

Rediscovering the Slavic Continuum in Representations Emerging from Neural Models of Spoken Language Identification

Badr M. Abdullah Jacek Kudera Tania Avgustinova
Bernd Möbius Dietrich Klakow

Department of Language Science and Technology (LST)
Saarland Informatics Campus, Saarland University, 66123 Saarbrücken, Germany
Corresponding author: babdullah@lsv.uni-saarland.de

Abstract

Deep neural networks have been employed for various spoken language recognition tasks, including tasks that are multilingual by definition such as spoken language identification. In this paper, we present a neural model for Slavic language identification in speech signals and analyze its emergent representations to investigate whether they reflect objective measures of language relatedness and/or non-linguists' perception of language similarity. While our analysis shows that the language representation space indeed captures language relatedness to a great extent, we find perceptual confusability between languages in our study to be the best predictor of the language representation similarity.

1 Introduction

The relationship between a group of human languages can be characterized across several dimensions of variation (Skirgård et al., 2017), including (1) the temporal dimension, wherein languages have diverged from a common historical ancestor as in the case of Romance languages; (2) the spatial dimension, wherein the speaker communities are geographically adjacent as in the case of the Indo-Aryan and Dravidian languages of India; and (3) the socio-political dimension, wherein languages have evolved under shared political and/or religious forces as in the case of Arabic and Swahili. Languages, or language varieties, can be related across all these dimensions, which often results in a dialect continuum. Speakers of languages that constitute a dialect continuum can usually communicate with each other efficiently using their own mother tongue. The degree of intercomprehensibility between speakers of different language varieties within a continuum is mainly determined by linguistic similarities. A notable case of this phenomenon is the mutual intelligibility among the Slavic languages, which we study in this paper.

One of the goals of linguistics is to study and categorize languages based on objective measures of linguistic distance. The degrees of similarity at different levels of the linguistic structural organization can be seen as preconditions for, as well as predictors of, successful oral intercomprehension. For closely-related languages, similarities at the pre-lexical, that is the acoustic-phonetic and phonological, level have been found to be better predictors of cross-lingual speech intelligibility than lexical similarities (Gooskens et al., 2008; Heeringa et al., 2009). In a different, yet relevant research direction, Skirgård et al. (2017) have investigated non-linguists' perception of language variation using data from the popular spoken language guessing game, the Great Language Game (GLG). By analyzing the confusion patterns of the GLG's human participants, the authors have shown that factors predicting players' confusion in the game correspond to objective measures of similarity established by linguists. For example, both phylogenetic relatedness and overlap in phoneme inventories have been identified as factors of perceptual confusability (and by implication, similarity) of languages in GLG.

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

The development of automatic systems that determine the identity of the language in a speech segment has received attention in the speech recognition community (see Li et al. (2013) for an overview). State-of-the-art approaches for automatic spoken language identification, henceforth LID, are based on multilayer deep neural networks (DNNs). DNN-based LID systems are parametric models that learn a mapping from spectral acoustic features of (untranscribed) speech to high-level feature representations in geometric space where languages are linearly separable. These models have shown tremendous success not only in discriminating between distant languages but also closely-related language varieties (Gelly et al., 2016; Shon et al., 2018; Mateju et al., 2018). Nevertheless, none of the previous works in spoken language recognition has analyzed the emerging representations from neural LID models for related languages. Thus, it is still unknown whether the distances in these representation spaces correspond to objective measurements of linguistic similarity and/or to non-linguists’ perception of language variation. In this paper, we aim to fill this gap and consider the family of Slavic languages as a case study. Our key contribution is two-fold:

- (1) We present an LID model for Slavic languages based on convolutional neural networks. Our model incorporates a domain-adversarial training strategy to improve its robustness against non-language sources of variability in speech signals. We show that our approach significantly improves generalization across datasets that differ in their recording conditions (§3 and §4).
- (2) We analyze and visualize the emergent representations from our robust LID model for 11 Slavic languages, five of which are not observed (held-out) during training. We show that the distances in the representation space correspond to measures of linguistic distance to a great extent (§5).

In this paper, we attempt to bridge different lines of research that have so far remained unconnected. On the one hand, we employ neural architectures from the field of spoken language recognition and build a robust model to identify languages in contemporary acoustic realizations of Slavic speech. On the other hand, we analyze the emerging language representations using techniques established by previous research in multilingual natural language processing (NLP). We consequently shed light on the speech modality and show how (untranscribed) speech signals can complement research done in computational studies of linguistic typology and language variation.

2 Background and Related Work

2.1 Slavic Languages

The Slavic language family is a branch of Indo-European languages that is conventionally divided into three subgroups: West-, East-, and South-Slavic. Apart from being related across the temporal dimension by sharing a common ancestor, Slavic languages form a spatial continuum of variation in a relatively connected geographic area across Europe and Northern Asia, except for the region where the Romance and Finno-Ugric wedge separates the South-Slavic from the West- and East-Slavic subgroups. Beside this traditional division (see Ethnologue, 23ed.), alternative classifications can be found in the Slavistics literature (cf. Bednarczuk, 2018; Pianka and Tokarz, 2000; Dalewska-Greń, 2020; Lehr-Spławiński et al., 1954; Nalepa, 1968; Mańczak, 2004). Nevertheless, and despite differences in taxonomies among various proposals, the development of contemporary Slavic languages from a common historical ancestor is uncontroversial. The supporting arguments are based on historical phonology and comparative studies of the phoneme inventories (Sawicka, 1991), as well as on studies of loanwords and Slavic toponyms. The high number of cognates as well as cross-linguistically shared features, such as lexical aspect, phonemic jotation and complex consonant clusters, provide strong evidence for common roots. In terms of diachronic phonology, the Common-Slavic era ends with the vocalization and reduction of the *yers* – the so-called “half-vowels” or “reduced vowels”, [ɤ] and [ɛ]. The outcomes of these alternations consistently define the most common division of Slavic. Similarly, the reflexes of *yat* [ĕ] provide a clear distinction between East-, West- and South-Slavic. The results of common phonological processes, such as liquid metathesis, palatalization and sibilization, also support the tripartite division. Moreover, these regularities of sound changes allow us to precisely trace the phonological development within the language family not only in the core standardized varieties but also in vernaculars and dialects. One of

the objectives of our work is to assess the extent to which neural models of spoken language learn to detect such regularities from acoustic realizations of contemporary Slavic speech.

2.2 Language Identification in Speech Signals

Research in automatic identification of the language in a speech signal (Li et al., 2013) is mainly concerned with the development of computational models that take an acoustic realization of a short utterance (usually a few seconds of speech) and predict the spoken language as output. Currently, end-to-end deep neural networks are the predominant paradigm for LID and have shown tremendous success in previous works (Lopez-Moreno et al., 2014; Gonzalez-Dominguez et al., 2014; Gelly and Gauvain, 2017). In this paradigm, the LID problem has usually been modelled as a temporal sequence classification problem in which a spectro-temporal representation of a spoken utterance (e.g., a sequence of spectral feature vectors) is transformed via a multi-layer neural network into a high-level vector representation that captures language-ID features. Other works in the literature have addressed LID for closely-related spoken language varieties including Arabic dialects (Shon et al., 2018; Gelly et al., 2016), Iberian languages (Gelly et al., 2016), and Slavic languages (Mateju et al., 2018; Abdullah et al., 2020).

At the intersection of speech recognition and linguistic typology, Gutkin et al. (2018) have trained a neural network on a large-scale multilingual speech database to predict typological features of the World Atlas of Language Structure (WALS) (Dryer and Haspelmath, 2013) for a language given a speech segment. The authors have shown that the speech modality contains enough signals to predict typological features of a held-out set of languages without explicit linguistic annotations. Their findings indicate that neural networks trained on multilingual speech could capture linguistic regularities and generalize beyond the languages observed in the training data.

2.3 Language Representations in Continuous Vector Spaces

Inspired by the advances in representation learning for NLP, multilingual neural models have been explored in the literature to induce real-valued language vectors, also known as *language representations* or *language embeddings*, where a single vector ($\mathbf{v} \in \mathbb{R}^d$) is associated with each language. Even though it has been motivated from different points of view, the main idea of this stream of research is to train a single NLP model on many languages whereby the language representation space is learned by exploiting the multilingual signal. For example, Johnson et al. (2017) introduced a multilingual neural machine translation (NMT) model in which the required target language of the translation was specified by the language embedding. Other works have either scaled this approach to a massively multilingual setting (Östling and Tiedemann, 2017; Malaviya et al., 2017) or explored other NLP tasks such as linguistic structure prediction (Bjerva et al., 2019) and grapheme-to-phoneme conversion (Peters et al., 2017). Furthermore, Rabinovich et al. (2017) and Bjerva et al. (2019) have analyzed the learned language representations and shown that the distance in the representation space reflects the phylogenetic distance between Indo-European languages. However, Bjerva et al. (2019) have argued that structural syntactic similarities between languages are a better predictor of the language representation similarities than phylogenetic similarities.

The most relevant analysis to ours is the recent work by Cathcart and Wandl (2020), in which the authors have trained a neural sequence-to-sequence model on a Slavic etymological dictionary. Their model was trained to consume a reconstructed Proto-Slavic word form and a language embedding, then emit a word form in the modern language specified by the language embedding. The authors have applied a clustering analysis on the learned language embeddings and successfully reconstructed the phylogenetic Slavic family tree. Our work complements this line of research with one fundamental difference: we perform our analysis on contemporary realizations of Slavic speech instead of the historically reconstructed phonological data without explicitly training our model to capture systematic sound changes.

3 Methodology

3.1 Slavic Speech Data

The data we use in this research are drawn from two different datasets:

- (1) **Radio Broadcast Speech (RBS)** A large collection of Slavic speech recordings were collected by crawling online radio stations in previous work (Nouza et al., 2016; Mateju et al., 2018). The dataset includes speech segments in 11 Slavic languages from the three subgroups: (1) South-Slavic: Bulgarian (BUL), Croatian (HRV), Serbian (SRP), Slovene (SLV), and Macedonian (MAC). (2) West-Slavic: Czech (CZE), Polish (POL), and Slovak (SLO). (3) East-Slavic: Russian (RUS), Ukrainian (UKR), and Belorussian (BEL). The audio recordings are either segments of professional news reports or of spontaneous speech during discussions. The recording conditions are diverse and the utterances occasionally include background music. We sample 8,000 and 500 utterances per language from the training split as our training and validation sets, respectively, and use the test set in Mateju et al. (2018) as our evaluation set.
- (2) **GlobalPhone Read Speech (GRS)** We also use the Slavic portion of the multilingual Global-Phone speech database (Schultz et al., 2013) which includes read speech recordings from native speakers of six Slavic languages: Bulgarian, Croatian, Czech, Polish, Russian, and Ukrainian. The utterances vary in length and quality across languages. Our final GRS training subset consists of 8,000 utterances per language.

3.2 Signal Representations of Speech

Human speech can be modelled with various signal representations. For automatic speech recognition (ASR), the conventional approach is to convert time-varying speech waveforms into time-frequency, or spectro-temporal, representations using a standard signal processing pipeline based on the short-time Fourier transform. An example of such representations is the mel-frequency spectral coefficients (MFSCs) representation, whose development has been inspired by the human auditory system. MFSCs describe the spectral envelope along the temporal dimension in a way that reflects the shape of the human vocal tract during speech production at each timepoint. In this paper, we use MFSCs for all presented experiments.

3.3 Neural LID Model

Our problem definition and LID models are based on the work of Abdullah et al. (2020). The LID task is defined as an instance of temporal sequence classification. A speech segment is first converted into a sequence of acoustic events $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$, where $\mathbf{x}_t \in \mathbb{R}^k$ is a spectral feature vector at timestep t . Then, the goal is to predict the spoken language \hat{y} given the sequence \mathbf{X} . This definition can be formalized using a deep neural network as a parameterization of the model as follows

$$\hat{y} = \arg \max_{y \in \mathcal{Y}} P(y | \mathbf{X}; \boldsymbol{\theta})$$

where \mathcal{Y} is the set of languages, $\boldsymbol{\theta}$ is the model’s parameters learned from a labelled dataset, and $P(y|\mathbf{X}; \boldsymbol{\theta})$ is the posterior probability of the language label y .

Baseline. We use an end-to-end convolutional neural network (CNN) with three convolutional layers followed by three feed-forward layers (see Fig. 1 for full description of the model). Our baseline LID model can be viewed as two components that are jointly trained: a *segment-level feature extractor* (F) and a *language classifier* (G), each associated with two sets of parameters $\boldsymbol{\theta}_F$ and $\boldsymbol{\theta}_G$, respectively. The parameters of the network $\boldsymbol{\theta}_F$ and $\boldsymbol{\theta}_G$ are learned in a supervised approach given a source dataset $\mathcal{D}_S = \{(\mathbf{X}_i, y_i)\}_{i=1}^{N_S}$ of N_S labelled samples and an optimization algorithm that minimizes cross-entropy loss.

Robust LID. To improve the robustness of our model against non-linguistic sources of variability in speech signals, we employ a well-established adversarial domain adaptation strategy (Ganin and Lempitsky, 2015), which has been shown to be effective for LID (Abdullah et al., 2020). Adversarial domain adaptation aims to minimize the discrepancy between the representations of the model given speech samples from two sources that differ in their recording conditions or spoken genre. This approach only requires a target dataset $\mathcal{D}_T = \{(\mathbf{X}_i)\}_{i=1}^{N_T}$ of N_T unlabelled samples, in addition to the labelled samples of the source dataset \mathcal{D}_S to train the model. To this end, an adversarial *domain classifier* (D) with

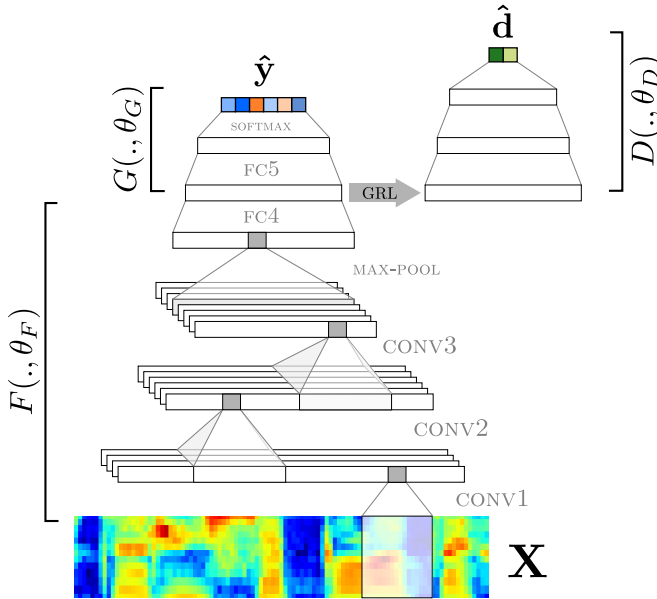


Figure 1: A schematic view of our LID model. A 3-layer 1-dimensional convolutional network followed by 3-layer fully-connected feed-forward network. The segment-level feature extractor $F(\cdot, \theta_F)$ maps the input sequence \mathbf{X} into a d -dimensional feature vector $\mathbf{f} \in \mathbb{R}^d$, i.e. $\mathbf{f} = F(\mathbf{X}; \theta_F)$. Then, the language classifier $G(\cdot, \theta_G)$ maps \mathbf{f} into a probability distribution over the language space, i.e. $\hat{\mathbf{y}} = G(\mathbf{f}; \theta_G)$. In the robust LID model, the domain classifier $D(\cdot, \theta_D)$ is connected to the network through a gradient reversal layer (GRL) to predict the domain given \mathbf{f} .

parameters θ_D is connected to the network with the objective of predicting the dataset of each speech sample in the training data. Now the goal is to encourage the segment-level feature extractor F is to produce representations that are language-discriminative but invariant with respect to the non-linguistic sources of variability in speech signals. We refer the reader to previous work to get a detailed overview of adversarial training for domain adaptation in the context of the LID task (Abdullah et al., 2020).

4 Experiments and Results

4.1 Training and Evaluation Data

We train our LID models to discriminate between the six Slavic languages that are shared by the RBS and GRS datasets, namely: Bulgarian, Czech, Croatian, Polish, Russian, and Ukrainian. The labelled speech samples are drawn from the RBS dataset (source dataset) while the unlabelled speech samples, which are used to improve our model robustness against non-language sources of variability, are drawn from the the GRS dataset (target dataset). For each language, we sample a balanced subset of 3-second 8,000 segments from each dataset to ensure our models are not effected by undesirable biases due to imbalanced conditions. We evaluate our models in two conditions: (1) in-domain evaluation, in which the evaluation samples come from the source dataset (i.e., the test split of RBS), and (2) cross-domain evaluation, in which the evaluation samples come from the target dataset whose labels are not observed during training (i.e., the development split of GRS).

4.2 Feature Extraction

In our experiments, we use the first 12 mel-frequency spectral coefficients (MFSCs) and frame-level averaged energy as low-level speech features. We extract frames of 25ms with 10ms overlap. Then, each speech sample is normalized with utterance-level mean and variance normalization.

4.3 Model Architecture and Hyperparameters

CNN Architecture. We employ three 1-dimensional convolutional layers over the temporal dimension with 128, 256, and 512 filters and filter widths of 5, 10, and 10 for each layer with strides of 1 step for each layer. Batch normalization and ReLU non-linearity are applied after each convolutional operation. We downsample the representation by applying a single max pooling operation at the end of the convolution block. For the language classifier, we use three fully-connected layers ($512 \rightarrow 512 \rightarrow 512 \rightarrow 6$) before the softmax for both the baseline and the robust LID models.

LID Model	In-domain				Cross-domain			
	1-sec	2-sec	3-sec	Full	1-sec	2-sec	3-sec	Full
Baseline LID	72.93	91.10	95.48	97.38	30.18	47.61	55.91	65.45
Robust LID	64.25	88.55	94.77	97.35	51.59	76.76	86.94	93.29
Δ	-11.9	-2.80	-0.74	-0.03	70.94	61.23	55.50	42.54

Table 1: In-domain and cross-domain evaluation of our LID models in balanced accuracy (%). Δ is the relative percentage difference in accuracy scores of baseline and robust LID models.

Adversarial Classifier. For our robust LID model, we add a 3-layer fully-connected feed-forward block ($512 \rightarrow 1024 \rightarrow 1024 \rightarrow 2$) to the network as the adversarial domain classifier D . The adversarial classifier takes the output of the segment-level feature extractor f (the output of layer FC4) and predicts the domain of the input sample. The adversarial loss is realized with a special layer, that is, a gradient reversal layer (GRL). The GRL behaves as an identity function during the forward pass but reverses the direction of the gradient signal during backpropagation.

Training Details. The cross-entropy loss is used for both the language classifier loss and the adversarial classifier loss. We use the Adam optimizer with a learning rate of 1×10^{-3} and train our models with a batch size of 256 samples. For the robust model, half of the samples within a mini-batch are drawn from the training split of the RBS dataset while the rest are drawn from the training split of the GRS dataset. Both the baseline and the robust models are trained for 50 epochs and the best models are selected based on the performance on the validation set. We do not use the early stopping criterion in our experiments.

4.4 Experimental Results

In this section, we report the evaluation results and show the effect of the adversarial training on the model’s robustness. Since the GRS evaluation data is imbalanced, we use balanced accuracy (Brodersen et al., 2010) as our evaluation metric to obtain a better estimate of the models’ performance. Table 1 shows our results for speech segments of various lengths. It can be observed from the cross-domain evaluation that the performance of the baseline model drops by a substantial factor. On the other hand, our robust model significantly improves cross-domain performance with little effect on the performance on the in-domain evaluation dataset, especially for longer utterances.

To get further insight into why adversarial training improves cross-domain performance, we analyze the predictions of the GRS evaluation samples made by the baseline and robust models by computing the F_1 score per language. The results of this analysis are shown in Table 2. For our baseline, we observe a much higher variance between languages compared to the robust model. The performance drop is more pronounced in the case of Ukrainian with a significant decrease in F_1 . However, our robust model boosts the F_1 score on Ukrainian from 14.66% to 94.49%.

LID Model	BUL	HRV	CZE	POL	RUS	UKR	macro Avg.	micro Avg.
Baseline LID	59.73	64.86	76.50	61.93	41.79	14.66	53.25	54.17
Robust LID	85.12	83.32	89.36	83.96	87.66	94.49	87.32	88.26
Δ	42.51	28.46	16.81	35.57	109.76	544.54	63.98	62.93

Table 2: F_1 score (%) per language. Predictions were obtained by feeding 3-second segments from GRS evaluation dataset to the LID models. Δ is the relative percentage difference in F_1 scores of baseline and robust LID models.

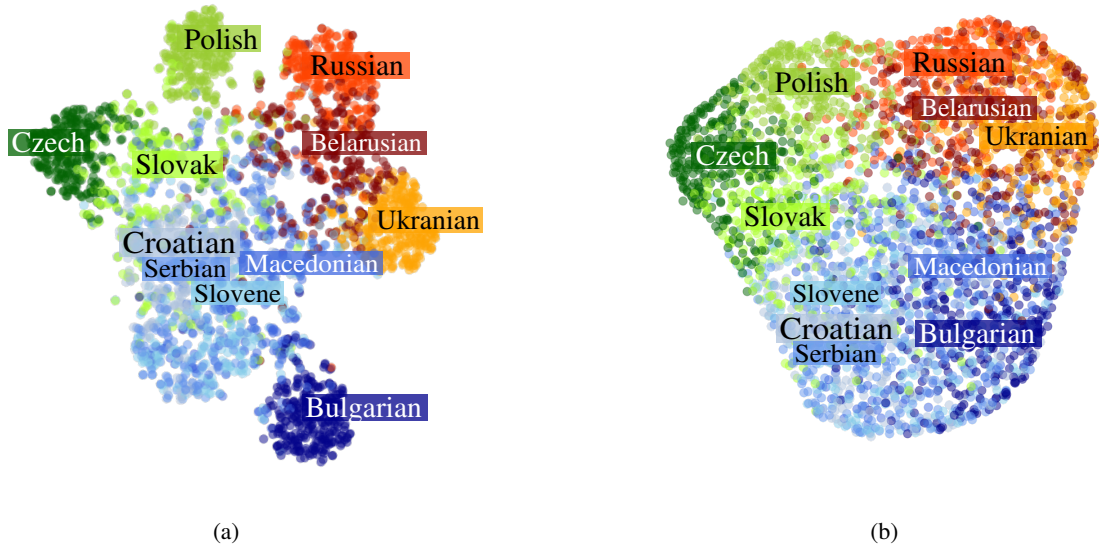


Figure 2: Two-dimensional visualization of representations of evaluation speech segments: (a) t-SNE projections, and (b) UMAP projections (best viewed in color).

5 Representation Similarity Analysis

In this section, we present our representation similarity analysis. From the robust LID model presented in the previous section, we obtain 512-dimensional representations for the 11 languages from the RBS evaluation set (500 5-second speech segments for each language). These representations are the output of the last fully-connected layer of the language classifier (i.e., layer FC5) before the softmax output layer.¹ It is worth pointing out that our LID model has been trained on a subset of only six languages out of the 11 languages that we analyze in this section.

5.1 Language Representation Visualizations

As our first analysis, we use dimensionality reduction techniques to obtain 2-dimensional projections from representations of the evaluation set and visualize the resulting data points. We use two dimensionality reduction techniques; t-SNE (Maaten and Hinton, 2008) and UMAP (McInnes et al., 2018). The resulting graphs are illustrated in Fig. 2. The motivation for using two different techniques in this analysis is that t-SNE and UMAP have different optimization objectives that complement each other. That is, the t-SNE algorithm preserves the local structure of the space; thus, it mainly reveals the cluster structure within the representation space. On the other hand, the UMAP algorithm preserves the global structure of the space. Nevertheless, both t-SNE and UMAP plots in Fig. 2 show very similar trends since the emerging subspaces shown in the figure correspond to the conventional sub-grouping of Slavic languages into East-, West-, and South-Slavic.

5.2 Correlation with Geographic Proximity

Following Bjerva et al. (2019), we investigate whether the language representations reflect the geographic proximity of their respective speaker communities. To perform this analysis, we first obtain a single prototypical vector representation for each Slavic language in our study by taking the average over the representations of the evaluation speech segments as

$$\mathbf{v}_L = \frac{1}{|\mathcal{E}_L|} \sum_{\mathbf{x} \in \mathcal{E}_L} \text{NN}(\mathbf{X})$$

¹In the speaker and language recognition community, these representations are usually referred to as \mathbf{x} -vectors, while the emerging geometric space is referred to as the \mathbf{x} -vector space.

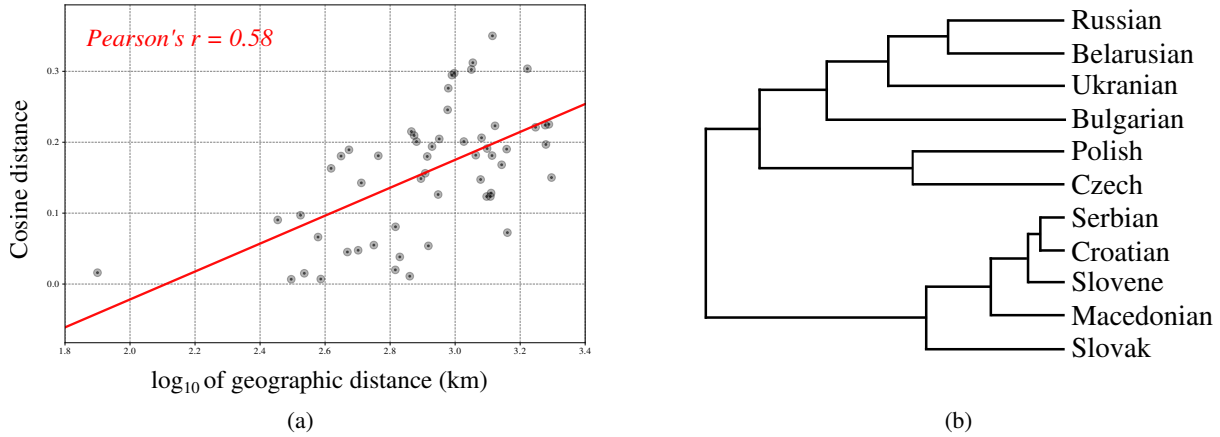


Figure 3: (a) Correlation between geographic distance and distances between prototype language representations measured by cosine similarity, and (b) a genetic tree generated from the pairwise distance matrix of the language prototype representations from our LID model.

where $\mathbf{v}_L \in \mathbb{R}^{512}$ is a prototypical vector representation for language L , \mathcal{E}_L is the evaluation speech segments for language L , and $\text{NN}(\cdot)$ is the output of the last non-linear layer (layer FC5) of the robust LID model. The distance in the representation space between two languages is computed using cosine distance. For geographic distance, we follow a similar approach as in Skirg ard et al. (2017). First, each language is characterized by a point location on the map given the latitude and longitude information in the ASJP linguistic database (which is intended to represent the cultural or historical center of the language). We then compute the pairwise distances between the points on the map (in kilometers) and convert them into \log_{10} scale. Fig. 3(a) shows a scatter plot between the data points in which the x -axis represents the geographic distance and the y -axis represents the cosine distance in the representation space. We observe a positive correlation between the two distance measures (Pearson’s $r = 0.58$). This clearly shows that the distance in the representation space does indeed reflect the geographic distance.

5.3 Genetic Signal in the Representation Space

Similar to the analysis in Bjerva et al. (2019) and Cathcart and Wandl (2020), we investigate the genetic signal in the representation space. To this end, the pairwise cosine distances computed in the previous section are first converted into a confusion matrix. Then, we generate a tree by performing hierarchical clustering on the confusion matrix using the Ward algorithm. The resulting tree is depicted in Fig. 3(b). We observe that the generated tree shows many similarities to phylogenetic trees that correspond to the widely accepted tripartite division of Slavic languages (that is, the 3-way categorization of Slavic languages into East-, West-, and South-Slavic). However, we observe two notable discrepancies between the widely-accepted Slavic tree and our generated tree: (1) the placement of Bulgarian (a South-Slavic language) within the East-Slavic group, and (2) the placement of Slovak (a West-Slavic language) within the South-Slavic group.

5.4 Fine-grained Analysis of Similarities

Although we have shown that the distances in the representation space of our LID model correlate with geographic distances and have adequately reconstructed the Slavic genetic tree, we seek to understand the factors that could better explain the similarities. To this end, we compare our generated tree with various genetic trees of Slavic languages presented in previous studies.²

- (1) **Levenshtein distance-based tree** Serva and Petroni (2008) have automatically generated a phylogenetic tree for 50 Indo-European languages using a renormalized Levenshtein distance based on a

²Although a correlation analysis on confusion matrices would have been more optimal, we instead apply tree similarity analysis because not all underlying confusion matrices are available.

	Tree	Distance
Levenshtein distance-based		0.270
Glottochronology-based		0.132
Geographic distance-based		0.140
GLG Confusion-based		0.084
	Random	0.371

Table 3: Tree distance evaluation. Lower values correspond to smaller distances.

Swadesh list of 200 words that have the same word meanings across languages. We take the Slavic branch from this tree for our similarity analysis.

- (2) **Glottochronology-based tree** Novotná and Blažek (2007) have used a classical glottochronological approach to generate a genetic tree for the Balto-Slavic languages. Their approach employs a manual calculation of pairwise distances between languages based on the recognition of cognates.
- (3) **Geographic distance-based tree** Using the geographic distance matrix we obtained from the ASJP database, we apply the same hierarchical clustering as we applied on our prototype language vectors to obtain a tree based on geographic distance.
- (4) **GLG Confusion-based tree** Skirgård et al. (2017) have applied hierarchical clustering on their player confusion data and their resulting tree shows that the Slavic languages form a pure cluster. We consider this tree to be a good approximation of non-linguist’s perception of language variation and similarity.
- (5) **Randomly generated tree** To give a reference point for a worst case scenario in our analysis, we generate a tree using the Ward hierarchical clustering on a random confusion matrix.

First, and to facilitate the measurement of distance between different trees, we consider Croatian and Serbian as a single node in our tree (i.e., Serbo-Croatian) and keep only the nodes that represent languages shared by all trees. This leaves us with eight languages: Bulgarian, Serbo-Croatian, Slovene, Czech, Polish, Slovak, Russian, and Ukrainian. Second, we compute the distances between the tree generated from the representation distances of our LID model and each of the aforementioned trees using the unweighted tree distance metric introduced by Rabinovich et al. (2017). The result of this evaluation is presented in Table 3. We observe that the most similar tree to our tree is the one based on the player confusion data from the GLG. These findings suggest that the factor that best explains the similarities within the emerging representations from our LID model is the perceptual similarity between Slavic languages, approximated by the confusability obtained from the GLG participants.

6 Discussion

Many recent works have shown that deep neural networks are good models of human perception. For example, Zhang et al. (2018) and Peterson et al. (2018) have shown that emerging representations from neural models trained on visual recognition tasks are predictive of human similarity judgments. For auditory recognition, neural speech recognition models have been shown to capture human-like behavior in cross-lingual phonetic perception (Schatz and Feldman, 2018). Following the same spirit, our objective is to investigate the extent to which neural models of spoken language identification capture language similarity. Nevertheless, and because of the complex space in which language variation can be realized, the similarity between two languages is a multidimensional phenomenon that cannot be expressed in a single number (Van Heuven, 2008). We therefore do not consider a single reference as a “ground truth” in our analysis, but consider several reference criteria of distance including genetic, geographic, and perceptual distance.

Our representation similarity analysis shows that emerging representations in our spoken LID model capture language similarity. The representation visualization illustrated in Fig. 2 demonstrates the gen-

eralization ability of our model to project speech segments of non-observed (held-out) languages into subspaces of their respective subgroups. Given that the data in our study constitute contemporary realizations of Slavic speech that do not explicitly encode diachronic sound changes, we first hypothesized that the geographic distances between the linguistic communities would be a good predictor of the distances in the representation space of the LID model. This turns out to indeed be the case as we observe a high positive correlation between geographic distances and cosine distances within prototype language representations.

On the other hand, we were less optimistic about our LID model capturing the genetic signal between Slavic languages (that is, whether the language representations encode the historical relationships between languages). Earlier works that have investigated computational approaches to generating genetic language trees have either employed historical etymological data capturing phonological sound changes (Cathcart and Wandl, 2020), sequences reflecting syntactic patterns in different languages (Rabinovich et al., 2017; Bjerva et al., 2019), or word lists reflecting lexical similarity (Serva and Petroni, 2008). Arguably, these sources of language data are more likely to preserve the relationship between languages across the temporal dimension than the contemporary Slavic speech we use in this study. Therefore, our initial intuition was that the resulting tree would reflect variation across the spatial dimension and not the temporal dimension. Nevertheless, the tree generated by our analysis is an adequate approximation of the Slavic genetic tree given the contemporary nature of the data sources.

Finally, it is striking to observe the similarity between our generated tree and the Slavic subtree constructed from player confusion patterns of GLG participants (Skirgård et al., 2017). There are two main factors that contribute to this similarity of our analysis to that of Skirgård et al. (2017): (1) West-Slavic and East-Slavic branches are clustered together before joining the South-Slavic branch to form a single Slavic cluster, and (2) the deviations from the widely accepted Slavic grouping are present in both analyses at approximately the same node locations since Bulgarian is grouped with East-Slavic and Slovak is grouped with South-Slavic. These findings provide evidence that the perceived language similarity is the factor that can best predict the geometric distance between languages in the emerging representation space from neural models of spoken language identification.

7 Conclusion

We have presented a convolutional neural model for Slavic language identification in speech signals and analyzed the extent to which its emerging representations reflect language similarity. Our analysis has shown that the distances in the emergent language representations reflect language variation across the temporal (phylogenetic) and spatial (geographic) dimensions. Moreover, by comparing our clustering analysis to the confusion patterns of the Great Language Game participants, we have shown that perceptual confusability is a better predictor of language representation similarities than phylogenetic and geographic distances.

Acknowledgements

We would like to thank Nicole Macher for assisting with the research presented in this paper. We extend our gratitude to anonymous reviewers for their insightful suggestions and comments. This research is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), Project ID 232722074, SFB 1102.

References

- Badr M. Abdullah, Tania Avgustinova, Bernd Mobius, and Dietrich Klakow. Cross-domain adaptation of spoken language identification for related languages: The curious case of slavic languages. In *Interspeech*, 2020.
- Leszek Bednarczuk. *Poczkatki i pogranicza polszczyzny*. Lexis, 2018.
- Johannes Bjerva, Robert Östling, Maria Han Veiga, Jörg Tiedemann, and Isabelle Augenstein. What do language representations really represent? *Computational Linguistics*, 45(2):381–389, 2019.

- Kay Henning Brodersen, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M Buhmann. The balanced accuracy and its posterior distribution. In *International Conference on Pattern Recognition*, pages 3121–3124. IEEE, 2010.
- Chundra Cathcart and Florian Wandl. In search of isoglosses: continuous and discrete language embeddings in Slavic historical phonology. In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, Online, July 2020. Association for Computational Linguistics.
- Hanna Dalewska-Greń. *Języki słowiańskie*. Wydawnictwo Naukowe PWN, Warszawa, second edition, 2020.
- Matthew S. Dryer and Martin Haspelmath, editors. *WALS Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig, 2013. URL <https://wals.info/>.
- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning*, pages 1180–1189, 2015.
- Gregory Gelly and Jean-Luc Gauvain. Spoken language identification using lstm-based angular proximity. In *Interspeech*, 2017.
- Gregory Gelly, Jean-Luc Gauvain, Lori Lamel, Antoine Laurent, Viet Bac Le, and Abdel Messaoudi. Language recognition for dialects and closely related languages. In *Odyssey*, volume 2016, pages 124–131, 2016.
- Javier Gonzalez-Dominguez, Ignacio Lopez-Moreno, Haşim Sak, Joaquin Gonzalez-Rodriguez, and Pedro J Moreno. Automatic language identification using long short-term memory recurrent neural networks. In *Interspeech*, 2014.
- Charlotte Gooskens, Wilbert Heeringa, and Karin Beijering. Phonetic and lexical predictors of intelligibility. *International journal of humanities and arts computing*, 2(1-2):63–81, 2008.
- Alexander Gutkin, Tatiana Merkulova, and Martin Jansche. Predicting the features of world atlas of language structures from speech. In *Proc. The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages*, pages 248–252, 2018.
- Wilbert Heeringa, Keith Johnson, and Charlotte Gooskens. Measuring norwegian dialect distances using acoustic features. *Speech Communication*, 51(2):167–183, 2009.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017.
- Tadeusz Lehr-Spławiński, Władysław Kuraszkiewicz, and Franciszek Sławski. *Przegląd i charakterystyka języków słowiańskich*. Państwowe Wydawn. Naukowe, 1954.
- Haizhou Li, Bin Ma, and Kong Aik Lee. Spoken language recognition: from fundamentals to practice. *Proceedings of the IEEE*, 101(5):1136–1159, 2013.
- Ignacio Lopez-Moreno, Javier Gonzalez-Dominguez, Oldrich Plchot, David Martinez, Joaquin Gonzalez-Rodriguez, and Pedro Moreno. Automatic language identification using deep neural networks. In *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5337–5341. IEEE, 2014.
- Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- Chaitanya Malaviya, Graham Neubig, and Patrick Littell. Learning language representations for typology prediction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.
- Witold Mańczak. *Przedhistoryczne migracje Słowian i pochodzenie języka staro-cerkiewno-słowiańskiego*. Nakładem Polskiej Akad. Umiejętności, 2004.

- Lukas Mateju, Petr Cerva, Jindrich Zdánský, and Radek Safarik. Using deep neural networks for identification of slavic languages from acoustic signal. In *Interspeech*, 2018.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. Umap: Uniform manifold approximation and projection. *Journal of Open Source Software*, 3(29):861, 2018.
- Jerzy Nalepa. *Słowiańszczyzna północno-zachodnia: podstawy jedności i jej rozpad*, volume 25. Państwowe Wydawn. Naukowe; Oddz. w Poznaniu, 1968.
- Jan Nouza, Radek Safarik, and Petr Cerva. Asr for south slavic languages developed in almost automated way. In *Interspeech*, 2016.
- Petra Novotná and Václav Blažek. Glottochronology and its application to the balto-slavic languages. *Baltistica*, 42(2):185–210, 2007.
- Robert Östling and Jörg Tiedemann. Continuous multilinguality with language vectors. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Valencia, Spain, April 2017. Association for Computational Linguistics.
- Ben Peters, Jon Dehdari, and Josef van Genabith. Massively multilingual neural grapheme-to-phoneme conversion. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.
- Joshua C Peterson, Jordan W Suchow, Krisha Aghi, Alexander Y Ku, and Thomas L Griffiths. Capturing human category representations by sampling in deep feature spaces. *Proceedings of the 40th Annual Conference of the Cognitive Science Society.*, 2018.
- Włodzimierz Pianka and Emil Tokarz. *Gramatyka konfrontatywna języków słowiańskich. 1 (2000)*. Śląsk, 2000.
- Ella Rabinovich, Noam Ordan, and Shuly Wintner. Found in translation: Reconstructing phylogenetic language trees from translations. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vancouver, Canada, July 2017. Association for Computational Linguistics.
- Irena Sawicka. Problems of the phonetic typology of the slavic languages. *Studies in the Phonetic Typology of the Slavic Languages*. Warszawa: Slawistyczny Ośrodek Wydawniczy przy Instytucie Słowianoznawstwa PAN, 1991.
- Thomas Schatz and Naomi H Feldman. Neural network vs. hmm speech recognition systems as models of human cross-linguistic phonetic perception. In *Proceedings of the Conference on Cognitive Computational Neuroscience*, 2018.
- Tanja Schultz, Ngoc Thang Vu, and Tim Schlippe. Globalphone: A multilingual text & speech database in 20 languages. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8126–8130. IEEE, 2013.
- Maurizio Serva and Filippo Petroni. Indo-european languages tree by levenshtein distance. *EPL (Europhysics Letters)*, 81(6):68005, 2008.
- Suwon Shon, Ahmed Ali, and James Glass. Convolutional neural network and language embeddings for end-to-end dialect recognition. In *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, pages 98–104, 2018.
- Hedvig Skirgård, Seán G Roberts, and Lars Yencken. Why are some languages confused for others? investigating data from the great language game. *PloS one*, 12(4), 2017.
- Vincent J Van Heuven. Making sense of strange sounds:(mutual) intelligibility of related language varieties. a review. *International journal of humanities and arts computing*, 2(1-2):39–62, 2008.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.