Community-Structured Decentralized Learning for Resilient EI

Harikrishna Kuttivelil
hkuttive@ucsc.edu
University of California, Santa Cruz

Katia Obraczka
katia@ucsc.edu
University of California, Santa Cruz

ACM Reference Format:
Harikrishna Kuttivelil and Katia Obraczka. 2021. Community-Structured Decentralized Learning for Resilient EI. In *First Workshop on Systems Challenges in Reliable and Secure Federated Learning (ResilientFL '21), October 25, 2021, Virtual Event, Germany.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3477114.3488764

1 INTRODUCTION

The symbiotic relationship at the intersection of edge computing [26] and artificial intelligence that powers many edge applications has resulted in the emergence of a new field, *edge intelligence* (EI) which has been gaining prominence in fields such as smart cities [28, 35], autonomous vehicle fleets [34], intelligent manufacturing and agriculture [37], and internet-of-things [16, 38]. EI has been classified [39] based on the levels of device, edge, and cloud involvement, thus also introducing the notion of an EI continuum ranging from cloud-edge cooperative intelligence down to on-device intelligence and includes EI that is conducted solely among edge devices, or in-edge collaborative intelligence. In this paper, we propose *community-structured decentralized edge learning*, a fully decentralized approach, in which edge devices self-organize into *learning communities* based on data and feature affinities and facilitate ML training and inference within those communities. While drawing on work from multi-agent systems and existing distributed learning paradigms [2, 10–12, 14, 15, 18], this new paradigm would fill an important gap in the EI continuum by considering learning within and across communities. We believe that the proposed paradigm could improve learning accuracy, especially where personalized or localized differentiation is desired, and communication efficiency while being resilient to cloud outages or communication disturbances; and that the challenges poised by the paradigm can be overcome by the continued development of enabling technologies.

Figure 1 shows different paradigms of evolution towards decentralized EI, including our proposed paradigm. The current state of the network edge in its diverse and varied ecosystem of personal and connected devices is characterized by limited resources (e.g. power, storage, consumption), heterogeneous communication infrastructure, and privacy and security considerations [24, 27, 30]. Federated learning [19] has emerged as an eminent paradigm of learning at the edge, and while it mostly takes place at the edge, it still relies on the cloud for coordination, aggregation, and dissemination. To circumvent dependence on centralization, many fully decentralized approaches, improvements, and implementations have been studied [3, 11, 13, 15, 17, 23]. Yet the current state of fully decentralized learning paradigms leaves two issues to be considered. First, while communication efficiency through selective collaboration has been introduced in centralized distributed learning schemes [12, 14] and gossip learning [13], it has not been fully accounted for in fully decentralized learning. Second, while methods for application-specific personalization have been introduced in other distributed learning schemes [14, 18], existing works on fully decentralized learning do not consider the potential of using data and feature affinity to enable localization of models for personalized applications.

ResilientFL '21, October 25, 2021, Virtual Event, Germany

ACM ISBN 978-1-4503-8708-8/21/10.
https://doi.org/10.1145/3477114.3488764

2 COMMUNITY-STRUCTURED DECENTRALIZED EDGE LEARNING

We propose *community-structured fully decentralized edge learning*, a new distributed learning paradigm where participating agents form *learning communities* – groups of agents connected by feature or data affinity. After agents train their models on local data, models are primarily shared within communities with the option of inter-community sharing as well. This process is illustrated in Figure 2 and explained in Section 2. Structures and selective collaboration in distributed learning are not new and have been implemented in centralized federated learning [12, 14, 18] and gossip training protocols [4, 21, 22, 27]. Our proposed notion of structure is a step further from the current state-of-the-art as it allows the dynamic formation of communities along application- as well as communication-specific feature affinities, allowing for better personalization, increased resiliency (in the face of cloud outages and communication disruptions) and communication efficiency. Below, we discuss the technologies needed to enable the proposed community-structured decentralized learning paradigm.

Decentralized Community Identification. Methods to identify and self-organize into communities are necessary to then conduct model sharing within those communities. Community identification methods have been proposed for social networks [9, 25, 29], large attribute networks [31], and in applications to smart cities and urban planning [7]. Feature-based clustering strategies have also been proposed specifically for federated learning [8, 32]. However, many of these feature-based community identifications are at least partly centralized, and so work must be done to fully decentralize the process. To this end, there are methods in decentralized cluster formation in ad-hoc networks [20, 33] and clustering in connected graphs [1]. These strategies can be adapted to organize communities based on communication of data and feature affinities. For example, phones running a next word prediction can coalesce into communities based on similar geographical and demographic features to deliver more relevant predictions.

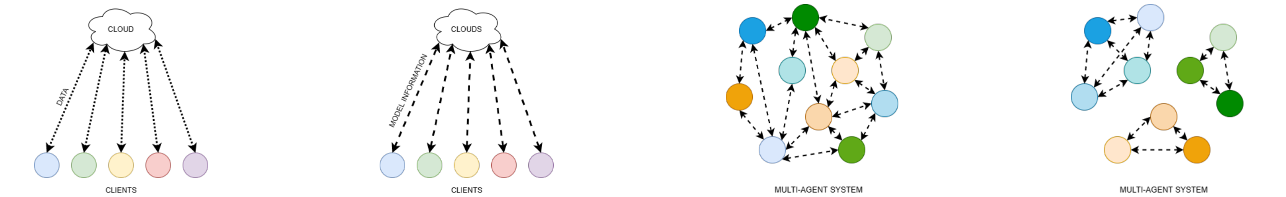

(a) Traditional learning. **(b) Federated learning.** **(c) Fully decentralized learning.** **(d) Community-structured.**

Figure 1: Evolution towards decentralized edge intelligence, including our proposed paradigm.

Decentralized Learning Algorithms. Decentralized learning algorithms are necessary for decentralized model training, local aggregation, model sharing, and local or collaborative inference. Reinventing these algorithms may not be necessary and so far, several decentralized learning paradigms [2, 4, 10, 11, 15] have been studied and improved [11, 12, 27]. The selection of these algorithms per application is a well-acknowledged problem generally in EI [6, 39], and, in pursuing community-structured decentralized learning specifically, needs to consider the more consistent exposure to fewer data sources when compared to general decentralized learning. Furthermore, work must also be done in adapting existing algorithms to the specific needs of edge applications, especially to facilitate personalization and localization. For example, many works on consensus strategies in decentralized paradigms [13, 17] are aimed at singular global convergence among the entire network, which is useful in some applications but not all. Many applications favor differentiated and more localized outcomes, such as predictive text in the NLP domain [14], and thus decentralized learning algorithms should be adapted to accommodate differentiated behavior among communities to produce useful differentiation of models.

Basic Building Blocks. While many EI frameworks exist [5, 11, 36], a framework for community-structured decentralized edge learning would need to include mechanisms for community identification and community-structured collaborative learning as discussed above. Figure 2 illustrates our vision of community-structured edge intelligence and its basic building blocks. During the initialization phase, agents securely exchange feature or data information with one another and measure affinity based on the shared information. Next, using specified similarity metrics and criteria, agents self-organize into communities according to their measured affinity. Similarly to clustering algorithms, the initialization and organization phases may have to iterate before reaching adequate *communitization*. Finally, agents train locally and exchange model information within their selected communities, with the option of sharing between communities as well. Community organization will trigger periodically due to changes in topology, network and device conditions, community membership, etc., in order facilitate dynamic and adaptive community structures in which learning occurs.

3 PRELIMINARY CASE STUDY

In our preliminary experiments, we used a custom Python-based decentralized federated learning simulator in which agents conduct decentralized learning tasks along a simple network graph. We

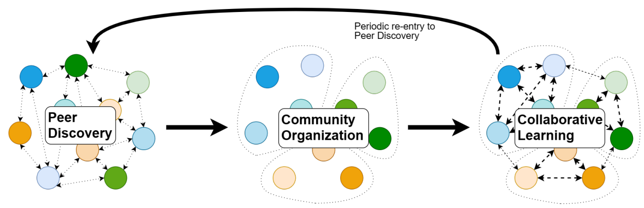

Figure 2: Community-structured decentralized edge intelligence and its basic building blocks.

used next-word prediction as the driving application where agents represent each of 12 Twitter users, organized into 3 topic-based communities, with each agent holding 14000 to 37000 local data points, split into local training and testing sets, and using LSTM networks on that data. Decentralized learning is conducted under three different modes – (1) a fully-connected network in which each agent shares its trained model information with every other agent, (2) using random partner selection in which each agent shares information with a random subset of 5 other agents, and (3) using community structure in which agents share information with other agents in their community. Each agent is evaluated on their own local test data. Figure 3 shows accuracy performance over 10 rounds, or cycles, of the different collaborative learning approaches, averaged over 5 runs. Community-structured collaboration performed 7% better in terms of accuracy while generating 73% fewer messages than fully-connected collaboration; and 9% better accuracy and 40% fewer messages than random partner collaboration. These preliminary results indicate that community-structured decentralized edge learning in applications that favor personalization and localization can provide higher accuracy and lower communication overhead.

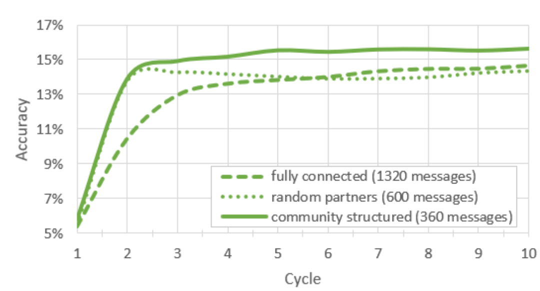

Figure 3: Preliminary results with next word prediction.

REFERENCES

[1] Amhmed Bhih, Princy Johnson, and Martin Randles. 2019. Decentralized iterative approaches for community clustering in the networks. *The Journal of Supercomputing* 75, 8 (2019), 4894–4917.

[2] Michael Blot, David Picard, Matthieu Cord, and Nicolas Thome. 2016. Gossip training for deep learning. *arXiv preprint arXiv:1611.09726* (2016).

[3] Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. 2018. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics* 112 (2018), 59–67.

[4] Jeff Daily, Abhinav Vishnu, Charles Siegel, Thomas Warfel, and Vinay Amatya. 2018. Gossipgrad: Scalable deep learning using gossip communication based asynchronous gradient descent. *arXiv preprint arXiv:1803.05880* (2018).

[5] Walter De Brouwer and Mason Borda. 2017. NeuRoN: decentralized artificial intelligence, distributing deep learning to the edge of the network.

[6] Shuiguang Deng, Hailiang Zhao, Weijia Fang, Jianwei Yin, Schahram Dustdar, and Albert Y Zomaya. 2020. Edge intelligence: the confluence of edge computing and artificial intelligence. *IEEE Internet of Things Journal* 7, 8 (2020), 7457–7469.

[7] Danielle L Ferreira, Bruno AA Nunes, Carlos Alberto V Campos, and Katia Obraczka. 2020. A deep learning approach for identifying user communities based on geographical preferences and its applications to urban and environmental planning. *ACM Transactions on Spatial Algorithms and Systems (TSAS)* 6, 3 (2020), 1–24.

[8] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. 2020. An efficient framework for clustered federated learning. *arXiv preprint arXiv:2006.04088* (2020).

[9] Tiantian He, Lun Hu, Keith CC Chan, and Pengwei Hu. 2018. Learning latent factors for community identification and summarization. *IEEE access* 6 (2018), 30137–30148.

[10] István Hegedűs, Gábor Danner, and Márk Jelasity. 2019. Gossip learning as a decentralized alternative to federated learning. In *IFIP International Conference on Distributed Applications and Interoperable Systems*. Springer, 74–90.

[11] Jiawen Kang, Zehui Xiong, Chunxiao Jiang, Yi Liu, Song Guo, Yang Zhang, Dusit Niyato, Cyril Leung, and Chunyan Miao. 2020. Scalable and Communication-Efficient Decentralized Federated Edge Learning with Multi-blockchain Framework. In *International Conference on Blockchain and Trustworthy Systems*. Springer, 152–165.

[12] Jakub Konečnỳ, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492* (2016).

[13] Lingjing Kong, Tao Lin, Anastasia Koloskova, Martin Jaggi, and Sebastian U Stich. 2021. Consensus control for decentralized deep learning. *arXiv preprint arXiv:2102.04828* (2021).

[14] Viraj Kulkarni, Milind Kulkarni, and Aniruddha Pant. 2020. Survey of personalization techniques for federated learning. In *2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*. IEEE, 794–797.

[15] Anusha Lalitha, Osman Cihan Kilinc, Tara Javidi, and Farinaz Koushanfar. 2019. Peer-to-peer federated learning on graphs. *arXiv preprint arXiv:1901.11173* (2019).

[16] En Li, Zhi Zhou, and Xu Chen. 2018. Edge intelligence: On-demand deep learning model co-inference with device-edge synergy. In *Proceedings of the 2018 Workshop on Mobile Edge Communications*. 31–36.

[17] Yuzheng Li, Chuan Chen, Nan Liu, Huawei Huang, Zibin Zheng, and Qiang Yan. 2020. A blockchain-based decentralized federated learning framework with committee consensus. *IEEE Network* (2020).

[18] Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. 2020. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619* (2020).

[19] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*. PMLR, 1273–1282.

[20] Elth Ogston, Benno Overeinder, Maarten Van Steen, and Frances Brazier. 2003. A method for decentralized clustering in large multi-agent systems. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. 789–796.

[21] Pitch Patarasuk and Xin Yuan. 2009. Bandwidth optimal all-reduce algorithms for clusters of workstations. *J. Parallel and Distrib. Comput.* 69, 2 (2009), 117–124.

[22] Amirhossein Reisizadeh, Saurav Prakash, Ramtin Pedarsani, and Amir Salman Avestimehr. 2019. Codedreduce: A fast and robust framework for gradient aggregation in distributed learning. *arXiv preprint arXiv:1902.01981* (2019).

[23] Abhijit Guha Roy, Shayan Siddiqui, Sebastian Pölsterl, Nassir Navab, and Christian Wachinger. 2019. Braintorrent: A peer-to-peer environment for decentralized federated learning. *arXiv preprint arXiv:1905.06731* (2019).

[24] Michael L Rustad and Thomas H Koenig. 2019. Towards a global data privacy standard. *Fla. L. Rev.* 71 (2019), 365.

[25] Daniel López Sánchez, Jorge Revuelta, Fernando De la Prieta, Ana B Gil-González, and Cach Dang. 2016. Twitter user clustering based on their preferences and the Louvain algorithm. In *International Conference on Practical Applications of Agents and Multi-Agent Systems*. Springer, 349–356.

[26] Mahadev Satyanarayanan. 2017. The emergence of edge computing. *Computer* 50, 1 (2017), 30–39.

[27] Yuanming Shi, Kai Yang, Tao Jiang, Jun Zhang, and Khaled B Letaief. 2020. Communication-efficient edge AI: Algorithms and systems. *IEEE Communications Surveys & Tutorials* 22, 4 (2020), 2167–2191.

[28] Bo Tang, Zhen Chen, Gerald Hefferman, Shuyi Pei, Tao Wei, Haibo He, and Qing Yang. 2017. Incorporating intelligence in fog computing for big data analysis in smart cities. *IEEE Transactions on Industrial informatics* 13, 5 (2017), 2140–2150.

[29] Chayant Tantipathananandh, Tanya Berger-Wolf, and David Kempe. 2007. A framework for community identification in dynamic social networks. In *Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*. 717–726.

[30] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10 (2017), 3152676.

[31] Xiao Wang, Di Jin, Xiaochun Cao, Liang Yang, and Weixiong Zhang. 2016. Semantic community identification in large attribute networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.

[32] Yuanli Wang, Joel Wolfrath, Nikhil Sreekumar, Dhruv Kumar, and Abhishek Chandra. 2021. Accelerated Training via Device Similarity in Federated Learning. In *Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking*. 31–36.

[33] Jane Yang Yu and Peter Han Joo Chong. 2005. A survey of clustering schemes for mobile ad hoc networks. *IEEE Communications Surveys & Tutorials* 7, 1 (2005), 32–48.

[34] Jun Zhang and Khaled B Letaief. 2019. Mobile edge intelligence and computing for the internet of vehicles. *Proc. IEEE* 108, 2 (2019), 246–261.

[35] Li Zhang, Jun Wu, Shahid Mumtaz, Jianhua Li, Haris Gacanin, and Joel JPC Rodrigues. 2019. Edge-to-edge cooperative artificial intelligence in smart cities with on-demand learning offloading. In *2019 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[36] Xingzhou Zhang, Yifan Wang, Sidi Lu, Liangkai Liu, Weisong Shi, et al. 2019. Openei: An open framework for edge intelligence. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 1840–1851.

[37] Yushu Zhang, Hui Huang, Lu-Xing Yang, Yong Xiang, and Ming Li. 2019. Serious challenges and potential solutions for the industrial Internet of Things with edge intelligence. *IEEE Network* 33, 5 (2019), 41–45.

[38] Yin Zhang, Xiao Ma, Jing Zhang, M Shamim Hossain, Ghulam Muhammad, and Syed Umar Amin. 2019. Edge intelligence in the cognitive internet of things: improving sensitivity and interactivity. *IEEE Network* 33, 3 (2019), 58–64.

[39] Zhi Zhou, Xu Chen, En Li, Liekang Zeng, Ke Luo, and Junshan Zhang. 2019. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proc. IEEE* 107, 8 (2019), 1738–1762.