

Computer-assisted pronunciation training—Speech synthesis is almost all you need

Daniel Korzekwa^{a,b,*}, Jaime Lorenzo-Trueba^a, Thomas Drugman^a, Bozena Kostek^b

^a Amazon Speech Research, Poland

^b Gdansk University of Technology, Faculty of ETI, Poland

ARTICLE INFO

Keywords:

Computer-assisted pronunciation training
Automated pronunciation error detection
Automated lexical stress error detection
Speech synthesis
Voice conversion
Deep learning

ABSTRACT

The research community has long studied computer-assisted pronunciation training (CAPT) methods in non-native speech. Researchers focused on studying various model architectures, such as Bayesian networks and deep learning methods, as well as on the analysis of different representations of the speech signal. Despite significant progress in recent years, existing CAPT methods are not able to detect pronunciation errors with high accuracy (only 60% precision at 40%–80% recall). One of the key problems is the low availability of mispronounced speech that is needed for the reliable training of pronunciation error detection models. If we had a generative model that could mimic non-native speech and produce any amount of training data, then the task of detecting pronunciation errors would be much easier. We present three innovative techniques based on phoneme-to-phoneme (P2P), text-to-speech (T2S) and speech-to-speech (S2S) conversion to generate correctly pronounced and mispronounced synthetic speech. We show that these techniques not only improve the accuracy of three machine learning models for detecting pronunciation errors, but also help establish a new state-of-the-art in the field. Earlier studies have used simple speech generation techniques such as P2P conversion, but only as an additional mechanism to improve the accuracy of pronunciation error detection. We, on the other hand, consider speech generation to be the first-class method of detecting pronunciation errors. The effectiveness of these techniques is assessed in the tasks of detecting pronunciation and lexical stress errors. Non-native English speech corpora of German, Italian, and Polish speakers are used in the evaluations. The best proposed S2S technique improves the accuracy of detecting pronunciation errors in AUC metric by 41% from 0.528 to 0.749 compared to the state-of-the-art approach.

1. Introduction

Language plays a key role in online education, giving people access to large amounts of information contained in articles, books, and video lectures. Thanks to spoken language and other forms of communication, such as a sign-language, people can participate in interactive discussions with teachers and take part in lively brainstorming with other people. Unfortunately, education is not available to everybody. According to the UNESCO report, 40% of the global population do not have access to education in the language they understand (UNESCO, 2016). ‘If you don’t understand, how can you learn?’ the report says. English is the leading language on the Internet, representing 25.9% of the world’s population (Statista, 2021). Regrettably, research by EF (Education First) (EF-Education-First, 2020) shows a large disproportion in English proficiency across countries and continents. People from regions of ‘very low’ language proficiency, such as the Middle East, are unable to navigate through English-based websites or communicate with people from an English-speaking country.

Computer-Assisted Language Learning (CALL) helps to improve the English language proficiency of people in different regions (Levy and Stockwell, 2013). CALL relies on computerized self-service tools that are used by students to practice a language, usually a foreign language, also known as a non-native (L2) language. Students can practice multiple aspects of the language, including grammar, vocabulary, writing, reading, and speaking. Computer-based tools can also be used to measure student’s language skills and their learning potential by using Computerized Dynamic Assessment (C-DA) test (Mehri Kamrood et al., 2019). CALL can complement traditional language learning provided by teachers. It also has a chance to make second language learning more accessible in scenarios where traditional ways of learning languages are not possible due to the cost of learning or the lack of access to foreign language teachers.

Computer-Assisted Pronunciation Training (CAPT) is a part of CALL responsible for learning pronunciation skills. It has been shown to help

* Corresponding author at: Amazon Speech Research, Poland.

E-mail address: daniel.korzekwa@gmail.com (D. Korzekwa).

people practice and improve their pronunciation skills (Neri et al., 2008; Golonka et al., 2014; Tejedor-García et al., 2020). CAPT consists of two components: an automated pronunciation evaluation component (Leung et al., 2019; Zhang et al., 2021; Korzekwa et al., 2021c) and a feedback component (Ai, 2015). The automated pronunciation evaluation component is responsible for detecting pronunciation errors in spoken speech, for example, for detecting words pronounced incorrectly by the speaker. The feedback component informs the speaker about mispronounced words and advises how to pronounce them correctly. This article is devoted to the topic of automated detection of pronunciation errors in non-native speech. This area of CAPT can take advantage of technological advances in machine learning and bring us closer to creating a fully automated assistant based on artificial intelligence for language learning.

The research community has long studied the automated detection of pronunciation errors in non-native speech. Existing work has focused on various tasks such as detecting mispronounced phonemes (Leung et al., 2019) and lexical stress errors (Ferrer et al., 2015). Researchers have given most attention to studying various machine learning models such as Bayesian networks (Witt and Young, 2000; Li et al., 2011) and deep learning methods (Leung et al., 2019; Zhang et al., 2021), as well as analyzing different representations of the speech signal such as prosodic features (duration, energy and pitch) (Chen and Wang, 2010), and cepstral/spectral features (Ferrer et al., 2015; Shahin et al., 2016; Leung et al., 2019). Despite significant progress in recent years, existing CAPT methods detect pronunciation errors with relatively low accuracy of 60% precision at 40%–80% recall (Leung et al., 2019; Korzekwa et al., 2021c; Zhang et al., 2021). Highlighting correctly pronounced words as pronunciation errors by the CAPT tool can demotivate students and lower the confidence in the tool. Likewise, missing pronunciation errors can slow down the learning process.

One of the main challenges with the existing CAPT methods is poor availability of mispronounced speech, which is required for the reliable training of pronunciation error detection models. We propose a reformulation of the problem of pronunciation error detection as a task of synthetic speech generation. Intuitively, if we had a generative model that could mimic mispronounced speech and produce any amount of training data, then the task of detecting pronunciation errors would be much easier. The probability of pronunciation errors for all the words in a sentence can then be calculated using the Bayes rule (Bishop, 2006). In this new formulation, we move the complexity to learning the speech generation process that is well suited to the problem of limited speech availability (Huybrechts et al., 2021; Shah et al., 2021; Fazel et al., 2021). The proposed method outperforms the state-of-the-art model (Leung et al., 2019) in detecting pronunciation errors in AUC metric by 41% from 0.528 to 0.749 on the GUT Isle Corpus of L2 Polish speakers.

To put the new formulation of the problem into action, we propose three innovative techniques based on phoneme-to-phoneme (P2P), text-to-speech (T2S) and speech-to-speech (S2S) conversion to generate correctly pronounced and mispronounced synthetic speech. We show that these techniques not only improve the accuracy of three machine learning models for detecting pronunciation errors, but also help establish a new state-of-the-art in the field. The effectiveness of these techniques is assessed in two tasks: detecting mispronounced words (replacing, adding, removing phonemes, or pronouncing an unknown speech sound) and detecting lexical stress errors. The results presented in this study are the culmination of our recent work on speech generation in pronunciation error detection task (Korzekwa et al., 2021c,b,a), including a new S2S technique.

In short, the contributions of the paper are as follows:

- A new paradigm for the automated detection of pronunciation errors is proposed, reformulating the problem as a task of generating synthetic speech.
- A unified probabilistic view on P2P, T2S, and S2S techniques is presented in the context of detecting pronunciation errors.

- A new S2S method to generate synthetic speech is proposed, which outperforms the state-of-the-art model (Leung et al., 2019) in detecting pronunciation errors.
- Comprehensive experiments are described to demonstrate the effectiveness of speech generation in the tasks of pronunciation and lexical stress error detection.

The outline of the rest of this paper is: Section 2 presents related work. Section 3 describes the proposed methods of generating synthetic speech for automatic detection of pronunciation errors. Section 4 describes the human speech corpora used to train the pronunciation error detection models in the experiments. Section 5 presents experiments demonstrating the effectiveness of various synthetic speech generation methods in improving the accuracy of the detection of pronunciation and lexical stress errors. Finally, conclusions and future work are presented in Section 6.

2. Related work

2.1. Pronunciation error detection

2.1.1. Phoneme recognition approaches

Most existing CAPT methods are designed to recognize the phonemes pronounced by the speaker and compare them with the expected (canonical) pronunciation of correctly pronounced speech (Witt and Young, 2000; Li et al., 2016; Sudhakara et al., 2019b; Leung et al., 2019). Any discrepancy between the recognized and canonical phonemes results in a pronunciation error at the phoneme level. Phoneme recognition approaches generally fall into two categories: methods that align a speech signal with phonemes (forced-alignment techniques) and methods that first recognize the phonemes in the speech signal and then align the recognized and canonical phoneme sequences. Aside these two categories, CAPT methods can be split into multiple other categories:

Forced-alignment techniques (Li et al., 2011, 2016; Sudhakara et al., 2019b; Cheng et al., 2020) are based on the work of Franco et al. (1997) and the Goodness of Pronunciation (GoP) method (Witt and Young, 2000). In the first step, GoP uses Bayesian inference to find the most likely alignment between canonical phonemes and the corresponding audio signal (forced alignment). In the next step, GoP calculates the ratio between the likelihoods of the canonical and the most likely pronounced phonemes. Finally, it detects mispronunciation if the ratio drops below a certain threshold. GoP has been further extended with Deep Neural Networks (DNNs), replacing the Hidden Markov Model (HMM) and Gaussian Mixture Model (GMM) techniques for acoustic modeling (Li et al., 2016; Sudhakara et al., 2019b). Cheng et al. (2020) improves GoP performance with the hidden representation of speech extracted in an unsupervised way. This model can detect pronunciation errors based on the input speech signal and the reference canonical speech signal, without using any linguistic information such as text and phonemes.

The methods that do not use forced-alignment recognize the phonemes pronounced by the speaker purely from the speech signal and only then align them with the canonical phonemes (Minematsu, 2004; Harrison et al., 2009; Lee and Glass, 2013; Plantinga and Fosler-Lussier, 2019; Sudhakara et al., 2019a; Zhang et al., 2020). Leung et al. (2019) use a phoneme recognizer that recognizes phonemes only from the speech signal. The phoneme recognizer is based on Convolutional Neural Network (CNN), a Gated Recurrent Unit (GRU), and Connectionist Temporal Classification (CTC) loss. Leung et al. report that it outperforms other forced-alignment (Li et al., 2016) and forced-alignment-free (Harrison et al., 2009) techniques in the task of detecting mispronunciations at the phoneme-level in L2 English.

There are two challenges with presented approaches for pronunciation error detection. First, phonemes pronounced by the speaker must be recognized accurately, which has been proved difficult (Zhang

et al., 2021; Chorowski et al., 2014, 2015; Bahdanau et al., 2016). Phoneme recognition is difficult, especially in non-native speech, as different languages have different phoneme spaces. Second, standard approaches assume only one canonical pronunciation of a given text, but this assumption is not always true due to the phonetic variability of speech, e.g., differences between regional accents. For example, the word ‘enough’ can be pronounced by native speakers in multiple ways: /ih n ah f/ or /ax n ah f/ (short ‘i’ or ‘schwa’ phoneme at the beginning). In our previous work, we solve these problems by creating a native speech pronunciation model that returns the probability of the sentence to be spoken by a native speaker (Korzekwa et al., 2021c).

Techniques based on phoneme recognition can be supplemented by a reference speech signal obtained from the speech database (Xiao et al., 2018; Nicolao et al., 2015; Wang et al., 2019) or generated from the phonetic representation (Korzekwa et al., 2021c; Qian et al., 2010). Xiao et al. (2018) use a pair of speech signals from a student and a native speaker to classify native and non-native speech. Nicolao et al. (2015) use the speech of the reference speaker to detect mispronunciation errors at the phoneme level. Wang et al. (2019) use Siamese networks to model the discrepancy between normal and distorted children’s speech. Qian et al. (2010) propose a statistical model of pronunciation in which they build a model that generates hypotheses of mispronounced speech.

In this work, we use the end-to-end method to detect pronunciation errors directly, without having to recognize phonemes as an intermediate step. The end-to-end approach is discussed in more detail in the next section.

2.1.2. End-to-end methods

The phoneme recognition approaches presented so far rely on phonetically transcribed speech labeled by human listeners. Phonetic transcriptions are needed to train a phoneme recognition model. Human-based transcription is a time-consuming task, especially with L2 speech, where listeners need to recognize mispronunciation errors. Sometimes L2 speech transcription may be even impossible because different languages have different phoneme sets, and it is unclear which phonemes were pronounced by the speaker. In our recent work, we have introduced a novel model (known as WEAKLY-S, i.e., weakly supervised) for detecting pronunciation errors at the word level that does not require phonetically transcribed L2 speech (Korzekwa et al., 2021b). During training, the model is weakly supervised, in the sense that in L2 speech, only mispronounced words are marked, and the data do not need to be phonetically transcribed. In addition to the primary task of detecting mispronunciation errors at the word level, the second task uses a phoneme recognizer trained on automatically transcribed L1 speech.

Zhang et al. (2021) employ a multi-task model with two tasks: phoneme-recognition and pronunciation error detection tasks. Unlike our WEAKLY-S model, they use the Needleman–Wunsch algorithm (Needleman and Wunsch, 1970) from bioinformatics to align the canonical and recognized phoneme sequences, but this algorithm cannot be tuned to detect pronunciation errors. The WEAKLY-S model automatically learns the alignment, thus eliminating a potential source of inaccuracy. The alignment is learned through an attention mechanism that automatically maps the speech signal to a sequence of pronunciation errors at the word level. Wang et al. (2019) propose to use a multi-task framework in which a neural network model is used to learn the joint space between the acoustic characteristics of adults and children. Additionally, Duan et al. (2019) propose a multi-task model for acoustical modeling with two tasks for native and non-native speech respectively.

The work of Zhang et al. (2021) and our recent work (Korzekwa et al., 2021b) are end-to-end methods of direct estimation of pronunciation errors, setting up a new trend in the field of automated pronunciation assessment. In this article, we use the end-to-end method as well, but we extend it by the S2S method of generating mispronounced speech.

2.1.3. Other trends

All the works presented so far treat pronunciation errors as discrete categories, at best producing the probability of mispronunciation. In contrast, Yan et al. (2020) propose a model capable of identifying phoneme distortions, giving the user more detailed feedback on mispronunciation. In our recent work, we provide more fine-grained feedback by indicating the severity level of mispronunciation (Korzekwa et al., 2021b).

Active research is conducted not only on modeling techniques but also on speech representation. Xu et al. (2021) and Peng et al. (2021) use the Wav2vec 2.0 speech representation that is created in an unsupervised way. They report that it outperforms existing methods and requires three times less speech training data. Lin and Wang (2021) use transfer learning by taking advantage of deep latent features extracted from the Automated Speech Recognition (ASR) acoustic model and report improvements over the classic GOP-based method.

In this work, we use a mel-spectrogram as a speech representation in the pronunciation error detection model. We also use a mel-spectrogram to represent the speech signal in the T2S and S2S methods of generating mispronounced speech.

2.2. Lexical stress error detection

CAPT usually focuses on practicing the pronunciation of phonemes (Witt and Young, 2000; Leung et al., 2019; Korzekwa et al., 2021c). However, there is evidence that practicing lexical stress improves the intelligibility of non-native English speech (Field, 2005; Lepage and Busà, 2014). Lexical stress is a phonological feature of a syllable. It is part of the phonological rules that govern how words should be pronounced in a given language. Stressed syllables are usually longer, louder, and expressed with a higher pitch than their unstressed counterparts (Jung et al., 2018). The lexical stress is related to the phonemic representation. For example, placing lexical stress on a different syllable of a word can lead to various phonemic realizations known as ‘vowel reduction’ (Bergem, 1991). Students should be able to practice both pronunciation and lexical stress in spoken language. We study both topics to better understand the potential of using speech generation methods in CAPT.

The existing works focus on the supervised classification of lexical stress using Neural Networks (Li et al., 2018; Shahin et al., 2016), Support Vector Machines (Chen and Wang, 2010; Zhao et al., 2011), and Fisher’s linear discriminant (Chen and He, 2007). There are two popular variants: (a) discriminating syllables between primary stress/no stress (Ferrer et al., 2015), and (b) classifying between primary stress/secondary stress/no stress (Li et al., 2013, 2018). Ramanathi et al. (2019) have followed an alternative unsupervised way of classifying lexical stress, which is based on computing the likelihood of an acoustic signal for a number of possible lexical stress representations of a word.

Accuracy is the most commonly used performance metric, and it indicates the ratio of correctly classified stress patterns on a syllable (Li et al., 2013) or word level (Chen and Wang, 2010). On the contrary, Ferrer et al. (2015), analyzed the precision and recall metrics to detect lexical stress errors and not just classify them.

Most existing approaches for the classification and detection of lexical stress errors are based on carefully designed features. They start with aligning a speech signal with phonetic transcription, performed via forced-alignment (Shahin et al., 2016; Chen and Wang, 2010). Alternatively, ASR can provide both phonetic transcription and its alignment with a speech signal (Li et al., 2013). Then, prosodic features such as duration, energy and pitch (Chen and Wang, 2010) and cepstral features such as Mel Frequency Cepstral Coefficients (MFCC) and Mel-Spectrogram (Ferrer et al., 2015; Shahin et al., 2016) are extracted. These features can be extracted on the syllable (Shahin et al., 2016) or syllable nucleus (Ferrer et al., 2015; Chen and Wang, 2010) level. Shahin et al. (2016) computes features of neighboring vowels,

and Li et al. (2013) includes the features for two preceding and two following syllables in the model. The features are often preprocessed and normalized to avoid potential confounding variables (Ferrer et al., 2015), and to achieve better model generalization by normalizing the duration and pitch on a word level (Ferrer et al., 2015; Chen and He, 2007). Li et al. (2018) adds canonical lexical stress to input features, which improves the accuracy of the model.

In our recent work, we use attention mechanisms to automatically derive areas of the audio signal that are important for the detection of lexical stress errors (Korzekwa et al., 2021a). In this work, we use the T2S method to generate synthetic lexical stress errors to improve the accuracy of detecting lexical stress errors.

2.3. Synthetic speech generation for pronunciation error detection

Existing synthetic speech generation techniques for detecting pronunciation errors can be divided into two categories: data augmentation and data generation.

Data augmentation techniques are designed to generate new training examples for existing mispronunciation labels. Badenhorst and De Wet (2017) simulate new speakers by adjusting the speed of raw audio signals. Eklund (2019) generates additional training data by adding background noise and convolving the audio signal with the impulse responses of the microphone of a mobile device and a room.

Data generation techniques are designed to generate new training data with new labels of both correctly pronounced and mispronounced speech. Most existing works are based on the P2P technique to generate mispronounced speech by perturbing the phoneme sequence of the corresponding audio using a variety of strategies (Lee et al., 2016; Komatsu and Sasayama, 2019; Fu et al., 2021; Yan et al., 2021; Korzekwa et al., 2021c). In addition to P2P techniques, in our recent work, we use T2S to generate synthetic lexical stress errors (Korzekwa et al., 2021b). Qian et al. (2010) introduce a generative model to create hypotheses of mispronounced speech and use it as a reference speech signal to detect pronunciation errors. Recently, we proposed a similar technique to create a pronunciation model of native speech to account for many ways of correctly pronouncing a sentence by a native speaker (Korzekwa et al., 2021c).

Synthetic speech generation techniques have recently gained attention in other related fields. Fazel et al. (2021) use synthetic speech generated with T2S to improve accuracy in ASR. Huang et al. (2016) use a machine translation technique to generate text to train an ASR language model in a low-resource language. At the same time, Shah et al. (2021) and Huybrechts et al. (2021) employ S2S voice conversion to improve the quality of speech synthesis in the data reduction scenario.

All the presented works on the detection of pronunciation errors treat synthetic speech generation as a secondary contribution. In this article, we present a unified perspective of synthetic speech generation methods for detecting pronunciation errors. This article extends our previous work (Korzekwa et al., 2021c,a,b) and introduces a new S2S method to detect pronunciation errors. To the best of our knowledge, there are no papers devoted to generating pronunciation errors with the S2S technique and using it in the detection of pronunciation errors.

3. Methods of generating pronunciation errors

To detect pronunciation errors, first, the spoken language must be separated from other factors in the signal and then incorrectly pronounced speech sounds have to be identified. Separating speech into multiple factors is difficult, as speech is a complex signal. It consists of prosody (F0, duration, energy), timbre of the voice, and the representation of the spoken language. Spoken language is defined by the sounds (phones) perceived by people. Phones are the realizations of phonemes—a human abstract representation of how to pronounce a

word/sentence. Speech may also present variability due to the recording channel and environmental effects such as noise and reverberation. Detecting pronunciation errors is very challenging, also because of the limited amount of recordings with mispronounced speech. To address these challenges, we reformulate the problem of pronunciation error detection as the task of synthetic speech generation.

Let s be the speech signal, r be the sequence of phonemes that the user is trying to pronounce (canonical pronunciation), and e be the sequence of probabilities of mispronunciation at the phoneme or word level. The original task of detecting pronunciation errors is defined by:

$$e \sim p(e|s, r) \quad (1)$$

where the formulation of the problem as the task of synthetic speech generation is defined as follows:

$$s \sim p(s|e, r) \quad (2)$$

The probability of pronunciation errors for all the words in a sentence can then be calculated using the Bayes rule (Bishop, 2006):

$$p(e|s, r) = \frac{p(e|r)p(s|e, r)}{p(s|r)} \quad (3)$$

From Eq. (3), one can see that there is no need to directly learn the probability of pronunciation errors $p(e|s, r)$, since the complexity of the problem has now been transferred to learning the speech generation process $p(s|e, r)$. Such a formulation of the problem opens the way to the inclusion of additional prior knowledge into the model:

- (1) Replacing the phoneme in a word while preserving the original speech signal results in a pronunciation error (P2P method).
- (2) Changing the speech signal while retaining the original pronunciation results in a pronunciation error (T2S method).
- (3) There are many variations of mispronounced speech that differ in terms of the voice timbre and the prosodic aspects of speech (S2S method).

To solve Eq. (3), we use Markov Chain Monte Carlo Sampling (MCMC) (Koller and Friedman, 2009). In this way, the prior knowledge can be incorporated by generating N training examples $\{e_i, s_i, r_i\}$ for $i = 1..N$ with the use of P2P (prior knowledge 1), T2S (prior knowledge 2), and S2S (prior knowledge 3) methods. Accounting for the prior knowledge, intuitively corresponds to an increase in the amount of training data, which contributes to outperforming state-of-the-art models for detecting pronunciation errors, as presented in Section 5. Eq. (3) can then be optimized with standard gradient-based optimization techniques. In the following subsections, we present the P2P conversion, T2S, and S2S methods of generating correctly and incorrectly pronounced speech in details.

3.1. P2P method

To generate synthetic mispronounced speech, it is enough to start with correctly pronounced speech and modify the corresponding sequence of phonemes. This simple idea does not even require generating the speech signal itself. It can be observed that the probability of mispronunciations depends on the discrepancy between the speech signal and the corresponding canonical pronunciation. This leads to the P2P conversion model shown in Fig. 1a.

Let $\{e_{\text{noerr}}, s, r\}$ be a single training example containing: the sequence of 0 s denoting correctly pronounced phonemes, the speech signal, and the sequence of phonemes representing the canonical pronunciation. Let r' be the sequence of phonemes with injected mispronunciations such as phoneme replacements, insertions, and deletions:

$$r' \sim p(r'|r) \quad (4)$$




Fig. 1. Probabilistic graphical models for three methods to generate pronunciation errors: P2P, T2S and S2S. Empty circles represent hidden (latent) variables, while filled (blue) circles represent observed variables. s —the speech signal, r —the sequence of phonemes that the user is trying to pronounce (canonical pronunciation), the superscript ' represents a variable with generated mispronunciations.

then the probability of mispronunciation for the j th phoneme is defined by:

$$e'_j = \begin{cases} 1 & \text{if } r'_j \neq r_j \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

The probabilities of mispronunciation can be projected from the level of phonemes to the level of words. A word is treated as mispronounced if at least one pair of phonemes in the word $\{r'_j, r_j\}$ does not match. At the end of this process, a new training example is created with artificially introduced pronunciation errors: $\{e_{\text{err}}, s, r'\}$. Note that the speech signal s in the new training example is unchanged from the original training example, and only phoneme transcription is manipulated.

Implementation

To generate synthetic pronunciation errors, we use a simple approach of perturbing phonetic transcription for the corresponding speech audio. First, we sample these utterances with replacement from the input corpora of human speech. Then, for each utterance, we replace the phonemes with random phonemes with a given probability.

3.2. T2S method

The T2S method expands on P2P by making it possible to create speech signals that match the synthetic mispronunciations. The T2S method for generating mispronounced speech is a generalization of the P2P method, as can be seen by the comparison of the two methods shown in Fig. 1a and b.

One problem with the P2P method is that it cannot generate a speech signal for the newly created sequence of phonemes r' . As a result, pronunciation errors will dominate in the training data containing new sequences of phonemes r' . Therefore, it will be possible to detect pronunciation errors only from the canonical representation r' , ignoring information contained in the speech signal. To mitigate this issue, there should be two training examples for the phonemes r' , one representing mispronounced speech: $\{e_{\text{err}}, s, r'\}$, and the second one for correct pronunciation: $\{e_{\text{noerr}}, s', r'\}$, where:

$$s' \sim p(s' | e_{\text{noerr}}, r') \quad (6)$$

Because we now have the speech signal s' , another training example can be created as: $\{e_{\text{err}}, s', r'\}$. In summary, T2S method extends a single training example of correctly pronounced speech to four combinations of correctly and incorrect pronunciations:

- $\{e_{\text{noerr}}, s, r\}$ – correctly pronounced input speech
- $\{e_{\text{err}}, s, r'\}$ – mispronounced speech generated by the P2P method
- $\{e_{\text{noerr}}, s', r'\}$ – correctly pronounced speech generated by the T2S method
- $\{e_{\text{err}}, s', r'\}$ – mispronounced speech generated by the T2S method

Implementation

The synthetic speech is generated with the Neural TTS described by Latorre et al. (2019). The Neural TTS consists of two modules.

The context-generation module is an attention-based encoder–decoder neural network that generates a mel-spectrogram from a sequence of phonemes. The Neural Vocoder then converts it into a speech signal. The Neural Vocoder is a neural network of architecture similar to Parallel Wavenet (Oord et al., 2018). The Neural TTS is trained using the speech of a single native speaker. To generate words with different lexical stress patterns, we modify the lexical stress markers associated with the vowels in the phonetic transcription of the word. For example, with the input of /r iy1 m ay0 n d/ we can place lexical stress on the first syllable of the word 'remind'.

3.3. S2S method

The S2S method is designed to simulate the diverse nature of speech, as there are many ways to correctly pronounce a sentence. The prosodic aspects of speech, such as pitch, duration, and energy, can vary. Similarly, phonemes can be pronounced differently. To mimic human speech, speech generation techniques should allow a similar level of variability. The T2S method outlined in the previous section always produces the same output for the same phoneme input sequence. The S2S method is designed to overcome this limitation.

S2S converts the input speech signal s in a way to change the pronounced phonemes (phoneme replacements, insertions, and deletions) from the input phonemes r to target phonemes r' while preserving other aspects of speech, including voice timbre and prosody (Eq. (7) and Fig. 1c). In this way, the natural variability of human speech is preserved, resulting in generating many variations of incorrectly pronounced speech. The prosody will differ in various versions of the sentence of the same speaker, while the same sentence spoken by many speakers will differ in the voice timbre.

$$s' \sim p(s' | e_{\text{noerr}}, r', s) \quad (7)$$

Similarly to the T2S method, the S2S method outputs four types of speech pronounced correctly and incorrectly: $\{e_{\text{noerr}}, s, r\}$, $\{e_{\text{err}}, s, r'\}$, $\{e_{\text{noerr}}, s', r'\}$, and $\{e_{\text{err}}, s', r'\}$.

Implementation

Synthetic speech is generated by introducing mispronunciations into the input speech, while preserving the duration of the phonemes and timbre of the voice. The architecture of the S2S model is shown in Fig. 2. The mel-spectrogram of the input speech signal s is forced-aligned with the corresponding canonical phonemes r to get the duration of the phonemes. The speaker id has to be provided together with the input speech to enable the source speaker's voice to be maintained. Mispronunciations are introduced into the canonical phonemes r according to the P2P method described in Section 3.1. Mispronounced phonemes r' along with phonemes duration and speaker id are processed by the encoder–decoder, which generates the mel-spectrogram s' . The encoder–decoder transforms the phoneme-level representation into frame-level features and then generates all mel-spectrogram frames in parallel. The mel-spectrogram is converted to an audio signal with Universal Vocoder (Jiao et al., 2021). Without the Universal Vocoder,

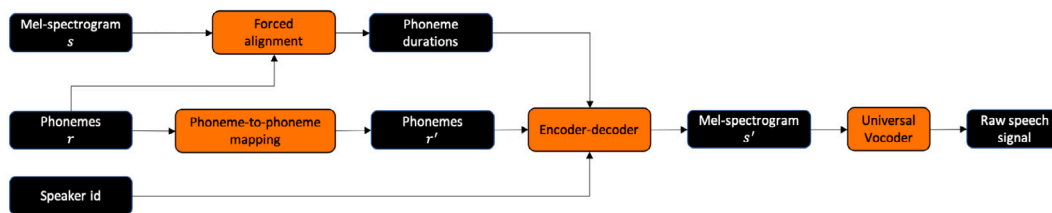


Fig. 2. Architecture of the S2S model to generate mispronounced synthetic speech while maintaining prosody and voice timbre of the input speech. The black rectangles represent the data (tensors) and the orange boxes represent processing blocks. This color notation is used in all machine learning model diagrams throughout the article. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

it would not be possible to generate the raw audio signal for hundreds of speakers included in the LibriTTS corpus. Details of the S2S method are shown in the works of Shah et al. (2021) and Jiao et al. (2021). The main difference between these two models and our S2S model is the use of the P2P mapping to introduce pronunciation errors.

3.4. Summary of mispronounced speech generation

Generation of synthetic mispronounced speech and detection of pronunciation errors were presented from the probabilistic perspective of the Bayes-rule. With this formulation, we can better understand the relationship between P2P, T2S and S2S methods, and see that the S2S method generalizes two simpler methods. Following this reasoning, we can argue that using the Bayes rule gives us a nice mathematical framework to potentially further generalize the S2S method, e.g. by adding a language variable to the model to support multilingual pronunciation error detection. There is another advantage of modeling pronunciation error detection from the probabilistic perspective—it paves the way for joint training of mispronounced speech generation and pronunciation error detection models. In the present work, we are training separate machine learning models for both tasks, but it should be possible to train both models jointly using the framework of Variational Inference (Jordan et al., 1999) instead of MCMC to infer the probability of mispronunciation in Eq. (3).

4. Speech corpora

4.1. Corpora of continuous speech

Speech corpora of recorded sentences is a combination of L1 and L2 English speech. L1 speech is obtained from the TIMIT (Garofolo et al., 1993) and the LibriTTS (Zen et al., 2019) corpora. L2 speech comes from the Isle (Atwell et al., 2003) corpus (German and Italian speakers) and the GUT Isle (Weber et al., 2020) corpus (Polish speakers). In total, we used 125.28 h of L1 and L2 English speech from 983 speakers segmented into 102 812 sentences. A summary of the speech corpora is presented in Table 1, whereas the details are presented in our recent work (Korzekwa et al., 2021b).

The speech data are used in all the pronunciation error detection experiments presented in Section 5. From the collected speech, we held out 28 L2 speakers and used them only to assess the performance of the systems in the mispronunciation detection task. It includes 11 Italian and 11 German speakers from the Isle corpus (Atwell et al., 2003), and 6 Polish speakers from the GUT Isle corpus (Weber et al., 2020). The human speech training data is extended with synthetic pronunciation errors generated by the methods presented in Section 3.

4.2. Corpora of isolated words

The speech corpora consist of human and synthetic speech. The data were divided into training and testing sets, with separate speakers assigned to each set. Human speech includes native (L1) and non-native (L2) English speech. L1 speech corpora are made of TIMIT (Garofolo et al., 1993) and Arctic (Kominek and Black, 2004). L2 corpora contain

Table 1

Summary of human speech corpora used in the pronunciation error detection experiments.

Native language	Hours	Speakers
English	90.47	640
Unknown ^a	19.91	285
German and Italian	13.41	46
Polish	1.49	12

^aAudiobooks read by volunteers from all over the world (Zen et al., 2019).

Table 2

Details of the training and test sets for the lexical stress error detection model.

Data set	Speakers (L2)	Words (unique)	Stress errors
Train set (human)	473 (10)	8223 (1528)	425
Train set (TTS)	1 (0)	3937 (1983)	2005
Test set (human)	176 (21)	2108 (378)	189

speech from L2-Arctic (Sudhakara et al., 2019a), Porzuczek (Porzuczek and Rojczyk, 2017), and our own recordings of 25 speakers (23 Polish, 1 Ukrainian and 1 Lithuanian). The synthetic data were generated using the T2S method and are only included in the training set. The data are summarized in Table 2. For a more detailed description of speech corpora, see Section 4 of our recent work (Korzekwa et al., 2021a). The speech corpora of isolated words are used in the lexical stress error detection experiment presented in Section 5.3.

5. Experiments

5.1. Generation of mispronounced speech

5.1.1. Experimental setup

The effect of using synthetic pronunciation errors based on the P2P, T2S and S2S methods is evaluated in the task of detecting pronunciation errors in spoken sentences at the word level. First, we analyze the P2P method by comparing it with the state-of-the-art techniques and measure the effect of adding synthetic pronunciation errors to the training data. We then compare P2P with T2S and S2S to assess the benefits of using more complex methods of generating pronunciation errors. The accuracy of detecting pronunciation errors is reported in standard Area Under the Curve (AUC), precision and recall metrics.

5.1.2. Overview of our WEAKLY-S model

We use the pronunciation error detection model (WEAKLY-S) recently proposed by us Korzekwa et al. (2021b). To train the model, the human speech training set is extended with 292,242 utterances of L1 speech with synthetically generated pronunciation errors. To generate pronunciation errors, the P2P, T2S, and S2S methods described in Section 3 are used.

The WEAKLY-S model produces probabilities of mispronunciation for all words, conditioned by the spoken sentence and canonical phonemes. Mispronunciation errors include phoneme replacement, addition, deletion, or an unknown speech sound. During training, the




Fig. 3. Architecture of the WEAKLY-S model for word-level pronunciation error detection trained in the multi-task setup. Task 1—to detect pronunciation errors e . Task 2—to recognize phonemes r_o .

model is weakly supervised, in the sense that only mispronounced words in L2 speech are marked by listeners and the data do not have to be phonetically transcribed. Due to the limited availability of L2 speech and the fact that it is not phonetically transcribed, the model is more likely to overfit. To solve this problem, the model is trained in a multi-task setup. In addition to the primary task of detecting mispronunciation error at the word level, the second task uses a phoneme recognizer which is trained on automatically transcribed L1 speech. Both tasks share components of the model, which makes the primary task less likely to overfit.

The architecture of the pronunciation error detection model is shown in Fig. 3. The model consists of two sub-networks. The Mispronunciations Detection Network (MDN) detects word-level pronunciation errors e from the audio signal s and canonical phonemes r , while the Phoneme Recognition Network (PRN) recognizes phonemes r_o pronounced by a speaker from the audio signal s . The detailed model architecture is presented in Section 2 of our recent work (Korzekwa et al., 2021b).

5.1.3. Results—P2P method

We conducted an ablation study to measure the effect of removing synthetic pronunciation errors from the training data. We trained four variants of the WEAKLY-S model to measure the effect of using synthetic data against other elements of the model. WEAKLY-S is a complete model that also includes synthetic data during training. In the NO-SYNTH-ERR model, we exclude synthetic samples of mispronounced L1 speech, significantly reducing the number of mispronounced words seen during training from 1,129,839 to just 5273 L2 words. The NO-L2-ADAPT variant does not fine-tune the model on L2 speech, although it is still exposed to L2 speech while being trained on a combined corpus of L1 and L2 speech. The NO-L1L2-TRAIN model is not trained on L1/L2 speech, and fine-tuning on L2 speech starts from scratch. This means that this model will not use a large amount of phonetically transcribed L1 speech data and ultimately no secondary phoneme recognition task will be used.

L2 fine-tuning (NO-L2-ADAPT) is the most important factor influencing the performance of the model (Fig. 4 and Table 3), with an AUC of 0.517 compared to 0.686 for the full model. Training the model on both L2 and L1 human speech together is not enough. This is because L2 speech accounts for less than 1% of the training data and the model naturally leans towards L1 speech. The second most important feature is training the model on a combined set of L1 and L2 speech (NO-L1L2-TRAIN), with an AUC of 0.565. L1 speech accounts for over 99% of training data. These data are also phonetically transcribed, and therefore can be used for the phoneme recognition task. The phoneme recognition task acts as a ‘backbone’ and reduces the effect of overfitting in the main task of detecting errors in the pronunciation of words. Finally, excluding synthetically generated pronunciation errors (NO-SYNTH-ERR) reduces an AUC from 0.686 to 0.615. Although, the synthetic data provides the least improvement to the model, it still increases the accuracy of the model by 11.5% in AUC, contributing to setting up a new state-of-the-art.

We compare the WEAKLY-S model with two state-of-the-art baselines. The Phoneme Recognizer (PR) model by Leung et al. (2019) is our




Fig. 4. Precision–recall curve for the ablation study on the GUT Isle corpus, illustrating the effect of using synthetic pronunciation errors generated by the P2P method.

first baseline. The PR is based on the CTC loss (Graves, 2012) and outperforms multiple alternative approaches of pronunciation assessment. The original CTC-based model uses a hard likelihood threshold applied to the recognized phonemes. To compare it with two other models, following our recent work (Korzekwa et al., 2021c), we have replaced the hard likelihood threshold with a soft threshold. The second baseline is PR extended by the pronunciation model (PR-PM model Korzekwa et al., 2021c). The pronunciation model takes into account the phonetic variability of the speech spoken by native speakers, which results in greater precision in detecting pronunciation errors. The results are shown in Table 4. It turns out that the WEAKLY-S model outperforms the second-best model in terms of an AUC by 30% from 0.528 to 0.686 and precision by 23% from 0.612 to 0.752 on the GUT Isle Corpus of Polish speakers. We are seeing similar improvements on the Isle Corpus of German and Italian speakers. The use of synthetic data is an important contribution to the performance of the WEAKLY-S model.

5.1.4. Results—T2S and S2S methods

The main limitation of the P2P method is that it does not generate a new speech signal. The method introduces mispronunciations by operating only on the sequence of phonemes for the corresponding speech. In this experiment, we demonstrate the T2S and S2S methods that can directly generate a speech signal to overcome this limitation. The S2S method introduces mispronunciations into the input native speech while preserving the prosody (phoneme durations) and timbre of the voice. Preserving speech attributes other than pronunciation increases speech variability during training and makes the pronunciation error detection model more reliable during testing. The T2S method can be considered as a simplified variant of the S2S method, in which there is only text as input.

The T2S and S2S methods are compared with the P2P method. Three WEAKLY-S models are trained, differing in the technique of generating mispronounced speech contained in the training data. The S2S method outperforms the P2P method by increasing an AUC score by 9% from 0.686 to 0.749 in the Gut Isle corpus of Polish speakers (Table 5). Additionally, an AUC increases from 0.815 to 0.834 for major pronunciation errors (Table 6), according to a similar experiment presented in Section 3.4 of Korzekwa et al. (2021b). Interestingly, the

Table 3

Ablation study for the GUT Isle corpus to show the effect of using synthetic data and other elements of the WEAKLY-S model.

Model	Description	AUC	Precision [%]	Recall [%]
NO-L2-ADAPT	No fine-tuning on L2 speech	0.517	57.89	40.11
NO-L1L2-TRAIN	No pretraining on L1&L2 speech	0.565	59.73	40.20
NO-SYNTH-ERR	No synthetically generated pronunciation errors in the training data	0.615	67.22	40.38
WEAKLY-S	Complete model	0.686	75.25	40.38

Table 4

Accuracy metrics of detecting word-level pronunciation errors. WEAKLY-S vs. baseline models.

Model	AUC	Precision [%;95%CI]	Recall [%;95%CI]
Isle corpus (German and Italian)			
PR	0.555	49.39 (47.59–51.19)	40.20 (38.62–41.81)
PR-PM	0.480	54.20 (52.32–56.08)	40.20 (38.62–41.81)
WEAKLY-S	0.678	71.94 (69.96, 73.87)	40.14 (38.56, 41.75)
GUT Isle corpus (Polish)			
PR	0.528	54.91 (50.53–59.24)	40.29 (36.66–44.02)
PR-PM	0.505	61.21 (56.63–65.65)	40.15 (36.51–43.87)
WEAKLY-S	0.686	75.25 (71.67–78.59)	40.38 (37.52–43.29)

Table 5

Comparison of the P2P, T2S and S2S methods in the task of pronunciation error detection assessed on the GUT Isle corpus.

Model	AUC	Precision [%]	Recall [%]
P2P	0.686	75.25 (71.67–78.59)	40.38 (37.52–43.29)
T2S	0.695	76.15 (72.59–79.36)	40.25 (37.44–43.22)
S2S	0.749	80.45 (76.94–83.47)	40.12 (37.12–43.02)

Table 6

Comparison of the P2P, T2S and S2S methods in the task of pronunciation error detection assessed on the GUT Isle corpus only for major pronunciation errors.

Model	AUC	Precision [%]	Recall [%]
P2P	0.815	91.67 (88.55–94.45)	40.31 (37.43–43.23)
T2S	0.819	92.11 (89.09–94.83)	40.21 (36.81–43.31)
S2S	0.834	93.54 (90.53–96.23)	40.15 (37.26–43.11)

T2S method is only slightly better than the P2P method, which suggests that the variability of the generated mispronounced speech provided by the S2S method is really important. The presented experiments show the potential of the S2S method in improving the accuracy of detecting pronunciation errors. The S2S method is able to control voice timbre, phoneme duration, and pronunciation, opening the door to transplanting all three properties from non-native speech and potentially further improving the accuracy of the model.

One downside of the S2S method is its complexity. Compared to the straightforward P2P method, the 9% improvement in an AUC is associated with high costs. The method involves training a complex multi-speaker S2S model to convert between input and output mel-spectrograms and requires training a Universal Vocoder model to convert a mel-spectrogram into a raw speech signal.

To better understand what prevents the model from achieving higher accuracy, we measure the performance of the model on synthetic pronunciation errors. We divide all synthetic pronunciation errors into four categories to reflect the severity of pronunciation errors. The ‘low’ category includes mispronounced words with only one mismatched phoneme between the canonical and pronounced phonemes of the word. The ‘medium’ category includes two mispronounced phonemes. The ‘high’ category gets three, and the ‘very high’ category includes four mispronounced errors. The AUC across different severity levels varies from 0.928 (low severity) to 1.00 (very high severity) as shown in Table 7. These AUC values are significantly higher than the results for non-native human speech, suggesting that making synthetic speech errors more similar to non-native speech may improve the accuracy of detecting pronunciation errors.

Table 7

Accuracy (AUC) in detecting pronunciation errors assessed in synthetic speech at different severity levels of mispronunciation for the best S2S method.

Severity	AUC
Low (phoneme distance = 1)	0.928
Medium (phoneme distance = 2)	0.974
High (phoneme distance = 3)	0.993
Very High (phoneme distance = 4)	1.00

5.2. Model of native speech pronunciation

5.2.1. Experimental setup

The P2P, T2S, and S2S are generative models that provide the probability of generating a particular output sequence. This probability can be used directly to detect pronunciation errors without generating the mispronounced speech and adding it to the training data. In this experiment, we show how to apply this approach in practice.

One of the challenges in detecting pronunciation errors is that a native speaker can pronounce a sentence correctly in many ways. The classic approach for detecting pronunciation errors is based on identifying the difference between pronounced and canonical phonemes. All pronunciations that do not correspond precisely to the canonical pronunciation will result in false pronunciation errors. One way to solve this problem is to use the P2P technique to create a native speech Pronunciation Model (PM) that determines the probability that a sentence is pronounced by a native speaker. A low likelihood value indicates a high probability of mispronunciation.

To evaluate the performance of the PM model, the pronunciation error detection model has been designed such that the PM model can be turned on and off. To disable the PM, we are modifying it so that it only takes into account one way of correctly pronouncing a sentence. In an ablation study, we measure whether the PM model improves the accuracy in detecting pronunciation errors at the word level. Note that in this experiment, synthetically generated pronunciation errors are not used explicitly. Instead, the native speech pronunciation model is used to implicitly represent the generative speech process.

5.2.2. Overview of the pronunciation error detection model

The design of the pronunciation error detection model consists of three subsystems: a Phoneme Recognizer (PR), a Pronunciation Model (PM), and a Pronunciation Error Detector (PED), shown in Fig. 5. First, the PR model estimates a belief over the phonemes produced by the student, intuitively representing the uncertainty in the student’s pronunciation. The PM model transforms this belief into a probability that a native speaker would pronounce the sentence this way, given the phonetic variability. Finally, the PED model decides which words were mispronounced in the sentence by processing three pieces of information: (a) what the student pronounced, (b) how likely it is that the native speaker would pronounce it that way, and (c) what the student was supposed to pronounce. Details of the entire model of pronunciation error detection are presented in Section 3 of our recent work (Korzekwa et al., 2021c). We will now only show the details of the PM model that are relevant to this experiment.




Fig. 5. Architecture of the system for detecting mispronounced words in a spoken sentence based on the native speech pronunciation model.

5.2.3. Overview of the native speech pronunciation model

PM is an encoder–decoder neural network, following Sutskever et al. (2014). Instead of building a text-to-text translation system between two languages, we use it for the P2P conversion. The sequence of phonemes \mathbf{r} that the native speaker was supposed to pronounce is converted to the sequence of phonemes \mathbf{r}' they had pronounced, denoted as $\mathbf{r}' \sim p(\mathbf{r}'|\mathbf{r})$. Once trained, PM acts as a probability mass function, computing the probability sequence π of the recognized phonemes \mathbf{r}_o pronounced by the student conditioned by the expected (canonical) phonemes \mathbf{r} . PM is denoted as in Eq. (8).

$$\pi = \sum_{\mathbf{r}_o} p(\mathbf{r}_o|\mathbf{o})p(\mathbf{r}' = \mathbf{r}_o|\mathbf{r}) \quad (8)$$

The PM model is trained on P2P speech data generated automatically by passing the speech of the native speakers through the PR. By using PR to annotate the data, we can make the PM model more robust against possible phoneme recognition inaccuracies in PR at the time of testing.

5.2.4. Results

The complete model with PM enabled is called PR-PM that stands for a Phoneme Recognizer + Pronunciation Model. The model with PM turned off is called PR-LIK that stands for Phoneme Recognizer outputting the likelihoods of recognized phonemes. PR-LIK is an extension of the PR-NOLIK model—the mispronunciation detection model proposed by Leung et al. (2019) that only returns the most likely recognized phonemes and does not use phoneme likelihoods to detect pronunciation errors. PR-NOLIK detects mispronounced words based on the difference between the canonical and recognized phonemes. Therefore, this system does not offer any flexibility in optimizing the model for higher precision by fine-tuning the threshold applied to the phoneme recognition probabilities.

Turning off PM reduces the precision between 11% and 18%, depending on the decrease in recall between 20% to 40%, as shown in Fig. 6. One example where the PM helps is the word ‘enough’ that can be pronounced in two similar ways: /ih n ah f/ or /ax n ah f/ (short ‘i’ or ‘schwa’ phoneme at the beginning.) The PM can take into account the phonetic variability and recognize both versions as correctly pronounced. Another example is coarticulation (Hieke, 1984). Native speakers tend to merge phonemes of adjacent words. For example, in the text ‘her arrange’ /hh er - er ey n jh/, two adjacent phonemes /er/ can be pronounced as one phoneme: /hh er ey n jh/. The PM model can correctly recognize multiple variations of such pronunciations.

Complementary to the precision–recall curve shown in Fig. 6, we present in Table 8 one configuration of the precision and recall scores for the PR-LIK and PR-PM systems. This configuration is chosen in a way to: (a) make the recall for both systems close to the same value, and (b) to illustrate that the PR-PM model has much greater potential to increase precision than the PR-LIK system. A similar conclusion can be drawn by checking various different precision and recall configurations in the precision and recall plots for both Isle and GUT Isle corpora.

5.3. Lexical stress error detection

5.3.1. Experimental setup

The full CAPT learning experience includes both the detection of pronunciation and lexical stress errors. To investigate the potential of

Table 8

Precision and recall of detecting word-level mispronunciations. CI—Confidence Interval. PR-PM—full model with the PM enabled. PR-LIK—the PR-PM model with the PM disabled.

Model	Precision [%;95%CI]	Recall [%;95%CI]
Isle corpus (German and Italian)		
PR-LIK	49.39 (47.59–51.19)	40.20 (38.62–41.81)
PR-PM	54.20 (52.32–56.08)	40.20 (38.62–41.81)
GUT Isle corpus (Polish)		
PR-LIK	54.91 (50.53–59.24)	40.29 (36.66–44.02)
PR-PM	61.21 (56.63–65.65)	40.15 (36.51–43.87)

speech generation in the lexical stress error detection task, we evaluate the T2S method, which is a simpler version of the S2S method evaluated in Section 5.1.4.

The lexical stress error detection model is trained to measure the benefits of employing synthetic mispronounced speech. The first model, denoted as Att_TTS is based on an attention mechanism and is trained on both human and synthetic speech with pronunciation errors. In this model, 1980 the most popular English words (Michel et al., 2011) were synthesized with correct and incorrect stress patterns using the method outlined in Section 3.2, and added to the speech corpora of isolated words presented in Section 4.2. The Att_NoTTS model is trained only on human speech. Each of the two models presented has its simpler version without the attention mechanism, marked as NoAtt_TTS and NoAtt_NoTTS. Both models will help to understand whether the benefits of using synthetic pronunciation errors depend on the model capacity.

The accuracy of detecting lexical stress errors is measured in terms of an AUC metric. To be comparable to the study by Ferrer et al. (2015), we use precision as an additional metric, while setting recall to 50%.

5.3.2. Overview of the lexical stress detection model

As shown in Fig. 7, the lexical stress error detection model consists of three subsystems: Feature Extractor, Attention-based Classification Model, and Lexical Stress Error Detector. The Feature Extractor extracts prosodic features and phonemes from the speech signal \mathbf{s} and the forced-aligned canonical phonemes \mathbf{r} . Prosodic features include: F0, intensity [dB SPL] and duration of phonemes. The F0 and intensity features are computed at the frame level. The Attention-based Classification Model uses the attention mechanism (Vaswani et al., 2017) to map frame-level and phoneme-level features to a syllable-level representation. It then produces lexical stress error probabilities at the syllable level. The Lexical Stress Error Detector reports a lexical stress error if the expected (canonical) and estimated lexical stress for a given syllable do not match and the corresponding probability is higher than the specified threshold. The detailed architecture of the model is presented in Section 3 of our recent work (Korzekwa et al., 2021a).

The NoAtt_TTS and NoAtt_NoTTS models do not have the attention mechanism. Instead, as a representation at the syllable level, they use the average acoustic feature values for the corresponding syllable nucleus. The hypothesis is that synthetic data will not be beneficial to a simpler model due to its limited capacity.




Fig. 6. Precision–recall curves for the evaluated systems to measure the effect of using the PM model in detecting pronunciation errors. PR-PM— full model with the PM enabled. PR-LIK—the PR-PM model with the PM disabled. PR-NOLIK—non-probabilistic variant of the PR-LIK model proposed by Leung et al. (2019).




Fig. 7. Attention-based model for the detection of lexical stress errors.




Fig. 8. Precision–recall curves for lexical stress error detection models.

5.3.3. Results

Enriching the training set with the incorrectly stressed words increases an AUC score from 0.54 to 0.62 (Att_NoTTS vs. Att_TTS in Fig. 8 and Table 9). Data augmentation helps because it increases the number of words with incorrect stress patterns in the training set. This prevents the model from using the strong correlation between phonemes and lexical stress in the correctly stressed words. Using data augmentation in the simpler model without the attention mechanism slightly reduced an AUC score from 0.45 to 0.44 (NoAtt_NoTTS vs. NoAtt_TTS). The NoAtt_TTS model has limited capacity due to not using the attention mechanism to model prosodic features, and thus is unable to benefit from synthetic speech.

We compare our results with the work of Ferrer et al. (2015). There were 46.4% (191 out of 411) of incorrectly stressed words in their corpus, well over 9.4% (189 out of 2109) words in our experiment. The fewer lexical stress errors that users make, the more difficult it is to detect them. Under these conditions, we can state that our lexical stress detection model based on T2S generated synthetic speech achieves higher scores in precision and recall compared to the work of Ferrer et al. (2015).

6. Conclusions

We propose a new paradigm for detecting pronunciation errors in non-native speech. Rather than focusing on detecting pronunciation errors directly, we reformulate the detection problem as a speech generation task. This approach is based on the assumption that it is easier to generate speech with specific characteristics than to detect those characteristics in speech with limited availability. In this way, we address one of the main problems of the existing CAPT methods, which is the low availability of mispronounced speech for reliable training of pronunciation error detection models.

We present a unified look at three different speech generation techniques for detecting pronunciation errors based on P2P, T2S and S2S conversion. The P2P, T2S, and S2S methods improve the accuracy of detecting pronunciation and lexical stress errors. The methods outperform strong baseline models and establish a new state-of-the-art. The best S2S method outperforms the baseline method (Leung et al., 2019) by improving the accuracy of detecting pronunciation errors in AUC metric by 41% from 0.528 to 0.749. The S2S method has the ability to control many properties of speech, such as voice timbre, prosody (duration), and pronunciation. This opens the door to the generation of mispronounced speech that can mimic certain aspects of non-native speech, such as voice timbre. The S2S method can be seen as a generalization of the simpler methods, T2S and P2P, providing a general framework for building a first-class models of pronunciation assessment. For better reproducibility, in addition to using publicly available speech corpora, we recorded the GUT Isle corpus of non-native English speech (Weber et al., 2020). The corpus is available to other researchers in the field.

In the future, we plan to extend the S2S method in order to generate synthetic speech as close as possible to non-native speech: (a) we will extract the voice timbre from the speech of non-native speakers and transfer it to native speech, following the paper of Merritt et al. on text-free voice conversion (Merritt et al., 2022), and (b) we will mimic the distribution of pronunciation errors in non-native speech. We expect both changes to increase the accuracy of detecting pronunciation errors in non-native speech. In the long run, we hope to demonstrate that “synthetic speech is all you need” by training the model with synthetic

Table 9
AUC, precision and recall [%, 95% Confidence Interval] metrics for lexical stress error detection models.

Model	Model with attention	Synthetic mispronunciations	AUC	Precision [%]	Recall[%]
Att_TTS	Yes	Yes	0.62	94.8 (89.18–98.03)	49.2 (42.13–56.3)
Att_NoTTS	Yes	No	0.54	87.85 (80.67–93.02)	49.74 (42.66–56.82)
NoAtt_TTS	No	Yes	0.44	44.39 (37.85–51.09)	50.26 (43.18–57.34)
NoAtt_NoTTS	No	No	0.45	48.98 (42.04–55.95)	50.79 (43.70–57.86)
Ferrer et al. (2015)	na	na	na	95.00 (na–na)	48.3 (na–na)

speech only and achieving state-of-the-art results in the pronunciation error detection task. This may revolutionize computer-assisted English L2 learning and CAPT. Moreover, such a paradigm may be transferred to the whole domain of computer-assisted foreign language learning.

CRedit authorship contribution statement

Daniel Korzekwa: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Jaime Lorenzo-Trueba:** Supervision, Writing – review & editing. **Thomas Drugman:** Writing – review & editing. **Bożena Kostek:** Resources, Supervision.

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.

References

Ai, R., 2015. Automatic pronunciation error detection and feedback generation for call applications. In: *International Conference on Learning and Collaboration Technologies*. Springer, pp. 175–186.

Atwell, E., Howarth, P., Souter, D., 2003. The ISLE corpus: Italian and German spoken learner's english. *ICAME J: Int. Comput. Arch. Mod. Med. Engl. J.* 27, 5–18.

Badenhorst, J., De Wet, F., 2017. The limitations of data perturbation for ASR of learner data in under-resourced languages. In: *2017 Pattern Recognition Association of South Africa and Robotics and Mechatronics (PRASA-RobMech)*. IEEE, pp. 44–49.

Bahdanau, D., Chorowski, J., Serdyuk, D., Brakel, P., Bengio, Y., 2016. End-to-end attention-based large vocabulary speech recognition. In: *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 4945–4949.

Bergem, D.R.v., 1991. Acoustic and lexical vowel reduction. In: *Phonetics and Phonology of Speaking Styles*.

Bishop, C.M., 2006. Pattern recognition. *Mach. Learn.* 128 (9).

Chen, N., He, Q., 2007. Using nonlinear features in automatic english lexical stress detection. In: *2007 Intl. Conference on Computational Intelligence and Security Workshops (CISW 2007)*. IEEE, pp. 328–332.

Chen, J.-Y., Wang, L., 2010. Automatic lexical stress detection for Chinese learners' of english. In: *2010 7th Intl. Symposium on Chinese Spoken Language Processing*. IEEE, pp. 407–411.

Cheng, S., Liu, Z., Li, L., Tang, Z., Wang, D., Zheng, T.F., 2020. ASR-free pronunciation assessment. *arXiv preprint arXiv:2005.11902*.

Chorowski, J., Bahdanau, D., Cho, K., Bengio, Y., 2014. End-to-end continuous speech recognition using attention-based recurrent NN: First results. *arXiv preprint arXiv:1412.1602*.

Chorowski, J.K., Bahdanau, D., Serdyuk, D., Cho, K., Bengio, Y., 2015. Attention-based models for speech recognition. In: *Advances in Neural Information Processing Systems*. pp. 577–585.

Duan, R., Kawahara, T., Dantsuji, M., Nanjo, H., 2019. Cross-lingual transfer learning of non-native acoustic modeling for pronunciation error detection and diagnosis. *IEEE/ACM Trans. Audio Speech Lang. Process.* 28, 391–401.

EF-Education-First, 2020. EF english proficiency index. Available at <https://www.ef/epi/>.

Eklund, V.-V., 2019. *Data Augmentation Techniques for Robust Audio Analysis* (Master's thesis). Tampere University.

Fazel, A., Yang, W., Liu, Y., Barra-Chicote, R., Meng, Y., Maas, R., Droppo, J., 2021. SynthASR: Unlocking synthetic data for speech recognition. In: *Proc. Interspeech 2021*. pp. 896–900.

Ferrer, L., Bratt, H., Richey, C., Franco, H., Abrash, V., Precoda, K., 2015. Classification of lexical stress using spectral and prosodic features for computer-assisted language learning systems. *Speech Commun.* 69, 31–45.

Field, J., 2005. Intelligibility and the listener: The role of lexical stress. *TESOL Quart.* 39 (3), 399–423.

Franco, H., Neumeier, L., Kim, Y., Ronen, O., 1997. Automatic pronunciation scoring for language instruction. In: *1997 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Vol. 2. IEEE, pp. 1471–1474.

Fu, K., Lin, J., Ke, D., Xie, Y., Zhang, J., Lin, B., 2021. A full text-dependent end to end mispronunciation detection and diagnosis with easy data augmentation techniques. *arXiv preprint arXiv:2104.08428*.

Garofolo, J.S., Lamel, L.F., Fisher, W.M., Fiscus, J.G., Pallett, D.S., 1993. DARPA TIMIT acoustic-phonetic continuous speech corpus CD-ROM. NIST speech disc 1-1.1. STIN 93, 27403.

Golonka, E.M., Bowles, A.R., Frank, V.M., Richardson, D.L., Freynik, S., 2014. Technologies for foreign language learning: A review of technology types and their effectiveness. *Comput. Assist. Lang. Learn.* 27 (1), 70–105.

Graves, A., 2012. Connectionist temporal classification. In: *Supervised Sequence Labelling with Recurrent Neural Networks*. Springer, pp. 61–93.

Harrison, A.M., Lo, W.-K., Qian, X.-j., Meng, H., 2009. Implementation of an extended recognition network for mispronunciation detection and diagnosis in computer-assisted pronunciation training. In: *Intl. Workshop on Speech and Language Technology in Education*.

Hieke, A.E., 1984. Linking as a marker of fluent speech. *Lang. Speech* 27 (4), 343–354.

Huang, G., Gorin, A., Gauvain, J.-L., Lamel, L., 2016. Machine translation based data augmentation for cantonese keyword spotting. In: *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 6020–6024.

Huybrechts, G., Merritt, T., Comini, G., Perz, B., Shah, R., Lorenzo-Trueba, J., 2021. Low-resource expressive text-to-speech using data augmentation. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 6593–6597.

Jiao, Y., Gabryś, A., Tinchev, G., Putrycz, B., Korzekwa, D., Klimkov, V., 2021. Universal neural vocoding with parallel wavenet. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 6044–6048.

Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K., 1999. An introduction to variational methods for graphical models. *Mach. Learn.* 37 (2), 183–233.

Jung, Y.-J., Rhee, S.-C., et al., 2018. Acoustic analysis of english lexical stress produced by Korean, Japanese and Taiwanese-Chinese speakers. *Phonetics Speech Sci.* 10 (1), 15–22.

Koller, D., Friedman, N., 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press.

Komatsu, S., Sasayama, M., 2019. Speech error detection depending on linguistic units. In: *Proceedings of the 2019 3rd International Conference on Natural Language Processing and Information Retrieval*. pp. 75–79.

Kominek, J., Black, A.W., 2004. The CMU arctic speech databases. In: *Fifth ISCA Workshop on Speech Synthesis*.

Korzekwa, D., Barra-Chicote, R., Zaporowski, S., Beringer, G., Lorenzo-Trueba, J., Serafinowicz, A., Droppo, J., Drugman, T., Kostek, B., 2021a. Detection of lexical stress errors in non-native (L2) english with data augmentation and attention. In: *Proc. Interspeech 2021*. pp. 3915–3919.

Korzekwa, D., Lorenzo-Trueba, J., Drugman, T., Calamaro, S., Kostek, B., 2021b. Weakly-supervised word-level pronunciation error detection in non-native english speech. In: *Proc. Interspeech 2021*. pp. 4408–4412.

Korzekwa, D., Lorenzo-Trueba, J., Zaporowski, S., Calamaro, S., Drugman, T., Kostek, B., 2021c. Mispronunciation detection in non-native (L2) english with uncertainty modeling. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 7738–7742.

Latorre, J., Lachowicz, J., Lorenzo-Trueba, J., Merritt, T., Drugman, T., Ronanki, S., Klimkov, V., 2019. Effect of data reduction on sequence-to-sequence neural tts. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 7075–7079.

Lee, A., Glass, J.R., 2013. Pronunciation assessment via a comparison-based system. In: *SLaTE*.

Lee, A., et al., 2016. *Language-Independent Methods for Computer-Assisted Pronunciation Training* (Ph.D. Diss.). Massachusetts Institute of Technology.

Lepage, A., Busà, M.G., 2014. Intelligibility of English L2: The effects of incorrect word stress placement and incorrect vowel reduction in the speech of French and Italian learners of English. In: *Proc. of the Intl. Symposium on the Acquisition of Second Language Speech*. 2014. pp. 387–400.

Leung, W.-K., Liu, X., Meng, H., 2019. CNN-RNN-CTC based end-to-end mispronunciation detection and diagnosis. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 8132–8136.

- Levy, M., Stockwell, G., 2013. CALL Dimensions: Options and Issues in Computer-Assisted Language Learning. Routledge.
- Li, H., Huang, S., Wang, S., Xu, B., 2011. Context-dependent duration modeling with backoff strategy and look-up tables for pronunciation assessment and mispronunciation detection. In: INTERSPEECH 2011, 12th Annual Conference of the International Speech Communication Association, Florence, Italy, August 27-31, 2011. ISCA, pp. 1133–1136.
- Li, K., Mao, S., Li, X., Wu, Z., Meng, H., 2018. Automatic lexical stress and pitch accent detection for L2 english speech using multi-distribution deep neural networks. *Speech Commun.* 96, 28–36.
- Li, K., Qian, X., Kang, S., Meng, H., 2013. Lexical stress detection for L2 english speech using deep belief networks. In: *Interspeech*. pp. 1811–1815.
- Li, K., Qian, X., Meng, H., 2016. Mispronunciation detection and diagnosis in L2 english speech using multidistribution deep neural networks. *IEEE/ACM Trans. Audio Speech Lang. Process.* 25 (1), 193–207.
- Lin, B., Wang, L., 2021. Deep feature transfer learning for automatic pronunciation assessment. In: *Proc. Interspeech 2021*. pp. 4438–4442.
- Mehri Kamrood, A., Davoudi, M., Ghaniabadi, S., Amirian, S.M.R., 2019. Diagnosing L2 learners' development through online computerized dynamic assessment. *Comput. Assist. Lang. Learn.* 1–30.
- Merritt, T., Ezzer, A., Biliński, P., Proszewska, M., Pokora, K., Barra-Chicote, R., Korzekwa, D., 2022. Text-free non-parallel many-to-many voice conversion using normalising flow. In: *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 6782–6786.
- Michel, J.-B., Shen, Y.K., Aiden, A.P., Veres, A., Gray, M.K., Pickett, J.P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., et al., 2011. Quantitative analysis of culture using millions of digitized books. *Science* 331 (6014), 176–182.
- Minematsu, N., 2004. Pronunciation assessment based upon the phonological distortions observed in language learners' utterances. In: *INTER_SPEECH*.
- Needleman, S.B., Wunsch, C.D., 1970. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.* 48 (3), 443–453.
- Neri, A., Mich, O., Gerosa, M., Giuliani, D., 2008. The effectiveness of computer assisted pronunciation training for foreign language learning by children. *Comput. Assist. Lang. Learn.* 21 (5), 393–408.
- Nicolao, M., Beeston, A.V., Hain, T., 2015. Automatic assessment of english learner pronunciation using discriminative classifiers. In: *2015 IEEE Intl. Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 5351–5355.
- Oord, A., Li, Y., Babuschkin, I., Simonyan, K., Vinyals, O., Kavukcuoglu, K., Driessche, G., Lockhart, E., Cobo, L., Stimberg, F., et al., 2018. Parallel wavenet: Fast high-fidelity speech synthesis. In: *International Conference on Machine Learning*. PMLR, pp. 3918–3926.
- Peng, L., Fu, K., Lin, B., Ke, D., Zhan, J., 2021. A study on fine-tuning wav2vec2.0 model for the task of mispronunciation detection and diagnosis. In: *Proc. Interspeech 2021*. pp. 4448–4452.
- Plantinga, P., Fosler-Lussier, E., 2019. Towards real-time mispronunciation detection in kids' speech. In: *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, pp. 690–696.
- Porzuczek, A., Rojczyk, A., 2017. English word stress in polish learners speech production and metacompetence. *Res. Lang.* 15 (4), 313–323.
- Qian, X., Meng, H., Soong, F., 2010. Capturing L2 segmental mispronunciations with joint-sequence models in computer-aided pronunciation training (CAPT). In: *2010 7th Intl. Symposium on Chinese Spoken Language Processing*. IEEE, pp. 84–88.
- Ramanathi, M.K., Yarra, C., Ghosh, P.K., 2019. ASR inspired syllable stress detection for pronunciation evaluation without using a supervised classifier and syllable level features. In: *INTER_SPEECH*. pp. 924–928.
- Shah, R., Pokora, K., Ezzer, A., Klimkov, V., Huybrechts, G., Putrycz, B., Korzekwa, D., Merritt, T., 2021. Non-autoregressive TTS with explicit duration modelling for low-resource highly expressive speech. In: *Proc. 11th ISCA Speech Synthesis Workshop (SSW 11)*. pp. 96–101.
- Shahin, M.A., Epps, J., Ahmed, B., 2016. Automatic classification of lexical stress in english and arabic languages using deep learning. In: *INTER_SPEECH*. pp. 175–179.
- Statista, 2021. Most common languages used on the internet as of January 2020, by share of internet users. Available at <https://www.statista.com/statistics/262946/share-of-the-most-common-languages-on-the-internet/>.
- Sudhakara, S., Ramanathi, M.K., Yarra, C., Das, A., Ghosh, P., 2019a. Noise robust goodness of pronunciation measures using teacher's utterance. In: *SLaTE*.
- Sudhakara, S., Ramanathi, M.K., Yarra, C., Ghosh, P.K., 2019b. An improved goodness of pronunciation (GoP) measure for pronunciation evaluation with DNN-HMM system considering HMM transition probabilities. In: *INTER_SPEECH*. pp. 954–958.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks. In: *Advances in Neural Information Processing Systems*. pp. 3104–3112.
- Tejedor-García, C., Escudero, D., Cámara-Arenas, E., González-Ferreras, C., Cardeñoso Payo, V., 2020. Assessing pronunciation improvement in students of english using a controlled computer-assisted pronunciation tool. *IEEE Trans. Learn. Technol.*
- UNESCO, 2016. If you don't understand, how can you learn? Available at <https://en.unesco.org/news/40-don-t-access-education-language-they-understand>.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. In: *Advances in Neural Information Processing Systems*. pp. 5998–6008.
- Wang, J., Qin, Y., Peng, Z., Lee, T., 2019. Child speech disorder detection with siamese recurrent network using speech attribute features. In: *INTER_SPEECH*. pp. 3885–3889.
- Weber, D., Zaporowski, S., Korzekwa, D., 2020. Constructing a dataset of speech recordings with lombard effect. In: *24th IEEE SPA*.
- Witt, S.M., Young, S.J., 2000. Phone-level pronunciation scoring and assessment for interactive language learning. *Speech Commun.* 30 (2–3), 95–108.
- Xiao, Y., Soong, F., Hu, W., 2018. Paired phone-posteriors approach to ESL pronunciation quality assessment. In: *Proc. Interspeech 2018*. pp. 1631–1635.
- Xu, X., Kang, Y., Cao, S., Lin, B., Ma, L., 2021. Explore wav2vec 2.0 for mispronunciation detection. In: *Proc. Interspeech 2021*. pp. 4428–4432.
- Yan, B.-C., Jiang, S.-W.F., Chao, F.-A., Chen, B., 2021. Maximum F1-score training for end-to-end mispronunciation detection and diagnosis of L2 english speech. *arXiv preprint arXiv:2108.13816*.
- Yan, B.-C., Wu, M.-C., Hung, H.-T., Chen, B., 2020. An end-to-end mispronunciation detection system for L2 english speech leveraging novel anti-phone modeling. In: *Proc. Interspeech 2020*. pp. 3032–3036.
- Zen, H., Dang, V., Clark, R., Zhang, Y., Weiss, R.J., Jia, Y., Chen, Z., Wu, Y., 2019. LibriTTS: A corpus derived from LibriSpeech for text-to-speech. In: *Proc. Interspeech 2019*. pp. 1526–1530.
- Zhang, Z., Wang, Y., Yang, J., 2021. Text-conditioned transformer for automatic pronunciation error detection. *Speech Commun.* 130, 55–63.
- Zhang, L., Zhao, Z., Ma, C., Shan, L., Sun, H., Jiang, L., Deng, S., Gao, C., 2020. End-to-end automatic pronunciation error detection based on improved hybrid ctc/attention architecture. *Sensors* 20 (7), 1809.
- Zhao, J., Yuan, H., Liu, J., Xia, S., 2011. Automatic lexical stress detection using acoustic features for computer assisted language learning. In: *Proc. APSIPA ASC*. pp. 247–251.