



# Federated Self-supervised Speech Representations: Are We There Yet?

Yan Gao<sup>1</sup>, Javier Fernandez-Marques<sup>2</sup>, Titouan Parcollet<sup>3</sup>, Abhinav Mehrotra<sup>2</sup>, Nicholas D. Lane<sup>1,2</sup>

<sup>1</sup>University of Cambridge, <sup>2</sup>Samsung AI, <sup>3</sup>Avignon University

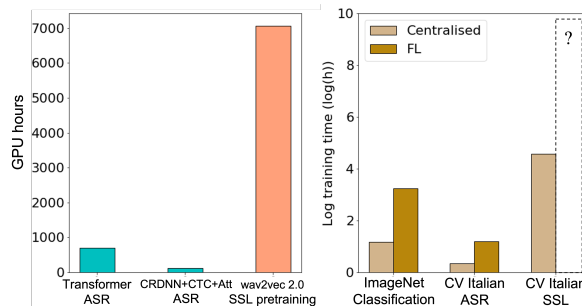
## Abstract

The ubiquity of microphone-enabled devices has led to large amounts of unlabelled audio data being produced at the edge. The integration of self-supervised learning (SSL) and federated learning (FL) into one coherent system can potentially offer data privacy guarantees while also advancing the quality and robustness of speech representations. In this paper, we provide a *first-of-its-kind* systematic study of the feasibility and complexities for training speech SSL models under FL scenarios from the perspective of algorithms, hardware, and systems limits. Despite the high potential of their combination, we find existing system constraints and algorithmic behaviour make SSL and FL systems nearly impossible to build today. Yet critically, our results indicate specific performance bottlenecks and research opportunities that would allow this situation to be reversed. While our analysis suggests that, given existing trends in hardware, hybrid SSL and FL speech systems will not be viable until 2027, we believe this study can act as a roadmap to accelerate work towards reaching this milestone much earlier.

## 1. Introduction

A large amount of audio data, often unlabelled and of private nature, is generated everyday from cellphones, tablets, personal assistants and other IoT devices [1, 2, 3, 4]. Being able to utilise this data to solve various speech related tasks has been of great interest to researchers for over a decade [5, 6, 7, 8, 9]. Self-supervised learning (SSL) allows the learning of representations from unlabelled data, which can later be used to solve specific downstream tasks, e.g., automatic speech recognition (ASR), speech translation, keyword spotting (KWS), and others. SSL shares limitations with other forms of unsupervised training favouring large-capacity models trained with vast amounts of data [5, 6]. Currently, such workloads can only be realised in data centers, where high-performance hardware is available. Such centralised SSL systems unavoidably raise concerns around privacy specially when dealing with speech data. Additionally, they also cause severe communication overheads due to the transferring of data from devices where it originates.


A natural way to mitigate these issues is to combine SSL with federated learning (FL) [14, 15, 16, 17]. In FL, a distributed population of devices collaboratively trains a single global model while keeping their local data private. A typical FL pipeline is comprised of three stages: first, a global model is initialised in a server to which all FL nodes (or *clients*) connect; then, a fraction of clients are chosen by the server, receive a copy of the global model and each performs on-device training using their own data; upon completion the clients send the updated model to the server, where a new global model is generated by aggregating the client updates. This process is repeated for a number of rounds. Deploying SSL workloads in a federated setting offers another benefit beyond privacy-preserving training: large-scale distributed feature learning from real-world data without costly and error-prone manual annotations. *Despite the combination offering significant benefits, no prior work has studied speech SSL in the context of FL.*



**Figure 1:** (left) The pretraining of a SOTA self-supervised model wav2vec 2.0 is several times more time-consuming than SOTA supervised ASR models (transformer [10, 11], CRDNN+CTC+Att [12, 11]). All of the models are trained with the LibriSpeech dataset using V100 GPUs. (right) FL leads to longer training time compared to centralised settings. The models are trained using V100 GPU in centralised settings and using Jetson Xavier NX in FL [13]. CV is Common Voice dataset.

The SSL models are notoriously costly to train. A SOTA SSL model for speech representation learning, wav2vec 2.0 [5], requires over 7000 GPU hours of pretraining (V100 GPUs). That is an order of magnitude more compute resources than a top-performing model trained with only supervision (Fig. 1 left). However, this leap in resource demands when transitioning from supervised to a SSL solution is not the only significant barrier to performing SSL under FL. It is paired with an empirical phenomena generally seen when migrating from centralised training to a federated version: due to the model update aggregation phase of FL, the overall global convergence is weakened and thus slower (Fig. 1 right). We again observe resource requirements (in this case raw training time) increasingly alarmingly – although this is now due to the act of federating the training process. While we know these two factors will drive the ultimate performance of systems that combine FL and SSL for speech, curiously, there is large gap in existing empirical knowledge as to how they will interact and translate into resource requirements. How long would a SSL speech model take to converge if federated? Specifically, what would the memory and compute requirements be for the participating FL clients? Would WER decrease in this decentralised setting? *All of these questions, ranging from those that are obvious down to questions specific in detail are virtually unknown presently.*

In this work, we conduct a systematic investigation as to the precise bottlenecks that cause SSL models to be so resource intensive to train, and further explore how these resource requirements (e.g., memory, compute) translate into the feasibility of SSL models on edge devices in FL environments. We find that the modest edge devices and existing training frameworks are not yet a match for the requirements of federating SSL speech. But more crucially, our empirical results provide for-the-first-time the much needed identification and quantification of the key limiting factors that underpin this situation. As an example of how our findings can accelerate research in this emerging area, we use them to propose a novel modification to FL aggregation that takes into account the training loss of in-



**Figure 2:** Per-layer breakdown of wav2vec 2.0 base model in terms of its memory consumption (left) and compute footprint (right). The data referred to in the legend field is obtained when training with inputs of length 5.5s (the average in CV Italian), batch size 4 and resampled at 16KHz. (Left) The CNN encoder block contributes to a significant memory utilization compared to the transformer layers that follow. At the end of the forward pass the memory utilization is at its maximum (shown with a red cross) since intermediate layer activations need to be kept in-memory in order to do backpropagation. In yellow we show the memory utilization for the same model but when using sequences of 12s and batch size 8. The solid yellow line shows the per-layer total memory utilization whereas the dashed one shows the cumulative memory needed for training. In this scenario training requires just under 10 GB of memory on top the memory usage of other processes responsible for data loading, checkpointing and the OS itself. (Right) Per-layer breakdown of the amount of compute (FLOPs) needed for training. At the back, yellow bars represent the total FLOPs when using with sequences of 12s and batch size 8. Early layers in the CNN block require over an order of magnitude more compute.

dividual FL clients. Our study estimates given existing FL and SSL technique and practices – along with historically guided expectation of hardware improvements [18, 19] – that hybrid self-supervised FL for speech will not be feasible until roughly half decade from now, 2027; even though the core principles are available today. Not only is this the first systematic estimate of its type for the speech community to reflect upon, we hope that it acts as a call to arms around which researchers will rally in an effort to reach this milestone much more quickly.

The scientific contributions of this study are: (1) We conduct a systematic analysis that discovers the core reasons why SSL models are resource intensive (§3.1 and §3.2); (2) We extend this analysis to explore the feasibility of SSL under FL conditions (§3.3); (3) We perform extrapolative analysis as to future hardware feasibility through faithful GPU simulations (§4).

## 2. Experimental Settings

We conduct all of the measurements using wav2vec 2.0 [5] on the CommonVoice Italian (CV Italian) dataset [20] trained on a variety of hardware platforms. The details are as follows:

**Model.** Our choice of SOTA SSL model for speech representation learning is wav2vec 2.0. This model is comprised of three modules: a convolution layers block that extract latent acoustic features from the raw waveform; a series of transformer layers that generate contextualised embeddings; and, a quantization module that discretises the output of the convolutional module to a finite set of speech representations. The model is jointly trained by a contrastive loss and a diversity loss. We consider both the *base* and *large* versions of wav2vec in our evaluation.

**Dataset.** The Common Voice Italian corpus (version 7.0) contains a total of 211K utterances (317 hours) obtained from volunteers recording sentences on smartphones, computers, tablets, etc. The validated set is used for SSL pretraining and consists of 195K utterances (288 hours) with an average sentence length of 5.5 seconds from more than 6K speakers. The training set is used for downstream ASR fine-tuning which contains 191 hours of speech, while the validation and test sets con-

**Table 1:** High-level breakdown of the number of parameters and FLOPs for inference when considering a sequence of 5.5 seconds for three main modules in wav2vec 2.0.

Modules	Parameters (M)		GFLOPs	
	Base	Large	Base	Large
CNN Encoder	4.60	4.73	27.20	27.28
Transformer	89.78	310.70	49.16	170.16
Quantization	0.41	0.57	0.32	0.94
Total	94.79	316.00	76.68	198.32

tain around 25 hours of speech. Due to their heterogeneity in terms of recordings and speaker diversity, CV datasets are seen as a perfect match to conduct FL-related experiments [15].


**Devices.** Our study considers different classes of devices for training speech SSL models. Server-grade NVIDIA A40 GPUs with 48GB of VRAM are used to represent centralised training, while MacBook Pro 2019, Raspberry Pi 4 and NVIDIA Jetson Xavier NX and AGX represent four tiers of FL devices. We also consider different mobile processors for memory comparison, including Snapdragon 870, A15 Bionic and Google Tensor.

## 3. SSL Speech System Resource Analysis

This section provides detailed analysis on the resource footprint of training wav2vec 2.0 (§3.1), and the sensitivity of resources to input audio length (§3.2). These results then enable a feasibility study of federated SSL under existing hardware (§3.3).

### 3.1. Architecture breakdown

Table 1 shows the number of parameters and floating operation points (FLOPs) for each of the three modules in wav2vec 2.0 as well as the total model size and FLOPs required for inference while considering an input audio of duration 5.5 seconds corresponding to the average length of the CV Italian dataset. An immediate observation is the asymmetry between the amount of parameters and FLOPs between the Encoder and Transformer blocks: while the latter has 19× more parameters, it only results



**Figure 3:** Memory utilization and FLOPs with sequence length mapped to relevant offline speech applications and batch of 4.


in  $1.8\times$  more FLOPs compared to the CNN Encoder. When considering the *large* wav2vec 2.0 model, this difference increases significantly: the Transformer has  $66\times$  more parameters but only accounts for  $6\times$  more FLOPs than the Encoder. This first-order analysis does not consider the training dynamics and therefore a more in-depth look into the impact of each sub-architectural module is needed to identify bottlenecks, both in terms of memory and compute, occurring at training time.

Figure 2 shows the per-layer memory consumption (left) and compute (right) of wav2vec 2.0 *base*. Immediately, the high memory utilization and compute needed for the CNN Encoder stands out compared to the rest of the model. These two dimensions are a function of the size of the input tensor passed through the layer – which is not captured in high-level analysis as in Table 1. When it comes to extending SSL models to constrained devices for FL, a key factor is the total amount of memory needed at training time. Figure 2 (left) shows a memory peak (red cross) of 2.54GB when training with input sequences of length 5.5s (avg. length) and batch size 4. When considering more desirable batch size of 8 and audio files of 12s, the total memory accumulated prior to backpropagation is 9.89GB. Such a delta between these training settings highlights how seemingly small changes can translate into having to consider an entirely new tier of devices to support such workloads.

### 3.2. On sequence lengths and numerical precision

Memory and FLOPs grow linearly with the length of the input sequence (Fig. 3). From a practical perspective, the sequence length can link to different types of offline speech applications that could be regarded as downstream tasks for the pre-trained SSL model. For instance, KWS may only require short utterances (e.g. 1-3s) [21, 22], while tasks such as speech synthesis or recognition may increase this to much longer sentences [23, 24]. However, currently released large-scale SSL models including wav2vec 2.0, HuBERT [25] or WavLM [26] are often trained with sentences longer than 15s, making them clearly intractable for current constrained devices. Hence, one could argue that the SSL stage could be tailored to a specific subset of downstream tasks and limit the length of the input sentences for pretraining, reducing the compute and memory footprints.

Parameter precision is another element that affects the memory and compute footprint of an architecture. Therefore, we compare the memory impact of full-precision FP32 with mixed precision (FP32&FP16) [27] (Fig. 4). From the latter results, it appears that the memory savings observed with mixed-precision is not significant and might not be critical while selecting a pool of devices to train SSL models. The impact on the training time, however, provides a different outcome (§3.3).



**Figure 4:** Memory usage for different sequences length with FP32 and mixed precision (FP32&FP16) and batch size of 4.

### 3.3. Federated SSL feasibility

With this section, we provide an analysis relying on memory consumption and training time answering the feasibility of applying federated learning to large-scale SSL pretraining.

**Memory cost.** Memory is a scarce resource in edge devices and limits the feasibility of on-device training in FL. For instance, some popular mobile processors and edge devices including the Raspberry Pi 4 (RPI), Jetson Xavier NX or the A15 Bionic chip are only accompanied with 8GB of memory while the Jetson Xavier AGX, Google Tensor and Snapdragon 870 chips commonly are given 16GB. According to our findings, the first batch of devices could train *base* wav2vec 2.0 models while the latest set could use *large* architectures with sequences of around 7 to 8 seconds at most (Fig. 3). However, if longer sentences are considered as it is typically the case with current SSL models, all the devices may simply be able to train the *base* model, or even none of them. Nevertheless, we should note that these devices, specially when it comes to smartphones, would likely be running other applications and processes in the background. As a result, the memory ceiling available for FL workloads is lower than the devices’ total memory. Therefore, it appears highly probable that even the *base* wav2vec 2.0 architecture exhibits a training memory cost already higher than what high-end devices have to offer.

**Training time.** Training time is another factor influencing the feasibility of SSL in FL. Typical large-scale SSL pretraining takes weeks or months even with hundreds of GPUs [28]. Training time might simply become intractable with limited compute devices and the slower convergence induced by FL.

Hence, we benchmark the observed training time (i.e. seconds per batch) with different FL devices and a modern GPU. As in prior experiments, sentences are sampled so that the average length is 5.5s. Table 2 reports the obtained measurements. When considering an unlikely yet simple batch size equal to 1, the 8-cores i9 of the MacBook Pro is  $30.3\times$  and  $39.5\times$  slower than a Nvidia A40 GPU when training on wav2vec 2.0 *base* and *large* respectively. In the case of the more modest 4-cores RPi, training  $4.4\times$  slower than the MacBook Pro and  $138\times$  slower than the A40. The Jetson boards on the other hand are equipped with embedded GPUs and therefore decrease this training time factor to  $2.7\times$  and  $5.6\times$  slower compared to the A40 for AGX and NX respectively on the *base* model. Once we increase the batch size to 4, which remains considerably low for SSL training that often rely on large batch size values (above 256 [28]), the *per-sequence* training times get reduced by 15%, 20%, 29% and 33% for the MacBook Pro, RPi, AGX and NX respectively thanks to a better hardware utilization. Finally, and conversely to the memory cost, mixed precision provides a substantial boost in training speed for devices supporting it:  $1.56\times$  and  $1.31\times$  faster than full precision for the Xavier NX and AGX with the *base* model. However, support for mixed precision

**Table 2: Training time (average seconds per batch) of wav2vec 2.0 models using sequences with average length 5.5 seconds. OOM stands for out of memory.**

Device	Batch	Base		Large	
		FP32	Mixed	FP32	Mixed
NVIDIA A40	1	0.12	0.11	0.23	0.21
	4	0.27	0.21	0.43	0.42
MacBook Pro 2019	1	3.76	—	9.05	—
	4	12.83	—	33.66	—
Raspberry Pi 4	1	16.60	—	OOM	—
	4	53.26	—	OOM	—
Jetson Xavier AGX	1	0.38	0.43	0.88	0.87
	4	1.08	0.82	OOM	1.72
Jetson Xavier NX	1	0.67	0.61	OOM	OOM
	4	1.78	1.14	OOM	OOM

training is currently rare in embedded devices.


In the more favorable scenario (viz. a A40 GPU and *base* model), we measure the total training time of one epoch for a FL setting with 10 clients (19.5K samples per client) to be 0.37 hours. This value is multiplied by the necessary 150 epochs (see §4) of the full FL procedure, hence reaching 2.31 days, or 55.5 hours of training. Extending this to our set of FL devices leads us to 110, 456, 9 and 15 days of training with the MacBook Pro, RPi, AGX and NX respectively i.e., considering that each client trains for the same amount of data concurrently. In practice, the total number of training rounds might be much higher, as it tends to increase with the number of available clients for FL (§4). Additionally, CV dataset is rather a small dataset for training SSL systems as it solely contains 290 hours of speech compared to the 10s of thousands of hours of recent pre-trained models. Consequently, and according to our analysis that highlights a training period of already two weeks for a small model on a small dataset, we claim that current experimental protocols leading to large-scale SSL models reaching state-of-the-art downstream performance are simply impractical for the current state of the real-world FL environment.

## 4. Federated SSL Speech Representations

With this last section, we estimate when FL/SSL speech systems will become feasible for edge hardware; before showing results that *to-the-best-of-our-knowledge* are the first attempt at training SOTA SSL speech under a simulated FL setting.

### 4.1. Feasibility under existing hardware trends

Compute capabilities have been doubling every 18 months for decades [18, 19, 29]. This implies a  $8\times$  acceleration of compute in 5 years. Based on our precise measurement of training time on the NX, which represents a typical FL device with an embedded GPU, we predict edge devices to exhibit the same compute performance than a A40 GPU by 2027. Thus, assuming the existing SSL models (i.e., no substantial growth in model complexity) and FL methodologies, we must wait for around 5 years for hybrid speech systems combining both technologies to become feasible. As discussed in §3.3, a large amount of memory is needed when considering longer input sequences. Currently, the AGX can be purchased with 32GB of memory, therefore we believe that by 2027 mid-tier edge systems (e.g. NX) would come with memory sizes similar to today’s A40 48GB. This will support performing SSL suitable to more challenging downstream tasks (e.g. ASR).



**Figure 5: (above) Validation loss for SSL pretraining with loss-based aggregation. (below) ASR downstream results on test set with SSL representations from FL and centralised settings.**

### 4.2. Simulated Federated SSL for Speech

We train federated wav2vec 2.0 *base* using SpeechBrain [11] and Flower [30] on two NVIDIA A40 GPUs.

**FL partitions.** Based on the natural partitioning of CV Italian we have two sets of experiments: (1) *10-client FL* and (2) *100-client FL*. In (1), we simulate a simple FL scenario where clients are generally few but with high availability during all rounds, and similar data distribution [31]. The dataset is randomly split into 10 partitions (i.e. one per client) without overlapping of speakers, each containing roughly the same amount of speech data. Setting (2) rises the number of clients to 100 making it a more challenging scenario as each device participates less frequently and has a smaller portion of data [31]. 20 randomly selected clients participate in each round.

**Federated SSL pretraining.** In each FL round, the SSL model is locally trained on each client with 1 local epoch, batch size of 4 in both FL settings. Then, the updated model weights are sent to the server for aggregation directly with FedAvg [14] or after re-weighting the model updates using the client loss as in [15]. The latter is better suited for heterogeneous data distributions across clients. The SSL model in 10-client setting, trained for 150 rounds, shows a faster converging behaviour than 100-client FL which is trained for 500 rounds (Fig. 5 above). This is expected as a 100-client setting is more challenging.

**ASR downstream evaluation.** The learned speech representations from SSL pretraining are evaluated with an ASR downstream task by adding 3 linear layers on the top. The entire network is fine-tuned for 80 epochs with CTC loss following the official SpeechBrain recipe. Figure 5 (bottom) shows the WER on the test set. Federated SSL representations provide competitive performance in both settings, demonstrating its feasibility while benefiting from unconstrained hardware environments (i.e., potentially unrealistic). We show that loss-based aggregation achieves better performance in both settings, which indicates that weakening the effects of low-quality clients can assist the aggregation process in federated SSL.

## 5. Conclusion

In this paper, we provided the first empirical study of the feasibility of combining two emerging techniques that will revolutionize the way speech systems are designed and operate: SSL and FL. Our results preview the future in which hybrid FL and SSL systems train universal speech representations. Our study highlights this future will not be easy to reach, with a number of technical challenges ahead of us. But we hope this work will act as a roadmap to the community towards this new milestone.

## 6. References

- [1] R. Arakawa, S. Takamichi, and H. Saruwatari, "Implementation of dnn-based real-time voice conversion and its improvements by audio data augmentation and mask-shaped device," in *Proc. SSW10*, 2019, pp. 93–98.
- [2] N. Nguyen, S. Sigg, A. Huynh, and Y. Ji, "Pattern-based alignment of audio data for ad hoc secure device pairing," in *2012 16th International Symposium on Wearable Computers*. IEEE, 2012, pp. 88–91.
- [3] R. S. Maloney and M. Paolisso, "What can digital audio data do for you?" *Field Methods*, vol. 13, no. 1, pp. 88–96, 2001.
- [4] S. K. Shah, Z. Tariq, and Y. Lee, "Audio iot analytics for home automation safety," in *2018 IEEE international conference on big data (big data)*. IEEE, 2018, pp. 5181–5186.
- [5] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [6] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *arXiv preprint arXiv:1904.05862*, 2019.
- [7] A. T. Liu, S.-w. Yang, P.-H. Chi, P.-c. Hsu, and H.-y. Lee, "Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6419–6423.
- [8] M. Ravanelli, J. Zhong, S. Pascual, P. Swietojanski, J. Monteiro, J. Trmal, and Y. Bengio, "Multi-task self-supervised learning for robust speech recognition," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6989–6993.
- [9] S. Pascual, M. Ravanelli, J. Serra, A. Bonafonte, and Y. Bengio, "Learning problem-agnostic speech representations from multiple self-supervised tasks," *arXiv preprint arXiv:1904.03416*, 2019.
- [10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [11] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, "SpeechBrain: A general-purpose speech toolkit," 2021, *arXiv:2106.04624*.
- [12] S. Kim, T. Hori, and S. Watanabe, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4835–4839.
- [13] X. Qiu, T. Parcollet, J. Fernandez-Marques, P. P. B. de Gusmao, D. J. Beutel, T. Topal, A. Mathur, and N. D. Lane, "A first look into the carbon footprint of federated learning," *arXiv preprint arXiv:2102.07627*, 2021.
- [14] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [15] Y. Gao, T. Parcollet, S. Zaiem, J. Fernandez-Marques, P. De Gusmao, D. Beutel, and N. Lane, "END-TO-END SPEECH RECOGNITION FROM FEDERATED ACOUSTIC MODELS," in *The International Conference on Acoustics, Speech, Signal Processing (ICASSP)*, Singapour, Singapore, May 2022. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03601224>
- [16] D. Leroy, A. Coucke, T. Lavril, T. Gisselbrecht, and J. Dureau, "Federated learning for keyword spotting," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6341–6345.
- [17] F. Granqvist, M. Seigel, R. van Dalen, Á. Cahill, S. Shum, and M. Paulik, "Improving on-device speaker verification using federated learning with privacy," *arXiv preprint arXiv:2008.02651*, 2020.
- [18] R. R. Schaller, "Moore's law: past, present and future," *IEEE spectrum*, vol. 34, no. 6, pp. 52–59, 1997.
- [19] J. Sevilla, L. Heim, A. Ho, T. Besiroglu, M. Hobbhahn, and P. Villedalobos, "Compute trends across three eras of machine learning," *arXiv preprint arXiv:2202.05924*, 2022.
- [20] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," *arXiv preprint arXiv:1912.06670*, 2019.
- [21] G. Chen, C. Parada, and G. Heigold, "Small-footprint keyword spotting using deep neural networks," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 4087–4091.
- [22] Y. Zhang, N. Suda, L. Lai, and V. Chandra, "Hello edge: Keyword spotting on microcontrollers," *arXiv preprint arXiv:1711.07128*, 2017.
- [23] X. Tan, T. Qin, F. Soong, and T.-Y. Liu, "A survey on neural speech synthesis," *arXiv preprint arXiv:2106.15561*, 2021.
- [24] D. Wang, X. Wang, and S. Lv, "An overview of end-to-end automatic speech recognition," *Symmetry*, vol. 11, no. 8, p. 1018, 2019.
- [25] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. 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.
- [26] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *arXiv preprint arXiv:2110.13900*, 2021.
- [27] P. Micikevicius, S. Narang, J. Alben, G. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu, "Mixed precision training," in *International Conference on Learning Representations*, 2018.
- [28] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino *et al.*, "Xls-r: Self-supervised cross-lingual speech representation learning at scale," *arXiv preprint arXiv:2111.09296*, 2021.
- [29] S. Hooker, "The hardware lottery," *Communications of the ACM*, vol. 64, no. 12, pp. 58–65, 2021.
- [30] D. J. Beutel *et al.*, "Flower: A friendly federated learning research framework," *arXiv preprint arXiv:2007.14390*, 2020.
- [31] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.