



## Research Article

## Neural representations for modeling variation in speech

Martijn Bartelds<sup>a,\*</sup>, Wietse de Vries<sup>a</sup>, Faraz Sanal<sup>b</sup>, Caitlin Richter<sup>b</sup>, Mark Liberman<sup>b</sup>,  
Martijn Wieling<sup>a,c</sup>

<sup>a</sup> Center for Language and Cognition, Faculty of Arts, University of Groningen, Groningen, The Netherlands

<sup>b</sup> Department of Linguistics, University of Pennsylvania, Philadelphia, PA, USA

<sup>c</sup> Haskins Laboratories, New Haven, CT, USA



## ARTICLE INFO

## Article history:

Received 25 November 2020

Received in revised form 26 January 2022

Accepted 3 February 2022

Available online 5 March 2022

## Keyword:

Acoustic distance

Acoustic embeddings

Neural networks

Pronunciation variation

Speech

Transformers

Unsupervised representation learning

## ABSTRACT

Variation in speech is often quantified by comparing phonetic transcriptions of the same utterance. However, manually transcribing speech is time-consuming and error prone. As an alternative, therefore, we investigate the extraction of acoustic embeddings from several self-supervised neural models. We use these representations to compute word-based pronunciation differences between non-native and native speakers of English, and between Norwegian dialect speakers. For comparison with several earlier studies, we evaluate how well these differences match human perception by comparing them with available human judgements of similarity. We show that speech representations extracted from a specific type of neural model (i.e. Transformers) lead to a better match with human perception than two earlier approaches on the basis of phonetic transcriptions and MFCC-based acoustic features. We furthermore find that features from the neural models can generally best be extracted from one of the middle hidden layers than from the final layer. We also demonstrate that neural speech representations not only capture segmental differences, but also intonational and durational differences that cannot adequately be represented by a set of discrete symbols used in phonetic transcriptions.

© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Past work in e.g., automatic speech recognition has found that variability in speech signals is often poorly modeled, despite recent advances in speech representation learning using deep neural networks (Huang, Slaney, Seltzer, & Gong, 2014; Huang, Dong, Liu, & Gong, 2014; Koenecke et al., 2020). This may be particularly true for monolingual as opposed to multilingual models (Želasko, Moro-Velázquez, Hasegawa-Johnson, Scharenborg, & Dehak, 2020). While acoustic variability may be caused by technical aspects such as microphone variability (Mathur, Isopoussu, Kawsar, Berthouze, & Lane, 2019), an important source of variation is the embedding of accent or dialect information in the speech signal (Hanani, Russell, & Carey, 2013; Najafian, DeMarco, Cox, & Russell, 2014). Non-native accents are frequently observed when a second language is spoken, and are mainly

caused by the first language background of non-native speakers. Similarly, regional accents are caused by the (first) dialect or regional language of the speaker. The accent strength of a speaker depends on the amount of transfer from their native language or dialect, and is generally influenced by a variety of characteristics, of which the age of learning the (second) language, and the duration of exposure to the (second) language are important predictors (Asher & García, 1969; Leather, 1983; Flege, 1988; Wieling, Bloem, Baayen, & Nerbonne, 2014).

However, accent and dialect variability are often overlooked in modeling languages using speech technology, and consequently high-resource languages such as English are often treated as homogeneous (Blodgett, Green, & O'Connor, 2016). Given that the number of non-native speakers of English is almost twice as large as the former group, this assumption is problematic (Vigliano, Motlicek, & Cernak, 2019). It is therefore important to accurately model pronunciation variation using representations of speech that allow accent and dialect variability to be adequately incorporated.

Traditionally, pronunciations are often represented by phonetically transcribing speech (Nerbonne & Heeringa, 1997;

\* Corresponding author at: Center for Language and Cognition, Faculty of Arts, University of Groningen, Oude Kijk in 't Jatstraat 26, 9712 EK Groningen, The Netherlands.

E-mail address: [m.bartelds@rug.nl](mailto:m.bartelds@rug.nl) (M. Bartelds).

Livescu & Glass, 2000; Gooskens & Heeringa, 2004; Heeringa, 2004; Wieling, Bloem, Mignella, Timmermeister, & Nerbonne, 2014; Chen, Wee, Tong, Ma, & Li, 2016; Jeszenszky, Stoeckle, Glaser, & Weibel, 2017). However, accurately transcribing speech using a phonetic alphabet is time consuming, labor intensive, and interference from transcriber variation might lead to inconsistencies (Hakkani-Tür, Riccardi, & Gorin, 2002; Bucholtz, 2007; Novotney & Callison-Burch, 2010). Additionally, phonetic transcriptions are not entirely adequate in representing how people speak, as fine-grained pronunciation differences that are relevant for studying accented speech (or dialect variation) may not be fully captured with a discrete set of symbols (Mermelstein, 1976; Duckworth, Allen, Hardcastle, & Ball, 1990; Cucchiari, 1996; Liberman, 2018).

Consequently, acoustic-only measures have been proposed for comparing pronunciations (Huckvale, 2007; Ferragne & Pellegrino, 2010; Strycharczuk, López-Ibáñez, Brown, & Leemann, 2020). Whereas these studies only considered limited segments of speech, or exclusively included speakers from a single language background, Bartelds, Richter, Liberman, and Wieling (2020) introduced a new method that did not have these limitations. Specifically, Bartelds et al. (2020) proposed an acoustic-only method for comparing pronunciations without phonetic transcriptions, including speakers from multiple native language backgrounds while using all information available within the speech signal. In their method, they represented accented speech as 39-dimensional Mel-frequency cepstral coefficients (MFCCs), which were used to compute acoustic-based non-native-likeness ratings between non-native and native speakers of English. They found a strong correlation of  $r = -0.71$  between their automatically determined acoustic-based non-native-likeness scores and previously obtained native-likeness ratings provided by human raters (Wieling et al., 2014). This result was close to, but still not equal to, the performance of an edit distance approach on the basis of phonetic transcriptions (which showed a correlation of  $r = -0.77$ ).

Bartelds et al. (2020) conducted several small-scale experiments to investigate whether more fine-grained characteristics of human speech were captured as compared to the phonetic transcription-based pronunciation difference measure. Their results showed that the acoustic-only measure captured segmental differences, intonational differences, and durational differences, but that the method was not invariant to characteristics of the recording device.

The quality of MFCC representations is known to be dependent on the presence of noise (Zhao & Wang, 2013). Recent work has shown that neural network models for self-supervised representation learning are less affected by noise, while being well-equipped to model complex non-linear relationships (Schneider, Baevski, Collobert, & Auli, 2019; Baevski, Schneider, & Auli, 2020; Ling, Liu, Salazar, & Kirchhoff, 2020; Baevski, Zhou, Mohamed, & Auli, 2020). Generally, neural models benefit from large amounts of labeled training data. However, self-supervised neural models learn representations of speech without the need for (manually) labeled training data. Therefore, these models can be trained using even larger amounts of data. Previous work has shown that fine-tuning these neural models using transcribed speech

resulted in representations that resembled phonetic structure, and offered significant improvements in downstream speech recognition tasks (van den Oord, Li, & Vinyals, 2019; Kahn et al., 2020). In contrast to previous methods for comparing pronunciations, these self-supervised (monolingual and multilingual) neural models are based on large amounts of data from a large group of (diverse) speakers and are therefore potentially robust against accent variation.

Consequently, in this paper, we employ and evaluate several of these self-supervised neural models in order to create a fully automatic acoustic-only pronunciation difference measure, which is able to quantify fine-grained differences between accents and dialects. Specifically, we compare and evaluate five self-supervised neural models, namely `wav2vec` (Schneider et al., 2019, subsequently denoted by `w2v`), `vq-wav2vec` with the BERT extension (Baevski, Schneider et al., 2020, subsequently denoted by `vqw2v`), `wav2vec 2.0` (Baevski, Zhou et al., 2020, subsequently denoted by `w2v2`), the multilingual `w2v2` model `XLSR-53` (Conneau, Baevski, Collobert, Mohamed, & Auli, 2020, subsequently denoted by `XLSR`), and `DeCoAR` (Ling et al., 2020). Each of these models learned speech representations by predicting short fragments of speech (e.g., approximately 300 ms on average in the case of `w2v2`) within spoken sentences from the training data. These predicted fragments therefore roughly correspond to one or more subsequent phonemes (including their transitions). These neural models were selected for this study as they achieved state-of-the-art speech recognition results on standard benchmarks such as the Wall Street Journal corpus (WSJ; Garofalo, Graff, Paul, & Pallett, 2007) and the LibriSpeech corpus (Panayotov, Chen, Povey, & Khudanpur, 2015), while differing with respect to their specific architecture.

There are several use cases in which adequately quantifying pronunciation differences automatically is important. First, the field of dialectometry (see e.g., Nerbonne & Heeringa, 1997; Wieling, Nerbonne, & Harald Baayen, 2011; Wieling & Nerbonne, 2015) investigates geographical (and social) dialect variation on the basis of pronunciation differences between different dialects. While there are several dialect (atlas) datasets containing phonetic transcriptions, differences in transcription practices (sometimes even within the same dataset; Wieling, Heeringa, & Nerbonne, 2007) limit the extent to which these pronunciations can be compared. An acoustic-only method would solve these compatibility issues, and would allow datasets that do not have phonetic transcriptions to be analyzed directly. Another use case is highlighted by a recent study of San et al. (2021). They automatically compare pronunciations acoustically to find pronunciations of a specific word from endangered languages in a large set of unannotated speech files. Such a system, if successful, directly impacts language maintenance and revitalisation activities.

To evaluate the quality of the pronunciation differences, we will use human perceptual judgements. Previous work has shown that human listeners can adequately assess and quantify differences between pronunciations (e.g., Preston, 1999; Gooskens, 2005; Scharenborg, 2007; ). To determine the relative performance of our methods, we compare the use of self-supervised neural models to the phonetic transcription-based approach of Wieling et al. (2014), and the MFCC-based

acoustic-only approach of Bartelds et al. (2020). More details about these methods are provided in Section 3.2.

To investigate the versatility and robustness of the various models, we use three different datasets for evaluation. The first is identical to the dataset used by Wieling et al. (2014) and Bartelds et al. (2020), and includes both acoustic recordings of native and non-native English speakers as well as human native-likeness judgements to compare against. The second is a new dataset which focuses on accented speech from a single group of (Dutch) non-native speakers, for which human native-likeness judgements are likewise available. As we would also like to evaluate the effectiveness of the neural models for a different type of data in another language, we additionally include a dataset with Norwegian dialect pronunciations and corresponding human native-likeness ratings. For reproducibility, we provide our code via <https://github.com/Bartelds/neural-acoustic-distance>.

To understand the phonetic information captured by the neural models, we introduce a visualization approach revealing the location of differences between two compared recordings, and conduct several additional small-scale experiments, in line with those conducted by Bartelds et al. (2020).

## 2. Materials

### 2.1. Datasets

Our acoustic data comes from three datasets in two different languages. We use two datasets that contain (mostly) non-native American-English pronunciations, and an additional dataset with Norwegian dialect pronunciations.

#### 2.1.1. Non-native American-English

Pronunciations from non-native speakers of American-English are obtained from the Speech Accent Archive (Weinberger & Kunath, 2011), as well as the Dutch speakers dataset described in Offrede et al. (2020). The Speech Accent Archive covers a wide variety of language backgrounds, while the Dutch speakers dataset is suitable for evaluating our method on a set of English pronunciations that have more fine-grained accent differences, as it only contains speakers with the same native (Dutch) language background.

The Speech Accent Archive contains over 2000 speech samples from native and non-native speakers of English. Each speaker reads the same 69-word paragraph that is shown in Example 1.

- (1) *Please call Stella. Ask her to bring these things with her from the store: Six spoons of fresh snow peas, five thick slabs of blue cheese, and maybe a snack for her brother Bob. We also need a small plastic snake and a big toy frog for the kids. She can scoop these things into three red bags, and we will go meet her Wednesday at the train station.*

Similar to past work of Wieling et al. (2014) and Bartelds et al. (2020), we use 280 speech samples from non-native American-English speakers as our target dataset (i.e. the non-native speakers for whom human native-likeness ratings are available), and 115 speech samples from U.S.-born L1 speakers as our reference native speaker dataset. As there is much regional variability in the pronunciation of the native American-English pronunciations, we use a set of reference

speakers (cf. Wieling et al., 2014) instead of a single reference speaker.

Among the 395 English samples from the Speech Accent Archive, 206 speakers are male and 189 speakers are female. From these speakers, 71 male and 44 female speakers belong to the native speaker (reference) set. The average age of the speakers in the entire dataset is 32.6 years ( $\sigma = 13.5$ ). Non-native speakers have an average age of onset for learning English of 10.5 years ( $\sigma = 6.6$ ). The 280 non-native American-English speakers have a total of 99 different native languages, with Spanish ( $N = 17$ ), French ( $N = 13$ ), and Arabic ( $N = 12$ ) occurring most frequently.

The Dutch speakers dataset includes recordings of native speakers of Dutch (with no other native languages) that all read the first two sentences of the same elicitation paragraph used for the Speech Accent Archive. These recordings were collected at a science event held at the Dutch music festival Lowlands, where Offrede et al. (2020) investigated the influence of alcohol on speech production in a native and non-native language. While the effect of alcohol on the pronunciation in the non-native language (English) was limited, we nevertheless only included the speech samples of all 62 sober participants (30 male and 32 female speakers). The average age of the speakers in this dataset is 33.4 years ( $\sigma = 10.3$ ). The average age of onset for learning English was not obtained, but generally Dutch children are exposed to English at an early age (i.e. the subject is mandatory in primary schools from the age of about 10 to 11 onwards, but children are usually exposed to English much earlier via mass media).

For each speaker in this dataset, we phonetically transcribed the pronunciations according to the International Phonetic Alphabet. These phonetic transcriptions were created by a single transcriber (matching the conventions used by Wieling et al., 2014), and used to obtain the transcription-based pronunciation distances (i.e. for comparison with the acoustic methods).

#### 2.1.2. Norwegian

This dataset consists of 15 recordings and phonetic transcriptions from Norwegian dialect speakers from 15 dialect areas (4 male and 11 female speakers). The average age of these speakers is 30.5 years ( $\sigma = 11$ ). Moreover, each speaker lived in the place where their dialect was spoken until the mean age of 20 years, and all speakers estimated that their pronunciations were representative of the dialect they speak.

Earlier work has used this dataset for comparing dialect differences on the basis of the Levenshtein distance (Gooskens & Heeringa, 2004) and formant-based acoustic features (Heeringa, Johnson, & Gooskens, 2009) to human perceptual dialect differences. We included this dataset and the perceptual ratings from Gooskens and Heeringa (2004) to specifically investigate whether the self-supervised neural models (even though these are, except for XLSR, based on the English language) are able to model differences for languages other than English.

The speakers in this dataset all read aloud 58 words from the fable 'The North Wind and the Sun'. The recordings were segmented in 58 samples corresponding to the words from the fable. For five dialects, one or two words were missing, as speakers were not always perfectly reading the text.

Phonetic transcriptions, which we use as input for the Levenshtein distance algorithm, were created by a single transcriber. The text, recordings, phonetic transcriptions, and transcription conventions are available online.<sup>1</sup>

## 2.2. Human accent and dialect difference ratings

Human accent ratings are widely used to evaluate accent-ness in speech (Koster & Koet, 1993; Munro, 1995; Magen, 1998; Munro & Derwing, 2001). Similarly, human ratings have been used to determine how different dialects are from each other (Gooskens & Heeringa, 2004). To evaluate our method, we report Pearson's correlation between the computed acoustic (or phonetic transcription-based) differences and the averaged human accent (or dialect difference) ratings. While we evaluated read as opposed to spontaneous speech, Munro and Derwing (1994) found that human accent ratings are not different for the two types of speech.

### 2.2.1. Non-native American-English

The perceptual data for the Speech Accent Archive speech samples were collected by Wieling et al. (2014). Native U.S.-born speakers of English were invited to rate the accent strength of a set of (at most) 50 samples through an online questionnaire. Accent strength ratings were given using a 7-point Likert scale ranging from 1 (very foreign sounding) to 7 (native English speaking abilities). While each speech sample contained the entire 69-word paragraph (average duration of the samples was 26.2 seconds), participants were allowed to provide their rating without having listened to the full sample. In total, the ratings of 1,143 participants were collected (57.6% male and 42.4% female) for a total of 286 speech samples, where each participant on average rated 41 speech samples ( $\sigma = 14$ ). The average amount of ratings per sample was 157 ( $\sigma = 71$ ). The mean age of the participants was 36.2 years ( $\sigma = 13.9$ ), and they most frequently lived in California (13.2%), New York (10.1%), and Massachusetts (5.9%). From the 286 samples, six were from native American-English speakers. These were also identified as such, as their average ratings ranged between 6.79 and 6.97 ( $0.22 \leq \sigma \leq 0.52$ ).

Human accent ratings of the second (Dutch speakers) dataset were provided by a different group of U.S.-born L1 speakers of English (Offrede et al., 2020). In this case, a questionnaire was created in which participants rated the accent strength of the speech samples on a 5-point Likert scale ranging from 1 (very foreign-sounding) to 5 (native English speaking abilities). Participants were not required to listen to the complete sample (average duration: 18.7 seconds) before providing their rating. A total of 115 participants (73.0% male, 25.2% female, and 1.8% other) rated an average of 17 speech samples each ( $\sigma = 9.2$ ). On average, each sample received 24 ratings ( $\sigma = 6.7$ ). The mean age of the participating raters was 47.9 years ( $\sigma = 16$ ). The participants most often originated from California (13.9%), New York (10.4%), and Pennsylvania (8.7%). As the samples were shorter than for the Speech Accent Archive, a less fine-grained rating scale was used.

The consistency of the ratings was assessed using Cronbach's alpha (Cronbach, 1951). For both studies, the ratings

were consistent, with alpha values of 0.85 and 0.92 for the Speech Accent Archive dataset and Dutch speakers dataset, respectively (Nunnally, 1978).

### 2.2.2. Norwegian

Gooskens and Heeringa (2004) carried out a listening experiment using the recordings of the Norwegian dataset. A total of 15 groups of raters (high school pupils, one group per dialect area) were asked to judge each speaker on a 10-point scale. A score of 1 was given when the pronunciation of the speaker was perceived to be similar to the rater's own dialect, while a score of 10 indicated that the pronunciation of the speaker was maximally dissimilar from the rater's own dialect. The average duration of the speech samples was about 31 seconds.

On average, each group consisted of 19 listeners (48% male and 52% female) with a mean age of 17.8 years. For the majority of their life (16.7 years, on average), raters had lived in the place where their dialect was spoken. Only 3% of the raters reported to never speak in their local dialect. About 81% of the raters reported to use their dialect often or always. The consistency of the ratings was not reported by Gooskens and Heeringa (2004).

## 3. Methods

### 3.1. Self-supervised neural models

We compare and evaluate five self-supervised pre-trained neural models (i.e. w2v, vqw2v, w2v2, XLSR, and DeCoAR). The self-supervised neural models have learned representations of acoustic recordings by training the models to predict upcoming speech frames, without using labeled data (Schneider et al., 2019; Ling et al., 2020; Baeviski, Schneider et al., 2020; Baeviski, Zhou et al., 2020). An important characteristic of these deep learning models is that they contain multiple hidden layers containing information about the underlying data. Architectures and training techniques of these models have typically been inspired by successful methods in natural language processing such as word2vec (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013), ELMo (Peters et al., 2018), and BERT (Devlin, Chang, Lee, & Toutanova, 2019).

All of the evaluated acoustic models, except XLSR, were pre-trained on the large unlabeled Librispeech dataset, which contains 960 hours of English speech obtained from audio books (LS960). This dataset is divided into two parts, namely a part which includes clean data (460 hours), and a part which includes noisy data (500 hours). Speakers with accents closest to American-English (represented by pronunciations from the Wall Street Journal-based CSR corpus (SI-84) described by Douglas & Baker, 1992) were included in the clean data part, while the noisy data part contained accents that were more distant from American-English (Panayotov et al., 2015). The XLSR model, instead, was trained on 56,000 hours of speech from a total of 53 languages, including European, Asian, and African languages. Note that the majority of the pre-training data for XLSR still consists of English speech (44,000 hours).

In addition to the pre-trained model variants, there are fine-tuned variants available for the w2v2 and XLSR models. These

<sup>1</sup> <https://www.hf.ntnu.no/nos/>.

models were fine-tuned on labeled data in a specific language to improve their performance on speech recognition tasks. However, the process of fine-tuning might have influenced the linguistic representations that are learned during pre-training. We therefore also include these fine-tuned model variants in our evaluation. For English, we evaluate the  $w2v2$  model that has been fine-tuned on 960 hours of English speech data (subsequently denoted by  $w2v2-en$ ), and the XLSR model that was fine-tuned on 1,686 hours of English speech data (further denoted by  $XLSR-en$ ). The  $w2v2-en$  model was chosen because it is the largest fine-tuned English model available, and Baevski, Zhou et al. (2020) showed that increasing the model size improved performance on all evaluated speech recognition tasks. For Norwegian, we included the XLSR model fine-tuned on 12 hours of Swedish (which was the closest language available to Norwegian with a fine-tuned model available; further denoted by  $XLSR-sv$ ).

The effectiveness of these self-supervised neural models was originally evaluated by using the learned representations for the task of automatic speech recognition. However, in this study we assess whether or not these acoustic models also capture fine-grained information such as pronunciation variation. As the investigated algorithms use multiple hidden layers to model the acoustic signal, we also evaluate (using a development set) which layers are most suitable for our specific task. More information about these and other aspects of the models can be found in Appendix A.1 and A.2.

### 3.2. Existing methods

For comparison with the self-supervised neural models, we also report the results on the basis of two existing approaches for quantifying pronunciation differences, namely the MFCC-based approach of Bartelds et al. (2020) and the phonetic transcription-based approach of Wieling, Margaretha, and Nerbonne (2012). Both methods are currently the best-performing automatic (acoustic- or transcription-based) algorithms for determining pronunciation differences that match human perceptual pronunciation differences well, and are explained in more detail below.

#### 3.2.1. Phonetic transcription-based distance calculation

The phonetic transcription-based distances are determined on the basis of the adjusted Levenshtein distance algorithm proposed by Wieling et al. (2012). The Levenshtein algorithm determines the cost of changing one phonetically transcribed pronunciation into another by counting the minimum amount of insertions, deletions, and substitutions (Levenshtein, 1966). The adjustment proposed by Wieling et al. (2012) extends the standard Levenshtein distance by incorporating sensitive segment differences (rather than the binary distinction of same vs. different) based on pointwise mutual information (PMI) (Church & Hanks, 1990). This data-driven method assigns lower costs to sound segments that frequently occur together, while higher costs are assigned to pairs of segments that occur infrequently together. These sensitive sound segment differences are subsequently incorporated in the Levenshtein distance algorithm. An example of a PMI-based Levenshtein alignment for two pronunciations of the word “afternoon” is shown in Fig. 1.

$\text{æ}$	$\text{ə}$	$\text{f}$	$\text{t}$	$\text{ə}$	$\text{r}$	$\text{n}$	$\text{u}$	$\text{n}$
$\text{æ}$		$\text{f}$	$\text{t}$	$\text{ə}$	$\text{r}$	$\text{n}$	$\text{u}$	$\text{n}$
.031		.030			.020			

Fig. 1. PMI-based Levenshtein alignment for two different pronunciations of the word “afternoon”. The total transcription-based pronunciation distance between the two pronunciations equals the sum of the costs of all edit operations (i.e. 0.081).

To obtain reliable segment distances using the PMI-based Levenshtein distance algorithm, it is beneficial if the number of words and segments is as large as possible. As the Dutch speakers dataset is relatively small, we instead used the sensitive segment differences obtained on the basis of the (larger) Speech Accent Archive dataset (i.e. the same as those used by Wieling et al., 2014).

After the Levenshtein distance algorithm (incorporating sensitive sound differences) is used to quantify the pronunciation difference between each word for a pair of speakers, the pronunciation difference between two speakers is subsequently determined by averaging all word-based pronunciation differences. Additionally, for the two English datasets, the difference between the pronunciation of a non-native speaker and native (American-English) speech (i.e. the non-native-likeness) is computed by averaging the pronunciation difference between the non-native speaker and a large set of native English speakers (the same for both datasets).

#### 3.2.2. MFCC-based acoustic distance calculation

For the Speech Accent Archive recordings, the MFCC-based differences between the individual non-native speakers and native English speakers were available from Bartelds et al. (2020). For the native Dutch speakers dataset, and the Norwegian dataset, we calculate these differences following the same approach. In short, this consists of comparing 39-dimensional MFCCs of pronunciations of the same word (by two speakers) to obtain the acoustic difference between the pronunciations. We use dynamic time warping to compare the MFCCs (Giorgino, 2009). This algorithm is widely used to compare sequences of speech features by computing the minimum cumulative distance (i.e. the shortest path) through a cost matrix that contains the Euclidean distance between every pair of points in the feature representations. To account for durational differences between the pronunciations, we normalize the minimum cumulative distance by the length of the feature representations. See Bartelds et al. (2020) for more details. Finally, the non-native-likeness is computed in the same way as for the Levenshtein distance algorithm, explained in the previous section.

## 4. Experimental setup

### 4.1. Non-native American-English pronunciation differences

Following Wieling et al. (2014) and Bartelds et al. (2020), we compute a measure of acoustic distance from native English speech by individually comparing the non-native target samples from both datasets to the 115 native reference samples. Neural representations of all samples are acquired by using the full samples as input to the neural models. The final output of these neural models should correspond with the original

input (including all frames), and will therefore not contain any new information. Because of this, we use the feature representations of hidden layers (discussed in Section 3.1) as acoustic embeddings. These representations are extracted by doing a forward pass through the model up to the target hidden layer. Specifically, we investigated for each neural model which layer performed best for our task, by evaluating the performance (i.e. the correlation with human ratings) using a held-out development set (25% of the data of the Speech Accent Archive dataset, and 50% of the data of the much smaller Dutch speaker dataset). As layers sometimes show very similar performance, we also evaluated which layers showed significant lower performance than the best-performing layer. For this, we used the modified z-statistic of Steiger (1980) for comparing dependent correlations. After selecting the best-performing layer, the performance is evaluated on the remaining data (and the full dataset, if the patterns of the development set and the other data are similar). Samples are cut into individual words after embedding extraction using time-alignments from the Penn Phonetics Lab Forced Aligner (Yuan & Liberman, 2008). For word pairs between a reference and target speaker, length normalized similarity scores between the embeddings are calculated using dynamic time warping.

Scores are averaged across all 69 words (Speech Accent Archive dataset) or 34 words (Dutch speakers dataset) to acquire a distance measurement between a target speaker and a reference speaker. To compute a single score of distance between a target speaker and native English speech, the distances between the target speaker and all reference native speakers are averaged.

We evaluate our algorithms on both datasets by calculating the Pearson correlation between the resulting acoustic distances and the averaged human native-likeness judgements for the target samples. Note, however, that the results on the basis of the Speech Accent Archive are likely more robust as this dataset contains a large amount of (longer) samples, a variety of native language backgrounds, and a larger amount of ratings per sample. We visualize the complete approach in Fig. 2.

#### 4.2. Norwegian pronunciation differences

For the Norwegian dataset, we measure acoustic distances by computing neural representations for the segmented word samples similar to the approach used for the non-native American-English samples. The selection of the best-performing layer for the neural methods was determined again using a validation set, containing a random sample of 50% of the data. Word-based neural representations of the same word are compared using dynamic time warping to obtain similarity scores, which are length normalized. These are subsequently averaged to obtain a single distance measure between two dialects (i.e. two speakers).

We evaluate our algorithms on the Norwegian dialects dataset by computing the Pearson correlation between the acoustic distances and perception scores provided by the dialect speakers, and compare this value to the correlation obtained by using phonetic transcription-based distances and MFCC-based distances instead of the self-supervised neural acoustic-only distances. As Gooskens and Heeringa (2004) found that dialect distances with respect to themselves erroneously increased the correlation with the perceptual distances, we excluded these distances from our analysis.

#### 4.3. Influence of sample

To obtain a better understanding of the influence of our reference sample, and the specific set of words on our results, we conduct several additional experiments on the (larger) Speech Accent Archive non-native English dataset using our best-performing model.

First, we investigate the effect of choosing a single reference speaker, as opposed to using the reference set of all 115 speakers. Second, we further examine the effect of speaker backgrounds on the correlation with human perception, by restricting the set of reference native speakers to speakers from the western half of the U.S. and the English-speaking part of Canada. We opt for this set, as these areas are characterized by less dialect variation compared to the

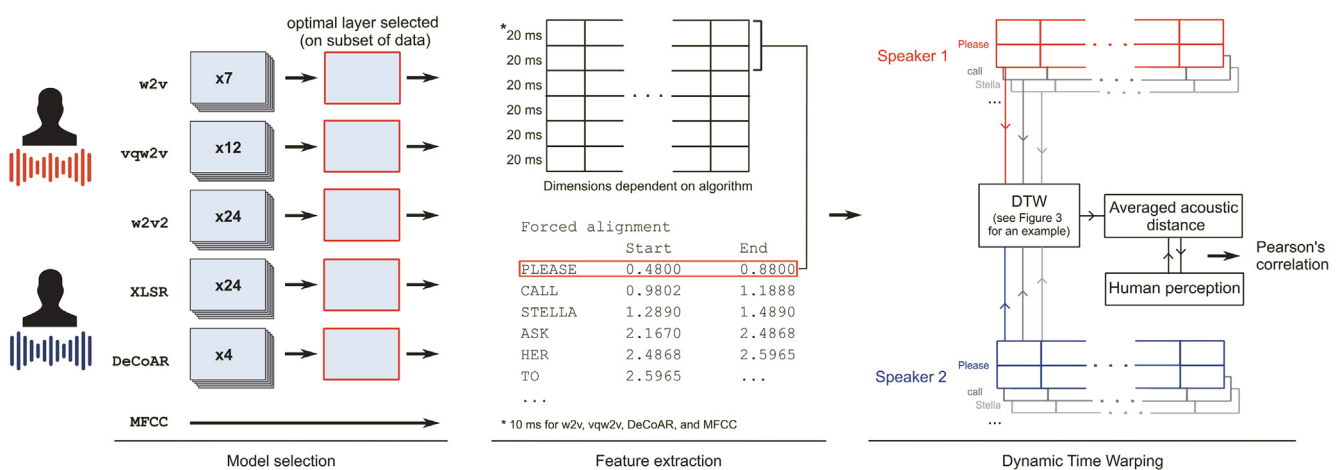


Fig. 2. Visualization of the acoustic distance measure where features are extracted using several acoustic-only methods. The output layer of the models is selected in a validation step. After feature extraction, the samples are sliced into individual words, which are subsequently compared using dynamic time warping. The word-based acoustic distances are then averaged and compared to human perception.

eastern half of the U.S. (Boberg, 2010). Third, as the gender distribution between the native and non-native speakers differed for our reference speaker set compared to the set of non-native speakers, we investigate the influence of gender by restricting the reference set to a single gender.

Finally, while the correlations are determined on the basis of an average over 69 words, we are also interested in the performance when only individual words are selected. This analysis may reveal which words are particularly informative when determining non-native-likeness.

#### 4.4. Understanding representations

To obtain a better understanding of the acoustic properties to which our final best-performing neural acoustic distance measure is sensitive, we conduct several additional experiments using the Speech Accent Archive recordings. We first evaluate how well the models are able to capture variation in specific groups of non-native speakers. By restricting the background (i.e. the native language) and thereby creating a more homogeneous sample (similar to the Dutch speakers dataset), human accent ratings may lie closer together. Strong correlations between human perception and acoustic distances when the range of scores is large (as in the full dataset), may not necessarily also imply strong correlations when there is less variation. Consequently, this experiment, together with the analysis of the Dutch speakers data, investigates whether or not our models also model human perception at a more fine-grained level.

In addition, to understand whether the acoustic distances comprise (linguistically relevant) aspects of pronunciation different from pronunciation distances computed using MFCCs or phonetic transcriptions, we fit multiple linear regression models. In those models, human accent ratings are predicted based on the acoustic distances of our best-performing self-supervised neural model, MFCC-based acoustic distances (Bartelds et al., 2020), and phonetic transcription-based differences (Wieling et al., 2014). We evaluate the contribution of each predictor to the model fit, and assess the model's explained variance to determine whether distinctive aspects of pronunciation are captured.

Finally, Bartelds et al. (2020) found that acoustic distances computed by using MFCCs not only captured segmental differences, but also intonational and durational differences between acoustically altered pronunciations of the same word. To assess whether this information is captured by our best-performing neural method as well, we replicate the experiment of Bartelds et al. (2020). Specifically, we compute acoustic distances between four series of recordings of the word "living" (ten repetitions per series) and compare the acoustic distances to those computed using MFCCs. The first two series of recordings were unmodified but recorded with a different recording device (the built-in microphone of a laptop, versus the built-in microphone of a smartphone). The third and fourth series were manipulated by changing the intonation ("living?") and relative duration of the first syllable ("li:ving"), respectively. To illustrate the results of this experiment, we have developed a visualization tool, which is discussed below and may help understand whether or not our best-performing (black box) neural method is able to distinguish aspects of speech that are linguistically relevant from those that are not.




Fig. 3. Visualization of neural acoustic distances per frame (based on  $w2v2$ ) with the pronunciation of /hy:d/ on the x-axis and distances to the pronunciation of /ho:d/ on the y-axis. The horizontal line represents the global distance value (i.e. the average of all individual frames). The blue continuous line represents the moving average distance based on 9 frames, corresponding to 180 ms. As a result of the moving average, the blue line does not cover the entire duration of the sample. Larger bullet sizes indicate that multiple frames in /ho:d/ are aligned to a single frame in /hy:d/.

##### 4.4.1. Visualization tool

For this study, we have developed a tool that visualizes the dynamic time warping alignments and the corresponding alignment costs to highlight where in the acoustic signal the differences between two pronunciations of the same word is most pronounced. As such, this tool may be helpful for interpreting the acoustic distances returned by our models, for example by highlighting that the acoustic differences between two pronunciations are most divergent at the end (or start) of a word. An illustration of the output (and interpretation) of this tool is shown in Fig. 3, which compares the pronunciation of a Dutch speaker pronouncing the two non-words /hy:d/ vs. /ho:d/. This example illustrates the relative influence of different phonemes on the acoustic distance within a word. The difference between the two pronunciations is lowest in the beginning of the word (/h/), whereas it is highest in the middle part (comparing [y:] and [o:]). The difference at the end (i.e. /d/) is higher than at the beginning (for /h/), which may reflect perseverative coarticulation, despite the transcriptions being identical. An online demo of this visualization tool can be used to generate similar figures for any pair of recorded pronunciations.<sup>2</sup>

## 5. Results

We first report on the performance of the non-native American-English speakers from the Speech Accent Archive and Dutch speakers dataset. Subsequently, we present the results on the Norwegian dataset to show how the self-supervised models perform on a language different from English. Finally, we discuss the phonetic information encoded in the pre-trained representations using visualizations of the acoustic distances, and report on the results from our additional experiments.

### 5.1. Non-native American-English pronunciation differences

Table 1 shows the correlations between the non-native-likeness scores and the average human native-likeness ratings for both datasets. The modified z-statistic of Steiger

<sup>2</sup> <https://bit.ly/visualization-tool>

**Table 1**

Pearson correlation coefficients  $r$  between acoustic-only or phonetic transcription-based distances and human native-likeness ratings, using  $w2v$ ,  $vqw2v$ ,  $w2v2$ ,  $XLSR$ ,  $w2v2-en$ ,  $XLSR-en$ ,  $DeCoAR$ , the PMI-based Levenshtein distance (LD), and MFCCs to compute distances on the Speech Accent Archive (SAA) dataset and native Dutch speakers dataset (DSD). All correlations are significant at the  $p < 0.001$  level. The values between parentheses show the selected layers of the neural models on the basis of the 25% validation set for the Speech Accent Archive dataset and the 50% validation set for the Dutch speakers dataset, respectively.

Model	SAA	DSD
$w2v$ (7, 5)	-0.69	-0.25
$vqw2v$ (11, 10)	-0.78	-0.67
$w2v2$ (17, 12)	-0.85	-0.70
$XLSR^3$ (16, 16)	-0.81	-0.47
$DeCoAR$ (2, 4)	-0.62	-0.40
$w2v2-en$ (10, 9)	<b>-0.87</b>	<b>-0.71</b>
$XLSR-en^3$ (8, 9)	-0.81	-0.63
LD (Wieling et al., 2014)	-0.77	-0.70
MFCC (Bartelds et al., 2020)	-0.71	-0.34

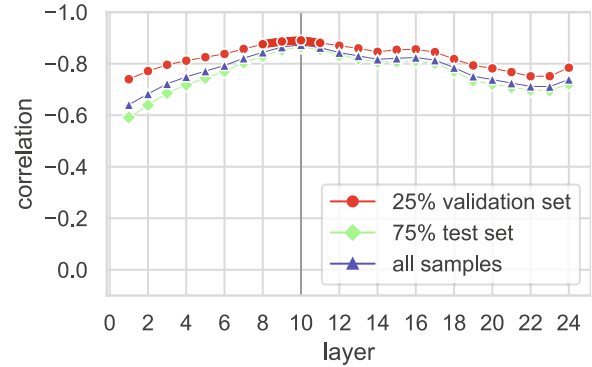
<sup>3</sup> We also computed correlation coefficients using the most recent  $XLS-R$  model (Babu et al., 2021), which is pre-trained on 436,000 hours of speech in 128 languages. To directly compare the results to  $XLSR$  and  $XLSR-en$ , we used the pre-trained model with the same number of parameters and fine-tuned this model on English labeled data available in the Common Voice dataset. However, the results of these newer models are not significantly better ( $p > 0.05$ ) from the results obtained using  $XLSR$  and  $XLSR-en$ . We therefore report those latter results.

(1980) shows that the  $w2v2-en$  model significantly outperforms all other models (including the Levenshtein distance approach, which was already reported to match human perception well; Wieling et al., 2014) when applied to the Speech Accent Archive dataset (all  $z$ 's  $> 3$ , all  $p$ 's  $< 0.001$ ). Similarly, for the Dutch speakers dataset, the  $w2v2-en$  is also the best-performing model. In this case, it significantly improved over  $w2v$ ,  $vqw2v$ ,  $DeCoAR$ ,  $XLSR$ , and MFCC (all  $z$ 's  $> 3$ , all  $p$ 's  $< 0.001$ ), but not over the other approaches ( $p > 0.05$ ).

For the neural models, the numbers between parentheses show the best-performing layer (on the basis of the performance on the validation set). As an example of how individual layers may show a different performance, Fig. 4 shows the performance for each layer for the best-performing  $w2v2-en$  model applied to the Speech Accent Archive dataset. It is clear that rather than selecting the final layer, the performance of an intermediate layer (10) is highest (and not significantly different from the performance of layers 8 to 11). Furthermore, there is a close match between the observed pattern for both the validation set and the test set. Appendix A.2 shows these graphs for all neural models and datasets.

## 5.2. Norwegian pronunciation differences

Table 2 shows the results for the Norwegian dialects dataset. In this experiment, we only include neural representations from the best-performing fine-tuned monolingual English and multilingual model in the previous section (i.e.  $w2v2-en$  and  $XLSR-sv$  as Swedish is more similar to Norwegian than English). Unfortunately, there is no monolingual Norwegian model available. In this case, the performance of the PMI-based Levenshtein distance is substantially (and significantly: all  $z$ 's  $> 3$ , all  $p$ 's  $< 0.001$ ) higher than both of the neural methods (which did not differ from each other in terms of performance, but did improve over the MFCC approach:



**Fig. 4.** Pearson correlation coefficients of acoustic distances compared to human accent ratings for different Transformer layers in the  $w2v2-en$  model. The vertical line marks the layer that was chosen as the best-performing layer based on the 25% validation set of the Speech Accent Archive dataset. Layers with a correlation that is not significantly different ( $p > 0.05$ ) from the optimal layer are indicated by the thick red line.

**Table 2**

Pearson correlation coefficients  $r$  between acoustic-only or phonetic transcription-based distances and human native-likeness ratings, using  $w2v2-en$ ,  $XLSR-sv$ , the PMI-based Levenshtein distance (LD), and MFCCs for computing pronunciation distances for the Norwegian dialect dataset. All correlations are significant at the  $p < 0.001$  level. The values between parentheses show the selected layers of the neural models on the basis of the 50% validation set.

Model	Mean $r$
$w2v2-en$ (3)	0.49
$XLSR-sv^4$ (7)	0.49
LD (Wieling et al., 2014)	<b>0.66</b>
MFCC (Bartelds et al., 2020)	0.22

<sup>4</sup> When using  $XLS-R$  fine-tuned on Swedish labeled data from the Common Voice dataset, the correlation coefficient is not significantly different ( $p > 0.05$ ) from  $XLSR-sv$ .

$z > 3$ ,  $p < 0.001$ ). Note that the correlations are positive, as higher perceptual ratings reflected more different dialects.

## 5.3. Influence of sample

In this section, we report on the influence of the specific sample of reference speakers and the included words across which we averaged. Table 3 reveals the influence of our specific sample of reference speakers by showing the averaged correlation coefficients (and the associated standard deviation) for the various methods applied to the Speech Accent Archive dataset. Instead of using the full set of 115 native speakers as reference set, in this analysis each individual native speaker was used once as the single reference speaker. Particularly of note is that only  $w2v2$ ,  $XLSR$  and their fine-tuned variants, as well as the PMI-based Levenshtein distance appear to be minimally influenced by individual reference speaker differences (i.e. reflected by the low standard deviations). Specifically,  $w2v2$  and  $w2v2-en$  yield the lowest standard deviations as well as the highest correlation ranges for individual reference speakers.

Additionally, we computed the correlation coefficient using our best-performing model (i.e.  $w2v2-en$ ) based solely on including reference native speakers from the western half of the U.S. and the English-speaking part of Canada. The resulting correlation of  $r = -0.87$  ( $p < 0.001$ ) was identical to the



**Table 3**

Averaged Pearson correlation coefficients  $r$ , with standard deviations and correlation ranges, between acoustic-only or phonetic transcription-based distances and human native-likeness ratings applied to the Speech Accent Archive dataset, using  $w2v$ ,  $vqw2v$ ,  $w2v2$  (pre-trained and fine-tuned), XLSR (pre-trained and fine-tuned), DeCoAR, the PMI-based Levenshtein distance (LD), and MFCCs to compute distances when individual U.S.-born native American-English speakers were treated as the single reference speaker. All correlation coefficients are significant at the  $p < 0.001$  level. The values between parentheses show the selected layer of the neural models on the basis of the validation set.

Model	Mean $r$	Std. Dev.	Range
$w2v$ (7)	-0.57	0.11	[-0.14, -0.73]
$vqw2v$ (11)	-0.69	0.08	[-0.16, -0.79]
$w2v2$ (17)	-0.83	0.02	[-0.73, -0.86]
XLSR (16)	-0.76	0.05	[-0.47, -0.83]
DeCoAR (2)	-0.49	0.08	[-0.22, -0.67]
$w2v2-en$ (10)	<b>-0.86</b>	<b>0.01</b>	<b>[-0.79, -0.88]</b>
XLSR-en (8)	-0.78	0.04	[-0.53, -0.83]
LD (Wieling et al., 2014)	-0.74	0.04	[-0.52, -0.79]
MFCC (Bartelds et al., 2020)	-0.45	0.10	[-0.20, -0.69]

correlation when including all reference speakers. The results were also similar when the reference speaker set was restricted to only men or women, with correlations of  $r = -0.87$  ( $p < 0.001$ ) and  $r = -0.87$  ( $p < 0.001$ ), respectively.

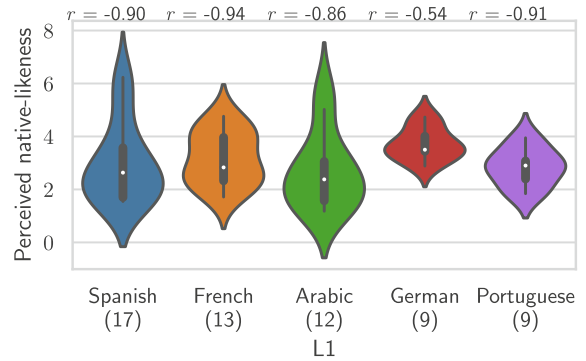
Finally, we calculated the correlation with human perception using  $w2v2-en$  when instead of the full 69-word paragraph individual words were selected. These correlations ranged from  $r = -0.50$  for the word “She” to  $r = -0.78$  for the word “Stella”. The average correlation was  $r = -0.67$  ( $p < 0.001$ ,  $\sigma = 0.06$ ). While the results on the basis of the full dataset show a higher correlation with human perception, it is noteworthy that some individual words also appear to correlate strongly with perception.

5.4. Understanding representations

To assess whether our best-performing model can also identify more fine-grained differences, we evaluate the model against several subsets of data consisting of non-native speakers from the same native language background. The spread in native-likeness ratings, as well as the correlations for the groups with the largest number of speakers are shown in Fig. 5. Except for the native speakers of German (with a relatively restricted range in native-likeness ratings), we observe strong correlations for all groups of speakers.

The low correlation for German speakers suggests that a restricted range of native-likeness ratings may negatively affect the correlation with human perceptual ratings. However, subsequent experiments using  $w2v2-en$  (not shown) revealed that the correlation when only including speakers who received average native-likeness ratings between, e.g., 5 and 6 was not lower than when increasing the range to include all speakers who received average native-likeness ratings between, e.g., 3 and 6.

To identify whether the acoustic distances computed using  $w2v2-en$  capture additional pronunciation characteristics compared to acoustic distances based on MFCCs or phonetic transcription-based distances, we fitted a multiple regression model predicting the human native-likeness ratings of the Speech Accent Archive dataset. Table 4 shows the estimated



**Fig. 5.** Violin plots visualizing the spread in native-likeness ratings for speakers of different native languages. The number of speakers is indicated between parentheses below the language. The correlation for each group is indicated above each violin plot.

coefficients (for standardized predictors), and summarizes the fit of the regression model. Acoustic distances computed using  $w2v2-en$  and phonetic transcription-based distances calculated by the PMI-based Levenshtein distance both contribute significantly to the model fit ( $p < 0.05$ ), whereas this is not the case for the MFCC-based distances. The contribution of  $w2v2-en$  is strongest as is clear from the standardized estimates. Overall, this model accounts for 77% of the variation in the human native-likeness assessments. A model fitted exclusively on the basis of the phonetic transcription-based distances explains 60% of the variation in the human native-likeness ratings. Given that a model fitted exclusively on the basis of the  $w2v2-en$ -based distances explains 76% of the variation in the human native-likeness ratings, these self-supervised neural models capture information that is not captured by phonetic transcriptions. Nevertheless, the abstractions provided by phonetic transcriptions do provide some (limited) additional information over the self-supervised neural models.

Table 5 shows how acoustic distances on the basis of the MFCC approach and the  $w2v2-en$  model are affected by into-

**Table 4**

Coefficients of a multiple regression model ( $R^2 = 0.77$ ) predicting human native-likeness judgements on the basis of phonetic transcription-based distances computed with the PMI-based Levenshtein distance (LD), and acoustic-only distances based on MFCCs and  $w2v2-en$ .

	Estimate (in z)	Std. Error	t-value	p-value
(Intercept)	2.98	0.03	86.56	< 0.001
LD (Wieling et al., 2014)	-0.15	0.06	-2.35	< 0.05
MFCC (Bartelds et al., 2020)	0.08	0.06	1.33	0.18
$w2v2-en$	-0.98	0.08	-11.75	< 0.001

**Table 5**

Normalized averaged acoustic distances of four variants of the word “living” (each repeated ten times) compared to the normal pronunciation of “living”, computed using  $w2v2-en$  and MFCCs. Standard deviations are shown between parentheses.

	$w2v2-en$	MFCC
Normal pronunciation	0.18 (0.10)	0.23 (0.13)
Normal pronunciation (different recording device)	0.29 (0.08)	0.88 (0.04)
Rising intonation	0.61 (0.07)	0.92 (0.03)
Lengthened first syllable	0.91 (0.05)	0.80 (0.03)

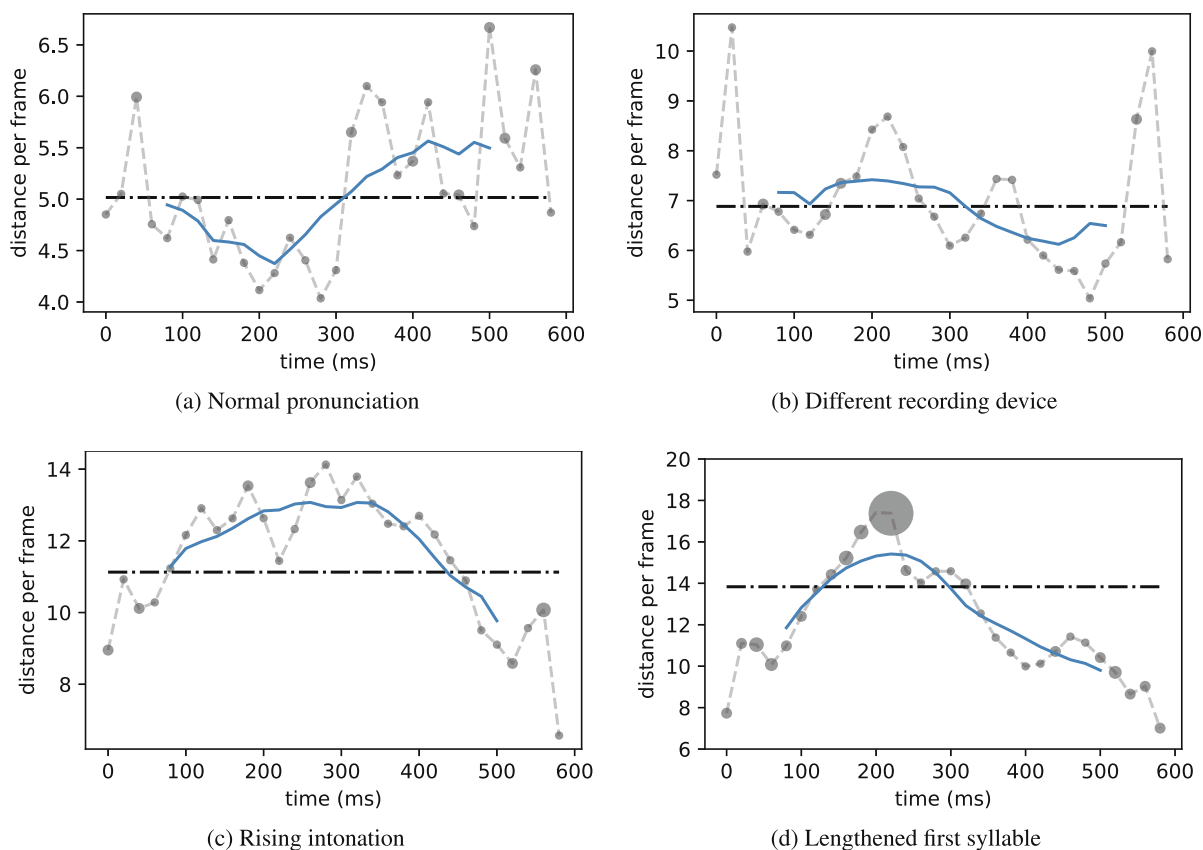
nation and timing differences, as well as by recording device. For each condition, ten repetitions were recorded. The recordings are the same as those used by Bartelds et al. (2020). To enable a better comparison, however, all obtained distances are scaled between 0 and 1. It is clear that the averaged distances from the repetitions of the same word (which may have differed slightly) are somewhat smaller for the  $w2v2-en$  model than for the MFCC approach. Importantly, whereas the MFCC approach does not cope well with a different recording device, the  $w2v2-en$  model appears to be much more robust (i.e. resulting in values closer to those of the normal pronunciation). Interestingly, whereas the MFCC approach appears to find larger differences between recordings differing in intonation compared to those with a lengthened first syllable, the opposite is true for the  $w2v2-en$  model. Both methods, however, appear to be sensitive to differences regarding these aspects.

To illustrate, Fig. 6 visualizes a comparison between a single normal pronunciation of “living” and four other pronunciations. Specifically, Fig. 6a shows a comparison with another normal pronunciation. Fig. 6b shows a comparison with the same pronunciation, but using a different recording device. Fig. 6c shows a comparison with a rising intonation pronunciation. Finally, Fig. 6d shows a comparison with a lengthened first syllable pronunciation. In line with Table 5, the values on the y-axis show that the distance between the two normal pronunciations is smaller than when using a different recording

device. Note that these distances were not normalized, as they simply compare two recordings. Both distances, however, are smaller than comparing against rising intonation (revealing a curvilinear pattern) and a lengthened first syllable (showing the largest difference at the beginning of the word; the lengthening is clear from the larger circle denoting an alignment with similar samples differing in duration).

## 6. Discussion and conclusion

In this study, we investigated how several self-supervised neural models may be used to automatically quantify pronunciation variation without needing to use phonetic transcription-based approaches. We used neural representations to calculate word-based pronunciation differences for English accents and Norwegian dialects, and compared the results to human perceptual judgements. While these ratings were provided on relatively crude (5 to 10-point) scales, and individual raters’ biases or strategies may have affected their ratings, averaging across a large number of raters for each sample likely yields an adequate estimate of native-likeness or perceived dialect distance. Our experiments showed that acoustic distances computed with Transformer-based models, such as  $w2v2-en$ , closely match the averaged human native-likeness ratings for the English datasets, and that performance greatly depended on the choice of layer. This finding not only demonstrates that



**Fig. 6.** Visualization of neural acoustic distances per frame (based on  $w2v2-en$ ) comparing each of the four variants of “living” to the same normal pronunciation. The horizontal line represents the global distance value (i.e. the average of all individual frames). The blue continuous line represents the moving average distance based on 9 frames, corresponding to 180 ms. As a result of the moving average, the blue line does not cover the entire duration of the sample. Larger bullet sizes indicate that multiple frames in the reference normal pronunciation are aligned to a single frame in the variant of “living” listed on the x-axis. Note the different scales of the y-axis, reflecting larger differences for the bottom two graphs compared to the top two graphs. See the text for further details.

these layers contain useful abstractions and generalizations of acoustic information, but also shows that the final layers represent information that is tailored to the target objective (which was speech recognition instead of our present goal of quantifying acoustic differences). This result is in line with findings in the field of natural language processing when using Transformer-based methods with textual data (Tenney, Das, & Pavlick, 2019; de Vries, van Cranenburgh, & Nissim, 2020). Furthermore, the  $w2v2$  and XLSR models appeared to be robust against the choice of reference speaker(s) to compare against. Even choosing a single reference speaker resulted in correlations that were not substantially different from those that used the full set. Interestingly, correlations on the basis of some words were not much lower than those on the basis of the full set of words, suggesting that a smaller number of words may already yield an adequate assessment of native-likeness.

Our newly-developed visualization tool helped us to understand these 'black box' models, as the visualization showed where the differences between two pronunciations were largest (i.e. the locus of the effect). This type of tool could potentially be used to provide visual feedback to learners of a second language or people with a speech disorder. However, the actual effectiveness of such an approach would need to be investigated.

Our results seem to indicate that phonetic transcriptions are no longer essential when the goal is to use these to quantify how different non-native speech is from native speech, and an appropriate Transformer-based model is available. This suggests that a time-consuming and labor intensive process can be omitted in this case. While our regression model showed that phonetic transcriptions did offer additional information not present in our neural acoustic-only approach, this information gain was very limited (an increase in  $R^2$  of only one percent). We furthermore showed that our neural method captures aspects of pronunciations (such as subtle durational or intonation differences) that are hard to capture by a set of discrete symbols used in phonetic transcriptions. Importantly, in contrast to a previous relatively successful acoustic approach (Bartelds et al., 2020), our present neural acoustic approach is relatively unaffected by non-linguistic variation (i.e. caused by using a different recording device). Nevertheless, further detailed research is needed to obtain a better view of what phonetic information is (not) captured by these models.

In contrast to the performance on the English datasets, we found that Transformer-based neural representations performed worse when applied to the Norwegian dialects dataset. However, pronunciations of the Norwegian dialects dataset were represented by a model which was trained exclusively or dominantly on English speech. Unfortunately, Norwegian was not among the pre-training languages included in the multilingual (XLSR) model, nor available for fine-tuning. We expect to see improved performance for a Norwegian  $w2v2$  model (when made available). Unfortunately, creating such a model is very costly in terms of required resources (generally based on hundreds of hours of speech) and computing power. At present, we estimate that pre-training (even without hyperparameter tuning and optimization) a new  $w2v2$  model for a different language takes about 150 days on a single state-of-the-art NVIDIA A100 GPU (costing approximately US\$ 10,000). Using

multiple GPUs in parallel reduces this duration, but also increases the required costs. Fortunately, the cost of these GPUs will usually decrease over time, and the speed of newly developed GPUs will increase.

Even though the evaluated architectures were originally designed for natural language processing, and the specific acoustic self-supervised neural models were created for improving performance in the domain of transforming speech to text, we have shown that the neural representations can also be successfully applied to an unrelated task in a different domain. Moreover, we have illustrated that Transformer-based speech representations are able to model fine-grained differences in homogeneous speaker groups (e.g., from the same language background), and adequately generalize over individual speaker differences, including gender, which makes them potentially useful for other tasks as well.

While our results are promising, the application of the  $w2v2$  approach for modeling pronunciation differences is only possible when an existing  $w2v2$  model is available for the language in question, or when sufficient data and computing resources are available to create a new  $w2v2$  model. In contrast to creating a new  $w2v2$  model for a new language, adjusting an existing model for a different task (such as quantifying differences with respect to pitch contours or timing patterns) is easier. This only requires an existing model to be fine-tuned on labeled examples. Generally, the amount of labeled data (and required GPU time) needed for fine-tuning is considerably lower compared to the large quantities of (unlabeled) data needed for pre-training. However, if such resources are not available, the monolingual English model appears to be a suitable alternative, generally outperforming the language-invariant acoustic-only method proposed by Bartelds et al. (2020). Future work, however, should be aimed at further investigating how existing high-resource language  $w2v2$  models may be exploited or extended when analyzing language variation in low-resource languages.

#### CRedit authorship contribution statement

**Martijn Bartelds:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing - original draft, Writing - review & editing. **Wietse de Vries:** Investigation, Software, Visualization, Writing - original draft, Writing - review & editing. **Faraz Sanal:** Investigation, Software. **Caitlin Richter:** Investigation. **Mark Liberman:** Supervision, Writing - review & editing. **Martijn Wieling:** Supervision, Methodology, Project administration, Writing - review & editing.

#### Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgments

The authors thank Hedwig Sekeres for creating the transcriptions of the Dutch speakers dataset, and Anna Pot for creating the visualization of the acoustic distance measure. We furthermore thank three anonymous

mous reviewers and the editor for their insightful comments, which have allowed us to improve this paper.

## Appendix A

Appendix A.1 provides all relevant technical details about the neural models used in this paper. Appendix A.2 visualizes the performance per layer for each of the neural models applied to different datasets.

### A.1. Technical details neural models

#### A.1.1. *wav2vec*

*wav2vec* ( $w2v$ ) is a self-supervised pre-trained neural model that has been developed for speech recognition (Schneider et al., 2019). This model consists of an encoder network and an aggregator network, and is trained in two model configurations, namely small and large. In this paper, we include the large model configuration to compute acoustic pronunciation distances, because the small model configuration is only trained on a subset of the Librispeech dataset, whereas the large model configuration uses the full Librispeech dataset.

The encoder network of the large model configuration consists of seven convolutional layers that create a dense representation of audio with a sliding window strategy (stride is 10 ms, window size is 30 ms). The dense output representations are aggregated by the aggregator network with 12 convolutional layers. The output of the encoder network is based on 30 ms of audio in steps of 10 ms, whereas the output of the aggregator network is based on windows of 810 ms.

$w2v$  is trained to predict upcoming audio frames of a speech utterance based on preceding frames. Inspired by *word2vec* (Mikolov et al., 2013), the model is trained with a contrastive loss objective, which is defined as the probability of distinguishing the actual frame from ten negative example frames sampled from the same utterance. To this end,  $w2v$  should be sensitive to content in the actual target frame. A regular loss objective for target frame prediction would let the model learn to replicate features that are consistent within their context, such as voice properties and noise. However, these features are not only undesirable, but they are also likely to be present as negative evidence in random negative samples from the same fragment. Therefore, a contrastive loss objective is more likely to reach better performance (Smith & Eisner, 2005).

During inference, speech features can be extracted from the encoder (512 dimensions) or the aggregator (512 dimensions). The encoder represents features within a 30 ms context window, whereas the aggregator outputs input reconstructions based on 810 ms of context. We select the features from the encoder, as initial experiments showed that this resulted in the highest performance.

#### A.1.2. *vq-wav2vec + BERT*

*vq-wav2vec* is an extension of  $w2v$  with the same architecture, except for the addition of a quantization layer between the encoder and the aggregator networks (Baevski, Schneider et al., 2020). This quantization layer creates a discrete representation of the dense encoder outputs. Quantization is done with either the Gumbel Softmax differentiable argmax

approach (Jang, Gu, & Poole, 2017), or with online K-means clustering (van den Oord, Vinyals, & Kavukcuoglu, 2017).

Discretization of  $w2v$  enables the use of algorithms that require discrete input, such as BERT. BERT is a non-recurrent neural network architecture, and training method, that can process sequential data (Devlin et al., 2019). Traditional neural methods require iterative processing of sequential data (i.e. recurrent neural networks), but the self-attention mechanism in the Transformer layers of BERT ensures that entire sequences can be processed at once (Vaswani et al., 2017). The self-attention mechanism works like a weighting mechanism for context in a sequence, and is based on content and position. Context-based representations are therefore only influenced by close context that is likely to be informative. The BERT model has shown to be highly scalable for text processing (Devlin et al., 2019).

Baevski, Schneider et al. (2020) applied a 12 layer BERT model to the discrete output of *vq-wav2vec*. The BERT model is trained by masking random spans of 100 ms of audio that have to be predicted as a pre-training objective, where each frame has a five percent chance of starting a masked sequence. The *vq-wav2vec* algorithm with the BERT extension was found to outperform the regular *vq-wav2vec* model on speech recognition tasks. Therefore, the BERT model may learn representations of speech that differentiate sounds in speech more clearly than the *vq-wav2vec* model itself. Moreover, Baevski, Schneider et al. (2020) show that, when applied to speech recognition, the full pipeline is slightly more effective with K-means quantization than Gumbel softmax quantization. We therefore use the *vq-wav2vec* algorithm with the BERT extension. We refer to this variant as *vqw2v*.

Speech representations can be extracted from multiple layers in the *vqw2v* pipeline. The *vqw2v* model itself can yield representations after encoding (512 dimensions), quantization (512 dimensions), and aggregation (512 dimensions). Additionally, the separately trained BERT model can provide representations after each of the 12 Transformer layers (768 dimensions). These Transformer layers can iteratively make embeddings more informative, but the final layers do not tend to be the most informative layers for downstream tasks (Tenney et al., 2019; de Vries et al., 2020). A likely explanation is that informative abstractions and generalizations in hidden layers are discarded in favor of actual target output. We choose the best-performing Transformer layer for our task based on a validation set from each dataset.

#### A.1.3. *wav2vec 2.0*

For the *vqw2v* algorithm, the Transformer layers in the BERT model are trained as a separate step after training *vqw2v*. For  $w2v2$ , the convolutional aggregator in *vqw2v* is replaced by a Transformer network (Baevski, Zhou et al., 2020). This has led to improved performance in speech recognition compared to *vqw2v* with BERT, suggesting that  $w2v2$  may contain better speech representations. Unlike *vqw2v* with the BERT extension,  $w2v2$  is trained as a single end-to-end model, and therefore the encoder outputs are optimized for use in the Transformer. The final pipeline of  $w2v2$  consists of a convolutional encoder, a quantizer, and a Transformer model. Gumbel softmax quantization is used in  $w2v2$ , and

the best-performing variant of  $w2v2$  in speech recognition contains a fixed amount of 24 Transformer layers.

Whereas the original  $w2v$  aggregator is trained to predict speech frames based on the preceding frames, the  $w2v2$  Transformer aggregator has to predict spans of randomly masked frames with the full fragment as context. The task of predicting single random frames is considered to be trivial, and therefore sequences of 10 consecutive frames are masked at each randomly sampled position with a probability of 6.5% for each frame to start a masked sequence. Effectively, during pre-training, 49% of all frames are masked in blocks with an average duration of 299 ms. Similarly to  $w2v$ , the  $w2v2$  model is trained with a contrastive loss function based on negative sampling.

Pre-trained models can be fine-tuned for speech recognition using labeled data. The models are augmented by adding a randomly initialized linear projection to the Transformer network. This projection contains an amount of classes that is similar to the size of the vocabulary of the labeled data. Using connectionist temporal classification (Graves, Fernández, Gomez, & Schmidhuber, 2006), probability scores of a textual output sequence can be obtained based on the vocabulary of the task.

Similar to  $vqw2v$ , embeddings can be extracted from the encoder (512 dimensions), the quantizer (768 dimensions), or the (fine-tuned) Transformer layers (1024 dimensions). We investigate the monolingual English  $w2v2$  model pre-trained on LS960, and a version that has subsequently been fine-tuned on 960 hours on labeled data from Librispeech (denoted by  $w2v2-en$ ). These models are chosen because they are the largest models available, and Baeviski, Zhou et al. (2020) showed that increasing the model size improved performance on all evaluated speech recognition tasks. We select the best-performing Transformer layer based on a validation set from each dataset.

#### A.1.4. XLSR

XLSR builds on  $w2v2$  by extending pre-training to 56,000 hours of speech from a total of 53 languages (Conneau et al., 2020). These 53 languages are obtained from the Common Voice dataset (which contains read speech from 36 European languages; Ardila et al., 2020), the BABEL dataset (conversational telephone speech from 17 Asian and African languages; Gales, Knill, Ragni, & Rath, 2014), and Multilingual Librispeech (audio books from 8 European languages; Pratap, Qiantong, Sriram, Synnaeve, & Collobert, 2020). Note that the majority of the pre-training data consists of English speech from Multilingual Librispeech dataset (44,000 hours), and that some languages occur in more than one dataset. Consequently, the total number of languages included during pre-training is 53. Due to the multilingual setup of the XLSR model, we expect improved task performance when using our Norwegian dataset, compared to the monolingual English models.

The architecture of XLSR is similar to  $w2v2$ , with the exception that a single set of discrete speech representations is learned by the quantizer on the basis of the encoder output. The discrete representations are subsequently shared across languages, creating connections between the different pre-training languages.

Similar to  $w2v2$ , embeddings can be extracted from the encoder (512 dimensions), the quantizer (768 dimensions), or the (fine-tuned) Transformer layers (1024 dimensions). We use the multilingual XLSR model pre-trained on 53 languages, and fine-tuned on languages from the Common Voice dataset (version 6.1) (Ardila et al., 2020). Specifically, we consider XLSR fine-tuned on English (1,686 hours, denoted by XLSR-en) and Swedish (12 hours, denoted by XLSR-sv), as they match (or, in the case of Swedish, is most similar to) the languages in our evaluation datasets. As before, we select the best-performing Transformer layer based on a validation set from each dataset.

#### A.1.5. DeCoAR

The  $w2v$  model uses convolutional layers to create representations of audio based on close context, whereas newer Transformer-based models use the entire audio fragment as context. DeCoAR uses an alternative method to process the audio sequences. Before BERT models were used in natural language processing, language models relied on recurrent neural models that process items in a sequence, one step at a time. Representations of each item are, in this case, based on the preceding representation. The most commonly used model in natural language processing that uses this method is ELMo (Peters et al., 2018). This model uses a stacked LSTM network for creating contextualized word embeddings.

Ling et al. (2020) apply the bi-directional LSTM method, that was proposed by Peters et al. (2018), to encode acoustic speech signals. The resulting DeCoAR model takes 40-dimensional log filterbank features as its input, and is trained to reconstruct the same features as its output. A filterbank transformation subsequently extracts frequency bands by dividing the frequency range into 40 triangular overlapping bins. These features are extracted with a 25 ms sliding window and a stride of 10 ms. DeCoAR consists of four bi-directional LSTM layers, each having 1024 cells. The output representation of DeCoAR is the concatenation of the forward and backward directions, and therefore consists of 2048 dimensions.

The novel DeCoAR model was shown to outperform  $w2v$  on a set of tasks, including phone classification (Ma, Ryant, & Liberman, 2021). Therefore, architectural differences of DeCoAR with  $w2v$ -based models may show performance differences when applied to other tasks, such as modeling speech variation.

#### A.2. Layer performance

The performance per layer for all neural models and datasets can be found in Fig. 7, Fig. 8, Fig. 9.

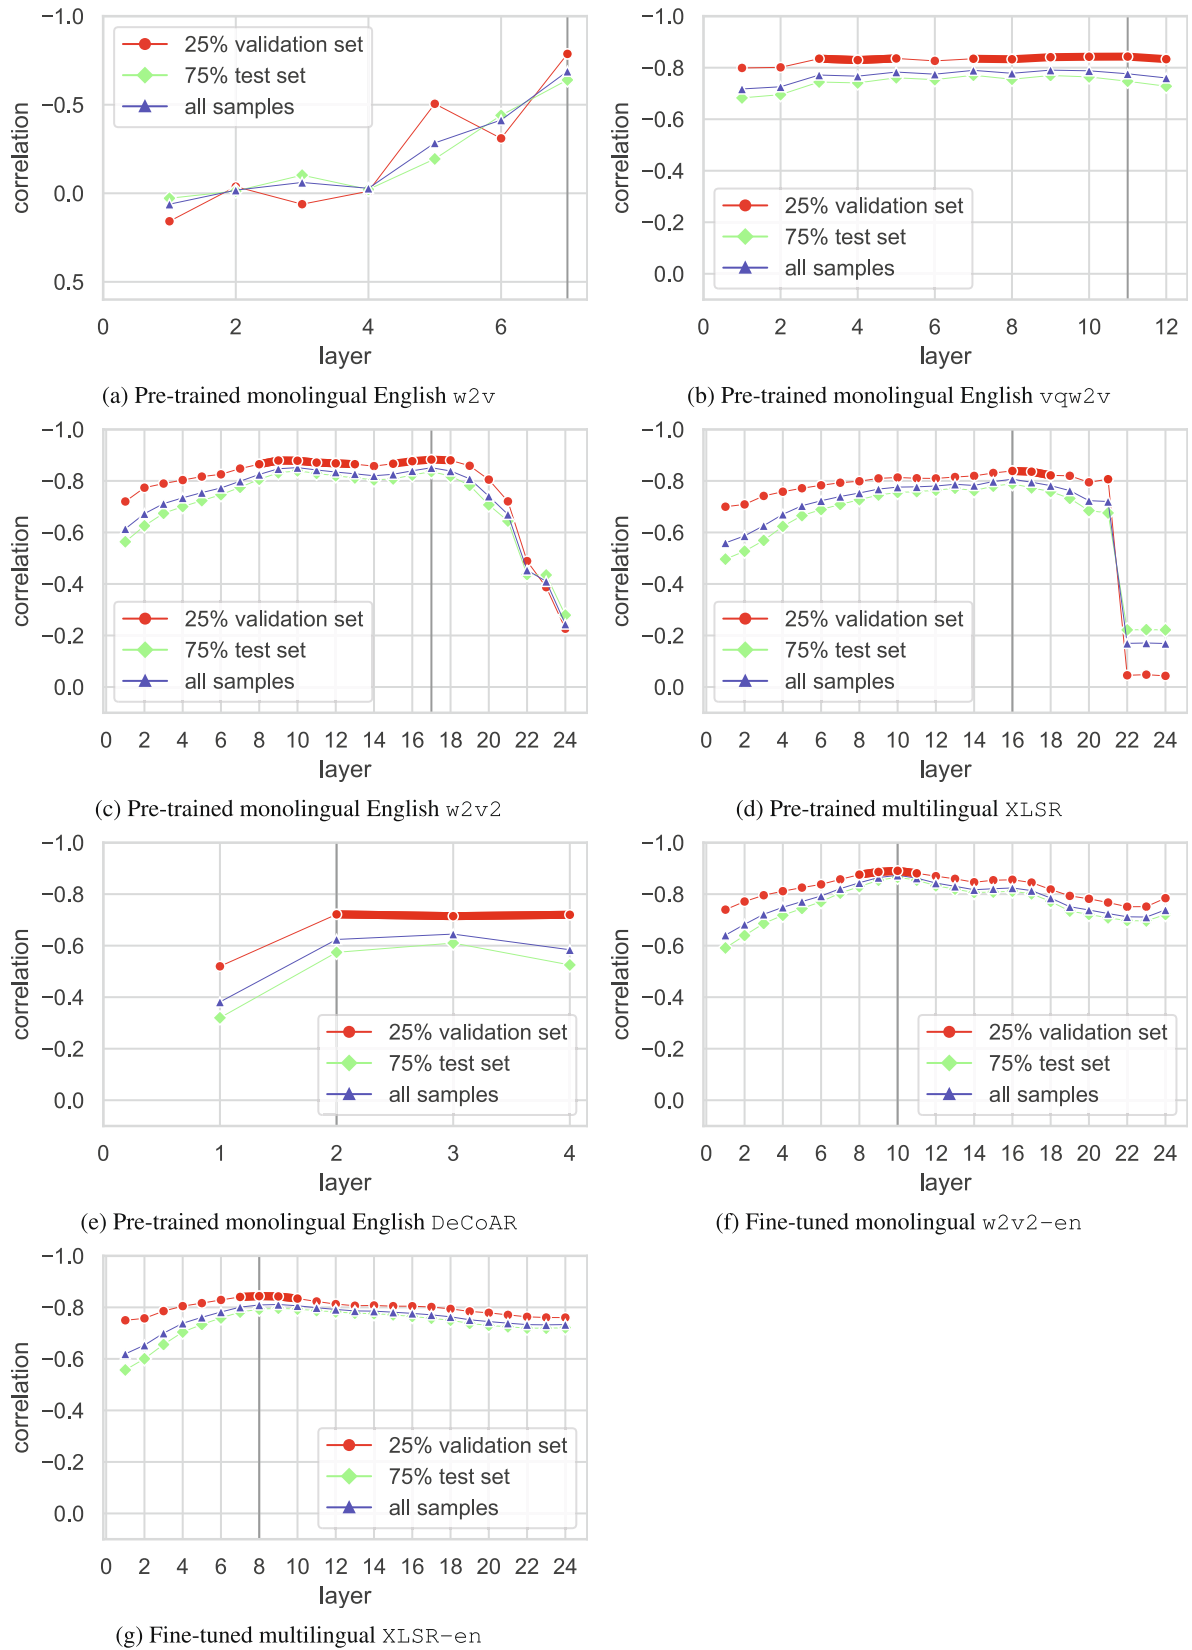
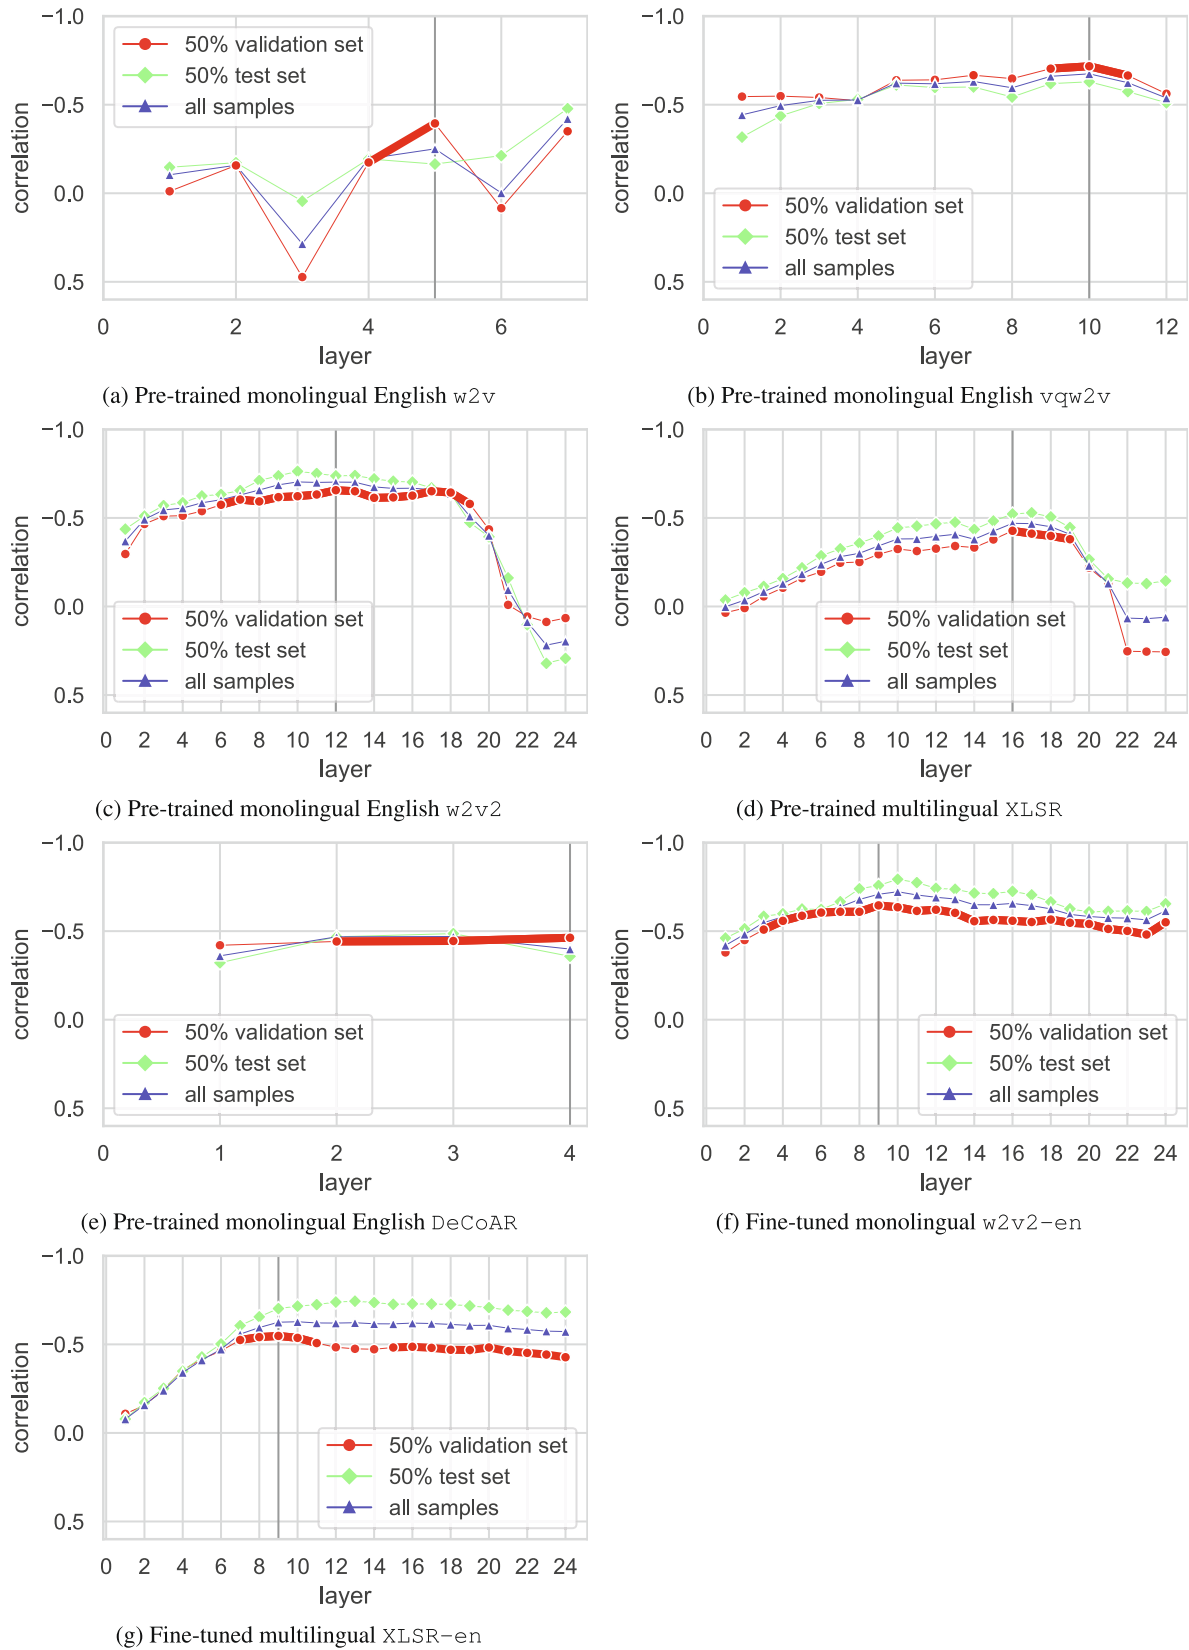
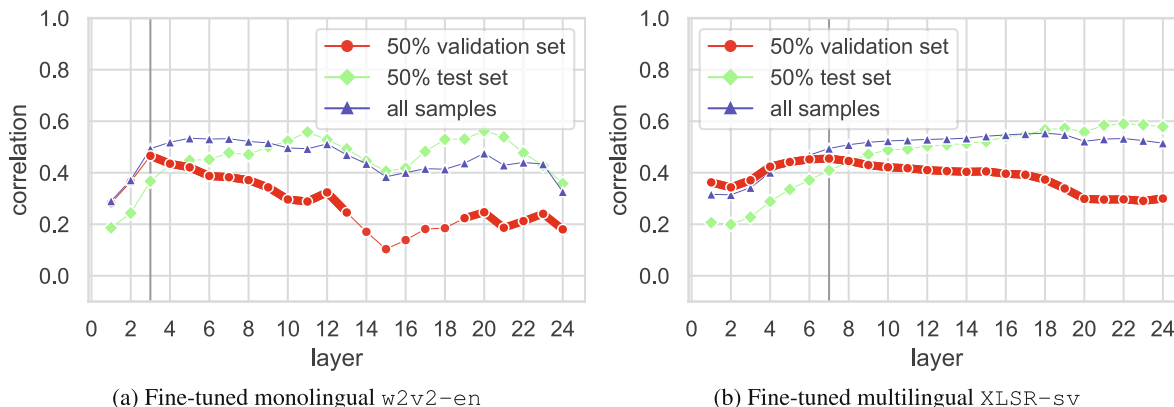


Fig. 7. Pearson correlation coefficients of acoustic distances compared to human accent ratings for different layers in w2v, vqw2v, DeCoAR, and the w2v2 and XLSR models. The vertical line marks the layer that was chosen as the best-performing layer based on the 25% validation set of the Speech Accent Archive dataset. Layers with a correlation that is not significantly different ( $p > 0.05$ ) from the optimal layer are indicated by the thick red line.



**Fig. 8.** Pearson correlation coefficients of acoustic distances compared to human accent ratings for different layers in  $w2v$ ,  $vqw2v$ , DeCoAR, and the  $w2v2$  and XLSR models. The vertical line marks the layer that was chosen as the best-performing layer based on the 50% validation set of the Dutch speakers dataset. Layers with a correlation that is not significantly different ( $p > 0.05$ ) from the optimal layer are indicated by the thick red line.



**Fig. 9.** Pearson correlation coefficients of acoustic distances compared to human accent ratings for different Transformer layers in the `w2v2-en` and `XLSR-sv` models. The vertical line marks the layer that was chosen as the best-performing layer based on the 50% validation set of the Norwegian dialects dataset. Layers with a correlation that is not significantly different ( $p > 0.05$ ) from the optimal layer are indicated by the thick red line.

## References

- Ardila, R., Branson, M., Davis, K., Kohler, M., Meyer, J., Henretty, M., et al. (2020). Common voice: A massively-multilingual speech corpus. In *Proceedings of the 12th language resources and evaluation conference, Marseille, France* (pp. 4218–4222). European Language Resources Association.
- Asher, J. J., & Garcia, R. (1969). The optimal age to learn a foreign language. *The Modern Language Journal*, 53(5), 334–341.
- Babu, A., Wang, C., Tjandra, A., Lakhotia, K., Xu, Q., Goyal, N., et al. (2021). XLS-R: Self-supervised cross-lingual speech representation learning at scale.
- Baevski, A., Schneider, S., & Auli M. (2020). vq-wav2vec: Self-supervised learning of discrete speech representations.
- Baevski, A., Zhou, Y., Mohamed, A., & Auli, M. (2020). wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems* (Vol. 33, pp. 12449–12460). Curran Associates Inc.
- Bartelds, M., Richter, C., Liberman, M., & Wieling, M. (2020). A new acoustic-based pronunciation distance measure. *Frontiers in Artificial Intelligence*, 3, 39.
- Blodgett, S. L., Green, L., & O'Connor, B. (2016). Demographic dialectal variation in social media: A case study of African-American English. In *Proceedings of the 2016 conference on empirical methods in natural language processing* (pp. 1119–1130). Austin, TX: Association for Computational Linguistics.
- Boberg, C. (2010). *The English language in Canada: Status, history and comparative analysis*. Studies in English language. Cambridge University Press.
- Bucholtz, M. (2007). Variation in transcription. *Discourse Studies*, 9(6), 784–808.
- Chen, N. F., Wee, D., Tong, R., Ma, B., & Li, H. (2016). Large-scale characterization of non-native Mandarin Chinese spoken by speakers of European origin: Analysis on icall. *Speech Communication*, 84, 46–56.
- Church, K. W., & Hanks, P. (1990). Word association norms, mutual information, and lexicography. *Computational Linguistics*, 16(1), 22–29.
- Conneau, A., Baevski, A., Collobert, R., Mohamed, A., Auli, M. (2020). Unsupervised cross-lingual representation learning for speech recognition.
- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3), 297–334.
- Cucchiarini, C. (1996). Assessing transcription agreement: Methodological aspects. *Clinical Linguistics & Phonetics*, 10(2), 131–155.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the north american chapter of the association for computational linguistics: Human language technologies, volume 1 (long and short papers)* (pp. 4171–4186). Minneapolis, MN: Association for Computational Linguistics.
- Duckworth, M., Allen, G., Hardcastle, W., & Ball, M. (1990). Extensions to the international phonetic alphabet for the transcription of atypical speech. *Clinical Linguistics & Phonetics*, 4(4), 273–280.
- Ferragne, E., & Pellegrino, F. (2010). Vowel systems and accent similarity in the British Isles: Exploiting multidimensional acoustic distances in phonetics. *Journal of Phonetics*, 38(4), 526–539.
- Flege, J. E. (1988). Factors affecting degree of perceived foreign accent in English sentences. *The Journal of the Acoustical Society of America*, 84(1), 70–79.
- Gales, M. J. F., Knill, K. M., Ragni, A., & Rath, S. P. (2014). Speech recognition and keyword spotting for low-resource languages: Babel project research at Cued. In *Fourth international workshop on spoken language technologies for under-resourced languages (SLTU-2014)* (pp. 16–23). International Speech Communication Association (ISCA).
- Garofalo, J. S., Graff, D., Paul, D., & Pallett, D. (2007). CSR-I (WSJ0) Complete LDC93S6A. Web Download. Philadelphia: Linguistic Data Consortium.
- Giorgino, T. (2009). Computing and visualizing dynamic time warping alignments in R: The dtw package. *Journal of Statistical Software*, 31(7), 1–24.
- Gooskens, C. (2005). How well can Norwegians identify their dialects? *Nordic Journal of Linguistics*, 28(1), 37–60.
- Gooskens, C., & Heeringa, W. (2004). Perceptive evaluation of levenshtein dialect distance measurements using norwegian dialect data. *Language Variation and Change*, 16(3), 189–207.
- Graves, Alex, Fernández, Santiago, Gomez, Faustino, & Schmidhuber, Jürgen (2006). Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd International Conference on Machine Learning, ICML '06* (pp. 369–376). New York, NY, USA: Association for Computing Machinery.
- Hakkani-Tür, D., Riccardi, G., & Gorin, A. (2002). Active learning for automatic speech recognition. *2002 IEEE international conference on acoustics, speech, and signal processing* (Vol. 4). IEEE, pages IV–3904.
- Hanani, A., Russell, M. J., & Carey, M. J. (2013). Human and computer recognition of regional accents and ethnic groups from British English speech. *Computer Speech & Language*, 27(1), 59–74. Special issue on Paralinguistics in Naturalistic Speech and Language.
- Heeringa, W., Johnson, K., & Gooskens, C. (2009). Measuring norwegian dialect distances using acoustic features. *Speech Communication*, 51(2), 167–183.
- Heeringa, W. J. (2004). *Measuring dialect pronunciation differences using levenshtein distance* (Ph.D. thesis). University of Groningen.
- Huang, Y., Slaney, M., Seltzer, M. L., & Gong, Y. (2014). Towards better performance with heterogeneous training data in acoustic modeling using deep neural networks. In *Fifteenth annual conference of the international speech communication association*.
- Huang, Y., Dong, Y., Liu, C., & Gong, Y. (2014). A comparative analytic study on the gaussian mixture and context dependent deep neural network hidden markov models. In *Fifteenth annual conference of the international speech communication association*.
- Huckvale, M. (2007). *ACCDIST: An accent similarity metric for accent recognition and diagnosis* (pp. 258–275). Berlin Heidelberg, Berlin, Heidelberg: Springer.
- Jang, E., Gu, S., & Poole, B. (2017). Categorical reparameterization with Gumbel-Softmax. arXiv:1611.01144 [cs, stat].
- Jeszczynski, P., Stoeckle, P., Glaser, E., & Weibel, R. (2017). Exploring global and local patterns in the correlation of geographic distances and morphosyntactic variation in Swiss German. *Journal of Linguistic Geography*, 5(2), 86–108.
- Kahn, J., Rivière, M., Zheng, W., Kharitonov, E., Xu, Q., Mazaré, P.-E., et al. (2020). Libri-Light: A benchmark for ASR with limited or no supervision. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 7669–7673. IEEE.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., et al. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14), 7684–7689.
- Koster, C. J., & Koet, T. (1993). The evaluation of accent in the English of Dutchmen. *Language Learning*, 43(1), 69–92.
- Leather, J. (1983). Second-language pronunciation learning and teaching. *Language Teaching*, 16(3), 198–219.
- Levenshtein, V. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Cybernetics and Control Theory*, 10(8), 707–710.
- Liberman, M. (2018). Towards progress in theories of language sound structure. In D. Brentari & J. L. Lee (Eds.), *Shaping Phonology*. University of Chicago Press.
- Ling, S., Liu, Y., Salazar, J., & Kirchhoff, K. (2020). Deep contextualized acoustic representations for semi-supervised speech recognition. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6429–6433).
- Livescu, K., & Glass, J. (2000). Lexical modeling of non-native speech for automatic speech recognition. In *2000 IEEE international conference on acoustics, speech, and signal processing*. Proceedings (Cat. No.00CH37100). Vol. 3, pp. 1683–1686.
- Ma, D., Ryant, N., & Liberman, M. (2021). Probing acoustic representations for phonetic properties. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 311–315).



- Magen, H. S. (1998). The perception of foreign-accented speech. *Journal of Phonetics*, 26(4), 381–400.
- Mathur, A., Isopoussu, A., Kawsar, F., Berthouze, N., & Lane, N. D. (2019). Mic2Mic: Using cycle-consistent generative adversarial networks to overcome microphone variability in speech systems. In *Proceedings of the 18th international conference on information processing in sensor networks*. New York, NY, USA: Association for Computing Machinery.
- Mermelstein, P. (1976). Distance measures for speech recognition, psychological and instrumental. *Pattern Recognition and Artificial Intelligence*, 116, 374–388.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems* (Vol. 26). Curran Associates Inc.
- Munro, M. J. (1995). Nonsegmental factors in foreign accent: Ratings of filtered speech. *Studies in Second Language Acquisition*, 17(1), 17–34.
- Munro, M. J., & Derwing, T. M. (1994). Evaluations of foreign accent in extemporaneous and read material. *Language Testing*, 11(3), 253–266.
- Munro, M. J., & Derwing, T. M. (2001). Modeling perceptions of the accentedness and comprehensibility of L2 speech the role of speaking rate. *Studies in Second Language Acquisition*, 23(4), 451–468.
- Najafian, M., DeMarco, A., Cox, S., & Russell, M. (2014). Unsupervised model selection for recognition of regional accented speech. In *Fifteenth annual conference of the international speech communication association*.
- Nerbonne, J., & Heeringa, W. (1997). Measuring dialect distance phonetically. In *Computational Phonology: Third meeting of the ACL special interest group in computational phonology*.
- Novotney, S., & Callison-Burch, C. (2010). Cheap, fast and good enough: Automatic speech recognition with non-expert transcription. In *Human Language Technologies: The 2010 annual conference of the North American Chapter of the Association for Computational Linguistics* (pp. 207–215). Los Angeles, CA: Association for Computational Linguistics.
- Nunnally, J. C. (1978). *Psychometric theory*. McGraw-Hill series in psychology. New York: McGraw-Hill.
- Offrede, T. F., Jacobi, J., Rebernik, T., de Jong, L., Keulen, S., Veenstra, P., et al. (2020). The impact of alcohol on I1 versus I2. *Language and Speech*. PMID: 32856992.
- van den Oord, A., Li, Y., & Vinyals, O. (2019). Representation learning with contrastive predictive coding.
- van den Oord, A., Vinyals, O., & Kavukcuoglu, K. (2017). Neural discrete representation learning. *Advances in neural information processing systems* (Vol. 30). Curran Associates Inc.
- Panayotov, V., Chen, G., Povey, D., & Khudanpur, S. (2015). Librispeech: An ASR corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 5206–5210).
- Paul, D. B., & Baker, J. M. (1992). The design for the Wall Street Journal-based CSR corpus. In *Speech and natural language: Proceedings of a workshop held at Harriman, New York, February 23–26, 1992*.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., et al. (2018). Deep contextualized word representations. In *Proceedings of the 2018 conference of the North American chapter of the Association for Computational Linguistics: Human language technologies, volume 1 (long papers)* (pp. 2227–2237). New Orleans, Louisiana: Association for Computational Linguistics.
- Pratap, V., Qiantong, X., Sriram, A., Synnaeve, G., & Collobert, R. (2020). MLS: A large-scale multilingual dataset for speech research. *Proceedings of Interspeech, 2020*, 2757–2761.
- Dennis Richard Preston (1999). *Handbook of perceptual dialectology*. John Benjamins Publishing Co.
- San, N., Bartelds, M., Browne, M., Clifford, L., Gibson, F., Mansfield, J., et al. (2021). Leveraging pre-trained representations to improve access to untranscribed speech from endangered languages.
- Scharenborg, O. (2007). Reaching over the gap: A review of efforts to link human and automatic speech recognition research. *Speech Communication*, 49(5), 336–347. Bridging the Gap between Human and Automatic Speech Recognition.
- Schneider, S., Baevski, A., Collobert, R., & Auli, M. (2019). wav2vec: Unsupervised pre-training for speech recognition. *Proceedings of Interspeech, 2019*, 3465–3469.
- Smith, N. A., & Eisner, J. (2005). Contrastive estimation: Training log-linear models on unlabeled data. In *Proceedings of the 43rd annual meeting of the Association for Computational Linguistics (ACL'05)* (pp. 354–362). Ann Arbor, MI: Association for Computational Linguistics.
- Steiger, J. H. (1980). Tests for comparing elements of a correlation matrix. *Psychological Bulletin*, 87(2), 245.
- Strycharczuk, P., López-Ibáñez, M., Brown, G., & Leemann, A. (2020). General northern English. Exploring regional variation in the north of England with machine learning. *Frontiers in Artificial Intelligence*, 3(48).
- Tenney, I., Das, D., & Pavlick, E. (2019). BERT rediscovers the classical NLP pipeline. In *Proceedings of the 57th annual meeting of the association for computational linguistics* (pp. 4593–4601). Florence, Italy: Association for Computational Linguistics.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. *Advances in neural information processing systems* (Vol. 30). Curran Associates, Inc.
- Vigino, T., Motlicek, P., & Cernak, M. (2019). End-to-end accented speech recognition. *Proceedings of Interspeech, 2019*, 2140–2144.
- de Vries, W., van Cranenburgh, A., & Nissim, M., (2020). What's so special about BERT's layers? A closer look at the NLP pipeline in monolingual and multilingual models. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4339–4350. Online. Association for Computational Linguistics.
- Weinberger, S. H., & Kunath, S. A. (2011). The speech accent archive: Towards a typology of english accents. In *Corpus-based studies in language use, language learning, and language documentation* (pp. 265–281). Brill Rodopi.
- Wieling, M., Jelke Bloem, R., Baayen, H., & Nerbonne, J. (2014). Determinants of English accents. In *Proceedings of the 6th conference on quantitative investigations in theoretical linguistics*.
- Wieling, M., Bloem, J., Mignella, K., Timmermeister, M., & Nerbonne, J. (2014). Measuring foreign accent strength in English: Validating levenshtein distance as a measure. *Language Dynamics and Change*, 4(2), 253–269.
- Wieling, M., Heeringa, W., & Nerbonne, J. (2007). An aggregate analysis of pronunciation in the goeman-taeldeman-van reenen-project data. *Taal en Tongval*, 59(1), 84–116.
- Wieling, M., Margaretha, E., & Nerbonne, J. (2012). Inducing a measure of phonetic similarity from pronunciation variation. *Journal of Phonetics*, 40(2), 307–314.
- Wieling, M., & Nerbonne, J. (2015). Advances in dialectometry. *Annual Review of Linguistics*, 1(1), 243–264.
- Wieling, M., Nerbonne, J., & Harald Baayen, R. (2011). Quantitative social dialectology: Explaining linguistic variation geographically and socially. *PLoS One*, 6(9), 1–14.
- Yuan, J., & Liberman, M. (2008). Speaker identification on the scotus corpus. *Journal of the Acoustical Society of America*, 123(5), 3878.
- Zhao, X., & Wang, D. L. (2013). Analyzing noise robustness of mfcc and gfcc features in speaker identification. In *2013 IEEE international conference on acoustics, speech and signal processing* (pp. 7204–7208).
- Želasko, P., Moro-Velázquez, L., Hasegawa-Johnson, M., Scharenborg, O., & Dehak, N. (2020). That sounds familiar: An analysis of phonetic representations transfer across languages. *Proceedings of Interspeech, 2020*, 3705–3709.