



HAL
open science

END-TO-END SPEECH RECOGNITION FROM FEDERATED ACOUSTIC MODELS

Yan Gao, Titouan Parcollet, Salah Zaiem, Javier Fernandez-Marques, Pedro de Gusmao, Daniel Beutel, Nicholas Lane

► **To cite this version:**

Yan Gao, Titouan Parcollet, Salah Zaiem, Javier Fernandez-Marques, Pedro de Gusmao, et al.. END-TO-END SPEECH RECOGNITION FROM FEDERATED ACOUSTIC MODELS. The International Conference on Acoustics, Speech, & Signal Processing (ICASSP), May 2022, Singapour, Singapore. ⟨hal-03601224⟩

HAL Id: hal-03601224

<https://hal.science/hal-03601224v1>

Submitted on 8 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

END-TO-END SPEECH RECOGNITION FROM FEDERATED ACOUSTIC MODELS

Yan Gao¹, Titouan Parcollet^{2,1}, Salah Zaiem³, Javier Fernandez-Marques⁴
Pedro P. B. de Gusmao¹, Daniel J. Beutel^{1,5}, and Nicholas D. Lane¹

¹University of Cambridge, ²Avignon University
³Telecom Paris, ⁴University of Oxford, ⁵Adap GmbH

ABSTRACT

Training Automatic Speech Recognition (ASR) models under federated learning (FL) settings has attracted a lot of attention recently. However, the FL scenarios often presented in the literature are artificial and fail to capture the complexity of real FL systems. In this paper, we construct a challenging and realistic ASR federated experimental setup consisting of clients with heterogeneous data distributions using the French and Italian sets of the CommonVoice dataset, a large heterogeneous dataset containing thousands of different speakers, acoustic environments and noises. We present the first empirical study on attention-based sequence-to-sequence End-to-End (E2E) ASR model with three aggregation weighting strategies – standard FedAvg, loss-based aggregation and a novel word error rate (WER)-based aggregation, compared in two realistic FL scenarios: *cross-silo* with 10 clients and *cross-device* with 2K and 4K clients. Our analysis on E2E ASR from heterogeneous and realistic federated acoustic models provides the foundations for future research and development of realistic FL-based ASR applications.

Index Terms— End-to-end ASR, federated learning

1. INTRODUCTION

Deep neural networks are now widely adopted in state-of-the-art (SOTA) ASR systems [1]. This success mostly relies on the centralised training paradigm where data needs first to be gathered in one single dataset before it can be used for training [2, 3, 4]. Such approach has a few clear benefits including fast training and, the ability to sample data in any preferred way due to the complete data visibility. However, recent concerns around data privacy along with the proliferation of both powerful mobile devices and low latency communication technologies (e.g. 5G), distributed training paradigms such as FL begin to receive more attention.

In FL, training happens at the source and training data is never sent to a centralised server. In a typical FL scenario each participating client/device receives a copy of an initial model and separately trains its model on its local data. This process generates a set of weight updates that are then sent to a server, where updates are aggregated. This process is

repeated for several rounds[5, 6, 7]. Being able to harvest information from thousands of mobile devices without having to collect users’ data makes federated and on-device training of ASR systems a feasible and an attractive alternative to traditional centralised training [6], whilst offering new opportunities to advance ASR quality and robustness given the unprecedented amount of user data directly available on-device. For example, such data could be leveraged to better adapt the ASR model to the users’ usage, or improve the robustness of models to realistic and low resources scenarios [7].

Despite the growing number of studies applying FL on speech-related tasks [8, 9, 10, 11, 12], very few of these have investigated its use for E2E ASR. Properly training E2E ASR models in a realistic FL setting comes with numerous challenges. First, it is notoriously complicated to train a deep learning model with FL on non independent and identically distributed data (non-IID) [7, 13, 14] and on-device speech data is extremely non-IID by nature (e.g. different acoustic environments, words being spoken, languages, microphones, amount of available speech, etc.). Second, state-of-the-art E2E ASR models are computationally intensive and potentially not suited for on-device training phases of FL. Indeed, the latest ASR systems rely on large Transformers [15, 16], Transducers [17, 18] or attention sequence-to-sequence (Seq2Seq) models [19, 20] that process high-dimensional acoustic features. Finally, E2E ASR training is difficult and very sensitive during early stages of optimisation due to the complexity of learning a proper alignment between the latent speech representation and the transcription. These three traits make it very difficult to train ASR models completely from scratch [21, 22].

To the best of our knowledge, existing works on FL for ASR typically approach these challenges by relinquishing few constraints of the environmental protocol. This in turn results in their experimental settings being still far away from the conditions in which a FL ASR would need to function. In fact, many works[23, 10] are evaluated on unrealistic datasets (*w.r.t* the FL scenario) such as LibriSpeech (LS) [24], which only contains recordings from selected speakers reading books in a controlled setting without background noise. Authors in [11] introduce a client-based adaptive training of a HMM-DNN based ASR system for *cross-silo* FL, a specific

setup which considers only a reduced number of clients with large amounts of homogeneous data, thus simplifying the complexity of real FL setups where data is non-IID. Then, [10] propose a federated transfer learning platform with improved performance using enhanced federated averaging and hierarchical optimization for E2E ASR, alleviating the aforementioned alignment issue but only evaluating it on LS.

In this work, we motivate the need to move away from clean speech corpora for evaluating FL-based ASR systems. We investigate FL models in a more realistic setting with the French Common Voice (CV) dataset [25], which provides a large, heterogeneous and uncontrolled set of speakers who used their own devices to record a given set of sentences; naturally fitting to FL with various users, acoustic conditions, microphones and accents. We evaluate both a *cross-silo* and a *cross-device* (i.e. large number of clients with few naturally non-IID data) FL setups while training a SOTA E2E ASR system. In particular, we compare three different weighting strategies during models aggregation. Our contributions are:

1. Quantitatively compare LibriSpeech to Common Voice towards a realistic FL setup to highlight the need for a shift in the evaluation of FL-based ASR models.
2. Present the first study on attention-based Seq2Seq E2E ASR model for FL scenarios. Our setup investigates challenges previously overlooked by others, such as extremely heterogeneous recording conditions.
3. Evaluate both *cross-silo* and *cross-device* FL with up to 4K clients on the naturally-partitioned and heterogeneous French and Italian subsets of Common Voice.
4. A first adapted aggregation strategy based on WER to integrate the specificity of ASR to FL.
5. Release the source code using Flower [26] and SpeechBrain [27] to facilitate replication and future research¹.

2. END-TO-END SPEECH RECOGNIZER

The considered E2E ASR system relies on the wide spread joint connectionist temporal classification (CTC) with attention paradigm [19]. This method combines a Seq2Seq attention-based model [28] with the CTC loss [29].

A typical ASR Seq2Seq model includes three modules: an encoder, a decoder and an attention module. Given a speech input sequence (i.e. speech signal or acoustic features) $\mathbf{x} = [x_1, \dots, x_{T_x}]$ with a length T_x , the *encoder* first converts it into a hidden latent representation $\mathbf{h}^e = [h_1^e, \dots, h_{T_x}^e]$. Then the *decoder* attends to the encoded representation \mathbf{h}^e combined with an attention context vector c_t from the attention module. This produces the different decoder hidden states $\mathbf{h}^d = [h_1^d, \dots, h_{T_y}^d]$, with T_y corresponding to the length of the target sequence \mathbf{y} . In a speech recognition scenario $T_x > T_y$.

¹github.com/yan-gao-GY/Flower-SpeechBrain

The standard training procedure of the joint CTC-Attention ASR pipeline is based on two different losses over a dataset S . First, the CTC loss is derived with respect to the prediction obtained from the encoder module of the Seq2Seq model:

$$\mathcal{L}_{CTC} = - \sum_S \log p(\mathbf{y}|\mathbf{h}^e), \quad (1)$$

Second, the attention-based decoder is optimised following the cross entropy (CE) loss:

$$\mathcal{L}_{CE} = - \sum_S \log p(\mathbf{y}|\mathbf{h}^d). \quad (2)$$

The losses are combined with a hyperparameter $\mu \in [0, 1]$ as:

$$\mathcal{L} = \mu \mathcal{L}_{CE} + (1 - \mu) \mathcal{L}_{CTC}. \quad (3)$$

In practice the CTC loss facilitates the early convergence of the system due its monotonic behavior while the attentional decoder needs to first figure out where to attend in the hidden representation of the entire input sequence.

3. FEDERATED TRAINING OF ACOUSTIC MODELS

The process of training an E2E acoustic model using federated learning follows four steps: 1) Following [10], model weights are initialised with a pre-training phase on a centralised dataset; 2) The centralised server samples K clients from a pool of M clients and uploads to them the weights of the model. 3) The clients train the model for t_{local} local epochs in parallel based on their local user data and send back the new weights to the server. 4) The server aggregates the weights and restart at step 2. This procedure is executed for T rounds until the model converges on a dedicated validation set (e.g. local to each client or centralised).

3.1. Federated Optimisation

For each training round, each client $k \in K$, containing n_k audio samples, runs $t \in [0, t_{local}]$ iterations with learning rate η_l to locally update the model based on Eq. 3,

$$w_{t+1}^{(k)} = w_t^{(k)} - \eta_l \tilde{g}_k, \quad (4)$$

with w_k the local model weights of client k , and \tilde{g}_k the average gradient over local samples. After training for t_{local} local epochs in the global round T , the updated weights $w_T^{(k)}$ of the client k are sent back to the server. Then, the local gradient $g_T^{(k)}$ is computed as:

$$g_T^{(k)} = w_T^{(k)} - w_{T-1}. \quad (5)$$

Then, the gradients from all clients are aggregated as follows:

$$\Delta_T = \sum_{k=1}^K \alpha_T^{(k)} g_T^{(k)}, \quad (6)$$

where $\alpha_T^{(k)}$ denotes different weighting strategies described in Section 3.2. The updated global model weights w_T are computed with a server learning rate η_s according to:

$$w_T = w_{T-1} - \eta_s \Delta_T, \quad (7)$$

During FL training, especially with heterogeneous data, the global model may deviate away from the original task or simply not converge [7, 13, 14], and therefore lead to performance degradation. To alleviate this issue, and motivated by [10], we propose an additional training iteration over a small batch of held-out data on the server, after the standard model update procedure with Eq.7.

3.2. Weighting Strategies

Federated Averaging (FedAvg) [5] is a popular [30, 31, 32, 8, 9, 23] aggregation strategy by which model updates from each client are weighted by $\alpha_T^{(k)}$, the ratio of data samples in each client over the total samples utilized in the round:

$$\alpha_T^{(k)} = \frac{n_k}{\sum_{k=1}^K n_k}, \quad (8)$$

In realistic FL settings with heterogeneous client data distribution, some clients may contain data that is skewed and not representative of the global data distribution (e.g. audio samples with different languages or multiple speakers). As a result, the aggregated model might simply not converge if such clients have proportionally more training samples than others. For instance, in our experiments, all attempts to train an ASR system from scratch failed due to this issue requiring a prior pre-training phase of the acoustic model. Second, clients containing low quality data would introduce unexpected noise into the training process (e.g. extreme noise in the background). Either scenario could lead to model deviation in the aggregation step, which can not be solved via the standard FedAvg weighting method (Eq. 8). A potential solution, instead, is to use the averaged training loss as a weighting coefficient, thus reflecting the quality of the locally trained model. Intuitively, higher loss would indicate that the global model struggles to learn from the client’s local data. More precisely, we compute the weighting with the *Softmax* distribution obtained from the training loss from Eq. 3. Eq. 8 is modified as follows:

$$\alpha_T^{(k)} = \frac{\exp(-\mathcal{L}_k)}{\sum_{k=1}^K \exp(-\mathcal{L}_k)}. \quad (9)$$

In the context of ASR, WER is commonly used as the final evaluation metric for the model instead of the training loss. We therefore propose a WER-based weighting strategy for aggregation. This approach utilizes the values $(1 - wer)$ obtained on the validation set as weighting coefficients $\alpha_T^{(k)}$:

$$\alpha_T^{(k)} = \frac{\exp(1 - wer_k)}{\sum_{k=1}^K \exp(1 - wer_k)}. \quad (10)$$

In this way, we directly optimise the model towards the relevant metric for speech recognition.

4. COMMON VOICE AS A REALISTIC FL SETUP

In this section we first present the Common Voice (CV) dataset used for the FL experiments. Then, we quantitatively demonstrate that CV is a much more adapted corpus to advance FL research than LibriSpeech (LS), motivating the need for a shift in the standard evaluation process.

4.1. Common Voice dataset

Both the French and Italian subsets of CV dataset (version 6.1) [25] are considered. Utterances are obtained from volunteers recording sentences all around the world, and in different languages, from smartphones, computers, tablets, etc. The French set contains a total of 328K utterances (475 hours in total) with diverse accents which were recorded by more than 10K French-speaking participants. The train set consists of 4212 speakers (425.5 hours of speech), while both validation and test sets contain around 24 hours of speech from 2415 and 4247 speakers respectively. The Italian set, on the other hand, is relatively small, containing 89, 21 and 22 hours of Italian training (748 speakers), validation (1219 speakers) and test (3404 speakers) data.

4.2. Setup analysis and LibriSpeech comparison

We argue that CV is closer to natural federated learning conditions than LS as much stronger variations are observed both intra- and inter-clients. While CV is a crowd-sourced dataset containing thousands of different acoustic conditions, microphones and noises, LS is a heavily controlled studio-quality corpus. The latter has been used by most research on FL ASR. We compare both datasets at three levels:

Low-level signal features. The selected features should be more descriptive of the background and recording conditions than speaker identity, as this is investigated when analysing clustering purity. Hence, we will consider: Loudness as it is highly linked to the microphone and the recording distance; the log of the Harmonicity to Noise Ratio (logHNR) as a proxy indicator of background noise; Permutation Entropy (PE) as it has been successfully used for microphone identification purposes[33].

The mean value of the signal feature is computed for every utterance by averaging the per-frame values. Then, for every client we compute the mean value and the standard deviation per client. The former distribution describes the inter-client variation while the latter describes the intra-client one. For the three considered features, the standard deviation of the mean value per client distribution is higher for Common Voice than for LibriSpeech, reaching 0.034, 11.466 and 0.053 for, respectively, Loudness, logHNR and Permutation

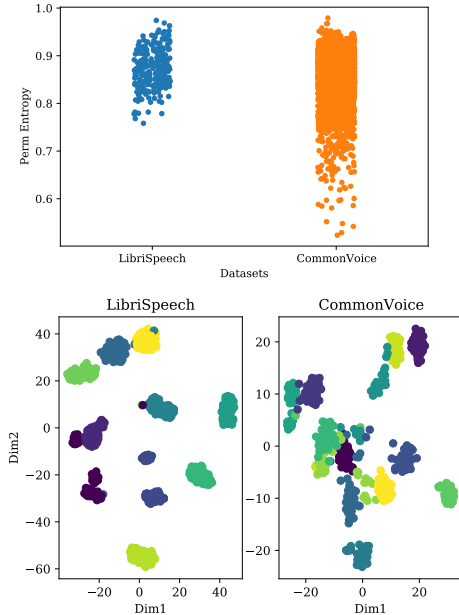


Fig. 1: (Above) Strip plot of the permutation entropy mean values per client in Librispeech and Common Voice. CV shows a heavy-tailed distribution as a consequence of the bigger diversity of recording settings. (Below) TSNE representation of embedded speech utterances. The colors correspond to the true clients (*i.e.* speakers). CV clients are clearly harder to separate than LS.

Entropy on CV compared to 0.017, 9.096 and 0.040 on LS. Concerning the intra-client variation, the standard deviation of the standard deviation per client distribution is also higher for CV than for LS reaching 0.009, 2.69 and 0.014 against 0.007, 2.31 and 0.007 respectively for loudness, logHNR and PE. It is also interesting to note the heavy tailed distribution obtained with the Permutation Entropy for CV, as depicted in Fig. 1. Indeed, the Kurtosis reaches 4.16 on CV versus -0.13 for LS. In practice, this mean that many clients may be outliers for CV, drastically impacting FL with conventional aggregation mechanisms.

Blind Signal-to-Noise ratios. We further inspect the noise difference between two datasets through computing a blind Signal to Noise Ratio estimation. First, a 10th-order LPC approximation is computed for every sample. Second, the voiced chunks are detected using the Probabilistic YIN algorithm for F0 estimation [34]. Finally, considering only the voiced chunks that are simpler to approach with an LPC estimate, the noise in the blind SNR estimation is defined as the difference between the real signal and the LPC approximation. Following the trend observed with the signal features, CV shows a higher standard deviation for the SNR mean values with $\sigma_{SNR_{CV}} = 18.47$ compared to $\sigma_{SNR_{LS}} = 10.32$ for LS. Then, a bigger variation within recordings of the same client is observed. Indeed, the standard deviation of the standard deviations obtained for each audio sample of the same

client is higher in CV than in LS, with 6.54 compared to 3.82. This suggests a higher variability in the recording conditions with respect to the same client. Common Voice speakers may contribute from different places and devices.

Clustering purity. We compare the overlap of speakers using pretrained speaker embeddings. For both datasets, speaker embeddings are computed on each utterance using the *Tristounet* model [35] open-sourced on *pyannote.audio* [36]. It is important to mention that *Tristounet* is not trained on LS or CV or audio book data. These embeddings are then clustered using the K-means algorithm with *kmeans++* initialization with the number of centroids equal to the number of clients. The purity of the clusters is defined as the proportion of points that belong to the same client as the majority of its computed cluster. Purity reaches 0.77 on LS and 0.62 on CV. Fig. 1 shows a TSNE representation of the utterance embeddings, and highlights the clustering difficulties in CV. This indicates that CV speakers are harder to separate using speaker embeddings. This confirms the two prior experiments using low-level audio features, as it suggests that varying signal features and recording conditions pollute the speech utterance which leads to harder speaker identification.

The analysis provided in this section evidences the drastic differences in corpora between LS and CV. The latter better captures the complexity that FL systems would face when deployed in the real world.

5. EXPERIMENTAL SETTINGS

This section first present the architecture of the E2E speech recognizer. Then, it describes the experimental setup of the FL environment alongside with key hyper-parameters.

5.1. E2E Speech Recognizer

The experiments are based on a Seq2Seq model trained with the joint CTC-attention objective [19]. The encoder follows the CRDNN architecture first described [27] (*i.e.* 2D CNN — LSTM — DNN). The decoder is a location aware GRU with a single hidden layer. The full set of parameters describing the model are given in the GitHub repository. Models are trained to predict sub-words units. No language model fusion is performed to properly assess the impact of the training procedure on the acoustic models. Data is augmented in the time-domain during training. The model has been implemented within SpeechBrain [27] and is therefore extremely easy to manipulate, customise and retrain.

5.2. Realistic Federated Learning

Based on the natural partitioning of the CV dataset we conduct two sets of experiments reflecting real usages of FL:

Cross-silo FL. In this scenario, clients are generally few, with high availability during all rounds and, often have similar




Fig. 2: Illustration of the sample distribution across the 2K-client and 4K-client FL settings from the French Common Voice dataset.

data distribution for training [7]. Shared data is often independent and identically distributed. Our implementation follows that of [10], the dataset is split in 10 random partitions (*i.e.* one per client) with no overlapping speakers each containing roughly the same amount of speech data.

Cross-device FL. This setup often involves thousands of clients having very different data distributions (non-IID), each participating in just one or a few rounds [7]. Here, we define two settings: First, we simulate a realistic scenario of single speaker using their individual devices. To reproduce this, we naturally divide the training sets based on users into 4095 and 649 partitions for French and Italian set, respectively. The second scenario allocates two users per device (e.g as in personal assistants or smart cars). For CV French, this lowers the number of clients to 2036. As depicted in Fig. 2, each setting drastically change the distribution of *low-resources* clients. The 4K setup offers a challenging scenario as most clients only contain very few samples.

5.3. Federated Learning for ASR: a hybrid approach

Training E2E ASR models in a FL setting is challenging. Jointly learning the alignment and the latent speech representation is a difficult task that commonly requires large datasets. Therefore, and as we experienced during our analysis, it is nearly impossible to train an E2E ASR model from scratch in a realistic FL setup. Table 1 shows that all the tested existing FL aggregation methods fail to converge without pre-training. This is due to the fact that most of the clients only contain few minutes of speech, resulting in an extremely noisy gradient to learn the alignment from. To overcome this issue we first pre-train the global model on half of the data samples. We do this by partitioning the original dataset into a small subset of speakers (with many samples) for centralised training (referred to subsequently as the *warm-up dataset*) and a much larger subset of speakers (having fewer samples each) for the FL experiment. For CV French, the small subset contains 117 speakers, leaving the remaining 4095 speakers for FL. Such statistics are reduced down to 99 and 649 speakers for Italian. We argue that this scenario remains realistic as, in practice,

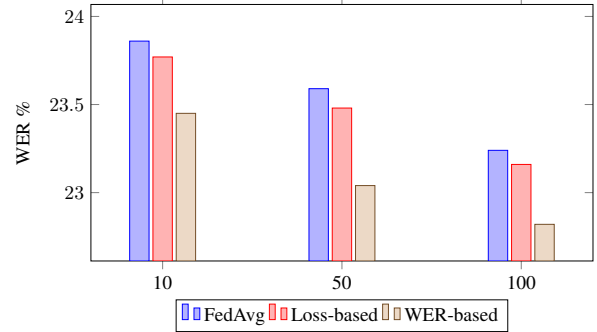


Fig. 3: Speech recognition performance when varying the number of sampled clients per round for the 4K-client setting on French CV.

centralised data is often available and can therefore be used to bootstrap the alignment.

The number of clients participating in each round influences the outcome of the experiments as well. To quantify this variation, we propose to vary the selected number of clients per round K from 10 to 100 for all weighting strategies on the 4K set. Then, we simply fix K with respect to the best WER obtained (*i.e.* 100) for the others setups. For the *cross-silo* environment, all clients are selected every round ($K = 10$).

In addition to setting the number of global rounds for the FL experiment, we must define as well the number of local epoch (*i.e.* on each client). This, however, is a non-trivial task [5]. In practice, we found that increasing the number of local epochs leads to clients over-fitting their own local data. Hence, clients are locally trained for only 5 epochs.

Depending on the available compute resources, training concurrently a large number of clients might become challenging. While models may be trained with CPUs or modest GPUs on real embedded hardware (e.g. RaspberryPi or NVIDIA Jetson), our simulated FL setup allows us to run these workloads on modern GPUs (e.g. Nvidia Tesla V100) running multiple clients concurrently on a single GPU and implemented with Flower [26] and SpeechBrain [27].

Models are finally evaluated both on a centralised test set and at the client-level with a small ensemble of local sentences. Indeed, for the French 4K setup, each client saves 10% (with a minimum of 2 samples) for testing purposes. To be more specific, centralised speakers are new ones, while local speakers have been seen at training time.

6. SPEECH RECOGNITION RESULTS

First, we compare the impact of selecting different numbers of clients K per round on the most challenging setup (4K clients in French, Fig. 3). Conversely to the literature, higher values of K tend to produce better WER. This is explained by the heterogeneity of the CV dataset, for which extremely noisy clients may perturb the averaging process with few clients per round. Indeed, few clients remain at more than 100% of WER

Table 1: Speech recognition results on the centralised test sets of French (“Fr”) and Italian (“It”) CV dataset for different scenarios and weighting strategies. “User-based” FL represents 4K clients for French and 649 for Italian.

Training Scenario		Fr WER (%)	It WER (%)
Centralised	All data (lower bound)	20.18	17.40
	1 st half (<i>warm-up</i>)	25.26	25.90
	2 nd half (post <i>warm-up</i>)	20.94	24.86
10-clients FL <i>Cross-silo</i>	FedAvg	21.26	20.97
	Loss-based	21.10	20.86
	WER-based	20.99	19.98
2K-clients FL <i>Cross-device</i>	FedAvg	22.83	—
	Loss-based	22.67	—
	WER-based	22.42	—
User-based FL <i>Cross-device</i>	FedAvg	23.24	24.32
	Loss-based	23.16	24.23
	WER-based	22.82	23.86
From scratch	FedAvg [5]	100+	100+
	FedProx [37], FedAdam [38]	100+	100+

even after the full training. For the remaining of the experiments, K will thus be fixed to 100.

Table 1 reports the results obtained across the different training setups. We notice that training on the entire dataset in a *centralised* way gives us the best WER with 20.18% and 17.40% for the French and Italian sets respectively, which is comparable to the current best literature [27]. This lower-bound is expected as the system has full visibility of the data and can sample the inputs in an almost IID fashion. On the other hand, when using only the *warm-up* dataset, we notice the effect of having fewer data points for training as the WER increases to 25.26% for the French set and 25.90% for the Italian set. This is expected as well as the system has now less data to learn from. This sheds some light on the inherent lower-bound limitations of FL, limited to partial data observations in each round. The third centralised scenario trains the warmed-up model on the 2nd half of data in an on-line training fashion. This model provides a slightly lower WER compared to all FL models in French set. However, we should note that this is an unrealistic setting as training models in a centralised way would void all the privacy guarantees that FL offers. In particular, this model only gains 0.14% improvement in Italian set compared with the *warm-up* model. This indicates the difficulty of training model on the second half data even in centralised fashion. The results on all FL settings exceed centralised training thanks to the centralised fine-tuning in between each round on the server side.

The effect of data visibility can indeed be seen in both *cross-silo* and *cross-device* scenarios, which do not have uniform access to data. However, since this problem is less severe in the former setup, with the correct choice of aggregation strategy we are still able to obtain a WER of 20.99% with the French set, which is very close to the centralised lower bound of 20.18%. The more challenging Italian set, on the other hand, obtains 19.98% WER with a 2.58% difference to the lower bound. As for the *cross-device* scenario, the effect of non-IID data distribution among devices leads to its best

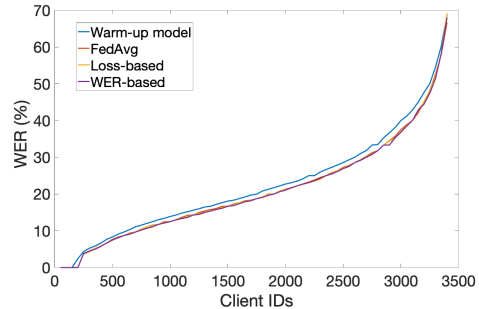


Fig. 4: Client test performance on the French set of Common Voice for different weighting strategies. The average WER for *warm-up* model, standard FedAvg, loss-based and WER-based aggregation are 23.76%, 22.13%, 22.11% and 21.91%. Clients are sorted w.r.t their WER. Clients with a WER higher than 100% are removed.

WER on French set being 22.43% and 22.82% in the 2K and 4K clients settings, even worse (23.86%) with the Italian set. These values are larger than the worst *cross-silo* result, showing the effects of the non-IID nature of the data partitioning.

Compared to different weighting strategies, WER-based and loss-based methods obtain a better performance in all settings, which indicates that weakening the effects of low-quality clients can assist the aggregation process in federated training with heterogeneous data distribution. Herein, we have two types of indicators reflecting the quality of clients. The results in Tab. 1 show that WER-based strategy obtain the lowest WER in both settings. This could be easily explained by the nature of the strategy which directly optimise the model toward the relevant metric for speech recognition.

Client level test performance is another concern in realistic FL. Fig. 4 shows the individual WER for each client on French set. All FL methods obtain better performance than the *warm-up* model (blue line), but the difference between the three aggregation strategies becomes less significant. WER-based method, however, obtains the best WER 21.91% when calculating the average performance over all the clients. As previously discussed, we can see that many clients still have a WER higher than 50% and 500 of them even have a local WER higher than 100%, clearly indicating the challenging nature of the CV dataset for FL.

7. CONCLUSION

In this paper, we presented the first study for realistic FL scenarios on attention-based Seq2Seq E2E ASR model with three aggregation weighting strategies – standard FedAvg, loss-based aggregation and a novel WER-based aggregation. We quantitatively compared LibriSpeech and Common Voice towards a realistic FL setup. All methods were evaluated with *cross-silo* and *cross-device* FL on two languages. Our work sets the foundations for future research of realistic FL ASR applications with an open source environment.

8. REFERENCES

- [1] Akshi Kumar, Sukriti Verma, and Himanshu Mangla, “A survey of deep learning techniques in speech recognition,” in *2018 International Conference on Advances in Computing, Communication Control and Networking (ICACCCN)*. IEEE, 2018, pp. 179–185.
- [2] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, et al., “Deep speech: Scaling up end-to-end speech recognition,” *arXiv preprint arXiv:1412.5567*, 2014.
- [3] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al., “Deep speech 2: End-to-end speech recognition in english and mandarin,” in *International conference on machine learning*, 2016, pp. 173–182.
- [4] Hagen Soltau, Hank Liao, and Hasim Sak, “Neural speech recognizer: Acoustic-to-word lstm model for large vocabulary speech recognition,” *arXiv preprint arXiv:1610.09975*, 2016.
- [5] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [6] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [7] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al., “Advances and open problems in federated learning,” *arXiv preprint arXiv:1912.04977*, 2019.
- [8] Andrew Hard, Kurt Partridge, Cameron Nguyen, Niranjan Subrahmanya, Aishanee Shah, Pai Zhu, Ignacio Lopez Moreno, and Rajiv Mathews, “Training keyword spotting models on non-iid data with federated learning,” *arXiv preprint arXiv:2005.10406*, 2020.
- [9] David Leroy, Alice Coucke, Thibaut Lavril, Thibault Gisselbrecht, and Joseph 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.
- [10] Dimitrios Dimitriadis, Kenichi Kumatani, Robert Gmyr, Yashesh Gaur, and Sefik Emre Eskimez, “A federated approach in training acoustic models,” in *Proc. Interspeech*, 2020.
- [11] Xiaodong Cui, Songtao Lu, and Brian Kingsbury, “Federated acoustic modeling for automatic speech recognition,” *arXiv preprint arXiv:2102.04429*, 2021.
- [12] Filip Granqvist, Matt Seigel, Rogier van Dalen, Áine Cahill, Stephen Shum, and Matthias Paulik, “Improving on-device speaker verification using federated learning with privacy,” *arXiv preprint arXiv:2008.02651*, 2020.
- [13] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra, “Federated learning with non-iid data,” *arXiv preprint arXiv:1806.00582*, 2018.
- [14] Felix Sattler, Simon Wiedemann, Klaus-Robert Müller, and Wojciech Samek, “Robust and communication-efficient federated learning from non-iid data,” *IEEE transactions on neural networks and learning systems*, vol. 31, no. 9, pp. 3400–3413, 2019.
- [15] Abdelrahman Mohamed, Dmytro Okhonko, and Luke Zettlemoyer, “Transformers with convolutional context for asr,” *arXiv preprint arXiv:1904.11660*, 2019.
- [16] Albert Zeyer, Parnia Bahar, Kazuki Irie, Ralf Schlüter, and Hermann Ney, “A comparison of transformer and lstm encoder decoder models for asr,” in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 8–15.
- [17] Mehryar Mohri, Fernando Pereira, and Michael Riley, “Weighted finite-state transducers in speech recognition,” *Computer Speech & Language*, vol. 16, no. 1, pp. 69–88, 2002.
- [18] Eric Battenberg, Jitong Chen, Rewon Child, Adam Coates, Yashesh Gaur Yi Li, Hairong Liu, Sanjeev Satheesh, Anuroop Sriram, and Zhenyao Zhu, “Exploring neural transducers for end-to-end speech recognition,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 206–213.
- [19] Suyoun Kim, Takaaki Hori, and Shinji Watanabe, “Joint ctc-attention based end-to-end speech recognition using multi-task learning,” in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 4835–4839.
- [20] Chung-Cheng Chiu, Tara N Sainath, Yonghui Wu, Rohit Prabhavalkar, Patrick Nguyen, Zhifeng Chen, Anjali Kannan, Ron J Weiss, Kanishka Rao, Ekaterina Gonnina, et al., “State-of-the-art speech recognition with

- sequence-to-sequence models,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4774–4778.
- [21] Andrew Rosenberg, Kartik Audhkhasi, Abhinav Sethy, Bhuvana Ramabhadran, and Michael Picheny, “End-to-end speech recognition and keyword search on low-resource languages,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5280–5284.
- [22] Sameer Bansal, Herman Kamper, Karen Livescu, Adam Lopez, and Sharon Goldwater, “Pre-training on high-resource speech recognition improves low-resource speech-to-text translation,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019, pp. 58–68.
- [23] Dhruv Guliani, Françoise Beaufays, and Giovanni Motta, “Training speech recognition models with federated learning: A quality/cost framework,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 3080–3084.
- [24] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [25] Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber, “Common voice: A massively-multilingual speech corpus,” *arXiv preprint arXiv:1912.06670*, 2019.
- [26] Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchu Qiu, Titouan Parcollet, and Nicholas D Lane, “Flower: A friendly federated learning research framework,” *arXiv preprint arXiv:2007.14390*, 2020.
- [27] Mirco Ravanelli, Titouan Parcollet, Peter Plantinga, Aku Rouhe, Samuele Cornell, Loren Lugosch, Cem Subakan, Nauman Dawalatabad, Abdelwahab Heba, Jianyuan Zhong, Ju-Chieh Chou, Sung-Lin Yeh, Szu-Wei Fu, Chien-Feng Liao, Elena Rastorgueva, François Grondin, William Aris, Hwidong Na, Yan Gao, Renato De Mori, and Yoshua Bengio, “Speech-Brain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.
- [28] Dzmitry Bahdanau, Jan Chorowski, Dmitriy Serdyuk, Philemon Brakel, and Yoshua Bengio, “End-to-end attention-based large vocabulary speech recognition,” in *2016 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2016, pp. 4945–4949.
- [29] Alex Graves and Navdeep Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *International conference on machine learning*, 2014, pp. 1764–1772.
- [30] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [31] Samuel Horvath, Stefanos Laskaridis, Mario Almeida, Ilias Leontiadis, Stylianos I. Venieris, and Nicholas D. Lane, “Fjord: Fair and accurate federated learning under heterogeneous targets with ordered dropout,” 2021.
- [32] Xinchu Qiu, Titouan Parcollet, Javier Fernandez-Marques, Pedro Porto Buarque de Gusmao, Daniel J. Beutel, Taner Topal, Akhil Mathur, and Nicholas D. Lane, “A first look into the carbon footprint of federated learning,” 2021.
- [33] Gianmarco Baldini and Irene Amerini, “An Evaluation of Entropy Measures for Microphone Identification,” *Entropy*, vol. 22, no. 11, 2020.
- [34] Matthias Mauch and Simon Dixon, “Pyin: A fundamental frequency estimator using probabilistic threshold distributions,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 659–663.
- [35] Hervé Bredin, “Tristounet: Triplet loss for speaker turn embedding,” 2017.
- [36] Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill, “pyannotate.audio: neural building blocks for speaker diarization,” 2019.
- [37] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith, “Federated optimization in heterogeneous networks,” *arXiv preprint arXiv:1812.06127*, 2018.
- [38] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan, “Adaptive federated optimization,” *arXiv preprint arXiv:2003.00295*, 2020.