

Integrating Human-in-the-loop into Swarm Learning for Decentralized Fake News Detection

Xishuang Dong, and Lijun Qian, *Senior Member, IEEE*,

Abstract—Social media has become an effective platform to generate and spread fake news that can mislead people and even distort public opinion. Centralized methods for fake news detection, however, cannot effectively protect user privacy during the process of centralized data collection for training models. Moreover, it cannot fully involve user feedback in the loop of learning detection models for further enhancing fake news detection. To overcome these challenges, this paper proposed a novel decentralized method, Human-in-the-loop Based Swarm Learning (HBSL), to integrate user feedback into the loop of learning and inference for recognizing fake news without violating user privacy in a decentralized manner. It consists of distributed nodes that are able to independently learn and detect fake news on local data. Furthermore, detection models trained on these nodes can be enhanced through decentralized model merging. Experimental results demonstrate that the proposed method outperforms the state-of-the-art decentralized method in regard of detecting fake news on a benchmark dataset.

Index Terms—Fake News Detection, Swarm Learning, Human-in-the-loop (HITL), Social Media

I. INTRODUCTION

The development of social media (e.g., Twitter) has significantly changed the way of information collection for human. Social media has dominated information generation and spreading, which results in that it is becoming more and more challenging for people to live without social media [1], [2], [3]. Although it reduces the cost of information retrieval, unfortunately, the absence of systematic and effective management on the information on social media platforms has led to that social media comes up with the hotbed of generation and spreading of fake news [4], [5], where fake news refers to the news that is intentionally and verifiably false [6], [7]. As a result, fake news causes many confusions and severe damages to the society. For example, during the COVID-19 pandemic, fake news makes it more difficult for people to find trustworthy and reliable information to combat the spreading and treatment of the virus [8]. Thus, preventing the spreading of fake news is imperative to decrease political polarization, increase trust in public institutions, and improve decision-making for everyone's life.

Machine learning techniques are effective to recognize potential fake news by building models on news features including content [9], [10], [11] and context [12], [13], [14]. Centralized detection methods dominated this field by

collecting big data to a cloud storage for building high-performance detection models, where deep learning techniques such as convolutional neural networks (CNN), recurrent neural networks (RNN), and deep graph models outperform other techniques [11], [15], [16], [17], [18], [19]. Unfortunately, one potential risk in the procedure of building these methods is to violate user privacy when collecting big data of news in the centralized manner.

To reduce this risk, decentralized methods are required to implement privacy preserving, which is to learn and infer on local data, not upload data to a centralized data storage for building detection models. Federated learning [20], [21], [22] is a distributed machine learning approach that enables training a high-quality centralized model while training data remains distributed over a large number of users. However, it relies on a model center to control the process of updating models for users, which increases the potentials of hacking regarding the communication between the model center and user models. On the contrary, swarm learning [23] is able to implement decentralized learning to maintain user privacy without the need for a central coordinator, thereby going beyond federated learning. It enables decentralized training without sharing the data through learning a set of nodes, where each node learns on training data locally and enhances the model collaboratively without sharing the training data. They share parameters (weights) derived from training the model on the local data. Thus, it allows users at the nodes to maintain the confidentiality and privacy of the raw data. Nevertheless, swarm learning is not designed to leverage valuable user feedback on the inference to further enhance learning performance since the user feedback has been approved to effective to improve data analysis in the loop of learning and inference [24].

In this paper, we propose a human-in-the-loop based based swarm learning (HBSL) to integrate user feedback into the learning and inference of swarm learning via human-in-the-loop (HITL) techniques [24]. HITL is to involve human activities in the process of building machine learning models to improve the model performance via human knowledge [24]. We applied HITL to generate the user feedback in the learning process to improve fake news detection. The detailed learning process is shown in the diagram shown in Figure 1. It consists of three stages in the learning process, namely, (a) Local Learning, (b) Model Updating, and (c) Human Feedback, which forms a loop until the learning is terminated when meeting stop criteria. In stage (a), all nodes will learn the detection models independently on the local data. Then these models in all nodes will be updated in the stage (b) by a

X. Dong and L. Qian are with the Center of Excellence in Research and Education for Big Military Data Intelligence (CREDIT Center), Department of Electrical and Computer Engineering, Prairie View A&M University, Texas A&M University System, Prairie View, TX 77446, USA. Email: xidong@pvamu.edu, liqian@pvamu.edu

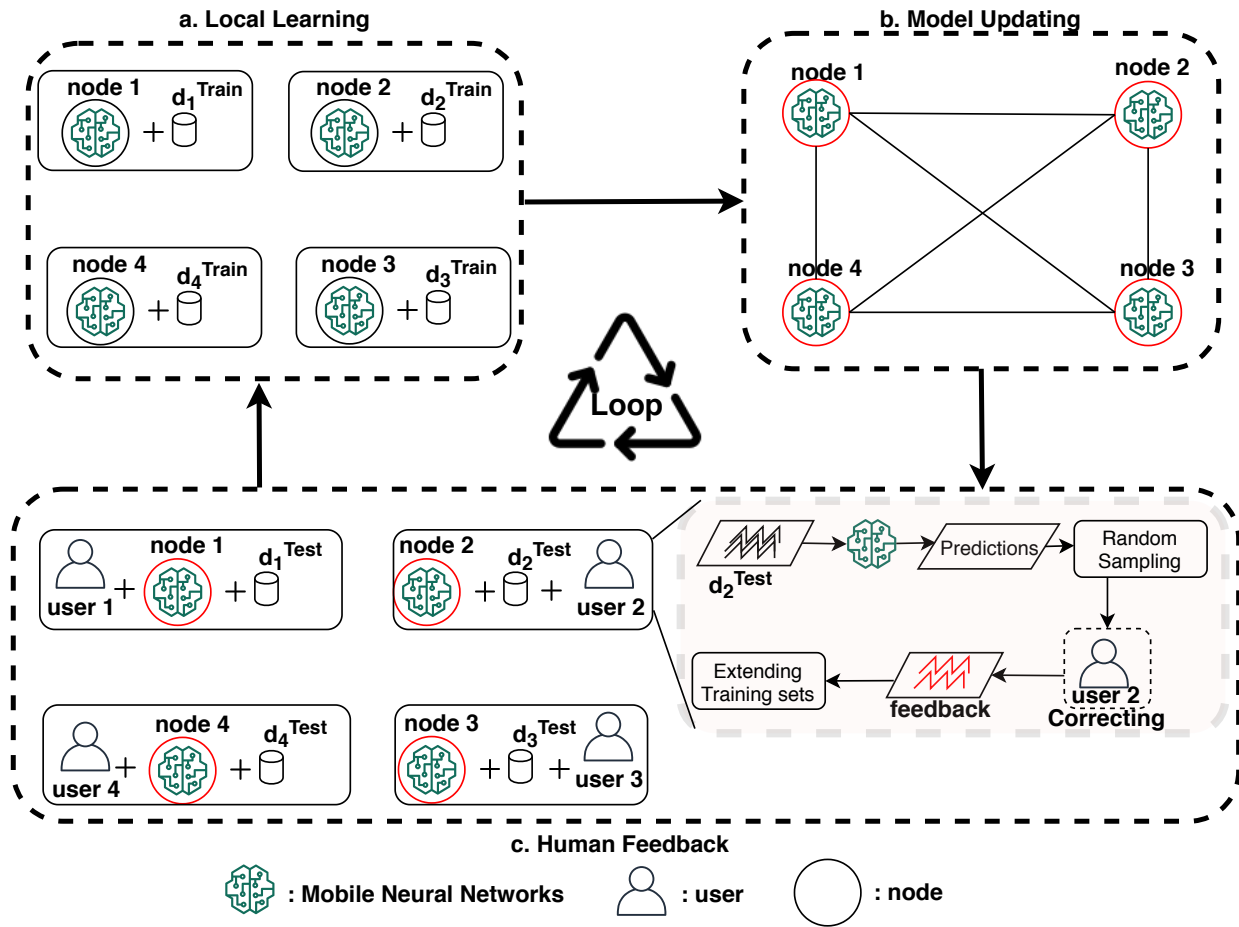


Fig. 1. Diagram of human-in-the-loop based swarm learning (HBSL) for fake news detection. It is composed of three stages, namely, (a) Local Learning, (b) Model Updating, and (c) Human Feedback, where the training set d_i^{Train} for node i will be different from d_j^{Train} for node j , and the testing sets d_i^{Test} and d_j^{Test} for nodes i and j , respectively, are different as well, $1 \leq i \leq 4$, and $1 \leq j \leq 4$. In stage (a), all nodes will learn the detection models on the local data. Then, in stage (b), these models will be updated by a master node selected through averaging model weights. Finally, in stage (c), all nodes will apply the model updated to accomplish fake news detection on the local testing data. Afterwards, users will correct a portion of predictions selected by random sampling as feedback to extend the training data. The models updated will be enhanced by fine-tuning on the training sets extended. All these three stages will form a loop of learning and inference to update their models until meeting the stop criteria.

master node through averaging model weights. In stage (c), all nodes will apply the model updated to accomplish fake news detection on the local testing data. Afterwards, users will provide feedback on the predictions of testing data, which is used to extend the training data to improve the detection. All models on these nodes will keep updating their models in a loop with these three stages until meeting the stop criteria of learning. Experimental results demonstrate that the proposed method outperformed swarm learning on decentralized fake news detection.

In summary, the contributions of this study are:

- We proposed a novel decentralized method through combing swarm learning and human-in-the-loop. In the learning and inference process, users in different nodes provides feedbacks on the inference results to generate feedback. Then nodes collect the feedback to extend training sets to fine-tune the model in order to enhance the inference performance within a loop of learning and inference.
- We validate our proposed model on a benchmark

LIAR [25]. Compared to swarm learning, the proposed model is able to significantly improve the performance for each node on detecting fake news by learning on local training data together with user feedbacks.

II. MODEL

The proposed decentralize method is to combine the advantages of swarm learning [23] and human-in the loop (HITL) [24], which is able to accomplish decentralized fake news detection via human feedback and model update in the swarm learning.

A. Swarm Learning

Swarm learning [23] is a decentralized machine-learning approach that unites edge computing, blockchain-based peer-to-peer networking and coordination while maintaining privacy without a central control. It implements learning on distributed nodes (edges) without sharing data to protect privacy in a local community. Compared to federated learning [20], [26], swarm

learning will be totally decentralized without parameter central control, which is illustrated as Figure 2.




Fig. 2. Comparison between swarm learning and federated learning. d_i is different from d_j , where $1 \leq i \leq 4$, and $1 \leq j \leq 4$. Federated learning keeps data with local data contributors and perform learning at the site of local data storage and availability, but control parameter updating by a central parameter server. On the contrary, swarm learning doesn't need parameter central control for parameter update since the data and parameters of a model will be stored locally.

In addition, swarm learning can update parameters without complex process to enhance the inference performance on nodes, only through merging parameters of models between nodes given by

$$P_M = \frac{\sum_{k=1}^n (w_k \times P_k)}{n \times \sum_{k=1}^n w_k}. \quad (1)$$

where P_M is the model parameter updated on a node, P_k is parameters from the k^{th} node, w_k is the weight of the k^{th} node, and n is the number of nodes participating in the merging process.

B. Human-in-the-Loop (HITL)

Human-in-the-loop can be applied to improve the performance of machine learning models by integrating human knowledge and experience for data analytics [24]. For example, human can significantly reduce algorithm bias in the training and inference in terms of human feedback for various tasks in the field of natural language processing (NLP) such as text classification [27], syntactic and semantic parsing [28], topic modeling [29], text summarization [30], and sentiment analysis [31]. The general framework is shown in Figure 3.




Fig. 3. Framework of human-in-the-loop for machine learning.

Human can provide feedback to the model training, data preprocessing, data collection regarding the model inference, even predictions to improve the model inference in a loop, where the feedback can be associate with inference results and its performance, inference time cost, and inference computation cost.

C. HITL Based Swarm Learning (HBSL)

This paper proposed a model to combine swarm learning and HITL to implement decentralized fake news detection.

Figure 1 presents the flow of building the proposed model with four nodes in three stages in the learning process, namely, (a) Local Learning, (b) Model Updating, and (c) Human Feedback, which forms a loop until the learning is terminated.

Local Learning: This stage is to independently learn the model on local data in the node. The local data is pre-processed by removing missing values, stemming, and one-hot representations of sentences. These representations are input to the model built via bidirectional recurrent neural networks (BRNN). This model will not be complex regarding the resource constrain like limited memory and thus contains one embedding layer, one forward layer, and one backward layers. The embedding layer represents the input $\langle x_1, x_2, x_3, \dots, x_t, \dots, x_n \rangle$ as an embedding vector $\langle e_1, e_2, e_3, \dots, e_t, \dots, e_n \rangle$, which is to resolve feature sparsity issues. Afterwards, these two layers generate two directional correlation features on the embedding vector. Then we combine these two features as the output z of this model, where z is a sequence $\langle z_1, z_2, z_3, \dots, z_t, \dots, z_n \rangle$ and z_t is given by

$$z_t = w_z^f h_t^f + w_z^b h_t^b + b_z, \quad (2)$$

where

$$h_t^f = a(w_h^f h_{t-1}^f + w_e^f e_t + b_h^f), \quad (3)$$

$$h_t^b = a(w_h^b h_{t+1}^b + w_e^b e_t + b_h^b), \quad (4)$$

z_t is the output of x_t for the input $\langle x_1, x_2, x_3, \dots, x_t, \dots, x_n \rangle$. $a(\cdot)$ is the activation function for hidden layers. w_z^f , w_h^f , and w_e^f are forward weights for three layers, namely, output layer, forward layer, and backward layer. w_z^b , w_h^b , and w_e^b are backward weights for these three layers, respectively. b_z , b_h^f , and b_h^b are bias for these three layers.

We utilize the output z to calculate the cross entropy loss given by

$$Loss = -y \times \log \phi(z) - (1 - y) \times \log(1 - \phi(z)), \quad (5)$$

where $\phi(z)$ is the sigmod function. These four nodes will learn these models with this identical architecture of BRNN on local data.

Model Updating: The master node in this stage will be selected in terms of predefined criterion. The criterion is that the node with the highest performance of detection is selected as the master node, which can be implemented by local communication between nodes. This node collects all models from other nodes and updates parameters w_{master} below.

$$w'_{master} = \frac{\sum_{i=1}^n w_i}{n}. \quad (6)$$

where $n = 4$ in this proposed method. Afterwards, the master node will share the updated model to replace parameters of models in other nodes.

Human Feedback: In the third stage, the model will perform inference on their local testing data in a parallel manner, where the dataset for this inference for different nodes will be various. For example, for the diagram in Figure 1, the dataset d_i^{Test} for node i will be different from d_j^{Test} for node

j , where $i \neq j$, $1 \leq i \leq 4$, and $1 \leq j \leq 4$. Then, each user i (Human) will generate feedback on the prediction results. In detailed, the user will randomly select a portion of predictions and correct them as feedback. This feedback will be integrated into corresponding training data to extend training sets. For instance, feedback on the predictions of $d_2^{T_{est}}$ will generate a set of corrected predictions d'_2 . Then d'_2 will be integrated into $d_2^{T_{rain}}$ to extend the training set for next round of local learning of local learning in the loop of learning.

Stages (a) to (c) form the loop to update node's model parameters to enhance detection performance. The details learning process is shown in Algorithm 1, where the portion of predictions is predefined.

Algorithm 1 Learning in the proposed method

- 1: **for** t in [1, rounds] **do**
 - 2: Learning on each node's data $d_{t,i}^{T_{rain}}$, $1 \leq i \leq 4$
 - 3: Updating the parameter $w_{t,i}$ of the model of the master node i with equation (6)
 - 4: Replacing the model parameter of other nodes with $w_{t,i}$
 - 5: Inferencing on each node's data $d_{t,i}^{T_{est}}$, $1 \leq i \leq 4$
 - 6: Generating feedback $d'_{t,i}$ by correcting a portion of predictions on each nodes, $1 \leq i \leq 4$
 - 7: Extending $d_{t,i}^{T_{rain}} = \{d_{t,i}^{T_{rain}}, d'_{t,i}\}$, $1 \leq i \leq 4$
- return** w_i for node i , $1 \leq i \leq 4$
-

III. EXPERIMENT

A. Datasets

We validate the effectiveness of the proposed model by detecting fake news on the benchmark LIAR [25]. It is a standard benchmark dataset for fake news detection. It includes 12,836 real-world short statements collected from a variety of occasions such as debate, campaign, Facebook, Twitter, interviews, ads, etc. Each statement is labeled with six-grade truthfulness, namely, true, false, half-true, part-fire, barely-true, and mostly-true. We reorganize the data as two classes by treating five classes including false, half-true, part-fire, barely-true, and mostly-true as **Fake** class and true as **True** class. Therefore, the fake news detection on this benchmark is converted to a binary classification task.

B. Experiment Setup

The key hyper-parameters for training the proposed model are shown in Table I and we employ Adam optimizer to complete the training.

TABLE I
HYPER-PARAMETERS FOR THE LEARNING OF THE MODEL ON NODES

Hyper-parameters	Values
Embedding size	40
Batch size	32
Number of epochs	10
Learning rate	1e-3
Feedback portion	0.2

C. Evaluation

We apply accuracy to evaluate the performance of fake news detection regarding the task features on the benchmark, where the accuracy is calculated by dividing the number of news detected correctly over the total number of news.

$$Accuracy = \frac{N_{correct}}{N_{total}}. \quad (7)$$

D. Results

We validate the proposed methods with two different configurations. One is to implement decentralized learning on four nodes with human feedback while the other is to implement learning on eight nodes.

1) *Learning on four nodes*: In that regard of limited data available for individual user in the real application, we distribute small number of samples (e.g. less than 1,000) for inference in different nodes, which is shown in Figure 4 with two classes: Fake and True. Specifically, we assume that all nodes share the same class distribution to simplify the task.

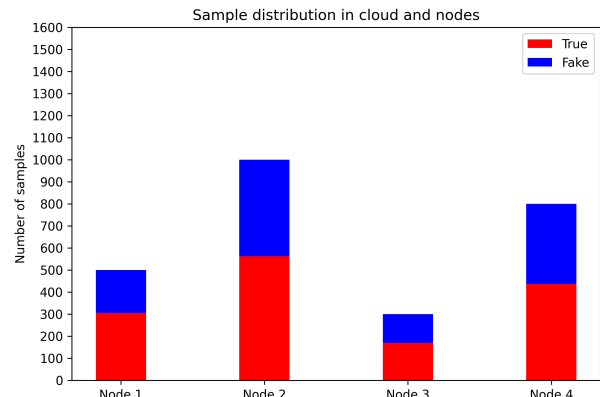


Fig. 4. Sample distribution on 4 nodes for fake news detection.

TABLE II
ACCURACY (%) COMPARISON ON FAKE NEWS DETECTION FOR LEARNING WITH 4 NODES.

SL	Node 1	Node 2	Node 3	Node 4	Average
Run 1	86.40%	86.40%	90.66%	91.50%	88.74%
Run 2	89.60%	89.60%	93.33%	85.50%	89.50%
Run 3	87.99%	89.60%	86.66%	86.00%	87.56%
Run 4	82.40%	85.60%	85.33%	85.50%	84.70%
Run 5	86.40%	89.60%	93.33%	85.50%	88.70%
Average	86.55%	88.16%	89.86%	86.80%	87.84%
HBSL	Node 1	Node 2	Node 3	Node 4	Average
Run 1	92.00%	89.99%	92.00%	94.99%	92.24%
Run 2	90.39%	92.79%	97.33%	90.49%	92.75%
Run 3	92.79%	94.40%	87.99%	91.00%	91.55%
Run 4	85.60%	89.20%	87.99%	89.99%	88.19%
Run 5	90.39%	89.20%	94.66%	91.00%	91.31%
Average	90.23%	91.11%	91.99%	91.49%	91.20%

We show the details of performance of 5 runs in Table II from two dimensions. One dimension is to examine the performance in terms of the accuracy for different runs. It appears that HBSL consistently performs better than SL in different runs. The other dimension is to check the performance for

different nodes. HBSL can improve detection performance for different nodes up to 5%. From these two dimensions, human feedback for swarm learning can consistently improve detection performance for 5 runs of 4 nodes. However, it can be observed that the performance of these nodes for different runs are various up to 7% and 10% for SL and HBSL, respectively. The unstable performance is caused by learning on small local data with the shallow RNNs on these nodes.

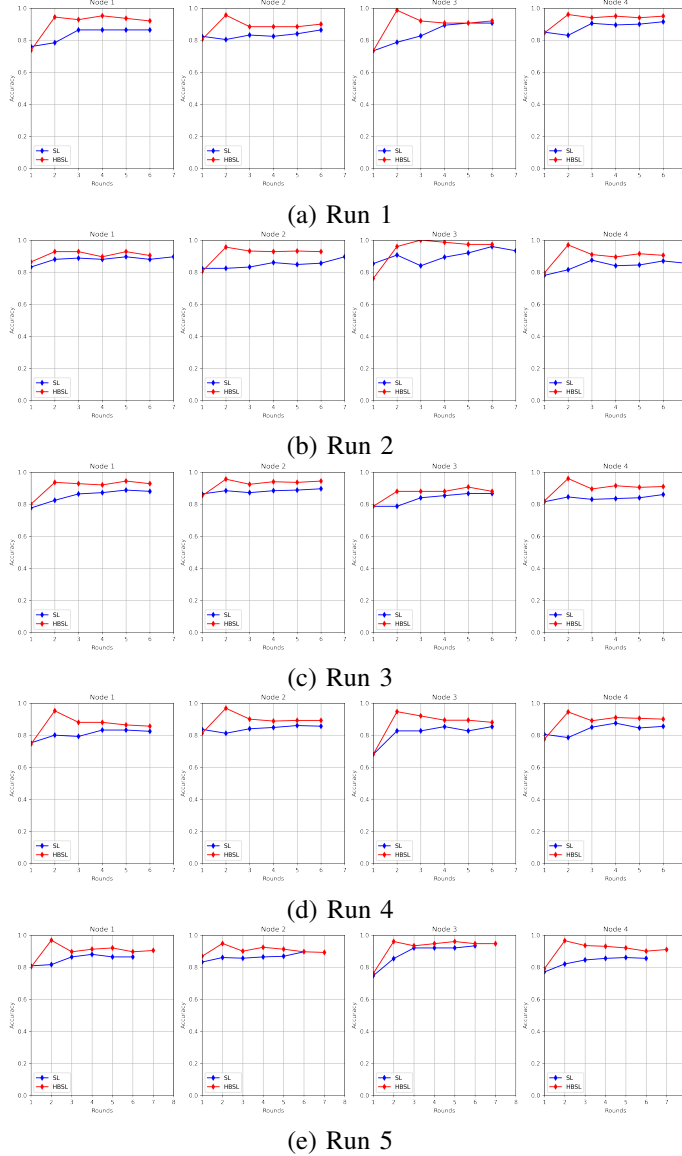


Fig. 5. Inference performance for 4 nodes in the learning process for HBSL in 5 runs. Accuracy refers to testing accuracy.

Moreover, Figure 5 presents the inference performance for different nodes in the learning process. It is observed that with human feedbacks, HBSL detection performance is higher than those of SL, which means that human feedback can contribute to improvement of detection performance. In addition, compared to SL, HBSL can converge to higher performance for most of nodes in different runs. Moreover, although human feedback enlarges the training sets by introducing more samples, it will not decrease learning speed by

increasing learning rounds.

2) *Learning on eight nodes:* We also examine if HBSL can perform better than SL with the condition of learning on 8 nodes. Similarly, we distribute small number of samples for learning in different nodes as well, which is shown in Figure 6. In addition to more data sets involved in this experiment, the size differences of samples for these data sets are larger compared to the case of learning on 4 nodes. For instance, the size difference between node 6 and node 3 is 1,700 (2,000 – 300) while that for the case of 4 node is 700 (1,000 – 300) between node 2 and node 3. Moreover, for node 8, the dataset is not balanced. These changes will introduce new challenges like increasing converge time for fake news detection.




Fig. 6. Sample distribution on 8 nodes for fake news detection.

Table III presents the comparison of the testing accuracy for the case with 8 nodes. Similar observation is obtained that HBSL outperforms SL on all nodes. However, the improvement is up to about 2%, which is less than 5% for the case of learning on 4 nodes. In other words, it seems that HBSL will be more suitable to learn on a local net with less nodes. Moreover, it enhances the detection performance for all runs in terms of average accuracy for different runs. Figure 7 shows the testing accuracy for 8 nodes in the learning process for 5 runs. It is observed that the curves of test accuracy is various since the sizes of data for fine-tuning models will be various in the loop of learning and inference. In addition, although some nodes performed worse in the learning of HBSL, they can still converge to higher or similar performance of SL. Finally, it uses more learning rounds to achieve model convergence as it involves more nodes to update models, which reduces the convergence speed.

IV. RELATED WORK

Fake news detection attracted lots of attentions to effectively preventing the dissemination of fake news. Traditional fake news detection based on machine learning is classified into three categories, namely, content feature based [32], [11], [33], [34], propagation feature based [14], and context feature based [35]. Recently, more dimensions have been exploited. The first dimension is to combine these features to enhance detection performance. Moti *et al.* represented news as a graph

TABLE III
ACCURACY (%) COMPARISON ON FAKE NEWS DETECTION FOR LEARNING WITH 8 NODES.

SL	Node 1	Node 2	Node 3	Node 4	Node 5	Node 6	Node 7	Node 8	Average
Run 1	87.99%	84.39%	93.33%	85.50%	89.14%	89.80%	88.66%	80.00%	87.35%
Run 2	87.19%	89.60%	85.33%	87.99%	89.14%	87.80%	88.66%	95.99%	88.96%
Run 3	81.59%	88.40%	90.66%	89.99%	87.42%	87.99%	93.33%	92.00%	88.92%
Run 4	88.80%	88.80%	85.33%	88.99%	89.14%	87.59%	93.99%	80.00%	87.83%
Run 5	89.60%	90.79%	90.66%	89.99%	86.28%	87.59%	88.66%	83.99%	88.45%
Average	87.03%	88.39%	89.06%	88.49%	88.22%	88.15%	90.66%	86.39%	88.30%
HBSL	Node 1	Node 2	Node 3	Node 4	Node 5	Node 6	Node 7	Node 8	Average
Run 1	89.60%	90.79%	90.66%	89.99%	86.28%	87.59%	88.66%	83.99%	88.45%
Run 2	89.60%	90.39%	87.99%	88.99%	89.14%	90.20%	89.99%	95.99%	90.29%
Run 3	85.60%	88.80%	92.00%	91.50%	91.42%	87.59%	95.99%	95.99%	91.11%
Run 4	91.29%	88.80%	83.99%	89.99%	89.71%	90.60%	93.33%	80.00%	88.45%
Run 5	93.59%	91.60%	92.00%	91.50%	90.28%	89.60%	89.33%	87.99%	90.73%
Average	89.93%	90.07%	89.32%	90.39%	89.36%	89.11%	91.46%	88.79%	89.80%

containing these three features and utilized graph convolutional networks (GCN) to classify the graph into fake or true [18]. In addition, it is resilience to adversarial attacks since generating adversarial samples on three features is extremely challenging. Wang *et al.* used multimodal features to detect fake news by developing event adversarial neural networks. It consists of three components: a multimodal feature extractor, a fake news detector, and an event classifier. The feature extractor recognized multimodal features on text and pictures related to news as input to the fake news detector and the event classifier. Then the fake news detector will beat the event classifier in the adversarial learning process, which is extract the multimodal features, that are more useful to fake news detection, through the learning process.

The second dimension is to fight against adversarial “fake news” generated by AI techniques. For instance, Fung *et al.* proposed *InfoSurgeon* to detect adversarial “fake news” that is represented as a fine-grained knowledge graph of news with external knowledge bases [36]. It utilized graph neural networks to determine if the knowledge graph of news is true or fake. Shu *et al.* developed *FactGen* to generate high-quality news by leveraging external facts to enrich the news and improve the consistency between input and output. Moreover, they proposed *FactGen_{def}* to detect these synthetic fake news with high performance [37].

The third dimension is to reduce efforts of labeling big data for training detection models. For example, Wang *et al.* proposed WeFEND, a reinforcement learning method, that is to leverage user’s reports as weak supervision to enlarge training sets for fake news detection [38]. It is composed of three components including annotator, reinforce selector, and detector. The annotator labeled data with weak labels and the reinforce selector chooses weakly labeled data to extend training sets for fine-tuning the detector.

The fourth dimension is to detect fake news from the perspective of psychology. For example, Karami *et al.* exploited fake news detection through combining psychology and data science. They extracted five features about motivation of spreading fake news that present user’s psychology, namely, uncertainty, emotions, lack of control, relationship enhancement, and rank, where the first three features are extracted with Linguistic Inquiry and Word Count (LIWC) [39], and the

other two are obtained based on user behaviors on social media like retweeting and following. These features are significantly different between users who spread fake news and users who spread true news. Moreover, these features can be combined with content features to improve detection performance [40].

Although these dimensions have broadened and enhanced fake news detection, few of work has exploited how to introduce human feedback into decentralized fake news detection to implement privacy preserving for users. This paper combined human-in-the-loop techniques with swarm learning to exploit decentralized fake news detection.

V. CONCLUSION

In this paper, a novel decentralized model is proposed for detecting fake news through combining swarm learning and human-in-the-loop. The learning and inference forms a loop until meeting stop criteria by learning and predicting on local data. Specifically, the human feedback is generated in this loop to extend training sets for improving detection performance. The proposed model is validated on a benchmark, LIAR dataset. Experimental results indicate that the proposed model could outperform swarm learning on fake news detection in a decentralized manner. In the future, we plan to extend this work by designing detection models according to node features.

REFERENCES

- [1] M. Imran, C. Castillo, F. Diaz, and S. Vieweg, “Processing social media messages in mass emergency: A survey,” *ACM Computing Surveys (CSUR)*, vol. 47, no. 4, pp. 1–38, 2015.
- [2] G. Pennycook and D. G. Rand, “Fighting misinformation on social media using crowdsourced judgments of news source quality,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 7, pp. 2521–2526, 2019.
- [3] S. Phuvipadawat and T. Murata, “Breaking news detection and tracking in twitter,” in *2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, vol. 3. IEEE, 2010, pp. 120–123.
- [4] H. Webb, P. Burnap, R. Procter, O. Rana, B. C. Stahl, M. Williams, W. Housley, A. Edwards, and M. Jirotko, “Digital wildfires: Propagation, verification, regulation, and responsible innovation,” *ACM Transactions on Information Systems (TOIS)*, vol. 34, no. 3, pp. 1–23, 2016.
- [5] L. Wu, F. Morstatter, K. M. Carley, and H. Liu, “Misinformation in social media: definition, manipulation, and detection,” *ACM SIGKDD Explorations Newsletter*, vol. 21, no. 2, pp. 80–90, 2019.

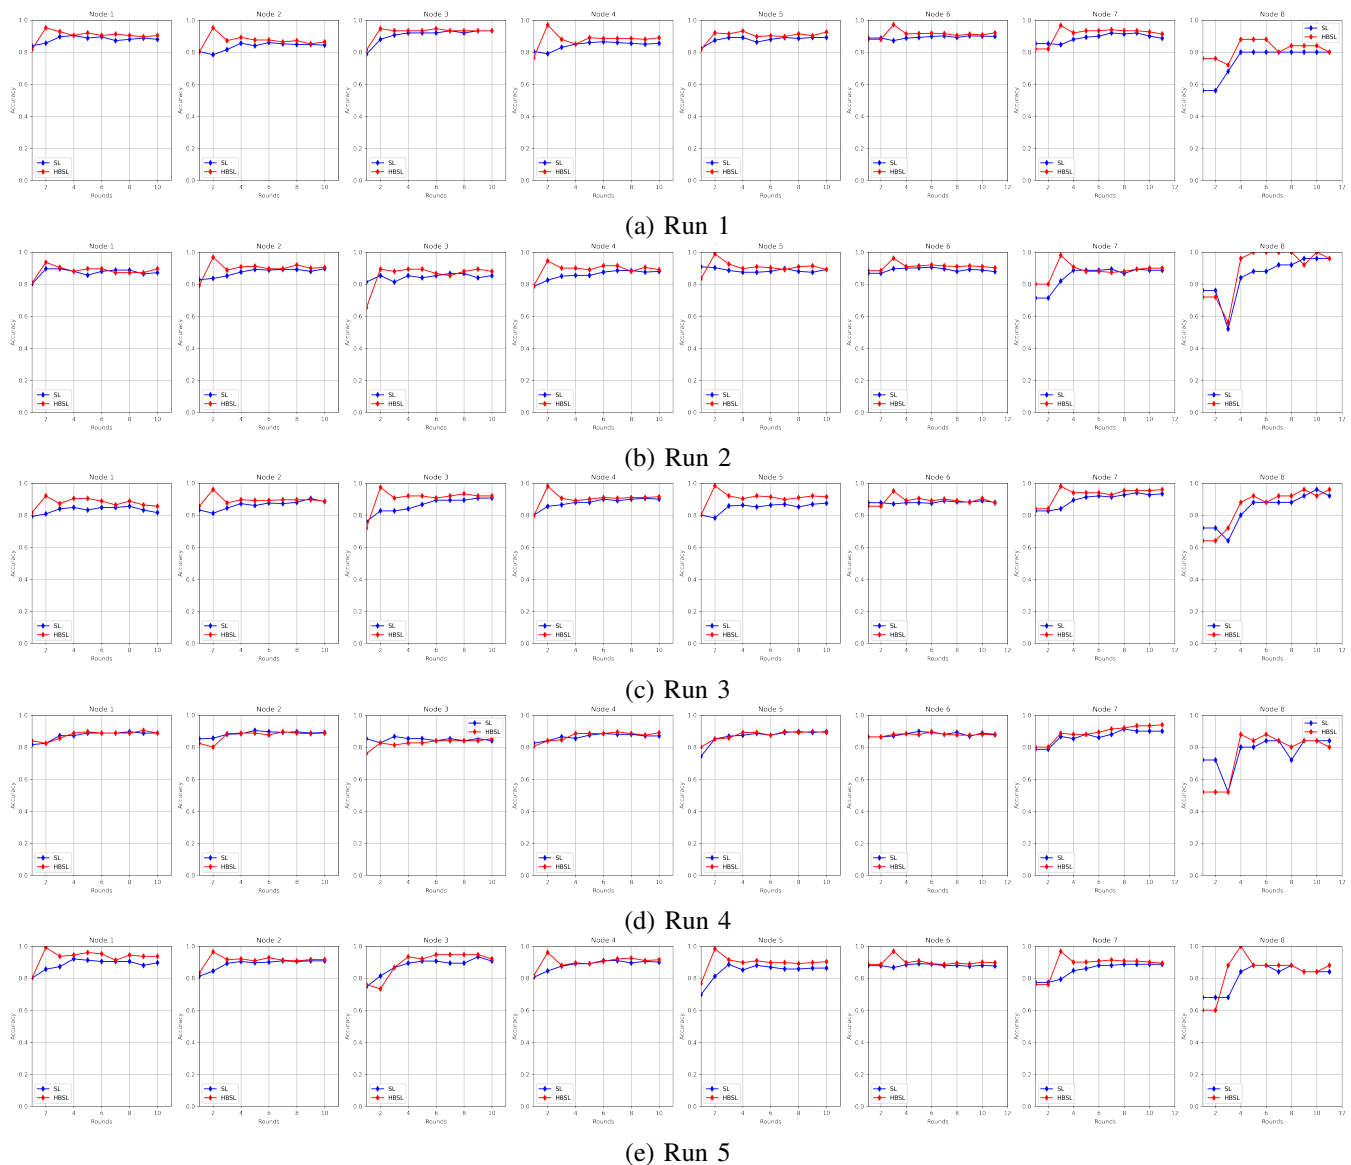


Fig. 7. Inference performance for 8 nodes in the learning process for HBSL in 5 runs. Accuracy refers to testing accuracy.

- [6] H. Allcott and M. Gentzkow, “Social media and fake news in the 2016 election,” *Journal of economic perspectives*, vol. 31, no. 2, pp. 211–36, 2017.
- [7] X. Zhou and R. Zafarani, “A survey of fake news: Fundamental theories, detection methods, and opportunities,” *ACM Computing Surveys (CSUR)*, vol. 53, no. 5, pp. 1–40, 2020.
- [8] J. Geleris, Y. Sun, J. Platt, J. Zucker, M. Baldwin, G. Hripcsak, A. Labella, D. K. Manson, C. Kubin, R. G. Barr *et al.*, “Observational study of hydroxychloroquine in hospitalized patients with covid-19,” *New England Journal of Medicine*, vol. 382, no. 25, pp. 2411–2418, 2020.
- [9] N. K. Conroy, V. L. Rubin, and Y. Chen, “Automatic deception detection: Methods for finding fake news,” *Proceedings of the association for information science and technology*, vol. 52, no. 1, pp. 1–4, 2015.
- [10] V. L. Rubin, N. Conroy, Y. Chen, and S. Cornwell, “Fake news or truth? using satirical cues to detect potentially misleading news,” in *Proceedings of the second workshop on computational approaches to deception detection*, 2016, pp. 7–17.
- [11] N. Ruchansky, S. Seo, and Y. Liu, “Csi: A hybrid deep model for fake news detection,” in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*. ACM, 2017, pp. 797–806.
- [12] E. Tacchini, G. Ballarin, M. L. Della Vedova, S. Moret, and L. de Alfaro, “Some like it hoax: Automated fake news detection in social networks,” *arXiv preprint arXiv:1704.07506*, 2017.
- [13] Y. Papanastasiou, “Fake news propagation and detection: A sequential model,” *Management Science*, vol. 66, no. 5, pp. 1826–1846, 2020.
- [14] K. Shu, D. Mahudeswaran, S. Wang, and H. Liu, “Hierarchical propagation networks for fake news detection: Investigation and exploitation,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 14, 2020, pp. 626–637.
- [15] O. Ajao, D. Bhowmik, and S. Zargari, “Fake news identification on twitter with hybrid cnn and rnn models,” in *Proceedings of the 9th international conference on social media and society*, 2018, pp. 226–230.
- [16] C. Song, K. Shu, and B. Wu, “Temporally evolving graph neural network for fake news detection,” *Information Processing & Management*, vol. 58, no. 6, p. 102712, 2021.
- [17] Y. Ren, B. Wang, J. Zhang, and Y. Chang, “Adversarial active learning based heterogeneous graph neural network for fake news detection,” in *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2020, pp. 452–461.
- [18] F. Monti, F. Frasca, D. Eynard, D. Mannoni, and M. M. Bronstein, “Fake news detection on social media using geometric deep learning,” *arXiv preprint arXiv:1902.06673*, 2019.
- [19] A. Benamira, B. Devillers, E. Lesot, A. K. Ray, M. Saadi, and F. D. Malliaros, “Semi-supervised learning and graph neural networks for fake news detection,” in *2019 IEEE/ACM International Conference on*

- Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 2019, pp. 568–569.
- [20] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, “Federated optimization: Distributed machine learning for on-device intelligence,” *arXiv preprint arXiv:1610.02527*, 2016.
- [21] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, and H. Yu, “Federated learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 13, no. 3, pp. 1–207, 2019.
- [22] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [23] S. Warnat-Herresthal, H. Schultze, K. L. Shastri, S. Manamohan, S. Mukherjee, V. Garg, R. Sarveswara, K. Händler, P. Pickkers, N. A. Aziz *et al.*, “Swarm learning for decentralized and confidential clinical machine learning,” *Nature*, vol. 594, no. 7862, pp. 265–270, 2021.
- [24] X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, and L. He, “A survey of human-in-the-loop for machine learning,” *arXiv preprint arXiv:2108.00941*, 2021.
- [25] W. Y. Wang, ““liar, liar pants on fire”: A new benchmark dataset for fake news detection,” *arXiv preprint arXiv:1705.00648*, 2017.
- [26] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [27] T. Karmakharm, N. Aletras, and K. Bontcheva, “Journalist-in-the-loop: Continuous learning as a service for rumour analysis,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): System Demonstrations*, 2019, pp. 115–120.
- [28] Y. Su, A. H. Awadallah, M. Khabsa, P. Pantel, M. Gamon, and M. Encarnacion, “Building natural language interfaces to web apis,” in *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 2017, pp. 177–186.
- [29] Y. Hu, J. Boyd-Graber, B. Satinoff, and A. Smith, “Interactive topic modeling,” *Machine learning*, vol. 95, no. 3, pp. 423–469, 2014.
- [30] N. Stiennon, L. Ouyang, J. Wu, D. M. Ziegler, R. Lowe, C. Voss, A. Radford, D. Amodei, and P. Christiano, “Learning to summarize from human feedback,” *arXiv preprint arXiv:2009.01325*, 2020.
- [31] Z. Liu, Y. Guo, and J. Mahmud, “When and why does a model fail? a human-in-the-loop error detection framework for sentiment analysis,” *arXiv preprint arXiv:2106.00954*, 2021.
- [32] Y. Long, Q. Lu, R. Xiang, M. Li, and C.-R. Huang, “Fake news detection through multi-perspective speaker profiles,” in *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, vol. 2, 2017, pp. 252–256.
- [33] R. Oshikawa, J. Qian, and W. Y. Wang, “A survey on natural language processing for fake news detection,” *arXiv preprint arXiv:1811.00770*, 2018.
- [34] K.-C. Yang, T. Niven, and H.-Y. Kao, “Fake news detection as natural language inference,” in *12th ACM International Conference on Web Search and Data Mining (WSDM-2019)(in Fake News Classification Challenge, WSDM Cup 2019)*, 2019.
- [35] K. Shu, S. Wang, and H. Liu, “Beyond news contents: The role of social context for fake news detection,” in *Proceedings of the twelfth ACM international conference on web search and data mining*, 2019, pp. 312–320.
- [36] Y. Fung, C. Thomas, R. G. Reddy, S. Polisetty, H. Ji, S.-F. Chang, K. McKeown, M. Bansal, and A. Sil, “Infosurgeon: Cross-media fine-grained information consistency checking for fake news detection,” in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 1683–1698.
- [37] K. Shu, Y. Li, K. Ding, and H. Liu, “Fact-enhanced synthetic news generation,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 15, 2021, pp. 13 825–13 833.
- [38] Y. Wang, W. Yang, F. Ma, J. Xu, B. Zhong, Q. Deng, and J. Gao, “Weak supervision for fake news detection via reinforcement learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 01, 2020, pp. 516–523.
- [39] Y. R. Tausczik and J. W. Pennebaker, “The psychological meaning of words: Liwc and computerized text analysis methods,” *Journal of language and social psychology*, vol. 29, no. 1, pp. 24–54, 2010.
- [40] M. Karami, T. H. Nazer, and H. Liu, “Profiling fake news spreaders on social media through psychological and motivational factors,” in *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, 2021, pp. 225–230.