

Unsupervised Clustered Federated Learning in Complex Multi-source Acoustic Environments

Alexandru Nelus, Rene Glitza, and Rainer Martin

Institute of Communication Acoustics, Ruhr-Universität Bochum

44780 Bochum, Germany

firstname.lastname@rub.de

Abstract—In this paper we introduce a realistic and challenging, multi-source and multi-room acoustic environment and an improved algorithm for the estimation of source-dominated microphone clusters in acoustic sensor networks. Our proposed clustering method is based on a single microphone per node and on unsupervised clustered federated learning which employs a light-weight autoencoder model. We present an improved clustering control strategy that takes into account the variability of the acoustic scene and allows the estimation of a dynamic range of clusters using reduced amounts of training data. The proposed approach is optimized using clustering-based measures and validated via a network-wide classification task.

Index Terms—clustered federated learning, clustering, privacy, acoustic sensor networks, autoencoder, unsupervised

I. INTRODUCTION

Acoustic sensor networks (ASNs) are gaining increasing popularity thanks to their capacity to exploit information from various (wireless) and sometimes ad-hoc acoustic sensors. ASNs have become hosts to many applications [1], [2], e.g., source localization [3], event classification [4], speech enhancement [5], etc. These, in turn, can benefit from knowledge of source-dominated clusters formed by microphone-carrying nodes in the vicinity of distributed acoustic sources [6], [7].

To support these emerging applications we propose an improved solution for unsupervised clustering of ad-hoc sensor nodes based on clustered federated learning (CFL) [8] which is derived from federated learning (FL) [9]. The latter can be considered as privacy-preserving, decentralized collaborative machine learning [10]. In this approach, ASN nodes (clients) locally train a neural network (NN) model and solely share locally learned NN weight updates with a central node (server). In this way, it is ensured that audio data is confined to the node level. Moreover, these NN weight updates are aggregated by the server and the new master model is downloaded by the clients. Clusters are then formed on the basis of weight update similarities and the CFL cycle can be further repeated inside each cluster to yield sub-clusters.

In our approach, the (*unsupervised*) light-weight autoencoder at node level provides the flexibility to adapt online to varying and unseen acoustic conditions, room configurations, speakers, noise sources etc. Exploiting this advantage, we propose pre-training of the autoencoder before CFL, freezing

all layers except the bottleneck layer, and only updating the latter during CFL. This allows a drastic reduction of the number of NN weight updates sent to and aggregated by the server, along with avoidance of overfitting issues that come along with small training sets.

For the evaluation of the proposed clustering method we introduce a complex acoustic environment, based on the apartment layout described in [11], with multiple rooms that have different reverberation conditions, open doors, furniture, and decorations along with multiple (≥ 2) sources of interest and without prior information about their number. In order to handle these challenging conditions, novel CFL control criteria are proposed that take into consideration the high acoustic variability of each cluster and allow the estimation of clusters using only 40 s of audio data. Moreover, the computation of cluster membership values is also updated in order to handle the variety of clusters generated. The evaluation methods consist of distance-based measures and an application in the form of a network-wide gender recognition task.

The remainder of this paper is structured as follows: we present the relation to both state-of-the-art and own prior work, after which we detail the proposed methods. This is followed by a description of the experimental results along with a discussion of results and conclusions.

II. RELATION TO PRIOR WORK

Microphone clustering in ASNs has been previously explored using e.g., coherence-based features [12], [13], eigenvectors [7], divergence of power spectral densities [14] or cepstral features [6], [15]. These clustering solutions and applications, although effective, do not explicitly incorporate privacy considerations and evaluations are confined to shoebox-type room scenarios.

CFL [16] has been mainly demonstrated in (semi-) supervised learning applications with (weak) classification labels and relatively large data sets [8], [17], [18]. Moreover, hard-clustering without the possibility of generating membership values has been used. We have recently addressed some of the aforementioned aspects and successfully explored the application of an unsupervised CFL clustering scheme in ASNs [19]. However, the latter work is also limited to a shoebox room with only two simultaneously active sources and, consequently, only two clusters. The number of sources was considered to be known and was used as a stopping

This work has been supported by the German Research Foundation (DFG) - Project Number 282835863.

criterion for the clustering algorithm. Estimation of clusters was performed on audio segments of 160 s.

III. UNSUPERVISED CLUSTERED FEDERATED LEARNING

Data submitted by distributed clients to a collaborative training system can be non-i.i.d. Therefore, a single global vector of model parameters θ^* may not be able to minimize the loss of all clients at the same time. This impediment has led to the development of the *clustered* FL approach [8]. CFL groups clients with *congruent* data distributions into individual clusters such that each cluster c learns its own model parameters θ_c . To this end, we define a matrix containing the L_2 -normalized weight update vectors of all M nodes, $\Delta\theta = (\Delta\theta_1/\|\Delta\theta_1\|, \dots, \Delta\theta_M/\|\Delta\theta_M\|)$ and compute the *cosine similarity matrix* $A = \Delta\theta^T \Delta\theta$ between the nodes' weight update vectors. Bi-partitioning the set of nodes results in clusters c_1 and c_2 , such that the minimum intra-cluster similarity is larger than the maximum inter-cluster similarity. CFL along with the generation of a new cosine similarity matrix A_c , are performed for each new cluster c . Bi-partitioning is further applied, if needed, and the process is repeated until the data congruence condition is satisfied. Clusters with only two nodes are no longer bi-partitioned.

In order to assess the congruence property within a cluster c , we make use of the mean and the maximum Euclidean norms of the clients' weight update vectors, defined as [8]

$$\Delta\bar{\theta}_c = \left\| \frac{1}{|c|} \sum_{i \in c} \Delta\theta_i \right\| \quad \text{and} \quad \Delta\hat{\theta}_c = \max_{i \in c} (\|\Delta\theta_i\|). \quad (1)$$

In the standard CFL approach, $\Delta\bar{\theta}_c \leq \varepsilon_1$ and $\Delta\hat{\theta}_c \geq \varepsilon_2$, where ε_1 and ε_2 are empirically set, indicate that the system has reached a stationary solution but this is not optimal for all clients (large individual gradients), thus prompting bi-partitioning. To account for varying cluster sizes and absolute values of weight update vectors caused by differing audio sources and room acoustics, we propose to use a dynamic initialization of ε_1 as a weighted sum of $\Delta\bar{\theta}_c$ and $\Delta\hat{\theta}_c$. This is done at cluster level in the CFL communication round $\tau = 0$. Additionally, we replace the second condition $\Delta\hat{\theta}_c \geq \varepsilon_2$ by the ratio between the mean and maximum norms $\Delta\bar{\theta}_c/\Delta\hat{\theta}_c \leq \varepsilon_2$ and thereby arrive at a more robust normalized test criterion.

The aforementioned congruence assessment is performed after a minimum of $\tau > \min_\tau$ communication rounds and the entire algorithm is stopped after \max_τ rounds. Furthermore, an additional stopping criterion $n_{\text{no-split}} > \varepsilon_3$ is proposed in order to account for the weight updates' divergence caused by training for too many consecutive communication rounds on congruent data. The CFL algorithm, along with the proposed modifications, is comprehensively detailed in Algorithm 1.

A. Autoencoder description

One of the necessary elements in adapting CFL to ASN scenarios that lack training labels is the incorporation of an unsupervised model, namely an autoencoder, thus leading to unsupervised CFL. Moreover, the short time intervals on the basis of which a clustering estimation is desired greatly limit

Algorithm 1: Unsupervised CFL for the estimation of source-dominated microphone clusters in ASNs.

Input: Pre-trained autoencoder h , thresholds ε_2 and ε_3 , lower (\min_τ) and upper (\max_τ) communication rounds bounds

Freeze all parameters of h except bottleneck subset θ

while *audio buffer* \neq *empty* **do**

 Read audio data D of M clients

 Initialize cluster list $C \leftarrow \{\{1, \dots, M\}\}$ with a single cluster element that contains all M clients

$\theta_c \leftarrow \theta \leftarrow$ random initialization, $\forall c \in C$

$n_{\text{no-split}} \leftarrow -1$, $\tau \leftarrow 0$

while $\tau \leq \max_\tau$ **and** $n_{\text{no-split}} \leq \varepsilon_3$ **do**

for $i \leftarrow 1, \dots, M$ **in parallel** **do**

Client i does:

$\theta_i^\tau \leftarrow \theta_c^\tau$, $i \in c$ and $c \in C$

$\Delta\theta_i^\tau \leftarrow \text{SGD}(\theta_i^\tau, D_i) - \theta_i^\tau$

end

Server does:

$C' \leftarrow \{\}$

$\text{split} \leftarrow \text{false}$

for $c \in C$ **do**

 Compute A_c and $\Delta\bar{\theta}_c, \Delta\hat{\theta}_c$ from (1)

if $\tau == 0$ **then** $\varepsilon_1 \leftarrow \Delta\bar{\theta}_c + 0.1\Delta\hat{\theta}_c$

if $\tau > \min_\tau$ **and** $\Delta\bar{\theta}_c \leq \varepsilon_1$ **and** $\frac{\Delta\bar{\theta}_c}{\Delta\hat{\theta}_c} \leq \varepsilon_2$ **then**

$\{c_1, c_2\} \leftarrow \text{Bi-partitioning}(c, A_c)$

$\theta_{c_x}^{\tau+1} \leftarrow \theta_c^\tau + \frac{1}{|c_x|} \sum_{i \in c_x} \Delta\theta_i^\tau$, $\forall c_x \in \{c_1, c_2\}$

$\text{split} \leftarrow \text{true}$

$n_{\text{no-split}} \leftarrow 0$

$C' \leftarrow C' + \{c_1, c_2\}$

else

$\theta_c^{\tau+1} \leftarrow \theta_c^\tau + \frac{1}{|c|} \sum_{i \in c} \Delta\theta_i^\tau$

$C' \leftarrow C' + \{c\}$

end

end

if not split **and** $n_{\text{no-split}} \geq 0$ **then**

$n_{\text{no-split}} \leftarrow n_{\text{no-split}} + 1$

end

$\tau \leftarrow \tau + 1$

$C \leftarrow C'$

end

end

the amount of training data clients can train on, potentially leading to overfitting issues. For this reason, [19] proposes pre-training an autoencoder h on more extensive data and then freezing all its layers except the bottleneck layer. The latter can be re-initialized with random values and further adapted with much smaller amounts of data by individual clients.

Table I summarizes the model, with the fifth layer being the bottleneck layer. To train the model, we use the mean squared error (MSE) loss between the input Y and reconstructed \hat{Y} log-mel band energy (LMBE) feature vectors.

TABLE I
NEURAL NETWORK ARCHITECTURE OF AUTOENCODER h .

Layer	Input	Operator	Out ch.	Stride	Kernel/Nodes	Activation
1	128 x 128	Conv2d	6	1	5 x 5	ReLu
2	6 x 124 x 124	MaxPool	-	2	2 x 2	-
3	6 x 62 x 62	Conv2d	16	1	5 x 5	ReLu
4	16 x 58 x 58	MaxPool	-	2	2 x 2	-
5	16 x 29 x 29	Dense	-	-	29	ReLu
6	16 x 29 x 29	Unpool	-	2	2 x 2	-
7	16 x 58 x 58	ConvTrans2d	6	1	5 x 5	ReLu
8	6 x 62 x 62	Unpool	-	2	2 x 2	-
9	6 x 124 x 124	ConvTrans2d	1	1	5 x 5	Sigmoid

B. Cluster membership values

As observed in fuzzy clustering approaches [15], determining cluster membership values (MVs) offers vital benefits to ASN-based applications. While standard CFL provides only hard clustering results, we have previously extended CFL to also compute MVs [19]. We now refine this method in order to cope with more than two sources and two clusters.

For each cluster $c \in C$, we first compute the mean intra- and inter-cluster similarities of each client $i \in c$ using A_c and stack them in vectors q and r , respectively, with

$$q_i = \frac{1}{|c|-1} \sum_{j \in c \setminus \{i\}} A_{i,j} \text{ and } r_i = \frac{1}{M-|c|} \sum_{k \in C^* \setminus c} A_{i,k}, \quad (2)$$

where $|\cdot|$ denotes the cardinality of a set, C is the total set of clusters, M is the total number of clients and $C^* = \{1, \dots, M\}$. Next, we apply min-max normalization to the intra- and inter-cluster similarities q and r , and balance them in vector p ,

$$p = \lambda q + (1 - \lambda)r, \quad (3)$$

with elements p_i . This is motivated by the observation that nodes closest to a source have not only a small inter-cluster similarity but also a small intra-cluster similarity. The node with the smallest p_i value is then selected as the cluster's reference node, assuming this is the most representative node for the source. The min-max normalized cosine similarities between the latter and all other nodes form the cluster's MVs vector μ . Thresholding with $\mu_i = 0, \forall \mu_i \leq v$ is applied in order to disregard nodes with low MVs. The impact of λ and the MV thresholding are further studied in Section IV.

IV. EXPERIMENTAL SETUP AND RESULTS

A. Database and rooms

We employ a subset of the LibriSpeech corpus [20], namely *train-clean-100*, with 251 speakers (125 female, 126 male) extracted from 16 kHz audiobook recordings. Voice activity

TABLE II
ROOM-SPECIFIC FLOOR AREA, SIMULATED REVERBERATION TIME (T_{60}) AND NUMBER OF RANDOMLY POSITIONED SOURCES AND NODES.

Room	Floor Area	T_{60}	# Sources	# Nodes
Living room	29.14 m ²	0.37 s	2	24
Bedroom	6.51 m ²	0.27 s	1	5
Toilet	1.22 m ²	0.33 s	0	2
Corridor	2.86 m ²	0.42 s	0	5
Bathroom	3.77 m ²	0.53 s	1	5

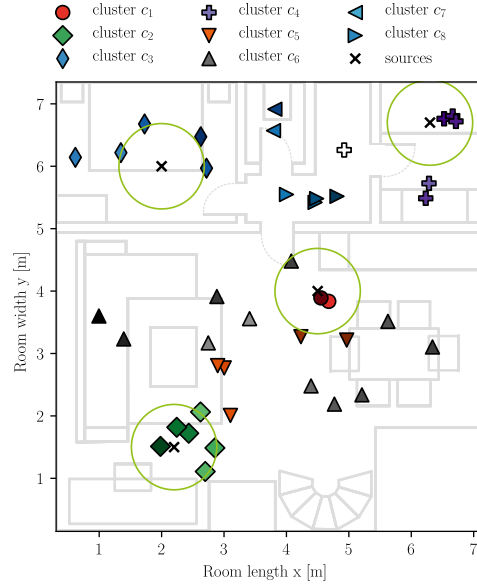


Fig. 1. Floor plan of simulated SINS apartment with cluster estimations for a single scenario. Color intensity is proportional to cluster membership values and green circles indicate the aggregated critical distance.

detection (VAD) is applied and the data is restructured into 25006 utterances of length 10 s each. The data is further split into *Libri-server* with 157 speakers (79 female, 78 male) for autoencoder and gender recognizer training, and *Libri-clients* (94 speakers) for clustering and speaker gender inference.

The apartment layout introduced in [11] and schematically illustrated in Fig. 1, is transformed into a 3D model. Auralization using CATT-Acoustic with cone-tracing [21] is then performed. Open doors, typical furniture, utensils, and decorations are included. The respective T_{60} reverberation times along with the rooms' floor areas, number of randomly positioned nodes, and simultaneously active sources are presented in Table II. The random positioning of nodes is made under the constraint that for every source in the living room, a minimum of three nodes are within critical distance, thus having higher direct component energy than reverberation energy. The audio signal $x_i(t)$ captured by each ASN node i is expressed as

$$x_i(t) = \sum_{z=0}^{N_S} s_z(t) * g_{z,i}(t), \quad (4)$$

where $g_{z,i}$ is the room impulse response (RIR) from source s_z to node i and N_S is the scenario's total number of sources.

In order to analyze the generalization of our proposed clustering method, we consider two types of ASN scenes. The first, *2SL*, only involves the living room which contains two simultaneously active sources (male and female). The second, *4SA*, involves the entire apartment and contains four simultaneously active sources (two male and two female). For each ASN scene, ten randomly positioned source-node constellations are created and for each constellation 20 gender-balanced speaker pairs are randomly selected from *Libri-clients* resulting in 200 simulation scenarios. We randomly select four utterances/speaker (10 s each) to perform CFL followed by gender recognition using the estimated clusters.

TABLE III

NORMALIZED CLUSTER-TO-SOURCE DISTANCE $\tilde{d}_{c_x}^{s_z}$ FROM CLUSTER c_x TO SOURCE s_z , AVERAGED OVER 200 SCENARIOS. RESULTS INCLUDE ENTIRE SINS APARTMENT (4SA) AND LIVING ROOM ONLY (2SL).

4SA	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}	c_{11}
s_1	0.20	0.84	0.81	0.97	0.44	0.71	0.77	0.72	0.72	0.53	0.69
s_2	0.81	0.24	1.19	1.11	0.92	0.69	1.12	1.12	1.07	0.92	1.65
s_3	0.81	1.04	0.11	1.01	0.68	0.75	0.46	0.46	0.77	0.82	0.79
s_4	0.98	1.01	1.00	0.08	0.83	0.80	0.63	0.63	0.42	0.82	1.01
2SL	c_1	c_2	c_3	c_4	c_5	c_6					
s_1	0.08	0.95	0.37	0.78	0.48	0.76					
s_2	0.94	0.09	0.68	0.28	0.67	0.30					

B. Autoencoder pre-training

The autoencoder h is pre-trained on the Libri-server set for 300 epochs using an SGD optimizer with a learning rate of $l_r = 0.1$. For each 10 s utterance, an LMBE feature vector \mathbf{Y} is extracted, as detailed in [22], using a short-time discrete Fourier transform (STFT) with window length $L_1 = 0.064$ s and step size $R_1 = 0.032$ s, along with $K = 128$ mel filters.

C. Clustering

The autoencoder h described above is used for unsupervised CFL as indicated in Algorithm 1. All the Θ parameters, except for the subset θ , that corresponds to the bottleneck layer, are frozen. The latter is randomly re-initialized before each clustering estimation which is now performed for four utterances (total of 40 s). This reduces the initial number of trainable parameters from $O_1 = 5999$ to $O_2 = 841$. The CFL control criteria are set to $\varepsilon_2 = 0.84$ and $\varepsilon_3 = 2$.

In order to evaluate clustering performance in relation to previous work [19] and state-of-the-art implementations [15], we extend the normalized cluster-to-source distance (CTS) $\tilde{d}_{c_x}^{s_z}$ of cluster c_x to source s_z [15] to multiple sources and clusters,

$$\tilde{d}_{c_x}^{s_z} = \frac{\|\rho_{s_z} - \bar{\rho}_{c_x}\|}{\bar{d}_S}, \quad \forall c_x \in C \text{ and } s_z \in S, \quad (5)$$

where ρ_{s_z} is the position of source s_z . The centroid $\bar{\rho}_{c_x}$ represents the average of geometric positions of nodes i assigned to cluster c_x weighted by their respective MVs, \bar{d}_S is the average of all unique source pair distances, S is the set of all sources, and C is the set of all estimated clusters.

Table III shows $\tilde{d}_{c_x}^{s_z}$ averaged over 200 simulation scenarios for each 4SA and 2SL scenes respectively. A small value of $\tilde{d}_{c_x}^{s_z}$ for $x = z$ indicates that $\bar{\rho}_{c_x}$ is very close to ρ_{s_z} , with $\tilde{d}_{c_x}^{s_z} = 0$ denoting an exact superposition. This goes along with a high value of $\tilde{d}_{c_x}^{s_z}$ for the other sources where $\tilde{d}_{c_x}^{s_z} > 1$ indicates that the distance between a source and a cluster centroid is larger than the average distance of sources. For both 4SA and 2SL scenes, we can indeed observe the aforementioned effects for the first four and two columns, respectively. When compared to [15] and [19], which both use unfurnished shoebox rooms, it can be observed that the results generated by the improved version of unsupervised CFL proposed in this work show great potential, especially for the 2SL scene which clearly outperforms [19].

Moreover, given the distinct characteristics of each combination of speaker groups and ASN constellations, our approach

TABLE IV

NUMBER OF SIMULATION SCENARIOS GENERATING CLUSTER c_x (N_{c_x}) AND MEAN NUMBER OF CLIENTS IN c_x ($\bar{N}_{\text{clients}, c_x}$). RESULTS INCLUDE ENTIRE SINS APARTMENT (4SA) AND LIVING ROOM ONLY (2SL).

4SA	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}	c_{11}
N_{c_x}	200	200	192	188	182	163	125	89	52	10	1
$\bar{N}_{\text{clients}, c_x}$	10.4	5.1	5.2	5.1	5.2	5.7	4.8	4.6	4.1	3.1	5
2SL	c_1	c_2	c_3	c_4	c_5	c_6					
N_{c_x}	200	200	79	32	8	4					
$\bar{N}_{\text{clients}, c_x}$	13.8	6.9	6.2	3.6	6.5	2.5					

generates a dynamic range of clusters. Since we do not limit the a priori number of clusters/sources and the acoustic environment is reverberant and diverse, CFL identifies clusters of similar weight update vectors not only in the vicinity of sources but also in areas where several sources mix, e.g., hallway (c_8) and toilette (c_7) clusters in Fig. 1. Thus the clustering results are able to reflect the scenario's acoustic diversity which is a good basis for advanced ASN applications. The number of simulation scenarios N_{c_x} where a cluster c_x is estimated along with its average number of nodes $\bar{N}_{\text{clients}, c_x}$ is presented in Table IV for both for 2SL and 4SA. This indicates good clustering reliability as at least N_S but no more than $3 * N_S$ clusters are frequently generated.

Systematic variations of λ from (3) were included in the MV calculation. For the 2SL and 4SA scene the best results were obtained using $\lambda = 0.5$ and $\lambda = 0$, respectively. This indicates that relying solely on the inter-cluster similarities is sufficient when this uses clusters from more than two sources, otherwise intra-cluster information is also helpful.

D. Gender recognition

Gender recognition was proposed as an objective measure for evaluating clustering utility and, indirectly, its performance in both [15] and [19]. To allow a direct comparison, we employ the gender recognition model e described in [19] which is trained on the Libri-server dataset along with RIR-based data augmentation for 13 epochs using a cross-entropy loss function and an SGD optimizer with $l_r = 0.01$. The input feature representation \mathbf{Y} is extracted for each 10 s utterance with $L_1 = 0.064$, $R_1 = 0.02$, and $K = 40$.

The evaluation metrics used are the Accuracy (A_{cc}) and F_1 -score (F_1) between the predicted and the ground truth gender labels of a cluster. These are averaged across all clusters in a scenario and averaged again over all 200 simulation scenarios. The predicted and ground truth gender labels of a cluster are given by the mode (majority decision) of the predicted and ground truth gender labels of its nodes. When using MVs, the predicted gender of a cluster is the MV-weighted average of its nodes. The ground truth gender of a node is given by the gender of the source with the shortest first peak delay of the RIR. The predicted gender of a node is the mode of the predicted gender across four utterances (40 s).

To gain a broader perspective on the capabilities and limitations of unsupervised CFL, we propose two meaningful prior-knowledge cases. In the first, we solely rely on the clustering results of the first N_S clusters from Table III. While this

TABLE V
GENDER RECOGNITION ACCURACY (A_{cc}) AND F_1 -SCORE (F_1) OF
ESTIMATED CLUSTERS, WITHOUT AND WITH MEMBERSHIP VALUE (MV)
WEIGHTING USING THRESHOLD v , AVERAGED OVER 200 SCENARIOS.
CLOSEST-TO-SOURCE (CTS) CLUSTERS OR CLUSTERS WITH HIGHEST
CONFIDENCE (CONFIDENCE-BASED) ARE SELECTED.

Full apartment (4SA)	no MV	$v = 0$	$v = 0.5$	$v = 0.8$	$v = 0.9$	$v = 0.95$
CTS-based A_{cc} (%)	83.4	91.5	93.4	93.7	93.2	92.7
CTS-based F_1 (%)	82.4	90.1	93.1	93.2	92.7	92.0
Confidence-based A_{cc} (%)	83.4	92.1	92.8	92.8	92.2	91.4
Confidence-based F_1 (%)	82.4	91.3	92.1	92.1	91.3	90.4
Living room (2SL)						
CTS-based A_{cc} (%)	88.1	98.5	99.0	99.8	99.8	99.5
CTS-based F_1 (%)	87.0	98.5	99.0	99.8	99.8	99.5
Confidence-based A_{cc} (%)	88.1	98.8	99.3	99.8	99.3	99.0
Confidence-based F_1 (%)	87.0	98.8	99.3	99.8	99.3	99.0

information is unlikely to be available in practice, it provides a good benchmark. In the second case, we select the N_S clusters with the highest gender classification NN output confidence. These might not necessarily be the closest to each source.

Results are presented in Table V for both 4SA and 2SL ASN scenes. As a general observation, it is clear that the proposed MV generation scheme has a great impact on improving the performance of cluster-level gender predictions. Moreover, an MV threshold of $v = 0.8$ appears to offer the best overall outcome as it discards some non-helpful nodes. Interestingly, much more aggressive thresholding does not offer better performance as too few nodes remain for gender decision aggregation, once more emphasizing the inherent benefits of clustering. Gender recognition achieves excellent performance for the 2SL ASN scene regardless of the prior knowledge type included. For the 4SA ASN scene, gender recognition is, unsurprisingly, more challenging, and using the N_S closest-to-each-source clusters does offer a small advantage. Nonetheless, it is interesting to see that even when including information from more reverberant clusters, the proposed clustering scheme offers significant advantages to the network-wide gender recognition task.

V. CONCLUSIONS

An improved version of unsupervised CFL was proposed in order to handle challenging acoustic environments with multiple simultaneously active sources and minimum training data requirements. The approach is robust and versatile, estimating not only the nearest source-dominated clusters but also those in the reverberant sound field. This, in turn, shows great potential for enhancing additional ASN-deployable tasks such as wake word detection or event classification in future works.

REFERENCES

- [1] A. Bertrand, "Applications and trends in wireless acoustic sensor networks: A signal processing perspective," in *2011 18th IEEE symposium on communications and vehicular technology in the Benelux (SCVT)*. IEEE, 2011, pp. 1–6.
- [2] S. Pasha, C. Ritz, and J. Lundgren, "A survey on ad hoc signal processing: Applications, challenges and state-of-the-art techniques," in *2019 IEEE International Symposium on Signal Processing and Information Technology (ISSPIT)*. IEEE, 2019, pp. 1–6.
- [3] G. F. Miller, A. Brendel, W. Kellermann, and S. Gannot, "Misalignment recognition in acoustic sensor networks using a semi-supervised source estimation method and markov random fields," *arXiv:2011.03432*, 2020.

- [4] J. Ebberts, L. Drude, R. Haeb-Umbach, A. Brendel, and W. Kellermann, "Weakly supervised sound activity detection and event classification in acoustic sensor networks," in *2019 IEEE 8th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2019, pp. 301–305.
- [5] S. Markovich-Golan, A. Bertrand, M. Moonen, and S. Gannot, "Optimal distributed minimum-variance beamforming approaches for speech enhancement in wireless acoustic sensor networks," *Signal Processing*, vol. 107, pp. 4–20, 2015.
- [6] S. Gergen, R. Martin, and N. Madhu, "Source separation by fuzzy-membership value aware beamforming and masking in ad hoc arrays," in *Speech Communication; 13th ITG-Symposium*, 2018, pp. 1–5.
- [7] M. H. Bahari, L. K. Hamaidi, M. Muma, J. Plata-Chaves, M. Moonen, A. M. Zoubir, and A. Bertrand, "Distributed multi-speaker voice activity detection for wireless acoustic sensor networks," *arXiv:1703.05782*, 2017.
- [8] F. Sattler, K.-R. Müller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–13, 2020.
- [9] B. McMahan, E. Moore, D. Ramage, and B. Arcas, "Federated learning of deep networks using model averaging," *arXiv:1602.05629*, 2016.
- [10] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.
- [11] G. Dekkers, S. Lauwereins, B. Thoen, M. W. Adhana, H. Brouckxon, T. van Waterschoot, B. Vanrumste, M. Verhelst, and P. Karsmakers, "The SINS database for detection of daily activities in a home environment using an acoustic sensor network," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*, November 2017, pp. 32–36.
- [12] I. Himawan, I. McCowan, and S. Sridharan, "Clustering of ad-hoc microphone arrays for robust blind beamforming," in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 2814–2817.
- [13] S. Pasha, J. Donley, and C. Ritz, "Blind speaker counting in highly reverberant environments by clustering coherence features," in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2017, pp. 1684–1687.
- [14] Y. Zhao, J. K. Nielsen, J. Chen, and M. G. Christensen, "Model-based distributed node clustering and multi-speaker speech presence probability estimation in wireless acoustic sensor networks," *The Journal of the Acoustical Society of America*, vol. 147, no. 6, pp. 4189–4201, 2020.
- [15] S. Gergen, A. Nagathil, and R. Martin, "Classification of reverberant audio signals using clustered ad hoc distributed microphones," *Signal Process.*, vol. 107, no. C, p. 21–32, Feb. 2015.
- [16] F. Sattler, K. Müller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multi-task optimization under privacy constraints," *arXiv:1910.01991*, 2019.
- [17] F. Sattler, S. Wiedemann, K. Müller, and W. Samek, "Sparse binary compression: Towards distributed deep learning with minimal communication," *International Joint Conference on Neural Networks, IJCNN 2019 Budapest, Hungary, July 14-19, 2019*, pp. 1–8, 2019.
- [18] F. Sattler, K. Müller, T. Wiegand, and W. Samek, "On the byzantine robustness of clustered federated learning," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. IEEE, 2020, pp. 8861–8865.
- [19] A. Nelus, R. Glitza, and R. Martin, "Estimation of microphone clusters in acoustic sensor networks using unsupervised federated learning," *arXiv:2102.03109*, 2021.
- [20] 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, South Brisbane, Queensland, Australia, April 19-24, 2015*. IEEE, 2015, pp. 5206–5210.
- [21] B.-I. Dalenbäck, *TUCT v2.0e:1*, CATT, Mariagatan 16A, SE-41471 Gothenburg, Sweden, 2019. [Online]. Available: <http://www.catt.se>
- [22] A. Nelus, J. Ebberts, R. Haeb-Umbach, and R. Martin, "Privacy-Preserving Variational Information Feature Extraction for Domestic Activity Monitoring versus Speaker Identification," in *Proc. Interspeech 2019*, 2019, pp. 3710–3714.