

PERSONALIZATION OF CTC SPEECH RECOGNITION MODELS

*Saket Dingliwal** *Monica Sunkara** *Srikanth Ronanki*
Jeff Farris *Katrin Kirchhoff* *Sravan Bodapati*

Amazon AWS AI
{skdin, sunkaral}@amazon.com

ABSTRACT

End-to-end speech recognition models trained using joint Connectionist Temporal Classification (CTC)-Attention loss have gained popularity recently. In these models, a non-autoregressive CTC decoder is often used at inference time due to its speed and simplicity. However, such models are hard to personalize because of their conditional independence assumption that prevents output tokens from previous time steps to influence future predictions. To tackle this, we propose a novel two-way approach that first biases the encoder with attention over a predefined list of rare long-tail and out-of-vocabulary (OOV) words and then uses dynamic boosting and phone alignment network during decoding to further bias the subword predictions. We evaluate our approach on open-source VoxPopuli and in-house medical datasets to showcase a 60% improvement in F1 score on domain-specific rare words over a strong CTC baseline.

Index Terms: CTC, personalization, long-tail recognition, contextual biasing.

1. INTRODUCTION

End-to-end (E2E) techniques such as CTC [1, 2], Transducer [3, 4] and Attention [5, 6] have been successfully deployed for training large vocabulary speech recognition applications. Among all the techniques, single-pass CTC inference is an attractive choice for industry applications due to its low speed and memory footprint [7, 8, 9]. Most recently, Conformer for speech recognition proposed in [10] combines Transformer and Convolution modules into “Conformer” blocks. This, in combination with CTC-Attention multi-task learning framework [11, 12] has become one of the popular recipes for training end-to-end speech recognition models [13]. While these models have achieved results on par with or better than conventional hybrid ASR systems [14, 15, 16], they still struggle to recognize rare domain specific words and OOV words that are unseen during training [17, 18]. We therefore propose the biasing approaches for personalization of CTC models with a predefined list of uncommon but important entities which typically include contact names of users, names of clients for a company, medical terms or airport names.

1.1. Relation to prior work in personalization

Several approaches have been proposed in the past for personalization, domain adaptation and contextual biasing of ASR systems. While some focused on adapting Language Models (LMs) in hybrid systems [19, 20, 21], others customized autoregressive architectures like Transducer and Attention-based encoder-decoder models [22, 23, 24, 25, 26]. We categorize these methods into (1) Early biasing, and (2) Decoder biasing. Methods in the former category focus

on directly improving the likelihood of predicting custom words by training additional modules along with the base model while techniques in the latter tries to bias the decoding process with additional context as a post-processing step. Early biasing methods like training an attention module over custom list of user-specific words [25, 27, 28], that induce a better coverage of rare WordPieces have been shown to significantly benefit autoregressive models. In parallel, post-processing methods like shallow fusion [29], deep fusion [23], leveraging grapheme-to-grapheme (G2G) [30] to produce additional pronunciation variants have also been shown to be useful for their simplicity.

While some of these approaches are beneficial, accurate recognition of rare and OOV words still presents a challenge to CTC models because of two main reasons: First, CTC models can not influence future predictions based on previous output tokens, and hence cannot use the contextual predictions to customize beam paths in the early biasing methods. Second, the outputs of CTC include blanks and duplicate tokens which makes decoder biasing further complex. Additionally, the output units of CTC (or any other E2E) models are typically either graphemes or subwords, whose sequences do not correspond well to how they are pronounced [23].

1.2. Novelty of this work

In this work, we propose a novel two-phase approach which includes both early and decoder biasing techniques to address the challenges that are typical to personalization of CTC models. In the first phase, we use an attention-based Adapter module that uses speech representations from the encoder to attend over a list of custom entity words, thereby learning when to copy these entities to the beam path. In the second phase, we propose a combination of techniques to further boost these entities. Specifically, we propose on-the-fly dynamic boosting of subwords in combination with a prefix beam search decoder. In addition to this, we also adopt a phone alignment network and bias the subword predictions at inference time by utilizing both phone and subword output predictions. The use of phone predictions alleviates the problem of misrecognition of words whose pronunciations do not match a corresponding subword sequence.

Overall, we showcase that our biasing approaches can significantly boost the performance of rare domain specific words and OOV words. Along with the improvements, the following characteristics make our joint model an important scientific and practical contribution. (1) Similar to [27, 28] and different from [25], our methods do not require any additional training data for personalization and the Adapter modules can be trained using the same data used for training the base model; (2) The parameters of the base model are not finetuned during personalization and hence the same model can be used on generic speech datasets with negligible effect; (3) Similar to [27, 25], the list of entities to the Adapter module can be changed at inference time as in a zero-shot setting and a common architecture

* Equal contribution

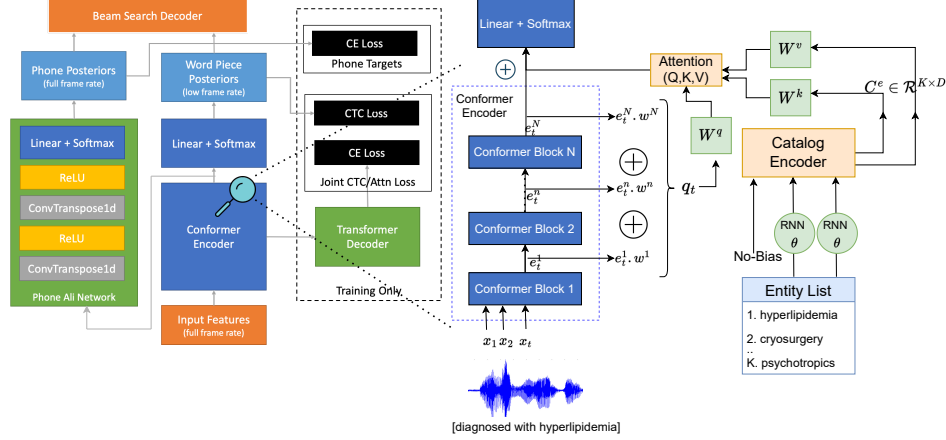


Fig. 1: Multi-task encoder for Conformer CTC model (left) with design for the Contextual Adapter (right) added to the encoder layers

fits all personalization use cases without the need for user or domain specific additional parameters; (4) We show that our joint approach can yield significant improvements regardless of the size of the base CTC model as well as the number of hours used for training; (5) We also provide analysis on the effect of biasing when the model is fed with a larger list of custom entities (more than 700). The main contributions of our work are summarized as follows:

- We propose an early biasing approach for Conformer-CTC architecture that leverages representations from multiple encoder layers to copy custom entities to the beam path.
- We also propose a combination of on-the-fly decoder biasing techniques for personalization of non-autoregressive CTC models during beam search decoding.
- We demonstrate how our proposed approaches work complementary to each other through extensive experimentation and ablation studies. On VoxPopuli dataset, our proposed joint model achieves 61.7% and 31.4% improvement in F1 scores of rare and OOV words respectively over the baseline CTC model.

2. MODEL ARCHITECTURE

2.1. Overview of Conformer CTC-Attention framework

The CTC-Attention framework [11], can be broken down into three different components: *Shared Encoder*, *CTC Decoder* and *Attention Decoder*. As shown in Figure 1, our *Shared Encoder* consists of multiple Conformer [10] blocks with context spanning a full utterance. Each Conformer block consists of two feed-forward modules that sandwich the Multi-Headed Self-Attention module and the Convolution module [10]. A Convolution frontend with Conv2D layers downsamples the input features by a factor of four before feeding them to a Conformer Encoder. The *CTC Decoder* consists of a linear layer with softmax activation, which transforms the *Shared Encoder* output to the final output distribution over the WordPiece units. The *Attention Decoder* consists of one Transformer decoder block and a softmax layer whose outputs are compared against reference outputs using a cross-entropy criterion. The intuition behind this joint modeling is to use a CTC objective function as an auxiliary task to train the *Shared Encoder* within the multitask learning (MTL) framework whose loss is given below:

$$L_{MTL} = \alpha L_{CTC} + (1 - \alpha) L_{Att} \quad (1)$$

where L_{MTL} , L_{CTC} , and L_{Att} represents multi-task learning loss, CTC-loss and *Attention Decoder* cross-entropy loss respectively. α

is a tunable parameter and represents the weight of CTC-loss whose value lies between 0 and 1. During inference time, only the *Shared Conformer Encoder* and linear layer is used and the *Attention Decoder* is discarded to ensure a low inference latency. Although *Attention Decoder* can be used in the second pass to give a more accurate result, we did not perform any experiments in this direction and will be explored as part of future work.

2.2. Phone alignment network with Conformer encoder

Typically, the CTC decoder outputs (subword units and blank symbols) are collapsed and concatenated to form final word-level output. However, accurate recognition of rare and OOV words from subword units is a challenging problem and can be complemented with a phone recognizer based on the work presented in [31]. During training, a joint phone alignment network as shown in Figure 1 is trained from *Shared Encoder* representations. This requires that we have phone alignments per frame derived from a hybrid system and a pronunciation lexicon. Since the Conformer architecture is low-frame-rate and the phone alignments are full-frame-rate, we use upsampling Convolutional layers (ConvTranspose1d) to transform the low-frame-rate encodings up to full-frame-rate. Per frame phone posteriors predicted from this network are optimized using a cross-entropy loss function and modified MTL loss is shown below:

$$L_{MTL} = \alpha((1 - \beta)L_{CTC} + \beta L_{Ali}) + (1 - \alpha)L_{Att} \quad (2)$$

where L_{Ali} represents alignment loss and α, β are tunable parameters and represent the weights of CTC-loss and alignment-loss respectively, whose value also lies between 0 and 1. If we already have a well trained CTC model, we could also directly add the phone recognizer and train it by freezing the *Shared Encoder*. In Section 3.2.2, we present further details on biasing the subword predictions by utilizing the phone alignment network.

3. PROPOSED APPROACH

3.1. Early Encoder Biasing

As shown in Figure 1, we propose to use Contextual Adapter to bias our pre-trained Conformer-CTC encoder. Similar to [25], our Contextual Adapter module consists of two components: Catalog Encoder and Biasing Adapter. The Catalog Encoder encodes a list of

custom entities to produce corresponding embeddings while the Biasing Adapter uses output representations from the Conformer Encoder blocks to attend over these entities and learns to copy them to the output, when observed in the input speech. Figure 1 summarizes the architecture we use for Contextual Adapters. Assume we have a set of $(K - 1)$ custom entity words i.e. $\{w^j\}_{j=1}^{j=(K-1)}$, we tokenize each word with the same tokenizer as the base CTC model, to form a sequence of tokens $\{c_1^j, c_2^j, \dots, c_n^j\}$. We then do an embedding lookup (parameterized by ϕ) to create a sequence of vectors $\{c_{1:n}^j\}^\phi$. We pass these sequences of subword embeddings through BiLSTM layers (parameterized by ψ) and use the final state as entity embedding i.e. C^j as defined in equation 3:

$$C^j = \text{BiLSTM}_\psi(\{c_{1:n}^j\}^\phi) \quad (3)$$

We also learn an embedding vector $C^{nb} \in \mathcal{R}^D$ for a special token $\langle \text{no-bias} \rangle$, so the model can attend to it when transcribing any word which is not present in the provided list of entities. Let $C^e \in \mathcal{R}^{K \times D}$ represent concatenated embeddings of dimension D of the entity items plus the no-bias token.

The Biasing Adapter then uses these embeddings as keys and values in an attention module to obtain a context vector that is added to the output of the last layer of the Conformer Encoder. More formally, let $\{e_t^i\}_{i=1}^{i=N}$ represent the output of each of the N Conformer blocks present in the Encoder at time step t . We define a set of N trainable weights $\{w^i\}_{i=1}^{i=N}$ to create our attention query q_t as described in Equation 4 and the output of the Biasing Adapter can be defined as b_t :

$$b_t = \text{Attn}(W^q q_t, W^k C^e, W^v C^e) \text{ where } q_t = \sum_{i=1}^N w^i e_t^i \quad (4)$$

where W^q, W^k, W^v are query, key and value weight matrices respectively and $\text{Attn}(\cdot)$ is an attention function as defined in [32]. The output of Biasing Adapter is added to final encoder output before passing through a softmax layer

$$e_t^N = e_t^N + b_t \quad (5)$$

The formulation of our proposed Contextual Adapters is adapted from [25], and modified for a non auto-regressive CTC model as follows. It leverages information available in different encoder layers, rather than just the last layer. This is essential specially for CTC architectures like ours where the last encoder layer representations are closer to output representations (word-pieces) as compared to other architectures that rely on multiple decoder layers. Also, note that for CTC models, the learnt catalog embeddings (C^e) are not specific to the content/semantics of the entity but are just a representation of the sub-word sequence of the entity. Different from [25], we do not use any additional data for training the adapters in our experiments. We introduce a simple trick to reuse the generic data used for training the base model to train our version of Contextual Adapters. From the training transcripts, we identify a list of rare words (any word that appears less than 13 times) and the corresponding utterances. Then for each of these utterance, the list of input custom entity words simply comprises of a random subset from all the rare words and the rare word present in that utterance. This helps train the Biasing Adapter such that it uses the speech signal to understand if it needs to copy the word-piece sequence of the rare word to the output or choose the no-bias for any time frame. This training method ensures that our personalization approaches are truly zero-shot and no additional labeled data is required for adopting them in practice. Formally, let γ represent all the parameters of the base Conformer CTC model obtained after minimizing loss L_{MTL} as defined in 1 on the dataset \mathcal{D} . Then we jointly optimize all the parameters in the Contextual Adapters using the subset of the dataset $\mathcal{D}^{rare} \subset \mathcal{D}$ as defined above, keeping γ fixed as defined in Eqn 6 below

$$\{\psi, \phi, C^{nb}, \{w^i\}_{i=1}^{i=N}, W^q, W^k, W^v\} = \underset{\gamma}{\text{argmin}} \sum_{\mathcal{D}^{rare}} L_{MTL}(\gamma, \cdot) \quad (6)$$

At inference time, we swap in a fixed list of custom entity words for all the utterances in the test dataset. Since the catalog of entities is fixed, we create the entity embeddings C^e only once and save it in cache. This ensures we do not add any observable change to the inference latency of the system. Further, if for any frame, the attention weight given to the no-bias term is the maximum among all other catalog entities, we explicitly set $b_t = 0$ so that there is no regression/change in performance of the system on generic dataset/speech segments without the entity words.

3.2. Decoder Biasing Methods

3.2.1. Adaptive subword boosting in beam search decoding

The process of tuning boosting probabilities may lead to inconsistent results and is considered a sub-optimal approach. Therefore, we implemented an adaptive on-the-fly boosting of the vocabulary of interest where biasing is done at subword level during beam search decoding before the pruning stage. As described in [33], our beam search decoder performs a simple beam-search with beam thresholding, histogram pruning, and language model smearing. The decoder accepts unnormalized acoustic scores ($g_{i,j}(\cdot)$) as input and attempts to maximize the following:

$$L(\theta) = \underset{\pi \in G_{lex}(\theta, T)}{\text{logadd}} \sum_{t=1}^T (f_{\pi_t}(x) + g_{\pi_{t-1}, \pi_t}(x) + h_{\pi_t}(x)) + \alpha \text{log} P_{lm}(\theta) + \beta |\{i | \pi_i = \#\}|, \quad (7)$$

where the “logadd” operation (also called “log-sum-exp”) is defined as $\text{logadd}(a, b) = \log(\exp(a) + \exp(b))$. $G_{lex}(\theta, T)$ is a graph constrained by base lexicon over T frames for a given transcription θ , and $\pi = \pi_1, \dots, \pi_t \in G_{lex}(\theta, T)$ is a path in this graph representing a valid subword sequence for this transcription. At each time step t , $f_{\pi_t}(x)$ denotes the log-probability by the CTC model (given an acoustic sequence x) and $h_{\pi_t}(x)$ denotes the additional boosting score. $P_{lm}(\theta)$ is the probability of the language model given a transcription θ , α and β , are hyper-parameters which control the weight of the language model, and silence (#) insertion penalty, respectively.

Let $\hat{y}(t, k)$ denote the log-probability distribution by the CTC model in sorted order at time t , with k representing the subword candidate index in the top-K list. To compute the boosting score, we use a *trie* (T_{vocab}) constructed from the list of custom entities and boost the rest of the sequence if the first subword (trie node) occurs in the top-K list. The boosting score for each subword candidate at time step t is determined dynamically by its difference in log-probability with the top-1 hypothesis.

$$h_{\pi_t}(x) = \begin{cases} \delta * (\hat{y}(t, k) - \hat{y}(t, 1)) & \pi_t \in T_{vocab} \\ 0 & \text{Otherwise} \end{cases} \quad (8)$$

where δ denotes the boosting scale. To reduce false positives, we follow an inverse *sigmoid* function for computing the boosting scale as shown below:

$$\delta = 1 / (1 + e^{(\hat{y}(t, k) - \hat{y}(t, 1) - 0.5 * k) / 0.1 * k}) \quad (9)$$

Unigram boosting: We also consider modifications to the external N-gram LM before performing fusion with a beam search decoder. In theory, one simple way of biasing is to add OOV words to the external LM with a fixed high unigram probability. This can also be achieved by adding an OOV class (instead of adding all words) in order to keep the LM unmodified during inference. In this work, we add not only OOV words, but also other common and rare words that are tested as part of experiments presented in the Section 4.

3.2.2. Phone alignment network for biasing

In this section, we describe two biasing techniques using the per-frame phone predictions generated from the phone alignment network described in the Section 2.2.

Phonetic distance based rescoring: Although unigram boosting helps to a certain extent, rare and OOV words are still considered unknown tokens in higher order N-grams. In our preliminary experiments, we observed that the words that are boosted during on-the-fly rescoring are less preferred due to shallow fusion with an external language model. We therefore propose to use phonetic distance based rescoring of the N-best lists to bias those hypotheses that contain rare and OOV words. Let $X = \{x_1, x_2, \dots, x_T\}$ be a sequence of per-frame phone posteriors generated from the phone alignment network. Let $Y^n = \{y_1^n, y_2^n, \dots, y_{L_n}^n\}$ be a phone sequence corresponding to the n^{th} -best hypothesis (denoted by Y_w^n) from the N-best lists generated from the *CTC decoder* by performing shallow fusion with an external language model. We then perform forced alignment of each hypothesized phone sequence from the N-best list to the per-frame phone posteriors generated from the phone alignment network using a dynamic time warping (DTW) algorithm. The X and Y^n sequences can be arranged to form a T -by- L^n grid, where each point (i, j) is the alignment between x_i and y_j^n . A warping path W maps the elements of X and Y^n to minimize the distance between them. The optimal path to (i_k, j_k) can be computed by:

$$D_{min}(i_k, j_k) = \min_{i_{k-1}, j_{k-1}} D_{min}(i_{k-1}, j_{k-1}) + d(i_k, j_k | i_{k-1}, j_{k-1}) \quad (10)$$

where d is the Euclidean distance. Let us define d^n as the DTW cost incurred to align j^{th} feature of Y^n with the i^{th} feature vector of X and the overall path cost (D^n) can be calculated as:

$$D^n = \sum_k d^n(i_k, j_k) \quad (11)$$

where $i_k = T$ and $j_k = L_n$. We insert a */silence/* phone between each word to allow for pauses between words during the forced alignment. We also use a fixed scale of 0.1 for DTW cost and the overall cost defined below is used for rescoring the N-best lists:

$$rescore = P_{CTC}(Y_w^n) + 0.1 * D^n \quad (12)$$

Use of pronunciation lexicon: Traditional hybrid systems use phonetic lexicons and/or clustered context-dependent acoustic targets to recognize long-tail words. Similarly, we hypothesize the phone predictions to be more accurate than subword predictions for rare and OOV words. In this approach, we take the 1-best hypothesis from the *CTC Decoder* after performing phonetic distance based rescoring and tokenize it into words. The word boundaries generated during forced alignment are used to retrieve corresponding per-frame phone predictions. We then perform window-based smoothing to replace any spurious predictions and collapse consecutive repetitive predictions into a single phone. The pronunciation obtained in this way is compared to the lexicon-derived pronunciation of each vocabulary from the user provided biasing list. If an exact match is found, we replace the word in the 1-best hypothesis with matched vocabulary. Although this is a heuristic based approach, it has been shown to generalize across datasets with no additional tuning.

3.2.3. Leveraging phoneme similarity

Previous work on personalization for RNN-T models [23] leveraged G2G to generate additional pronunciation variants for domain-specific entities during decoding. Similarly, we generate few variants

using phoneme similarity and decompose them into subwords while maintaining the word-level label. With this approach, the probability of predicting the actual word slightly improves with an increase in the number of unique G2G variants.

4. EXPERIMENTS

4.1. Data

Training data: We train our base Conformer CTC model on two different corpora in size i.e., a large 50k+ hour English corpus and a smaller subset of 5k hour corpus, sampled from in-house paired audio and text data. These two data regimes are representative of a wide range of end-to-end ASR systems for various different speech applications. Both corpora include audio files with a good mix of accents, speakers, sampling rates and background noise.

VoxPopuli: For evaluation, we use a 5 hours test subset of publicly available VoxPopuli English data [34] with an average recording length of 9.6 seconds. We extract out entities from the reference text based on the frequency of each token in our larger training corpus. Any token that appears in the reference text but occurs less than 50 times in the training transcripts is called rare while OOVs are the words which are unseen during training. Our evaluation subset of the dataset contains 44407 tokens in total, of which 574 (96 unique) are rare and 55 (42 unique) are OOVs.

Medical: We also evaluated on an in-house medical dataset consisting of 609 recordings with each audio spanning an average of 147 seconds (25 hours in total). For preparation of the dataset, speakers of different genders, accents and ages read written conversational style utterances in different acoustic settings (background noise). These utterances contain a desired list of 755 medical entities (716 rare and 39 OOVs) like names of medicines and diseases eg. ‘acetabulofemoral’, ‘ecchymoses’, etc which are long (average 9.2 characters per word) and hard to be recognized by any generic ASR system and hence used to gauge the performance of our proposed methods. This dataset contains a total of 189380 tokens, of which 3625 are rare tokens and 198 OOV tokens. We use this dataset to demonstrate the efficacy of our proposed methods while scaling to a larger list of medical terms. For experimentation, we also trained a medical domain-specific language model on an internal medical text corpus with 3.9 million words (21k unique tokens).

4.2. Experimental Setting

Baseline: We train two different mixed-bandwidth 16k models of the multi-task Conformer CTC-Attention architecture as described in Section 2.2. For the large model, we stacked 20 layers of Conformer blocks with 8 attention heads and roughly 142.5 million parameters and trained on our large English corpus, while for the small model, we stacked 16 blocks with 4 attention heads and 30.2 million parameters and trained on 5k hour subset. We train a sentence-piece tokenizer with token size of 2048 and 1024 for large and small model respectively. Both baselines are trained with ADAM optimizer with a learning rate of 0.001. Further, we train an external 4-gram language model (LM) for shallow fusion. During inference, we use a beam size of 50 and LM weight of 0.6 in all our experiments. All our work is implemented in the open-source ESPnet tool [35].

Encoder-biasing: We train the Contextual Adapters on the same subset of data used for training the generic CTC model. To save compute costs, we only use the 10k hour subset from our large corpus for training the adapters. For both the models, $D = 128$ (catalog embedding dimension) and we vary k (catalog size) from 30 to 250

Table 1: Performance comparison of personalization methods on VoxPopuli dataset

	Model	WERR (%age)	Rare			OOVs		
			F1	P	R	F1	P	R
1	Baseline Conformer-CTC (Large)	-	49.4	97.7	33.1	59.0	100.0	41.8
2	Row 1 + Contextual Adapter (top layer)	-1	74.4	86.5	65.4	68.5	69.8	67.3
3	+ Enforce no-bias token	-0.5	71.9	86.7	61.4	68.0	70.6	65.5
4	+ Intermediate Layers	2	77.0	95.3	64.6	69.3	76.1	63.6
5	Row 1 + Adaptive subword boosting	-1	72.5	87.4	62	68.2	79.5	59.8
6	+ Phonetic distance rescoring	-0.8	74.6	89.5	63.9	69.1	80.1	60.7
7	+ Pronunciation lexicon	0	76.2	89.1	66.5	71.6	78.5	65.7
8	+ G2G	0.2	76.7	88.6	67.6	72.5	78.1	67.6
9	+ Shallow fusion with boosted LM	0.4	77.4	90.1	67.9	73.2	78.6	68.5
10	Proposed joint model (Row 4 + Row 9)	1.2	79.9	92.6	70.2	77.5	81	74.3

during training (starting with 30 and increasing it by 4 every epoch to a maximum of 250). We only use 1 attention head for the Biasing Adapter and attention dimension of 128. We freeze all other parameters of the base model as we update the Adapter parameters (1.32 million, which is < 1% of the parameters of the base large model). We train two versions of the Adapters for the larger model, one where attention query q_i is with weighted sum of 6, 12 and 20th encoder layer (called 'intermediate layers') and the other with just the last layer [25] (called 'top layer'). Further, we run inference with and without enforcing $b_i = 0$ for frames with maximum weight to no-bias token (denoted as 'enforce no-bias'). We use the best version for the 5k-hr model ablation.

Decoder-Biasing: At each step of beam search decoding we consider a top 10 candidates for expanding our decoding tree. Since the BPE model typically generates a unique segmentation of words, subwords into which rare words are split have poor probability. However, we noticed that often a different subword sequence which corresponds to a rare word occurs in the beam paths. To overcome this issue, we can use a dropout in the BPE model to generate multiple variations of segmentation of rare words. In our experiments, we generated up to 10 variations of segmentation for the rare words that we boost in prefix based adaptive subword boosting. We use a unigram boosting score of -0.2 in our domain language model after careful tuning on the dev dataset.

Joint model: When combining encoder and decoder biasing methods, we do not tune any hyper-parameters of the individual techniques again. All the results presented in this paper are obtained by direct combination, without changing the individual methods.

5. RESULTS

Tables 1 and 2 present the results of our proposed personalization methods when biased with rare and OOV words found in the VoxPopuli test dataset. In Table 1, we compare both encoder and decoder biasing methods as well as their joint combination against a strong Conformer-CTC baseline. All the methods used an external 4-gram language model in fusion with a beam search decoder. The numbers reported in the first column are Word Error Rate Reduction %age (WERR) over the baseline 9.1 WER. Overall, the proposed approaches outperform the baseline in both WER and F1 score.

Encoder-biasing: In the encoder biasing, the addition of the Contextual Adapter to the top layer alone significantly improved the F1 score of both rare words and OOVs, but also resulted in a small increase in WER. This is likely due to the high attention weight from the biasing adapter forcing the encoder to copy a word and not allow-

Table 2: Performance comparison with smaller subset of training data (5000 hours) on VoxPopuli dataset

Model	WERR %age	F1 Rare	F1 OOV
Baseline Conformer-CTC (Small)	-	32.9	47.4
Encoder Biasing (Row 4 from Table 1)	0.6	53.6	64.4
Decoder Biasing (Row 9 from Table 1)	0.9	54.2	67.5
Proposed Joint Model	1.6	56.1	69.8

ing the model to recover from any false positives. However, when adapters are introduced in intermediate layers with learnable weights and the no-bias token is enforced, the overall WER is improved over the baseline model along with F1 score.

Decoder-biasing: In decoder biasing, we applied a combination of methods on top of the baseline beam search decoding to recover domain-specific vocabulary as summarized in the Table 1. First we implemented a subword level boosting on all candidates in the beam paths before pruning as described in Section 3.2.1 which led to significant improvements across all categories of vocabulary using this technique. We noticed that even though the subword boosting technique was very effective, the correct hypothesis with domain-specific vocabulary doesn't surface in the 1-best but remains in the N-best which accounted for 10% of errors. To alleviate this, we applied phonetic distance based rescoring which led to a relative improvement of 3% in rare vocabulary recognition. However, we only noticed a 1.3% improvement in OOV. This is probably due to the model assigning low scores to subwords that make up the OOV words, resulting in their absence in the N-best hypotheses. We then applied a pronunciation lexicon to further improve the 1-best hypothesis. As hypothesized, the phone predictions of the model were more accurate than subword predictions in certain cases, thereby leading to improvements in both rare and OOV words. Finally, we applied G2G and a simple technique of shallow fusion with a boosted language model as described in 3.2 which improved performance across both rare and OOV words due to boosting at word level.

Joint model: When we experimented with the combined model containing both the encoder and decoder biasing methods, we observe a significant improvement in F1 scores. From our error analysis, we noticed that the main challenge for poor recognition of rare and OOV words is either beam paths containing rare subword sequences getting pruned due to low probability or the rare subwords not being present in the top-k candidates. Our joint model is very effective in addressing these issues due to the complimentary nature of the proposed approaches. The encoder biasing method helps very

Table 3: Performance comparison of proposed methods with larger catalog size (755) on medical dataset

Model	WERR %age	F1 Rare	F1 OOV
Baseline Conformer-CTC (Large)	-	36.4	15.7
+ Shallow Fusion (SF) with medical LM	11.5	39.2	15.4
SF + Encoder Biasing	13.2	58.1	29.8
SF + Decoder Biasing	4.9	57.9	22.4
SF + Proposed Joint Model	5.6	62.1	27.2

Table 4: Example utterances from VoxPopuli dataset with outputs from our proposed joint model

baseline: <i>We will be able to start our trio negotiations</i>
joint model: <i>We will be able to start our trialogue negotiations</i>
baseline: <i>Resignation of professor galeo as chairman of the snc</i>
joint model: <i>Resignation of professor ghalioun as chairman of the snc</i>
baseline: <i>president taani has stated that industry is at the heart of europa</i>
joint model: <i>president tajani has stated that industry is at the heart of europa</i>
baseline: <i>Mediation of pat cox and alexander KSSK</i>
joint model: <i>Mediation of pat cox and alexander KWANIEWSKI</i>
baseline: <i>How is it possible for jan marle to completely disappear</i>
joint model: <i>How is it possible for jan marsalek to completely disappear</i>
baseline: <i>Fellow creations and other nations</i>
joint model: <i>Fellow croatians and other nations</i>
baseline: <i>A significant number of tanous national minorities</i>
joint model: <i>A significant number of autochthonous national minorities</i>

effectively in generating higher probability scores of the rare subwords it copies which helps in preventing the rare subwords from getting pruned in the beam path or top-k candidates. This is further improved by applying the decoder biasing techniques in boosting the rare word candidate paths to 1-best which led to improvements in F1 across both rare and OOV words. We also observed that the Contextual Adapter biasing approach performs better than only subword boosting which gave the most gains in decoder biasing methods across both categories. However, when combined with other approaches such as phonetic distance based rescoring, pronunciation based lexicon lookup and G2G, decoder biasing has better recall on rare and OOVs. We hypothesize that encoder-biasing method tends to mis-recognize words whose pronunciations differ from the spellings as it does not have access to phonetic information of the entity word. However, the proposed decoder-biasing methods leverage phone sequences from phone decoder and lexicon based pronunciation of the word to boost or correct the rare, OOV words in the n-best list and the final 1-best. Decoder based biasing methods also provide the flexibility to not rely on the underlying model performance by providing several knobs to finetune for boosting weights and nbest list correction methods. Table 4, summarizes the qualitative examples of the joint model outputs.

5.1. Ablation study with model trained on smaller corpus

We evaluated our model performance when trained on a smaller corpus of 5k hour. Table 2 summarizes the results comparing different proposed approaches against the baseline (with 12.5 WER). We observe that both methods improved WER and F1 of rare and OOV entities with decoder biasing methods performing better on OOV words. This is because the decoder biasing methods such as lexicon lookup and G2G work effectively even if the subword sequences related to OOV are absent in the top candidates of the decoding. Our joint proposed model has the best performance with 70.5% and

Table 5: Performance comparison of our methods customized for VoxPopuli entity list on general test set

Model	WERR (%age)
Baseline Conformer-CTC (Large)	-
Encoder Biasing	-0.34
Decoder Biasing	-0.39
Proposed joint model	-0.42

47.3% relative improvements in F1 of rare and OOV words respectively over the baseline.

5.2. Effect of personalization on general dataset

Personalization or adaptation of models to a specific use case or domain often leads to degradation in performance on general datasets. To demonstrate that our proposed approaches are generic, we evaluated the different methods personalized for VoxPopuli dataset on a general test set not containing entities (comprising of 1671 audio files) sampled from an original VoxPopuli data distribution with a baseline 8.8 WER. As summarized in table 5, we noticed a very small degradation with individual encoder and decoder biasing approaches. The joint model had a slightly higher degradation of 0.42% than individual approaches due to overbiasing. However, this degradation (less than 0.5%) is negligible compared to the significant improvements of rare and OOV words on the overall dataset.

5.3. Efficacy of proposed approaches on large catalog size

Table 3 presents the results on a large catalog size typically found in the medical domain. We find that our Contextual Adapter method scales very well to a larger custom list. We observed a significant 48% relative improvement in F1 score on rare words along with a small improvement in WER. Although, the Adapters were exposed to a maximum of 250-sized entity list while training, they were observed to be performing well for a list that is roughly three times in size. However, we see a regression in WER as compared to a simple shallow fusion approach on using decoder-biasing methods on such a large list. This is not surprising because if we boost a large number of words while decoding, it is likely that many similar sounding words will match, thereby causing a drop in precision. On the other hand, an encoder-biasing method like Adapters rely on the speech signal to selectively choose when to copy a word, rather than simply copying it every time. Particularly, in our technique 'enforce no bias', we explicitly ensure that system output does not change for the frames where no-bias token has the maximum attention weight. This analysis can be useful for the research community to choose among various proposed personalization methods based on their use case, size of custom entities and trade-offs.

6. CONCLUSIONS

In this paper, we proposed a combination of techniques to bias a conformer-CTC model to significantly improve the recognition of domain specific vocabulary, especially rare (long-tail) and OOV words. We explored both encoder biasing and on-the-fly decoder biasing techniques and compared their performance under two different data settings with base models trained on both large and small corpora. We demonstrated the efficacy of our proposed approaches with significant improvements in F1 score when biased with a large catalog. Our proposed model showed significant improvements of 60.3% and 30% in relative F1 score on rare and OOV words respectively on the VoxPopuli dataset with negligible degradation (less than 0.5%) in WER on general utterances.

7. REFERENCES

- [1] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. ICASSP*, 2013, pp. 6645–6649.
- [2] Julian Salazar, Katrin Kirchhoff, and Zhiheng 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.
- [3] Alex Graves, "Sequence transduction with recurrent neural networks," *arXiv preprint arXiv:1211.3711*, 2012.
- [4] Yan Zhang He, Tara N Sainath, Rohit Prabhavalkar, Ian McGraw, Raziq Alvarez, Ding Zhao, David Rybach, Anjali Kannan, Yonghui Wu, Ruoming Pang, et al., "Streaming end-to-end speech recognition for mobile devices," in *Proc. ICASSP*, 2019, pp. 6381–6385.
- [5] William Chan, Navdeep Jaitly, Quoc V Le, and Oriol Vinyals, "Listen, attend and spell," *arXiv preprint arXiv:1508.01211*, 2015.
- [6] Liang Lu, Changliang Liu, Jinyu Li, and Yifan Gong, "Exploring transformers for large-scale speech recognition," *arXiv preprint arXiv:2005.09684*, 2020.
- [7] Ian McGraw, Rohit Prabhavalkar, Raziq Alvarez, Montse Gonzalez Arenas, Kanishka Rao, David Rybach, Ouais Alsharif, Haşim Sak, Alexander Gruenstein, Françoise Beaufays, et al., "Personalized speech recognition on mobile devices," in *Proc. ICASSP*, 2016, pp. 5955–5959.
- [8] Ye Bai, Jiangyan Yi, Jianhua Tao, Zhengkun Tian, Zhengqi Wen, and Shuai Zhang, "Listen attentively, and spell once: Whole sentence generation via a non-autoregressive architecture for low-latency speech recognition," 2020.
- [9] Chanwoo Kim, Dhananjaya Gowda, Dongsoo Lee, Jiyeon Kim, Ankur Kumar, Sungsoo Kim, Abhinav Garg, and Changwoo Han, "A review of on-device fully neural end-to-end automatic speech recognition algorithms," in *2020 54th Asilomar Conference on Signals, Systems, and Computers*, 2020, pp. 277–283.
- [10] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang, "Conformer: Convolution-augmented transformer for speech recognition," 2020.
- [11] Suyoun Kim, Takaaki Hori, and Shinji 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.
- [12] Shinji Watanabe, Takaaki Hori, Suyoun Kim, John R Hershey, and Tomoki Hayashi, "Hybrid CTC/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [13] Pengcheng Guo, Florian Boyer, Xuankai Chang, Tomoki Hayashi, Yosuke Higuchi, Hirofumi Inaguma, Naoyuki Kamo, Chenda Li, Daniel Garcia-Romero, Jiatong Shi, et al., "Recent developments on ESPNET toolkit boosted by conformer," in *Proc. ICASSP*, 2021, pp. 5874–5878.
- [14] Jinyu Li, Yu Wu, Yashesh Gaur, Chengyi Wang, Rui Zhao, and Shujie Liu, "On the comparison of popular end-to-end models for large scale speech recognition," *arXiv preprint arXiv:2005.14327*, 2020.
- [15] Jinyu Li, Rui Zhao, Zhong Meng, Yanqing Liu, Wenning Wei, Sarangarajan Parthasarathy, Vadim Mazalov, Zhenghao Wang, Lei He, Sheng Zhao, et al., "Developing RNN-T models surpassing high-performance hybrid models with customization capability," *arXiv preprint arXiv:2007.15188*, 2020.
- [16] Jinyu Li, "Recent advances in end-to-end automatic speech recognition," *arXiv preprint arXiv:2111.01690*, 2021.
- [17] Tara N Sainath, Rohit Prabhavalkar, Shankar Kumar, Seungji Lee, Anjali Kannan, David Rybach, Vlad Schogol, Patrick Nguyen, Bo Li, Yonghui Wu, et al., "No need for a lexicon? evaluating the value of the pronunciation lexica in end-to-end models," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5859–5863.
- [18] Antoine Bruguier, Fuchun Peng, and Françoise Beaufays, "Learning personalized pronunciations for contact name recognition," in *INTERSPEECH*, 2016.
- [19] Ashish Shenoy, Sravan Bodapati, Monica Sunkara, Srikanth Ronanki, and Katrin Kirchhoff, "Adapting long context NLM for ASR rescoring in conversational agents," in *Interspeech 2021*. aug 2021, ISCA.
- [20] Saket Dingliwal, Ashish Shenoy, Sravan Bodapati, Ankur Gandhe, Ravi Teja Gadde, and Katrin Kirchhoff, "Prompt tuning gpt-2 language model for parameter-efficient domain adaptation of asr systems," 2021.
- [21] Christophe Van Gysel, Mirko Hannemann, Ernest Pusateri, Youssef Oualil, and Ilya Oparin, "Space-efficient representation of entity-centric query language models," 2022.
- [22] Mahaveer Jain, Gil Keren, Jay Mahadeokar, Geoffrey Zweig, Florian Metze, and Yatharth Saraf, "Contextual RNN-T for open domain ASR," in *Proc. Interspeech*, 2020, pp. 11–15.
- [23] Duc Le, Gil Keren, Julian Chan, Jay Mahadeokar, Christian Fuegen, and Michael L Seltzer, "Deep Shallow Fusion for RNN-T Personalization," in *IEEE Spoken Language Technology Workshop (SLT)*, 2021, pp. 251–257.
- [24] Aditya Gourav, Linda Liu, Ankur Gandhe, Yile Gu, Guitang Lan, Xiangyang Huang, Shashank Kalmane, Gautam Tiwari, Denis Filimonov, Ariya Rastrow, et al., "Personalization Strategies for End-to-End Speech Recognition Systems," in *Proc. ICASSP*, 2021, pp. 7348–7352.
- [25] Kanthashree Mysore Sathyendra, Thejaswi Muniyappa, Feng-Ju Chang, Jing Liu, Jinru Su, Grant P. Strimel, Athanasios Mouchtaris, and Siegfried Kunzmann, "Contextual adapters for personalized speech recognition in neural transducers," 2022.
- [26] Nilaksh Das, Duen Horng Chau, Monica Sunkara, Sravan Bodapati, Dhanush Bekal, and Katrin Kirchhoff, "Listen, know and spell: Knowledge-infused subword modeling for improving asr performance of oov named entities," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7887–7891.
- [27] Golan Pundak, Tara N. Sainath, Rohit Prabhavalkar, Anjali Kannan, and Ding Zhao, "Deep context: end-to-end contextual speech recognition," 2018.

- [28] Antoine Bruguier, Rohit Prabhavalkar, Golan Pundak, and Tara N. Sainath, “Phoebe: Pronunciation-aware contextualization for end-to-end speech recognition,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 6171–6175.
- [29] Anjali Kannan, Yonghui Wu, Patrick Nguyen, Tara N. Sainath, Zhifeng Chen, and Rohit Prabhavalkar, “An analysis of incorporating an external language model into a sequence-to-sequence model,” 2017.
- [30] Duc Le, Thilo Koehler, Christian Fuegen, and Michael L Seltzer, “G2G: TTS-driven pronunciation learning for graphemic hybrid ASR,” in *Proc. ICASSP*, 2020, pp. 6869–6873.
- [31] Rui Zhao, Jian Xue, Jinyu Li, Wenning Wei, Lei He, and Yifan Gong, “On Addressing Practical Challenges for RNN-Transducer,” *arXiv preprint arXiv:2105.00858*, 2021.
- [32] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” 2017.
- [33] Vitaliy Liptchinsky, Gabriel Synnaeve, and Ronan Collobert, “Letter-based speech recognition with gated convnets,” *arXiv preprint arXiv:1712.09444*, 2017.
- [34] Changhan Wang, Morgane Rivière, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux, “Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation,” 2021.
- [35] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplín, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai, “ESPnet: End-to-end speech processing toolkit,” in *Proc. Interspeech*, 2018, pp. 2207–2211.