



**HAL**  
open science

# End-to-end acoustic modelling for phone recognition of young readers

Lucile Gelin, Morgane Daniel, Julien Pinquier, Thomas Pellegrini

► **To cite this version:**

Lucile Gelin, Morgane Daniel, Julien Pinquier, Thomas Pellegrini. End-to-end acoustic modelling for phone recognition of young readers. *Speech Communication*, 2021, 134, pp.71-84. 10.1016/j.specom.2021.08.003 . hal-03373156

**HAL Id: hal-03373156**

**<https://hal.science/hal-03373156>**

Submitted on 16 Oct 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# End-to-end acoustic modelling for phone recognition of young readers

Lucile Gelin<sup>a,b</sup>, Morgane Daniel<sup>b</sup>, Julien Pinquier<sup>a</sup> and Thomas Pellegrini<sup>a</sup>

<sup>a</sup>IRIT, Paul Sabatier University, CNRS, Toulouse, France

<sup>b</sup>Lalilo, France

## ARTICLE INFO

### Keywords:

child speech  
phone recognition  
transformer  
connectionist temporal classification  
transfer learning  
low-resource

## ABSTRACT

Automatic recognition systems for child speech are lagging behind those dedicated to adult speech in the race of performance. This phenomenon is due to the high acoustic and linguistic variability present in child speech caused by their body development, as well as the lack of available child speech data. Young readers' speech additionally displays peculiarities, such as slow reading rate and presence of reading mistakes, that hardens the task. This work attempts to tackle the main challenges in phone acoustic modelling for young child speech with limited data and improve understanding of strengths and weaknesses of a wide selection of model architectures in this domain. We find that transfer learning techniques are highly efficient on end-to-end architectures for adult-to-child adaptation with a small amount of child speech data. Through transfer learning, a Transformer model complemented with a Connectionist Temporal Classification (CTC) objective function, reaches a phone error rate of 28.1%, outperforming a state-of-the-art DNN-HMM model by 6.6% relative, as well as other end-to-end architectures by more than 8.5% relative. An analysis of the models' performance on two specific reading tasks (isolated words and sentences) is provided, showing the influence of the utterance length on attention-based and CTC-based models. The Transformer+CTC model displays an ability to better detect reading mistakes made by children, which can be attributed to the CTC objective function effectively constraining the attention mechanisms to be monotonic.

## 1. Introduction

Speech recognition systems, enabled by the advent of high computational capacities, are nowadays widely used in various applications: human-machine interaction, home assistantship, education, tutoring, entertainment. Being more and more exposed to numerical technologies, children are becoming an important part of the public using these systems. However, although the speech recognition accuracy is approaching perfection on adult speech, the performance on child speech is far from it (Potamianos and Narayanan, 2003; Shivakumar and Georgiou, 2020). Tackling the challenges brought by child speech for speech recognition thus becomes essential to level up global accuracy levels.

Numerous research projects have been conducted on the differences between adult and child speech (Potamianos and Narayanan, 2007; Mugitani and Hiroya, 2012; Lee et al., 1999; Gerosa et al., 2006). They have shown important acoustic, linguistic and prosodic variability in child speech. Acoustic variability is caused by developmental changes in the early years of life, in particular in the speech production apparatus. The small size of the vocal tract induces shifted fundamental and formant frequencies in high frequencies, slowly transferring to lower frequencies during growth and reaching adult level around 15 years old (Mugitani and Hiroya, 2012). This evolution causes high interspeaker variability across age groups. Research shows high intra-speaker acoustic variability in the spectral domain, particularly for young children, due to the stabilisation of pitch control occurring only around age 8 (Lee et al., 1999). Linguistic and pronunciation variabilities can be attributed to the slow development of articulatory mechanisms, which are not fully developed at ages 5-7. Bad positioning of the tongue and lips can also cause phonological errors among

young children, which show to disappear with age (Fringi et al., 2015). Finally, prosodic variability is caused by developing and limited knowledge of the language: prosodic events, such as lower duration of phones, false starts, hesitations and breathing have been found more frequently from 8 to 10-year-old children's speech in comparison with older age groups (Potamianos and Narayanan, 1998).

A common application of speech recognition for children has educational purposes: reading tutors. Learning to read is indeed a challenge for young children, who might need extra help to master it. Several projects have been implemented over the years (Mostow and Aist, 2001; Bolaños et al., 2011; Proença, 2018; Godde et al., 2017), applied to different languages, age groups and reading tasks. Lalilo<sup>1</sup> provides an online reading assistant for 5-8-year-old children, featuring reading aloud exercises where children record themselves reading and get feedback on their reading. It aims at enabling extra practice in reading aloud, which is often not allowed a lot of time in class because of practical issues, while providing feedback and support in areas of difficulty for the student. Speech recognition for children learning to read is an arduous task: non-proficient readers' speech contains many disfluencies and reading mistakes that can be laborious to detect automatically. The usage of reading assistants in classroom environments adds up difficulties due to the common presence of babble noise on the recordings, sometimes concealing the target student's speech.

## 2. Related works

Prior studies on ASR for child speech have shown that the performance is below that of the latest systems dedicated to adult speech (Potamianos and Narayanan, 2003;

<sup>1</sup><https://www.lalilo.com/>

Shivakumar and Georgiou, 2020). ASR for child speech is a relatively recent research domain, that has expanded with the creation of dedicated events, for instance, the Workshop on Child Computer Interaction (WOCCI)<sup>2</sup> and the Child Speech Recognition Challenge organised during the IEEE Spoken Language Technology Workshop (Yu et al., 2020). A study on grade-specific ASR performance shows significant gaps of word error rate (WER): 26.9%, 14.6%, 10.5% and 5.1% among children of age 5, 6, 7 and 8+, confirming that the younger the children, the lower the accuracy (Yeung and Alwan, 2018). The same study shows promising results by training children-specific and even grade-specific ASR models, as long as sufficient data are available to learn despite the important acoustic variability in child speech.

Due to limited available children data in most languages, deep neural network-based systems only recently started to be exploited. Serizel and Giuliani (2014) train a hybrid Deep Neural Network - Hidden Markov Model (DNN-HMM) system on adult and children mixed data, that serves as initialisation to train a children-adapted model with only child speech. An amount of seven hours of child speech training data, ages 7-13, enables them to reach a 12.4% phone error rate. For a children language learner application, Metallinou and Cheng (2014) present a DNN-HMM that, even trained on less data, surpasses Gaussian Mixture Model - Hidden Markov Model (GMM-HMM) systems. Their best-performing model reaches a 19.3% WER with 430 hours of child speech training data. Shivakumar and Georgiou (2020) provide valuable insights on acoustic modelling for child speech recognition with DNN-HMM, by training models on adult speech, then adapting them with a dataset of 91 hours of 6-14-year-old children's speech, reaching a 17.8% WER. They also carry a deep analysis of the influence of several parameters such as the age of the speakers and the quantity of child data.

End-to-end architectures have shown to equal (or outperform) hybrid architectures on a broad range of ASR datasets (Karita et al., 2019b). Their major advantage is to discard the need for pre-segmenting training data and post-processing outputs, directly predicting the final label sequence, which considerably simplifies the process of decoding utterances. This unification of the whole training process in a single step prevents potential behavioural mismatches between modules that have been trained separately. In the same objective of simplicity, end-to-end systems mostly act at the character level, since it eliminates the process of converting the word-level transcriptions into phones, thus the need for a pronunciation dictionary. A first step towards end-to-end ASR was made by Graves et al. (2006) with the Connectionist Temporal Classification (CTC) objective function, used with success with a model based on recurrent neural networks (RNN) on the TIMIT phone recognition task. Sequence-to-sequence (seq2seq) architectures were then presented, with encoder-decoder structures (Sutskever et al., 2014; Chorowski et al., 2014, 2015; Lu et al., 2015; Chiu et al., 2018). The Listen, Attend and Spell (LAS) ar-

chitecture (Chan et al., 2016) put forward innovative methods, using an attention mechanism to link the encoder and decoder modules, based on RNNs, and bypass the conditional independence assumption made by the CTC. Transformer architectures (Vaswani et al., 2017; Dong et al., 2018) made the most of attention mechanisms by replacing the RNNs usually contained in both encoder and decoder with multi-head self-attention modules, enabling better versatility and faster computing. All seq2seq architectures (LAS, Transformer...) can also be trained with a combination of cross-entropy (CE) and CTC objective functions, as presented in Watanabe et al. (2017) and Karita et al. (2019a). This method aims at constraining the attention to be more monotonic, as well as helping the system to converge faster. A recent study demonstrated the superior ASR performance of the Transformer trained with CE+CTC objectives on diverse individual corpora and multilingual tasks, in comparison with LAS-like and hybrid DNN-HMM models (Karita et al., 2019b).

Due to the limited available child speech data, very few studies present child speech recognition with end-to-end models. Methods though exist, like Transfer Learning (TL) that consists of training an acoustic model on a large out-of-domain dataset, then retraining it with a small in-domain corpus to retain part of the acquired knowledge and adapt it to the application domain. It showed significant performance improvement with hybrid DNN-HMM approaches for adult-to-children transferring applications (Shivakumar and Georgiou, 2020; Tong et al., 2017a; Qian et al., 2016). TL was also used on end-to-end architectures for low resource languages (Cho et al., 2018; Tong et al., 2017b).

As for end-to-end models for child ASR, Andrew et al. (2015) show improvement on child speech with a CTC-based system jointly trained on very large quantities of mixed adult and child speech data. Usage of seq2seq models for child speech recognition is a new research subject, as show the extremely recent communication of technical reports on this matter (Ng et al., 2020; Chen et al., 2020). Both studies use the same dataset, containing 59 hours of Mandarin speech from 7 to 11-year-old children. The former uses transfer learning for adult-to-child adaptation and reports a 23.6% character error rate (CER) with a CE+CTC multi-objective RNN-based seq2seq model. On the other hand, the latter mixes adult and child speech and announces 18.8% CER with a convolutional-augmented Transformer model and data augmentation techniques. An even more recent study reports a comparison of diverse end-to-end systems for automatic child speech recognition, but they use a substantial amount of data (223 hours) from children that are quite older (9-11 years old) than in our study, and focus on conversational speech over read speech (Shivakumar and Narayanan, 2021). They obtain a 9.2% CER and 16.0% WER on matched test data with a self-attention-based encoder trained with a CTC objective function. On test data with a wider age range (6-16), they obtain a 33.6% WER with a Time-Depth Separable Convolution Network, with significant performance degradation for 5-8 years old chil-

<sup>2</sup><https://sites.google.com/view/wocci>

dren due to acoustic differences across ages and higher intra-speaker variability in young children’s speech. While these recent works do character recognition, as end-to-end systems were designed for, in this work we stick to modelling phones instead, in agreement with our final application: pronunciation assessment of children learning to read. Furthermore, these studies lack a detailed evaluation of the strengths and weaknesses of seq2seq models in comparison with state-of-the-art hybrid DNN-HMM approaches, when applied to child speech in a low-resource scenario.

Our work attempts to furnish valuable insights in this new research domain. We apply a wide catalogue of end-to-end ASR systems to child speech, to:

- Compare recent end-to-end models (RNN, LAS, Transformer) to a baseline hybrid DNN-HMM model for phone recognition;
- Show the efficiency of transfer learning on end-to-end models for adult-to-child adaptation, even with very low amounts of child data;
- Analyse the influence of different components of end-to-end architectures, such as the attention mechanism or the connectionist temporal classification (CTC) objective function, on the performance;
- Evaluate the models on different reading tasks: isolated words and sentences;
- Assess the impact of CE+CTC multi-objective training for evaluation of utterances containing reading mistakes;

This article is organised as follows. Section 3 presents the adult and child speech material used in this study, and offers a detailed description of young readers’ typical mistakes. Section 4 provides comprehensive explanation of our transfer learning strategy. Section 5 presents the acoustic models that we apply to child speech in this paper: the baseline, a hybrid DNN-HMM model, as well as several end-to-end models. Section 6 shows the validation of our models on adult speech, and studies the results obtained on child speech. Finally, section 7 offers analyses that provide further insights on end-to-end acoustic models’ behaviour when confronted with different reading tasks and reading levels.

### 3. Speech material

We use two datasets of French speech: the *Common Voice* adult corpus<sup>3</sup>, and an in-house child speech corpus, hereinafter called *Lalilo*. Tables 1 and 2 display general information on these data.

#### 3.1. Adult dataset: Common Voice

The Common Voice corpus is created through a participatory online platform, where everyone can record himself

<sup>3</sup>Corpus available at: <https://voice.mozilla.org/fr>

**Table 1**  
Information on Common Voice (CV) adult speech dataset

Dataset	Train	Valid	Test
Duration (h)	148.9	2.4	7.2
Speakers	1276	372	1113
Avg time per spk (s)	420.1	23.5	23.4
SNR (dB)			
Mean	34.4	34.3	34.3
Std	14.7	14.5	14.7

reading sentences. Being composed of recordings with different equipment in different environments, it is thus particularly adapted to our reading speech task in classrooms. Nonetheless, since speakers usually record themselves alone in their computer room, the data does not contain babble noise and has a high mean signal-to-noise ratio (SNR). In French, the Train, Valid and Test sets we used for these experiments contain respectively 148.9, 2.4 and 7.2 hours of speech (see Table 1). The Test set has been designed to maximise the number of distinct speakers while keeping the same average time per speaker as in the Valid set. Each recording is validated by two annotators, thereby the corpus contains a limited quantity of miscues.

#### 3.2. Child dataset: Lalilo

The Lalilo corpus contains recordings of Kindergarten to second-grade children, aged 5-8, reading aloud isolated words and sentences. These two tasks are commonly given by teachers to beginning readers, according to their reading level. Detailed information is displayed in Table 2. The recordings have been gathered either in person directly in schools or through a reading-aloud exercise in the Lalilo web platform. In the first case, the environmental conditions are somewhat clean: a good-quality microphone is used, and the noise level is controlled. In the second case, however, teachers usually let a small group of students play on the platform, with reduced supervision, which inevitably implies a highly variable level of babble noise on the recordings. Because schools do not always benefit from headsets, computer built-in microphones are sometimes used, which record a wide range of ambient noise. Table 2 displays the mean SNR for each train and test sets and shows that the Lalilo corpus contains more noise than the Common Voice corpus (see Table 1).

The training set contains sentences and isolated words, with an equal proportion of isolated words and words from sentences. Similarly, we use a validation set that contains both isolated words and sentences. The reading assistant offers children to read either isolated words or sentences, according to their reading level, thus we constructed a test set that can be split in two, each corresponding to one task. They are named Test W (for words) and Test S (for sentences) in Table 2.

Phonetically transcribing with exactitude what has been read by non-reader children is a challenging task, due to ar-

**Table 2**

Information on Lalilo child speech dataset. The "Test W" and "Test S" respectively designate the isolated words and sentences test sets. The "% correct read" designates the percentage of phones belonging to an utterance that contains no reading mistake.

Dataset	Train	Valid	Test	Test W	Test S
Duration (h)	13.0	0.41	1.32	0.84	0.48
Speakers	3014	459	685	425	262
Avg time per spk (s)	15.2	3.2	6.9	7.0	6.7
% correct read	100	100	57.6	48.5	64.0
SNR (dB)					
Mean	20.9	20.1	20.6	22.2	22.8
Std	13.1	12.7	12.6	13.2	12.1

tication mechanisms in the process of being acquired, potential hesitations and stammering, or even the presence of acoustically non-existing phones. Although it is not an impossible task, it is very time-consuming and costly, and we have chosen not to invest in phone-level transcriptions for the whole corpus at the moment. Accordingly, only correctly pronounced and fluently read utterances that are identical to the prompted text, were included in the training and validation data for both the acoustic and language models. Only the test data, containing words and sentences with reading mistakes, has been phonetically transcribed to enable the computation of performance scores. The transcriptions have been made manually by two human judges, and recordings have been discarded in case of disagreement.

### 3.3. Reading learners' speech

Reading learners have heterogeneous reading levels, according to their age, grade, time spent reading and other factors. Their speech hence contains a wide variety of reading rates and amount of reading mistakes.

Students' reading rates are measured by teachers with a Word Correct Per Minute (WCPM), i.e. the number of correctly read words in a minute interval. They compare this score to national standards to assess their students' progression. Very slow and very fast readings can severely damage the recognition by causing insertions and deletions.

Mistakes done by reading learning children are diverse and sometimes unique, which makes the manual annotation of data quite difficult for human experts, and can cause all the more confusion for an ASR system. Reading mistakes can be categorised into two main categories, word-level and sentence-level, then into several subcategories. Understanding the different types of reading mistakes can help analyse the system's behaviour when encountering these.

Reading mistakes at the word level can usually enter in one (or both) of the following categories:

- **Hesitation:** presence of intra-word silence(s) due to the child hesitating when reading;
- **Substitution:** a phone is substituted by another.

Example: *cat* is read *cut*;

- **Insertion:** a phone is inserted inside a word.

Example: *cat* is read *cart*;

- **Deletion:** a phone is deleted from a word.

Example: *cat* is read *at*.

For substitutions, insertions and deletions, the resulting word can either exist or not.

Reading mistakes can also occur at the sentence level, like repetitions of one or several words, false starts (the child attempts to read the word but the word is not read entirely), or hesitations between words.

Word-level mistakes are more difficult to detect (for both humans and machines) than sentence-level mistakes, as they usually concern one or two phones: in particular, substitutions and insertions of phones are easy to miss. Furthermore, the data we gathered comes from the Lalilo reading assistant that adapts the difficulty of exercises to each child's reading level: children who were asked to read isolated words were more susceptible to having a lower global reading level than children who had to read sentences, causing the isolated word recordings to contain very laborious readings. The Lalilo Test W set thus offers a particularly arduous task.

A word is considered as erroneously read when it contains at least one of the reading mistakes listed above. A sentence is considered as such when it contains at least one erroneous word or a sentence-level reading mistake. Table 2 displays the percentage of phones included in correctly-read utterances in Lalilo Test W and S sets: respectively 48.5% and 64.0% of phones belong to utterances that contain no reading mistake. These percentages are in agreement with the recordings gathered on the Lalilo platform: children reading at lower levels are given words to read, thus more than half of isolated words contain a reading mistake, while more advanced students are given sentences to read and statistically make fewer reading mistakes.

## 4. Tackling the low-resource configuration

Transfer learning (TL) can be done on many ASR applications, such as adaptation between languages (Abad et al., 2020; Tong et al., 2017b; Cho et al., 2018), between different speech types, such as broadcast news and conversational speech in Abad et al. (2020), or between native and non-native speech (Duan et al., 2020). According to the application, the available quantity of target data, or the similarity between the source and target domains, it will be applied differently. In our case, adult-to-child adaptation, the source model is a model trained on adult speech, while the adaptation data is child speech data. This particular adaptation is very sensitive to the amount of target data and the target children's age, since their vocal apparatus and speech quality are so different than adults, and vary greatly during children's growth (Shivakumar and Georgiou, 2020). We follow the findings of Shivakumar and Georgiou (2020): overall, they find that the best adaptation configuration is either to adapt the whole network or to adapt two layers at the bottom

and two layers at the top, both strategies obtaining the same score. However, for young (6-8 years old) children, they show that the high complexity and variability in acoustic and prosodic characteristics of their speech necessitate more parameters (hence layers) to be fully captured, suggesting that the first option should be chosen. A minimum amount of adaptation data is nevertheless required (~10 hours): below, adapting the whole network would hurt performance due to noise introduced by high variability. When sufficient adaptation data are provided, the network is able to learn young child speech variability. In these conditions, adapting all the layers of the model performs better than only some layers. Our training child corpus containing approximately 13 hours of speech from very young children (5-8 years old), we follow the advice and choose to apply transfer learning [by re-training all the layers of the source model](#) in our experiments.

For each acoustic model architecture presented in the next section, we trained a model on adult speech to serve as a source model, then used TL to adapt this model with our child speech data.

## 5. Phone-level acoustic models

This study aims at comparing the performance of state-of-the-art end-to-end and baseline hybrid DNN-HMM ASR systems when applied to the challenging task that is child speech recognition at the phone level, with a low quantity of data and a noisy environment. This section will present the different acoustic models that we review in this work. We will train, with each architecture, two separate models:

- Adult model, trained on *Common Voice* data with almost 150 hours;
- TL model, using the adult model and adapting it with *Lalilo* child speech data (~13 hours).

### 5.1. TDNNF-HMM: the baseline

Acoustic models for speech recognition were originally built from a Hidden Markov Model (HMM), linked to either a single Gaussian distribution or a more refined Gaussian Mixture Model (GMM). The latter models the state-output distribution in an unsymmetrical and multi-modal manner, in opposition to the former, enabling the accurate modelling of speaker, accent or gender differences in speech (Gales and Young, 2008). With the advent of deep neural networks (DNN), the GMM-HMM models were progressively replaced by hybrid DNN-HMM acoustic model architectures. This change of paradigm introduces a slight change in probability computing: while GMM directly provide the posterior probability  $P(x_t|s)$  of an observation  $x$  at time  $t$  given a state  $s$ , used by the HMM, the DNN computes  $P(s|x_t)$ , from which  $P(x_t|s)$  can be computed with Bayes' rule. Although GMM-HMM models are not anymore used for acoustic modelling, they keep the purpose of generating the alignments on which DNN-HMM models are trained.

Time-Delay Neural Networks (TDNN), introduced for phone recognition by Waibel et al. (1989), showed to be

particularly adapted for ASR, with their ability to represent relationships between acoustic events in time while providing time-invariance of the features learnt by the network. It introduces delays in the weighted sum of inputs which, in practice, are implemented as spatially expanded units over a certain number of frames. Context width varies depending on the layers: bottom layers learn short duration acoustic-phonetic characteristics, while top layers learn more complex features of longer duration. For instance, TDNNs have been used with success for vowel recognition on children's speech in a small-resource language (Yong and Ting, 2011).

We use as a baseline a Factorised Time-Delay Neural Network (TDNNF) (Povey et al., 2018), a refinement of TDNN models, with a Lattice-Free Maximum Mutual Information (LF-MMI) criterion (Povey et al., 2016). These models are currently state-of-the-art in speech recognition and have shown to reach an 11.7% WER on 55 hours of child speech data (Wu et al., 2019).

We use the Kaldi toolkit (Povey et al., 2011) to train our TDNNF acoustic models, which follow available recipes<sup>4,5</sup>. A 3-state GMM-HMM monophone model provides alignments to train the TDNNF: it gave better results than using triphone GMM models, due to a small number of occurrences of each triphone in the child speech training dataset, as well as the presence of non-existing phones (reading mistakes) in the testing set. The network has a chain model architecture (Povey et al., 2016), and is trained with the LF-MMI criterion, similarly to Veselý et al. (2013). The network contains 12 TDNNF layers of dimension 1024 as hidden layers, each comprised of one TDNN layer of dimension 1024 and two bottleneck layers of dimension 128. [A smaller number of layers led to worse results, both on adult and child speech.](#) The input features are 40-dimensional Mel-frequency cepstral coefficients (MFCC) with Cepstral Mean and Variance Normalisation (CMVN), as do Bayerl and Riedhammer (2019), as well as classical Kaldi recipes. [Chain models operate a 3-frame sub-sampling for computation acceleration, which additionally enables to perform frame-shifting data augmentation.](#) Adult and TL models are trained on 990 and 89 epochs respectively, both with a learning rate of  $5e-4$  and a  $l_2$  regularisation rate of  $1e-2$ . The adult and TL models have a size of 7.6M parameters, and respectively took 46 and 2.1 hours to train on a single GTX 2080 Ti GPU.

### 5.2. End-to-end systems

Hybrid approaches have made possible the use of recurrent neural networks –which are highly efficient for modelling time series (Graves et al., 2013)– for sequence labelling, by combining them with HMMs that provide pre-segmentation of the input sequence for training and post-transformation into the output label sequence. However, the training process of DNN-HMM models is quite complex, with the need of generating alignments of each speech utterance and then using lattices to infer a label sequence.

<sup>4</sup><https://colibris.link/script-TDNNF>

<sup>5</sup><https://colibris.link/script-TDNNF-TL>

Furthermore, its functioning relies on several parts (acoustic model, language model, pronunciation model), which are trained independently with different objectives, causing behavioural mismatches between components.

End-to-end architectures aim at simplifying and unifying the training process of a speech recognition system, by encompassing all components into a single neural network. In this objective, end-to-end models are usually trained on characters, which can be directly derived from words, instead of phones that can only be obtained with a pronunciation dictionary. In this work, we will stick with phones, since it suits better for modelling the grapheme-phoneme correspondence that children learn when starting reading.

Given a sequence of input speech features  $X = (x_1, \dots, x_T)$ , and the corresponding phone-label output sequence  $Y = (y_1, \dots, y_N)$ , with  $T$  the number of frames in the speech utterance and  $N$  the length of the phone sequence, the objective is to learn the conditional probability of  $y_i$  given previous outputs  $y_{<i}$  and input  $X$ :

$$P(Y|X) = \prod_{i=1}^N P(y_i|X, y_{<i}) \quad (1)$$

During the inference process, these probabilities are computed from the input test audio frames and the previously predicted phones (not having access to the reference output phones). A beam search algorithm is then applied, with a beam size of 5 and a maximum hypothesis length of 30 phone labels for the Lalilo word set, and 130 for the Lalilo sentence set and the adult set. No language model is used, as we found that it tends to reconstruct existing words and cover up children's reading mistakes, while we aim at detecting as accurately as possible the phones uttered by children, may they form existing words or not.

Among the existing end-to-end systems, two paradigms stand out in the literature: the CTC-based and the seq2seq architectures. Different methods have been implemented with the latter, using recurrent neural networks and/or attention mechanisms. Combinations of the two paradigms have also shown great performance. The following sections present the different end-to-end architectures we explored and applied to child speech recognition.

### 5.2.1. RNN-CTC model

The CTC paradigm, introduced by Graves et al. (2006), discards the obligation of having an HMM by learning automatically alignments between the input and output sequences. In phone recognition applications, it aligns the speech frames  $X = (x_1, \dots, x_T)$  and their phone sequence  $Y = (y_1, \dots, y_N)$ , with the condition that  $N \leq T$ , where  $T$  and  $N$  are respectively the number of frames in the speech utterance and the length of the phone sequence. The length difference between  $T$  and  $N$  is coped with the addition of a "blank" label, represented by "-", which activation corresponds to the probability of observing no label. During alignment, the system learns a set of possible paths that are phone-label sequences of length  $T$  and which probability is the product of all probabilities  $p_k^t$  of observing a label  $k$  at

time  $t$ . Considering a phone-label sequence  $Y$ , it can be obtained through a set  $\theta_Y$  of paths  $\pi$ , the total probability of  $Y$  being defined as follows:

$$P(Y|X) = \sum_{\pi \in \theta_Y} \prod_{t=1}^T p_{\pi_t}^t \quad (2)$$

CTC-based networks are trained with gradient descent and the maximum likelihood function, which maximises the log probabilities of the target labels. The probability  $P(L|X_{1,\dots,T})$  of each individual labelling  $L$  is computed with the CTC forward-backward algorithm, which divides the sum over  $T$ -long paths  $\pi_{1,\dots,T} \in \theta_{L_{1,\dots,T}}$  into an iterative sum over  $t$ -long paths  $\pi_{1,\dots,t} \in \theta_{L_{1,\dots,t}}$ . The iterations are then dynamically computed with forward and backward propagation, which makes possible the computation of the many possible paths.

We use, in this work, an end-to-end CTC architecture, named RNN-CTC and shown in Figure 1, which is composed of a simple encoder with recurrent neural networks. The input features are 40-dimensional MFCC with CMVN, as do Bayerl and Riedhammer (2019). The RNNs are composed of Bidirectional Gated Recurrent Unit layers (BiGRU) (Chung et al., 2015). They contain one Bi-GRU input layer with 120 cells, then four Bi-GRU hidden layers with  $2 \times 160$  cells in each layer, and finally a linear output layer of dimension 34, corresponding to 33 French phones and the CTC blank label. For an increased speed and better concentration of information, we reduced the time resolution of the input by concatenating pairs of consecutive input frames. On the output of each hidden layer is applied a 10% dropout rate. The output layer uses the log-softmax activation function. Models were trained on up to 100 epochs, with an early stopping mechanism linked to the validation loss. Training was done with the Adam optimizer. A grid search on the hyper-parameters was done. The adult RNN-CTC models were trained with a batch size of 100 and a  $9e-5$  learning rate. TL models were trained with a batch size of 50 and a  $1e-4$  learning rate. We used a learning rate scheduler that divided by 10 the learning rate after 2 epochs with no improvement in the validation loss. The adult and TL models took an average of 24.3 and 4 hours to train on a single GTX 2080 Ti GPU. They each contain a total of 2.1M trainable parameters.

### 5.2.2. Seq2seq LAS model

At first presented for machine translation tasks, where an input sequence of words in a language needs to be translated to an output sequence of words in another language, seq2seq architectures aim at addressing the problem of performing sequence labelling with variable-length input and output sequences without relying on HMMs like in hybrid DNN-HMM approaches (Sutskever et al., 2014). This framework has shown to be generalisable to many applications: image captioning (Vinyals et al., 2015; Xu et al., 2015), conversational modelling (Vinyals and Le, 2015), and, of course, acoustic modelling for speech recognition. ASR tasks are

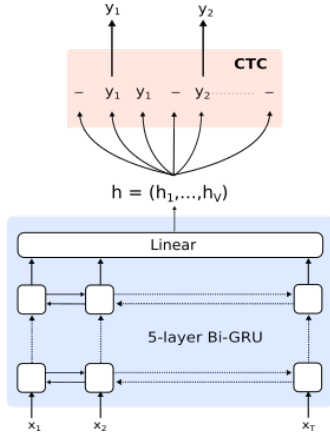


Figure 1: Architecture of the RNN-CTC model

characterised by long input sequences (speech frames) in comparison with other tasks. Unlike CTC-based models, seq2seq models do not assume that the label outputs are conditionally independent of each other, which enables greater freedom but can also cause more frequent confusions.

Seq2seq models are usually composed of an encoder that extracts information from the audio input and maps it into a variable-length vector, and a decoder that takes this vector and generates an output sequence of tokens (characters, phones or any other speech units), one at a time. The use of an attention mechanism (Bahdanau et al., 2015), which searches through the encoder output to locate parts of the sequence that are valuable for predicting target tokens, greatly improves the performance. This mechanism generates at each decoder step an attention vector, which conveys some context information from the encoder to the decoder.

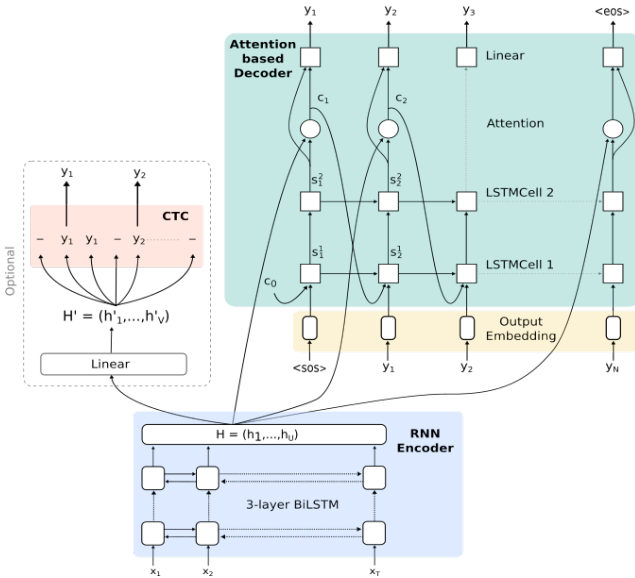


Figure 2: Architecture of the LAS and LAS+CTC models

One of the seq2seq models used in this paper is sim-

ilar to a LAS architecture, recently introduced by Chan et al. (2016). Although our architecture, shown in Figure 2, slightly differs from the initial one, we will in the next sections borrow the name "LAS" for our systems.

The input features to the encoder are 40-dimensional filterbanks with first and second derivatives. We use identical inputs for LAS and LAS+CTC models, following Watanabe et al. (2017). The encoder contains three layers of bidirectional Long-Short Term Memory (bLSTM) layers, which output a matrix  $H = (h_1, \dots, h_U)$ , where each  $h$  has a dimension 256 and  $U$  corresponds to the number of audio frames. A dropout rate of 0.2 is applied to the encoder layers. The decoder takes as input, at each step  $i$ , the 512-dimension embeddings of the label output  $y_i$ , and passes them through a LSTM layer that outputs a state  $s_i^{(1)}$  following Equation 3. A second LSTM layer outputs a state  $s_i^{(2)}$  from the output of the previous layer  $s_i^{(1)}$  and the previous state of the current layer  $s_{i-1}^{(2)}$  (see Equation 3).

$$\begin{aligned} s_i^{(1)} &= \text{bLSTM}(s_{i-1}^{(1)}, y_{i-1}, c_{i-1}) \\ s_i^{(2)} &= \text{bLSTM}(s_i^{(1)}, s_{i-1}^{(2)}) \end{aligned} \quad (3)$$

The  $c_i$  component is the context vector of dimension 512, computed by the attention mechanism from the current state  $s_i^{(2)}$  and the encoder output  $H$  with Equation 4. We use here a simple dot-product attention mechanism, which considers the decoder state  $s_i^{(2)}$  as the query, and the encoder output  $H$  as both the keys and values.

$$\begin{aligned} c_i &= \text{DotProductAttention}(s_i^{(2)}, H) \\ &= \text{Softmax}(s_i^{(2)} H^T) H \end{aligned} \quad (4)$$

Finally, a Multi-Layer Perceptron (MLP), composed of two linear layers (size 512) separated by a tanh activation layer, provides the probability  $p(y_i|X, y_{<i})$  of equation 1, taking as input the decoder state  $s_i^{(2)}$  and the attention context  $c_i$ :

$$p(y_i|X, y_{<i}) = \text{MLP}(s_i, c_i) \quad (5)$$

Our LAS models are trained with the Adam optimizer, using  $l_2$  regularisation at a rate of  $1e-5$ , and a learning rate of  $1e-4$ , the latter being halved after each epoch the validation loss did not improve. Scheduled sampling (Bengio et al., 2015) is used to train the LAS models. This method, which consists of feeding the decoder either the reference label or the lastly predicted label, aims at helping the model converge faster and providing more robustness to inference mistakes. The system chooses to input the lastly predicted label instead of the reference label with a fixed probability of 10%. All models are trained on 50 epochs, where only the best-validated model is kept. The adult and TL models take an average of 50.7 and 3.9 hours to train on a single GTX 2080 Ti GPU. The final models each contain a total of 7.6M parameters.



### 5.2.3. Seq2seq LAS+CTC model

The simple attention mechanism used in the LAS is too flexible for ASR tasks, because it allows non-sequential alignments, while speech recognition inputs and outputs are generally sequential (Watanabe et al., 2017). Another characteristic of speech recognition is that the input and output sequences can vary greatly in length, especially with reading learners with different reading levels, which makes the alignments more difficult to track. Since CTC enables to compute constrained monotonic alignments, it was proposed by Watanabe et al. (2017) to combine both paradigms into a hybrid CTC/attention-based system. During training, a CE+CTC multi-objective learning method combines the two losses. Two outputs, one from the encoder CTC and the other from the attention-based decoder, can generate output sequences.

The CTC-attention hybrid models follow the same architecture and use the same hyper-parameters as the LAS models described in the previous section. It also uses the same input features, filterbanks of dimension 40, with first and second derivatives, as do Watanabe et al. (2017). Scheduled sampling is applied with a 10% rate. The CTC loss is combined with the attention loss following Equation 6, with  $\lambda = 0.2$ . Other values of  $\lambda$  (0.5, 0.8) gave worse results, as found by Watanabe et al. (2017).

$$\text{loss} = \lambda \times \text{loss}_{CTC} + (1 - \lambda) \times \text{loss}_{CE} \quad (6)$$

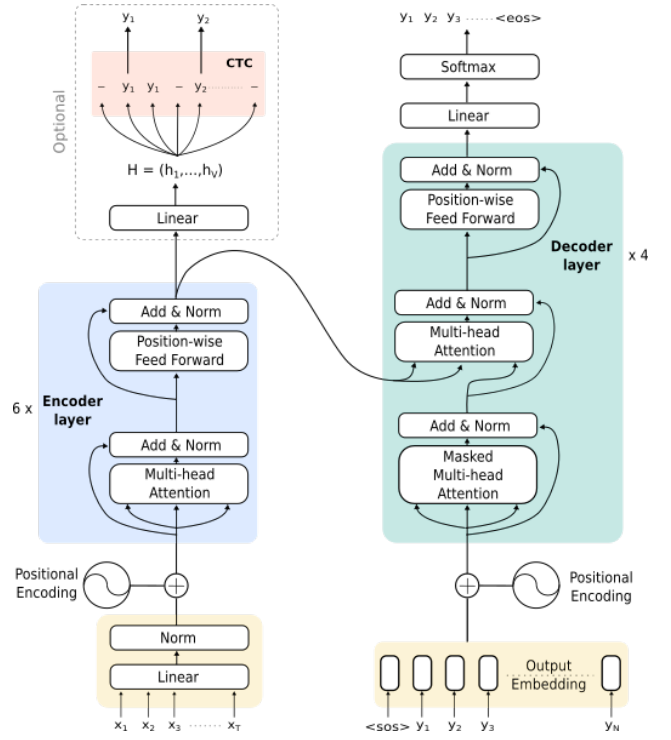
The attention-based decoder and encoder CTC outputs are taken separately to generate two sets of hypothesised label sequences. Two scores will thus be presented for the LAS+CTC model, which will be denoted as *LAS+CTC enc* and *LAS+CTC dec* for the encoder CTC and attention-based decoder outputs, respectively.

### 5.2.4. Seq2seq Transformer model

Presented by Vaswani et al. (2017) and adapted to speech recognition by Dong et al. (2018), the Transformer model follows a Seq2seq encoder-decoder architecture, but relies solely on attention mechanisms, instead of recurrent neural networks in classical Seq2seq systems. The recurrence, essential to extract position information in speech frames, is replaced by positional encodings concatenated to the input encodings, as well as multi-head self-attention mechanisms and position-wise feed-forward neural networks in the encoder and decoder blocks. Discarding the need for recurrent neural networks enables to compute dependencies between each pair of positions at once, instead of one by one. It allows for faster training and more parallelisation in comparison with LAS systems presented in previous sections.

Figure 3 presents the architecture of our Transformer, with the optional CTC module (see section 5.2.5).

The Transformer models' encoder takes as input the audio sequence  $X = (x_1, \dots, x_T)$ , where  $x_i$  is a 80-dimensional filterbank vector. The inputs are processed with linear and normalisation layers of size  $d_{\text{model}} = 256$ . Positional encodings give information about the relative position of the tokens in the sequences, replacing the recurrence or convo-



**Figure 3:** Architecture of the Transformer and Transformer+CTC models

lution used in classical models. These encodings have the same dimension (256) as input/output embeddings to ease the summation between the two vectors. As in literature (Vaswani et al., 2017; Dong et al., 2018; Karita et al., 2019b), we use sinusoidal positional encodings, defined as:

$$\begin{aligned} \text{PosEnc}_{(pos, 2i)} &= \sin(pos/10000^{2i/d_{\text{model}}}) \\ \text{PosEnc}_{(pos, 2i+1)} &= \cos(pos/10000^{2i/d_{\text{model}}}) \end{aligned} \quad (7)$$

where  $pos$  is the audio frame or label position, and  $i$  the  $i$ -th dimension of the positional encodings, with  $0 \leq i < d_{\text{model}}$ .

The encoder contains six layers, each layer being composed of two sub-layers: a multi-head self-attention module, and a position-wise feed-forward fully connected neural network module. Each of these modules is followed by a normalisation layer with residual connection (He et al., 2016). The first sub-layer contains a multi-head scaled dot-product attention (Vaswani et al., 2017), which relates different positions of inputs to create more valuable representations. Each of the  $i \in \{1..h, h = 4\}$  heads follows Eq. 8, where  $Q_i$ ,  $K_i$  and  $V_i$  represent respectively the  $i$ -th linear projection of the query, keys and values, and  $d_k = 256$  is the dimension of  $Q$ ,  $K$  and  $V$ . The scaling by  $1/\sqrt{d_k}$  restrains the dot product from growing too large when  $d_k$  is large, which would give very small gradients with the softmax function. The final output of this module is a linear projection of the concatenation of these 4 heads. The second sub-layer is composed of a position-wise feed-forward neural network that consists of two linear layers separated by a ReLU activation and is

applied to each position separately.

$$\begin{aligned} \text{head}_i &= \text{ScaledDotProductAttention}(Q_i, K_i, V_i) \\ &= \text{Softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right) V_i \end{aligned} \quad (8)$$

The inputs of the decoder, which are the reference phone labels during training, and the previously decoded phone labels during inference, are offset by one position with the start-of-sequence (< sos >) token, and go through an embedding layer of dimension 256. As for the encoder inputs, positional encodings are computed and summed to the decoder's inputs. Unlike LAS models, which decode token by token, the Transformer's decoder takes as input the whole utterance at the same time, enabling it to parallelise and decode faster. This characteristic however prevents from using scheduled sampling as defined by Bengio et al. (2015), where the previously predicted token is needed to predict the current token.

The four layers of the decoder each contain three sub-layers. The first is the same self-attention module as in the encoder, with the difference that it is masked. This mask, combined with the offset by one position of the phone-label embeddings serving as inputs to the decoder, ensures that for a position  $i$ , the predictions depend only on prior  $i - 1$  positions, sole known positions in inference mode. The second sub-layer is another multi-head attention module, which functions as the attention mechanisms in other seq2seq systems: the encoder stack outputs the keys and values, while the query is provided by the prior sub-layer of the decoder. The third sub-layer is a position-wise feed-forward neural network, which is identical to the second sub-layer of the encoder.

The decoder output is finally processed through a classification network composed of a linear projection and a softmax operation, which yield the final phone-label sequence.

The Transformer models are trained with the Adam optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$  and  $\epsilon = 1e-9$ . We use the same specific learning rate scheduler as Vaswani et al. (2017) for training, defined in Equation 9, with `warmup_steps = 4000`.

$$\text{lr} = d_{\text{model}}^{-0.5} \cdot \min(\text{step}^{-0.5}, \text{step} \cdot \text{warmup\_steps}^{-1.5}) \quad (9)$$

Models are trained on 100 epochs and we score with the best-validated model. The adult and TL models respectively take on average 30 and 3 hours to train on a single GTX 2080 Ti GPU. The final model contains a total of 14.3M parameters.

### 5.2.5. Seq2seq Transformer+CTC model

In the same objective as with the LAS+CTC system, we can optionally use the CTC objective function to improve the Transformer model training and inferring processes, as has been shown by Karita et al. (2019a). In the same way, the CTC and attention objectives can be combined with Equation 6.

**Table 3**

PER (%) obtained with adult-trained acoustic models tested on adult speech

Model	Test CV
TDNNF-HMM ( <i>baseline</i> )	23.5
RNN-CTC	16.1
LAS	12.6
LAS+CTC enc	16.7
LAS+CTC dec	11.9
Transformer	<b>7.5</b>
Transformer+CTC enc	11.9
Transformer+CTC dec	8.0

The Transformer+CTC models use the same inputs, filterbanks of dimension 80, as do Karita et al. (2019a,b). The models were trained with the same process as the Transformer, and the combination of objectives worked the best with  $\lambda = 0.3$  as in Karita et al. (2019a). The adult model trained during 100 epochs, while 10 epochs were sufficient to train the TL model, resulting in a very short training time for the latter (below one hour).

In the same way as the LAS+CTC, the Transformer+CTC model has two outputs, one from the encoder with the CTC function and one from the decoder with the attention mechanism. These outputs are denoted *Transformer+CTC enc* and *Transformer+CTC dec*, respectively, in the following.

## 6. Evaluation

With the implementation of a phone recognition system for a reading assistant application, we do not aim at reconstructing and correcting words based on the detected phones, but at transcribing accurately what the child has read, including potential phone-level reading mistakes. Therefore, we do not measure performance with the classical WER but with a Phone Error Rate (PER). The PER metric considers all mismatches between the recogniser hypothesis and the manual phone-level annotated reference, and is defined in Equation (10), with  $C$ ,  $I$ ,  $S$ ,  $D$  respectively referring to the number of correct detections, insertions, substitutions and deletions.

$$\text{PER} = \frac{I + S + D}{C + S + D} \quad (10)$$

### 6.1. Comparison of models on adult speech

Adult models were trained with about 150 hours (see Table 1). We score their performance with the PER metric on the Test CV adult speech set, as displayed in Table 3.

We observe that the models' performances follow their chronological appearance in the literature: the worse is the baseline, the TDNNF-HMM model with a PER of 23.5%, and the best is the Transformer with a PER of 7.5%. In addition to being simpler to train, end-to-end models all perform significantly better than the baseline hybrid TDNNF-

**Table 4**  
PER (%) obtained with acoustic models without (adult-trained) and with TL, tested on child speech

Model	Without TL	With TL
TDNNF-HMM ( <i>baseline</i> )	<b>45.9</b>	30.1
RNN-CTC	59.4	32.4
LAS	77.9	33.9
LAS+CTC enc	63.5	32.7
LAS+CTC dec	72.9	30.7
Transformer	76.6	28.5
Transformer+CTC enc	60.0	29.3
Transformer+CTC dec	70.9	<b>28.1</b>

HMM model (from 29% to 68% relative improvement on the PER). Attention mechanisms show to be particularly efficient: the LAS, LAS+CTC dec and Transformer+CTC dec systems give lower PER than the RNN-CTC or the TDNNF-HMM, and the Transformer, which is solely based on attention, reaches the lowest PER. The CTC objective function brings valuable information for training (LAS+CTC dec is better than LAS) but inferring hypothesis sequences with the CTC output is not as accurate as using the attention-based decoder output (LAS+CTC enc is significantly worse than LAS+CTC dec). We did, however, not succeed in exploiting the previously shown potential of the CTC objective function for the Transformer, resulting in a slightly worse performance for model Transformer+CTC dec in comparison with the Transformer. As seen on LAS+CTC models, using the encoder CTC output of the Transformer+CTC for inferring sequences does not provide comparable performance with using the decoder output.

## 6.2. Comparison of models on child speech

In this section, we test the adult-trained model on child speech and compare their performance with models trained with transfer learning. We train TL models from adult models, so they benefit from the 150 hours of adult speech, as well as the 13 hours of child speech. Table 4 displays the PER scores for all architectures when trained only on adult speech (first column) or adapted with child speech (second column).

The six systems can also be trained directly on the child training set to obtain child-adapted acoustic models. **While the TDNNF manages to learn from this small amount of data, reaching a PER of 32.5%, end-to-end models, which necessitate a bigger amount of data to correctly learn speech representations, all obtain PER scores above 50%.**

### 6.2.1. Adult models on child speech

The difference in score for adult models on adult speech (Table 3) and child speech (Table 4) is drastic: all models lose on average 51 points of PER. These extreme gaps in performance may be explained by the acoustic and prosodic difference between adult and child speech.

Contrary to what has been observed for adult models when tested on adult speech, the TDNNF-HMM performs

significantly better than its end-to-end counterparts when tested on child speech: its PER score is doubled, while end-to-end systems' scores suffer from multiplicative factors that go from 3.8 (LAS+CTC enc) to 10.2 (Transformer). **Out of curiosity we tested the TL models back on adult speech, and the end-to-end systems performed very badly, suggesting that these mechanisms adapt very specifically to acoustic characteristics of data they are trained on.** The CTC objective function helps the attention module during training, as show the comparisons between LAS and LAS+CTC dec models, and between Transformer and Transformer+CTC dec models. Additionally, inferring phone sequences with the encoder CTC output shows better performance than with the attention-based decoder output for both LAS and Transformer models. Finally, the RNN-CTC model obtains a significantly better score than attention-based models. These observations suggest that attention-based systems necessitate matched training and testing sets, while CTC-based systems can cope with slight mismatches.

### 6.2.2. TL models on child speech

At first sight, we observe that TL models perform drastically better than adult-trained models, showing the positive effect of transfer learning even with a quantity of data as small as 13 hours. We observe a mean relative gain of 50.9% between scores in the first and second column. **TL models were found to perform better than models trained solely on child speech. The improvement is already significant for the TDNNF (7.3% relative), and is much more important for end-to-end architectures, since child models displayed scores above 50%.**

The attention-based systems (LAS and Transformers) benefit the most from the transfer learning, and the most spectacular improvement can be seen in the Transformer's performance, corresponding to a relative reduction of 62.8% of the PER. Models that do not rely on attention, i.e. TDNNF-HMM and RNN-CTC, benefit from TL to a lesser extent (*resp.* 34.4% and 45.5% relative improvements). The TL models obtain PER scores that are included in a relatively small range (28% - 34%), which is interesting considering the very wide range of results for models without TL (45% - 78%). This suggests that independently of the architecture, and due to the small amount of target data, the efficiency of transfer learning might reach a limit when confronted with arduous tasks like speech recognition for children learning to read.

The attention-based systems, whose attention mechanisms trained on adult sentences were not able to adapt to the diverse utterance lengths and mismatched prosodic characteristics of child speech, show to reach, thanks to the transfer learning, comparable performance with the non-attention-based systems. The child speech training data, containing equal proportions of isolated words and sentences, empowers the attention mechanisms to adjust to its specificities. In this way, the models based on the Transformer architecture outperform the TDNNF-HMM baseline with relative improvements up to 6.6%. The best architecture is

the Transformer+CTC with the decoder output, reaching a 28.1% PER. The systems that use RNNs are less efficient than TDNNF-HMM and Transformers, suggesting that they need more data to gain precision on extracting relevant information in child speech. These results concur with the global ranking between hybrid DNN-HMM, RNN-based seq2seq and Transformer architectures presented in Karita et al. (2019b).

If we focus on the influence of the CTC, we observe the same patterns for LAS and Transformers architectures, although to various extents. It shows highly useful for training the LAS+CTC, improving by a relative 9.4% the score of the LAS+CTC dec model in comparison with the LAS model, and brings a lower improvement between Transformer and Transformer+CTC dec models (only 1.4% relative).

## 7. Discussion

We presented different architectures for acoustic modelling and compared their performance when applied to the specific case of child speech with limited data. Although PER scores obtained in the previous section may still seem high, it is important to note that the child speech recordings can contain high levels of babble noise, typical in children environments. The mean SNRs of the Lalilo training, validation and testing sets are 20.9 dB, 20.1 dB and 20.6 dB, respectively (see Table 2), including recordings with SNR values down to 0 dB. These values are much lower than those of the Common Voice dataset (~34.3 dB, see Table 1), which has been gathered in relatively calm environments. While the best model obtains a 28.1% PER on all recordings, it performs significantly better on clean recordings (i.e. SNR  $\geq 20$  dB), reaching a 20.0% PER. The presence of babble noise on recordings highly impacts the performance: the best model obtains PER values of 33.0% and 44.8% on recordings with  $10 \text{ dB} \leq \text{SNR} \leq 20 \text{ dB}$  and recordings with  $\text{SNR} \leq 10 \text{ dB}$ . Babble noise has a high effect on the model's amount of insertions and substitutions, which are doubled, and of deletions, which are tripled. The robustness of the system to babble noise could be improved with data augmentation and multi-condition training techniques (Gibson et al., 2018; Airaksinen et al., 2019): it would most certainly improve the global performance by providing additional child speech data, as well as better training the model to recognise speech in noise.

In the scope of phone recognition on young readers' speech, we will in this section further explore the use of these models for evaluating specific reading tasks, i.e. isolated words and sentences. An analysis of the behaviour of the Transformer+CTC dec model in the presence of reading mistakes will also be carried, with a focus on the beneficial influence of the CE+CTC multi-objective training for detecting common sentence-level reading mistakes.

### 7.1. Application to specific reading tasks

Children learning to read are typically offered different reading tasks, based on their reading level or on the skill they

have to improve. We focus here on two reading tasks: isolated words and sentences. Knowing that acoustic model architectures might behave differently depending on the length of the utterance or the number of words, we take an interest in detailing each architecture's behaviour on these two tasks. Figures 4 and 5 display the models' performance on the child speech test sets, respectively Test S (sentences) and Test W (isolated words).

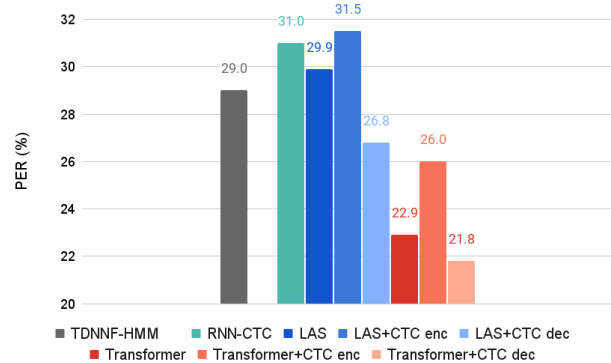


Figure 4: PER (%) of all TL models on Test S (sentences)

On sentences, Transformers obtain far better scores than the others, outperforming the TDNNF-HMM baseline by relative improvements from 10.3% to 24.8%. The Transformer+CTC dec obtains the lowest score of 21.8%. A clear tendency can be observed in Fig. 4 for models using multi-objective CE+CTC training. It shows to bring improvement: the LAS+CTC dec model shows a relative gain of 10.4% over the LAS model, the Transformer+CTC dec of 4.8% over the Transformer model. For both architectures, the CTC encoder output seems to bring significantly worse results than the decoder output that relies on attention. This observation, combined with the poor score of the RNN-CTC on sentences, shows that the CTC objective on its own has difficulties for recognising sentences, probably due to high complexity in the alignment of sentences uttered by reading learners, i.e. including reading miscues such as repetitions.

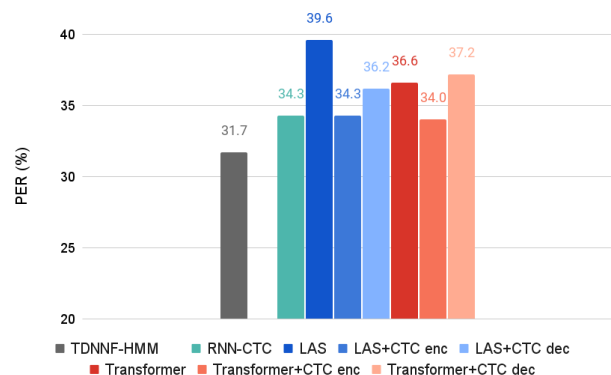


Figure 5: PER (%) of all TL models on Test W (isolated words)

The scores on isolated words (Test W) are significantly higher than on sentences, suggesting that the task is quite arduous. This difficulty is most probably brought by the short length of utterance, as well as the higher proportion of reading mistakes (in Test W, 51.5% of phones belong to an utterance that contains at least one reading mistake, versus 36.0% in Test S, see Table 2). On isolated words, the baseline hybrid TDNNF-HMM architecture is far from being outranked, reaching a PER of 31.7%, which lies 2.3% absolute below the second-best model. Very clear patterns can be observed in Figure 5, with well-defined groups of PER. The three models that use the CTC as output (LAS+CTC enc, RNN-CTC and Transformer+CTC enc) form a first group with PERs between 34.0% and 34.3%, which correspond to 7.3% to 8.2% relative degradation in comparison to the baseline. We can infer from this that the CTC takes over the attention mechanisms when the utterances are too short, and manages to align the phones properly with a very small context. We then see a second group, which includes the two models that are trained with the CTC objective and use the decoder output (Transformer+CTC dec and LAS+CTC dec), as well as the Transformer model. This floor corresponds to 14.2% to 17.4% relative augmentation over the baseline. Contrary to sentences, the CE+CTC multi-objective training has mitigated effect on LAS and Transformer models. While it is still highly efficient for LAS models (8.6% relative improvements between LAS and LAS+CTC dec models), it slightly degrades the performance of the Transformer (-1.6% relative). The worse result is obtained by the LAS model, with a 24.9% relative augmentation of the PER in comparison with the baseline. These four models ranking last on isolated words are all using an attention-based decoder to infer phone sequences. It shows that attention mechanisms have difficulties inferring labels from very short audio sequences. This result echoes the results of Chan et al. (2016), where their error rate significantly rises for 1-word utterances.

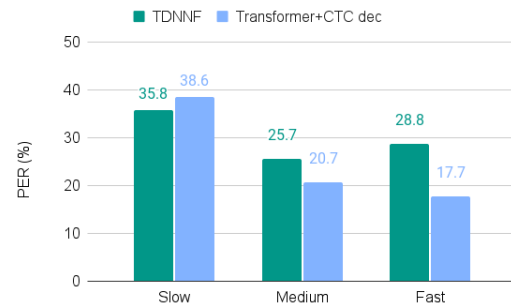
## 7.2. Application to non-proficient readers' speech

Aiming at phone recognition for young readers, our system will inevitably encounter **diverse reading rates** and many reading mistakes and must correctly recognise the uttered phones, independently from the number of reading mistakes and the child's reading level. However, **very slow reading rates** and the presence of reading mistakes can severely hurt the model's accuracy. In this section, we evaluate the Transformer+CTC dec performance when confronted with slow reading rates and reading mistakes.

### 7.2.1. Influence of the reading rate

Following teacher practices, we measure the children's reading rate with a Word Correct Per Minute (WCPM) metric. It is computed by taking the number of correctly read words in an utterance and divide it by the time taken to read that utterance. Computing the WCPM on short utterances, while it is usually computed on a whole minute of reading aloud, results in overestimating it, since the child does not have to breathe, or change line for example. To avoid too important biases, we compute it only on sentences. We define

three ranges based on national standards<sup>6</sup>: "slow" readings have a WCPM below 50 (level expected in first grade, age 6), and "fast" readings a WCPM above 90 (level expected in third grade, age 8), and "average" readings a WCPM in-between those values. The Test S has a mean WCPM of 96.7, with a standard deviation of 47.0.



**Figure 6:** PER (%) of the TDNNF and Transformer+CTC dec TL models on Test S as a function of the estimated reading speed

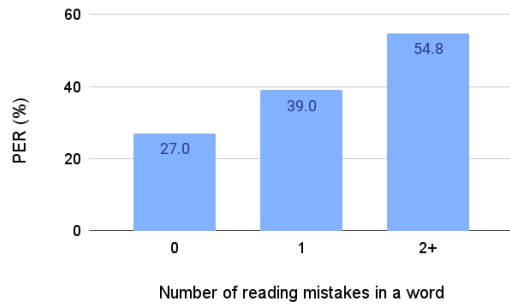
Figure 6 displays the PER obtained by the TDNNF and Transformer+CTC dec models on each WCPM range. We observe that the Transformer+CTC dec model struggles with slow readings, for which the PER is doubled in comparison with fast readings. This can be explained by the difficulties to recognise slow readers' speech, containing disfluencies (intra-word pauses, phone extensions...). It could be partly resolved by including more slow readers' speech in the training set, which currently has a mean WCPM of 114.3 ( $\pm 40.8$ ). The TDNNF seems to be affected by the slow reading rates to a lesser extent. However, important WCPM values degrade its performance: it could be due to the 3-frame sub-sampling of chain hybrid models (Povey et al., 2016), that could keep the TDNNF from addressing correctly very short phone durations. Interestingly, we observe the opposite for the Transformer+CTC dec, which suggests the high precision of attention mechanisms to extract information from fast read speech.

### 7.2.2. Influence of reading mistakes

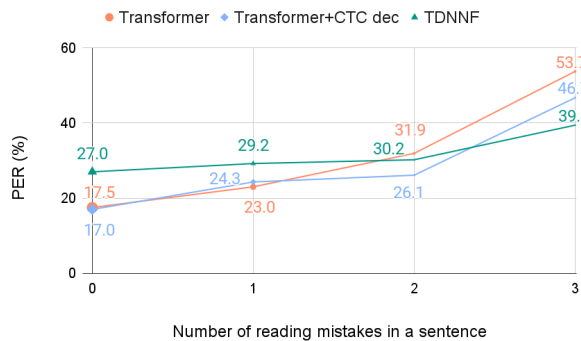
Figures 7 and 8 display the evolution of the PER scores obtained by the Transformer+CTC dec model in function of the number of reading mistakes contained in the utterances, for Test W and Test S, respectively. Figure 8 additionally includes the results of the TDNNF and Transformer models, which serve as a comparison to study the influence of the attention mechanisms and the CTC function.

In Test W, we count the number of mistakes at the word-level, in terms of phones: for example, a word with a substituted phone is counted as one mistake, a word with a hesitation and an inserted phone belongs in the 2-and-more mistakes category, and a false start or a repeated word where the child pronounces more than two phones is also classified as 2+ mistakes. As indicated in Table 2, 48.5% of phones ut-

<sup>6</sup><https://colibris.link/national-standards-WCPM>



**Figure 7:** PER (%) of the Transformer+CTC dec TL model on Test W as a function of the number of reading mistakes in the word



**Figure 8:** PER (%) of the TDNNF, Transformer and Transformer+CTC dec TL models on Test S as a function of the number of reading mistakes in the sentence.

tered in Test W belong in the no-mistake category, while 1 mistake and 2-and-more mistake categories include respectively 26.0% and 25.5% of the phones. In Test S, on the other hand, we count the number of mistakes at the sentence-level, in terms of words: each misread word counts as a mistake, a false start, a repeated word or a word containing a hesitation as well. Test S contains 64% of phones belonging to correctly read utterances (see Table 2), and other categories contain (in ascending order) 20%, 6.8%, and 9.2% of the phones uttered in Test S. In Figure 8, the size of the points is proportional to the number of reference phones in each category.

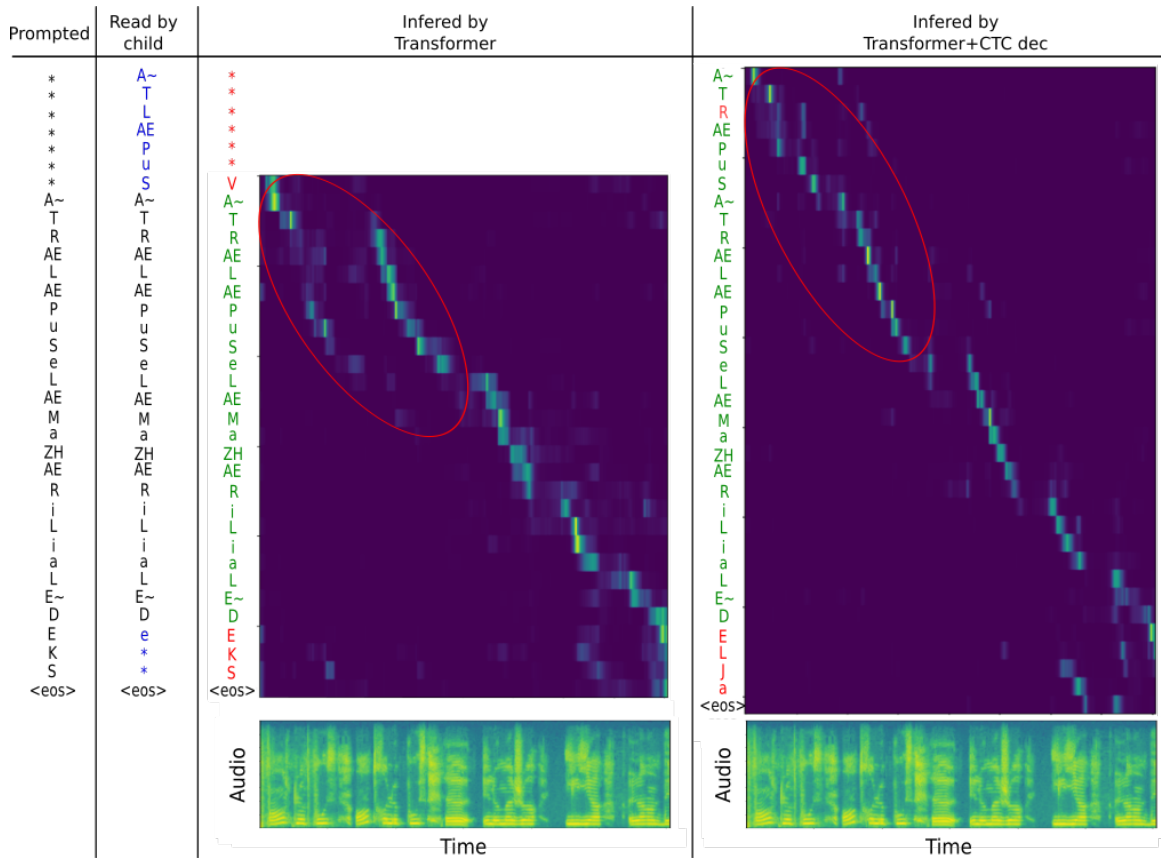
We can easily notice the tendency following which the more reading mistakes are present in the utterance, the more degraded is the PER, for both isolated words and sentences. On Test W, the presence of a single reading mistake augments the PER by 44.4% relative in comparison with correct readings. More mistakes occurring in a word degrade the PER to the same significant extent: a relative degradation of 40.5% is observed between 1 and 2+ mistakes. On Test S, the evolution of the PER for Transformer and Transformer+CTC dec models is also striking: it drastically increases when the utterance contains reading mistake(s). In comparison, the TDNNF's scores increase only slightly until two mistakes per utterance, then more importantly but still

to a lesser extent.

Analysing further the Transformer+CTC dec performance, we see that going from correct reading to reading with one mistake brings a 42.9% relative degradation. The score for utterances with two mistakes is only slightly higher than with a single mistake, while it highly augments for utterances with more than three reading mistakes (78.9% relative augmentation over category 2-mistakes). These drastic degradations can be explained by the characteristics of children's typical reading mistakes (detailed in section 3.3). Beginner readers tend to read with a slow speech rate and prosody: hesitations inter- or intra-word and phones of longer duration, for instance. These peculiarities can cause the attention mechanisms to lose focus and miss out on some phones. Additionally, substitutions of phones can sometimes give unusual combinations of phones that do not exist or are rarely present in the language, which the model is not trained to recognise. The decoder module in particular acts like a language model that favours known sequences of phones. Finally, repetitions are usual reading mistakes (in sentences in particular), as students often start back from the beginning of the sentence when they have difficulties, and can cause great confusion. The attention mechanisms risk in this case to attend the wrong occurrence of the word, or to consider two occurrences as a single one. This is problematic in the common case where the child reads a word a first time with a mistake, then a second time self-correcting, because the attention could confuse the two occurrences and miss the mistake.

In the end, while the TDNNF is significantly less efficient than the Transformer+CTC dec model on correctly read or slightly erroneous utterances (up to 2 mistakes), it handles better severe mistakes. This phenomenon can be attributed to the fact that it uses neither attention mechanisms nor a language model-like decoder module. The Transformer+CTC dec model shows to perform better on reading mistakes than the Transformer, suggesting a beneficial influence of the CTC multi-objective training.

Fig. 9 displays, for an utterance that well represents the type of reading mistakes a young reader can do, the phone-level prompted text and actual text read by the child, as well as the transcriptions obtained with Transformer and Transformer+CTC dec models. Additionally, it presents the weights extracted from the attention module that links the encoder and decoder of the two systems. The spectrogram of the utterance is also shown for time reference. By comparing the prompted text and the text read by the child, we can see that the child repeats the first three words (the repeated phones are in blue on Fig. 9) and makes a mistake when reading the first word ( $A \sim T$  instead of  $A \sim TRAE$ ). The objective of our acoustic models is to transcribe exactly what the child has read, and be particularly accurate on potential reading mistakes, to help the child learn. A first observation is that the Transformer does not detect the first attempt of the child, while the Transformer+CTC dec does, although with a substitution of a phone ( $R$  instead of  $L$ , in red on Fig. 9). The parallel diagonals that are pointed out by a red ellipse on the



**Figure 9:** Phone-level prompted text, actual text read by the child and transcriptions obtained with models Transformer and Transformer+CTC dec of an example utterance. Child’s reading mistakes (repetitions, phone substitutions and phone deletions) are in blue, and correct detections and errors in models’ transcriptions in green and red, respectively. Attention weights captured between the encoder and decoder of Transformer and Transformer+CTC dec models are displayed, as well as the spectrograms for time reference.

Transformer attention weights show that the model confuses the two occurrences of the repeated words into one. On the contrary, the Transformer+CTC dec attention weights form a single diagonal, correctly detecting the two occurrences as separate. This success may be explained by the influence of the CTC objective function, which constrains the attention to be monotonic, enabling the model to correctly detect repetitions.

Another observation is that the child makes several word-level reading mistakes in the last word of the sentence: a phone substitution and two phone deletions ( $E \sim D e$  instead of  $E \sim D E K S$ ). The word read by the child does not exist in the French language and is not represented in the training data. The Transformer unexpectedly detects the correct reading, which means it misses the child’s mistake. This is due to the presence of this same content in the child speech training data, but read correctly by other students, which has caused the Transformer to predict the learnt existing word. On the other hand, the Transformer+CTC dec model does detect a reading mistake in the last word, forming a non-existing word, although not the one uttered by the child. We can infer that since it is able to form non-existing words, it might show more robustness to children reading mistakes,

as shown in Fig. 8.

The issues that attention mechanisms face when confronted with repetition or non-existing words can cause problematic errors when transcribing young readers’ mistakes. The child training data, containing only correct readings, prevents the models to be trained to handle readings with mistakes. A perspective for improving our model’s performance in the presence of reading mistakes is to add some representations of these mistakes in the child training data. It would enable the model to better detect incongruous substitutions and insertions of phones, as well as repetitions, false starts and hesitations. However, this method would require manual annotation at the phone level of a significant quantity of data, which is an arduous and costly task. Synthetic errors could partly remedy the issue: although word-level mistakes might be thorny to create, utterances with sentence-level mistakes such as repetitions, false starts and hesitations could be automatically created and added to the training data. Our Transformer models could also gain robustness on child speech peculiarities during inference by using scheduled sampling methods, adapted to the specific Transformer structure, during training (Mihaylova and Martins, 2019; Zhou et al., 2019).

## 8. Conclusion

Speech technologies are nowadays widening their usage domains, including numerous applications designed for children. However, the speech recognition systems' performance is significantly lower on child speech than on adult speech. In particular, educational numerical resources for children learning to read could immensely benefit from accurate speech recognition to detect reading mistakes.

End-to-end architectures have proved their ability to outperform hybrid DNN-HMM approaches for ASR. In this work, we apply end-to-end architectures to child speech with a limited amount of child data and show that, with the help of transfer learning strategies, a Transformer+CTC can reach a 28.1% PER, and outmatch a TDNNF-HMM model by 6.6% relative, as well as other end-to-end architectures (RNN, LAS) by 8.5% to 17.1% relative. We specifically study CE+CTC multi-objective training on diverse end-to-end architectures, which shows to bring significant improvements. We find degradation of performance for attention-based models when the utterance length is reduced, while TDNNF-HMM and CTC-based models seem to better handle very short utterances. **Finally, we compare the performance of our best models on diverse reading rates and in presence of reading mistakes. We show that CE+CTC multi-objective training indeed constrains the attention to be more monotonic, which enables the Transformer+CTC model to better detect common young readers' mistakes. The TDNNF, however, stays slightly ahead for severely low-level readers' speech.** Our detailed studies provide valuable insights on the end-to-end architectures remaining weaknesses for transcribing young readers' speech. Using different speech units, such as syllables or characters, on their own or in combination with phones, could bring complementary information on reading learners' speech. Having seen that the small amount of child data could limit the potential of transfer learning, diverse data augmentation methods could be investigated. The problem of noise could in particular be tackled with data augmented with child-adapted babble noise. Creating synthetic child training data with common reading mistakes, such as repetitions, false starts, or hesitations would enable the models to better handle children mistakes.

## References

Abad, A., Bell, P., Carmantini, A., Renals, S., 2020. Cross lingual transfer learning for zero-resource domain adaptation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6909–6913. doi:10.1109/ICASSP40776.2020.9054468.

Airaksinen, M., Juvela, L., Alku, P., Räsänen, O., 2019. Data augmentation strategies for neural network F0 estimation, in: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6485–6489.

Andrew, Sak, H., de Chaumont Quitry, F., Sainath, T., Rao, K., 2015. Acoustic modelling with CD-CTC-SMBR LSTM RNNS, in: *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, pp. 604–609. doi:10.1109/ASRU.2015.7404851.

Bahdanau, D., Cho, K., Bengio, Y., 2015. Neural machine translation by joint learning to align and translate. *CoRR abs/1409.0473*.

Bayerl, S.P., Riedhammer, K., 2019. A comparison of hybrid and end-to-end models for syllable recognition, in: *Text, Speech, and Dialogue*, pp. 352–360.

Bengio, S., Vinyals, O., Jaitly, N., Shazeer, N., 2015. Scheduled sampling for sequence prediction with recurrent neural networks, in: *Proc. of the International Conference on Neural Information Processing Systems (NIPS) - Volume 1*, MIT Press. p. 1171–1179. doi:10.5555/2969239.2969370.

Bolaños, D., Cole, R., Ward, W., Borts, E., Svirsky, E., 2011. FLORA: Fluent oral reading assessment of children's speech. *ACM Trans. Speech Lang. Process.* 7, 16. doi:10.1145/1998384.1998390.

Chan, W., Jaitly, N., Le, Q., Vinyals, O., 2016. Listen, Attend and Spell: A neural network for large vocabulary conversational speech recognition, in: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4960–4964. doi:10.1109/ICASSP.2016.7472621.

Chen, G., Na, X., Wang, Y., Yan, Z., Zhang, J., Ma, S., Wang, Y., 2020. Data augmentation for children's speech recognition – the "Ethiopian" system for the SLT 2021 Children Speech Recognition Challenge. *arXiv:2011.04547*. working paper.

Chiu, C.C., Sainath, T., Wu, Y., Prabhavalkar, R., Nguyen, P., Chen, Z., Kannan, A., Weiss, R.J., Rao, K., Gonina, K., Jaitly, N., Li, B., Chorowski, J., Bacchiani, M., 2018. State-of-the-art speech recognition with sequence-to-sequence models, in: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4774–4778. doi:10.1109/ICASSP.2018.8462105.

Cho, J., Baskar, M.K., Li, R., Wiesner, M., Mallidi, S.H., Yalta, N., Karafiát, M., Watanabe, S., Hori, T., 2018. Multilingual sequence-to-sequence speech recognition: Architecture, transfer learning, and language modeling, in: *IEEE Spoken Language Technology Workshop (SLT)*, pp. 521–527. doi:10.1109/SLT.2018.8639655.

Chorowski, J., Bahdanau, D., Cho, K., Bengio, Y., 2014. End-to-end continuous speech recognition using attention-based recurrent NN: First results, in: *Proc. of the International Conference on Neural Information Processing Systems (NIPS): Workshop on Deep Learning*, pp. 1–10.

Chorowski, J., Bahdanau, D., Serdyuk, D., Cho, K., Bengio, Y., 2015. Attention-based models for speech recognition, in: *Proc. of the International Conference on Neural Information Processing Systems (NIPS)*, MIT Press. p. 577–585. doi:10.5555/2969239.2969304.

Chung, J., Gulcehre, C., Cho, K., Bengio, Y., 2015. Gated feedback recurrent neural networks, in: *Proc. of the International Conference on Machine Learning (ICML)*, p. 2067–2075.

Dong, L., Xu, S., Xu, B., 2018. Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition, in: *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5884–5888. doi:10.1109/ICASSP.2018.8462506.

Duan, R., Kawahara, T., Dantsuji, M., Nanjo, H., 2020. Cross-lingual transfer learning of non-native acoustic modeling for pronunciation error detection and diagnosis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28, 391–401. doi:10.1109/TASLP.2019.2955858.

Fringi, E., Lehman, J.F., Russell, M.J., 2015. Evidence of phonological processes in automatic recognition of children's speech, in: *Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH)*, Dresden, pp. 1621–1624.

Gales, M., Young, S., 2008. The application of hidden markov models in speech recognition. *Foundations and Trends® in Signal Processing* 1, 195–304. URL: <http://dx.doi.org/10.1561/2000000004>, doi:10.1561/2000000004.

Gerosa, M., Giuliani, D., Narayanan, S., 2006. Acoustic analysis and automatic recognition of spontaneous children's speech, in: *Ninth International Conference on Spoken Language Processing*, pp. 1886–1889.

Gibson, M., Plahl, C., Zhan, P., Cook, G., 2018. Multi-condition deep neural network training. *Studenttexte zur Sprachkommunikation: Elektronische Sprachsignalverarbeitung*, 77–84.

Godde, E., Bailly, G., Escudero, D., Bosse, M.L., Estelle, G., 2017. Evaluation of reading performance of primary school children: Objective measurements vs. subjective ratings, in: *Proc. of the International Workshop on Child Computer Interaction (WOCCI)*, pp. 23–27. doi:10.21437/WOCCI.2017-4.



- Graves, A., Fernández, S., Gomez, F., Schmidhuber, J., 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks, in: Proc. of the International Conference on Machine Learning (ICML), pp. 369–376. doi:10.1145/1143844.1143891.
- Graves, A., Mohamed, A., Hinton, G., 2013. Speech recognition with deep recurrent neural networks, in: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6645–6649. doi:10.1109/ICASSP.2013.6638947.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778. doi:10.1109/CVPR.2016.90.
- Karita, S., Soplín, N.E.Y., Watanabe, S., Delcroix, M., Ogawa, A., Nakatani, T., 2019a. Improving Transformer-Based End-to-End Speech Recognition with Connectionist Temporal Classification and Language Model Integration, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Graz, pp. 1408–1412. doi:10.21437/Interspeech.2019-1938.
- Karita, S., Wang, X., Watanabe, S., Yoshimura, T., Zhang, W., Chen, N., Hayashi, T., Hori, T., Inaguma, H., Jiang, Z., Someki, M., Soplín, N.E.Y., Yamamoto, R., 2019b. A Comparative Study on Transformer vs RNN in Speech Applications. IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU), 449–456. doi:10.1109/ASRU46091.2019.9003750.
- Lee, S., Potamianos, A., Narayanan, S.S.Y., 1999. Acoustics of children's speech: developmental changes of temporal and spectral parameters. The Journal of the Acoustical Society of America 105, 1455–1468. doi:10.1121/1.426686.
- Lu, L., Zhang, X., Cho, K., Renals, S., 2015. A study of the recurrent neural network encoder-decoder for large vocabulary speech recognition, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Dresden, pp. 3249–3253.
- Metallinou, A., Cheng, J., 2014. Using deep neural networks to improve proficiency assessment for children english language learners, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Singapore, pp. 1468–1472.
- Mihaylova, T., Martins, A.F.T., 2019. Scheduled sampling for transformers, in: Proc. of the Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, Association for Computational Linguistics, pp. 351–356. doi:10.18653/v1/P19-2049.
- Mostow, J., Aist, G., 2001. Evaluating tutors that listen: An overview of Project LISTEN, in: Smart machines in education: The coming revolution in educational technology. The MIT Press, pp. 169–234. doi:10.5555/570950.570957.
- Mugitani, R., Hiroya, S., 2012. Development of vocal tract and acoustic features in children. The Journal of the Acoustical Society of Japan 68, 234–240. doi:10.1250/ast.33.215.
- Ng, S.I., Liu, W., Peng, Z., Feng, S., Huang, H.P., Scharenborg, O., Lee, T., 2020. The CUHK-TUDELFT system for the SLT 2021 Children Speech Recognition Challenge. arXiv:2011.06239. working paper.
- Potamianos, A., Narayanan, S., 1998. Spoken dialog systems for children, in: IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 197–200 vol.1. doi:10.1109/ICASSP.1998.674401.
- Potamianos, A., Narayanan, S., 2003. Robust Recognition of Children's Speech. IEEE Transactions on Speech and Audio Processing 11, 603–616. doi:10.1109/TSA.2003.818026.
- Potamianos, A., Narayanan, S., 2007. A review of the acoustic and linguistic properties of children's speech, in: IEEE 9th Workshop on Multimedia Signal Processing, pp. 22–25. doi:10.1109/MMSP.2007.4412809.
- Povey, D., Cheng, G., Wang, Y., Li, K., Xu, H., Yarmohammadi, M., Khudanpur, S., 2018. Semi-orthogonal low-rank matrix factorization for deep neural networks, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Hyderabad, pp. 3743–3747. doi:10.21437/Interspeech.2018-1417.
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., Hannemann, M., Motlicek, P., Qian, Y., Schwarz, P., Silovsky, J., Stemmer, G., Vesely, K., 2011. The Kaldi speech recognition toolkit, in: IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU), pp. 1–4.
- Povey, D., Peddinti, V., Galvez, D., Ghahremani, P., Manohar, V., Na, X., Wang, Y., Khudanpur, S., 2016. Purely sequence-trained neural networks for ASR based on lattice-free MMI, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), San Francisco, pp. 2751–2755. doi:10.21437/Interspeech.2016-595.
- Proença, J.D.L., 2018. Automatic Assessment of Reading Ability of Children. Ph.D. thesis. University of Coimbra.
- Qian, Y., Wang, X., Evanini, K., Suendermann-Oeft, D., 2016. Improving DNN-based automatic recognition of non-native children speech with adult speech, in: Workshop on Child Computer Interaction, pp. 40–44. doi:10.21437/WOCCI.2016-7.
- Serizel, R., Giuliani, D., 2014. Deep neural network adaptation for children's and adults' speech recognition, in: Proc. of the Italian Computational Linguistics Conference (CLiC-it), pp. 137–140.
- Shivakumar, P.G., Georgiou, P., 2020. Transfer learning from adult to children for speech recognition: Evaluation, analysis and recommendations. Computer Speech & Language 63, 101077. doi:10.1016/j.cs1.2020.101077.
- Shivakumar, P.G., Narayanan, S., 2021. End-to-end neural systems for automatic children speech recognition: An empirical study. arXiv:2102.09918.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks, in: Proc. of the International Conference on Neural Information Processing Systems (NIPS), Cambridge, MA, USA. p. 3104–3112.
- Tong, R., Wang, L., Ma, B., 2017a. Transfer learning for children's speech recognition. Proc. of the International Conference on Asian Language Processing (IALP), 36–39. doi:10.1109/IALP.2017.8300540.
- Tong, S., Garner, P., Boulard, H., 2017b. Multilingual training and cross-lingual adaptation on CTC-based acoustic model. Speech Communication 104. doi:10.1016/j.specom.2018.09.001.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, u., Polosukhin, I., 2017. Attention is all you need, in: Proc. of the International Conference on Neural Information Processing Systems (NIPS). Curran Associates Inc., Red Hook, NY, USA. p. 6000–6010.
- Vesely, K., Ghoshal, A., Burget, L., Povey, D., 2013. Sequence-discriminative training of deep neural networks, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Lyon, pp. 2345–2349.
- Vinyals, O., Le, Q., 2015. A neural conversational model. Proc. of the International Conference on Machine Learning (ICML): Deep Learning Workshop.
- Vinyals, O., Toshev, A., Bengio, S., Erhan, D., 2015. Show and tell: A neural image caption generator, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3156–3164. doi:10.1109/CVPR.2015.7298935.
- Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., Lang, K.J., 1989. Phoneme recognition using time-delay neural networks. IEEE Transactions on Acoustics, Speech, and Signal Processing 37, 328–339. doi:10.1109/29.21701.
- Watanabe, S., Hori, T., Kim, S., Hershey, J.R., Hayashi, T., 2017. Hybrid CTC/Attention architecture for end-to-end speech recognition. IEEE Journal of Selected Topics in Signal Processing 11, 1240–1253. doi:10.1109/JSTSP.2017.2763455.
- Wu, F., Garcia, P., Povey, D., Khudanpur, S., 2019. Advances in automatic speech recognition for child speech using factored time delay neural network, in: Proc. of the Annual Conference of the International Speech Communication Association (INTERSPEECH), Graz, pp. 1–5. doi:10.21437/Interspeech.2019-2980.
- Xu, K., Ba, J.L., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R.S., Bengio, Y., 2015. Show, Attend and Tell: Neural image caption generation with visual attention, in: Proc. of the International Conference on Machine Learning (ICML), JMLR.org. p. 2048–2057. doi:10.5555/3045118.3045336.
- Yeung, G., Alwan, A., 2018. On the difficulties of automatic speech recognition for kindergarten-aged children, in: Proc. of the Annual Conference of the International Speech Communication Association (IN-

- TERSPEECH), Hyderabad, pp. 1661–1665. doi:10.21437/Interspeech.2018-2297.
- Yong, B.F., Ting, H.N., 2011. Speaker-independent vowel recognition for malay children using time-delay neural network, in: 5th Kuala Lumpur International Conference on Biomedical Engineering 2011, Springer Berlin Heidelberg, pp. 565–568.
- Yu, F., Yao, Z., Wang, X., An, K., Xie, L., Ou, Z., Liu, B., Li, X., Miao, G., 2020. The SLT 2021 children speech recognition challenge: Open datasets, rules and baselines. *arXiv:2011.06724*. working paper.
- Zhou, P., Fan, R., Chen, W., Jia, J., 2019. Improving generalization of transformer for speech recognition with parallel schedule sampling and relative positional embedding. *CoRR abs/1911.00203*. *arXiv:1911.00203*.