

# Fast End-to-End Speech Recognition via Non-Autoregressive Models and Cross-Modal Knowledge Transferring from BERT

Ye Bai, *Student Member, IEEE*, Jiangyan Yi, *Member, IEEE*, Jianhua Tao, *Senior Member, IEEE*  
Zhengkun Tian, *Student Member, IEEE* Zhengqi Wen, *Member, IEEE* Shuai Zhang, *Student Member, IEEE*

**Abstract**—Attention-based encoder-decoder (AED) models have achieved promising performance in speech recognition. However, because the decoder predicts text tokens (such as characters or words) in an autoregressive manner, it is difficult for an AED model to predict all tokens in parallel. This makes the inference speed relatively slow. In contrast, we propose an end-to-end non-autoregressive speech recognition model called LASO (Listen Attentively, and Spell Once). The model aggregates encoded speech features into the hidden representations corresponding to each token with attention mechanisms. Thus, the model can capture the token relations by self-attention on the aggregated hidden representations from the whole speech signal rather than autoregressive modeling on tokens. Without explicitly autoregressive language modeling, this model predicts all tokens in the sequence in parallel so that the inference is efficient. Moreover, we propose a cross-modal transfer learning method to use a text-modal language model to improve the performance of speech-modal LASO by aligning token semantics. We conduct experiments on two scales of public Chinese speech datasets AISHELL-1 and AISHELL-2. Experimental results show that our proposed model achieves a speedup of about 50× and competitive performance, compared with the autoregressive transformer models. And the cross-modal knowledge transferring from the text-modal model can improve the performance of the speech-modal model.

**Index Terms**—speech recognition, fast, end-to-end, non-autoregressive, attention, BERT, cross-modal, transfer learning

## I. INTRODUCTION

DEEP learning has significantly improved the performance of automatic speech recognition (ASR). Conventionally, an ASR system consists of an acoustic model (AM), a pronunciation lexicon, and a language model (LM). The deep neural network (DNN) is used to model observation probabilities of the hidden Markov models (HMMs) [1], [2]. This DNN-HMM hybrid approach has achieved success in ASR. However, the pipeline of a DNN-HMM hybrid system usually requires training Gaussian mixture model based HMMs (GMM-HMM) for

generating frame-level alignments and tying states. Building the pronunciation lexicon requires the knowledge of experts in phonetics. This complexity of the building pipeline limits the development of an ASR system. Moreover, the different building procedures of the AM and the LM make the system difficult to be optimized jointly. Thus, the possible error accumulation in the pipeline influences the performance of a hybrid ASR system.

Pure neural network based end-to-end (E2E) ASR systems attract interests of researchers these years [3], [4], [5], [6], [7]. Different from the hybrid ASR systems, these systems use one DNN to model acoustic and language simultaneously so that the network can be optimized with back-propagation algorithms in an E2E manner. In particular, attention-based encoder-decoder (AED) models have achieved promising performance in ASR [8], [9]. The AED models first encode the acoustic feature sequence into latent representations with an encoder. With the latent representations, the decoder predicts the text token sequence step-by-step. The attention mechanism queries a proper latent vector from the outputs of the encoder for the decoder to predict. However, even with the non-recurrent structure which can be implemented in parallel [10], [11], the recognition speed limits the deployment of AED models in real-world applications. Two main reasons influence inference speed:

- First, the encoder encodes the whole utterance, so the decoder starts inference after the user speaks out the whole utterance;
- Second, multi-pass forward propagation of the decoder costs much time during beam-search.

Several work focus on the first problem to make the system can generate the token sequence in a streaming manner. Monotonic attention mechanism [12], [13] enforces the attention alignments to be monotonic. Therefore, the encoder only encodes a local chunk in the acoustic feature sequence, and the decoder predicts the next token without the future context in the acoustic feature sequence. Triggered-attention systems [14], [15] use connectionist temporal classification (CTC) spikes to segment the acoustic feature sequence adaptively, and the decoder predicts the next token when it is activated by a spike. Transducer-based models [7], [16], [17] use an extra blank token and marginalize all possible alignments, so the model can immediately predict the next token when the encoder accumulates enough information. All these models

Ye Bai, Zhengkun Tian, Jianhua Tao and Shuai Zhang are with the University of Chinese Academy of Sciences, Beijing 100190, China, and with NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China. (e-mail: baiye2016@ia.ac.cn; zhengkun.tian@nlpr.ia.ac.cn; shuai.zhang@nlpr.ia.ac.cn).

Jiangyan Yi and Zhengqi Wen are with NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China (e-mail: jiangyan.yi@nlpr.ia.ac.cn; zqwen@nlpr.ia.ac.cn).

Jianhua Tao is also with NLPR, Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China, and the CAS Center for Excellence in Brain Science and Intelligence Technology, Beijing 100190, China (e-mail: jhtao@nlpr.ia.ac.cn). (Corresponding authors: Jiangyan Yi; Jianhua Tao.)

can generate a token sequence in a steaming manner and show promising results. However, they limit the models using local information in the speech sequence. It potentially ignores global semantic relationships in the speech sequence. The global semantic relationships contain not only the relationships among acoustic frames but also the relationships among tokens [18], as shown in Fig. 1.

In this paper, we aim to address the second problem, i.e., we would like to generate the token sequence without beam-search. We propose an attention-based feedforward neural network model for non-autoregressive speech recognition called “LASO” (Listen Attentively, and Spell Once). The LASO model first uses an encoder to encode the whole acoustic sequence into high-level representations. Then, the proposed *position dependent summarizer* (PDS) module queries the latent representation corresponding to the token position from the outputs of the encoder. It bridges the length gap between the speech and the token sequence. At last, the decoder further refines the representations and predicts a token for each position. Because the prediction of one token does not depend on another token, beam-search is not used. And because the network is a non-recurrent feedforward structure, it can be implemented in parallel. To further improve the ability to capture the semantic relationship of LASO (especially for the decoder), we propose to use a cross-modal knowledge transferring method. Specifically, we align the semantic spaces of the hidden representations of the LASO and the pre-trained large-scale language model BERT [19]. We use the teacher-student learning based knowledge transferring method [20] to leverage knowledge in BERT. We conduct experiments on two publicly available Chinese Mandarin datasets to evaluate the proposed methods with different data sizes. The experiments demonstrate that our proposed method achieves a competitive performance and efficiency.

The contributions of this work are summarized as follows.

- 1) We propose a non-autoregressive attention-based feedforward neural network LASO for speech recognition. A PDS module is proposed to convert a variant-length speech sequence to a fixed-length sequence. Each representation in the fixed-length sequence corresponds to the position of a token so that the ASR is considered as a position-wise classification. A decoder, which plays a role of bidirectional LM, further captures the relationship among these representations with self-attention. LASO leverages the whole context of the speech and generates each token in the sequence in parallel. The experiments demonstrate that our proposed model achieves high efficiency and competitive performance.
- 2) We propose a cross-modal knowledge transferring method from BERT for improving the performance of LASO. The experiments demonstrate the effectiveness of the knowledge transferring from BERT. The results also show that the speech signals have similar internal structures with corresponding text so that knowledge from BERT can benefit the non-autoregressive ASR model.
- 3) We present detailed visualization of the model. The visualization results show that the proposed PDS module




Fig. 1. A spectrogram of an example utterance. A word corresponds to a segment in the speech signal. The relationships among the segments can be seen as the relationships among the corresponding tokens, which are referred to as language semantics in this paper.

can attend to specific encoded acoustic representations and generates meaningful hidden representations corresponding to tokens from speech. And the decoder can capture token relationships from the aggregated hidden representations.

Compared with the preliminary version [21], the new content in this paper includes leveraging pre-trained BERT models to further improve the performance, graph-based view of the proposed model, more detailed experiments on large-scale datasets, detailed visualization and analysis for better understanding the models. The rest of the paper is organized as follows. Section II briefly compares the autoregressive AED models and the non-autoregressive AED models. Section III re-formulates the speech recognition as a position-wise classification problem. Section IV describes the proposed LASO model. Section V describes how to train the model and how we distill the knowledge from BERT. Section VI introduces how the model generates a sentence. Section VII compares this work with previous related work. Section VIII and Section IX present setup and results of experiments, respectively. Section X discusses this paper. At last, Section XI concludes this paper and presents future work.

## II. BACKGROUND: AUTOREGRESSIVE MODELS VS. NON-AUTOREGRESSIVE MODELS

In this section, we introduce the background of autoregressive AED models (ARM) and non-autoregressive AED models (NARM). We compare these two paradigms to better introduce the proposed method in the rest sections.

### A. Autoregressive AED Models

The ARM predicts the next token based on the previously generated tokens. That is, for a speech-text pair  $(X, Y)$ , the ARM factorizes the conditional probability  $P(Y|X)$  with the chain rule:

$$P_{\text{ARM}}(Y|X) = P(y_1|X) \prod_{j=2}^L P(y_j|y_{<j}, X), \quad (1)$$

where  $X = [x_1, \dots, x_T]$  is the acoustic feature sequence, each  $x$  denotes a feature vector,  $Y = [y_1, \dots, y_L]$  is the text token sequence, and  $y_{<j} = [y_1, \dots, y_{j-1}]$  denotes the previous context of the token  $y_j$ . The token can be a word or a sub-word (phone, character, or word piece). This can be seen as a conditional language model, i.e., the model estimates

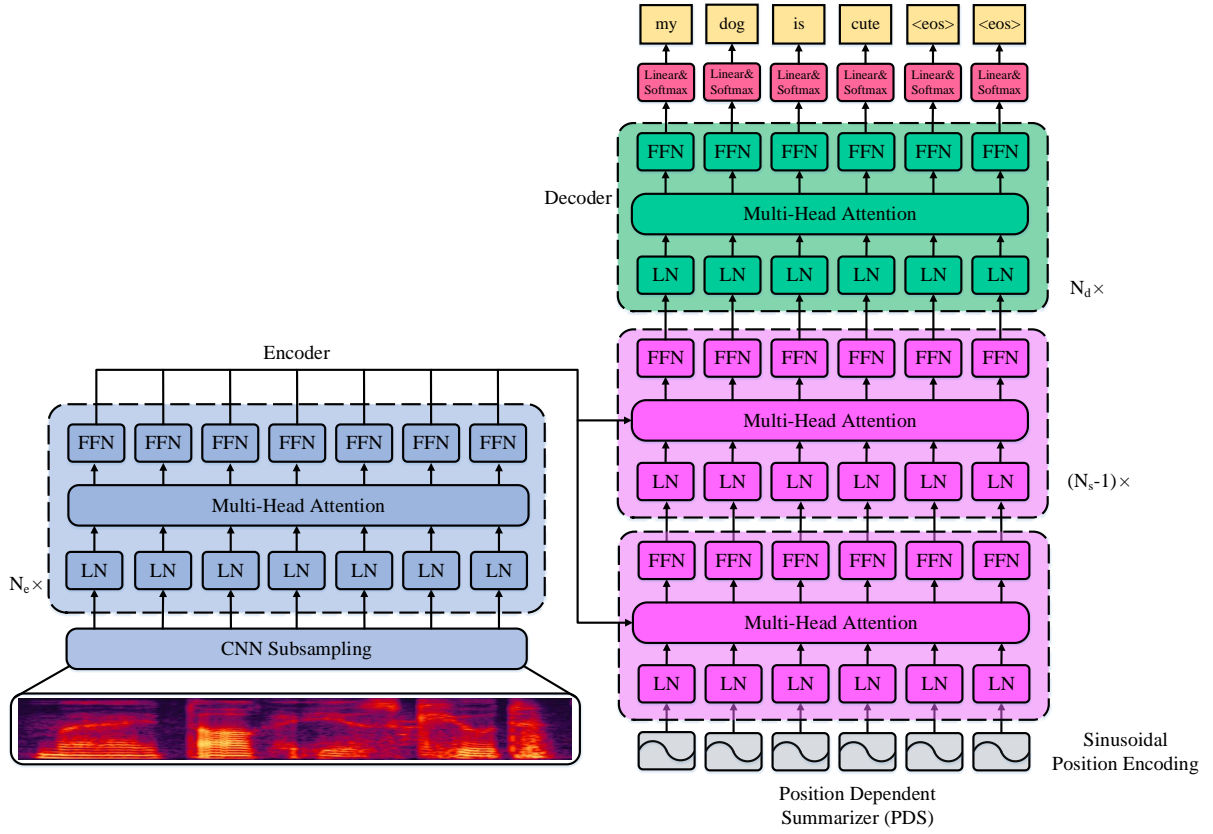


Fig. 2. An illustration of the proposed LASO model. The LASO consists of an encoder, a position-dependent summarizer (PDS), and a decoder. All three modules are composed of basic attention blocks, which consist of multi-head attention and a position-wise feedforward network (FFN in the figure). We first use a CNN to subsample the acoustic feature sequence. Then, the encoder extracts high-level representations from the subsampled sequence. The PDS queries the high-level acoustic representations corresponding to each token position. Then, the decoder further refines the language semantics. For each position, a probability distribution over the vocabulary is computed with a softmax function. And during inference, we select the most likely token at each position to form the token sequence. The extra tokens  $\langle \text{eos} \rangle$  are removed in the token sequence. Each attention block includes layer normalization (LN in the figure) and a residual connection. And we add sinusoidal position encodings to the subsampled acoustic feature sequence. The whole architecture is non-recurrent, so it can be implemented in parallel for fast inference.

the probability of the token sequence given the acoustic feature sequence. Typically, Eq. (1) is implemented with a neural network, which is an encoder-decoder architecture. The encoder encodes the acoustic feature sequence, and the decoder computes  $P(y_j|y_{<j}, X)$  step-by-step.

To find the token sequence which has the highest probability approximately, a beam-search algorithm is used during inference. The decoder maintains a beam of potential candidates of the token sequence which have high probabilities. Thus, the decoder has to forward propagate these candidates. In addition, the generation of each token depends on the previously generated tokens, so it is difficult to implement parallel generation.

### B. Non-Autoregressive AED Models

Different from the ARM, the NARM predicts each token without dependence on other tokens. Specifically, the NARM assumes the conditional independence on tokens:

$$P_{\text{NARM}}(Y|X) = \prod_{j=1}^L P(y_j|X). \quad (2)$$

Because each probability does not depend on the other tokens, parallel implementation is possible. Another view of Eq. (2)

is to process each token independently rather than to process the production. The details are described in Section III.

Conventionally, the relationships among tokens are considered an important factor for the token sequence generation. We refer to this relationship as *language semantics* in this paper. The previous non-autoregressive model CTC assumes conditional independence on token sequence [22]. However, to achieve good performance, the CTC-based systems use n-gram LMs for modeling the language semantics [23]. And the advanced version transducer models [7] use a neural network to model the language semantics. An AED model captures the language semantics by the decoder [3], [4], [5]. Different from them, in this paper, we propose to use a self-attention mechanism to model the implicit language semantics, which compensates for the loss of the explicit autoregressive language model. The details are described in Section IV.

### III. ASR AS POSITION-WISE CLASSIFICATION

In this paper, we give a new perspective on the speech recognition problem. The basic idea is that the language semantics is implicitly contained in the speech signal. Fig. 1 shows a spectrogram of an utterance “my dog is cute”. Each segment corresponds to a word in the utterance. We observe

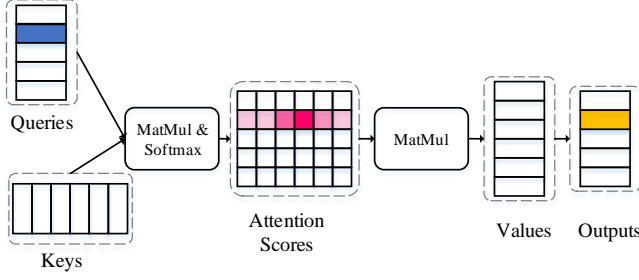


Fig. 3. An illustration of dot-product attention. The queries and the keys are used to compute the attention scores with matrix multiplication and the softmax function. The attention scores are used as weights to fuse values.

that the language semantics, i.e., the relationships among the tokens, is expressed among the segments implicitly. Thus, if the speech signal of the whole utterance is available, we can leverage this implicit language semantic to improve the performance of ASR.

Based on the above observations, we consider the ASR problem as position-wise classification, given the whole speech utterance. Namely, we use the whole acoustic feature sequence, including explicit acoustic characteristics and implicit language semantics, to predict one token, but not estimate the probability of the token sequence. When all the tokens are predicted, we simply put them together as the recognition result. Formally, we predict the following probability:

$$P(y_j|X) = f(X), \quad j = 1, \dots, L, \quad (3)$$

where  $X$  denotes the whole acoustic feature sequence,  $y_j$  denotes a token in the token sequence,  $L$  is the length of the token sequence, and  $f$  is some non-linear function. In this paper, we use the proposed feedforward neural network LASO as the non-linear function  $f$ . Usually, the length of the token sequence is unknown in advance. We use a simple way to tackle this. We set  $L$  as a big enough number, and the tail of the token sequence will be predicted to the filler token.

#### IV. THE PROPOSED LASO MODEL

In this section, we introduce the proposed LASO model. The architecture is shown in Fig. 2. The encoder encodes the acoustic feature sequence into high-level representations. The PDS queries the high-level representation corresponding to each token position. Another purpose of the PDS is to bridge the length gap between the speech and the token sequence. The decoder further captures language semantics from the outputs of the PDS. At last, the probability distributions over the vocabulary are computed with the linear transformations and the softmax functions. For inference, the most likely token at each position is selected. For the token sequence whose length is shorter than  $L$ , the tail is filled with the filler token  $\langle \text{eos} \rangle$ . This also predicts the length of the token sequence automatically. These filler tokens are easily removed.

The whole network is a feedforward structure so that it can be implemented in parallel to make the inference fast. We introduce the details of the basic attention block and each

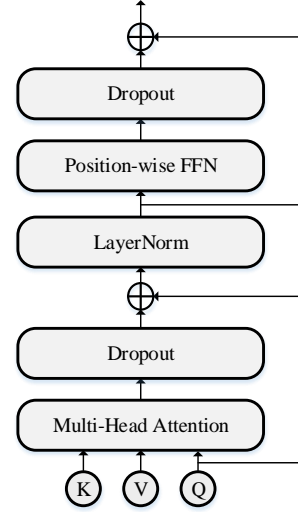


Fig. 4. An illustration of an attention block.

module in the rest of this section. We also introduce a graph-based view of the whole model.

##### A. Attention Block

The feedforward attention structure [24] captures the global relationship in a sequence. Different from recurrent neural networks which encode history context into latent vectors, the feedforward attention mechanism uses a weighted sum to fuse the input sequence. Because of its feedforward structure, it can be computed in parallel. In this work, we use the scaled dot-product attention and position-wise feedforward network as the basic submodule, following [24]. But we use “pre-norm” [25] for stable training. The structure is shown in Fig. 4.

The scaled dot-product attention is computed by

$$\text{Atten}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{D_k}}\right)V, \quad (4)$$

where  $Q \in \mathbb{R}^{T_q \times D_k}$  denotes the queries,  $K \in \mathbb{R}^{T_k \times D_k}$  denotes the keys, and  $V \in \mathbb{R}^{T_k \times D_v}$  denotes values. As shown in Fig. 3, the attention scores are computed with the dot products of queries and keys, then are normalized with softmax functions. The normalized attention scores will be sharp at some positions, and others are small. Then, by matrix multiplication, the values are fused to the corresponding position. This procedure can be seen as that the query queries keys and fetches out a corresponding value from values.

To make the attention scores various, it can be extended to a multi-head version:

$$\begin{aligned} \text{MHA}(Q, K, V) &= \text{Concat}(h_1, \dots, h_H)W^o, \\ h_i &= \text{Atten}(QW_i^q, KW_i^k, VW_i^v), i = 1, \dots, H. \end{aligned} \quad (5)$$

The queries, keys, and values are transformed into subspaces with parameter matrices  $W_i^q, W_i^k, W_i^v$ , where  $i$  is the index of a head. Then, the scaled dot-product attention is computed for the transformed inputs. At last, the outputs are concatenated together and multiplied with  $W^o$ .  $h_i$  is one attention head.  $H$  is the number of heads.



A position-wise feedforward neural network (FFN) transforms the output of the attention at each position:

$$\text{FFN}(u) = W_2 \text{Activate}(W_1 u + b_1) + b_2, \quad (6)$$

where  $u$  is a vector at one position,  $W_1$ ,  $W_2$ ,  $b_1$ , and  $b_2$  are learnable parameters, “Activate” is a nonlinear activation function. In this work, gated linear units (GLUs) [26] are used.

Residual connection [27] and layer normalization (LN) [28] are also used. Different from [24], we set the LN layer before the residual connection, following [25] to make training stable and effective. The structure of the attention block is shown in Fig. 4.

### B. Encoder

The encoder extracts high-level representations from the acoustic feature sequence. We first use a two-layer convolutional neural network (CNN) to subsample the acoustic feature sequence. We set the stride on the time axis to 2, so that the frame rate is reduced to 1/4. Then the outputs of the CNN are flattened to a  $T$ -by- $D_m$  matrix, where  $T$  is the length of the subsampled feature sequence, and  $D_m$  is the dimensionality. Another purpose of the CNN is to capture the locality of the acoustic feature sequence.

Then, the encoder has a stack of  $N_e$  attention blocks, as shown in Fig. 2. The keys, queries, and values are all the same, so it is a self-attention mechanism. That is, the attention scores are obtained by computing dot-product between every two vectors of the inputs. Therefore, the long-term dependency is captured.

### C. Position Dependent Summarizer

The PDS module is the core of the proposed LASO model. The PDS module leverages queries, which depend on positions of the token sequence, to query the high-level representations from the encoder. It can be seen that the module “summarizes” the acoustic features, so we name it as “summarizer”. As the result, it bridges the gap between the length of the acoustic feature sequence and the length of the token sequence.

The PDS module consists of  $N_s$  attention blocks, as shown in Fig. 2. For the first block, the queries are position encodings. And the queries of the other blocks are the outputs of the previous block. Each query represents the token position in the token sequence. The keys and the values are all the outputs of the encoder. The length of  $L$  of the position encoding sequence is pre-set by counting the length of utterances in the training set and add some tolerance. For example, if the maximum length of token sequences in the training set is 90, we can set this  $L$  to 100.

We use sinusoidal position encodings [24]:

$$\begin{aligned} \text{pe}_{i,2j} &= \sin(i/10000^{2j/D_m}), \\ \text{pe}_{i,2j+1} &= \cos(i/10000^{2j/D_m}), \end{aligned} \quad (7)$$

where  $i = 1, \dots, L$  denotes the  $i$ -th position and  $j$  is the index of an element. A benefit of using sinusoidal position encodings is that it would allow the model to easily learn relativity of positions. That is, the positional encoding of a fixed distance between two positions  $k$  can be represented as a linear function of the two positions [24].

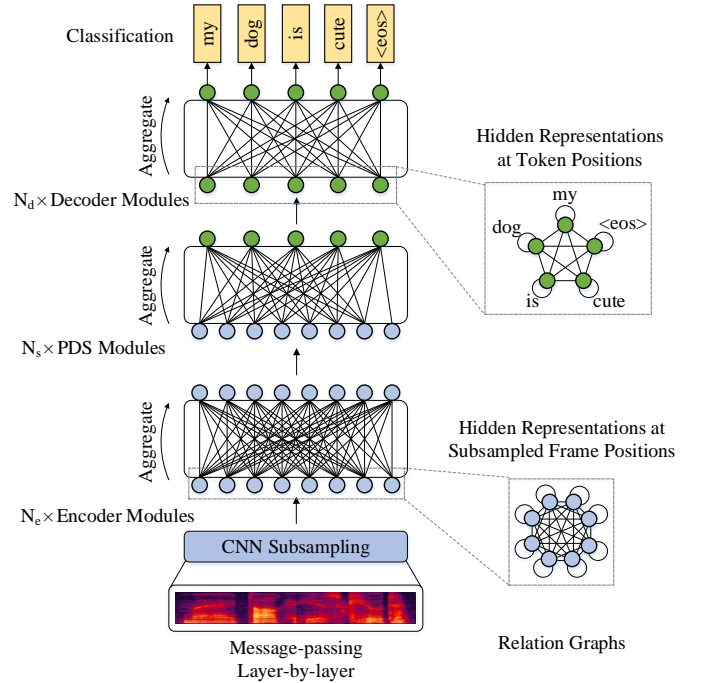


Fig. 5. An illustration of the LASO in a graph-based view. The left part shows that the messages are passed layer-by-layer. Each module (i.e. encoder module, PDS module, decoder module) aggregates the previous representations by the attention mechanism. The right part shows the relationships among representations in one layer.

### D. Decoder

The decoder further refines the representations of the PDS module. Similar to the encoder, it is a self-attention module, which consists of  $N_d$  attention blocks. It captures the relationships in the sequence, i.e., implicit language semantics, which is queried by the PDS. The inputs of the decoder are the outputs of the PDS. This decoder leverages the whole context of the utterance.

After the decoder, a linear transformation and a softmax function are used to compute the probability distribution over the token vocabulary.

### E. The Graph-Based View of the Model

We introduce a graph-based view of LASO to further illustrate the proposed model, as shown in Fig. 5. This view can be divided into two aspects: a relation graph [29] in one layer and message-passing of the forward procedure [30], [31].

For the aspect of the relations of the representations in one layer, the hidden representations can be seen as a fully connected graph with self-loops. Specifically, for the encoder which processes the acoustic features, the hidden representation at each position is related to all the hidden representations (including itself). Thus, the representation sequence can be viewed as a fully connected graph (the graph with blue nodes in the right part of Fig. 5). With the PDS module, the representations are converted and are related to token positions. These representations correspond to each token position and consist of another fully connected graph (the graph with green nodes in the right part of Fig. 5), i.e., one representation is related to all the representations in the sequence. The weights on the

edges are computed with the attention mechanism. This is different from the autoregressive models. For an autoregressive model, the representation of a token is related to the previous representations, but not all the representations in the sequence.

For the aspect of the forward procedure, the forward propagation can be seen as a message-passing procedure in each module. Specifically, the outputs and the inputs of a module consist of a bipartite graph, as shown in the left part of Fig. 5. The weights of the edges are attention scores. All representations in the inputted sequence are aggregated to the corresponding node of the output sequence by the weighing sum. Therefore, the messages in the previous relation graphs are passed to the next one in the forward propagation.

The graph-based view shows how LASO processes the acoustic feature sequence as a whole and achieves hidden representations at token positions. It also shows why the hidden representations at token positions are related to all the other hidden representations, which is different from the autoregressive model.

#### F. Formulation

The LASO model can be formulated as follows:

$$\begin{aligned} Z &= \text{Enc}(X), \\ q_i &= \text{Summarize}(Z, \text{pe}_i), \quad i = 1, 2, \dots, L, \\ Q &= [q_1, \dots, q_L] \\ P(y_i|X) &= \text{Dec}(Q), \quad i = 1, 2, \dots, L, \end{aligned} \quad (8)$$

where  $X = [x_1, \dots, x_T]$  is the feature sequence,  $Z$  denotes the high-level representations encoded with the encoder, and the probability over the vocabulary at each position  $P(y_i|X)$  is computed with the PDS and the decoder. Function “Summarize” represents the PDS module, which attends the outputs of the encoder in terms of the positional encodings. Because the positional encodings, which are deterministic but not random, are a part of the whole model, they are not written in the probability expression.

#### V. LEARNING

In this section, we introduce the learning procedure of the LASO model.

##### A. Maximum Likelihood Estimation

We use *maximum likelihood estimation* (MLE) criterion to train the parameters of the LASO model. We minimize the following negative log-likelihood (NLL) loss.

$$\text{NLL}(\theta) = -\frac{1}{NL} \sum_{n=1}^N \sum_{i=1}^L \log P_{\theta}(y_i^{(n)}|X^{(n)}), \quad (9)$$

where  $X^{(n)}, Y^{(n)}$  is the  $n$ -th speech-text pair in the corpus. The total number of the pairs is  $N$ .  $y_i^{(n)}$  is  $i$ -th token in  $Y^{(n)}$ .  $L$  is the length of the token sequence, which is preset. If the length of a text data is shorter than  $L$ ,  $\langle \text{eos} \rangle$  is used to pad it, as shown in Fig. 2. Thus, the length of the token sequence is estimated automatically.  $\theta$  denotes trainable parameters of the model.

This training procedure is end-to-end and does not depend on any frame-level label generated by some GMM-HMM system.

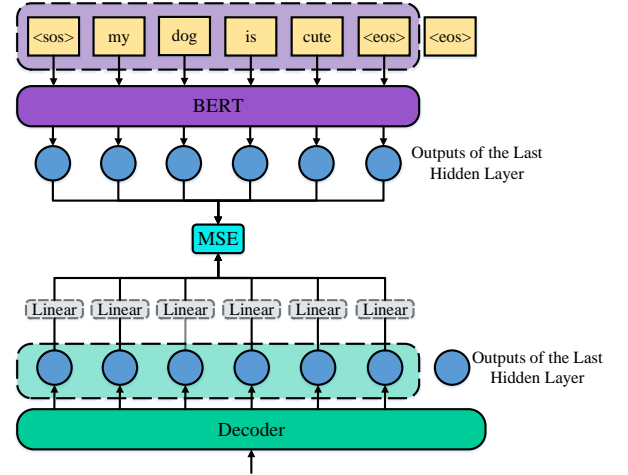


Fig. 6. An illustration of semantic refinement from BERT. The valid part of the token sequence is inputted into BERT. And the MSE loss between the last hidden layers of decoder and BERT is minimized. The optional linear transformation is used when the dimensionalities of the decoder and BERT are different.

##### B. Semantic Refinement from BERT

To further improve the performance, we use teacher-student learning [32], [33], [34], [35] to refine knowledge from pre-trained LM. BERT, a kind of denoising autoencoder LM trained on very large-scale text, has shown the powerful ability of language modeling and achieved state-of-the-art performance on many NLP tasks [19]. Inspired by our previous paper [20], we transfer the knowledge from BERT to the LASO model. Another potential advantage of using BERT as the teacher model is that both our proposed LASO and BERT are bidirectional models, i.e., the model predicts a token using both the left context and the right context.

The basic idea is that the BERT can provide a good semantic representation for each token. And we consider the outputs of the decoder also provide token-level representation. Therefore, we make the outputs of the decoder approximate the BERT. We minimize the mean squared error (MSE) between their last hidden layers [35], as shown in Fig. 6. To match the training procedure of BERT, we add a  $\langle \text{sos} \rangle$  token at the head of the token sequence.

We input the token sequence into BERT model to fetch out the outputs of the last hidden layer. Note that we only input the valid part of the token sequence, but not the padding token  $\langle \text{eos} \rangle$  at the tail.  $\langle \text{sos} \rangle$  and  $\langle \text{eos} \rangle$  are converted to  $[\text{CLS}]$  and  $[\text{SEP}]$ , which are two special tokens added at the head and the tail of a sentence in BERT [19].

We also use the valid part of the outputs of the decoder, i.e., the outputs corresponding to the subsequence from  $\langle \text{sos} \rangle$  to the first  $\langle \text{eos} \rangle$ . Then, we compute the MSE:

$$\text{MSE}(\theta) = \frac{1}{N} \sum_{n=1}^N \frac{1}{L_v^{(n)}} \sum_{i=1}^{L_v^{(n)}} \sum_{d=1}^{D_B} (h_{i,d}^{(n)} - s_{i,d}^{(n)})^2, \quad (10)$$

where  $h_{i,d}^{(n)}$  is the  $d$ -th element of the  $i$ -th vector in the decoder output sequence for the  $n$ -th data, and  $s_{i,d}^{(n)}$  is the  $d$ -th element of the  $i$ -th vector in the BERT output sequence for the  $n$ -th

data,  $D_B$  is the dimensionality of BERT output,  $L_v^{(n)}$  is the valid length of  $n$ -th data, and  $N$  is the total number of the data.  $\theta$  represents all parameters of the model.

If the dimensionality of the decoder is different from BERT, we can simply use an optional linear transformation to make the dimensionality matching, as shown in Fig. 6.

Another benefit of transferring knowledge with the hidden layer rather than probability is that it makes the construction of vocabulary more flexible. We can use only a part vocabulary of BERT rather than all tokens.

At last, we combine the NLL loss and the MSE loss as the final loss:

$$L(\theta) = \text{NLL}(\theta) + \lambda \text{MSE}(\theta), \quad (11)$$

where  $\lambda$  is a coefficient to balance the values of the two losses. The typical value is 0.005.

The BERT model is only used at the training stage. It does not add any extra complexity during inference.

## VI. INFERENCE

The inference of LASO is simple. We just select the most likely token at each position:

$$\hat{y}_i = \arg \max_{y_i} P(y_i|X), \quad i = 1, \dots, L, \quad (12)$$

Then, the special tokens  $\langle \text{eos} \rangle$  (or including  $\langle \text{sos} \rangle$  if the model is trained with BERT) are removed. Note that the positional encodings ( $[pe_1; \dots; pe_L]$ ) are a part of the whole model so that they are inputted into the model as a whole.

This inference procedure does not depend on beam-search, so multi-pass forward propagation is not needed. Thus, the inference time cost is much reduced.

## VII. RELATED WORK

In this section, we review and compare the previously related work in two main aspects. One is the non-autoregressive AED model. And the other is the utilization of LMs for AED models.

### A. Non-Autoregressive AED Models

Non-autoregressive AED models are first used in machine translation (MT). Gu *et al.* first proposed non-autoregressive machine translation and introduced fertility to tack the multimodality problem [36]. Auxiliary regularization [37] and enhanced decoder input [38] are proposed to improve the performance. Lee *et al.* proposed an iterative refinement algorithm for MT [39]. MA *et al.* proposed a flow model for sequence generation [40]. These models showed promising results on both performance and efficiency. However, these models are used for MT but not ASR. Speech signal has its specific property, for example, the monotonic alignment with the token sequence, and the implicit language semantics in speech. This motivated us to propose a simpler non-autoregressive AED model for ASR. The non-autoregressive transformer, which completes the masked tokens iteratively, is proposed for ASR [41]. MASK-CTC uses decoding results as the initialization and refines the results with a masked language model [42]. Different from the previous work, we propose a

non-autoregressive model LASO, which forward propagates for one-pass. We reformulate speech recognition as a position-wise classification problem. We propose the PDS to extract token-level representation from speech. The PDS bridges the length gap between the speech and the token sequence.

### B. The Utilization of LMs

Fusion methods, such as shallow fusion, deep fusion, and cold fusion, integrate external LMs to improve the performance [43], [44]. However, these methods add extra complexity during inference. And these methods can only use unidirectional LMs, so recent powerful bidirectional LMs such as ELMo [45], BERT [19], are not applicable. BERT was used to rescore the n-best results [46]. However, rescoring increases computation during inference. And it does not use the representation ability of BERT. Bai *et al.* proposed LST approach to transfer knowledge from an LM to an ARM with teacher-student learning [20]. [47] distilled knowledge from BERT to an ARM. In this paper, we transfer the knowledge from BERT to improve the LASO model. LASO and BERT both capture the global token-level relationship. With this method, the LASO model can benefit from the representation ability of BERT but does not add any extra computation during inference.

### C. Cross-Modal Semantic Alignment

The semantic refinement from BERT is also related to cross-modal semantic alignment, i.e., aligning the spaces of LASO and BERT. The concept of cross-modal semantic alignment is first used in cross-modal retrieval [48], [49], [50], [51], [52], [53]. Especially the recent deep learning work learn a shared semantic space between the images and the text-based on neural networks [50], [51], [52], [53]. The basic idea of these work is to map the features from different modalities so that the system can easily compute their similarities. Recent ASR-free approach to text query-based keyword search from a speech also based on this idea, i.e., learning a shared semantic space between text query and speech [54]. The motivation of our proposed semantic refinement from BERT is also to align semantic of LASO and BERT. However, different from the retrieval work, we do not train the models of the two modalities but only train the LASO model. Namely, the BERT model, which has been confirmed as a powerful language model, is an auxiliary model to train the LASO model and is not used during inference. In addition, this work is in a cross-model knowledge transferring setting, i.e., transferring knowledge from text-modal model BERT to speech-modal model LASO<sup>1</sup>.

## VIII. EXPERIMENTAL SETUP

In this section, we introduce the used datasets and the experimental setup. All the experiments are implemented with deep learning toolkit PyTorch [55] with python programming language.

<sup>1</sup>Different from AED models, the inputs of LASO are only speech features but not text embeddings. Therefore, it is a unimodal model.

TABLE I  
THE DESCRIPTION OF THE DATASETS

		#Utter.	#Hours	#Speaker	Duration (Sec.)			#Token Per Sentence		
					Min.	Max.	Avg.	Min.	Max.	Avg.
AISHELL-1	Training	120,098	150	340	1.2	14.5	4.5	1.0	44.0	14.4
	Dev.	14,326	18	40	1.6	12.5	4.5	3.0	35.0	14.3
	Test	7,176	10	20	1.9	14.7	5.0	3.0	21.0	14.0
AISHELL-2	Training (iPhone)	1,009,223	1000	1347	0.5	19.3	3.6	1.0	53.0	10.9
	Dev. (iPhone)	2,500	2	5	1.1	9.4	2.9	1.0	25.0	9.9
	Dev. (Android)	2,500	2	5	1.1	9.4	2.9	1.0	25.0	9.9
	Dev. (HiFi Mic.)	2,500	2	5	1.1	9.4	2.9	1.0	25.0	9.9
	Test (iPhone)	5,000	4	10	1.1	8.5	2.9	1.0	25.0	9.9
	Test (Android)	5,000	4	10	1.1	8.5	2.9	1.0	25.0	9.9
	Test (HiFi Mic.)	5,000	4	10	1.1	8.5	2.9	1.0	25.0	9.9

TABLE II  
THE SYMBOLS OF THE HYPER-PARAMETERS OF THE ARCHITECTURE

Symbol	Description
$D_m$	The dimensionality of the inputs of the multi-head attention.
$D_{in}$	The inner dimensionality of the position-wise FFN.
Activation	The type of activation function of the position-wise FFN.
#Enc.	The number of blocks of the encoder.
#PDS	The number of blocks of PDS.
#Dec.	The number of blocks of the decoder.

### A. Datasets

We conduct experiments on public Chinese speech datasets AISHELL-1<sup>2</sup> [56] and AISHELL-2<sup>3</sup> [57]. These two datasets have different scales of data so that we can evaluate the generalization on both a small dataset and a large dataset. The reason to select these two datasets is that character-based tokenization of the officially pre-trained Google’s BERT for Chinese is proper for the experiment of semantic refinement from BERT<sup>4</sup>.

AISHELL-1 contains 178 hours of Mandarin speech. The speech is recorded by 400 speakers. All audio is recorded with high fidelity microphones in 44.1 kHz, then subsampled to 16 kHz. The content of the datasets covers 5 domains including “Finance”, “Science and Technology”, “Sports”, “Entertainments”, and “News”.

AISHELL-2 contains about 1000 hours of Mandarin Speech for training. The training set is recorded by 1991 speakers with iPhone smartphones. The content covers voice commands, digital sequence, places of interest, entertainment, finance, technology, sports, English spellings, and free speaking without specific topics. The development sets and the test sets are recorded with different equipment to evaluate generalization for different equipment.

The details of the two datasets are shown in Table I.

<sup>2</sup><http://www.openslr.org/33/>

<sup>3</sup>[http://www.aishelltech.com/aishell\\_2](http://www.aishelltech.com/aishell_2)

<sup>4</sup>We did some initial experiments on English dataset LibriSpeech. However, the tokenization method of English BERT makes the vocabulary very large (27563 English word pieces) so that it is not easy for the model to converge. We noticed similar phenomena of MASK-CTC [42] on Latin-alphabet-based experiments, while it performs very well in the Japanese task. In the future, we will try to select proper modeling units for Latin-alphabet-based data and train our own BERT models. We provide discussion on English dataset LibriSpeech in the supplemental materials.

### B. Setup

**Basic settings.** We first evaluate the models on the small-scale (150 hours) dataset AISHELL-1. Then we extend the experiments to the large-scale (1000 hours) dataset AISHELL-2. We use 80-dimension Mel-filter bank features (FBANK) as the inputs, which are extracted every 10ms with 25ms of frame length. For AISHELL-1, the token vocabulary contains 4231 characters in the training set and three special symbols, i.e., “<sos>” for the start of the sentence, “<unk>” for unseen characters, and “<eos>” as the filler of the tail of a token sequence. The vocabulary size for AISHELL-2 is 5252.

**Baseline settings.** We build an autoregressive AED model Speech-Transformer [10], [11] and a non-autoregressive SAN-CTC model [58] as two baselines. The two models use the same settings as LASO, including the above-mentioned features, vocabulary, and basic attention blocks (Fig. 4). Both two models use the whole input signal, which is the same with LASO. Because we would like to compare the performance of SAN-CTC in the end-to-end non-autoregressive setting, the decoding process is greedy, i.e., directly select the most likely token and remove the blanks, without beam-search on an N-gram LM based searching graph.

For Speech-Transformer, both the encoder and the decoder have 6 attention layers. The model dimensionality is 512. The number of heads of the attention is 8. The dimensionality of the intermediate FFN is 2048. And the activation function is GLU. It also uses the same CNN subsampling layer as the LASO models, as shown in Fig. 2. We refer to this model as *Transformer*. For SAN-CTC, because it only has an encoder part, we set the number of the attention layer to 12. The model size is comparable to *Transformer* and other models. We refer to this model as *SAN-CTC*.

**LASO settings.** We compare different architectures of LASO. The symbols of network configuration are shown in Table II. The basic attention blocks are shown in Fig. 4, which are the same as the two baselines. The lengths of positional encodings of the PDS module are set to 60 ( $L$  in Eq. (8)).

**Training settings.** We use Adam algorithm [59] to optimize the models. We use the warm-up learning rate schedule [24]:

$$\alpha = D^{-0.5} \cdot \min(step^{-0.5}, step \cdot warmup^{-1.5}). \quad (13)$$

The warm-up step is set to 12000. The dropout rate is set to 0.1. Each batch contains about 100 seconds of speech, and



we accumulate gradients of 12 steps to simulate a big batch [60] for stabilizing training. We train the models until they converge. The typical number of epochs for LASO is 130. And the typical number of epochs for Speech-Transformer and SAN-CTC is 80.

We use SpecAugment [61] for data augmentation. The frequency masking width is 27. The time masking width is 40. Both frequency masking and time masking are employed twice. But we do not use time warping. We leverage label smoothing with 0.1 for over-confidence problems during training. We average parameters of the models which are saved at the last 10 epochs as the final model.

We use the Google’s pre-trained Chinese BERT model<sup>5</sup> for semantic refinement. This model has 12 transformer layers. The model dimensionality is 768. The total number of parameters is 110M. The vocabulary of the BERT model contains 21128 tokens. During training, the coefficient  $\lambda$  in Eq. (11) is set to 0.005.

**Performance evaluation.** For performance evaluation, we use the standard edit-distance based error rate, i.e. character error rate (CER). For speed evaluation, we use both real-time factor (RTF) and average processing time (APT) [36]. RTF is a standard metric to evaluate the processing time cost of an ASR system. It is the average time cost to process one-second speech:

$$\text{RTF} = \frac{\text{Total Processing Time}}{\text{Total Duration}}. \quad (14)$$

This metric is dimensionless and independent of utterance duration. To consider the impact to the processing time of the utterance duration, we compute APT:

$$\text{APT} = \frac{\text{Total Processing Time}}{\text{Total Number of Utterance}}. \quad (15)$$

The unit of APT is second.

The reason to use ATP is to show the efficiency to process one utterance. It considers the waiting time of a user and is suitable for both online applications (such as speech interaction systems) and offline applications (such as voice document transcription). Specifically, for online applications, the waiting time of a user is “the time between the stop-point at which the user stops speaking and the appearance on the screen of the ASR result”. And for offline applications, the waiting time of a user is “the time between the point at which the user inputs the utterance and the appearance on the screen of the ASR result”. These metrics ignore the uncontrollable factors of a speech engineer, e.g., data transmission speed on the Internet.

RTF may give an *overestimated sense* of whole-utterance ASR systems (AED, Speech-Transformer, or bidirectional AM based hybrid models), because the whole-utterance ASR systems starts processing speech after receiving the whole utterance, unlike some streaming models. For streaming models, the waiting time of a user of an online application can be estimated as

$$\text{RTF} \times \text{the length of the last data package}. \quad (16)$$

However, for whole-utterance ASR systems, the length of whole-utterance is needed to be considered. To address this issue, we compute APT, which directly average the processing time cost of an utterance. We compute these two practical values on commonly used deep learning devices (GPUs) to show the time cost of the ASR systems. We also mention that the time cost of feature extraction is included in our experimental implementation.

## IX. EXPERIMENTAL RESULTS

In this section, we introduce the experimental results.

### A. Comparing Model Architectures on AISHELL-1

First, we comparing the performance of the different model configurations. Table III shows the results of the different configurations. From Table III, we conclude statements as follows.

- 1)  $D_m$ : A large number of the dimensionality of the attention block makes the model have more powerful representation ability and achieve better performance.
- 2) Activation: We find that using GLU as the activation function is more effective than using ReLU consistently. The GLU will introduce more parameters to the model. However, comparing the deeper model with ReLU and the shallower model with GLU, whose model sizes are relatively similar, the model with GLU can achieve better performance.
- 3) #Enc. and #Dec.: Increasing the number of the blocks of the encoder and the decoder improves the performance of the model.
- 4) #PDS: We compare the model with one-layer PDS and the one with two-layer PDS. We find the difference between the two is not very significant.
- 5) Semantic Refinement from BERT: We can see that with semantic refinement from BERT, the performance improvements are significant for the  $D_m = 512$ . However, for the models with  $D_m = 256$ , the performance improvements are subtle. We analyze that it is because the discrepancy of the dimensionality of the LASO and BERT is too large so that it is hard for the LASO models to learn the representations of BERT.

In the rest of the paper, we refer to model 2, model 12, and model 16 in Table III as LASO-small, LASO-middle, and LASO-big, respectively, to compare performance with previous models.

### B. Comparisons with Other Methods on AISHELL-1

We compare our proposed LASO with baseline autoregressive model Transformer and non-autoregressive CTC model SAN-CTC. We also compare it with previous work. The results are shown in Table IV. We can see that our proposed LASO models achieve competitive performance, compared with the hybrid models, the CTC models, and autoregressive transformer models. In particular, with semantic refinement from BERT, LASO-middle and LASO-big outperform the

<sup>5</sup>[https://storage.googleapis.com/bert\\_models/2018\\_11\\_03/chinese\\_L-12\\_H-768\\_A-12.zip](https://storage.googleapis.com/bert_models/2018_11_03/chinese_L-12_H-768_A-12.zip)

<sup>7</sup><https://github.com/kaldi-asr/kaldi/blob/master/egs/aishell/s5/RESULTS>

TABLE III  
THE CHARACTER ERROR RATES ON AISHELL-1 WITH DIFFERENT HYPER-PARAMETERS

Model	Dm	Din	Activation	#Enc.	#PDS	#Dec.	Model size	w/o BERT		w/ BERT	
								CERs on Dev.	CERs on Test	CERs on Dev.	CERs on Test
1	256	2048	ReLU	4	1	4	15.9M	7.9	8.8	7.7	8.6
2	256	2048	GLU	4	1	4	20.6M	7.1	8.1	7.0	7.8
3	256	2048	ReLU	6	1	6	21.2M	7.2	8.2	7.1	8.0
4	256	2048	GLU	6	1	6	28.0M	6.6	7.5	6.6	7.4
5	256	2048	ReLU	6	2	6	22.8M	7.4	8.3	7.3	8.1
6	256	2048	GLU	6	2	6	30.1M	6.6	7.5	6.3	7.1
7	256	2048	ReLU	8	2	6	25.4M	7.3	8.4	7.3	8.1
8	256	2048	GLU	8	2	6	33.8M	6.7	7.5	6.5	7.4
9	512	2048	ReLU	4	1	4	37.0M	7.0	7.8	6.6	7.4
10	512	2048	GLU	4	1	4	46.5M	6.4	7.4	5.9	6.7
11	512	2048	ReLU	6	1	6	49.6M	6.6	7.5	6.0	6.7
12	512	2048	GLU	6	1	6	63.3M	6.2	7.0	5.4	6.2
13	512	2048	ReLU	6	2	6	53.9M	6.6	7.5	6.6	7.5
14	512	2048	GLU	6	2	6	68.6M	6.1	6.9	5.4	6.1
15	512	2048	ReLU	8	2	6	60.2M	6.6	7.5	6.2	6.9
16	512	2048	GLU	8	2	6	80.0M	5.9	6.6	5.2	5.8
17	768	2048	ReLU	6	1	6	85.5M	6.5	7.3	5.9	6.6
18	768	2048	GLU	6	1	6	105.8M	5.9	6.9	5.3	6.1

TABLE IV  
COMPARISONS WITH BASELINES ON AISHELL-1

Model	#Param.	CERs		RTF / APT
		Dev.	Test	
KALDI(nnet3) * † ‡	-	-	8.6	-
KALDI(chain) * † ‡	-	-	7.4	-
LAS [62]	-	9.4	10.6	-
ESPNet (Transformer) † ‡ [63]	-	6.0	6.7	-
A-FMLM [41]	-	6.2	6.7	-
Fan et al. (Transformer) [64]	-	-	6.7	-
AGS CTC ‡ [65]	-	7.0	7.9	-
Transformer (baseline1)	67.5M	6.1	6.6	0.19 / 961ms
SAN-CTC (baseline2)	56.4M	7.2	7.8	0.0033 / 16ms
LASO-small	20.6M	7.1	8.1	0.0027 / 13ms
LASO-small w/ BERT	20.6M	7.0	7.8	0.0027 / 13ms
LASO-middle	63.3M	6.2	7.0	0.0035 / 17ms
LASO-middle w/ BERT	63.3M	5.4	6.2	0.0035 / 17ms
LASO-big	80.0M	5.9	6.6	0.0040 / 20ms
LASO-big w/ BERT	80.0M	<b>5.2</b>	<b>5.8</b>	0.0040 / 20ms

\* from the KALDI official repository<sup>6</sup>.

† with speed perturbation based data augmentation.

‡ with an extra language model at the inference stage.

autoregressive Transformer model. Moreover, the inference speed of the non-autoregressive models is much faster than the autoregressive models.

We implemented a competitive baseline autoregressive model Transformer, which achieves a CER of 6.6% on the test set of AISHELL-1. The performance of LASO-middle, which has a similar scale of the number of the parameters with Transformer, is comparable to Transformer. And the bigger model LASO-big outperforms Transformer. Both LASO-middle and LASO-big outperform the CTC based non-autoregressive model SAN-CTC.

Table IV also lists RTF and APT of the models. RTF is the ratio of the total inference time to the total duration of

the test set. APT is the averaged time cost for decoding one utterance (including the time of feature extraction) on the test set. The inference is done utterance by utterance on an NVIDIA RTX 2080Ti GPU. We can see that the inference speed of the non-autoregressive models is much faster than the autoregressive transformer model Transformer. The APT of the non-autoregressive models is reduced by 50×. And we can see that even the model size of LASO-big is larger than Transformer, the APT of LASO-big is much reduced. With the non-autoregressive and the feedforward structure, the inference of the model can be implemented very efficiently.

The inference speed of the baseline non-autoregressive model SAN-CTC is also very fast. However, compared with the similar scale LASO models, the performance is degraded. We analyze that the LASO models use the decoder, which plays the role of an autoencoder LM, to capture language semantics more effectively. In addition, for CTC models, the blank symbols inserted in the token sequence may influence the model to capture language semantics.

For LASO-middle and LASO-big, semantic refinement from BERT much reduces CERs by 11% to 12% on the test set. This demonstrates that teacher-student learning with BERT can improve the ability of the LASO to capture the language semantics. However, for the small-size model LASO-small semantic refinement from BERT does not improve the performance. We analyze that the dimensionality of the representations of LASO-small and BERT are very different (256 vs. 768), which influences the LASO model to learn knowledge from the BERT model.

### C. Comparisons with Other Methods on AISHELL-2

We then extend the experiments to the larger scale dataset AISHELL-2. AISHELL-2 contains about 1000 hours of training data. And the covered topics of AISHELL-2 are more diverse than AISHELL-1. Furthermore, the training set is

<sup>7</sup><https://github.com/espnet/espnet/blob/master/egs/aishell2/asr1/RESULTS.md>

TABLE V  
COMPARISONS WITH BASELINES ON AISHELL-2

Model	#Param.	CERs							
		Dev. (iPhone)	Dev. (Android)	Dev. (HiFi Mic.)	Dev. (Avg.)	Test (iPhone)	Test (Android)	Test (HiFi Mic.)	Test (Avg.)
KALDI (chain) [57] † ‡	-	9.1	10.4	11.8	10.4	8.8	9.6	10.9	9.8
LAS [62]	-	-	-	-	-	9.2	9.7	10.3	9.7
ESPNet (transformer) * † ‡	-	-	-	-	-	7.5	8.9	8.6	8.3
Transformer (baseline1)	67.5M	6.4	7.2	7.7	7.1	7.1	8.0	8.2	7.8
SAN-CTC (baseline2)	56.4M	8.3	8.9	8.8	8.6	8.0	9.0	8.9	8.7
LASO-small	20.6M	8.2	9.5	9.7	9.1	8.5	9.5	9.5	9.2
LASO-small w/ BERT	20.6M	8.9	10.0	10.3	9.7	8.8	9.8	10.5	9.7
LASO-middle	63.3M	6.6	7.5	7.6	7.2	6.8	7.4	7.3	7.2
LASO-middle w/ BERT	63.3M	6.5	7.2	7.4	7.0	6.6	<b>7.2</b>	<b>7.1</b>	7.0
LASO-large	80.0M	6.4	7.3	7.3	7.0	6.7	7.4	7.4	7.1
LASO-large w/ BERT	80.0M	<b>6.2</b>	<b>7.2</b>	<b>7.3</b>	<b>6.9</b>	<b>6.5</b>	<b>7.2</b>	<b>7.1</b>	<b>6.9</b>

\* from the ESPnet official repository<sup>7</sup>.

† with speed perturbation based data augmentation.

‡ with an extra language model at the inference stage.

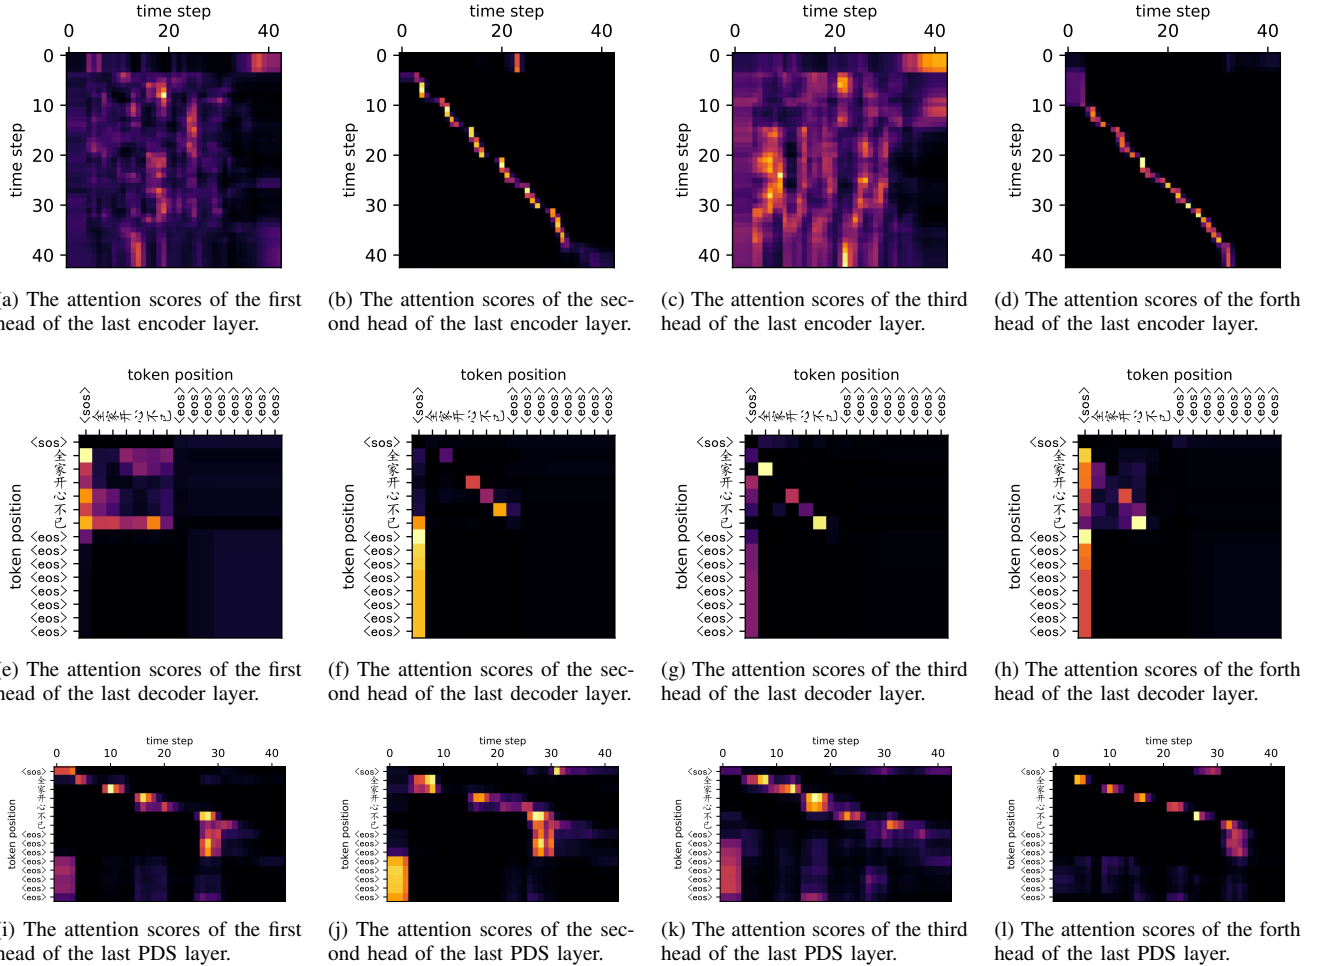


Fig. 7. The visualization of the attention scores of the model LASO-big. Here, we show the first four heads of the last layer of each module. And because the token sequence is long to visualize, we truncate the first 15 tokens. All visualization is listed in the supplemental materials.

recorded with iPhone, and the test sets of AISHELL-2 cover three different channels, i.e., iPhone, Android smartphones, and Hi-Fi microphones. So we can evaluate the generalization of the models on AISHELL-2 more in detail. Because AISHELL-2 is larger than AISHELL-1 and training models on AISHELL-2 costs much more time, we directly use the se-

lected architectures from previous experiments on this dataset.

The experimental results are shown in Table V. We can see that the proposed LASO models also achieve a promising performance. Compared with the hybrid systems, LAS systems, and the CTC model, all LASO models achieve better performance. And with the similar and larger scale model

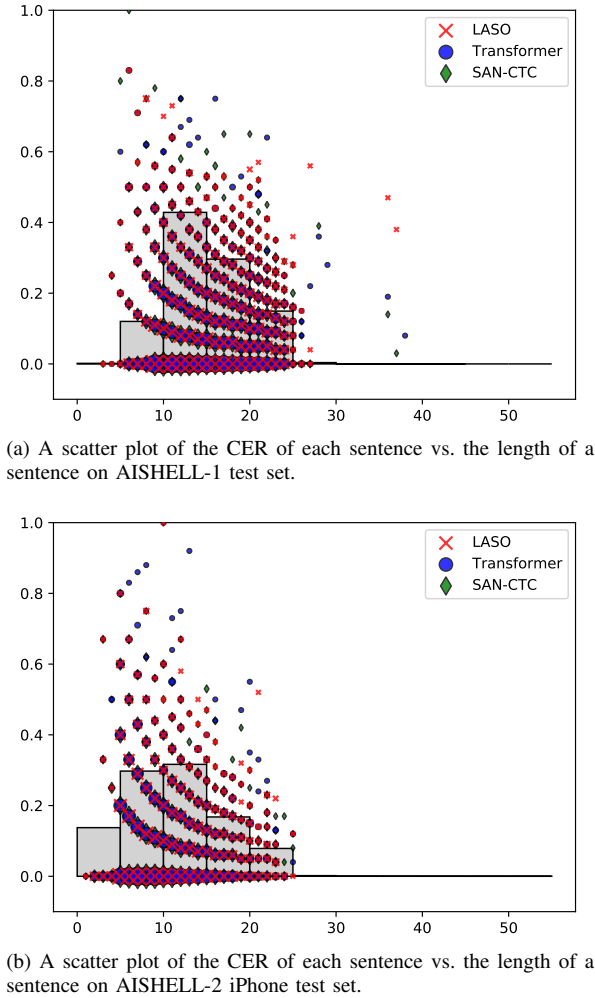


Fig. 8. Scatter plots for analyzing the impact of the lengths of the sentences. The grey histograms represent the ratios of lengths of the sentence (i.e., the number of the tokens) on the training set. The area of a scatter shows the number of the examples at that point. We use baseline1 Transformer, baseline2 SAN-CTC, and LASO-middle since they have similar model sizes. From the plots, we can see no significant difference between the three models. The three models have the same CER for some sentences so that the scatters are overlapped. Note that an outlier in the second figure which CER is 100% since the reference of this sentence is wrong.

sizes, LASO-middle and LASO-big outperform previous transformer models. We find that with more data and a larger model, LASO can achieve better performance than the well-trained transformer models. With semantic refinement from BERT, the performances are further improved. However, the improvements are not significant like the experiments on AISHELL-1. And for LASO-small, the performance degrades. The possible reasons are 1) the model capacity is very large between the small model LASO-small and the big model BERT; 2) AISHELL-2 has more data and is more complex than AISHELL-1, which influences the effectiveness of knowledge distillation [66], [67].

## X. DISCUSSION

In this section, we analyze the attention patterns and discuss the impact of the sentence lengths. And we show some decoded examples of the models.

### A. Visualization of Attention Patterns

To better understand the behaviors of the LASO model, we visualize the attention scores of an utterance with LASO-big. We show the first four heads of attention scores of the last layers of the encoder, the PDS, and the decoder. The detailed visualization is listed in the supplemental materials. Fig. 7 shows the visualization results. We summarize the observations as follows.

- 1) Different heads in one layer have different attention patterns. This implies that one representation of the sequence in different head attends different representations.
- 2) For the encoder, some attention patterns show the top-left to bottom-right alignments (Fig. 7b and Fig. 7d). This meets the expectation: the matched representations are around the corresponding representation in the speech sequence. And some heads do not show obvious patterns.
- 3) For the decoder, we can see that some hidden representation at a token position attends the hidden representation at the previous token position (Fig. 7g), and some hidden representation at a token position attends the hidden representation at the next token position (Fig. 7f). Most hidden representations corresponding to filler tokens  $\langle \text{eos} \rangle$  attend the hidden representations corresponding to  $\langle \text{sos} \rangle$ .
- 4) For the PDS, the position corresponding to a token attends a small range of representations, and the overall patterns are from the top-left to the bottom-right (Fig. 7i, Fig. 7j, Fig. 7k, and Fig. 7l). For the first several positions corresponding to the filler  $\langle \text{eos} \rangle$  tokens, the attention pattern is like a vertical line. However, the lines are corresponding to different positions of the speech in different head (Fig. 7j and Fig. 7l). And most other  $\langle \text{eos} \rangle$  tokens attend the first several representations. We analyze that these representations denote fillers.

From the above observations, we conclude that 1) different attention patterns make the model fuse the representations from various aspects; 2) the attention mechanism can learn meaningful alignments in terms of the positional encodings; 3) the two special filler tokens  $\langle \text{eos} \rangle$  and  $\langle \text{sos} \rangle$  absorb the meaningless representations from the encoder; 4) specific attention patterns exist in the different heads of the decoder. These demonstrate that the PDS module attends specific acoustic representations based on positions and the decoder captures the token relationship based on the self-attention mechanism.

### B. The Impact of the Lengths of Sentences

Position parameters exist in the PDS module. This may cause that the performance relies on the length of a sentence. To check this point, we plot scatters of the CER of each sentence vs. the length of a sentence in Fig. 8. We can see no significant difference among the three models.

Fig. 8 also provides histograms of ratios of the different lengths of the training set. We can see that the distributions approximate Gaussian distribution as expected. And the trend of CERs is similar to the length distribution of the training set, i.e., the CERs of more sentences are zero in the middle

part of the figures. This is because the models are trained with more data with the middle lengths. All three models can recognize the sentences with unseen lengths (or few-shot lengths) in the training set. However, the error rates are relatively higher than the sentences with seen lengths. This phenomenon is as expected. Because the three models (Transformer, SAN-CTC, LASO) are all whole-utterance ASR models, the length of the utterance is an implicit factor to train the models. This is different from the local-window models based conventional hybrid models, i.e., time-delay neural networks (TDNN), CNNs, or latency-control BLSTM. LASO is more sensitive with the unseen lengths than Transformer and SAN-CTC. Because the PDS module needs to be trained with various lengths.

In engineering practice, two methods can be applied to address the unseen-length issue of the whole-utterance models: 1) select data with various lengths to train the models; 2) using a voice activity detection (VAD) system to cut long utterances into segments.

### C. Examples of Semantic Refinement from BERT

We show some recognized results of the models with and without semantic refinement from BERT in Table VI. The recognized results are from AISHELL-1 test set. The model architecture is LASO-big. The grey background is the wrongly recognized words. And the red words are correctly recognized by the model with semantic refinement from BERT.

From Table VI, we can see that in terms of the semantic of the words, the model with semantic refinement from BERT correctly recognized the words which were not recognized by the model without semantic refinement from BERT. Specifically, for case 1 to case 4, the model w/o BERT recognized the word as another word which has similar pronunciation but mismatched meaning (e.g., “一场”, “骤”). Case 5 is an interesting example. The model w/o BERT recognized a word as “爷泳”, but it is wrong in grammar. The results of the model w/ BERT at the corresponding position is “爷爷” which is a correct word in grammar, but it does not match the reference. These results show that the proposed semantic refinement from BERT impacts the model in semantic and can improve semantic representation performance for these examples.

## XI. CONCLUSIONS AND FUTURE WORK

This paper proposes a feedforward neural network based non-autoregressive speech recognition model called LASO. The model consists of an encoder, a position dependent summarizer, and a decoder. The encoder encodes the acoustic feature sequence into high-level representations. The PDS converts the acoustic representation sequence to the token-level sequence. And the decoder further captures the token-level relationship. Because the prediction of each token does not rely on other tokens, and the whole model is feedforward, the parallelization of the whole sentence prediction is realizable. Thus, the inference speed is much improved, compared with the autoregressive attention-based end-to-end models. Furthermore, we propose to refine semantics from a large-scale pre-trained language model BERT to improve the performance. Experimental results show that LASO can

TABLE VI  
EXAMPLES OF RECOGNIZED RESULTS

1	Reference	而二零零八年举办夏季奥运会所留下的宝贵遗产	/yi2 chan3/
	w/o BERT	而二零零八年举办夏季奥运会所留下的宝贵一场	/yi4 chang3/
	w/ BERT	而二零零八年举办夏季奥运会所留下的宝贵遗产	/yi2 chan3/
2	Reference	当月住宅类商品房成交套数 骤 跌	/zhou4/
	w/o BERT	当月住宅类商品房成交套数 周 跌	/zhou1/
	w/ BERT	当月住宅类商品房成交套数 骤 跌	/zhou4/
3	Reference	数十名市民赶到越秀区 一 酒 家维权	/yi4/ /jia1 wei2 quan2/
	w/o BERT	数十名市民赶到越秀区 以 酒 酒未钱	/yi3/ /jiu3 wei4 qian2/
	w/ BERT	数十名市民赶到越秀区 一 酒 家维权	/yi4/ /jia1 wei2 quan2/
4	Reference	尽管她努力 瘦身	/shou4 shen1/
	w/o BERT	尽管她努力 受存	/shou4 cun2/
	w/ BERT	尽管她努力 瘦身	/shou4 shen1/
5	Reference	圆 圆 的脸蛋非常的可爱	/yuan2 yuan2/
	w/o BERT	爷泳 的脸蛋非常的可爱	/ye2 yong3/
	w/ BERT	爷爷 的脸蛋非常的可爱	/ye2 ye5/

/-/ is pinyin for labeling the pronunciation of the Chinese characters. The number is the tone of the character.

achieve a competitive performance and high efficiency. In the future, we will improve the performance of LASO by the architecture and loss functions. We find that LASO is trained with more epochs during training. We will try to find strategies to speed up training. And we will try to find more proper modeling units for Latin-alphabet-based data.

## XII. ACKNOWLEDGMENT

The authors are grateful to the anonymous reviewers for their invaluable comments that improve the completeness and readability of this paper.

## REFERENCES

- [1] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, “Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups,” *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [2] D. Yu and L. Deng, *AUTOMATIC SPEECH RECOGNITION*. Springer, 2016.
- [3] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, “Attention-based models for speech recognition,” in *Advances in neural information processing systems*, 2015, pp. 577–585.
- [4] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio, “End-to-end attention-based large vocabulary speech recognition,” *international conference on acoustics, speech, and signal processing*, pp. 4945–4949, 2016.
- [5] 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)*. IEEE, 2016, pp. 4960–4964.
- [6] S. Kim, T. Hori, and S. Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4835–4839.



- [7] A. Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [8] C.-C. Chiu, T. N. Sainath, Y. Wu, R. Prabhavalkar, P. Nguyen, Z. Chen, A. Kannan, R. J. Weiss, K. Rao, E. Gonina *et al.*, "State-of-the-art speech recognition with sequence-to-sequence models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4774–4778.
- [9] C. Lüscher, E. Beck, K. Irie, M. Kitza, W. Michel, A. Zeyer, R. Schlüter, and H. Ney, "Rwth asr systems for librispeech: Hybrid vs attention," *Proc. Interspeech 2019*, pp. 231–235, 2019.
- [10] L. Dong, S. Xu, and B. Xu, "Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5884–5888.
- [11] S. Zhou, L. Dong, S. Xu, and B. Xu, "Syllable-based sequence-to-sequence speech recognition with the transformer in mandarin chinese," in *Proc. Interspeech 2018*, 2018, pp. 791–795. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1107>
- [12] C. Raffel, M.-T. Luong, P. J. Liu, R. J. Weiss, and D. Eck, "Online and linear-time attention by enforcing monotonic alignments," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2837–2846.
- [13] C.-C. Chiu and C. Raffel, "Monotonic chunkwise attention," 2018. [Online]. Available: <https://openreview.net/pdf?id=Hko85pICW>
- [14] N. Moritz, T. Hori, and J. Le Roux, "Triggered attention for end-to-end speech recognition," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5666–5670.
- [15] T. Hori, N. Moritz, C. Hori, and J. L. Roux, "Transformer-Based Long-Context End-to-End Speech Recognition," in *Proc. Interspeech 2020*, 2020, pp. 5011–5015. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2928>
- [16] Z. Tian, J. Yi, J. Tao, Y. Bai, and Z. Wen, "Self-attention transducers for end-to-end speech recognition," *Proc. Interspeech 2019*, pp. 4395–4399, 2019.
- [17] Z. Tian, J. Yi, Y. Bai, J. Tao, S. Zhang, and Z. Wen, "Synchronous transformers for end-to-end speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7884–7888.
- [18] Y.-A. Chung and J. Glass, "Speech2vec: A sequence-to-sequence framework for learning word embeddings from speech," in *Proc. Interspeech 2018*, 2018, pp. 811–815. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-2341>
- [19] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," Jun. 2019.
- [20] Y. Bai, J. Yi, J. Tao, Z. Tian, and Z. Wen, "Learn spelling from teachers: Transferring knowledge from language models to sequence-to-sequence speech recognition," *Proc. Interspeech 2019*, pp. 3795–3799, 2019.
- [21] Y. Bai, J. Yi, J. Tao, Z. Tian, Z. Wen, and S. Zhang, "Listen attentively, and spell once: Whole sentence generation via a non-autoregressive architecture for low-latency speech recognition," *Proc. Interspeech 2020*, 2020.
- [22] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 369–376.
- [23] Y. Miao, M. Gowayyed, and F. Metze, "Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 167–174.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [25] T. Q. Nguyen and J. Salazar, "Transformers without tears: Improving the normalization of self-attention," *arXiv preprint arXiv:1910.05895*, 2019.
- [26] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 933–941.
- [27] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [28] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [29] J. You, J. Leskovec, K. He, and S. Xie, "Graph structure of neural networks," in *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, ser. Proceedings of Machine Learning Research, vol. 119. PMLR, 2020, pp. 10881–10891. [Online]. Available: <http://proceedings.mlr.press/v119/you20b.html>
- [30] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, ser. Proceedings of Machine Learning Research, D. Precup and Y. W. Teh, Eds., vol. 70. PMLR, 2017, pp. 1263–1272. [Online]. Available: <http://proceedings.mlr.press/v70/gilmer17a.html>
- [31] Q. Guo, X. Qiu, P. Liu, Y. Shao, X. Xue, and Z. Zhang, "Star-transformer," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 1315–1325.
- [32] C. Bucilua, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2006, pp. 535–541.
- [33] J. Li, R. Zhao, J.-T. Huang, and Y. Gong, "Learning small-size dnn with output-distribution-based criteria," in *Fifteenth annual conference of the international speech communication association*, 2014.
- [34] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.
- [35] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "Fitnets: Hints for thin deep nets," *arXiv preprint arXiv:1412.6550*, 2014.
- [36] J. Gu, J. Bradbury, C. Xiong, V. O. Li, and R. Socher, "Non-autoregressive neural machine translation," in *International Conference on Learning Representations*, 2018.
- [37] Y. Wang, F. Tian, D. He, T. Qin, C. Zhai, and T.-Y. Liu, "Non-autoregressive machine translation with auxiliary regularization," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 5377–5384.
- [38] J. Guo, X. Tan, D. He, T. Qin, L. Xu, and T.-Y. Liu, "Non-autoregressive neural machine translation with enhanced decoder input," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 3723–3730.
- [39] J. Lee, E. Mansimov, and K. Cho, "Deterministic non-autoregressive neural sequence modeling by iterative refinement," 2018.
- [40] X. Ma, C. Zhou, X. Li, G. Neubig, and E. Hovy, "Flowseq: Non-autoregressive conditional sequence generation with generative flow," pp. 4273–4283, 2019.
- [41] N. Chen, S. Watanabe, J. Villalba, and N. Dehak, "Listen and fill in the missing letters: Non-autoregressive transformer for speech recognition," *arXiv preprint arXiv:1911.04908*, 2019.
- [42] Y. Higuchi, S. Watanabe, N. Chen, T. Ogawa, and T. Kobayashi, "Mask CTC: Non-Autoregressive End-to-End ASR with CTC and Mask Predict," in *Proc. Interspeech 2020*, 2020, pp. 3655–3659. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2404>
- [43] C. Gulcehre, O. Firat, K. Xu, K. Cho, L. Barrault, H. Lin, F. Bougares, H. Schwenk, and Y. Bengio, "On using monolingual corpora in neural machine translation," *arXiv: Computation and Language*, 2015.
- [44] A. Sriram, H. Jun, S. Satheesh, and A. Coates, "Cold fusion: Training seq2seq models together with language models," in *Proc. Interspeech 2018*, 2018, pp. 387–391. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1392>
- [45] M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, "Deep contextualized word representations," Jun. 2018.
- [46] J. Shin, Y. Lee, and K. Jung, "Effective sentence scoring method using bert for speech recognition," in *Asian Conference on Machine Learning*, 2019, pp. 1081–1093.
- [47] H. Futami, H. Inaguma, S. Ueno, M. Mimura, S. Sakai, and T. Kawahara, "Distilling the knowledge of bert for sequence-to-sequence asr," *Proc. Interspeech 2020*, 2020.
- [48] V. Lavrenko, R. Manmatha, J. Jeon *et al.*, "A model for learning the semantics of pictures," in *Nips*, vol. 1, no. 2. Citeseer, 2003.
- [49] J. Jeon, V. Lavrenko, and R. Manmatha, "Automatic image annotation and retrieval using cross-media relevance models," in *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval*, 2003, pp. 119–126.
- [50] Q.-Y. Jiang and W.-J. Li, "Deep cross-modal hashing," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 3232–3240.

- [51] M. Fan, W. Wang, P. Dong, L. Han, R. Wang, and G. Li, "Cross-media retrieval by learning rich semantic embeddings of multimedia," in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 1698–1706.
- [52] L. Zhen, P. Hu, X. Wang, and D. Peng, "Deep supervised cross-modal retrieval," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10394–10403.
- [53] Z. Yang, Z. Lin, P. Kang, J. Lv, Q. Li, and W. Liu, "Learning shared semantic space with correlation alignment for cross-modal event retrieval," *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, vol. 16, no. 1, pp. 1–22, 2020.
- [54] K. Audhkhasi, A. Rosenberg, A. Sethy, B. Ramabhadran, and B. Kingsbury, "End-to-end asr-free keyword search from speech," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1351–1359, 2017.
- [55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019, pp. 8026–8037. [Online]. Available: <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>
- [56] 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)*. IEEE, 2017, pp. 1–5.
- [57] J. Du, X. Na, X. Liu, and H. Bu, "Aishell-2: transforming mandarin asr research into industrial scale," *arXiv preprint arXiv:1808.10583*, 2018.
- [58] J. Salazar, K. Kirchhoff, and Z. Huang, "Self-attention networks for connectionist temporal classification in speech recognition," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 7115–7119.
- [59] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [60] M. Ott, S. Edunov, D. Grangier, and M. Auli, "Scaling neural machine translation," pp. 1–9, Oct. 2018. [Online]. Available: <https://www.aclweb.org/anthology/W18-6301>
- [61] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition," *Proc. Interspeech 2019*, pp. 2613–2617, 2019.
- [62] S. Sun, P. Guo, L. Xie, and M. Hwang, "Adversarial regularization for attention based end-to-end robust speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1826–1838, 2019.
- [63] S. Karita, N. Chen, T. Hayashi, T. Hori, H. Inaguma, Z. Jiang, M. Someki, N. E. Y. Soplin, R. Yamamoto, X. Wang, S. Watanabe, T. Yoshimura, and W. Zhang, "A comparative study on transformer vs rnn in speech applications," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, Dec 2019, pp. 449–456.
- [64] Z. Fan, S. Zhou, and B. Xu, "Unsupervised pre-training for sequence to sequence speech recognition," *arXiv preprint arXiv:1910.12418*, 2019.
- [65] F. Ding, W. Guo, L. Dai, and J. Du, "Attention-based gated scaling adaptive acoustic model for ctc-based speech recognition," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7404–7408.
- [66] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *International Conference on Learning Representations*, 2017.
- [67] J. H. Cho and B. Hariharan, "On the efficacy of knowledge distillation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 4794–4802.