



THE UNIVERSITY *of* EDINBURGH

## Edinburgh Research Explorer

### Exploration of A Self-Supervised Speech Model: A Study on Emotional Corpora

**Citation for published version:**

Li, Y, Mohamied, Y, Bell, P & Lai, C 2023, Exploration of A Self-Supervised Speech Model: A Study on Emotional Corpora. in *Proceedings of the 2022 IEEE Spoken Language Technology Workshop*. Institute of Electrical and Electronics Engineers (IEEE), pp. 868-875, The IEEE Spoken Language Technology Workshop, 2022, Doha, Qatar, 9/01/23. <https://doi.org/10.1109/SLT54892.2023.10023428>

**Digital Object Identifier (DOI):**

[10.1109/SLT54892.2023.10023428](https://doi.org/10.1109/SLT54892.2023.10023428)

**Link:**

[Link to publication record in Edinburgh Research Explorer](#)

**Document Version:**

Peer reviewed version

**Published In:**

Proceedings of the 2022 IEEE Spoken Language Technology Workshop

**General rights**

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

**Take down policy**

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [openaccess@ed.ac.uk](mailto:openaccess@ed.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# EXPLORATION OF A SELF-SUPERVISED SPEECH MODEL: A STUDY ON EMOTIONAL CORPORA

Yuanchao Li, Yumnah Mohamied, Peter Bell, Catherine Lai

Centre for Speech Technology Research, University of Edinburgh  
{yuanchao.li, ymohamie, c.lai, peter.bell}@ed.ac.uk

## ABSTRACT

Self-supervised speech models have grown fast during the past few years and have proven feasible for use in various downstream tasks. Some recent work has started to look at the characteristics of these models, yet many concerns have not been fully addressed. In this work, we conduct a study on emotional corpora to explore a popular self-supervised model – wav2vec 2.0. Via a set of quantitative analysis, we mainly demonstrate that: 1) wav2vec 2.0 appears to discard paralinguistic information that is less useful for word recognition purposes; 2) for emotion recognition, representations from the middle layer alone perform as well as those derived from layer averaging, while the final layer results in the worst performance in some cases; 3) current self-supervised models may not be the optimal solution for downstream tasks that make use of non-lexical features. Our work provides novel findings that will aid future research in this area and theoretical basis for the use of existing models.

**Index Terms**— wav2vec 2.0, self-supervised learning, speech emotion, speech recognition, paralinguistics

## 1. INTRODUCTION

Choosing the right features is a priority in machine learning-based speech tasks. How much target information the features contain fundamentally determines how well a model will work. There are a large number of features to represent and explain the complexity and variability of speech signals in extensive multi-disciplinary studies [1, 2, 3]. Task specific features have been widely and effectively used in various speech tasks. For example, cepstral features such as Mel-Frequency Cepstral Coefficients (MFCC), Linear-Frequency Cepstral Coefficients (LFCC), and Perceptual Linear Prediction (PLP) cepstral coefficients dominated Automatic Speech Recognition (ASR) for many years [4, 5]. Similarly, other speech tasks have their own preferred feature sets. In Speech Emotion Recognition (SER), suprasegmental features, such as pitch, energy, speaking rate [6, 7], have proven more helpful than information about phonetic segments. Aspects of speech which are often discarded in automatic transcription, such as disfluencies, are also known to be helpful in tasks such

as SER [8]. The same situation is true for other tasks, such as dialog act detection [9], which leads to handcrafted engineering to understand the contributions of various features.

On the other hand, directly learning feature mappings from speech signals without handcrafted engineering has emerged as a trend during the past decade. Such End-to-End (E2E) approaches benefit from the success of deep learning technologies and have proven useful in many speech tasks, including ASR, SER, speaker verification, and disorder classification [10, 11, 12, 13]. The E2E approach eliminates the separate step of feature extraction and enables joint training of multiple tasks due to shared representations. This can allow the models to learn feature spaces that are more representative of the actual task than handcrafted features.

Inspired by the success of Self-Supervised Learning (SSL) in natural language processing [14, 15], work on addressing the general lack of task-specific labeled speech data has accelerated in the past few years. Most of these approaches can be divided into generative modeling approaches and discriminative modeling approaches [16, 17, 18]. SSL utilizes information extracted from the input data itself as the label to learn to encode general-purpose representations. These pre-trained upstream models have proven effective for downstream speech tasks, including speaker verification [19, 20] and SER [21]. However, what these models are actually learning is still understudied and questions and concerns remain about why and how these models benefit downstream tasks: *Are the generated representations optimal for every task? How to utilize them for different purposes?*

With these questions in mind, we study wav2vec 2.0 [17] on emotional corpora, demonstrating how this type of self-supervised model can be explored for downstream tasks. Our experiments show that: 1) wav2vec 2.0 appears to discard some paralinguistic information that is less useful for word recognition purposes and does not treat all emotions and paralinguistic features equally; 2) for SER, representations from the final layer could result in the worst performance in some cases; 3) current self-supervised models need to be carefully fine-tuned to adapt to downstream tasks that make use of non-lexical features. We hope our findings can provide the research community with a new perspective to look at the effectiveness and usage of self-supervised models.

## 2. RELATED WORK

There is no doubt that large-scale speech models using SSL are becoming integral in speech processing tasks. Most of them can be divided into generative or discriminative approaches. The generative approaches generate future frames from past frames, or masked frames from unmasked frames by learning to minimize reconstruction loss [16, 22, 23]. On the other hand, the discriminative approaches discriminate positive samples from negative samples while minimizing contrastive prediction loss [24, 17, 18]. These self-supervised models are generally trained on Librispeech [25], a corpus based on public domain audio books primarily used for ASR research. Although self-supervised objectives are general, the design of popular SSL models has been primarily driven by the goal of improving automatic transcription.

Unlike traditional speech modeling approaches that have been extensively researched, these SSL models have just started to be explored in very recent years, with wav2vec 2.0 (W2V2) attracting the most attention for its wide application potential. For example, Pasad et al. [26] conducted layer-wise analysis of W2V2 using a suite of tools and found 1) acoustic and linguistic properties are encoded in different layers; 2) the pre-trained model follows an autoencoder-style behavior; 3) the model encodes some non-trivial word meaning information. Fan et al. [20] showed that W2V2 has the ability to discriminate between speakers and also languages, and this distinction is more obvious in lower layers. They hence proposed multi-task learning of speaker verification and language identification, and verified its feasibility. Li et al. [27] noticed the recognition of longer emotional utterances that contain more contextual information benefits from the contextual characteristic of W2V2. They proposed a joint training scheme by hierarchically fusing multiple W2V2 outputs for SER. Yang et al. [28] set up benchmark performance using self-supervised speech models on a range of tasks.

Nevertheless, these self-supervised speech models are still understudied and the above-mentioned works have limitations. For example, [26] did not extend their exploration to downstream tasks. In [20] and [27], only a portion of the layer difference was shown, so misses a thorough layer-wise analysis. In [28], they presented downstream task performance without further explanation. Furthermore, none of those studies investigated paralinguistic characteristics in W2V2 representations. As such, in the current work, we build on previous work while adding new perspectives from detailed quantitative analysis on emotional corpora.

## 3. CORPORA AND MODEL DESCRIPTION

### 3.1. Corpora

**IEMOCAP** (IEM) [29] has five dyadic sessions with ten actors (five male and five female), each with a scripted and improvised multimodal interaction. The corpus consists of ap-

proximately 12 hours of speech that has been annotated by three annotators with ten emotion classes. Following prior research [27], we combined *Happy* and *Excited*, and removed utterances that do not have transcripts, bringing the total number of utterances used in this study to 5,500, each with one label from four classes: *Angry*, *Happy*, *Neutral*, and *Sad*.

**RAVDESS** (RAV) [30] contains a speech set and a song set. We only use the speech set, which has 1,440 utterances from 24 actors (12 female, 12 male) in eight emotions: *Calm*, *Happy*, *Sad*, *Angry*, *Fear*, *Surprise*, *Disgust*, and *Neutral*. Ratings were provided by untrained individuals. In the process of collecting data, the actors spoke two fixed sentences with different classes of emotion, so the corpus has a good balance of emotions. The actors were given two trials for each utterance and asked to produce 60 speech clips in total.

The major reason that we choose to use RAV is that, even though other corpora may have a larger size, it provides fixed sentences with different emotional expressions. Such a setting excludes the lexical influence by “forcing” different emotions to have the same linguistic content, thus helping us to better explore the effects of the acoustic properties of W2V2 by eliminating the effects raised by lexical content (e.g., word pronunciation causing prosody variation).

### 3.2. Model

We look at W2V2 [17], a SSL framework comprised of three major components: a CNN-based local encoder that extracts a sequence of embeddings from raw audio as latent speech representation  $Z$ , a Transformer network for obtaining context representation  $C$ , and a quantization module for discretizing  $Z$  into  $Q$ . Following previous work [26], we focus our attention on the latent representations learned by the Transformer module of W2V2. In this work, we use *wav2vec2-base*, *wav2vec2-base-100h*, and *wav2vec2-base-960h* models, which are the pre-trained and fine-tuned models (on 100h and 960h of Librispeech) respectively. We refer to them as *PT*, *FT100*, and *FT960*. We choose W2V2 because it is the most widely used SSL speech model, with the expectation that the exploratory approach can be generalized to similar SSL models.

## 4. EXPERIMENTS AND RESULTS

We perform a set of probing experiments, including the following quantitative measures:

**Probing SER performance.** We first implement a layer-wise analysis by using the output of every individual layer within the Transformer network to demonstrate how information encoded by W2V2 contributes to SER. Next, as there is no common practice of how to utilize W2V2 representations as input features for downstream tasks, we compare the performance of three commonly used approaches of using W2V2 representations as input features: 1) taking the last layer output [31, 32, 33]; 2) taking the average of all layer outputs [34];

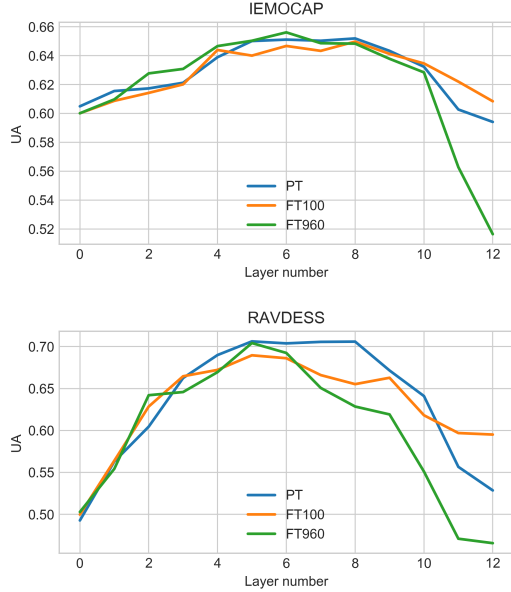


Fig. 1. SER accuracy comparison using different models.

3) taking the weighted average of all layer outputs (assigning a trainable weight to each layer output) [21, 28]. We also propose a fourth approach which excludes the last two layers from averaging as they generally underperform other layers. We evaluate the performance using Unweighted Accuracy (UA). Like most downstream tasks, we use W2V2 models as frozen feature extractors. Since our goal is to explore information in W2V2 representations, we build a simple downstream model comprising only two dense layers (128 and 16 neurons, respectively) with *ReLU* activation and one output layer (four neurons for IEM and eight neurons for RAV) with *Softmax* activation. The learning rate is set as  $1e-4$  and  $2e-4$  for IEM and RAV with the *AdamW* optimizer, respectively, and the weight decay is set as  $1e-5$ . The batch size is 64, and we train the models until validation loss converges, as different layer outputs converge at different steps. For IEM, we implement 5-fold cross-validation in accordance with prior works. For RAV, we randomly divide 24 speakers into four groups and implement 4-fold cross validation.

Fig. 1 depicts the trends of layer-wise UA on the two corpora. We include layer 0 (the output of the CNN encoder right before the Transformers) in accordance with prior works. We see that: **1)** Before the best middle layer (layer 6 for IEM and layer 5 for RAV), all three models (*PT*, *FT100*, and *FT960*) show the same trend: accuracies go up and are relatively close, but then start dropping after the middle layer. This upward-downward trend is possibly related to the acoustic-linguistic property of W2V2 [26]. Raw frame-level inputs are encoded by the Transformers until the middle layer. At this point, the representations encode phonetic information but have not yet lost much of the original acoustic properties. This makes the middle layers contain the most useful information for SER. In subsequent layers, the represen-

Table 1. UA (%) using different inputs and models.

Input	Model	Corpus	
		IEM	RAV
Best layer	<i>PT</i>	65.19	70.62
	<i>FT100</i>	64.97	68.96
	<i>FT960</i>	65.61	70.42
Last layer	<i>PT</i>	59.41	52.85
	<i>FT100</i>	60.84	59.51
	<i>FT960</i>	51.64	46.56
Average	<i>PT</i>	64.93	67.26
	<i>FT100</i>	64.72	67.40
	<i>FT960</i>	65.51	64.20
Avg. w/o last two	<i>PT</i>	65.11	67.36
	<i>FT100</i>	64.90	67.50
	<i>FT960</i>	65.87	65.56
Weighted average	<i>PT</i>	65.28	68.47
	<i>FT100</i>	64.94	68.89
	<i>FT960</i>	65.67	65.11

tations gradually encode word identity and word meaning more strongly. At this stage, potential ASR errors with the loss of the original acoustic information have been appeared to lead to drops in SER accuracy. **2)** On IEM, there are barely any differences among the UAs until layer 11, while on RAV, the differences after the middle layer are more dramatic. This phenomenon is plausible as RAV only has two fixed statements, yet IEM contains various sentences, allowing fine-tuned W2V2 models to make more use of linguistic information, which makes up for the acoustic loss. In RAV, however, every sentence is repeated with each emotion, which means linguistic information has no contribution to emotion discrimination. **3)** In general,  $PT > FT100 > FT960$  from the middle layer but *FT100* clearly outperforms the other two on the last two layers. We assume that moderate fine-tuning enables W2V2 to achieve a good acoustic-linguistic balance, as word information has been found encoded by the last two layers in fine-tuned models [26], and such a balance helps *FT100* achieve better performance on the last two layers.

Table 1 compares the SER accuracies using different inputs and models, and yields the following findings: **1)** UAs of the best layer of *PT* and *FT960* are close and higher than *FT100*. So, while moderate fine-tuning enables the model to capture both acoustic and linguistic properties for *FT100*, neither of them is fully encoded causing a decrease in accuracy for the best layer. **2)** The situation is reversed on the last layer. Compared to *FT100*, *PT* lacks linguistic information while *FT960* relies too much on imperfect linguistic information due to “over” fine-tuning. **3)** Word-level information does not help SER on RAV, as mentioned before, which makes the deeper layers of *FT960* the worst. Hence, it is reasonable that *FT960* generates better performance on IEM yet worse performance on RAV by either taking the average on all layers, the average without the last two layers, or the weighted average on all layers as input. **4)** Corpora like RAV, whose emotions depend only on acous-




Fig. 2. CCA similarity comparison for paralinguistic property.

Table 2. Extracted paralinguistic features.

Feature set	Low-level descriptors
Energy	Loudness; Harmonics-to-noise ratio
Frequency	Pitch; Formant 1; Formant 1, 2, 3 frequency
Spectral	Alpha ratio; Hammarberg index; Formant 1, 2, and 3 relative energy; Spectral slope 0-500 Hz, 500-1500 Hz; Harmonic difference H1-H2 and H1-A3
Voice quality	Jitter; Shimmer

tics, may only need the best middle layer as input for SER. However, it is hard to say which input is obviously the best for corpora like IEM, whose emotions are also affected by text. 5) Except for the “best” layer inputs, *FT960* and *FT100* produce the best results on IEM and RAV, respectively. This differs from the patterns in Fig. 1 from which we would expect *FT960* to perform the worst on IEM and *PT* the best on RAV. It means that the performance obtained by averaging layer outputs does not equal the average of all layer performance, which demonstrates that representations of different layer contain different information contributing to SER.

**Probing paralinguistic information.** In this experiment, we measure the similarities between each layer’s output and different types of paralinguistic features to see how W2V2 retains well-known acoustic correlates of speech perception. We evaluate the similarity using Canonical Correlation Analysis (CCA) [35]. The selected paralinguistic features are listed in Table 2, which are mainly based on eGeMAPS [36], commonly used as a minimal set of features for SER. We also extract MFCC as linguistic (phone) features for comparison. We downsample them to make their sequence lengths comparable to W2V2 representations as required by CCA.

Fig. 2 shows the layer-wise CCA on IEM and RAV using three W2V2 models. 1) The curves of paralinguistic fea-

tures are flatter than those of MFCC, indicating that W2V2 does not focus on as much paralinguistic information. In particular, voice quality features seem barely taken into the encoding process. 2) When looking at *PT* models, we can see a reverse trend from the middle layer, which reinforces a view that the *PT* model follows an autoencoder style behavior where deeper layers “reconstruct” the input [26]. The reverse trend of similarity with MFCC on IEM is weaker than that on RAV, possibly demonstrating that the linguistic complexity makes the reconstruction process harder and more error-prone. Hence, the deeper representations are more accurate and more similar to MFCC on RAV than on IEM. The peculiar pattern on the last two layers is due to the training objective of masked segment prediction (cf. [26]). 3) When looking at the fine-tuned models, we note that the similarities keep decreasing since the models have been fine-tuned towards ASR and learn to compute speech information from frame to phoneme, and then to word level with layer depth [26]. This phenomenon reinforces our explanation of the accuracy drop in Fig. 1 that acoustic properties are being replaced by linguistic ones that contain errors. 4) The graphs indicate that the overall similarity variation on IEM is larger than on RAV. This is again, likely due to the fact that RAV has much less linguistic variation overall, which we in turn see as less change in CCA through the layers. Moreover, since how we say a sentence is affected by its content, the CCA variation of paralinguistic features is larger on IEM than on RAV. Finally, since the layer outputs contain more complex linguistic information, the overall CCA values for paralinguistic features on IEM are lower than those on RAV, no matter the starting values or overall values. However, the starting values of similarities with MFCC on both corpora are almost the same, further suggesting that W2V2 focuses on learning linguistic information rather than paralinguistic.




Fig. 3. Pair-wise correlations of layer representations.

Table 3. Correlations between the last two and prior layers.

L.	Model	L. 6	L. 7	L. 8	L. 9	L. 10
11	PT	0.25	0.24	0.24	0.31	0.45
	FT100	0.29	0.30	0.31	0.38	0.53
	FT960	0.24	0.26	0.27	0.32	0.43
12	PT	0.26	0.25	0.26	0.31	0.43
	FT100	0.27	0.28	0.29	0.34	0.44
	FT960	0.22	0.22	0.23	0.26	0.32

**Probing layer correlation.** To better understand how different layer outputs are correlated with each other before and after fine-tuning, and how W2V2 encodes information and contributes to SER, we calculate pair-wise CCA similarities of W2V2 representations from every layer and plot the similarities using heat maps to visualize the correlations. We only discuss IEM, as the same patterns are found on RAV.

Fig. 3 shows the pair-wise correlations. **1)** In *PT* model, we can notice an arrow-like shape from layers 0–10. Specifically, layer 9 and 10 have higher correlations with shallow layers compared to those in *FT100* and *FT960*. This shows that representations start becoming more general in those two layers, indicating pre-trained W2V2 is indeed getting previous information back as assumed by [26], and also explains the reverse trend in *PT* model. However, such a pattern seems specific to these two layers as it is not obvious in prior deep layers, even layer 8, which also accounts for the reverse trend (Fig. 2 and Fig. 3 in [26]). After fine-tuning, this phenomenon disappears. The two layers become the same as shallow layers that have high correlations with nearby layers, and the correlations get weaker with distance. **2)** In *PT* model, the last two layers are highly correlated with each other, even more obvious than the prior adjacent layers. However, the correlation gets weaker as fine-tuning goes. The color becomes dimmer in *FT100* and further dimmer in *FT960*. Nevertheless, an interesting phenomenon appears in *FT100*: the similarities between last two layers (especially layer 11) and the prior deep layers (layer 6 to 10) become higher. We present the CCA values in Table 3. It is obvious that the correlations in *FT100* are the highest, but they decrease in *FT960*. Moreover, the lower right part is slightly brighter in *FT100* than in *FT960* (values are skipped as limited space) indicating the deeper layers are more correlated with each other, which means low-level linguistic information (e.g., phonetics) generally exists. These phenomena validate our assumption that moderate fine-tuning enables W2V2 to achieve a good acoustic-linguistic balance but over fine-tuning “forces” the model to concen-





Fig. 4. Hierarchical CCA similarity differences.

trate on learning high-level linguistic properties (e.g., word meaning) towards ASR.

**Probing hierarchical property.** Since SSL enables frames to capture context information, the representations are expected to contain higher-level meanings. To verify this, we prepare the extracted paralinguistic features at frame, phoneme, and word levels and measure their similarities with W2V2 representations using CCA, respectively. As our purpose is only to verify whether W2V2 features contain high-level speech information, we do not use force alignment to determine the perfect boundaries. Instead, we adopt a less accurate yet efficient approach to compute hierarchical features: we constitute a phoneme by averaging five consecutive frames (we also tried to add overlap, but the results didn’t make much difference): a word by averaging five consecutive phonemes, based on the fact that frame length is set as 25ms when being extracted, and phoneme length varies from 50ms to 200ms (five frames on average) and word length from 250ms to 1,000ms (five phonemes on average) in IEM [37]. We use all the paralinguistic features provided by eGeMAPS and implement the composition of hierarchical features. Then we downsample the paralinguistic features or the W2V2 representations depending on their lengths to make them comparable. Finally, we compute the CCA differences ( $CCA_{phoneme} - CCA_{frame}$  and  $CCA_{word} - CCA_{phoneme}$ ) which represent how similar the higher-level features with W2V2 representations are compared to the lower-level ones.

From Fig 4, we can note that **1)** higher-level paralinguistic features do have higher similarities with W2V2 representations as the difference values are all positive. Besides, the value of  $CCA_{word} - CCA_{phoneme}$  is even higher than  $CCA_{phoneme} - CCA_{frame}$ , which means W2V2 representations are more similar to word-level paralinguistic information. **2)** The CCA differences barely change until layer 11 and become larger in the last two layers, which is due to the masked segment prediction enabling them to capture more context information (which is high level), especially on layer 11. **3)** The curves of fine-tuned models are flatter because the last two layers become more coherent with the other layers by fine-tuning. Note, since the paralinguistic property is affected





**Fig. 5.** Discriminative analysis for emotion bias.

by linguistics in IEM, the patterns are not as clear as RAV, yet we observed similar trends, so do not discuss it here.

**Probing emotion bias.** Different emotions have different paralinguistic patterns [38, 39]. For example, hot angry and happy emotions usually have high intensity and pitch, while sad and calm emotions have low intensity [40]. Hence, we calculate CCA similarities between paralinguistic features with W2V2 representations of every emotion for discriminative analysis. We also use all the paralinguistic features in eGeMAPS as in the previous task.

As illustrated in Fig 5: **1)** Higher similarities between paralinguistic features with W2V2 representations are found in *Neutral* emotion for both the *PT* model and the *FT* models, pointing to interesting observations: a) *Neutral* is likely more frequently represented within Librispeech, as it is a corpus of read audiobooks where most emotional cues arise only within speech of fictional characters, i.e. bias in the data during pre-training of W2V2 consequently results in learned representations which are emotion-agnostic; b) the pre-training pretext task in W2V2 (predicting masked segments) is not sufficient to learn a truly generalized representation in which different emotions are captured effectively. We also see that the curves converge after the middle layer on *PT* model. This again indicates that the deeper layers (except the last two) of *PT* model reconstruct the acoustic input. **2)** The curves become even less distinguishable at the last two layers, indicating again the autoencoder type of learning resulting from the masked segment prediction does not help distinguish emotions. This may be because the paralinguistic information of a masked segment is difficult to predict from unmasked segments, as paralinguistic information is more spontaneous and less contextual compared to linguistic information. The masked segment prediction discriminates linguistic information while blurring the paralinguistic difference among frame segments, which

makes the paralinguistic properties of every emotion become similar, resulting in the close curves. **3)** For the *FT960* model, the distances between the curves increase with depth. It seems that W2V2 not only avoids encoding paralinguistic information, but consistently discards some paralinguistic features as learning proceeds, especially in speech that contains rich paralinguistic information, e.g., voice quality, which leads to increasing differences in the similarities between W2V2 representations and paralinguistic features across emotions (otherwise, the distances should not change with layer depth).

## 5. DISCUSSION

1) Fine-tuning affects W2V2 by transforming it from an acoustic-aware model into a linguistic-aware model. Layers of the first half of the Transformer are responsible for encoding acoustic information, as all three models show almost the same patterns. The latter half starts encoding linguistic information as pattern differences occur, but an exception is that the last two layers of *PT* model reconstruct the input.

2) W2V2 should be used with caution on downstream tasks because it potentially loses important paralinguistic information. As information that is not helpful to ASR is discarded with layer depth on fine-tuned models, *PT* model is a better choice for tasks that are largely paralinguistic-dependent. Moreover, the best layer outperforms layer averaging for SER, while the last layer could be the worst choice.

3) While W2V2 (and possibly other similar SSL models) is a universal solution, it is not suitable for all downstream tasks. It does not, for example, outperform previous SER works that take raw signals as input but use less sophisticated end-to-end structures [12]. Besides, as some paralinguistic information is largely involved in pragmatics such as turn-taking and backchanneling [41], the deep layers of W2V2 may not be able to model these dialog-level functions.

## 6. CONCLUSIONS

In this work, we study W2V2 by conducting a set of quantitative analysis on emotional corpora. We found W2V2 lacks the ability to capture paralinguistic information. We also contribute to understanding the types of representations W2V2 learns by thoroughly comparing layer outputs in their correlations and SER. The hierarchy and emotion analysis pave the path for better usage of W2V2 on downstream tasks. Our results confirm the assumptions of previous work and strengthen their conclusions, as well as provide novel findings towards better leveraging of SSL speech models.

## 7. ACKNOWLEDGMENTS

We would like to thank the authors of [26] for their kind help and valuable discussion.

## 8. REFERENCES

- [1] Gunnar Fant, “Speech sounds and features,” *The MIT Press*, 1973.
- [2] Jacob Benesty, M Mohan Sondhi, Yiteng Huang, et al., *Springer handbook of speech processing*, vol. 1, Springer, 2008.
- [3] Oudeyer Pierre-Yves, “The production and recognition of emotions in speech: features and algorithms,” *International Journal of Human-Computer Studies*, vol. 59, no. 1-2, pp. 157–183, 2003.
- [4] Steven Davis and Paul Mermelstein, “Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences,” *IEEE transactions on acoustics, speech, and signal processing*, vol. 28, no. 4, pp. 357–366, 1980.
- [5] Hynek Hermansky, “Perceptual linear predictive (PLP) analysis of speech,” *the Journal of the Acoustical Society of America*, vol. 87, no. 4, pp. 1738–1752, 1990.
- [6] Robert W Frick, “Communicating emotion: The role of prosodic features,” *Psychological bulletin*, vol. 97, no. 3, pp. 412, 1985.
- [7] Jaime C Acosta and Nigel G Ward, “Achieving rapport with turn-by-turn, user-responsive emotional coloring,” *Speech Communication*, vol. 53, no. 9-10, pp. 1137–1148, 2011.
- [8] Johanna D Moore, Leimin Tian, and Catherine Lai, “Word-level emotion recognition using high-level features,” in *International Conference on Intelligent Text Processing and Computational Linguistics*. Springer, 2014, pp. 17–31.
- [9] Raul Fernandez and Rosalind W Picard, “Dialog act classification from prosodic features using support vector machines,” in *Speech prosody 2002, International conference*, 2002.
- [10] Alex Graves and Navdeep Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *International conference on machine learning*. PMLR, 2014, pp. 1764–1772.
- [11] David Snyder, Pegah Ghahremani, Daniel Povey, Daniel Garcia-Romero, Yishay Carmiel, and Sanjeev Khudanpur, “Deep neural network-based speaker embeddings for end-to-end speaker verification,” in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 165–170.
- [12] Yuanchao Li, Tianyu Zhao, and Tatsuya Kawahara, “Improved End-to-End Speech Emotion Recognition Using Self Attention Mechanism and Multitask Learning,” in *Interspeech*, 2019, pp. 2803–2807.
- [13] Yilin Pan, Bahman Mirheidari, Zehai Tu, Ronan O’Malley, Traci Walker, Annalena Venneri, Markus Reuber, Daniel Blackburn, and Heidi Christensen, “Acoustic feature extraction with interpretable deep neural network for neurodegenerative related disorder classification,” in *Proceedings of Interspeech 2020*. International Speech Communication Association (ISCA), 2020, pp. 4806–4810.
- [14] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer, “Deep Contextualized Word Representation,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, New Orleans, Louisiana, June 2018, Association for Computational Linguistics.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [16] Yu-An Chung, Wei-Ning Hsu, Hao Tang, and James Glass, “An unsupervised autoregressive model for speech representation learning,” *arXiv preprint arXiv:1904.03240*, 2019.
- [17] Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 12449–12460, 2020.
- [18] Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [19] Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Zhuo Chen, Peidong Wang, Gang Liu, Jinyu Li, Jian Wu, Xiangzhan Yu, et al., “Why does Self-Supervised Learning for Speech Recognition Benefit Speaker Recognition?,” *arXiv preprint arXiv:2204.12765*, 2022.
- [20] Zhiyun Fan, Meng Li, Shiyu Zhou, and Bo Xu, “Exploring wav2vec 2.0 on speaker verification and language identification,” *arXiv preprint arXiv:2012.06185*, 2020.
- [21] Leonardo Pepino, Pablo Riera, and Luciana Ferrer, “Emotion recognition from speech using wav2vec 2.0 embeddings,” *arXiv preprint arXiv:2104.03502*, 2021.



- [22] Yu-An Chung, Hao Tang, and James Glass, “Vector-quantized autoregressive predictive coding,” *arXiv preprint arXiv:2005.08392*, 2020.
- [23] Alexander H Liu, Yu-An Chung, and James Glass, “Non-autoregressive predictive coding for learning speech representations from local dependencies,” *arXiv preprint arXiv:2011.00406*, 2020.
- [24] Aaron Van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” *arXiv e-prints*, pp. arXiv–1807, 2018.
- [25] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an ASR corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [26] Ankita Pasad, Ju-Chieh Chou, and Karen Livescu, “Layer-wise analysis of a self-supervised speech representation model,” in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 914–921.
- [27] Yuanchao Li, Peter Bell, and Catherine Lai, “Fusing asr outputs in joint training for speech emotion recognition,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7362–7366.
- [28] Shu-wen Yang, Po-Han Chi, Yung-Sung Chuang, Cheng-I Jeff Lai, Kushal Lakhotia, Yist Y Lin, Andy T Liu, Jiatong Shi, Xuankai Chang, Guan-Ting Lin, et al., “Superb: Speech processing universal performance benchmark,” *arXiv preprint arXiv:2105.01051*, 2021.
- [29] Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan, “IEMOCAP: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [30] Steven R Livingstone and Frank A Russo, “The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English,” *PLoS one*, vol. 13, no. 5, pp. e0196391, 2018.
- [31] Li-Wei Chen and Alexander Rudnicky, “Exploring Wav2vec 2.0 fine-tuning for improved speech emotion recognition,” *arXiv preprint arXiv:2110.06309*, 2021.
- [32] Mayank Sharma, “Multi-lingual multi-task speech emotion recognition using wav2vec 2.0,” in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6907–6911.
- [33] Xingyu Cai, Jiahong Yuan, Renjie Zheng, Liang Huang, and Kenneth Church, “Speech Emotion Recognition with Multi-Task Learning,” in *INTERSPEECH*, 2021.
- [34] Jonathan Boigne, Biman Liyanage, and Ted Östrem, “Recognizing more emotions with less data using self-supervised transfer learning,” *arXiv preprint arXiv:2011.05585*, 2020.
- [35] David R Hardoon, Sandor Szedmak, and John Shawe-Taylor, “Canonical correlation analysis: An overview with application to learning methods,” *Neural computation*, vol. 16, no. 12, pp. 2639–2664, 2004.
- [36] Florian Eyben, Klaus R Scherer, Björn W Schuller, Johan Sundberg, Elisabeth André, Carlos Busso, Laurence Y Devillers, Julien Epps, Petri Laukka, Shrikanth S Narayanan, et al., “The Geneva minimalistic acoustic parameter set (GeMAPS) for voice research and affective computing,” *IEEE transactions on affective computing*, vol. 7, no. 2, pp. 190–202, 2015.
- [37] Weidong Chen, Xiaofen Xing, Xiangmin Xu, Jianxin Pang, and Lan Du, “SpeechFormer: A Hierarchical Efficient Framework Incorporating the Characteristics of Speech,” *arXiv preprint arXiv:2203.03812*, 2022.
- [38] Xi Li, Jidong Tao, Michael T Johnson, Joseph Soltis, Anne Savage, Kirsten M Leong, and John D Newman, “Stress and emotion classification using jitter and shimmer features,” in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP’07*. IEEE, 2007, vol. 4, pp. IV–1081.
- [39] Marko Lugger and Bin Yang, “The relevance of voice quality features in speaker independent emotion recognition,” in *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP’07*. IEEE, 2007, vol. 4, pp. IV–17.
- [40] Iker Luengo, Eva Navas, and Inmaculada Hernáez, “Feature analysis and evaluation for automatic emotion identification in speech,” *IEEE Transactions on Multimedia*, vol. 12, no. 6, pp. 490–501, 2010.
- [41] Nigel G Ward, Yuanchao Li, Tianyu Zhao, and Tatsuya Kawahara, “Interactional and pragmatics-related prosodic patterns in Mandarin dialog,” in *Speech prosody*, 2016.