

# LEARNING TO COUNT WORDS IN FLUENT SPEECH ENABLES ONLINE SPEECH RECOGNITION

George Sterpu, Christian Saam, Naomi Harte

SigmaMedia Lab, ADAPT Centre, Trinity College Dublin, Ireland

## ABSTRACT

Sequence to Sequence models, in particular the Transformer, achieve state of the art results in Automatic Speech Recognition. Practical usage is however limited to cases where full utterance latency is acceptable. In this work we introduce Taris, a Transformer-based online speech recognition system aided by an auxiliary task of incremental word counting. We use the cumulative word sum to dynamically segment speech and enable its eager decoding into words. Experiments performed on the LRS2, LibriSpeech, and Aishell-1 datasets of English and Mandarin speech show that the online system performs comparable with the offline one when having a dynamic algorithmic delay of 5 segments. Furthermore, we show that the estimated segment length distribution resembles the word length distribution obtained with forced alignment, although our system does not require an exact segment-to-word equivalence. Taris introduces a negligible overhead compared to a standard Transformer, while the local relationship modelling between inputs and outputs grants invariance to sequence length by design.

*Index Terms*— Online ASR, word segmentation

## 1. INTRODUCTION

Having a natural conversation with a computer has fascinated humankind for a long time. A key ingredient of this ambition is granting computers the ability to recognise spoken words with minimum latency. This allows a more interactive communication, where the computer is able to interrupt a speaker to acknowledge or ask for clarifications.

Despite the remarkable progress in end-to-end automatic speech recognition technology based on sequence to sequence neural network architectures [1], an unresolved issue is reducing the latency from full utterances down to a few words. This sentence-level, or offline conditioning, is a fundamental barrier to online decoding.

Humans develop the ability to segment words in continuous speech from the earliest stages of life [2]. There is evidence that we integrate a set of acoustic, phonetic, prosodic, and statistical cues in order to segment words in fluent speech [3]. This leads us to ask whether the ability to segment speech into *word* units with a neural network offers

the potential to help crack the challenge of decoding online. This approach would take advantage of the monotonicity of speech, allow the network focus on local properties, and remove the offline conditioning.

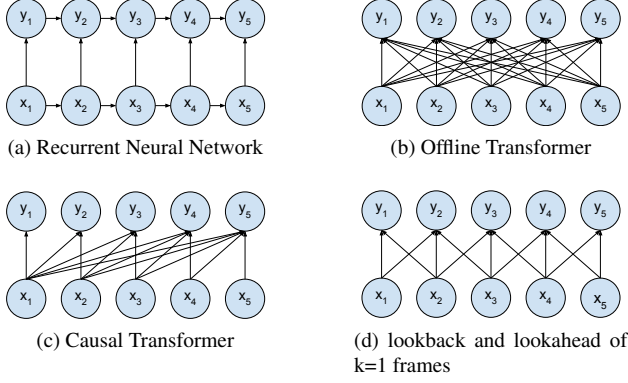
To this end, we introduce Taris, a Transformer-based system for online speech recognition that learns to model the local relationships between text and audio in speech, relaxing the global conditioning constraint of the original model. We achieve this through self-supervision by introducing an auxiliary word counting task which facilitates the segmentation of speech. Taris allows efficient minibatch training and introduces a negligible overhead compared to the original Transformer model, without trading off the recognition accuracy. We make our software implementation publicly available<sup>1</sup>.

## 2. BACKGROUND

A major technical challenge in online speech decoding is formulating the problem in a fully differentiable framework. Previous attempts include the Recurrent Neural Network Transducer [4, 5, 6, 7], Neural Transducer [8, 9], segmental conditional random fields [10, 11], hard monotonic attention [12, 13], segment attention [14, 15, 16], or triggered attention [17]. However the models made use of dynamic programming, training in expectation, or policy gradients, and the authors report training difficulties. Our work retains the segment attention design, but tackles the problem of speech segmentation from a different angle. By learning to count words through self-supervision, we introduce a mechanism that allows end-to-end training using only backpropagation.

Recent proposals in online speech recognition address this challenge by assuming one sub-word unit per segment [18, 19], or discover an inventory of sub-word units [20], a concept previously explored in machine translation [21]. Our focus in this work is on word units. In English, words allow a monotonic and bijective mapping between their acoustic and symbolic representations, however these properties do not hold at the sub-word level due to the highly complex spelling rules in English orthography. Moreover, words can be counted in a deterministic way, which allows us to introduce a self-supervision word counting task without requiring new labels.

<sup>1</sup><https://github.com/georgesterpu/Taris>



**Fig. 1.** Typical connectivity patterns at the sequence level for representation learning

The sequence to sequence (seq2seq) architecture was proposed in [22, 23]. An Encoder transforms a variable length input sequence into a sequence of latent representations, and a Decoder maps the latent sequence onto a target sequence of a different length, aiming to establish a soft-alignment between elements of the inputs and the targets. In attention-based seq2seq networks, the conditional dependency of each output token on the entire input sequence prohibits online decoding. Yet, it has been shown that, once convergence is reached, there are predominantly local relationships between the output tokens and the audio representations in speech [24, 25]. Therefore, potentially incurring no loss in accuracy, a local conditioning of the outputs on the inputs would break the offline limitation and reduce the algorithmic latency. The new goal is to learn robust associations between input and output subsequences which stand for the same linguistic concepts.

The Transformer [26] is a good seq2seq candidate for this task and we choose it as a foundation for our system Taris. Unlike the recurrent neural network that uses causal connections between timesteps (Figure 1a), the Transformer allows feature contextualisation at the sequence level through self-attention, illustrated in Figure 1b. This offline modelling strategy provides a theoretical upper limit of the segmentation performance. Furthermore, the self-attention connections in the Transformer block can be adjusted to allow causal modelling (Figure 1c) or non-causal modelling with a window (Figure 1d). The window length is directly linked to the algorithmic latency of Taris and its accuracy, and we investigate this trade-off in Section 5.3.

### 3. MODEL ARCHITECTURE

#### 3.1. Encoding

Taris takes as input a variable length sequence of audio vectors  $\mathbf{a} = \{a_1, a_2, \dots, a_N\}$  and applies the Encoder stack of the Transformer model defined in [26]. Because of latency considerations, instead of the original full connectivity in Fig-

ure 1b, we use the type displayed in Figure 1d, with controlled look-back  $e_{LB}$  and look-ahead  $e_{LA}$  frames. We denote the outputs of the encoder as:

$$\mathbf{o}_A = \text{Encode}(\mathbf{a}, e_{LB}, e_{LA}) \quad (1)$$

Next, we apply a sigmoidal gating unit on each encoder output  $o_{A_i}$  to obtain a scalar score for each frame:

$$\alpha_i = \text{sigmoid}(o_{A_i} W_G + b_G) \quad (2)$$

$$\text{where } \text{sigmoid}(x) = \frac{1}{1 + \exp(-x)}, W_G \in \mathbb{R}^{h \times 1}, b_G \in \mathbb{R}^1$$

We assign to every single input frame  $i$  a segment index  $\hat{w}_i$  by taking the *cumulative sum* of  $\alpha$  and applying the *floor* function on the output:

$$\hat{w}_i = \left\lfloor \sum_{j=1}^i \alpha_j \right\rfloor \quad (3)$$

Namely, the first predicted segment is delimited by a cumulative sum of  $\alpha$  between 0 and 1, the second segment by the same quantity between 1 and 2, and so on.

#### 3.2. Decoding

During training, the Decoder stack receives the labelled grapheme sequence  $\mathbf{y} = \{y_1, y_2, \dots, y_L\}$ , made of English letters and the unique word delimiter SPACE. We assign to every grapheme  $k$  a word index  $w_k$  by leveraging the SPACE tokens in the labelled sequence:

$$w_k = \sum_{j=1}^k (y_j == \text{SPACE}) \quad (4)$$

Thus, whereas symbolic segmentation of speech uses a unique SPACE token to separate words, acoustic segmentation flags word boundaries by tracking the frame locations where the partial sum of the word counting signal  $\alpha_i$  passes to the next integer value.

We modify the decoder-encoder connectivity of the Attention layer of [26] to allow our decoder to perform soft-alignment over a *dynamic* window of segments estimated by the encoder. More precisely, we only allow those connections for which the following condition is met:

$$\mathbf{V} = \widehat{W}_{ik} \leq (W_{ik} + d_{LA}) \text{ and } \widehat{W}_{ik} \geq (W_{ik} - d_{LB}) \quad (5)$$

In (5),  $d_{LA}$  and  $d_{LB}$  denote the number of segments the decoder is allowed to look-ahead and look-back respectively. The  $W$  and  $\widehat{W}$  matrices are obtained from the  $w$  and  $\hat{w}$  arrays by applying the tile operation, which repeats one sequence for a number of times equal to the length of the other one. In more detail,  $\mathbf{V}$  is a 2D matrix  $\in \mathbb{R}^{N \times L}$  that defines the admissible connections between any decoder timestep and any

encoder timestep, acting as a bias on the decoder-encoder attention. Setting  $\mathbf{V}$  as a matrix of ones recovers the original Transformer model. The extension to 3D tensors that include the batch dimension is straightforward, offering Taris efficient minibatch training and inference.

The decoder implements a traditional character level auto-regressive language model that predicts the next grapheme in the sequence conditioned on all the previous characters and the dynamic audio context vector  $c_k$ :

$$c_k = \text{Attention}(\text{keys} = o_A, \text{query} = o_{D_{k-1}}, \text{mask} = V) \quad (6)$$

$$o_{D_k} = \text{Decode}(y, c_k) \quad (7)$$

$$p_k \equiv P(y_k | c_k, y_{1:k-1}) = \text{softmax}(o_{D_k} W_v + b_v) \quad (8)$$

where  $W_v \in \mathbb{R}^{h \times v}$ ,  $b_v \in \mathbb{R}^v$

In (8),  $v$  is the vocabulary size of 28 tokens. We measure the difference between the estimated word sum  $\Sigma \hat{w} = \sum_i \alpha_i$  and the true word count  $|w| = \sum_k (y_k == \text{SPACE})$  as:

$$\text{Word Loss} = (|w| - \Sigma \hat{w})^2 \quad (9)$$

We define the training loss as:

$$\text{CE Loss} = \frac{1}{L} \sum_k -y_k \log(p_k) \quad (10)$$

$$\text{Loss} = \text{CE Loss} + \lambda \text{Word Loss} \quad (11)$$


In our experiments we used a scale factor  $\lambda = 0.01$  found empirically. The self attention connections of the auto-regressive Decoder are causal as depicted in Figure 1c.

Taris requires a negligible overhead in parameters (given by the  $W_G$  and  $b_G$  variables in (2)) and operations (equations (2)-(5)) over the original Transformer.

#### 4. WHY LEARN TO COUNT WORDS

Proper lexical segmentation of speech depends on context and semantics, as commonly illustrated by the example *how to wreck a nice beach* sounding similar to *how to recognise speech*. Thus, strategies incrementally scanning for hard boundaries [18, 16, 19] are less suited to word units, prompting [19] to perform beam search on the entire sequence of sub-word tokens estimated from each segment. Instead, Taris has to develop intrinsic word counting mechanisms. One plausible strategy is to incrementally gather lexical evidence at the sub-word level, and learn to represent boundary-informative acoustic cues on a manifold where they can be accumulated.

We conjecture that learning the ability to count words facilitates the segmentation of speech into words, and we discuss below our intuition behind it. In Figure 2 we illustrate the word counting sub-problem to be solved by the network. Starting in the bottom left corner, the network predicts scores



**Fig. 2.** Illustration of the word counting problem. Given a spoken phrase *something like a garden compost*, the network has to reach the correct word count 5 taking any path starting in the bottom left corner. The blue path allows an easy segmentation of speech into words. The green and dark orange paths are possible, but do not facilitate segmentation.

for every audio frame in the sentence, and the cumulative sum is promoted get as close as possible to the total word count, shown with a red circle. There is a very large number of paths that can be taken to reach the target count. However, when trained on large amounts of naturally distributed speech, we predict that Taris converges towards genuine word segmentation by having the cumulative sum cross all the intermediate word boundaries shown with yellow circles. In other words, the network may learn to self-normalise the accumulated probabilities for each word regardless of their length or cued structure.

We believe it suffices to train a system with the right amount of speech data, with the following intuition. As words appear in multiple contexts throughout a dataset, learning to count words may then have a normalisation effect on the fraction of  $\Sigma \hat{w}$  allocated to each word in a sentence. Each word unit will approach a unitary mass allocation as its acoustic realisation is seen more often in multiple contexts. For the less frequent words, the correct allocation may happen by marginalisation if the sentences they appear in contain relatively more frequent words. Loosely speaking, it is the task of solving a system of linear equations where the variables are the partial sums corresponding to the acoustic frames between two consecutive estimated boundaries.

Since we do not explicitly model the pauses between words, and the convergence towards the segmental behaviour is a mathematical conjecture without analytic proof for now, it is likely to observe deviations in practice on learnt solutions. However, Taris does not require a very strict approximation of word boundaries to function correctly. Instead, it is sufficient to just avoid frequent under- and over-segmentation, as it directly impacts the model's latency.

## 5. EXPERIMENTS AND RESULTS

We first conduct our experiments on the audio part of the unconstrained speech dataset LRS2 [27] for rapid prototyping, and on the 100h partition of LibriSpeech [28] for empirical validation at a larger scale. To extract audio features  $\mathbf{a}$  in Equation 1, we apply the log scale Short-time Fourier Transform on the waveform inputs, following same procedure as in [29] for noise corruption at 10, 0, and -5 db.

Our implementation of Taris is forked from the official Transformer model in TensorFlow 2 [30]. We train our LibriSpeech models for a total of 500 epochs at an initial learning rate of 0.001, decayed to 0.0001 after 400 epochs. The training time is approximately 200 seconds for a single epoch of LibriSpeech 100h on an Nvidia Titan XP GPU. The LRS2 models were trained with the same learning rates for 100 and 20 epochs respectively, on each noise level.

### 5.1. Neural network details

Our models use 6 layers in the Encoder and Decoder stacks, a model size  $d_{model} \equiv h = 256$ , a filter size  $d_{ff} = 256$ , one attention head, and 0.1 dropout on all attention weights and feedforward activations. The models occupy 25 MB on disk.

### 5.2. The End-of-sentence (EOS) token

During our initial experiments, we noticed that traditional evaluation and training strategies in neural speech recognition are commonly misusing the EOS token, making it difficult to evaluate online systems. The commonly used ASR datasets are a collection of variable-length utterances, and the system’s accuracy is computed for each utterance using an edit distance based algorithm. These utterances are often fragments from full spoken sentences, such as the one illustrated in Figure 2, that were cropped using voice activity detection algorithms (e.g. in LRS2), and sometimes the fragmentation includes the ending and the start of two consecutive sentences, with the punctuation removed from the ground truth transcription (e.g. in LibriSpeech). In other words, the ASR system does not receive full sentence units, and cannot develop the linguistic notion of an end of sentence. In our experiments it became obvious that one way the ASR model differentiates between an EOS token and a word delimiter (SPACE) likely comes from the apriori knowledge of the sentence length, and that EOS becomes more likely as the decoder-encoder alignment distribution advances towards the last remaining audio frames in the sentence. The aspect above becomes problematic in an online setting, as the decoder is fed with a limited acoustic context. Given the nature of the dataset utterances, an online decoder does not know when to stop the decoding process, as EOS cannot be estimated even spuriously anymore. Online decoding would often stop after just a few words in an utterance, biasing the accuracy on longer sentences.

To circumvent this problem, we made two important changes to the traditional model. First, we replaced the EOS token in the labels, which cannot be predicted reliably, with the SPACE token. Second, we modified the stopping condition of the beam search inference decoder as follows: instead of stopping when all beams reach the EOS token, it now stops when the decoder predicts as many words as there were estimated by the audio encoder. This new strategy is mostly beneficial to the evaluation procedure, but should also be useful in practice as it allows the decoder to emit a controllable number of words. With this change, we are able to evaluate the error rate of Taris on full test sentences for which we lack any label alignments.

### 5.3. Learning to count words

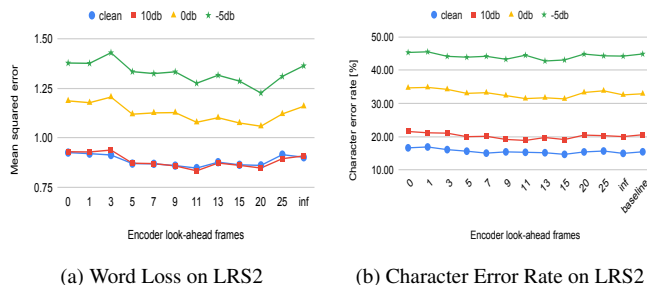


Fig. 3. Offline system evaluation for an increasing length of feature contextualisation in the encoder


We first investigate to what extent a sequence to sequence Transformer model can learn to count the number of words from audio data on LRS2. We train multiple models and gradually increase the number of encoder look-ahead frames  $e_{LA}$  to measure the variation of the *Word Loss* as more future context becomes available. We see in Figure 3a that the mean squared word count error is sub-unitary in clean speech and 10db noise, i.e. the estimated count is less than one word away from truth. This suggests that words can be counted relatively well from acoustic speech. In addition, using a future context length of 11 frames offers the lowest counting error under all noise conditions. In Figure 3b we plot the mean Character Error Rate achieved by all our systems, including the offline Transformer baseline without the auxiliary Word Loss, and we observe no significant difference, with the 95% confidence intervals of the mean errors between 1% and 1.4%. Therefore, this auxiliary task is not detrimental to the original accuracy obtained on LRS2 using only the cross-entropy loss.

### 5.4. Online ASR decoding

The decoder in our previous experiment had access to the entire encoder memory. For our online model we opt for an encoder lookahead  $e_{LA}$  of 11 frames and infinite lookback  $e_{LB} = \infty$ , as we showed in Section 5.3 that there are dimin-

ishing gains beyond this threshold. This roughly corresponds to an encoding latency of 330 msec for each encoder layer.

In this experiment we evaluate the error rate of Taris on LRS2 for an increasing number of decoder look-ahead segments  $d_{LA}$ , while setting the look-back value  $d_{LB} = \infty$ . For a practical online model it may be a good trade-off to limit the decoder look-back context to a single sentence when transcribing continuously. We plot the Character Error Rate in Figure 4 for an increasing number of acoustic segments that the decoder is allowed to attend to. We notice that there are di-



**Fig. 4.** Online decoding error rate on LRS2. We fix  $e_{LB} = \infty$ ,  $e_{LA} = 11$  frames,  $d_{LB} = \infty$  and we only allow  $d_{LA}$  to vary.


minishing returns after a context look-ahead  $d_{LA}$  of 4 words. The overall accuracy beyond this threshold is comparable to the offline systems shown in Figure 3b.

### 5.5. Evaluation on longer sentences

In the previous experiments we have used the LRS2 dataset for rapid prototyping. However, since it contains many short sentences, the potentially higher decoding error rate of Taris on the longer sentences might have little effect on the reported average error rate. We re-train and evaluate our models on the 100 hour clean partition of the LibriSpeech dataset, displaying the mean error and 95% confidence interval (CI) around the mean in Table 1.

First, we notice that the systems achieve an error rate similar to the one obtained on LRS2, despite the increased amount of data, suggesting that further gains are possible for larger model sizes. We also notice that the word loss can be slightly detrimental to the overall accuracy for the same network capacity, particularly for the models with unbounded attention span. This prompts a deeper investigation into the interplay between the cross entropy and word counting losses, as our constant scale factor  $\lambda$  is likely a less than optimal solution to this multitask problem.

Next, we compare the distributions of the segment lengths estimated by Taris and those estimated with the pre-trained Montreal forced aligner [31], both plotted in Figure 5. Not only are the histograms highly overlapped, but the one produced by Taris is in line with the average speaking rate of



**Fig. 5.** Segmentation length distribution (in milliseconds) of Taris compared to the reference provided by the Montreal forced aligner

read speech. The small differences between the reference and hypothesis are likely owed to the short silences between words which were excluded from the reference, whereas Taris does not explicitly model silences and includes them into segments. Latency has not received sufficient consideration in prior work to facilitate a direct comparison, as systems were trained with offline encoders [17] or large receptive fields [16], relied on beam search over the output distribution [19], or used phoneme units [8]. Very recent work combining a weaker online model with an offline rescorer [32] that allows to revise online hypotheses with a final hypothesis introduces the notion of end-pointing latency. The offline rescorer is triggered after the utterance has been determined to be finished, that is when a threshold period has elapsed after a suspected end of utterance without further speech activity. Since our model is fully online and does not have to wait to rescore, this metric is not applicable to our system.

### 5.6. Evaluation on Mandarin speech

Since the word segmentation strategy in Taris is tailored for English, we are interested in extending the principle to Mandarin speech. Unlike English, Mandarin is characterised by a low number of morphemes per word and has almost no inflectional affixes, being considered a highly analytic language. In addition, the commonly used writing system belongs to the *scriptio continua* style, with no delimitation between words. On these grounds, we make a structural change in Taris: instead of learning to count words (spaces between them), we let the system count the number of characters, similar to the quantity loss used in [19]. This would drive the system to segment the acoustics associated with each character, which is almost always equivalent to a syllable.

For our experiment we use the Aishell-1 dataset [33], which contains 165 hours of fluent speech recordings from 400 speakers coming mostly from the Northern area of China, and covers a broad range of topics. The transcription file comprises an inventory of 4333 characters, which will determine the final output size of the decoder. Since the



**Table 1.** System evaluation on LibriSpeech 100h clean partition.

Model	parameters				CER		Word Loss
	$e_{LB}$	$e_{LA}$	$d_{LB}$	$d_{LA}$	mean [%]	95% CI [%]	
Transformer [26]	$\infty$	$\infty$	$\infty$	$\infty$	13.37	0.444	<i>N/A</i>
Transformer + Word Loss	$\infty$	$\infty$	$\infty$	$\infty$	14.64	0.451	0.92
Taris : infinite look-back	$\infty$	11	$\infty$	5	15.70	0.451	1.12
Taris : finite look-back	11	11	5	5	13.83	0.451	0.76

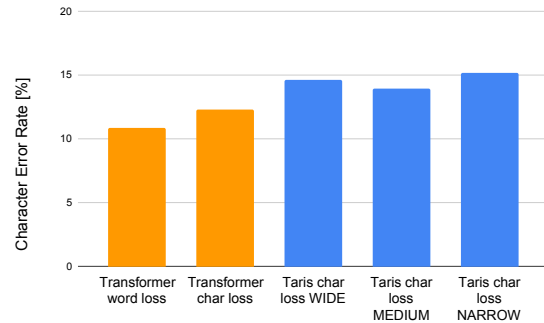
labels also include candidate blank spaces between words, we also evaluate Taris at the word level as we did on English. Despite the larger dataset size, we maintain the same Transformer size as before for faster prototyping. We label the different parametrisations of Taris as follows: WIDE:  $e_{LB} = e_{LA} = 11$  frames,  $d_{LB} = d_{LA} = 5$  segments, MEDIUM:  $e_{LB} = e_{LA} = 3$ ,  $d_{LB} = d_{LA} = 5$ , NARROW:  $e_{LB} = e_{LA} = 2$ ,  $d_{LB} = d_{LA} = 2$ .

Figure 6 shows the error rates of the offline and online systems on the Aishell-1 corpus. Despite the small model size, the absolute decoding accuracy is comparable with the baseline results in [33] obtained with the Kaldi toolkit. We notice that Taris does a much better job at learning to count the number of characters than the number of words, with our NARROW model obtaining a counting error of 0.2538. Since a Chinese character almost always corresponds to a single syllable, this result suggests that syllables may be easier to segment in fluent speech than words. Furthermore, the syllable level segmentation allows both the encoder and the decoder of Taris to use relatively low context lengths and further reduce the overall latency. Not shown in the figure, the error rate of a *word* counting Taris model is approximately 40%, implying that a good unit segmentation is essential for online decoding.

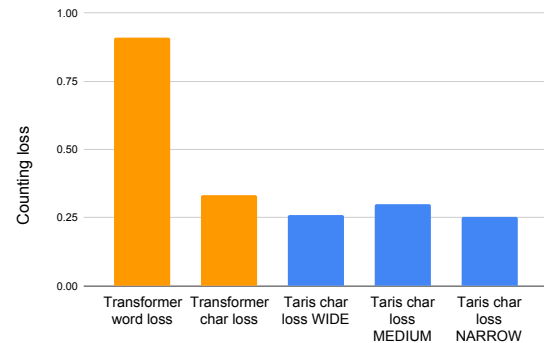
## 6. CONCLUSION

We have proposed a simple, efficient, and fully differentiable solution for online speech recognition that does not require additional labels. Taris is inspired from early language acquisition in infants, and aims to segment a speech stream by learning to count the number of words therein. We show that our method matches the accuracy of an offline system once it listens to 5 dynamic segments. Lowering this latency remains a topic for exploration, e.g. by gradually reducing the look-ahead parameter  $d_{LA}$  later in training, explicitly modelling silences, or investigating the role of context, grammar, and semantics in lexical recognition.

Generalising to sentences of different lengths from the ones seen in training has recently been identified as a major problem for neural online speech recognition systems [34, 35]. By modelling only the local relationships in speech through finite look-back and look-ahead, we preserve the same property of the Neural Transducer [8] to effectively decouple the sentence length from the learnt representations,



(a) Offline and online decoding error rate



(b) Unit counting loss

**Fig. 6.** System Evaluation on Aishell-1

while allowing adaptive segments and simpler training.

It can be argued that Taris exploits human knowledge of the speech signal structure and embeds the concept of words and the local acoustic relationships, instead of being a more generic, self-organising neural network. Yet, the local processing of speech is merely the one dimensional equivalent of local convolutions applied to images, where the objects are replaced by words. Moreover, one-dimensional convolutions are commonly used in speech recognition [36, 37, 38, 39]. Given their major impact in research despite their lack of invariance to orientation, scaling, or even small perturbations, there is still much to be learned from engineered models in the pursuit of artificial general intelligence.

The ADAPT Centre for Digital Content Technology is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund.

## 7. REFERENCES

- [1] C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, E. Gonina, N. Jaitly, B. Li, J. Chorowski, and M. Bacchiani, “State-of-the-Art Speech Recognition with Sequence-to-Sequence Models,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4774–4778.
- [2] P.W. Jusczyk and R.N. Aslin, “Infants’ detection of the sound patterns of words in fluent speech,” *Cognitive Psychology*, vol. 29, no. 1, pp. 1 – 23, 1995.
- [3] Elizabeth K. Johnson and Peter W. Jusczyk, “Word segmentation by 8-month-olds: When speech cues count more than statistics,” *Journal of Memory and Language*, vol. 44, no. 4, pp. 548 – 567, 2001.
- [4] Alex Graves, “Sequence transduction with recurrent neural networks,” *ICML Representation Learning Workshop*, 2012.
- [5] A. Graves, A. Mohamed, and G. Hinton, “Speech recognition with deep recurrent neural networks,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, 2013, pp. 6645–6649.
- [6] K. Rao, H. Sak, and R. Prabhavalkar, “Exploring architectures, data and units for streaming end-to-end speech recognition with rnn-transducer,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2017, pp. 193–199.
- [7] E. Battenberg, J. Chen, R. Child, A. Coates, Y. G. Y. Li, H. Liu, S. Satheesh, A. Sriram, and Z. Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2017, pp. 206–213.
- [8] Navdeep Jaitly, Quoc V Le, Oriol Vinyals, Ilya Sutskever, David Sussillo, and Samy Bengio, “An online sequence-to-sequence model using partial conditioning,” in *Advances in Neural Information Processing Systems 29*, 2016, pp. 5067–5075.
- [9] T. N. Sainath, C. Chiu, R. Prabhavalkar, A. Kannan, Y. Wu, P. Nguyen, and Z. Chen, “Improving the performance of online neural transducer models,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5864–5868.
- [10] Eugen Beck, Mirko Hannemann, Patrick Dötsch, Ralf Schlüter, and Hermann Ney, “Segmental encoder-decoder models for large vocabulary automatic speech recognition,” in *Proc. Interspeech*, 2018, pp. 766–770.
- [11] H. Tang, L. Lu, L. Kong, K. Gimpel, K. Livescu, C. Dyer, N. A. Smith, and S. Renals, “End-to-end neural segmental models for speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1254–1264, 2017.
- [12] Y. Luo, C. Chiu, N. Jaitly, and I. Sutskever, “Learning online alignments with continuous rewards policy gradient,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 2801–2805.
- [13] Colin Raffel, Minh-Thang Luong, Peter J. Liu, Ron J. Weiss, and Douglas Eck, “Online and linear-time attention by enforcing monotonic alignments,” in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 2017, ICML’17, p. 2837–2846.
- [14] Chung-Cheng Chiu and Colin Raffel, “Monotonic chunkwise attention,” in *International Conference on Learning Representations*, 2018.
- [15] Ruchao Fan, Pan Zhou, Wei Chen, Jia Jia, and Gang Liu, “An Online Attention-Based Model for Speech Recognition,” in *Proc. Interspeech*, 2019, pp. 4390–4394.
- [16] Junfeng Hou, Wu Guo, Yan Song, and Li-Rong Dai, “Segment boundary detection directed attention for online end-to-end speech recognition,” *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2020, no. 1, pp. 3, 2020.
- [17] N. Moritz, T. Hori, and J. L. Roux, “Triggered attention for end-to-end speech recognition,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 5666–5670.
- [18] M. Li, M. Liu, and H. Masanori, “End-to-end Speech Recognition with Adaptive Computation Steps,” in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6246–6250.
- [19] L. Dong and B. Xu, “Cif: Continuous integrate-and-fire for end-to-end speech recognition,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6079–6083.
- [20] J. Drexler and J. Glass, “Learning a subword inventory jointly with end-to-end automatic speech recognition,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6439–6443.
- [21] Julia Kreutzer and Artem Sokolov, “Learning to segment inputs for nmt favors character-level processing,” in *Proceedings of the International Workshop on Spoken Language Translation*, 2018.

- [22] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015*.
- [23] Ilya Sutskever, Oriol Vinyals, and Quoc V Le, “Sequence to Sequence Learning with Neural Networks,” in *Advances in Neural Information Processing Systems 27*, 2014, pp. 3104–3112.
- [24] Jan Chorowski, Dzmitry Bahdanau, Dmitriy Serdyuk, Kyunghyun Cho, and Yoshua Bengio, “Attention-based models for speech recognition,” in *Proceedings of the 28th International Conference on Neural Information Processing Systems*. 2015, pp. 577–585, MIT Press.
- [25] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 4960–4964.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems 30*, 2017, pp. 5998–6008.
- [27] BBC and University of Oxford, “The Oxford-BBC Lip Reading Sentences 2 (LRS2) Dataset,” [http://www.robots.ox.ac.uk/~vgg/data/lip\\_reading/lrs2.html](http://www.robots.ox.ac.uk/~vgg/data/lip_reading/lrs2.html), 2017.
- [28] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5206–5210.
- [29] G. Sterpu, C. Saam, and N. Harte, “How to teach dnns to pay attention to the visual modality in speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1052–1064, 2020.
- [30] The TensorFlow Model Garden, “Transformer Translation Model,” <https://github.com/tensorflow/models/tree/r2.1.0/official/nlp/transformer>, 2020.
- [31] Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldif,” in *Proc. Interspeech 2017*, 2017, pp. 498–502.
- [32] T. N. Sainath, Y. He, B. Li, A. Narayanan, R. Pang, A. Bruguier, S. Chang, W. Li, R. Alvarez, Z. Chen, C. Chiu, D. Garcia, A. Gruenstein, K. Hu, A. Kannan, Q. Liang, I. McGraw, C. Peyser, R. Prabhavalkar, G. Pundak, D. Rybach, Y. Shangguan, Y. Sheth, T. Strohmaier, M. Visontai, Y. Wu, Y. Zhang, and D. Zhao, “A streaming on-device end-to-end model surpassing server-side conventional model quality and latency,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6059–6063.
- [33] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, “Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline,” in *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*, 2017, pp. 1–5.
- [34] C. Chiu, W. Han, Y. Zhang, R. Pang, S. Kishchenko, P. Nguyen, A. Narayanan, H. Liao, S. Zhang, A. Kannan, R. Prabhavalkar, Z. Chen, T. Sainath, and Y. Wu, “A comparison of end-to-end models for long-form speech recognition,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 889–896.
- [35] A. Narayanan, R. Prabhavalkar, C. Chiu, D. Rybach, T. N. Sainath, and T. Strohmaier, “Recognizing long-form speech using streaming end-to-end models,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019, pp. 920–927.
- [36] O. Abdel-Hamid, A. Mohamed, H. Jiang, L. Deng, G. Penn, and D. Yu, “Convolutional neural networks for speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 10, pp. 1533–1545, 2014.
- [37] V. Pratap, A. Hannun, Q. Xu, J. Cai, J. Kahn, G. Synnaeve, V. Liptchinsky, and R. Collobert, “Wav2letter++: A fast open-source speech recognition system,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6460–6464.
- [38] Jason Li, Vitaly Lavrukhin, Boris Ginsburg, Ryan Leary, Oleksii Kuchaiev, Jonathan M. Cohen, Huyen Nguyen, and Ravi Teja Gadde, “Jasper: An End-to-End Convolutional Neural Acoustic Model,” in *Proc. Interspeech*, 2019, pp. 71–75.
- [39] S. Kriman, S. Beliaev, B. Ginsburg, J. Huang, O. Kuchaiev, V. Lavrukhin, R. Leary, J. Li, and Y. Zhang, “Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6124–6128.