

# Don't Be So Sure! Boosting ASR Decoding via Confidence Relaxation

Tomer Wullach, Shlomo E. Chazan

OriginAI, Ramat-Gan, Israel  
tomew, shlomi@originai.co

## Abstract

Automatic Speech Recognition (ASR) systems frequently use a search-based decoding strategy aiming to find the best attainable transcript by considering multiple candidates. One prominent speech recognition decoding heuristic is beam search, which seeks the transcript with the greatest likelihood computed using the predicted distribution. While showing substantial performance gains in various tasks, beam search loses some of its effectiveness when the predicted probabilities are highly confident, i.e., the predicted distribution is massed for a single or very few classes. We show that recently proposed Self-Supervised Learning (SSL)-based ASR models tend to yield exceptionally confident predictions that may hamper beam search from truly considering a diverse set of candidates. We perform a layer analysis to reveal and visualize how predictions evolve, and propose a decoding procedure that improves the performance of fine-tuned ASR models. Our proposed approach does not require further training beyond the original fine-tuning, nor additional model parameters. In fact, we find that our proposed method requires significantly less inference computation than current approaches. We propose aggregating the top  $M$  layers, potentially leveraging useful information encoded in intermediate layers, and relaxing model confidence. We demonstrate the effectiveness of our approach by conducting an empirical study on varying amounts of labeled resources and different model sizes, showing consistent improvements in particular when applied to low-resource scenarios.

## Introduction

Self-Supervised Learning (SSL) has been widely adopted in a variety of tasks (Devlin et al. 2018; Liu et al. 2019, 2020b; Baevski et al. 2020; Liu et al. 2020a) and has shown to be capable of producing representation that can be effectively utilized for downstream tasks. Models trained via SSL leverage a massive amount of unlabeled data during pre-training, and are further tuned for specific tasks using a significantly smaller amount of labeled data. In recent years, SSL has been employed for speech processing tasks (Baevski et al. 2020; Hsu et al. 2021; Liu et al. 2020a; Van den Oord, Li, and Vinyals 2018), utilizing transformer encoders layers (Vaswani et al. 2017) creating context-aware representations.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

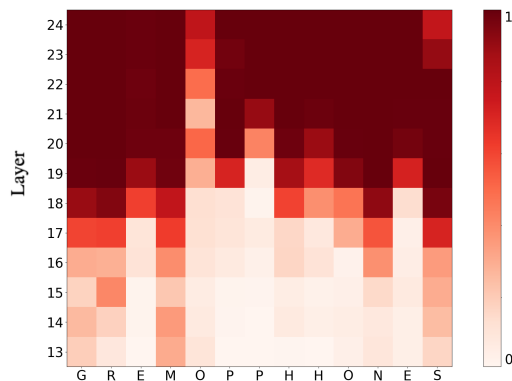


Figure 1: The evolution of confidence throughout the top 12 transformer layers of Wav2vec 2.0 Large fine-tuned on 960 hours of Librispeech. The confidence increases at the top layers, ultimately reaching exceptionally high levels.

Recently proposed SSL ASR models commonly employ a greedy decoding strategy, selecting the token with the highest predicted probability at each time step (Baevski et al. 2020; Hsu et al. 2021). This approach is potentially sub-optimal as sequences of high quality might be disregarded. Consequently, researchers often maintain a search-based decoder such as beam search (Baevski et al. 2020; Chen et al. 2021; Hsu et al. 2021), showing substantial performance gains, in particular when fused with a Language Model (LM) (Hsu et al. 2021; Baevski et al. 2020; Chen et al. 2021; Gulati et al. 2020).

Beam search considers several alternatives by maintaining the  $B$  partial sequences with the highest conditional probability at each decoding step, where  $B$  is commonly referred to as Beam Width. Although enabling easy, straightforward decoding, beam search might be prone to biases when predictions are highly confident. In that case, the predicted probability distribution at each time step is massed to very few tokens, thus, reducing the relevancy of the beam width, as a small number of candidates dominate all others. Consequently, highly confident predictions might prevent the algorithm from truly considering multiple alternatives.

From another perspective, previous works have shown

that different layers encode different aspects of the training data (Pasad, Chou, and Livescu 2021; Chang et al. 2021) such as semantic and acoustic information. Thus, predictions based solely on the top layer might neglect informative features or give excessive weight to others.

In this study, we monitor and analyze the confidence levels of several high-performing ASR models, which are composed of a CNN-based encoder followed by Transformer Encoder layers that produce context-aware predictions. In experiments using Wav2vec 2.0 (Baevski et al. 2020) and HuBERT (Hsu et al. 2021), we find that high confidence levels widely emerge at the top transformer layers, while intermediate layers tend to be less confident and more diverse. Furthermore, we evaluate the adverse influence of these findings on the performance of a beam search decoder.

Recent studies demonstrated the merits of aggregating deep transformer layers (Yang et al. 2020), showing that aggregating layers encoding different features create enriched representations that boost performance (Dou et al. 2020; Wang et al. 2020). Different from previous works that employed the concept of leveraging intermediate representations (Chang et al. 2021; Dou et al. 2018; Chang, Yang, and Lee 2022; Wang et al. 2020), Our method *does not require additional training nor trainable parameters*. Furthermore, we find that the relaxed distribution formed by the aggregation operation enables the use of smaller beam widths without compromising performance. Hence, reducing beam search’s computation costs (Meister, Vieira, and Cotterell 2020).

Our contributions can be summarized as follows:

1. We show that state-of-the-art SSL ASR models produce highly confident predictions, and illustrate how confidence progressively accumulates across layers.
2. We demonstrate the negative impact of high confidence on speech recognition when using a beam search decoder.
3. We suggest a layer aggregation approach for relaxing the predicted confidence, harnessing information encoded in intermediate layers, and improving computation costs. We show the benefits of our approach using different amounts of labeled resources and models of sizes.

## Related Work

**Self-Supervised Speech Recognition** Recently proposed SSL speech recognition models follows the large body of work showing tremendous success with employing transformer-based encoders for extracting informative features from raw input. In most cases, these models are first pre-trained using vast amounts of unlabeled data and later fine-tuned using labeled data by solving a supervised training objective. The recently proposed Wav2vec 2.0 model (Baevski et al. 2020) is composed of a feature encoder that consists of stacked 1-dimensional convolution layers followed by multiple transformer encoder layers. The feature encoder maps raw audio signals into latent speech representations, which are then fed to the first (lowest) transformer encoder layer. The output of each transformer encoder layer is used as input to the next-in-line transformer

encoder layer, utilizing the self-attention mechanism to create contextualized speech representations. The model is pre-trained using a contrastive objective, mapping masked spans of a continuous signal into a discrete set of latent representations using a Gumbel-softmax operation (Jang, Gu, and Poole 2016). The loss is computed by contrasting true quantized latent representations of masked time steps from a set of distractors. Once pre-trained, the model is fine-tuned to speech recognition by projecting each time step’s contextualized representation, created by the top transformer encoder layer, into  $C$  classes, where  $|C|$  is the target vocabulary size. The fine-tuning uses Connectionist Temporal Classification (CTC) loss (Graves et al. 2006) to optimize the model parameters. The CTC loss enables getting around the alignment mismatch of the speech representations and textual labels, as the input audio length is much longer than the target text sequence and most datasets do not provide an accurate alignment of each audio frame and its corresponding text label. The conditional probability which is used to compute CTC loss is defined as:

$$P_{CTC}(Y | X) = \sum_{A \in A_{X,Y}} \prod_{t=1}^T p_t(a_t | X), \quad (1)$$

where  $A$  is a set of valid alignments for input sequence  $X$  and target sequence  $Y$ , and  $p_t(a_t | X)$  is the probability of observing label  $a_t$  at time  $t$ .

Different from Wav2vec 2.0 pre-training, Hidden Unit BERT (HuBERT) (Hsu et al. 2021) aims to predict the masked speech representation’s cluster, assigned using K-means applied to the input’s mel-frequency cepstral coefficient features (MFCC). HuBERT uses an offline quantization step to predict, thus, can directly infer target classes rather than relying on a contrastive objective.

**Intermediate Transformer Layers Information** Many studies made efforts to shed light on transformer encoder-based models (BERT-like) mechanisms and the information captured within their intermediate layers (Tenney, Das, and Pavlick 2019; Clark et al. 2019), an area of interest dubbed BERTology (Rogers, Kovaleva, and Rumshisky 2020). A key finding is that some features are better represented in certain layers (Shah et al. 2021) and that different layers may encode different input properties. Building on these observations, previous works pointed out the merits of incorporating features encoded in lower layers (Xiong et al. 2018) and the benefits of fusing multiple layers instead of the traditional approach of using the output produced by the single last layer (Kondratyuk and Straka 2019; Su and Cheng 2020). To summarize the information encoded across layers, Ji et al. (2021) propose fusing layer representations by sequentially feeding them to a recurrent layer such as LSTM (Hochreiter and Schmidhuber 1997), and using its final state as a global representation. In previous work, Wang et al. (2020) propose using fusion and pooling methods for Neural Machine Translation. They use an aggregation method similar to ours, however, they aggregate layer representations rather than layer logits, and employ the aggregation on the transformer decoder module.




Figure 2: Example from Librispeech test-clean: "...people are not so different from gramophones..." (wrongly predicted as *gre-mophones*). Top 4 probabilities for each time step, predicted using wav2vec 2.0 Large fine-tuned using 960 hours of Librispeech data. The model tends to produce exceptionally confident predictions (upper figure). The proposed layer aggregation approach (lower figure) correctly predicted "*gramophones*". The aggregation relaxed the model's confidence and diversified the decoded candidates.




Figure 3: The Prediction evolution across the top 10 transformer layers, using Wav2vec 2.0 Large with a beam search decoder fused with LM. The baseline approach wrongly predicted "break", while the proposed approach was able to produce the correct word "brake".

Different from previous works, our method does not require further training nor additional parameters, and applied on models which have already been pre-trained and fine-tuned. To the best of our knowledge, this is the first work that addresses speech recognition confidence related issues via layer aggregation, and exhibits its effectiveness in both performance and computational costs.

### Methodology

We investigate SSL-based speech recognition models (denoted as Acoustic Models) fine-tuned to predict a transcript sequence of length  $T$ ,  $Y = [y_1, \dots, y_T]$ , while aiming to minimize a CTC loss (Graves et al. 2006). The models are fed with an input audio sequence  $X$ , and output a logit matrix of shape  $R^{T \times C}$ . The resulting conditional probability can be defined as  $P(Y | X) \in R^{T \times C}$ .

Let  $G$  be a speech recognition model, pre-trained using SSL and fine-tuned for speech recognition on dataset  $D$ .

Comprising  $G$  are a feature extraction module, that extracts latent representations forming a sequence of speech units  $Z$  of length  $T$ , and a contextual module with  $N$  transformer encoder layers. The contextual module is fed with  $Z$  and produces contextual representations (Baevski et al. 2020; Hsu et al. 2021). We denote  $H_n^t$  as the representation produced by transformer layer  $n \in 1, \dots, N$  at time step  $t \in 1, \dots, T$ . Let projection\_head( $H_N$ ) be a linear projection layer trained to project the emissions of the top transformer layer,  $H_N$ , into  $C$  classes, where  $|C|$  is the vocabulary size. Note that  $C$  is composed of alphabet characters as well as special CTC tokens representing blanks and separators denoted as  $\epsilon$  and  $|$ , respectively.

**Decoding Strategies** A greedy decoding heuristic would be observing each position as if it is isolated from the rest of the sequence, and selecting the token that maximizes the probability defined as:

$$P_{greedy}(Y | X) = \arg \max_Y \prod_{t=1}^T p_t(y_t | X). \quad (2)$$

However, such paradigm might neglect candidates with higher quality, as previous predictions cannot be altered given newly decoded information. For this purpose, a beam search decoder is frequently employed, aiming to maximize the combination of both logits and LM scores:

$$P_G(Y | X) + \alpha_1 P_{LM}(Y) + \alpha_2 |Y|, \quad (3)$$

where  $Y$  is the predicted transcript sequence,  $\alpha_1$  denotes the weight of the language model probability and  $\alpha_2$  is the word score  $|Y|$  is the length of the sequence.

**Confidence Levels** A network is over-confident when the probability distribution is massed to a single class (Pereyra et al. 2017). In other words, highly confident predictions occurs when the softmax probabilities computed by employing the projection head on  $H_n$  have low entropy.

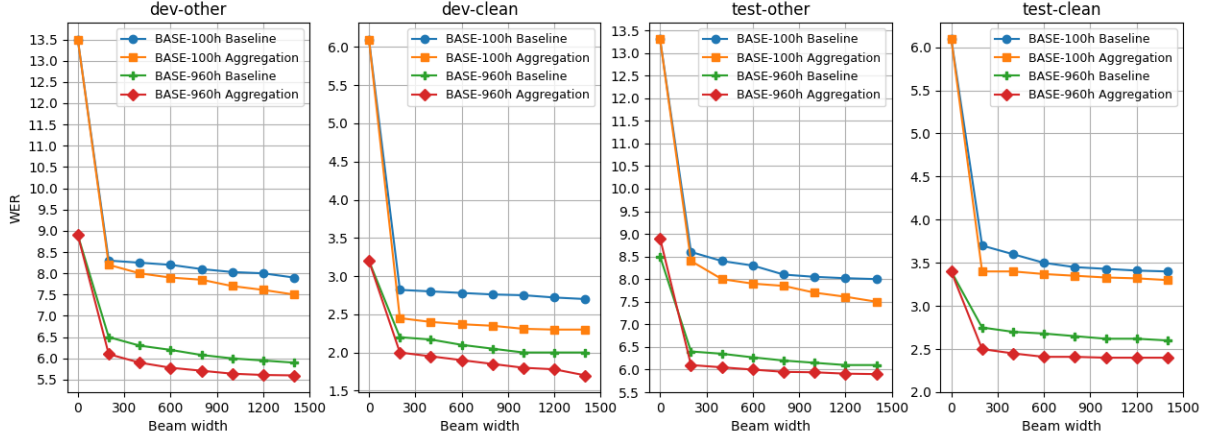


Figure 4: WER as a function of beam width, computed on Librispeech splits using Wav2vec 2.0 BASE fine-tuned using 100h and 960h. The proposed aggregation method is able to perform as well as the standard method, using significantly smaller beam size.

To investigate how confidence evolves throughout the transformer layers, we employ the same projection\_head followed by a Softmax operation on the representations emitted by *each transformer layer*.

While projection\_head is trained to project the model’s final representations into the vocabulary space, previous works demonstrated that the same projection layer can be employed on intermediate representations as well (Geva et al. 2021).

Our findings indicate that there is a confidence increase at the top transformer layer, as illustrated in Figure 1. We hypothesize that excessive confidence might prevent beam search from considering promising, yet less confident, predictions and subsequently compromise its performance.

Figure 2 (upper) illustrate the top 4 predicted probabilities at each time step. It is worth noting that the predicted probability distribution at the 5-th position is massed almost entirely at token "E".

Completing the picture, the predicted token evolution is illustrated in Figure 3. For each transformer layer at each time step, we chose the token with the largest probability, computed using projection\_head. This illustrates that a correct prediction may live in intermediate layers rather than the top layer. Thus, utilizing intermediate layers can provide beam search with additional valuable information.

**Layer Aggregation.** Building on these observations, we propose aggregating the logits of the top  $M < N$  transformer layers, forming relaxed and better-informed logits to be used by the beam search decoder. The aggregation of the projected logits can be viewed as a projection of the sum of the  $M$  layers. Intuitively, this can be accomplished using the distributive property of the matrices, where  $A \times B + A \times C = A \times (B + C)$ .

As a consequence of the confidence increase across the layers, a scaling mechanism is required to prevent the top layers from dominating the aggregation product. We apply

$L2$  normalization on each layer’s predicted logits, regularizing the impact of overly confident layers.

Formally, the layer aggregated logits can be described as follows:

$$\widehat{logits}(X) = \sum_{n=N-M}^N projection\_head\left(\frac{H_n}{\|H_n\|_2}\right). \quad (4)$$

Finally, the aggregated logits and the top layer’s logits are interpolated, forming the logits used for beam search:

$$\beta \cdot logits(X) + (1 - \beta) \cdot \widehat{logits}(X), \quad (5)$$

where  $\beta$  is the aggregation trade-off coefficient.

The following sections provide a comparative study of the aggregation approach effectiveness applied on SSL-ASR models.

## Experimental Setup

Our study includes recently proposed SSL-based models, namely, Wav2vec 2.0 and HuBERT. We experiment with several models variants that differ in size and pre-training data. The model size varies between 95M and 964M parameters, and pre-training was performed using either 960 hours of Librispeech ('LS-960') (Panayotov et al. 2015) or 60k hours of Libri-Light ('LL-60K') (Kahn et al. 2020).

We conduct our experiments using pre-trained models downloaded from the HuggingFace platform (<https://huggingface.co/models>), which were fine-tuned using 10m, 1h, 10h, 100h, 360h and 960h 'train' subsets of Librispeech.

The results for all of the experiments were attained using a beam search decoder with a 4-gram LM provided by pyctcdecode (<https://github.com/kensho-technologies/pyctcdecode>). The beam search decoder parameters,  $\alpha_1$  and  $\alpha_2$  (see eq. 3), were searched and optimized

Model	Pre-training	Layers	10-min		1-hour		10-hour		100-hour		360-hour		960-hour	
			$\beta$	$M$	$\beta$	$M$	$\beta$	$M$	$\beta$	$M$	$\beta$	$M$	$\beta$	$M$
W2v BASE	LS-960	12	0.25	4	0.25	4	0.3	4	0.5	5	0.7	4	0.75	6
W2v LARGE	LS-960	24	0.5	6	0.5	6	0.65	6	0.7	6	0.7	6	0.75	12
HUBERT BASE	LS-960	12	0.5	4	0.5	5	0.6	5	0.65	5	0.7	5	0.7	6
HUBERT LARGE	LL-60k	24	-	-	-	-	-	-	0.7	10	0.75	12	0.75	13
HUBERT XL	LL-60k	48	-	-	-	-	-	-	-	-	-	-	0.8	24

Table 1: The Number of aggregated layers ( $M$ ) and aggregation trade-off coefficient ( $\beta$ ) used by the fine-tuned models.

Model	Baseline				Layer Aggregation				Baseline				Layer Aggregation			
	test-clean		test-other		test-clean		test-other		dev-clean		dev-other		dev-clean		dev-other	
	W	C	W	C	W	C	W	C	W	C	W	C	W	C	W	C
<b>10-min labeled</b>																
W2v BASE	9.1	4.2	15.6	5.9	8.3	4.0	13.5	5.9	8.9	4.5	15.7	6.0	8.1	4.0	14.2	5.7
W2v LARGE	8.9	4.1	13.1	5.5	8.1	4.0	12.2	5.2	8.6	4.4	12.9	5.4	7.9	3.9	11.7	5.1
HuB BASE	9.7	4.2	15.3	5.8	8.6	4.2	13.3	5.8	9.1	4.5	15.0	5.6	8.4	4.3	7.5	3.7
<b>1-hour labeled</b>																
W2v BASE	5.5	2.5	11.3	4.9	5.0	2.1	9.9	4.2	5.0	2.4	10.8	4.8	4.7	1.9	10.1	4.3
W2v LARGE	5.1	2.2	9.4	4.3	4.8	1.8	9.0	4.0	4.8	1.9	8.5	4.1	4.3	1.8	8.2	4.0
HuB BASE	6.1	2.8	11.3	4.8	5.7	2.4	9.8	4.3	5.6	2.3	10.9	4.8	5.1	2.4	9.2	4.0
<b>10-hour labeled</b>																
W2v BASE	4.3	1.8	9.5	4.1	4.0	1.7	9.0	3.9	3.8	1.5	9.1	4.0	3.7	1.5	8.9	3.8
W2v LARGE	3.8	1.5	7.3	3.5	3.6	1.6	6.9	3.2	3.4	1.4	6.9	3.3	3.2	1.3	6.8	3.2
HuB BASE	4.3	1.7	9.4	4.0	4.2	1.7	9.2	3.9	3.9	1.6	9.0	4.0	3.7	1.5	3.6	1.5
<b>100-hour labeled</b>																
W2v BASE	3.4	1.5	8.0	3.7	3.3	1.4	7.5	3.5	2.7	1.2	7.9	3.7	2.3	1.0	7.5	3.6
W2v LARGE	2.8	1.2	6.0	2.7	2.5	1.0	5.7	2.4	2.3	1.0	5.7	2.5	2.2	0.9	5.3	2.2
HuB BASE	3.4	1.4	8.1	3.7	3.4	1.4	8.0	3.6	2.7	1.1	7.8	3.6	2.6	1.0	7.6	3.4
HuB LARGE	3.0	1.3	7.3	3.3	2.8	1.2	6.8	2.8	2.5	1.2	5.0	2.3	2.1	1.0	4.5	2.2
<b>360-hour labeled</b>																
W2v BASE	3.1	1.4	7.5	3.3	2.8	1.2	7.1	3.1	2.6	1.1	7.7	3.5	2.4	1.0	7.2	3.2
W2v LARGE	2.7	1.1	5.7	2.6	2.5	0.8	5.3	2.4	2.2	0.9	5.6	2.5	2.0	0.9	5.1	2.0
HuB LARGE	2.0	0.9	6.9	3.0	1.9	0.9	6.6	3.0	2.4	1.1	4.8	2.2	2.0	0.8	4.0	2.0
<b>960-hour labeled</b>																
W2v BASE	2.6	0.76	6.1	2.5	2.4	0.75	5.9	2.4	2.0	0.7	5.9	2.8	1.7	0.5	5.6	2.3
W2v LARGE	2.3	0.64	5.0	2.0	2.1	0.63	4.9	1.9	1.7	0.2	4.6	2.0	1.5	0.2	4.3	2.0
HuB BASE	2.4	0.7	6.0	2.5	2.2	0.65	5.8	2.3	1.9	0.7	5.4	2.4	1.6	0.5	5.0	2.2
HuB LARGE	1.7	0.4	3.5	1.6	1.7	0.4	3.1	1.1	1.5	0.4	3.2	1.2	1.5	0.4	3.1	1.1
HuB XL	1.6	0.4	3.0	1.1	1.4	0.4	2.9	0.9	1.7	0.4	2.6	1.0	1.4	0.3	2.3	1.0

Table 2: WER(W) and CER(C) results on Librispeech test and dev sets, computed using the Baseline and the proposed Layer Aggregation approach. The Pre-trained models, Wav2vec 2.0 (W2v) and HuBERT (HuB), were fine-tuned using varying sizes of Librispeech data. The results were obtained using beam search fused with a 4-gram LM.

using the Ax toolkit (<https://github.com/facebook/Ax>). We chose using a 4-gram LM in all of our experiments, which is widely used for speech recognition decoding (Baevski et al. 2020; Hsu et al. 2021), although our proposed aggregation method is not constrained to a specific LM and can be applied using other LM.

We tune the aggregation hyperparameters, namely, the number of aggregated layers ( $M$ ) and aggregation trade-off coefficient ( $\beta$ ) using grid search employed on the 'dev-clean' set. Table 1 summarizes the aggregation hyperparameters for each of the experimented models.

**Speech Recognition Performance.** We conduct two experimental setups: *Baseline* and *Layer Aggregation*. The baseline setup employs the logits produced using the top layer, while the layer aggregation employs the aggregated logits using the corresponding aggregation parameters. Both setups utilize the same decoder, namely, beam search fused with a 4-gram LMs.

We evaluate both setups using the Librispeech 'test-clean', 'test-other', 'dev-clean', 'dev-other' splits, and report the results using the Word Error Rate(WER) and Character Error Rate(CER) metrics.

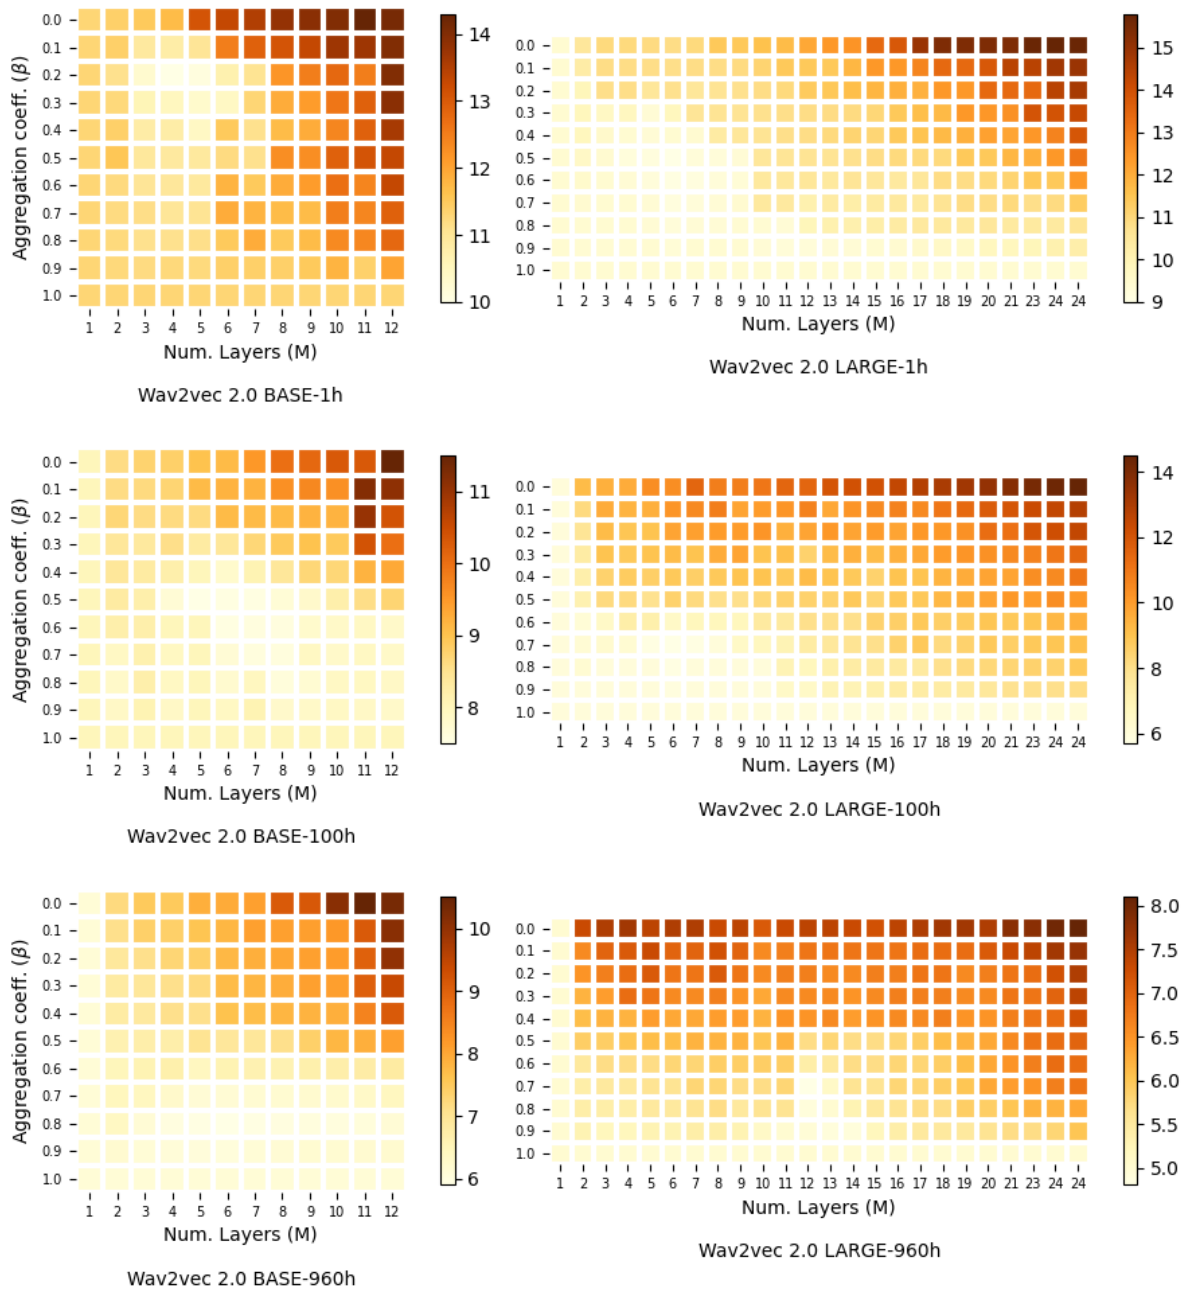


Figure 5: Layer aggregation WER results on Librispeech 'test-other' dataset. The results indicate that low-resourced models (1h and 100h) tend to rely on intermediate layers ( $\beta \leq 0.5$ ), as opposed to models fine-tuned with a large amount of labeled data (960h) that leans towards the top layer.

## Results

**Speech Recognition Performance on Librispeech.** Table 2 reports the results of the Baseline and proposed Layer Aggregation methods. The results indicate that the proposed approach either matches or improves WER and CER in all experiments. Layer aggregation is particularly effective in lower-resourced settings, i.e. when the amount of labeled

data is 10 hours or less. We find that the optimal number of aggregated layers and aggregation trade-off coefficient depends on the size of the model and the amount of labeled data used for fine-tuning, as demonstrated in Table 1. It can be observed that models fine-tuned on low-resourced labeled data rely more on the aggregation product rather than the logits produced by the top layer.

Model	Fine-tuning Data	Temp.	Agg.
Wav2vec 2.0 BASE	100h	7.9	7.5
	960h	6.0	5.9
Wav2vec 2.0 Large	100h	6.0	5.7
	960h	5.0	4.9

Table 3: WER results on Librispeech ‘test-other’, comparing temperature scaling and the layer aggregation method.

Figure 5 exhibits the relationship between the number of aggregated layers ( $M$ ) and the weight given to intermediate layers ( $1 - \beta$ ) by plotting WER values attained using two Wav2vec 2.0 variants, fine-tuned on different amounts of labeled data. We observed a recurring pattern in the experiments conducted, where models fine-tuned using a small amount of labeled data mostly rely on intermediate layers ( $\beta \leq 0.3$ ), while models fine-tuned using a larger portion of labeled data rely heavily on the top layer’s representations.

**Computation Efficiency.** We study the performance of the baseline and the layer aggregation approach using increasing beam width levels, as it is well known that computation costs and beam width are tied (Meister, Vieira, and Cotterell 2020). To this end, we experiment with two Wav2vec 2.0 BASE variants, fine-tuned using 100 and 960 hours of Librispeech. We set the aggregation parameters, namely,  $\beta$  and  $M$  according to those detailed in Table 1.

Figure 4 demonstrates the effect of the beam width on both setups. We observe that while larger beam width typically yields better performance, the aggregation method can match the baseline’s WER levels while using a significantly smaller beam width, hence, reducing inference computational costs. For example, when applied to the ‘dev-other’ set, the aggregation method can match the baseline’s performance using a width of 1500, while using a beam width of 400.

**Confidence Relaxation.** To further validate our proposed approach, we compare our method with another relaxation method, namely, Temperature scaling. For this purpose, we added a temperature scaling step (Guo et al. 2017) to the softmax operation of the top layer ( $H_N$ ) and calibrated the temperature parameter using Librispeech ‘dev-clean’ data. We conducted our experiments using two Wav2vec 2.0 BASE variants, fine-tuned using 100 and 960 hours of Librispeech. Table 3 shows a comparison of the layer aggregation method with standard Temperature scaling. Note, that besides for a single case, the temperature scaling did not show any significant performance gains. Layer aggregation outperformed temperature scaling in all of the examined cases, suggesting that harnessing intermediate layers is beneficial.

**Ablation.** Examining the errors of models fine-tuned using different sizes of labeled data, we find that the lower-resourced models’ errors (fine-tuned using 10m, 1h, 10h) are characterized by spelling errors such as dropping characters or transcribing words exactly the way they sound. The

errors of resource-rich models mostly occur with rare words. We noted that the aggregation method successfully assisted with both types of errors. Taking rare words errors for example, using the aggregation method “Phedrus” was corrected to “Phaedrus”, “Credius” to “Critias”, and “Chelsey” corrected to “Chelsea”.

Next, we wish to analyze the choices of the number of aggregated layers and the weight given to intermediate layers. The trends shown in Figure 5 show that models fine-tuned using smaller resources place their aggregation weight on intermediate layers regardless of model size, while the weight is shifted towards the top layers as the amount of fine-tuning labeled data increases. These observations are on par with the findings of Pasad *et al.* (Pasad, Chou, and Livescu 2021) which demonstrated that: (a) semantic and acoustic properties are well encoded at the top layers of models fine-tuned using large amounts of labeled data. (b) low-resource fine-tuning is insufficient for diverging the same top layers from the pre-training objective to speech recognition. The latter gives additional motivation to utilize intermediate layers in low-resource setups and further explains why layer aggregation works.

Furthermore, the proportion of the optimal number of aggregated layers remains the same across model sizes and fine-tuning resources, stretching from 33-41% of the layers in low-resource models to  $\sim 50\%$  of the layers in resource-rich models.

## Conclusions

In this study, we found that the predictions of SSL-based ASR systems are extremely confident. We then assessed the negative impact of this property on a beam search decoder, while analyzing and visualizing the prediction evolution throughout the transformer layers. Finally, we proposed an aggregation method that aggregates the top transformer layers. Our proposed approach was able to reduce WER and CER with no additional training or parameters, while reducing computational costs.

## References

- Baevski, A.; Zhou, Y.; Mohamed, A.; and Auli, M. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*.
- Chang, H.-J.; Yang, S.-w.; and Lee, H.-y. 2022. DistilHuBERT: Speech representation learning by layer-wise distillation of hidden-unit BERT. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7087–7091. IEEE.
- Chang, X.; Maekaku, T.; Guo, P.; Shi, J.; Lu, Y.-J.; Subramanian, A. S.; Wang, T.; Yang, S.-w.; Tsao, Y.; Lee, H.-y.; et al. 2021. An exploration of self-supervised pretrained representations for end-to-end speech recognition. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 228–235. IEEE.
- Chen, S.; Wang, C.; Chen, Z.; Wu, Y.; Liu, S.; Chen, Z.; Li, J.; Kanda, N.; Yoshioka, T.; Xiao, X.; et al. 2021. Wavlm:

- Large-scale self-supervised pre-training for full stack speech processing. *arXiv preprint arXiv:2110.13900*.
- Clark, K.; Khandelwal, U.; Levy, O.; and Manning, C. D. 2019. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dou, Z.-Y.; Tu, Z.; Wang, X.; Shi, S.; and Zhang, T. 2018. Exploiting deep representations for neural machine translation. *arXiv preprint arXiv:1810.10181*.
- Dou, Z.-Y.; Wang, X.; Shi, S.; and Tu, Z. 2020. Exploiting deep representations for natural language processing. *Neurocomputing*.
- Geva, M.; Schuster, R.; Berant, J.; and Levy, O. 2021. Transformer Feed-Forward Layers Are Key-Value Memories. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*.
- Gulati, A.; Qin, J.; Chiu, C.-C.; Parmar, N.; Zhang, Y.; Yu, J.; Han, W.; Wang, S.; Zhang, Z.; Wu, Y.; et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On calibration of modern neural networks. In *International conference on machine learning*, 1321–1330. PMLR.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hsu, W.-N.; Bolte, B.; Tsai, Y.-H. H.; Lakhota, K.; Salakhutdinov, R.; and Mohamed, A. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Ji, J.; Luo, Y.; Sun, X.; Chen, F.; Luo, G.; Wu, Y.; Gao, Y.; and Ji, R. 2021. Improving image captioning by leveraging intra-and inter-layer global representation in transformer network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 1655–1663.
- Kahn, J.; Rivière, M.; Zheng, W.; Kharitonov, E.; Xu, Q.; Mazaré, P.-E.; Karadayi, J.; Liptchinsky, V.; Collobert, R.; Fuegen, C.; et al. 2020. Libri-light: A benchmark for asr with limited or no supervision. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 7669–7673. IEEE.
- Kondratyuk, D.; and Straka, M. 2019. 75 languages, 1 model: Parsing universal dependencies universally. *arXiv preprint arXiv:1904.02099*.
- Liu, A. T.; Yang, S.-w.; Chi, P.-H.; Hsu, P.-c.; and Lee, H.-y. 2020a. Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6419–6423. IEEE.
- Liu, Y.; Gu, J.; Goyal, N.; Li, X.; Edunov, S.; Ghazvininejad, M.; Lewis, M.; and Zettlemoyer, L. 2020b. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Meister, C.; Vieira, T.; and Cotterell, R. 2020. Best-first beam search. *Transactions of the Association for Computational Linguistics*, 8: 795–809.
- Panayotov, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 5206–5210. IEEE.
- Pasad, A.; Chou, J.-C.; and Livescu, K. 2021. Layer-wise analysis of a self-supervised speech representation model. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 914–921. IEEE.
- Pereyra, G.; Tucker, G.; Chorowski, J.; Kaiser, Ł.; and Hinton, G. 2017. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*.
- Rogers, A.; Kovaleva, O.; and Rumshisky, A. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8: 842–866.
- Shah, J.; Singla, Y. K.; Chen, C.; and Shah, R. R. 2021. What all do audio transformer models hear? probing acoustic representations for language delivery and its structure. *arXiv preprint arXiv:2101.00387*.
- Su, T.-C.; and Cheng, H.-C. 2020. Sesamebert: Attention for anywhere. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, 363–369. IEEE.
- Tenney, I.; Das, D.; and Pavlick, E. 2019. BERT re-discovers the classical NLP pipeline. *arXiv preprint arXiv:1905.05950*.
- Van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv e-prints*, arXiv:1807.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*.
- Wang, Q.; Li, F.; Xiao, T.; Li, Y.; Li, Y.; and Zhu, J. 2020. Multi-layer representation fusion for neural machine translation. *arXiv preprint arXiv:2002.06714*.
- Xiong, H.; He, Z.; Hu, X.; and Wu, H. 2018. Multi-channel encoder for neural machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.



Yang, M.; Wang, R.; Chen, K.; Wang, X.; Zhao, T.; and Zhang, M. 2020. A novel sentence-level agreement architecture for neural machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.