a resource available, and had to acquire additional, unrelated material.

The best system in the closed submission for test set A and

⁴In the 2015 edition, organizers did not use language group names as in the 2014 edition. We use them for both editions in this paper for the sake of clarity and consistency.

B was MAC (Malmasi and Dras, 2015b) which proposed an ensemble of SVM classifiers for this task. Two other SVM-based approaches were tied in 2nd for test set A, one by the NRC team (Goutte and Léger, 2015) and MMS (Zampieri et al., 2015a), which experimented with three different approaches and obtained the best results combining TF-IDF and an SVM classifier previously used for native language identification (Gebre et al., 2013). The NRC team included members of NRC-CNRC, winners of the DSL closed submission track in 2014. Both in 2014 and in 2015 they used a two-stage classification approach to predict first the language group and then the language within the predicted group. Two other teams used two-stage classification approaches: NLEL (Fabra-Boluda et al., 2015) and BRUniBP (Ács et al., 2015).

A number of computational techniques have been explored in the DSL 2015 including token-based backoff by SUKI team (Jauhiainen et al., 2015), prediction by partial matching (PPM) by BOBICEV (Bobicev, 2015), and word and sentence vectors by PRHLT (Franco-Salvador et al., 2015).

3. Methods

In the next subsections we describe the methodology behind our 4 experiments as well as the data used.

3.1. Data

All experiments reported here are performed on the DSL Corpus Collection (DSLCC) versions 1.0. and 2.0. (Tan et al., 2014). Both versions cover five groups of two to three languages or varieties each (groups A-E, Table 3). The 2015 collection adds Bulgarian and Macedonian (group G) plus sentences from “Other” languages. In a couple of experiments (Sections 3.3. & 4.2.) we use the output of the 22 entries submitted to the 2015 Shared Task.⁵

Grp	Language/variety	2014		2015	
		Train	Test	Train	Test
A	Bosnian	20k	1000	20k	1000
	Croatian	20k	1000	20k	1000
	Serbian	20k	1000	20k	1000
B	Indonesian	20k	1000	20k	1000
	Malaysian	20k	1000	20k	1000
C	Czech	20k	1000	20k	1000
	Slovak	20k	1000	20k	1000
D	Brazil Portuguese	20k	1000	20k	1000
	European Portuguese	20k	1000	20k	1000
E	Argentine Spanish	20k	1000	20k	1000
	Peninsular Spanish	20k	1000	20k	1000
G	Bulgarian	-	-	20k	1000
	Macedonian	-	-	20k	1000
X	Others	-	-	20k	1000

Table 3: Number of sentences in the DSLCC v. 1.0 and 2.0.

3.2. Progress Test

In our first experiment, we evaluate the improvements achieved from one shared task to the other. For that purpose, we measure the performance, on the 2014 and 2015 test data, of three systems representative of the top performance in both years:

- the top 2014 system, NRC-closed-2014 (Goutte et al., 2014);
- the top 2015 closed task system, MAC-closed-2015 (Malmasi and Dras, 2015b);
- the top 2015 open task system, NRC-open-2015 (Goutte and Léger, 2015).

The 2015 shared task had two key additions: a new group of close languages (Bulgarian/Macedonian, group G) and data from other languages (group X). The 2015 results were measured on all groups, but the 2014 system was not trained to recognize either group G or group X data. As a consequence, in addition to the full 2014 and 2015 test sets, we evaluated performance on the subset of the 2015 test set that contains the groups in the 2014 shared task, i.e. groups A to E (5 groups and 11 variants).

3.3. Ensemble and Oracle

An interesting research question for this task is to measure the upper-bound on accuracy. This can be measured by treating each shared task submission as an independent system and combining the results using ensemble fusion methods such as a plurality voting or oracle. This type of analysis has previously been shown to be informative for the similar task of Native Language Identification (Malmasi et al., 2015b). Moreover, this analysis can also help reveal interesting error patterns in the submissions.

Following the approach of Malmasi et al. (2015b), we apply the following combination methods to the data.

Plurality Voting: This is the standard combination strategy that selects the label with the highest number of votes, regardless of the percentage of votes it received (Polikar, 2006). This differs from a *majority* vote combiner where a label must obtain over 50% of the votes.

Oracle: An oracle is a type of multiple classifier fusion method that can be used to combine the results of an ensemble of classifiers which are all used to classify a dataset. The oracle will assign the correct class label for an instance if at least one of the constituent classifiers in the system produces the correct label for that data point. This method has previously been used to analyze the limits of majority vote classifier combination (Kuncheva et al., 2001). It can help quantify the *potential* upper limit of an ensemble’s performance on the given data and how this performance varies with different ensemble configurations and combinations.

Accuracy@N: To account for the possibility that a classifier may predict the correct label by chance (with a probability determined by the random baseline) and thus exaggerate the oracle score, an Accuracy@N combiner has been proposed (Malmasi et al., 2015b). This method is inspired by the “Precision at *k*” metric from Information Retrieval (Manning et al., 2008) which measures precision at fixed low levels of results (*e.g.* the top 10 results). Here, it is an extension of the Plurality vote combiner where instead of selecting the label with the highest votes, the labels are ranked by their vote counts and an instance is correctly classified if the true label is in the top *N* ranked candidates.⁶

⁵<https://github.com/Simdiva/DSL-Task>.

⁶In case of ties we choose randomly from the labels with the same number of votes.

Another way to view it is as a more restricted version of the Oracle combiner that is limited to the top N ranked candidates in order to minimize the influence of a single classifier having chosen the correct label by chance. In this study we experiment with $N = 2$ and 3. We also note that setting $N = 1$ is equivalent to the Plurality voting method. Results from the above combiners are compared to a random baseline and to the best system in the shared task.

3.4. Learning Curves

Learning curves are an important tool to understand how statistical models learn from data. They show how the models behave, in terms of performance, with increasing amounts of data. In order to compute learning curves for the DSL task, we picked a simple model that is easy to train and performs close to the top systems.

From the full training set, we subsample data at various sample sizes. In order to keep the training data balanced, we sample the same amount N_s of examples from each language variant. In our setup, we use $N_s = 20,000$ (full training set), 10,000, 5000, 2000, 1000, 500, 200 and 100. For each subsample, we train a statistical model, and test it on the official 2015 test set. We replicate this experiment 10 times at each sample size, except for the full training set. This helps us estimate the expected performance at each sample size, as well as error bars on the expectation.

3.5. Manual Annotation

Finally, to make this evaluation even more comprehensive, we also conducted a human evaluation experiment on some of the misclassified instances. We asked human annotators to assign the correct language or language variety of each sentence for the most difficult language groups, namely group A (Bosnian, Croatian, and Serbian), D (Brazilian and European Portuguese), and E (Argentinian and Peninsular Spanish). In this experiment we included all instances that were misclassified by the Oracle (i.e. no submission got right). For groups D and E, we added sentences that were incorrectly classified by the plurality vote method as well amounting to twelve instances per group.

Such analyses of misclassifications can provide further insights and help better understand the difficulties of the task. Ács et al. (2015) showed that for 52 misclassified Portuguese instances, only 22 have been labeled correctly by the annotators with low inter-annotator agreement. We report the results obtained by the manual evaluation step in Section 4.4..

4. Results

4.1. Progress Test

Table 4 displays the results of the progress test. The 2014 system is evaluated on the 2014 test set and 2015 progress set containing 5 groups and 11 languages. The 2015 systems are evaluated on the full 2014 and 2015 test sets and the 2015 progress set.

Table 4 shows that, when measured on the 2014 test set, the increase in performance of the 2015 systems, although significant, is modest (+3-.6%). It should be noted however that the 2014 system is the only one for which the training data exactly matches the test data. Somewhat surprisingly,

System	2014	2015	
	test	progress	test
NRC-closed-2014	95.70	90.15	-
MAC-closed-2015	96.31	94.33	95.54
NRC-open-2015	96.04	94.48	95.65

Table 4: Progress test for 2014 and 2015 systems. “progress” is the 2015 test set without group G and X.

MAC-closed-2015, the best 2015 closed track submission, which is trained only on the 2015 training set, performs slightly better than NRC-open-2015, which was trained on both 2014 and 2015 data, and would therefore be expected to perform better on 2014 data.

Year-to-year improvements are more evident from the results obtained on the progress test. We note that average performance on the progress test is lower than on the full 2015 test set. This is due to the omission of groups G and X, on which 2015 systems performed very well. Looking at the performance of the best 2014 system, it is apparent that the 2015 progress set was harder than the 2014 test set. This is most likely due to the much shorter sentences providing less evidence for ngram statistics. On the progress set, the 2015 systems outperform the 2014 system by more than 4%. This indicates that the 2014 system suffers from the mismatch in data, and suggests a large year-to-year improvement in performance on shorter sentences.

4.2. Ensemble and Oracle

The 22 entries in the shared task (normal test set) were combined in various ensembles and the results are shown below in Table 5. We observe that a plurality vote among all the entries yields only a very small improvement over the best single system (Malmasi and Dras, 2015b).

	Accuracy (%)
Random Baseline	7.14
Shared Task Best	95.54
Plurality Vote	96.04
Oracle	99.83
Accuracy@3	99.83
Accuracy@2	99.47

Table 5: Ensemble results on the DSL 2015 shared task systems. The ensemble consists of all 22 submission from all teams.

The oracle results, however, are substantially higher than the voting ensemble and close to 100% accuracy. The accuracy@2 and accuracy@3 results are almost identical to the full oracle, suggesting that almost all of the errors are the result of a confusion between the top 2-3 results. This is due to the fact that DSL errors are almost always within a group, i.e. between 2 or 3 variants. As shown in the learning curves in Figure1, group prediction reaches perfect performance using relatively few examples, so the remaining confusions are always within a group of languages or variants. This differs from results observed for Native Language Identification where there is a large difference between the oracle and accuracy@2 results.

4.3. Learning Curves

For simplicity, learning curves were obtained for a simple system trained only on character 6-grams. This is essentially NRC’s first 2015 run, which performed 0.7% below the top system, trained using various training set sizes.

Figure 1 shows learning curves for the average group and language classification performance (top) and within each group (bottom). The group-level curve (dashed, top) shows that predicting the group is done perfectly from around 1000 examples per language. Average language-level performance is lower and still increasing at 20k examples per language.

Looking closer at discrimination performance for each group, we see that performance for groups C, G is essentially perfect as early as 100-500 examples. This suggests that discriminating Czech from Slovak and Bulgarian from Macedonian is easy. In those cases, it may be more challenging to investigate side issues such as robustness to changes in source, genre, or regional proximity. Group B was clearly harder to learn, but there is little room for improvement above ~10,000 examples/language.

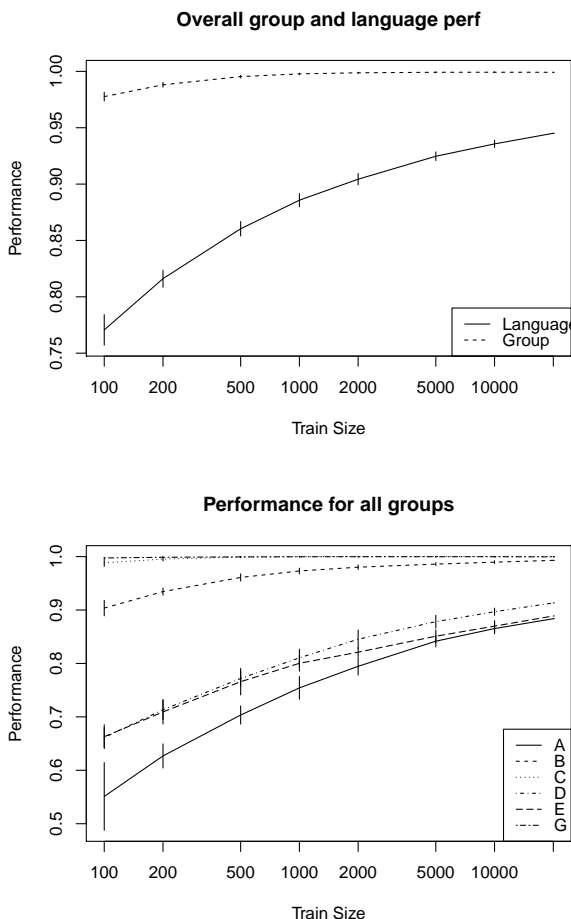


Figure 1: Learning curves for average group and language prediction (top), and within group performance (bottom).

Groups A, D and E display more interesting learning curves. Learning still takes place for the full training set. The rate of progress slows down, but doesn’t seem to plateau. This is reflected on the average performance:

the last doubling of the training data (right of the curves) brings around 1.6-1.9 additional percent of accuracy, while the first brought 4.6-7.6% increase in performance. This suggests that, despite diminishing returns, bringing more data covering languages in these three groups would still improve prediction accuracy.

4.4. Manual Annotation

As noted in Section 4.2., the oracle achieves an accuracy of 99.83%, leaving only 24 misclassified sentences: sixteen from group A, three from group D, and five from group E. We included most of these instances in a manual annotation experiment providing twelve instances to several native speaker of these languages and asking them to assign the correct language or language variety of each text.

In our experimental setting, we made sure that the annotators were not exclusively speakers of one of the languages or varieties of the group. However, a perfect balance in participation between languages was very difficult to be obtained. Within a group of two (or three) languages or varieties, native speaker’s perception of whether a given text belongs to his own language or not may vary according to the person’s own language. We discuss this issue later in this section taking group D (Brazilian and European Portuguese) an example.

As to the participants, for group A we asked six annotators (one Bosnian, two Croatians and three Serbians); for group D we had ten annotators participating (eight Brazilians and two Portuguese); and finally for Group E, seventeen annotators participated (fifteen Argentinians and two Spaniards). In Table 6 we report the percentage of times that instances 1 to 12 from which group were correctly annotated by the native speakers along with the maximum, minimum and mean performance of the annotators.

Group	A	D	E
Classes	3	2	2
Annotators	6	10	17
% Correct Inst. ID 1	16.66	30.00	35.30
% Correct Inst. ID 2	0	80.00	29.41
% Correct Inst. ID 3	16.66	80.00	58.82
% Correct Inst. ID 4	0	60.00	47.05
% Correct Inst. ID 5	33.33	50.00	35.29
% Correct Inst. ID 6	0	30.00	76.47
% Correct Inst. ID 7	0	70.00	76.47
% Correct Inst. ID 8	66.66	80.00	11.76
% Correct Inst. ID 9	16.66	90.00	35.29
% Correct Inst. ID 10	16.66	60.00	100
% Correct Inst. ID 11	33.33	80.00	82.35
% Correct Inst. ID 12	0	100	70.59
Best Annotator Accuracy	25.00	91.66	83.33
Mean Annotator Accuracy	16.66	67.50	54.90
Worst Annotator Accuracy	8.33	25.00	16.66

Table 6: Manual Evaluation Results for Groups A, D and E.

The group A (Bosnian, Croatian and Serbian) was the most difficult for the annotators. The language of five out of twelve instances was not correctly assigned by any of the six annotators. This is firstly explained because group A

is the only group containing three languages. However, we also noted that classifiers showed very high degree of confusion when discriminating between Bosnian and Croatian texts. All instances from this group misclassified by the oracle, and therefore include in this experiment, were either Bosnian or Croatian. The shared task organizers and our paper assume that all gold labels are correct. It might be the case, however, that a few labels in group A contain errors. In case there are incorrect reference labels, these errors would surely negatively impact both classifiers' and humans' performance. This is a possibility that we cannot confirm nor disregard. For pragmatic reasons we rely on the sources that were compiled for the DSLCC.⁷

Albeit still challenging, the task proved to be more feasible for group E than for group A. The average performance of the annotators was slightly above the 50% baseline. One of the seventeen annotators was able to correctly assign the language of the texts ten out of twelve times, achieving 83.33 accuracy. For all the three groups studied, the only two cases in which all annotators correctly assigned the language of an instance was group D ID 12 and group E ID 10. To exemplify, the latter is the following Peninsular Spanish text:

- (1) Entonces lo entiendo todo. La prensa no tiene razón. No estamos en guerra, ni falta que hace. Esta especie de bronca continua es una ilusión, una imagen grotesca que proyectan los medios, pero no es real. Por el amor de Dios, que no os pase como a mí, que meriendo paellas de orfidales todos los días. En ocasiones veo Truebas, sí, pero mis amigos me ponen en mi sitio. No nos dejemos engañar.

The best results were obtained by the annotators of group D (Brazilian and European Portuguese). The average performance of the annotators was 17.50 percentage points above the baseline. This corroborates the findings of Zampieri and Gebre (2012) who showed that due to differences in spelling and lexical variation, Brazilian and Portuguese texts can be discriminated automatically with almost perfect performance (researchers report 99.8% accuracy).

One interesting finding of this experiment is that the competence of identifying whether a given text comes from Brazilian or European Portuguese does not seem to be the same for speakers of these two varieties. Performance highly depends on the lexical variation included. To exemplify, next we present the text included as group D ID 3:

- (2) O médio Bruno Neves tem 25 anos e jogou a última época no Grêmio, do Brasil, tendo também alinhado já no Fluminense e no Cruzeiro.

The eight Brazilians were unanimous in assigning this sentence as European Portuguese, whereas the two European Portuguese speakers assigned them as Brazilian Portuguese.

⁷As mentioned in Section 2.1. incorrect tags in the gold data were present in the American and British English dataset of the DSLCC v 1.0 (see Zampieri et al. (2014) for a discussion).

This is a particularly interesting example because it talks about three Brazilian football clubs: *Grêmio*, *Fluminense*, and *Cruzeiro*. Thematically, this sentence is much more likely to be published in a Brazilian newspaper than in a Portuguese one, and due to the influence of these named entities the sentence was misclassified. However, it features terms that are exclusively used in European Portuguese such as *médio* (BR: *meia* or *meio-campista*, EN: *midfielder*) and *época* (BR: *temporada*, EN: *season*).⁸ For Brazilians, these words were probably an indication that the text was not written in Brazilian Portuguese. On the other hand, European Portuguese speakers who were not aware that these two words are not used in Brazil, probably were influenced by the club names to assume that the sentence is from Brazil.

To sum up, manual annotation for this task is by no means trivial which also explains the difficulty that algorithms have in discriminating between similar languages. We confirm that named entities play an important role in this task and that they can influence not only the performance of algorithms but also the performance of human annotators. Finally, a general tendency we observed is that it is easier to identify an instance that is not from the speaker's own language than the opposite. Our results indicate that humans are better in telling what is not a text written in their own language or variety than telling what it is. We would like to investigate this phenomenon in the future using more annotated data.

5. Conclusion and Future Work

This paper presents a comprehensive evaluation of state-of-the-art language identification systems trained to recognize similar languages and language varieties using the results of the first two DSL shared tasks. We evaluate the progress made from one edition of the shared task to the next. Using plurality voting and oracle, we estimate an upper bound on the achievable performance, and identify some particularly challenging sentences. We show learning curves that help us identify how the task is learned and which groups of languages may need more attention.

Finally, we propose an experiment with native speakers of the three most challenging language groups, group A (Bosnian, Croatian, and Serbian), group D (Brazilian and European Portuguese), and group E (Argentinian and Peninsular Spanish). Our results suggest that humans also find it difficult discriminating between similar languages and language varieties. In future work we would like to investigate human performance in this task focusing on two aspects: 1) how native speakers identify language variation; 2) which words, expressions or syntactic structures are the most discriminating features of a given language or variety according to the speakers of that language. Both of these aspects will provide us new insights into language variation that can be used for linguistic analysis as well as to improve computational methods to discriminate between similar languages.

⁸See Soares da Silva (2010) for a study on lexical variation involving Brazilian and European Portuguese football terms.

Acknowledgements

We would like to thank the human annotators who participated in the manual evaluation experiment. We also thank the DSL shared task organizers for providing the data we used in this paper.

Bibliographical References

- Ács, J., Grad-Gyenge, L., and de Rezende Oliveira, T. B. R. (2015). A two-level classifier for discriminating similar languages. In *Proceedings of the LT4VarDial Workshop*.
- Baldwin, T. and Lui, M. (2010). Multilingual language identification: ALTW 2010 shared task data. In *Proceedings of Australasian Language Technology Workshop*.
- Bobicev, V. (2015). Discriminating between similar languages using ppm. In *Proceedings of the LT4VarDial Workshop*.
- Brown, R. (2013). Selecting and weighting n-grams to identify 1100 languages. In *Proceedings of TSD*.
- Brown, R. D. (2014). Non-linear mapping for improved identification of 1300+ languages. In *Proceedings of EMNLP*.
- Ciobanu, A. M. and Dinu, L. P. (2016). A computational perspective on Romanian dialects. In *Proceedings of LREC*.
- Dunning, T. (1994). Statistical identification of language. Technical report, Computing Research Lab - New Mexico State University.
- Elfardy, H. and Diab, M. T. (2014). Sentence level dialect identification in Arabic. In *Proceedings of ACL*.
- Fabra-Boluda, R., Rangel, F., and Rosso, P. (2015). NLEL UPV autoritas participation at Discrimination between Similar Languages (DSL) 2015 shared task. In *Proceedings of the LT4VarDial Workshop*.
- Franco-Salvador, M., Rosso, P., and Rangel, F. (2015). Distributed representations of words and documents for discriminating similar languages. In *Proceedings of the LT4VarDial Workshop*.
- Gebre, B. G., Zampieri, M., Wittenburg, P., and Heskens, T. (2013). Improving native language identification with tf-idf weighting. In *Proceedings of the 8th BEA workshop*.
- Goutte, C. and Léger, S. (2015). Experiments in discriminating similar languages. In *Proceedings of the LT4VarDial Workshop*.
- Goutte, C., Léger, S., and Carpuat, M. (2014). The NRC system for discriminating similar languages. In *Proceedings of the VarDial Workshop*.
- Grefenstette, G. (1995). Comparing two language identification schemes. In *Proceedings of JADT 1995, 3rd International Conference on Statistical Analysis of Textual Data*, Rome.
- Huang, C. and Lee, L. (2008). Contrastive approach towards text source classification based on top-bag-of-word similarity. In *Proceedings of PACLIC*.
- Jauhiainen, T., Jauhiainen, H., and Lindén, K. (2015). Discriminating similar languages with token-based backoff. In *Proceedings of the LT4VarDial Workshop*.
- King, B., Radev, D., and Abney, S. (2014). Experiments in sentence language identification with groups of similar languages. In *Proceedings of the VarDial Workshop*.
- Kuncheva, L. I., Bezdek, J. C., and Duin, R. P. (2001). Decision templates for multiple classifier fusion: an experimental comparison. *Pattern Recognition*, 34(2):299–314.
- Ljubešić, N. and Kranjčić, D. (2015). Discriminating between closely related languages on twitter. *Informatica*, 39(1).
- Ljubešić, N., Mikelic, N., and Boras, D. (2007). Language identification: How to distinguish similar languages? In *Proceedings of the 29th International Conference on Information Technology Interfaces*.
- Lui, M. and Cook, P. (2013). Classifying English documents by national dialect. In *Proceedings of Australasian Language Technology Workshop*.
- Lui, M., Lau, J. H., and Baldwin, T. (2014a). Automatic detection and language identification of multilingual documents. *Transactions of the Association for Computational Linguistics*, 2:27–40.
- Lui, M., Letcher, N., Adams, O., Duong, L., Cook, P., and Baldwin, T. (2014b). Exploring methods and resources for discriminating similar languages. In *Proceedings of VarDial*.
- Maier, W. and Gómez-Rodríguez, C. (2014). Language variety identification in Spanish tweets. In *Proceedings of the LT4CloseLang Workshop*.
- Malmasi, S. and Dras, M. (2014). Language Transfer Hypotheses with Linear SVM Weights. In *Proceedings of EMNLP*.
- Malmasi, S. and Dras, M. (2015a). Automatic Language Identification for Persian and Dari texts. In *Proceedings of PACLING 2015*, pages 59–64.
- Malmasi, S. and Dras, M. (2015b). Language identification using classifier ensembles. In *Proceedings of the LT4VarDial Workshop*.
- Malmasi, S. and Dras, M. (2015c). Multilingual Native Language Identification. In *Natural Language Engineering*.
- Malmasi, S., Refaee, E., and Dras, M. (2015a). Arabic Dialect Identification using a Parallel Multidialectal Corpus. In *Proceedings of PACLING 2015*, pages 209–217, Bali, Indonesia, May.
- Malmasi, S., Tetreault, J., and Dras, M. (2015b). Oracle and human baselines for native language identification. In *Proceedings of the BEA workshop*.
- Manning, C. D., Raghavan, P., and Schütze, H. (2008). Evaluation in information retrieval. In *Introduction to Information Retrieval*, pages 151–175. Cambridge university press Cambridge.
- Polikar, R. (2006). Ensemble based systems in decision making. *Circuits and systems magazine, IEEE*, 6(3):21–45.
- Porta, J. and Sancho, J.-L. (2014). Using maximum entropy models to discriminate between similar languages and varieties. In *Proceedings of the 1st Workshop on Applying NLP Tools to Similar Languages, Varieties and Dialects (VarDial)*, Dublin, Ireland.
- Purver, M. (2014). A simple baseline for discriminating similar language. In *Proceedings of the 1st Workshop on*

- Applying NLP Tools to Similar Languages, Varieties and Dialects (VarDial)*, Dublin, Ireland.
- Ranaivo-Malançon, B. (2006). Automatic identification of close languages - case study: Malay and Indonesian. *ECTI Transactions on Computer and Information Technology*, 2:126–134.
- Simões, A., Almeida, J. J., and Byers, S. D. (2014). Language identification: a neural network approach. *Proceedings of Slate*.
- Soares da Silva, A. (2010). Measuring and parameterizing lexical convergence and divergence between European and Brazilian Portuguese: endo/exogeneous influence and foreign and normative influence. *Advances in Cognitive Sociolinguistics*.
- Solorio, T., Blair, E., Maharjan, S., Bethard, S., Diab, M., Ghoneim, M., Hawwari, A., AlGhamdi, F., Hirschberg, J., Chang, A., and Fung, P. (2014). Overview for the first shared task on language identification in code-switched data. In *Proceedings of the First Workshop on Computational Approaches to Code Switching*.
- Tan, L., Zampieri, M., Ljubešić, N., and Tiedemann, J. (2014). Merging comparable data sources for the discrimination of similar languages: The DSL corpus collection. In *Proceedings of The BUCC Workshop*.
- Tiedemann, J. and Ljubešić, N. (2012). Efficient discrimination between closely related languages. In *Proceedings of COLING*.
- Tillmann, C., Mansour, S., and Al-Onaizan, Y. (2014). Improved sentence-level Arabic dialect classification. In *Proceedings of the VarDial Workshop*, pages 110–119, Dublin, Ireland, August.
- Zaidan, O. F. and Callison-Burch, C. (2014). Arabic dialect identification. *Computational Linguistics*.
- Zampieri, M. and Gebre, B. G. (2012). Automatic identification of language varieties: The case of Portuguese. In *Proceedings of KONVENS*.
- Zampieri, M., Gebre, B. G., and Diwersy, S. (2013). N-gram language models and POS distribution for the identification of Spanish varieties. In *Proceedings of TALN*.
- Zampieri, M., Tan, L., Ljubešić, N., and Tiedemann, J. (2014). A report on the DSL shared task 2014. In *Proceedings of the VarDial Workshop*.
- Zampieri, M., Gebre, B. G., Costa, H., and van Genabith, J. (2015a). Comparing approaches to the identification of similar languages. In *Proceedings of the LT4VarDial Workshop*.
- Zampieri, M., Tan, L., Ljubešić, N., Tiedemann, J., and Nakov, P. (2015b). Overview of the DSL shared task 2015. In *Proceedings of LT4VarDial*.
- Zubiaga, A., San Vicente, I., Gamallo, P., Pichel, J. R., Alegria, I., Aranberri, N., Ezeiza, A., and Fresno, V. (2014). Overview of TweetLID: Tweet language identification at SEPLN 2014. In *Proceedings of SEPLN*.
- Zubiaga, A., San Vicente, I., Gamallo, P., Pichel, J. R., Alegria, I., Aranberri, N., Ezeiza, A., and Fresno, V. (2015). Tweetlid: a benchmark for tweet language identification. *Language Resources and Evaluation*, pages 1–38.