

Asynchronous Decentralized Federated Lifelong Learning for Landmark Localization in Medical Imaging

Guangyao Zheng
Rice University
Houston, Texas, USA
tz30@rice.edu

Vladimir Braverman
Rice University
Houston, Texas, USA

Michael A. Jacobs
Department Of Diagnostic And Interventional Imaging,
McGovern Medical School, UTHHealth
Houston, Texas, USA

Vishwa S. Parekh
The University of Maryland Medical Intelligent Imaging
(UM2ii) Center
Baltimore, Maryland, USA

ABSTRACT

Federated learning is a recent development in the machine learning area that allows a system of devices to train on one or more tasks without sharing their data to a single location or device. However, this framework still requires a centralized server to consolidate individual models into one synchronously or have inefficient or frail peer-to-peer communication, which are potential bottlenecks for the use of federated learning. In this paper, we propose a novel method of asynchronous decentralized federated lifelong learning (ADFL) method that inherits the merits of federated learning and can train on multiple tasks simultaneously without the need for a central node or synchronous training, or less-than-desirable peer-to-peer communication. Thus, overcoming the potential drawbacks of conventional federated learning. We demonstrate excellent performance on the brain tumor segmentation (BRATS) dataset for localizing the left ventricle on multiple image sequences and image orientation. Our framework allows agents to achieve the best performance with a mean distance error of 7.81, better than the conventional central aggregation agent's mean distance error of 11.78, and significantly ($p=0.01$) better than a conventional lifelong reinforcement learning (LL) agent with a distance error of 15.17 after eight rounds of training. In addition, all ADFL agents have better performance than a conventional reinforcement learning (RL) agent with no LL implementation. In conclusion, we developed an ADFL framework with excellent performance and speed-up compared to conventional LL agents.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent reinforcement learning.**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD FL4Data-Mining '23, August 7, 2023, Long Beach, CA, USA

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

KEYWORDS

Federated learning, Lifelong learning, Deep reinforcement learning, Landmark localization

ACM Reference Format:

Guangyao Zheng, Vladimir Braverman, Michael A. Jacobs, and Vishwa S. Parekh. 2023. Asynchronous Decentralized Federated Lifelong Learning for Landmark Localization in Medical Imaging. In *KDD FL4Data-Mining '23, August 7, 2023, Long Beach, CA, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Medical imaging, MRI (Magnetic Resonance Imaging), PET (Positron Emission Tomography), CT (Computerized Tomography), X-ray, and Ultrasound, play a critical role in the diagnosis, prognosis, and preventative care of patients. The use of machine learning methods in medical imaging, such as classification, segmentation, noise reduction, and landmark localization, has been used in different completing complicated environments and settings [4, 8, 13, 14, 20]. However, these methods are usually done on single tasks, without the ability to generalize to other tasks. They often require a full dataset on a device for training, which may cause privacy concerns about patient data and computational constraints for the device specifications [24].

To address these challenges, Federated Learning (FL) has emerged as a promising approach that enables multiple agents to collaboratively train a model without sharing their data [1]. A federated learning system aims to protect data privacy and reduce computational costs at the local agent level by distributing the computation to multiple agents to train the model on their local data and sharing only the model updates with a central server. Federated learning implementations have shown promising results in various medical applications [6, 18, 23]. However, federated learning frameworks often rely on synchronized learning schedules, meaning all participating agents start training at the same time. They also require agents to have the same architecture in order for the central server to aggregate the model weights. Data and agent heterogeneity influence the training speed which greatly reduces the efficiency and challenges the robustness of these approaches [16]. Additionally, Federated learning approaches cannot perform Lifelong Learning (LL), which is an important aspect of machine learning applied to medical imaging. Works have shown the ability to improve accuracy, have excellent performance on multiple tasks, and generalize

[7]. With medical imaging tasks involving constant and rapid altering in imaging environments, such as new imaging sequences or abnormal patient conditions, a federated learning framework that is trained in an older environment may not perform well when evaluated in an unseen environment. One way to address this issue is to train the model again on the new dataset. However, new environment data can potentially be scarce, and this approach may lead to catastrophic forgetting, where the model loses its ability to operate effectively in the older environment. Federated lifelong learning implementations such as Huang et. al. exist [5] in the medical field, but it is limited by their ring-type structure, which means the communication delay can potentially be very high due to unresponsive or failed agents that bottleneck the entire system.

To address all the limitations mentioned above, we propose an asynchronous decentralized federated lifelong learning (ADFLL) approach to landmark localization in medical imaging. This framework leverages Federated learning’s ability to protect data privacy and reduce computational constraints, while also permitting data and agent heterogeneity to be in the system. This framework does not require a central node, and can together lifelong learn multiple tasks without catastrophic forgetting. We provide a flexible, efficient, and robust framework that can be deployed in real-world applications. This paper presents experimental results demonstrating the efficacy of our framework on the 2017 brain tumor segmentation (BraTS) dataset consisting of 8 different image environments and imaging sequences, highlighting its potential to revolutionize landmark localization in the medical imaging field while also maintaining data privacy and reducing computational costs.

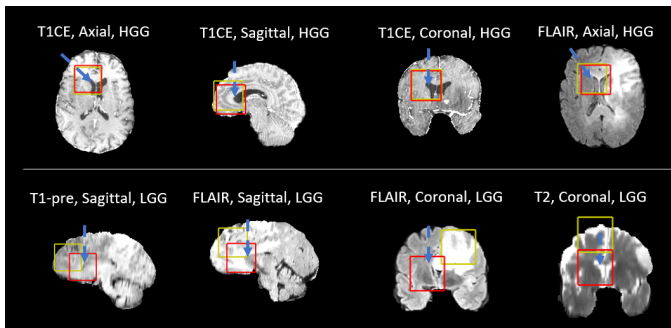


Figure 1: Illustration of the 8 task-environment pairs. The red boxes indicate the true landmark location of the top left ventricle. The yellow box is a predicted location from ADFLL agents during their training progressions

2 EXPERIMENT AND RESULT

2.1 Dataset and Experimental Setup

2.1.1 Clinical Data. For evaluation of our ADFLL framework, we utilized the 2017 brain tumor segmentation (BraTS) dataset [10]. This dataset consisted of 285 patients and included pre-contrast T1-weight, post-contrast T1-weighted, T2-weighted, and Fluid Attenuated Inversion Recovery (FLAIR) sequences in the axial orientation. From this dataset, we randomly sampled a subset of 100 patients to use as our experiment dataset. 60 patients have high-grade glioma

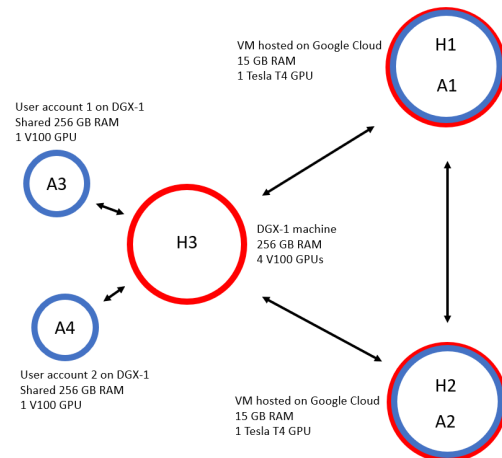


Figure 2: Illustration of the 4-agent decentralized federated lifelong learning framework of our experiment.

(HGG) and 40 patients have low-grade glioma (LGG). We split the 100 patients into two parts 80:20, 80 were used for training and 20 for evaluation, with the training set consisting of 48 HGG and 32 LGG tumors, and the test set consisting of 12 HGG and 8 LGG tumors. We reconstructed the dataset to include all three imaging orientations (coronal, sagittal, and axial). As a result, we obtained a total of twenty-four imaging environments with combinations of two pathologies, 4 imaging sequences, and 3 image orientations. The top left ventricle was chosen as the task for this experiment, and 8 task-environment pairs were sampled as shown in Fig. 1.

2.1.2 Functionality Experimental Setup. Every agent implements a multi-task lifelong reinforcement learning algorithm for localizing landmarks across the human anatomy. The federated lifelong learning component is implemented by integrating experience replay buffers from previous experiences shared by agents across the network for training. There are four agents in this experimental system: we implemented two on an NVIDIA DGX-1 each with an NVIDIA V100 and two on Google Cloud each with an NVIDIA T4. The topology of the system is shown in Fig. 2. The two agents A1, and A2 on Google Cloud have their individual hubs H1 and H2. The two agents A3, and A4 on the DGX-1 are connected to the third hub H3 with a total of three hubs for 4 agents. Since the GPUs on DGX-1 are much more powerful than the GPUs on Google Cloud, A3 and A4 will run significantly faster than A1 and A2. We also implemented asynchronous learning, meaning when the agent finishes training on a task, as long as there are new ERBs that they have not learned from, they will start a new round and learn from those ERBs. Each agent will also get a different image training dataset each round. This process is continued until all four agents complete three rounds of training, guaranteeing all 8 sampled tasks for this experiment will be learned by the system.

All-knowing agent and partially-knowing agent: To better compare our framework with non-lifelong learning ones, we ran two different deep reinforcement learning agents. Agent X is the all-knowing agent, with all 8 datasets available to it at the start

Patient Characteristics	AgentX	Agent Y	AgentM	Agent1	Agent2	Agent3	Agent4
Coronal_LGG_t1	10.05	8.94	10.68	8.06	8.06	10.49	6.08
Coronal_LGG_t2	9.22	8.31	8.94	9.49	7.35	8.25	33.79
Sagittal_LGG_flair	10.77	60.42	14.59	11.75	8.12	8.31	10.49
Axial_LGG_t1ce	7.07	89.91	16.67	31.19	6.16	6.4	3.74
Axial_HGG_flair	4.47	90.05	22.16	66.56	4.24	12.04	22.09
Sagittal_HGG_t1	31.19	65.49	13.15	10.05	6.71	11.75	7.55
Sagittal_HGG_t2	10.25	68.61	11.58	24.7	12.33	22.67	12.37
Coronal_HGG_t1ce	11.22	44.11	23.58	40.47	9.54	39.71	13.19
Mean	11.78	54.48	15.17	25.28	7.81	14.95	13.66
Std. dev	8.16	32.05	5.32	20.46	2.40	11.17	9.87
Ttest (vs. Agent X)		0.01	0.40	0.18	0.22	0.54	0.73
Ttest (vs. Agent M)	0.40	0.01		0.12	0.01	0.95	0.72
Ttest (vs. Agent Y)	0.01		0.01	0.01	0.00	0.01	0.02

Table 1: Comparison of distance error between our agents (Agent 1-4) after round 3, all-knowing deep reinforcement learning agent (Agent X) after round 1, partially-knowing deep reinforcement learning agent (Agent Y) after round 1, and traditional lifelong deep reinforcement learning agent (Agent M) after round 8.

of training for the functionality experiment and 24 datasets for the ablation study. It trained on the available data for one round. Agent Y is the partially-knowing agent, which only has access to one dataset and can therefore only train for one round.

Traditional lifelong deep reinforcement learning agent: To better compare our framework with the traditional lifelong learning framework, we ran an Agent M that has access to the dataset sequentially and is therefore trained for eight rounds to account for learning eight different environments for the functionality experiment and two rounds for the ablation study.

Experimental Metric: The performance metric was set as the terminal Euclidean distance between the agent’s prediction and the target landmark. We performed paired t-tests to compare the performance of the decentralized federated lifelong learning framework with the traditional lifelong learning framework and all-knowing deep reinforcement learning agent and partial-knowing deep reinforcement learning agent. The p-value for statistical significance was set to $p \leq 0.05$.

2.1.3 Ablation Study. We conducted two simulation experiments to evaluate the scalability, flexibility and robustness of our framework. We initialized the same type of the agents in both experiments as in the functionality experiment. For both experiments, we evaluated the average performance of all the agents for the task of localizing the top left ventricle across all 24 imaging environments. Additionally, since it is prohibitively expensive to experiment on 24 different machines, these systems were simulated on the NVIDIA DGX-1, with a synchronous training protocol.

Addition of agents experiment: We initialized a system with four agents, as previously described in the functionality experiment. We subsequently increased the number of agents in the system from

4 to 16 agents over the progression of 4 rounds (4,8,12,16). We further simulated a communication dropout of 75% to account for network communication issues in the real world leading information loss while transmitting ERBs across agents. The goal of this experiment was to demonstrate how newer agents joining the system at different points in time can take advantage of the available within the system to learn the collective knowledge available in the system within just one round.

Deletion of agents experiment: In the deletion experiment, we gradually decreased The number of agents in the system from 24 to 1 agent over the progression of 5 rounds (24,12,6,3,1). The communication for this experiment was also simulated with a 75% dropout. The goal of this experiment was to demonstrate how the proposed ADFLL system preserves the collective knowledge in a lifelong learning manner across all the tasks even as the agents contributing the knowledge leave the system.

2.2 Results

We conducted a functionality experiment based on 8 sub-task-environment pairs: Axial HGG t1ce, Sagittal HGG t1ce, Coronal HGG t1ce, Axial HGG flair, Sagittal LGG flair, Coronal LGG flair, Coronal LGG t2, Sagittal LGG t1. We sampled one image from each task to test the performance of our model and baseline models. Each round the four federated lifelong learning will receive a new task. They will begin the next round when there is also ERB to train from. Since the agents’ training speeds are very different, A1 and A2 will finish their tasks slower, allowing them to learn from more ERBs at once. As shown in Table 1, after three rounds of training, A2 was able to achieve a mean distance error of 7.81 on all 8 tasks, compared to the all-knowing agent’s 11.78 ($p=0.22$), but significantly lower compared to partially-knowing agent’s 54.58

($p < 0.001$), and the traditional lifelong learning agent's 15.17 ($p = 0.01$) after eight rounds of training. Since this is a real-world experiment, other agents trained faster than A2, meaning that they did not have all the ERBs available to them when they started their last round of training, resulting in their performances being worse than A2. The possible reasons for Agent 1 to have lower performance than Agent X and Agent M can be because it did not have all ERBs or the training was stuck at a local minimum. This can be easily solved by sharing the model parameters of the latest agent. Note that the all-knowing agent and the partially-knowing agent only train for 1 round for this experiment because they have no lifelong learning capability.

Compared to the All-knowing agent shown in 3, our framework is able to achieve an excellent performance boost. Compared to the traditional lifelong learning agent, our framework is able to achieve a significant performance boost and an outstanding speedup.

In our two ablation studies, our framework showed scalability of up to 24 agents, robustness against network dropout, and flexibility in system topology. As shown in Figure 4, we see that the average Euclidean distance error across all agents decreases as more agents are added to the system, with an average Euclidean distance error of 16.89 at the end of 4 rounds. As shown in Figure 5, we also see that the average Euclidean distance error across all agents decreases, while half of the agents are deleted every round, resulting in an average Euclidean distance error of 8.55 after 5 rounds. This shows that the knowledge agents learned, captured in ERBs are not lost when agents are removed from the system. And when agents are being added, the new agents can catch up to existing agents in one round. Moreover, the 75% dropout rate that is applied to every round of both experiments shows the robustness of our framework against network failures, a major bottleneck for federated learning frameworks.

3 CONCLUSION

Previous works have explored the application of federated learning to the medical field [12, 19, 22]. They have shown decentralized federated learning system setups, each demonstrating good performance in their experiment tasks. But because of their system topology implementation, one node failure can potentially collapse the entire system. Moreover, the learning tasks examined were limited, binary classifications or MNIST dataset classification, resulting in limited potential applications. Additionally, their implementation offers a synchronous training procedure, which means in a real application scenario, users of their framework will have to coordinate the training process.

Asynchronous federated learning has also been explored in other areas [3]. They offer the ability to deal with nodes with different computational power but lack the decentralization that allows the system to be more flexible.

Asynchronous decentralized federated learning has also been explored [5, 9]. However, they are still limited in their system implementation. The cost of removing a central node is a quadratic complexity communication scheme in that every node communicates with every node.

We have demonstrated a privacy-aware, asynchronous decentralized federated learning system with robust and efficient system topology. We have demonstrated excellent performance on landmark localization tasks on the BraTS image dataset. Our framework performs better than all-knowing deep reinforcement learning agents and traditional lifelong learning agents. Moreover, in our ablation study, our framework demonstrated excellent scalability, flexibility, and robustness. In the future, we will optimize our approach, expand the system further, and increase computational efficiency.

ACKNOWLEDGMENTS

To Robert, for the bagels and explaining CMYK and color spaces.

REFERENCES

- [1] Mohammed Adnan, Shivam Kalra, Jesse C. Cresswell, Graham W. Taylor, and Hamid R. Tizhoosh. 2022. Federated learning and differential privacy for medical image analysis. *Scientific Reports* 12, 1 (Dec. 2022). <https://doi.org/10.1038/s41598-022-05539-7>
- [2] Amir Alansary, Loic Le Folgoc, Ghislain Vaillant, Ozan Oktay, Yuanwei Li, Wenjia Bai, Jonathan Passerat-Palmbach, Ricardo Guerrero, Konstantinos Kamnitsas, Benjamin Hou, et al. 2018. Automatic view planning with multi-scale deep reinforcement learning agents. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 277–285.
- [3] Yujing Chen, Yue Ning, and Huzefa Rangwala. 2019. Asynchronous Online Federated Learning for Edge Devices. *ArXiv abs/1911.02134* (2019).
- [4] Junlong Cheng, Shengwei Tian, Long Yu, Chengrui Gao, Xiaojing Kang, Xiang Ma, Weidong Wu, Shijia Liu, and Hongchun Lu. 2022. ResGANet: Residual group attention network for medical image classification and segmentation. *Medical Image Analysis* 76 (2022), 102313. <https://doi.org/10.1016/j.media.2021.102313>
- [5] Yixing Huang, Christoph Bert, Stefan Fischer, Manuel Schmidt, Arnd Dörfler, Andreas Maier, Rainer Fietkau, and Florian Putz. 2022. Continual Learning for Peer-to-Peer Federated Learning: A Study on Automated Brain Metastasis Identification. *arXiv:arXiv:2204.13591*
- [6] Meirui Jiang, Zirui Wang, and Qi Dou. 2022. HarmoFL: Harmonizing Local and Global Drifts in Federated Learning on Heterogeneous Medical Images. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 1 (Jun. 2022), 1087–1095. <https://doi.org/10.1609/aaai.v36i1.19993>
- [7] Neerav Karani, Krishna Chaitanya, Christian Baumgartner, and Ender Konukoglu. 2018. A Lifelong Learning Approach to Brain MR Segmentation Across Scanners and Protocols. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger (Eds.). Springer International Publishing, Cham, 476–484.
- [8] M.O. Khairandish, M. Sharma, V. Jain, J.M. Chatterjee, and N.Z. Jhanjhi. 2022. A Hybrid CNN-SVM Threshold Segmentation Approach for Tumor Detection and Classification of MRI Brain Images. *IRBM* 43, 4 (2022), 290–299. <https://doi.org/10.1016/j.irbm.2021.06.003>
- [9] Qi Liu, Bo-Jun Yang, Zhaojian Wang, Dafeng Zhu, Xinyi Wang, Kai Ma, and Xiping Guan. 2022. Asynchronous Decentralized Federated Learning for Collaborative Fault Diagnosis of PV Stations. *IEEE Transactions on Network Science and Engineering* 9 (2022), 1680–1696.
- [10] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. 2014. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE transactions on medical imaging* 34, 10 (2014), 1993–2024.
- [11] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [12] TV Nguyen, MA Dakka, SM Diakiw, MD VerMilyea, M Perugini, JMM Hall, and D Perugini. 2022. A Novel Decentralized Federated Learning Approach to Train on Globally Distributed, Poor Quality, and Protected Private Medical Data. <https://doi.org/10.21203/rs.3.rs-1371143/v1>
- [13] Julia M. H. Noothout, Bob D. De Vos, Jelmer M. Wolterink, Elbrich M. Postma, Paul A. M. Smeets, Richard A. P. Takx, Tim Leiner, Max A. Viergever, and Ivana Išgum. 2020. Deep Learning-Based Regression and Classification for Automatic Landmark Localization in Medical Images. *IEEE Transactions on Medical Imaging* 39, 12 (2020), 4011–4022. <https://doi.org/10.1109/TMI.2020.3009002>
- [14] Liangrui Pan, Hetian Wang, Lian Wang, Boya Ji, Mingting Liu, Mitchai Chongcheawchannan, Jin Yuan, and Shaoliang Peng. 2022. Noise-reducing attention cross fusion learning transformer for histological image classification

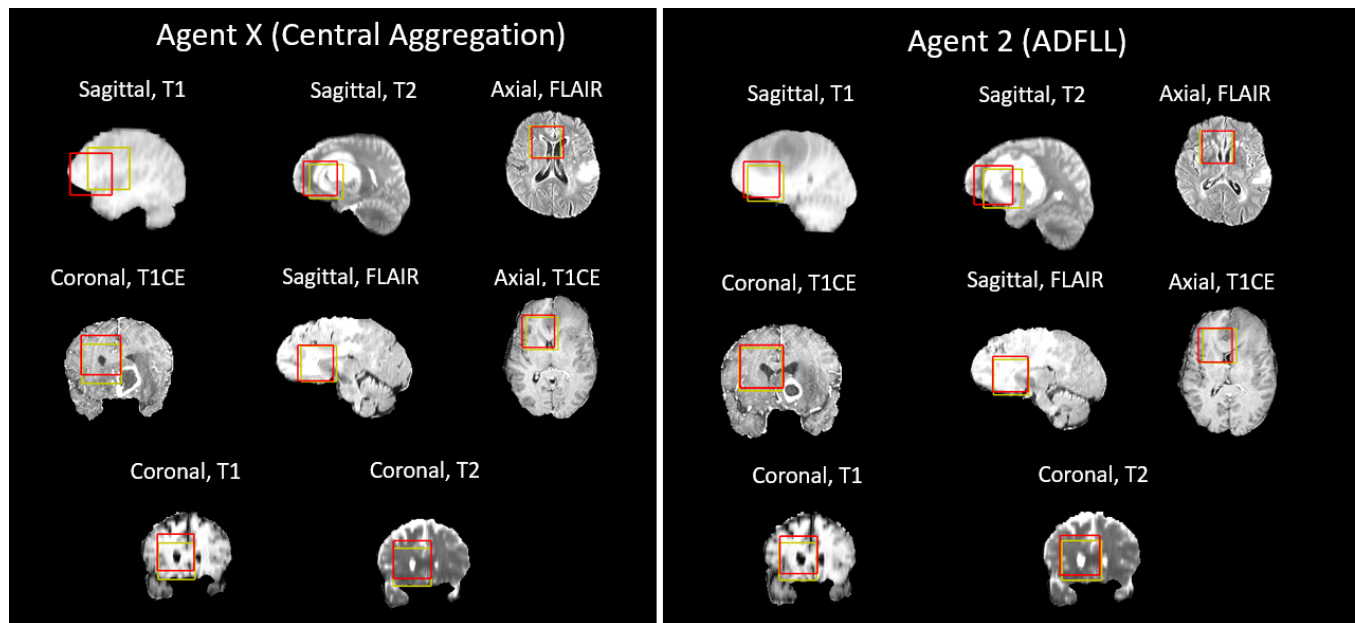


Figure 3: Comparison of distance error across 8 different tasks between Agent X (Central Aggregation) and Agent 2 (ADFLL).




Figure 4: Comparison of distance error of all agents in the system across 4 rounds of training

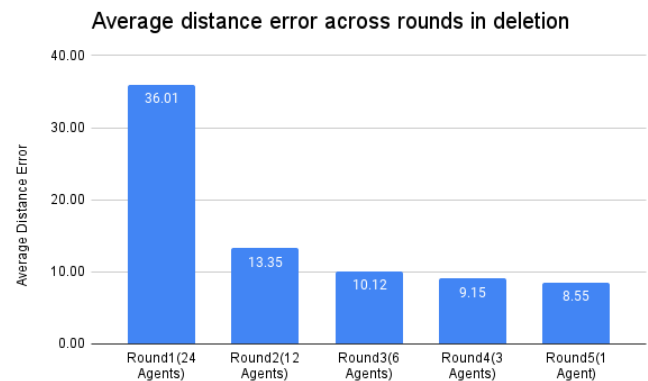
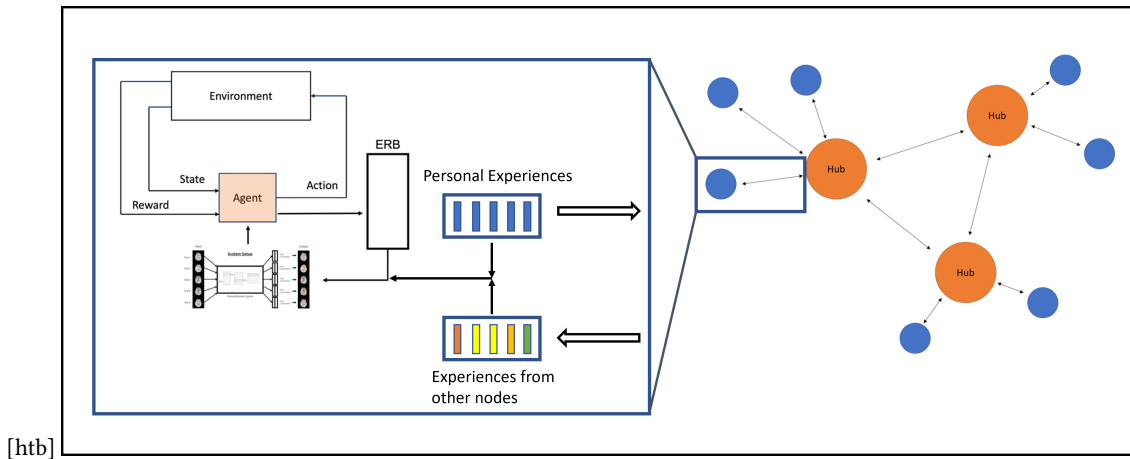


Figure 5: Comparison of distance error of all agents in the system across 4 rounds of training

of osteosarcoma. *Biomedical Signal Processing and Control* 77 (2022), 103824. <https://doi.org/10.1016/j.bspc.2022.103824>

- [15] Vishwa S Parekh, Vladimir Braverman, Michael A Jacobs, et al. 2020. Multitask radiological modality invariant landmark localization using deep reinforcement learning. In *Medical Imaging with Deep Learning*. PMLR, 588–600.
- [16] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R. Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus Maier-Hein, Sébastien Ourselin, Micah Sheller, Ronald M. Summers, Andrew Trask, Daguang Xu, Maximilian Baust, and M. Jorge Cardoso. 2020. The future of digital health with federated learning. *npj Digital Medicine* 3, 1 (1 Dec. 2020). <https://doi.org/10.1038/s41746-020-00323-1>
- [17] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. 2019. Experience replay for continual learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [18] Holger R. Roth, Ken Chang, Praveer Singh, Nir Neumark, Wenqi Li, Vikash Gupta, Sharut Gupta, Liangqiong Qu, Alvin Ihsani, Bernardo C. Bizzo, Yuhong Wen, Varun Buch, Meesam Shah, Felipe Kitamura, Matheus Mendonça, Vitor Lavor, Ahmed Harouni, Colin Compass, Jesse Tetreault, Prema Dogra, Yan Cheng,

- Selnur Erdal, Richard White, Behrooz Hashemian, Thomas Schultz, Miao Zhang, Adam McCarthy, B. Min Yun, Elshaimaa Sharaf, Katharina V. Hoebel, Jay B. Patel, Bryan Chen, Sean Ko, Evan Leibovitz, Etta D. Pisano, Laura Coombs, Daguang Xu, Keith J. Dreyer, Ittai Dayan, Ram C. Naidu, Mona Flores, Daniel Rubin, and Jayashree Kalpathy-Cramer. 2020. Federated Learning for Breast Density Classification: A Real-World Implementation. In *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning: Second MICCAI Workshop, DART 2020, and First MICCAI Workshop, DCL 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4–8, 2020, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 181–191.
- [19] Abhijit Guha Roy, Shayan Siddiqui, Sebastian Pölsterl, Nassir Navab, and Christian Wachinger. 2019. BrainTorrent: A Peer-to-Peer Environment for Decentralized Federated Learning. *ArXiv abs/1905.06731* (2019).
- [20] Sumit Tripathi, Taresh Sarvesh Sharan, Shiru Sharma, and Neeraj Sharma. 2022. An Augmented Deep Learning Network with Noise Suppression Feature for Efficient Segmentation of Magnetic Resonance Images. *IETE Technical Review* 39, 4 (2022), 960–973. <https://doi.org/10.1080/02564602.2021.1937349>



[htb]

Figure 6: Illustration of decentralized federated learning setup. Blue circles represent individual agents, and orange circles represent hubs.

- [21] Athanasios Vlontzos, Amir Alansary, Konstantinos Kamnitsas, Daniel Rueckert, and Bernhard Kainz. 2019. Multiple Landmark Detection using Multi-Agent Reinforcement Learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 262–270.
- [22] Zhao Wang, Yifan Hu, Shiyang Yan, Zhihao Wang, Ruijie Hou, and Chao Wu. 2022. Efficient Ring-Topology Decentralized Federated Learning with Deep Generative Models for Medical Data in eHealthcare Systems. *Electronics* 11, 10 (2022). <https://doi.org/10.3390/electronics11101548>
- [23] Zengqiang Yan, Jeffrey Wicaksana, Zhiwei Wang, Xin Yang, and Kwang-Ting Cheng. 2021. Variation-Aware Federated Learning With Multi-Source Decentralized Medical Image Data. *IEEE Journal of Biomedical and Health Informatics* 25, 7 (2021), 2615–2628. <https://doi.org/10.1109/JBHI.2020.3040015>
- [24] Joo Hun Yoo, Hyejun Jeong, Jaehyeok Lee, and Tai-Myoung Chung. 2021. Federated Learning: Issues in Medical Application. In *Future Data and Security Engineering: 8th International Conference, FDSE 2021, Virtual Event, November 24–26, 2021, Proceedings*. Springer, Berlin, Heidelberg, 3–22.

A RESEARCH METHOD

A.1 Deep Reinforcement Learning

We created a deep reinforcement learning framework that utilizes the deep Q-network (DQN) algorithm, which is depicted in Figure 6. The 3D DQN model we used in this paper was adapted from existing works [2, 11, 15, 21]. The environment was represented by a 3D imaging volume with x , y , and z dimensions. The agent was represented by a 3-dimensional bounding box with six possible actions: moving in the positive or negative in the x , y , or z axis. The state was defined by the current location (or a chain of locations) of the agent, each represented by a 3-dimensional bounding box. The reward was calculated by the change in distance to the target landmark location before and after the agent takes an action. The agent’s exploration within the environment generated state-action-reward-resulting state $[s, a, r, s']$ tuples, which were recorded in the experience replay buffer (ERB) over multiple episodes. The information contained in the ERBs are non-sensitive information, as the action and reward are numbers regarding the DRL model, and the state and resulting states are small fractions of the total 3D image, roughly 0.3%.

A.2 Lifelong Learning

We implemented lifelong learning using selective experience replay [17]. The goal of selective experience replay is to avoid catastrophic forgetting by focusing on selected experiences from previous tasks. Additionally, this technique is agnostic to the model being used and enables the sharing of experiences across different models. To achieve lifelong learning, we utilized a selective experience replay buffer that collects a sequence of experience samples throughout the model’s training process. In order to learn a generalized representation of both current and past tasks, the model selects a batch of experiences from both the ERB of its current task and from the replay buffers of previous tasks during training.

A.3 Asynchronous Decentralized Federated Lifelong Learning

We developed the Asynchronous Decentralized Federated lifelong Learning (ADFLL) by constructing a network of lifelong deep reinforcement learning agents. Each agent shares their database of personal experiences with each other to facilitate learning from each other experiences. More specifically, once an agent finishes training with a dataset and an ERB, the resulting experience from the training is shared with the network. Furthermore, we modified the training setup for each agent to sample experiences from the current dataset ERB, the agent’s personal experiences and the incoming experiences from other agents, as shown in Fig. 6. As a result, every agent in the network can learn from each other’s experiences, thereby integrating federated lifelong learning capability.

In a naive setup, every agent would communicate their experiences with every other agent in the network. However, such an all-to-all communication setup is highly inefficient and not scalable as it would require a large amount of communication bandwidth. To address this issue, we implemented a homogeneous distributed database system as illustrated in Fig. 6. As shown in Fig. 6, our network consists of a predefined set of hub nodes that act as communication hubs for spatially adjacent nodes in the network. Subsequently, every agent in the network exclusively communicates with their

ERB Id	Modality	Landmark	Pathology	Hub Agent IDs that currently have this ERB
ERB_XXXXX	T1	Ventricle	LGG	XXXXX

Figure 7: Snapshot of the shared database maintained by the hub nodes

nearest hub node at the end of each personal training round. The experience sharing between an agent and a hub node is bidirectional. Finally, every hub node maintains a shared experience database on the network as shown in Fig. 7. The hub nodes periodically communicate with each other to synchronize their databases. The agents in the system are not required to have standardized training

speed or start training at the same time. The hub will regulate and preserve the experiences in the system and agents in the system can train on different tasks. An example of this system is demonstrated in Fig. 6.

The advantage of our system setup is that it is robust against node or hub failures. When a node fails, the only loss is the training information from that node, and when a hub fails, the loss is the ERBs it contains but other hubs do not. Moreover, the communication complexity is linear with respect to the number of nodes, each node only needs to communicate with its respective hub, and hubs sync periodically. Compared to other federated learning systems, centralized or not, they either are prone to system-wide failure caused by a node failure, or sacrifice communication complexity to prevent system-wide failures.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009