

Do VSR Models Generalize Beyond LRS3?

Yasser Abdelaziz Dahou Djilali^{1,2} Sanath Narayan¹ Eustache Le Bihan¹
 Haithem Boussaid¹ Ebtesam Almazrouei¹ Merouane Debbah¹
¹Technology Innovation Institute, UAE ²Dublin City University, Ireland

Abstract

The Lip Reading Sentences-3 (LRS3) benchmark has primarily been the focus of intense research in visual speech recognition (VSR) during the last few years. As a result, there is an increased risk of overfitting to its excessively used test set, which is only one hour duration. To alleviate this issue, we build a new VSR test set named WildVSR, by closely following the LRS3 dataset creation processes. We then evaluate and analyse the extent to which the current VSR models generalize to the new test data. We evaluate a broad range of publicly available VSR models and find significant drops in performance on our test set, compared to their corresponding LRS3 results. Our results suggest that the increase in word error rates is caused by the models' inability to generalize to slightly "harder" and in the wild lip sequences than those found in the LRS3 test set. Our new test benchmark is made public in order to enable future research towards more robust VSR models.

1. Introduction

The primary objective of machine learning revolves around training models that are capable of better generalization. Typically, the quantification of generalization occurs through evaluating a model's performance on a held-out test set. The question then arises: what does satisfactory performance on this test set indicate? At the very least, it is desirable that the model exhibits similar performance on a new test set derived from the same data creation process. In this work, we study these questions for the problem of Visual Speech Recognition (VSR).

Indeed, most perception problems are interpolative in their nature [5] and satisfy the manifold hypothesis [10]. These tasks are intuitive for humans, and are usually solved in the early layers of the visual cortex in a matter of milliseconds (*i.e.*, classification, recognition, *etc.*) [17, 39]. For

such problems, deep learning is a perfect fit with its ability to perform non-linear interpolation in a complex high-dimensional manifold, enabling arbitrary complex behavior [5, 38], allowing for better generalization. However, lip-reading experts allude to a high-level step-wise and iterative reasoning to solve the task. This likely suggests that VSR has some higher level of extrapolation as compared to the common perception tasks. Thus, extensive model's generalization study is highly needed for the field. However, the one hour LRS3 test set is the main focus for evaluation state-of-the-art models.

In this paper, we follow the LRS3 [1] creation process to build a new set from the wild. As expected, we observe that all VSR SoTA models fail to reach their reported Word Error Rate (WER) from LRS3 on the new test set. Nevertheless, the WER drops by an average of 30 points, our experiments reveal that the comparative ranking of models remains remarkably consistent when evaluated on our fresh test sets. Specifically, the models that exhibit the highest accuracy on the original test sets also demonstrate the highest accuracy on the new test sets. This suggests that the WER drops are not a result of extensive hyper-parameters tuning that fit the particular lip sequences found in the initial test set. We study why this phenomenon arises, specifically, with the following contributions:

- We present a new VSR test set, WildVSR, incorporating higher visual diversity, and spoken vocabulary.
- We benchmark existing models on the new test set, and find a clear performance drop.
- Nevertheless, we show a diminishing return in performance *vs.* compute, where self-supervised approaches consume significantly higher training compute budget for a moderate performance.
- We propose a new metric that accounts for a model's confidence, which improves the WER based ranking.

2. Related Work

State-of-the-art approaches: The work of [22] proposed a curriculum learning approach, where shorter sequences are

Code: https://github.com/YasserdahouML/WildVSR_test_set

initially used for training followed by progressively adding longer ones. Differently, VTP [27] proposed sub-words learning scheme using frame-word boundaries to crop out training samples for a better convergence. These training strategies are computationally demanding and hard to scale to larger datasets. The recent works of [2, 23] proposed to exploit the audio latent representations as part of a auxiliary task, where the latent features from the encoder are optimized to predict pretrained ASR representations, in addition to decoding the target text. This extra supervision through the ASR representations makes the optimization more stable and improves the convergence. Intuitively, if the transformer encoder is able to match the audio features statistics in earlier layers, it is easier to adjust the attention weights in the later layers for improved text decoding.

Another line of research leverages cross-modal pretraining on large datasets in a self-supervised way (SSL), followed by finetuning on labeled video-text pairs [13, 21, 35, 36, 41]. The work of AV-HuBERT [35] fuses the masked audio-visual representations to predict the cluster assignments created from the audio features, thus, distilling knowledge from the audio stream features to model visual inputs. VATLM [41] extends this by attempting to unify the modalities using a one tower design, where a single network is optimized to construct a common representation space for video, audio and text. This is achieved by setting a unified tokenizer for all modalities, and then performing the masked prediction task over the unified tokens. The works of [24, 34] designed cross-modal self-supervised learning frameworks by adopting contrastive learning [14] to learn discriminative visual representations that appear to improve VSR performance and generalization. RAVen [13] designed an asymmetric SSL framework to induce a one-way knowledge distillation, where the audio networks predict both audio and video representations, whereas the visual network is restricted to predict the audio features only. This forces the audio network to serve as a strong teacher, as it would adjust to both modalities at the same time.

More recently, Auto-AVSR [20] proposed to obtain text transcripts for unlabelled datasets using pre-trained ASR models. These auto-labeled video-text pairs along with manually labeled datasets were then utilized to train conformer-based AVSR models in a fully-supervised manner. Furthermore, SynthVSR [19] first learned a speech-driven lip animation model on unlabeled audio-visual datasets and generated a synthetic dataset, which was utilized along with labeled and auto-labeled datasets, similar to [20], for training the VSR model.

Datasets: Typically, most VSR approaches exploit the popular LRS3 [1] dataset for training the models. While the use of LRW [7] and LRS2 [37] is lesser due to their restrictive licences, different strategies are commonly employed to increase the training data. *E.g.*, using unlabeled/machine-

labeled (AVSpeech [9], VoxCeleb2 [6]), generating synthetic videos (3.6k hours in [19]), collecting large-scale non-public datasets (YT31k [25], YT90k [31]), *etc.* However, for evaluating the VSR models, the aforementioned works extensively utilize the publicly-available LRS3 test set of one hour duration, thereby increasing the risk of models overfitting to LRS3 test set distribution. While the works of [4, 25, 31] additionally evaluate on YTDEV18 [25] or MEET360 [4], these test benchmarks are private. In this work, we set out to gather a challenging evaluation benchmark for VSR and analyse the performance of available models in-depth. Our proposed benchmark will be made public for aiding future research in VSR to build more robust models that generalize better to videos in the wild.

3. Building the WildVSR Test Set

Utilizing YouTube as a data source has become a popular approach for constructing audio and/or visual speech recognition datasets since this platform offers access to a vast amount of audio-visual content. However, developing a quality VSR test set requires careful processing of raw videos to create pairs of visual utterances, namely visual lip movement, and transcriptions. Given the computational cost associated with this procedure, it is crucial to meticulously filter and extract relevant content from the vast YouTube repository before processing it. Thus, the overall approach should consist of two stages: 1) select relevant YouTube videos, and 2) process the selected videos to construct the aforementioned pairs. Regarding the first stage, previous works have employed two different approaches.

First, in the work by Makino *et al.* [25], there is no mention of their video selection method. They extract snippets from YouTube videos where audio and transcripts match and then process these snippets to extract face tracks that align with the transcript. However, this approach does not guarantee that the initially extracted snippets from YouTube videos will display lip movements at all, as no strategy to select those videos is used. As a result, all extracted snippets have to be processed in search of face tracks while they may not be present at all, *e.g.*, a screen recorded tutorial might not have a face track although the audio matches with the corresponding transcript. This can impose a significant computational burden. Still, they managed to build a 25 hours test set of 20k utterances called YTDEV18 but it is worth highlighting that this test set is not publicly available.

A second approach [1], employed to create the publicly available and thus widely used LRS3 dataset, involves restricting the content sources to high-quality videos, focusing on TED and TEDx talks. These talks have reliable transcripts and visuals very likely correspond to what is targeted, maximizing the input-output rate of a processing pipeline intended to build a VSR dataset. However, this approach drastically reduces the available content from

YouTube and may limit the dataset’s ability to represent “in the wild” performances, as TED talks are often filmed in similar visual context, with underrepresented individuals and a formal language register.

Our approach tends to blend both strategies by profiting from the diversity of available data on YouTube as in [25] while targeting quality content likely to match our requirements as in [1]. It also has the advantage of being scalable and language-oriented, meaning that the aim is to be computationally efficient and usable to target languages other than English. Moreover, as done in [1], it frees from legal issues encountered with former public datasets LRW [7] and LRS2 [37], by targeting only free-to-use content.

3.1. Data Gathering

We leverage the power of YouTube search API that offers multiple filters solutions: (i) Search using keywords, (ii) target most relevant videos in regard of a specified language (here English) and (iii) target only videos published under creative commons license.

We start from a list of 21 keywords, including interview or discussion for example, in order to get the most relevant content, expanding the content-oriented filtering explored in [1]. Furthermore, we ensure that these videos do not overlap with LRS3 by cross-verifying the YouTube ids. The resulting videos are then processed as described next.

3.2. Data Processing

Each video is divided into multiple shots using a scene change detection [12] method based on changes in three-dimensional histograms and faces in each frame of the video are detected using the YOLOv5n0.5-Face [28] for its unmatched precision-computation cost ratio. The detected faces are then matched and tracked across the frames to obtain multiple face tracks. Afterwards, we utilize SyncNet [8] for filtering these tracks through active speaker detection and obtain face track segments with active speakers corresponding to the audio of the track. We employ the ASR model Whisper [29] to detect language for discarding non-English tracks and to obtain pseudo-transcripts. Using the word timestamps provided in Whisper results, tracks are then segmented into clips with durations ranging from 0.5 to 16 seconds. These clips, along with their transcripts, are further verified manually to ensure high-quality clip-text pairs spanning a total of 4.8 hours in duration, and a total of 2854 utterances coming from 478 YouTube videos.

3.3. Test Set Quality Evaluation

Vocabulary size and visual content variability are two aspects to consider when evaluating a VSR test set quality. As shown in Table 1, our WildVSR test set achieves a clear increase in vocabulary size, encompassing 71% of the vocabulary from LRS3, along with an increase in number

Table 1. **Comparison of statistics between LRS3 and our proposed test sets.** Our proposed test set has higher number of utterances along with $1.5\times$ unique speakers, $4.6\times$ word instances, $3\times$ vocabulary coverage and $5.3\times$ the duration.

Test set	# Spk.	# Utt.	Word inst.	Vocab	Hours
LRS3	412	1,321	9,890	1,997	0.9
WildVSR	618	2,854	45,182	6,040	4.8




Figure 1. **Comparison between LRS3 and proposed test sets in terms of gender and race attributes.** Compared to the LRS3 test set, we observe a marginal improvement for race attribute (in the center) and relatively better diversity in terms of gender (on the left) for our test set.

of unique speakers, utterances, word instances, and total duration. We use VGG-Face [32] to identify the different speakers present across the test set. Afterward, we manually review and verify all speakers, resulting in 618 distinct identities distributed across 478 YouTube videos. Furthermore, we employ Deepface framework [33] to obtain coarse-level demographic attribute metrics, namely gender and race. These attributes are then verified manually. Additionally, we rely on subjective analysis to determine whether a speaker’s accent is native or not based solely on fluency. Compared to LRS3 test set Deepface predictions (not verified manually) as shown in Figure 1, our proposed test set demonstrates a slightly more balanced distribution across all attributes, indicating improved diversity. However, it is important to note that biases inherent to the online platform may be present in the dataset.

4. Experiments

Here, we evaluate a broad range of VSR models spanning five years of progress in a highly active area of research. The models include the fully-supervised models: Ma *et al.* [23], VTP [27], Auto-AVSR [20] and a set of self-supervised models fine-tuned on LRS3 with the different pretraining regimes such as RAVen [13] and AV-HuBERT [35]. The VSR models generally comprise a ResNet-3D frontend followed by a transformer encoder-decoder. The frontend encodes the video lip sequence into a temporal sequence of features, which is fed into the encoder. The decoder then autoregressively decodes the text

Table 2. **Performance comparison on our proposed test set in low-resource and high-resource settings.** The performance is reported in terms of WER and proposed $Rank_{wer}$ metrics. ‘Base’ and ‘Large’ denote the size of the self-supervised video encoder employed. Performance of supervised approaches are also reported. The LRS3 test set performance is also shown for reference. We also report the performance of ASR (audio only) and AVSR (audio-visual) models on both LRS3 and our WildVSR test sets, in addition to the compute budget required (in ExaFLOPs) for training the respective AVSR and VSR models.

	Method	Unlabeled AV data	Labeled data	Decoding	Compute (ExaFLOPs)	WER		$Rank_{wer}$	
						LRS3	WildVSR	LRS3	WildVSR
ASR Models	Wav2vec2.0 [3]	–	–	CTC	–	6.1	17.7	6.9	18.0
	Whisper [29]	–	–	CE	–	1.1	4.2	1.3	4.8
AVSR Models	Base AV-HuBERT [35]	1759h	30h	CE	39.7	4.1	24.4	4.3	26.1
	Base AV-HuBERT [35]	1759h	433h	CE	39.9	1.8	13.7	2.1	14.2
	Large AV-HuBERT [35]	1759h	30h	CE	106.7	3.4	22.8	3.6	24.3
	Large AV-HuBERT [35]	1759h	433h	CE	107.3	1.5	12.9	1.7	13.2
<i>Low-resource setting</i>									
Self-Supervised Base	AV-HuBERT [35]	433h	30h	CE	39.7	51.8	79.4	67.0	83.3
	RAVen [13]	433h	30h	CTC+CE	3.7	48.1	75.4	64.2	80.4
	AV-HuBERT [35]	1759h	30h	CE	39.7	46.1	73.2	62.2	77.5
	RAVen [13]	1759h	30h	CTC+CE	14.9	40.2	66.9	54.1	72.1
Self-Supervised Large	AV-HuBERT [35]	433h	30h	CE	106.7	44.8	75.7	59.5	81.8
	AV-HuBERT [35]	1759h	30h	CE	106.7	32.5	61.9	41.1	68.0
	AV-HuBERT [35] w/ self-training	1759h	30h	CE	106.7	28.6	52.4	37.7	58.6
	RAVen [13]	1759h	30h	CTC+CE	96.8	32.5	57.7	41.6	63.4
	RAVen [13] w/ self-training	1759h	30h	CTC+CE	131.7	23.8	48.4	31.5	52.0
<i>High-resource setting</i>									
Supervised	Ma <i>et al.</i> [23]	–	1459h	CTC+CE	2.1	32.3	58.4	42.4	63.6
	Prajwal <i>et al.</i> [27]	–	698h	CE	†	40.6	75.6	57.3	86.9
	Prajwal <i>et al.</i> [27]	–	2676h	CE	†	30.7	68.7	43.7	82.1
	Auto-AVSR [20]	–	661h	CTC+CE	6.7	32.7	62.3	42.5	67.1
	Auto-AVSR [20]	–	1759h	CTC+CE	17.8	25.1	49.3	31.7	53.2
	Auto-AVSR [20]	–	3448h	CTC+CE	34.9	19.1	38.6	24.8	41.8
Self-Supervised Base	AV-HuBERT [35]	433h	433h	CE	39.9	44.0	71.6	58.1	76.2
	RAVen [13]	433h	433h	CTC+CE	5.0	39.1	69.9	52.8	77.6
	AV-HuBERT [35]	1759h	433h	CE	39.9	34.8	58.1	50.5	63.9
	RAVen [13]	1759h	433h	CTC+CE	16.3	33.1	60.0	42.6	65.7
Self-Supervised Large	AV-HuBERT [35]	433h	433h	CE	107.3	41.6	69.4	56.1	73.5
	AV-HuBERT [35]	1759h	433h	CE	107.3	28.6	51.7	37.4	55.9
	AV-HuBERT [35] w/ self-training	1759h	433h	CE	107.3	26.9	48.7	34.9	53.0
	RAVen [13]	1759h	433h	CTC+CE	105.3	27.8	52.2	36.6	55.3
	RAVen [13] w/ self-training	1759h	433h	CTC+CE	131.7	23.1	46.7	30.8	49.8

tokens by cross-attending to the encoder output features. Among the above models, Ma *et al.* and Auto-AVSR replace the standard transformer layers with conformer layers, where convolutions are additionally interleaved with self-attention layers. We rely on the publicly available source code and their pre-trained weights for evaluation. Overall, the average WER ranges between 19 and 40 on the LRS3 test set. Here, Table 2 shows the main results on both LRS3 and our proposed test sets. Next, we describe the main trends that arise from the experiments.

Significant drops in WER: All models see a clear drop in WER from the LRS3 to the WildVSR test sets. For the best model on LRS3 (*i.e.*, Auto-AVSR: 19.1 WER), it loses nearly 20 points to achieve 38.6 WER on WildVSR. Also, the self-supervised models pre-trained on 1759h of VoxCeleb2-en + LRS3 are slightly more robust to pre-training on the 433h of LRS3 only (average 21 vs. 32 WER drops). In fact, the WER on the WildVSR test set can be linearly approximated given the WER on the

LRS3 test set. Let’s denote a pair of scores for all models $(wer^{LRS3}, wer^{WildVSR})$, then calculate the respective means as: $\mu_{LRS3} = \frac{1}{M} \sum wer_i^{LRS3}$ and $\mu_{WildVSR} = \frac{1}{M} \sum wer_i^{WildVSR}$, where M denotes the number of all the models considered. Each score is normalized as:

$$b = \mu_{WildVSR} - (m \cdot \mu_{LRS3}), \quad (1)$$

where $m = \sigma/\gamma$. Here, σ and γ are computed as

$$\sigma = \sum_{i=1}^M (\Delta wer_i^{LRS3} \cdot \Delta wer_i^{WildVSR}), \quad (2)$$

$$\gamma = \sum_{i=1}^M (\Delta wer_i^{LRS3})^2, \quad (3)$$

where $\Delta wer_i^k = |wer_i^k - \mu^k|$ and $k \in \{LRS3, WildVSR\}$. Using the WER pairs from Table 2, the new WER of a model is approximately given by the following formula:

$$wer^{WildVSR} \approx 1.31 \cdot wer^{LRS3} + 14.05. \quad (4)$$

Since $m > 1$, it indicates that models with lower WER on LRS3 experience a comparatively smaller decline when tested on WildVSR. This suggests that better performance on LRS3 yields improved models robustness. We hypothesize that not much overfitting is happening in LRS3, as we do not observe any WER diminishing returns. Hence, there is minimal change in models' ordering across both test sets.

Training compute vs. performance: To estimate the training compute in FLOPs for each model, we utilize the methodologies proposed by [15]. This estimation considers various factors including: model size, batch size, training data, and the specifics of the training procedure. As illustrated in Figure 2, supervised models, notably Auto-AVSR, demonstrate an impressive balance between performance and computational efficiency. On the other hand, self-supervised methods (AV-HuBERT and RAVen) demand a significantly higher compute budget ($\approx 3.6\times$ the compute of Auto-AVSR), while achieving only a moderate performance in terms of WER. Further details on the compute calculations are provided in Section C of supplementary.

Mode collapse: Whilst Wav2vec2.0 fails in matching the exact target speech, the predictions are closer and represent reasonable errors, like homophones (*e.g.*, PARATON *vs.* PERITON, LIKE POOR *vs.* LUDPORE). This suggests that Wav2vec2.0 robustly detect the input patterns, but might miss on prediction due to entangled representations. Clearly however, the failure of the state-of-the-art VSR model, Auto-AVSR, reveals a different nature of error compared to Wav2vec2.0. In the case of Auto-AVSR failure, the predictions deviate significantly from the target speech and often appear to be random. For instance, it produces degenerated predictions, *e.g.*, SPORTING BUSINESS *vs.* THERE IS BOARDING BUSES.

4.1. Model Consistency

The word error rate (WER) [16] is a standard metric used for comparing different VSR models. However, given that the test set videos have varying target lengths, weighted average WER (μ_{wer}) across the test set might not be sufficient for comparing different approaches, *e.g.*, a model might fit precisely to some samples while having poor predictions for others. Furthermore, we observe that the WER distribution on the LRS3 test set is non-symmetric with more mass around 0-20, while the weighted standard deviation (σ_{wer}) is in the order of the mean. Thus, we combine both mean and standard deviation in a unified rank metric, as $\mu_{wer}(1 + \sigma_{wer})$, to compare the models. Such a metric correctly penalizes models that achieve lower μ_{wer} at the cost of higher σ_{wer} . Let us denote y_N and \bar{y}_N as the set of ground-truth labels and predictions respectively, where N is the number of samples in the test set. The WER is calculated between each pair $wer_i = (y_i, \bar{y}_i)$, then the set of all scores is denoted as wer_N . To account for the variable

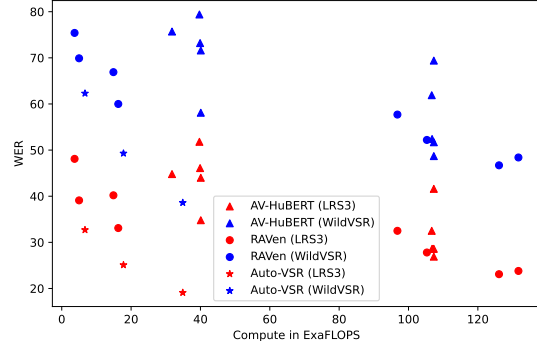


Figure 2. **Training compute (in exaFLOPs) vs. performance (in WER).** The best performing models of AV-HuBERT and RAVen that employ pretraining + finetuning achieve only a moderate performance while requiring $\approx 3.6\times$ training compute, compared to Auto-AVSR that is trained in a fully-supervised manner.

length targets, the average WER is given by:

$$\mu_{wer} = \frac{\sum_{i=1}^N w_i \alpha_i}{\sum_{i=1}^N \alpha_i} = \sum_{i=1}^N p_i w_i, \quad (5)$$

where $p_i = \frac{\alpha_i}{\sum_{i=1}^N \alpha_i}$ with α_i denoting the number of words in y_i . We define the variance as follows:

$$\sigma_{wer} = \sum_{i=1}^N (w_i - \mu_{wer})^2 p_i. \quad (6)$$

We observe that the WER distribution on the LRS3 test set is non-symmetric, with more mass around 0, and the standard deviation is in the order of the mean, thus, we propose the following metric to rank the models, based on their standard deviation and mean WER:

$$Rank_{wer} = \mu_{wer}(1 + \sigma_{wer}). \quad (7)$$

The goal is to have an increasing function in both μ_{wer} and σ_{wer} , to penalize models that might have lower mean WER but higher standard deviation.

Ranking models: As shown in Table 2, models have lower weighted standard deviation on our WildVSR test set ($\sigma_{wer} \approx 0.10$), hence, closer $Rank_{wer}$ compared to wer . Surprisingly however, on LRS3, the σ_{wer} is in the order of the mean WER for the models ($\sigma_{wer} \approx 0.30$). This suggests that there is a significant amount of variations and inconsistencies in the performance of the models, where the predictions match exactly certain samples from the LRS3 test set, which is not the case for the WildVSR test set. This highlights that VSR models can be sensitive to various factors, such as diverse lip sequences with varying acoustic conditions, accents, vocabulary, or speaking styles. The high σ_{wer} suggests that approaches in Table 2 are not robust enough to handle such variability, and their performance fluctuate significantly across samples.




Figure 3. **Visualization of the dominant spatial modes of the Tucker decomposition over time** of the encoder representations on four sample videos each from LRS3 (in blue) and WildVSR (in red) test sets. The representations are obtained from the best variants of RAVen (on the left) and Auto-AVSR (on the right) models. The detected patterns on LRS3 are salient, whereas fewer modes are detected on WildVSR.

4.2. Analysis of Representation Variability

With a VSR model denoted as f_θ , we sample a batch of video sequences X_{LRS3} and X_{WildVSR} from LRS3 and our test sets, respectively. We employ f_θ to map each batch to the representations space z_{LRS3} and z_{WildVSR} both $\in \mathbb{R}^{B \times T \times D}$, where B is batch size, T is temporal dimension and D denotes spatial features dimension. To investigate the variability of representations across the two test sets, we employ a combination of power iteration and Tucker decomposition. First, we calculate the power iteration for each representation tensor. The power iteration [26] process allows us to identify the dominant eigenvector in each tensor, which represents the principal direction of variability. The power iteration is performed iteratively until convergence, and then we normalize each representation tensor to be in the same scale. Next, we employ the Tucker decomposition [40] to extract the underlying factors of each normalized representation tensor. The Tucker decomposition is a tensor factorization technique that decomposes a tensor into a set of core tensors and factor matrices. For z_{LRS3} and z_{WildVSR} , the Tucker decomposition is performed with a rank of (r, r, r) , resulting in

$$z_k = C_k \times_1 F_k^{(1)} \times_2 F_k^{(2)} \times_3 F_k^{(3)}, \quad (8)$$

where $k \in \{\text{LRS3}, \text{WildVSR}\}$. The factor matrices obtained from Tucker decomposition represent the latent factors of variability in the representation tensors. Let $F_{\text{LRS3}}^{(i)}$ and $F_{\text{WildVSR}}^{(i)}$ denote the factor matrices corresponding to z_{LRS3} and z_{WildVSR} , respectively. To project each representation tensor onto the factor matrices, we perform a multi-mode dot product. This operation aligns the representation tensor with the corresponding factor matrices, capturing the influence of each factor on the tensor. With $k \in \{\text{LRS3}, \text{WildVSR}\}$, the projection is done as follows:

$$\text{proj}_k = z_k \times_1 F_k^{(1)T} \times_2 F_k^{(2)T} \times_3 F_k^{(3)T}. \quad (9)$$




Figure 4. **Visualization of Tucker decomposition's spatial mode over time** for RAVen (on the left) and AV-HuBERT (on the right) encoder representations on sample videos from LRS3, before (in red) and after (in blue) finetuning. We see that RAVen encoder before finetuning fails to capture salient modes indicating that VSR representations are better learned only during its finetuning stage.

Finally, we visualize the projected tensors $\text{proj}_{\text{WildVSR}}$ and $\text{proj}_{\text{LRS3}}$ to examine the variability of representations across the two test sets. This visualization provides insights into the differences in the learned representations and helps assess the effectiveness of our model. The entire procedure enables a comprehensive analysis of the variability in representations across different test sets, facilitating the evaluation and interpretation of model performance.

Evaluation: We select a batch of 128 examples from both LRS3 and our WildVSR test sets, and use the best variants of RAVen and Auto-AVSR for the analysis. Figure 3 shows the projected tensors over time on both test sets. Observing the visualizations, it becomes evident that both models exhibit a considerably stronger response to the sequences presented in the LRS3 test set. These models successfully capture the underlying modes and patterns inherent in the input signals, enabling accurate detection and subsequent generation of precise transcriptions. The rich and diverse set of features recognized by the models within the LRS3 test set contributes to their superior predictive capabilities. Conversely, when examining the representations derived from our own test set, a notable difference emerges. The features extracted from our test set display a distinct lack of variability, indicating a reduced number of patterns being recognized by the models. Consequently, the predictive performance of the models on our test set is compromised, leading to sub-optimal transcription outcomes.

Which SSL is better? We notice fundamental differences between RAVen and AV-HuBERT self-supervised procedures, where the former learns two separate encoders for video and audio modalities and optimizes to match the latents. Differently, the latter fuses them in a single encoder and learns to predict pre-assigned cluster memberships. Hence, we use the same Tucker decomposition to compare their representations before and after finetuning. As shown in Figure 4, SSL RAVen encodes much less information explaining the need for requiring 75 epochs at for finetuning stage. In contrast, AV-HuBERT detects modes

Table 3. **Performance comparison on multiple folds of our test set.** The folds are the intersection of where all models are below 30 WER for *top-k*, and more than 50 WER for *bottom-k*. We observe that the models match their LRS3 performance on *top-k* examples.

Model	Video folds		
	Top-k	Bottom-k	All
Auto-AVSR [20]	18.6	71.4	38.6
AV-HuBERT [35] w/ self-training	24.8	77.3	48.7
RAVEN [13] w/ self-training	23.5	75.0	46.7




Figure 5. **Performance (WER) comparison across five different data folds.** Fold 1 is an overlap subset where all models obtain a WER score higher than 50. We keep progressively adding data from the remaining test set to create Folds 2 to 5, where Fold 5 contains all the test set videos.

even without seeing any labelled data, suggesting the reason behind the small finetuning budget. We hypothesize that coarse-grained clustering as done in AV-HuBERT better suits the VSR task, as it pushes the encoder to learn phonemes that can be lightly mapped to word labels.

5. Potential Causes for WER drop

Following [30], the error difference between the respective test sets can be decomposed into three parts:

$$\mathcal{L}_{WildVSR} - \mathcal{L}_{LRS3} = \underbrace{(\mathcal{L}_{WildVSR} - \mathcal{L}_{LRS3})}_{\text{Adaptivity gap}} + \underbrace{(\mathcal{L}_{WildVSR} - \mathcal{L}_{LRS3})}_{\text{Distribution Gap}} + \underbrace{(\mathcal{L}_{WildVSR} - \mathcal{L}_{LRS3})}_{\text{Generalization gap}} \quad (10)$$

The *Adaptivity gap* quantifies the extent to which adjusting (adapting) a model to fit the specific test set \mathcal{L}_{LRS3} leads to an underestimation of the test error. The *Distribution gap* measures how our new data distribution and generation process affects the model performance. The *Generalization gap* is influenced by the inherent random sampling error. These components are hard to track, and distinguish in practice.

Linear proportional gains: Eq. 4 suggests that no diminishing returns is apparent. This strongly indicates that the reduction in Word Error Rate (WER) could primarily be attributed to the *Distribution Gap*. Despite our best efforts to emulate the original LRS3 dataset creation procedure, the distribution gap remains the main explanation for the observed decreases in WER. Upon examination of our data

Table 4. **Performance comparison on our proposed test set with varying the attributes.** We report the best variant per model.

Method	Accent		Gender	
	Native	Non-Native	Male	Female
Auto-AVSR [20]	35.1	46.9	38.3	37.4
AV-HuBERT [35]	47.5	55.3	49.2	47.7
RAVEN [13]	45.2	55.0	47.5	45.3




Figure 6. **Variation of WER as duration of video clips is varied.** VSR models decode better with longer context (video duration).

Table 3, we observe that a subset of our test set exhibits similar complexity levels to those found in LRS3, thus models achieve their reported LRS3 WER scores. We postulate that our test set includes a larger number of examples from difficult LRS3 modes, in addition to new modes not found in LRS3. While current VSR models showcase impressive WER scores on the original LRS3 test set, they still encounter difficulties in effectively generalizing from ‘easy’ lip sequences to more challenging ones.

Hard samples: As shown in Figure 5, we selected a data subset where all models obtain a WER score higher than 50. The subsets contain 110/639 samples for LRS3/WildVSR respectively. On both LRS3 and WildVSR, models obtain a $WER \geq 75$ for these challenging examples. We checked these samples manually, and noticed a high variability in head poses along with shorter videos. There might be other confounding variables from the accents/vocabulary that are harder to assess. Figure 6 shows the performance comparison on different folds obtained by partitioning the test sets based on the lip sequence length. We observe that shorter videos (less than 2 seconds, *i.e.*, 50 frames) present a bottleneck, which results in performance degradation of the approaches from their corresponding average WER on the whole test sets. This is likely due to the lack of rich contextual features in shorter video sequences, which leads to sub-optimal temporal modeling in the video encoder. Furthermore, we kept progressively adding samples from the remaining test set, and observed a faster decay on LRS3. We hypothesize that the higher amount of ‘hard samples’ (23%) in WildVSR (*vs.* 08% for LRS3) is responsible for this behavior as the final WER is averaged across samples.

Speaker characteristics: As demonstrated in Table 4, the models show an average decrease of 8 points when evaluated on non-native speakers compared to native speakers. Additionally, there is a relatively similar performance between male and female speakers across the models. Also see Table. A.3 for the scores on more attributes.

Head poses: Here, we create two subsets \mathcal{S}_1 and \mathcal{S}_2 with 1010 and 639 videos, respectively. \mathcal{S}_1 contains the best-performing videos across different models, while \mathcal{S}_2 contains videos that are hard to decode for all the models. Next, we detect sequences with frontal and extreme poses by first recovering the 3D head pose using [11] and then regressing the 3D-model parameters that best fit to each image frame using a parametric 3D model [18] learned from 3D scans of human faces. Frontal and extreme poses are considered based on predefined face angles. We perform this analysis and find that extreme poses represent 31% and 52% on \mathcal{S}_1 and \mathcal{S}_2 , respectively. This suggests that extreme head poses adversely affect the model performance for VSR task.

Vocabulary: As demonstrated in Table 3, state-of-the-art (SoTA) models reproduce their reported scores on a subset of our test set (specifically, *top-k*, comprising 1010 videos). We conducted an in-depth analysis of the vocabulary utilized in this fold and found that it comprises a total of 2315 words. Interestingly, this word set exhibits a 75% similarity to the LRS3 vocabulary, which consists of 1969 words. This observation may suggest that VSR models tend to perform optimally when operating within a constrained vocabulary set, analogous to the one present in the LRS3 dataset. This underlines the significant influence that vocabulary consistency and word choice have on the performance of VSR models. The restricted vocabulary environment may simplify the model’s task, leading to a higher performance. Understanding these limitations can provide crucial insights for future development of VSR models.

Unified model architectures: Despite the many approaches in VSR literature, they mostly share the same architecture while differing in the training procedures and objectives. The inductive bias-free nature of transformers helps preventing the explicit definition of *overfitting*. It should be noted that while these methods do eliminate the most salient aspects of overfitting, they do not eliminate all possibilities. There exists a scenario where any degree of test set adaptivity results in a consistent decrease in accuracy across all models. In such a case, the diminishing returns phenomenon would not be observed since subsequent models could still exhibit improved performance. Detecting and studying this specific manifestation of adaptive overfitting would likely necessitate the use of a novel test set that truly adheres to independent and identically distributed properties, rather than being derived from a distinct data collection effort. Identifying an appropriate dataset for conducting such an experiment remains extremely challenging.

6. Discussion

Recommendations: As observed earlier in compute-performance trade-off section, although self-supervised learning followed by fine-tuning paradigm might seem exciting, it emerges as a less optimal approach for VSR. The inherent complexity of video data intensifies the computational demands, making it an expensive avenue to explore. Moreover, as these learned encoders are utilized for a singular downstream task of VSR, the potential benefits of SSL are somewhat neutralized. Additionally, the finetuning phase can consume a budget comparable to that of direct supervised training, as in RAVen [13]. Figure 2 shows the effectiveness of fully-supervised learning when supplemented with state-of-the-art ASR models, such as Whisper [29], serving as automatic labelers. This methodology, as proved by Auto-AVSR [20], provides a good balance between performance and computational efficiency, while still achieving superior performance.

Ethical considerations: Due to the data collection process focusing on YouTube, biases inherent to the platform may be present in the test set. Also, while measures are taken to ensure diversity in content, the dataset might still be skewed towards certain types of content due to the filtering process. Furthermore, we have taken specific steps to ensure that WildVSR respects individual privacy. Firstly, since most VSR approaches utilize only the 96×96 cropped region around the mouth as input, we make available the cropped sequence to reduce the potential for individual identification, emphasizing only the pertinent region for VSR model’s input. Also, we provide a mechanism for individuals who recognize themselves in the test set to opt-out. Should someone wish to have their data removed, they can contact us and we will promptly exclude their content.

Conclusion: In this work, we have highlighted the lack of generalization within the field of Visual Speech Recognition (VSR) due to an excessive focus on the Lip Reading Sentences-3 (LRS3) test set. To mitigate this, we have proposed a new VSR test set, named WildVSR, incorporating a higher visual diversity and spoken vocabulary. Indeed, the benchmarking of a wide range of publicly available VSR models on this new test set revealed significant drops in performance compared to the LRS3 test set. This outcome underlines the models’ difficulties in generalizing to “harder” lip sequences present in our test set. Interestingly, the comparative ranking of models remained consistent across the original and new test sets, indicating that these performance drops are not merely a result of over-tuning for specific lip sequences on LRS3 test set. We also introduced a novel metric that combines the mean and standard deviation of Word Error Rates (WER), better capturing a model’s consistency across various test samples. It is our hope that this will stimulate the development of more robust VSR models, furthering advancements in this challenging field.

References

- [1] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. Lrs3-ted: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*, 2018. 1, 2, 3
- [2] Triantafyllos Afouras, Joon Son Chung, and Andrew Zisserman. ASR is all you need: cross-modal distillation for lip reading. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2020-May:2143–2147, 11 2019. 2
- [3] 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*, 33:12449–12460, 2020. 4
- [4] Oscar Chang, Hank Liao, Dmitriy Serdyuk, Ankit Shah, and Olivier Siohan. Conformers are all you need for visual speech recognition. *arXiv preprint arXiv:2302.10915*, 2023. 2
- [5] François Chollet. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019. 1
- [6] Joon Son Chung, Arsha Nagrani, and Andrew Zisserman. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*, 2018. 2
- [7] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *Computer Vision—ACCV 2016: 13th Asian Conference on Computer Vision, Taipei, Taiwan, November 20–24, 2016, Revised Selected Papers, Part II 13*, pages 87–103. Springer, 2017. 2, 3
- [8] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *Computer Vision—ACCV 2016 Workshops: ACCV 2016 International Workshops, Taipei, Taiwan, November 20–24, 2016, Revised Selected Papers, Part II 13*, pages 251–263. Springer, 2017. 3
- [9] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*, 2018. 2
- [10] Charles Fefferman, Sanjoy Mitter, and Hariharan Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049, 2016. 1
- [11] Panagiotis P Filntisis, George Retsinas, Foivos Paraperas-Papantoniou, Athanasios Katsamanis, Anastasios Roussos, and Petros Maragos. Visual speech-aware perceptual 3d facial expression reconstruction from videos. *arXiv preprint arXiv:2207.11094*, 2022. 8
- [12] Igor S. Gruzman and Anna S. Kostenkova. Algorithm of scene change detection in a video sequence based on the threedimensional histogram of color images. In *2014 12th International Conference on Actual Problems of Electronics Instrument Engineering (APEIE)*, pages 1–1, 2014. 3
- [13] Alexandros Haliassos, Pingchuan Ma, Rodrigo Mira, Stavros Petridis, and Maja Pantic. Jointly learning visual and auditory speech representations from raw data. *arXiv preprint arXiv:2212.06246*, 2022. 2, 3, 4, 7, 8
- [14] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020. 2
- [15] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 5
- [16] Dietrich Klakow and Jochen Peters. Testing the correlation of word error rate and perplexity. *Speech Communication*, 38(1-2):19–28, 2002. 5
- [17] Hirotaka Kosaka, Masao Omori, Tetsuya Idaka, Tetsuhito Murata, T Shimoyama, Tomohisa Okada, Norihiro Sadato, Yoshiharu Yonekura, and Yuji Wada. Neural substrates participating in acquisition of facial familiarity: an fmri study. *Neuroimage*, 20(3):1734–1742, 2003. 1
- [18] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics (TOG)*, 36(6), 11 2017. 8
- [19] Xubo Liu, Egor Lakomkin, Konstantinos Vougioukas, Pingchuan Ma, Honglie Chen, Ruiming Xie, Morrie Doulaty, Niko Moritz, Jachym Kolar, Stavros Petridis, et al. Synthvsr: Scaling up visual speech recognition with synthetic supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18806–18815, 2023. 2
- [20] Pingchuan Ma, Alexandros Haliassos, Adriana Fernandez-Lopez, Honglie Chen, Stavros Petridis, and Maja Pantic. Auto-avs: Audio-visual speech recognition with automatic labels. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 2, 3, 4, 7, 8
- [21] Pingchuan Ma, Rodrigo Mira, Stavros Petridis, Björn W Schuller, and Maja Pantic. Lira: Learning visual speech representations from audio through self-supervision. *arXiv preprint arXiv:2106.09171*, 2021. 2
- [22] Pingchuan Ma, Stavros Petridis, and Maja Pantic. End-to-end audio-visual speech recognition with conformers. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7613–7617. IEEE, 2021. 1
- [23] Pingchuan Ma, Stavros Petridis, and Maja Pantic. Visual speech recognition for multiple languages in the wild. *Nature Machine Intelligence*, pages 1–10, 2022. 2, 3, 4
- [24] Shuang Ma, Zhaoyang Zeng, Daniel McDuff, and Yale Song. Contrastive Learning of Global-Local Video Representations. *Advances in Neural Information Processing Systems*, 9:7025–7040, 4 2021. 2
- [25] Takaki Makino, Hank Liao, Yannis Assael, Brendan Shillingford, Basilio Garcia, Otavio Braga, and Olivier Siohan. Recurrent neural network transducer for audio-visual speech recognition. In *2019 IEEE automatic speech recognition and understanding workshop (ASRU)*, pages 905–912. IEEE, 2019. 2, 3
- [26] RV Mises and Hilda Pollaczek-Geiringer. Praktische verfahren der gleichungsauflösung. *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, 9(1):58–77, 1929. 6

- [27] KR Prajwal, Triantafyllos Afouras, and Andrew Zisserman. Sub-word level lip reading with visual attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5162–5172, 2022. [2](#), [3](#), [4](#)
- [28] Delong Qi, Weijun Tan, Qi Yao, and Jingfeng Liu. Yolo5face: Why reinventing a face detector. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V*, pages 228–244. Springer, 2023. [3](#)
- [29] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. *arXiv preprint arXiv:2212.04356*, 2022. [3](#), [4](#), [8](#)
- [30] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, pages 5389–5400. PMLR, 2019. [7](#)
- [31] Dmitriy Serdyuk, Otavio Braga, and Olivier Siohan. Audio-visual speech recognition is worth 32x32x8 voxels. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 796–802. IEEE, 2021. [2](#)
- [32] Sefik Ilkin Serengil and Alper Ozpinar. Lightface: A hybrid deep face recognition framework. In *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pages 23–27. IEEE, 2020. [3](#)
- [33] Sefik Ilkin Serengil and Alper Ozpinar. Hyperextended lightface: A facial attribute analysis framework. In *2021 International Conference on Engineering and Emerging Technologies (ICEET)*, pages 1–4. IEEE, 2021. [3](#)
- [34] Changchong Sheng, Matti Pietikäinen, Qi Tian, and Li Liu. Cross-modal self-supervised learning for lip reading: When contrastive learning meets adversarial training. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 2456–2464, 2021. [2](#)
- [35] Bowen Shi, Wei-Ning Hsu, Kushal Lakhotia, and Abdelrahman Mohamed. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*, 2022. [2](#), [3](#), [4](#), [7](#)
- [36] Bowen Shi, Wei-Ning Hsu, and Abdelrahman Mohamed. Robust self-supervised audio-visual speech recognition. *arXiv preprint arXiv:2201.01763*, 2022. [2](#)
- [37] Joon Son Chung, Andrew Senior, Oriol Vinyals, and Andrew Zisserman. Lip reading sentences in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6447–6456, 2017. [2](#), [3](#)
- [38] Christopher Summerfield. *Natural General Intelligence: How understanding the brain can help us build AI*. Oxford University Press, 2022. [1](#)
- [39] Alexander Todorov. The role of the amygdala in face perception and evaluation. *Motivation and Emotion*, 36:16–26, 2012. [1](#)
- [40] Ledyard R Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966. [6](#)
- [41] Qiushi Zhu, Long Zhou, Ziqiang Zhang, Shujie Liu, Binxiang Jiao, Jie Zhang, Lirong Dai, Daxin Jiang, Jinyu Li, and Furu Wei. Vatlm: Visual-audio-text pre-training with unified masked prediction for speech representation learning. *arXiv preprint arXiv:2211.11275*, 2022. [2](#)