Privacy Preserving Loneliness Detection: A Federated Learning Approach

Malik Muhammad Qirtas
School of Computer Science and Information Technology
University College Cork
Cork, Ireland
malik.qirtas@ieee.org

Dirk Pesch
School of Computer Science and Information Technology
University College Cork
Cork, Ireland
dirk.pesch@ucc.ie

Evi Zafeiridi
School of Computer Science and Information Technology
University College Cork
Cork, Ireland
EZafeiridi@ucc.ie

Eleanor Bantry White
School of Applied Social Studies
University College Cork
Cork, Ireland
E.BantryWhite@ucc.ie

***Abstract*—Today's smartphones have sensors that enable monitoring and collecting data on users' daily activities, which may be converted into behavioral indicators of users' health and well-being. Although previous research has used passively sensed data through smartphones to identify users' mental health state, including loneliness, anxiety, depression, and even schizophrenia, the issue of user data privacy in this context has not been well addressed. Here we focus on the feeling of loneliness, which, if persistent, is associated with a number of negative health outcomes. While modern artificial intelligence technology, specifically machine learning, can assist in detecting loneliness or depression, current approaches have applied machine learning to centrally collected user data at a single location with the potential to compromise user data privacy. To address the issue of privacy, we investigated the feasibility of using federated learning on single user data to identify loneliness collected by different smartphone sensors. Federated learning can help protect user privacy by avoiding the transmission of sensitive data from mobile devices to a central server location. To evaluate the federated method's performance in detecting loneliness, we also trained models on all user data using a centralised machine learning approach and compared the results. The results indicate that federated learning has considerable promise for detecting loneliness in a binary classification problem while maintaining user data privacy.**

***Index Terms*—mHealth, sensing, wearables, privacy, loneliness, federated learning**

I. Introduction

Loneliness has been a global concern over the last several decades, with many seeing it as the primary source of unhappiness in their life [1]. Loneliness is stated as a condition in which a person has a quantitative or qualitative absence of social connections [2]. Loneliness has a variety of negative health impacts, ranging from insomnia to increased anxiety, depression, and weakened immune function. It is a common problem that most people experience at some stage in their lives; however, it can be very distressing, especially when this problem becomes chronic [3]. Additionally, it is associated with poor cardiovascular and cognitive health. Loneliness has been linked to a 40% increase in the risk of dementia [4]. These negative consequences led the United States' 17th Surgeon General to warn that one of the major health concerns confronting Americans is an "epidemic of loneliness" [5].

Identifying that a person is lonely is crucial for reducing loneliness. However, loneliness detection can be challenging. In recent years, smartphones have shown potential as ecologically viable instruments for predicting one's behavior and mental well-being. Sensors in these smartphones have created the possibility of transforming them into a robust health tracking system. Our scoping review on loneliness detection through passive sensing [6] considered the insight that the sensors' data streams provided by smartphones can be utilized to monitor users' daily lives and behaviors, which can then be used as markers of mental health.

This has opened up new avenues for researchers to use smartphones as passive sensing devices to determine an individual's mental well-being. Passive sensing refers to the practice of gathering data in the background using the multiple sensing capabilities of smartphones and other wearable devices without the user's active engagement. Smartphones can collect large volumes of data about a user's behavioral patterns, which may be turned into bio-indicators of the user's mental and physical health.

Prior research has used machine learning approaches in combination with a number of clinical loneliness scales to detect loneliness in individuals using passively collected data through their smartphone sensors. The majority of studies employed a traditional machine learning model training process, in which all data from multiple participants is collected at a single server, which trains the model on the whole data set [7, 8, 9]. The primary drawback of this technique is that sensitive data about users is collected at a single location, jeopardizing

This work has been funded by SFI Centre for Research Training in Advanced Networks for Sustainable Societies (ADVANCE CRT), which is a part of Science Foundation Ireland.

their privacy in the event of a data breach or by potentially allowing others to access this data. In contracts, we applied a federated learning technique in this work to train individual machine learning models using each user's data locally (e.g. on their smartphone). The key benefit of this approach is that it safeguards a user's privacy by sharing just a subset of training results with a central server and transmitting no private or sensitive data to the server, such as data about a user's identity or smartphone usage.

We discuss past research on loneliness detection through passive sensing, present our federated learning method and architecture, and the results from different algorithms used for loneliness detection through federated learning. We empirically evaluate the efficacy of using federated learning in detecting loneliness using data acquired passively via smartphone sensors. A subset of the StudentLife dataset [10] has been used for this purpose with 4 different binary classifiers in a federated learning environment using the Flower framework [11]. Flower is a user-friendly framework designed for implementing federated learning systems. We also trained the same models in a centralised machine learning environment to evaluate the performance of federated versus centralised methods for loneliness detection using passive sensing data. Random Forest performed the best in centralised machine learning models, with an accuracy of 85%, followed by the XGBoost (83%) support vector machine (78.5%) and logistic regression (73%). In the case of the federated learning approach, XGBoost performed best with an accuracy of 77.5%, followed by random forest (72%), SVM (68%), and logistic regression (51%). While in this initial study centralised learning achieves better detection accuracy than federated learning, the difference is not large, suggesting that the privacy preserving federated learning approach bears the potential to close the gap to centralised learning performance with further development.

II. Related Work

Past research using passive sensing to detect loneliness did not focus sufficiently on data privacy. We have examined the existing literature for loneliness detection through passive sensing in our review article with various aspects of these studies, especially population, privacy, and validation issues [6]. Among the works we reviewed, the following are examples that did not specifically consider user privacy. Sanchez et al. [12] utilized machine learning to infer the level of loneliness in 12 older persons who spent one week using a mobile app. The phones' call logs and global positioning system (GPS) coordinates were gathered. The authors developed four centralised machine learning models, for family loneliness, spousal loneliness, societal loneliness, and existential crises, with estimated accuracy of 91.6 %, 83.3 %, 66.6 %, and 83.3 %, respectively. But since they used a centralised data collection approach, data privacy remains a problem with this approach.

In a similar way, Pulekar et al. [13] adopted a machine learning strategy to detect loneliness in 9 students by analyzing data gathered passively via smartphones. They gathered communication data from study participants (call logs, SMS logs, browsing data, emails, and social media activity) as well as social interaction data (through Bluetooth and WIFI encounters). Additionally, they explored the relationship between the Big Five Personality traits and feelings of loneliness. However, the sample size was quite small, making generalization to a wider population difficult. Additionally, as a result of centralised machine learning model training, all data was gathered at a single server via participants' smartphones, jeopardizing their privacy in case of a data breach.

Doryab et al. [13] passively collected data from 160 users via smartphone and fitness bands in order to detect loneliness. They analyzed SMS and call logs, location data, Bluetooth and WiFi addresses from smartphones, and health data from fitness trackers. They employed machine learning models to train and infer individuals' levels of loneliness. Three methods of analysis were presented to determine the feasibility of passively detecting loneliness via smart devices: statistical analysis (based on UCLA responses), data mining analysis (presented behavioral patterns using smartphone and fitness band data), and machine learning analysis (loneliness detection and change in loneliness level over an academic semester using smartphone and fitness band data). They detected loneliness with a high accuracy of 76%. Since the models were trained on all students' data at a centralised server, it compromises the privacy of each student's personal data.

Other research used passive sensing in a smart home context to detect loneliness. A smart home can include a variety of sensors, depending on the application. For instance, numerous studies have used in-home surveillance through video cameras, while others have used body-worn tags [14, 15]. On the other hand, inexpensive ambient sensors allow for a more unobtrusive method of monitoring behaviours in the home without the user's involvement. Such ambient in-home sensing has been used for human behaviour learning over the last few years in which emotions, daily life patterns, or personality could be related to loneliness levels. Those studies have also used machine learning approaches to infer different behavioural patterns of users through the data collected with ambient sensors. However, the data privacy issue has not been well addressed in those studies either since all data was also collected at a central server for detecting loneliness.

All previous research studies relied on centralised machine learning techniques, which compromises users' privacy by transferring data to a central server. The main objective of this work is to assess if a federated learning strategy, where learning takes place on a user's own device that also collects the data, can achieve similar performance to a centralised learning strategy that learns on all user's data at a central location.

III. Methodology

A. Federated Learning Architecture

In Federated Learning (FL), each client trains their model locally on an individual basis. In other words, each client

TABLE I
SUBSET OF STUDENTLIFE DATASET

Sensors	Personality	Survey
Activity	Openness	UCLA
Conversation	Consciousness	
Location	Agreeableness	
Calls and SMS Logs	Neuroticism	
Bluetooth	Extraversion	

undergoes a unique model training process. Only learnt local model parameters are forwarded to a trusted server for aggregation into an aggregated global model. The server then returns to these clients the aggregated global model, and this process is repeated. In our experiment, we simulate this process by training local models on each student's mobile device using only that device's data, and then sharing those local model parameters with a central server for the global model's aggregation process.

FL does not store data centrally; rather, it enables users to train a shared model collectively using their local data [16]. As a result, it inherently ensures data privacy and reduces transmission costs. The learning activity is often accomplished through a loose federation of participants (clients) overseen by a central coordinator (server). Rather than uploading local data to the server, the clients calculate changes to the shared model maintained by the server, and only this update is sent. Below are the steps for a single epoch in FL

1) Due to the fact that the parameters of the global model and the local models in the clients are randomly initialized, all of these parameters of local and global models will be distinct. As a result, the global model communicates its parameters to the clients prior to the clients training their local models.
2) Using these parameters, clients begin training their local models using their own local data.
3) While each client is training its own local model, it updates its parameters simultaneously. Following the training process, each client transmits its local model weights to the global model at a central server. No private data will leave the client device during this step.
4) The global model takes the average of these parameters shared by all the clients and sends them back to the clients for the next iteration as its new weight parameters.

These steps complete one epoch. We can set the number of epochs to repeat all the above steps again and again to improve the performance of the global model.

We utilized the Flower framework [11] to construct a federated learning environment. Aggregation algorithms on the server side are crucial in FL because they must generate a global model from the weights of each client without access to the users' private data and then share this global model with each individual client. Flower offers a variety of federation algorithms for model aggregation on the server side, and

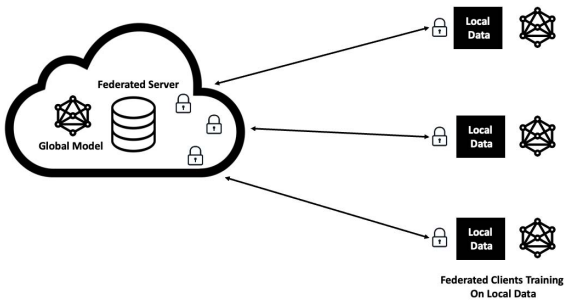

Fig. 1. Federated Learning Architecture: The central server communicates with the clients without exchanging any private data from clients. Still, the global model on the server uses the model weights from each client's local models to achieve a higher overall performance.

we chose the FedAvg algorithm, which has been frequently utilized in the literature for classification problems and it has shown excellent results for different types of data [17].

FedAvg manages the training using a central server that maintains the shared global model θ_t, where t is the communication round. FedAvg contains five hyperparameters: the proportion of clients C for model training, the size of the local mini-batch B, the number of local epochs on clients E, a learning rate, and a learning rate decay λ. The FedAvg algorithm is presented in Algorithm 1, where m = C X k is the total number of clients participating in the process. We have a single global model θ_t and multiple local client level models θ^k. All the models have the same model structure but different parameter values. So based on this, direct model aggregation can be implemented as described in line 9 of Algorithm 1, in which all uploaded local models θ^k are weighted and averaged based on the ratio n_k/n, which is proportional to the amount of data on each client to generate the global model θ_t. The communication between server and clients only consists of the global model parameters θ_t and local model parameters θ^k.

The process starts by initializing the global model θ_0 randomly. A FedAvg communication round thus consists of the following: The server picks a subset of m clients, $m = C \cdot K \geq 1$, and sends the current global model θ_t to all m clients. Clients then update their local models θ_k with the new shared model from the server. After that, each client trains the local model for E epochs by splitting the local data into batch size B. After local training, clients then share their locally trained models θ_k with the central server. The server then computes a weighted sum of all received models and generates a new global model θ_t. This new global model is then shared again with all clients.

Algorithm 1 FedAvg Algorithm. K is the total number of clients; B is the size of mini-batches, T is the total number of communication rounds, E is the total local training epochs, and η is the learning rate

Server:
Initialise global model θ_0
for each communication round t=1,2,...T **do**
  Select $m = C \cdot K$ clients, where $C \in (0, 1)$
  for each **Client** k=1,2,...m in parallel **do**
    Download θ_t to **Client** k
    Do **Client** k update and receive θ^k
  end for
  Update global model $\theta_t \leftarrow \sum_{k=1}^{m} \frac{n_k}{n} \theta^k$
end for

Client k update:
Replace local model $\theta^k \leftarrow \theta_t$
for local epoch from 1 to E **do**
  for batch $b \in (1, B)$ **do**
    $\theta^k \leftarrow \theta^k$ - $\eta\Delta\ L_k$ (θ^k,b)
  end for
end for
Return θ^k

B. Dataset

We have used a subset of the Student Life dataset [10] from Dartmouth College. This data was collected from 48 students over a ten-week period. The StudentLife project collected sensing data from smartphones, including accelerometer, microphone, light sensor, GPS and Bluetooth. Additionally, the dataset contains the answers of a mental health survey administered to students to ascertain their levels of depression, stress, and loneliness.

In terms of loneliness measures, the dataset contains scores on the UCLA loneliness scale [18], a 20-item questionnaire designed to assess subjective feelings of loneliness. Ten of these items are positive, while the remainder are negative. This scale measures loneliness on a scale of 20 to 100, with higher scores indicating greater levels of loneliness. Scores greater than 44 indicate a strong sense of loneliness. The lowest score in the sample data was 25, while the highest score was 64.

C. Data Pre-processing

During data pre-processing, we removed data from students who did not complete the second loneliness questionnaire, resulting in data from 41 students for our study. We converted the UNIX timestamps of each sensor data into a human-readable local date and time format using each participant's timezone data. Students engage in a variety of activities throughout the day, which is why we divided a 24-hour period into three sessions — day session (9am – 6pm), evening session (6pm – 12am), and night (12am – 9am). We addressed the training dataset's class imbalance using the synthetic minority oversampling approach (SMOTE) [19], which creates synthetic data for the minority class, resulting in a balanced training data set. This approach has been applied to each student's data to handle data imbalance per client basis. To handle missing values, we first deleted all records containing outliers and then imputed continuous missing data for each participant using the median of a particular feature. For categorical data, we used the mode of that particular feature. Since the feature scaling has no effect on tree-based algorithms, we scaled the numerical features for the rest of the algorithms. We utilized the "standard scaler" in our models since it transforms the data in such a manner that it has a mean of zero and standard deviation of one, and it normalizes the data.

Additionally, we used the Reproducible Analysis Pipeline for Data Streams (RAPIDS) to calculate behavioral features for each student [20]. By quantifying the per-participant per-epoch behavioral patterns (i.e. routines, irregularity, and variability) in these student data sets using basic counts, standard deviations, entropy, and regularity index measures, we generated digital biomarkers (features). For each participant, all computed features were merged on per epoch level (day, evening and night). We retrieved 117 features in total from sensor data; explanations of these features can be found in the RAPIDS documentation [20, 21]. Categorical features were converted into integer representation via one-hot encoding.

D. Loneliness Detection Using Centralised Machine Learning

We treated the problem of detecting loneliness as a binary classification; 1 for lonely and 0 for not lonely. We have categorized lonely and non-lonely students based on their UCLA survey scores. Any student with a UCLA overall score of 44 or higher is considered lonely. We trained four different binary classifiers using all of the data from 41 students for centralised machine learning. We employed the algorithms of logistic regression, random forest, support vector machine, and XGBoost. We used the 10-cross validation technique for model evaluation and to calculate accuracy, precision, recall, and F1 Score. The results of the model with the final feature set are presented in the evaluation section.

E. Loneliness Detection Using Federated Learning

We have evaluated a number of different binary classifiers, including logistic regression, support vector machine, random forest and XGBoost [22]. Following our Federated Learning approach, these models are trained on local clients using each student's data set, and then the weights of trained models are shared with the server for aggregation into a global model. A total of 41 clients are used for training local models. Each model is trained for 10 epochs. Initially, model weights are randomly initialised with a pre-training phase on the centralised server and then the server shares the model weights with each of the 41 clients. The clients then train the model for local epochs in parallel based on their local user data and send back the new weights to the server. The server then aggregates the weights and restarts the next epoch.

We assessed accuracy, as measured by the percentage of correctly classified samples; precision, as measured by the percentage of classified samples that actually belonged to a class; recall, as measured by the percentage of class samples

that were accurately classified, and F1 score as measured by the harmonic mean of precision and recall [22].

IV. Evaluation

We formulated the loneliness detection problem as a binary classification problem with 2 classes based on students' scores on the UCLA loneliness scale. Class 0 represents individuals who are not lonely, while class 1 represents lonely individuals.

Random forest performed the best in terms of accuracy for the centralised machine learning models, with a score of 85%. While XGBoost has a slightly lower accuracy of 83%, it has the highest precision, recall, and F1 score of the any algorithm. The table II summarizes the performance of all classifiers used in the centralised machine learning technique.

The performance of different classifiers trained using the federated learning approach is listed in table III. As shown in the table, non-linear classifiers XGBoost and random forest have better performance compared to linear classifiers, such as logistic regression and support vector machine. In terms of all evaluation metrics; accuracy, precision, recall and F1 score, XGBoost is the best performing classifier in FL approach followed by the random forest, support vector machine and logistic regression.

The preliminary results from our experiments indicate that centralised machine learning models outperform federated learning in general, but the performance of tree-based algorithms in FL is close to that of centralised machine learning models. The reason for the low performance by FL could be due to the data heterogeneity problem which may exist in some clients who are having only single class representation. Although FL models perform less well overall than centralised models, they can provide data privacy and a lower communication cost as they do not transfer raw sensing data from client devices to the central server, which could be more important than the best accuracy in certain sensitive scenarios.

TABLE II
Performance of centralised machine learning models classifying loneliness using a subset from the StudentLife dataset.

	Accuracy	**Precision**	**Recall**	**F1 Score**
XGBoost	83%	93%	89%	91.5%
RF	85%	89%	82%	87%
SVM	78.5%	77%	83%	81%
Logistic Regression	73%	68%	79%	75%

TABLE III
Performance of federated learning models classifying loneliness using a subset from the StudentLife dataset.

	Accuracy	**Precision**	**Recall**	**F1 Score**
XGBoost	77.5%	81%	84%	81.5%
RF	72%	61.5%	76%	68%
SVM	68%	58%	71%	67%
Logistic Regression	51%	43%	61%	56%

V. Conclusion and Future work

This paper presented a comparison between centralised machine learning and federated learning for detecting student loneliness based on passively acquired mobile sensor data. The paper's primary objective was to investigate the feasibility of utilizing federated learning to identify loneliness based on smartphone sensor data while maintaining privacy. The results indicate that while centralised machine learning models perform better than FL models, federated learning has great potential for loneliness detection via passive sensing due to its ability to train models on a user's device without requiring the user to share data, thus maintaining privacy. Tree-based algorithms showed the best performance for both centralised and FL approaches.

In future work, we will apply various hyperparameter tuning strategies to FL models in order to increase their performance. We will also use deep learning methods for model training and will compare results. Additionally, we will investigate the importance and relevance of features in federated learning methods for effectively detecting loneliness. Moreover, we plan to incorporate other sensor data; phone lock/unlock data, audio data and app usage data to extract additional features for loneliness detection.

References

[1] John T Cacioppo and William Patrick. *Loneliness: Human nature and the need for social connection*. WW Norton & Company, 2008.

[2] Jenny de Jong-Gierveld. "Developing and testing a model of loneliness." In: *Journal of personality and social psychology* 53.1 (1987), p. 119.

[3] Marilyn Campbell. "Loneliness, social anxiety and bullying victimization in young people: A literature review". In: *Psychology and Education* 50.3-4 (2013), pp. 1–10.

[4] Angelina R Sutin et al. "Loneliness and risk of dementia". In: *The Journals of Gerontology: Series B* 75.7 (2020), pp. 1414–1422.

[5] V Murthy. *The Surgeon General's prescription of happiness. TEDMED*. 2016.

[6] Malik Muhammad Qirtas et al. "Loneliness and Social Isolation Detection Using Passive Sensing Techniques: Scoping Review". In: *JMIR mHealth and uHealth* 10.4 (2022), e34638.

[7] Afsaneh Doryab et al. "Identifying behavioral phenotypes of loneliness and social isolation with passive sensing: statistical analysis, data mining and machine learning of smartphone and fitbit data". In: *JMIR mHealth and uHealth* 7.7 (2019), e13209.

[8] Alicia Martinez et al. "A Predictive Model for Automatic Detection of Social Isolation in Older Adults". In: *2017 International Conference on Intelligent Environments (IE)*. IEEE. 2017, pp. 68–75.

[9] Congyu Wu et al. “Improving prediction of real-time loneliness and companionship type using geosocial features of personal smartphone data”. In: *Smart Health* 20 (2021), p. 100180.

[10] Rui Wang et al. “StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones”. In: *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*. 2014, pp. 3–14.

[11] Daniel J Beutel et al. “Flower: A friendly federated learning research framework”. In: *arXiv preprint arXiv:2007.14390* (2020).

[12] Wendy Sanchez et al. “Inferring loneliness levels in older adults from smartphones”. In: *Journal of Ambient Intelligence and Smart Environments* 7.1 (2015), pp. 85–98.

[13] Gauri Pulekar and Emmanuel Agu. “Autonomously sensing loneliness and its interactions with personality traits using smartphones”. In: *2016 IEEE Healthcare Innovation Point-Of-Care Technologies Conference (HI-POCT)*. IEEE. 2016, pp. 134–137.

[14] Tamara L Hayes et al. “Distributed healthcare: Simultaneous assessment of multiple individuals”. In: *IEEE Pervasive Computing* 6.1 (2007), pp. 36–43.

[15] Joshua R Smith et al. “RFID-based techniques for human-activity detection”. In: *Communications of the ACM* 48.9 (2005), pp. 39–44.

[16] Chen Zhang et al. “A survey on federated learning”. In: *Knowledge-Based Systems* 216 (2021), p. 106775.

[17] Brendan McMahan et al. “Communication-efficient learning of deep networks from decentralized data”. In: *Artificial intelligence and statistics*. PMLR. 2017, pp. 1273–1282.

[18] Dan Russell, Letitia A Peplau, and Carolyn E Cutrona. “The revised UCLA Loneliness Scale: concurrent and discriminant validity evidence.” In: *Journal of personality and social psychology* 39.3 (1980), p. 472.

[19] Nitesh V Chawla et al. “SMOTE: synthetic minority over-sampling technique”. In: *Journal of artificial intelligence research* 16 (2002), pp. 321–357.

[20] J Vega et al. “Rapids: Reproducible analysis pipeline for data streams collected with mobile devices”. In: *J Med Internet Res Preprints. URL: https://preprints. jmir. org/preprint/23246 [accessed 2020-08-18]* (2020).

[21] *RAPIDS*. URL: https://www.rapids.science. (accessed 2021-02-21).

[22] Giuseppe Bonaccorso. *Machine learning algorithms*. Packt Publishing Ltd, 2017.