

Holonic Learning: A Flexible Agent-based Distributed Machine Learning Framework

Ahmad Esmaeili
Purdue University
West Lafayette, United States
aesmaei@purdue.edu

Zahra Ghorrati
Purdue University
West Lafayette, United States
zghorrat@purdue.edu

Eric T. Matson
Purdue University
West Lafayette, United States
ematson@purdue.edu

ABSTRACT

Ever-increasing ubiquity of data and computational resources in the last decade have propelled a notable transition in the machine learning paradigm towards more distributed approaches. Such a transition seeks to not only tackle the scalability and resource distribution challenges but also to address pressing privacy and security concerns. To contribute to the ongoing discourse, this paper introduces Holonic Learning (HoL), a collaborative and privacy-focused learning framework designed for training deep learning models. By leveraging holonic concepts, the HoL framework establishes a structured self-similar hierarchy in the learning process, enabling more nuanced control over collaborations through the individual model aggregation approach of each holon, along with their intra-holon commitment and communication patterns. HoL, in its general form, provides extensive design and flexibility potentials. For empirical analysis and to demonstrate its effectiveness, this paper implements HoloAvg, a special variant of HoL that employs weighted averaging for model aggregation across all holons. The convergence of the proposed method is validated through experiments on both identically and independently distributed (IID) and Non-IID settings of the standard MNIST dataset. Furthermore, the performance behaviors of HoL are investigated under various holarchical designs and data distribution scenarios. The presented results affirm HoL's prowess in delivering competitive performance particularly, in the context of the Non-IID data distribution.

KEYWORDS

Distributed learning; Holonic Learning; Collaborative Learning; Edge Computing

ACM Reference Format:

Ahmad Esmaeili, Zahra Ghorrati, and Eric T. Matson. 2024. Holonic Learning: A Flexible Agent-based Distributed Machine Learning Framework. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 9 pages.

1 INTRODUCTION

Today's interconnected world, marked by the proliferation of information and computational capabilities across vast digital landscapes, has led to critical challenges concerning the scalability,

adaptability, speed, and support for diversity in the traditional centralized learning models. As a solution to such limitations and harnessing the collective power of the networked devices—whether servers or edge devices—recent decade has witnessed a substantial growth in the rise of distributed learning methodologies that not only expedite training process but also enhances robustness and privacy.

Federated Learning (FL) [17] stands out as one of the prominent distributed learning approaches attracting attentions from both academia and industry [26]. By empowering end users to participate in the collaborative learning process while retaining ownership of their data, FL ensures both data security and privacy on one hand, and leveraging the potential of computationally constrained devices on the other. Despite its apparent simplicity in its original format, FL introduces a series of intricate open problems and challenges tackled by a plethora of research work. Communication efficiency, convergence rates, privacy preservation, addressing potential adversarial attacks, robustness across heterogeneous devices and data distributions, non-identically and independently distributed (non-IID) data, mitigating potential aggregation bias, and ensuring fairness in model updates are among the nuanced and critical open problems that arise within the realm of Federated Learning [13].

Emerging real-world scenarios and applications have given rise to new trends in methodologies that focus on relaxing some of FL's core assumptions. Some of the noteworthy examples include peer-to-peer Learning [12, 22], edge-based learning [19, 24], and split learning [9, 23]. In fully decentralized Peer-to-Peer (P2P) learning the objective revolves around replacing the central orchestration server with a peer-to-peer communications topology between the clients to mitigate the risk of single-point of failure and address the potential communication bottlenecks in large-scaled systems. In edge-based FL system, on the other hand, the emphasis lies in relocating the parameter server to the vicinity of edge devices to enhance latency and strike a balance between computation and communication. Unlike the previous approaches, split learning introduces another dimension of parallelism and shifts the focus from data partitioning to splitting the execution of the learning model into layers across multiple clients and the server.

Hierarchical approaches have consistently been instrumental in shaping the development and structure of intricate systems, and FL is no exception to this pattern. In Hierarchical FL approaches, the clients are clustered into groups to reduce communication overheads and add an extra privacy layer through periodic model updates within and across groups [18]. The studies reported in [1, 11, 16, 25] exemplify noteworthy contributions that center on



This work is licensed under a Creative Commons Attribution International 4.0 License.

enhancing energy consumption, latency, and communication efficiency. In [1], a two-tier hierarchical FL architecture is employed for wireless heterogeneous cellular networks. This structure clusters mobile users into localized cells, leveraging local FL within micro-cell clusters and aggregating model updates at the macro-level through base stations. The HierFAVG model introduced in [16] presents a three-layer hierarchy—client-edge-cloud—extending FL with a two-level average-based model aggregation approach. In [25], a dual-level framework is proposed, comprising mobile devices and edge/cloud servers. This framework optimizes resource allocation and aggregation intervals adaptively. Finally, the work outlined in [11], unlike the preceding models, employs an arbitrary number of levels in its design. Their proposed tree-like structure has a fixed hierarchical topology and is constructed based on the physical proximity and connectivity of devices. Notably, their model utilizes a device-to-device (D2D) consensus method to update the local models at the lowest level, followed by a sequence of recursive aggregations through the upstream path, culminating in a central server.

This paper presents a holonic solution for the distributed learning problem. The multi-level architecture of holonic systems embodies an arrangement of self-contained entities known as holons, forming a self-similar structure rooted in the whole/part relationship. Benefiting from their balance between local autonomy and coordinated decision-making, holonic systems adeptly tackle challenges such as adaptability, fairness, and scalability. Additionally, their decentralized authority and decision-making approach mitigate the risk of a single point of failure, further enhancing their resilience. Holonic solutions have already been employed in the design and management of machine learning systems. As an example, the Hierarchical Agent-based Machine LEarning PlaTform (HAMLET) presented in [5] leverages holons to organize machine learning resources distributed across an array of networked devices and orchestrates automatic parallel training and testing tasks. Our paper, in contrast, primarily focuses on distributed learning and introduces a novel holonic collaborative learning methodology.

Figure 1 portrays an illustrative multi-level holonic model built over a set of connected devices through grouping them together at different levels of granularity. The construction of such structure indeed constitutes a highly challenges phase, potentially contingent upon various factors ranging from physical constraints to trust and security considerations. To maintain our main focus and streamline the analysis, we assume that the holonic construction is preexisting and accessible. We will return to this figure in the subsequent section as we succinctly review the fundamentals of holonic concepts.

The Holonic Learning (HoL) model presented here offers a novel approach to distributed learning. It views holons at each level as self-contained agents collaborating to optimize model parameters, granting autonomy to each holon while aligning with the superordinate goal. Unlike fully decentralized peer-to-peer models, HoL allows diverse policy enforcement across levels, distinguishing it from hierarchical approaches with central authority. Key contributions include: (1) introducing a collaborative learning framework that combines vertical and horizontal interactions, drawing inspiration from holonic structures for sophisticated collaboration patterns while maintaining an intuitive design; and (2) demonstrating how

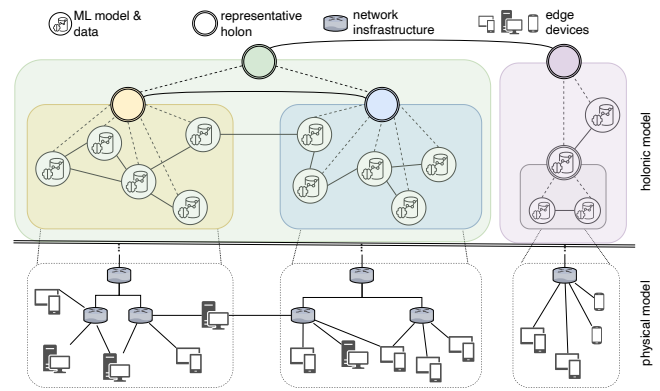


Figure 1: An example HoL model built over an IoT network

the HoL can generalize the current distributed ML frameworks, empirically showcasing its consistent convergence behavior across holarchies and data distributions. The presented performance evaluations based on training loss and test accuracy highlight the advantages of holarchical structures over plain collaboration networks.

It is worth mentioning that HoL is positioned not as a competitor to FL but as a more generalized and versatile framework for distributed learning, suitable for heterogeneous and dynamic environments. As it is shown later, classical FL can be assumed as a special case within HoL, where HoL’s design is confined to a single level without horizontal inter-agent interactions. Additional key distinctions include but are not limited to: (1) incorporating aggregation at different levels of HoL, enabling diverse, flexible, and personalized processes for both inter- and intra-holon collaborations; (2) leveraging horizontal collaborations to address the single point of failure issue; (3) preserving device/organization agency, allowing decisions based on peer interactions, in contrast to FL’s centralized authoritative process; and (4) supporting multi-layered security and privacy measures through a modular design, enhancing overall robustness, especially in scenarios with numerous devices where computationally intensive encryption methods may be impractical.

The paper is organized as follows: Section 2 reviews the fundamental principles of holonic systems and the holonic multi-agent concepts employed in this study. Section 3 delves into the formulation of our holonic learning model and presents the details of a more specific use-case analyzed in this paper. Section 4 discusses how HoL generalizes some of existing distributed models. Section 5 elaborates on the experiments conducted and examines the model’s characteristics from multiple perspectives. Finally, Section 6 concludes the paper and offers suggestions for future research directions.

2 THE PREAMBLE

2.1 Holonic Structures

Formally introduced by Arthur Koestler, the term “holon” refers to a philosophical concept describing entities that are simultaneously whole in themselves and parts of larger systems [14]. In other words, a self-contained holon embodies both individuality and integration,

which helps to model complex systems and understand how they evolve in response to changing circumstances. Although it was originally used to describe the self-similar hierarchical structure of biological and social system, it has found its way into many disciplines such as computer science [5], engineering [3], physics [4], and manufacturing [21], to name a few. This section highlights the key definitions related to the holonic systems that are heavily used in the subsequent context.

As stated above, the key idea of holons is their dual nature, or as described in [15], how they behave like a two-faced Janus: the face representing the whole looks inward and the other looks up, or outward. That is at the same time that each holon is independent with its own characteristics and functions, it is also a relatively more concrete component of larger systems. Applying such a whole-part relationship to different entities of a system create a multi-level interconnected structure of holons, where the behavioral variations of each holon at any level is characterized by operational variation of its subordinate holons in a lower level. In the holonic terminology, such a hierarchical architecture is called “holarchy”, and the higher-level and lower-level holons are respectively referred to as “super-holons” and “sub-holons”. In this paper, we also use the terms “terminal holons” and “non-terminal” holons to refer to holons at the lowest level of the holarchy and the ones in the intermediate layers, respectively.

When incorporating holonic principles into the realm of computer science and the design of distributed autonomous systems, the “Holonic Multi-Agent Systems (HMAS)” emerges as an exceptionally promising and broadly applicable model to embrace. As originally proposed in [8], a holon (holonic-agent) is formed by a group of agents giving up part of their autonomy to the super-holon and through the designation of one agent as the representative called “head”. With administrative and/or directive responsibilities, the head plans for the holon on basis of its subordinates. With that said, a holon is formally characterized by the tuple $\langle head, subholons, C \rangle$, where C denote commitments, i.e., the itera-holon relationships agreed by all members [7]. In this formalization, every agent (a) within a Multi-Agent System (MAS) is considered as an atomic (terminal) holon denoted as $\langle a, a, \emptyset \rangle$, positioned at the lowest tier of the holarchy. Additionally, there are numerous ways to construct each higher-level holon by combining its subordinate elements, resulting in an exponential diversity of configurations.

There are multitude of research efforts, spanning both domain-specific and generic contexts, that have delved into the dynamic formation of holarchies, such the ones reported in [2, 6, 10] and [20] to name a few. Due to our focus on learning, this paper opts for a predetermined, static structure for the utilized holarchy, while acknowledging that relaxing this assumption could potentially pave the way for novel research trajectories within the realm of distributed learning. To emphasize both inter- and intra-holon interactions, we define a “holarchical network” as depicted in Fig. 2. In constructing a holarchy, super-holons are virtual concepts overseen by head holons. These head holons can be specialized terminal holons designated solely for this role, or they can be existing holons selected (temporarily or permanently) to serve as heads while concurrently fulfilling their duties as subordinates at higher levels. To simplify the discussion in this paper, we assume the former cases in our HoL framework.

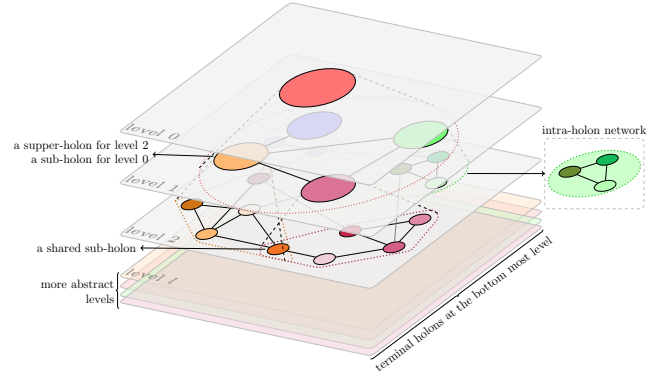


Figure 2: An example holarchical network. The levels are numbered from top to bottom, and different color shades represent the heterogeneity of the holons.

2.2 Notations

We use the following notations throughout this paper. Please note that we use the term “holon” to refer to the nodes of the HoL model, not only to comply with the holonic terminologies, but also due to the fact that it facilitates referring to the both physical and abstract nodes residing at different level of the proposed HoL model. At the physical level, one can safely assume each holon represent an edge device.

- $h_i^{(l)}$: the i -th terminal holon at level l .
- $\mathbf{h}_i^{(l)}$: the i -th super-holon at level l . This notations is also used to denote the head of the holon.
- $h_i^{(l)}$: the i -th holon at level l regardless of specifying if it is terminal or non-terminal.
- $n_i^{(l)}$: the number of subordinates in the holon $\mathbf{h}_i^{(l)}$.
- $\mathcal{M}_i^{(l)}$: the learning model of holon $h_i^{(l)}$.
- $\theta_i^{(l)}(t)$: the model parameters of holon $h_i^{(l)}$ at the end of iteration t .
- $\mathcal{F}(\theta_i^{(l)}; \mathbf{x}^{(j)}, y^{(j)})$: the loss function of holon $h_i^{(l)}$ for data sample $(\mathbf{x}^{(j)}, y^{(j)})$. For the sake of brevity we write it as $\mathcal{F}_i^{(l)}(\theta; \xi^{(j)})$ or $\mathcal{F}_i^{(l)}(\theta)$ depending on the detail level.
- $\mathbb{X}_i^{(l)}$: the private dataset used for training holon $h_i^{(l)}$'s model.
- $C_i^{(l)}$: the local update stopping criteria of holon $h_i^{(l)}$.
- $\mathcal{G}_i^{(l)}$: the communication graph of super-holon $\mathbf{h}_i^{(l)}$.
- $\mathcal{U}_i^{(l)}$: the superior(s) of holon $h_i^{(l)} \in \{h_i^{(l)}, \mathbf{h}_i^{(l)}\}$ in the holarchy.
- $\mathcal{B}_i^{(l)}$: the subordinate(s) of holon $h_i^{(l)} \in \{h_i^{(l)}, \mathbf{h}_i^{(l)}\}$ in the holarchy.
- $\mathcal{N}_i^{(l)}$: the neighbor(s) of holon $h_i^{(l)} \in \{h_i^{(l)}, \mathbf{h}_i^{(l)}\}$ according to the graph $\mathcal{G}_j^{(l-1)}$ s.t. $h_j^{(l-1)} \in \mathcal{U}_i^{(l)}$.
- $\mathcal{A}_i^{(l)}$: the aggregation function used by holon $h_i^{(l)}$.
- $[M_1|M_2]$: the column-wise concatenation of two matrices M_1 and M_2 .

3 HOLONIC LEARNING MODEL

This section delves into the details of the proposed distributed holonic learning model. As stated in section 2.1 the holons of a holarchy can be seen as reference points, representing different development stages within the system. In other words, each level of the holarchy is self-contained and consists of autonomous entities with specific interactions and problems, depending on the problem's abstraction.

3.1 Problem Formulation

The presented holonic learning employs a level-based learning model. In the example depicted in Fig. 1, for instance, the learning problem on the cloud holon is viewed as the process of optimizing model parameters based on the information that is shared by its two subordinate collaborators without worrying about the details of the other levels. With that said, formally, each individual holon locally seeks the solution to the minimization of its empirical risk function $F_i^{(l)}$. For the terminal holon $h_i^{(l)}$, the problem is defined as:

$$\min_{\theta} \left[F_i^{(l)}(\theta) := \frac{1}{|\mathbb{X}_i^{(l)}|} \sum_{j=1}^{|\mathbb{X}_i^{(l)}|} \mathcal{F}_i^{(l)}(\theta; \xi^{(j)}) \right] \quad (1)$$

and for the higher level holon $h_i^{(l)}$, we have:

$$\min_{\theta} \left[F_i^{(l)}(\theta) := \sum_{h_k^{(l+1)} \in \mathcal{B}_i^{(l)}} \phi_k^{(l+1)} F_k^{(l+1)}(\theta) \right] \quad (2)$$

where $\phi_k^{(l+1)}$ denote the weight the holon applies to the empirical risk value of the subordinate holon $h_k^{(l+1)}$.

The proposed HoL model tries to find the near-optimal solution for the above-mentioned optimization problem in a sequence of local model updates and aggregation of the feedback it receives from its neighbors, subordinates, superiors. Let $\theta_i^{(l)} = [\theta_{i,1}^{(l)}, \theta_{i,2}^{(l)}, \dots, \theta_{i,d}^{(l)}]^\top \in \mathbb{R}^d$ specify the model parameter vector of holon $h_i^{(l)}(t)$. For n , b , and u updates that the holon has received from its neighbors, internal/subordinates, and superiors, respectively, we define the corresponding matrices based the received contributions as follows:

$$\Theta_{\mathcal{N}_i}^{(l)}(t) = [\theta_j^{(l)} | \dots | \theta_k^{(l)}] \in \mathbb{R}^{d \times n} \quad (3)$$

and

$$\Theta_{\mathcal{B}_i}^{(l)}(t) = \begin{cases} [\theta_j^{(l+1)} | \dots | \theta_m^{(l+1)}] \in \mathbb{R}^{d \times b} & \text{if } h_i^{(l)}(t) \text{ is non-terminal} \\ [\theta_i^{(l)}] \in \mathbb{R}^d & \text{if } h_i^{(l)}(t) \text{ is terminal} \end{cases} \quad (4)$$

$$\Theta_{\mathcal{U}_i}^{(l)}(t) = \begin{cases} [\theta_j^{(l-1)} | \dots | \theta_o^{(l-1)}] \in \mathbb{R}^{d \times u} & \text{if } l > 0 \\ \text{none} & \text{if } l = 0 \end{cases} \quad (5)$$

Additionally, each holon uses the column vectors $\mathbf{w}_{\mathcal{N}_i}^{(l)}(t) \in \mathbb{R}^n$, $\mathbf{w}_{\mathcal{B}_i}^{(l)}(t) \in \mathbb{R}^b$, and $\mathbf{w}_{\mathcal{U}_i}^{(l)}(t) \in \mathbb{R}^u$ to weigh the contributions of its neighbors and subordinates, respectively. With that said, each holon

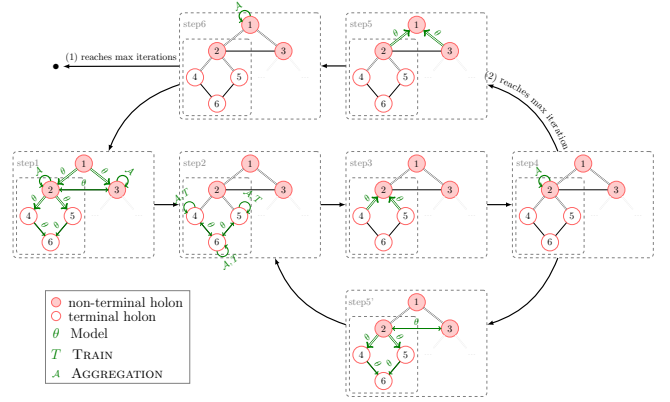


Figure 3: An example demonstrating the flow of communications in Algorithm 1.

uses its own aggregation function to update the value of its model parameters for the next step as follows:

$$\theta_i^{(l)}(t+1) = \mathcal{A}_i^{(l)}(\Theta_i^{(l)}(t), \mathbf{W}_i^{(l)}(t)) \quad (6)$$

where $\Theta_i^{(l)}(t)$ and $\mathbf{W}_i^{(l)}(t)$ are referred to as contribution matrix and contribution weights vector, respectively, and are constructed as follows:

$$\Theta_i^{(l)}(t) := [\Theta_{\mathcal{B}_i}^{(l)}(t) | \Theta_{\mathcal{N}_i}^{(l)}(t) | \Theta_{\mathcal{U}_i}^{(l)}(t)] \in \mathbb{R}^{d \times (n+b+u)} \quad (7)$$

$$\mathbf{W}_i^{(l)}(t) := [\mathbf{w}_{\mathcal{B}_i}^{(l)}(t)^\top | \mathbf{w}_{\mathcal{N}_i}^{(l)}(t)^\top | \mathbf{w}_{\mathcal{U}_i}^{(l)}(t)^\top]^\top \in \mathbb{R}^{(n+b+u)} \quad (8)$$

The proposed holonic learning employs a criteria-based iterative learning process that dedicates two separate procedures for horizontal and vertical contributions as presented in Algorithm 1. During its lifetime, holon $h_i^{(l)}$ uses lines 2 to 10 to continuously listen to new messages from the peers, including its subordinates (if non-terminal), neighbors, and superiors, in parallel and properly updates its knowledge matrices and weight vectors accordingly. In case the holon is training in sync with its peers, it also forwards any sent model from the super-holon to the neighbors. Concurrently, the holon follows a set of aggregation, updating, sharing steps until its termination criteria is satisfied. Function RECEIVEDENOUGHPEER-MODELS checks whether it has received expected number of models from its peers and if it has, it runs aggregation on the models and updates its knowledge base and initiates its training process if it is terminal. Next, through the steps between lines 22 and 25, the holon determines proper peers to share its train results. Choosing the peers to send the model depends on the training stage of the holon, its position in the holarchy, and the internal messaging policy of the super-holon it is a member of. Fig. 3 illustrates the steps of the algorithm in a very simple holarchy. It is worth noting that the synchronization of the communications in this example are merely for the sake of clarity and is not the requirement of the proposed model.

The above-mentioned formulation is defined generically in such a way that it supports the autonomy of holons and not enforcing any specific aggregation or communication schemes, it supports the dynamic aspects of the system, including communication

Algorithm 1: The generic learning procedure run by each holon $h_i^{(l)}$ in parallel.

```

1 Procedure Learn()
2   while  $C_i^{(l)} = \text{false}$  non-blocking in parallel do
3      $\theta_j, w_j, \text{sender} \leftarrow \text{RECEIVE}(\tilde{\theta})$ 
4      $\Theta_{\mathcal{U}_i}^{(l)}, w_{\mathcal{U}_i}^{(l)}, \Theta_{\mathcal{B}_i}^{(l)}, w_{\mathcal{B}_i}^{(l)}, \Theta_{\mathcal{N}_i}^{(l)}, w_{\mathcal{N}_i}^{(l)} \leftarrow \text{PROCESS}(\theta_j, w_j)$ 
5     if  $\text{sender} \in \mathcal{U}_i^l$  and SYNCHRONIZED then
6       for  $n \in \mathcal{N}_i^l$  in parallel do
7          $\text{SENDMODEL}(n, \theta_j)$ 
8       end
9     end
10  end
11  while  $C_i^{(l)} = \text{false}$  do
12    if RECEIVEDENOUGHPEERMODELS() then
13       $\theta_i^{(l)}(t) \leftarrow \mathcal{A}_i^{(l)}(\Theta_i^{(l)}, W_i^{(l)})$ 
14       $\text{UPDATEKB}(\Theta_{\mathcal{U}_i}^{(l)}, w_{\mathcal{U}_i}^{(l)}, \Theta_{\mathcal{B}_i}^{(l)}, w_{\mathcal{B}_i}^{(l)}, \Theta_{\mathcal{N}_i}^{(l)}, w_{\mathcal{N}_i}^{(l)})$ 
15    else
16       $\text{WAIT}()$ 
17    end
18    if  $h_i^{(l)}$  is terminal then
19       $\theta_i^{(l)}(t) \leftarrow \text{TRAIN}(\theta_i^{(l)}(t), \lambda)$ 
20       $\Theta_{\mathcal{B}_i}^{(l)} \leftarrow [\theta_i^{(l)}(t)]$ 
21    end
22     $\mathcal{P}_i^{(l)} \leftarrow \text{WHOTOSENDMODEL}()$ 
23    for  $p \in \mathcal{P}_i^{(l)}$  in parallel do
24       $\text{SENDMODEL}(p, \theta_i^{(l)}(t))$ 
25    end
26  end
27 end

```

graphs, memberships, peer weighting, and other evolving factors. To demonstrate the utilization of the proposed model, we present Holonic Averaging as a homogeneous special embodiment.

3.2 Holonic Averaging Learning

There is a wide spectrum of learning models that can be designed given the proposed generic HoL framework. This section delves into the details of a specific instance that exemplifies the principles of the presented model and scrutinizes its corresponding algorithms. The presented Holonic Averaging Learning (HoAL) mimics the aggregation mechanism used in the FAVG approach [17] and provides a distributed mechanism for a supervised learning problem. Below, we provide the breakdown of the details of HoAL's design components.

Holonic Structure. The holons in HoAL are heterogeneous with respect to their private data and homogeneous in terms of their learning models and aggregation functions. Additionally, we assume that each holon belongs to a unique super-holon. To put it

formally, $\forall i \neq j, \forall l, l' \in \mathbb{N}$ we have:

$$\mathcal{B}_i^{(l)} \cap \mathcal{B}_j^{(l')} = \emptyset, \quad \mathcal{A}_i^{(l)} = \mathcal{A}_j^{(l')}, \quad \mathcal{M}_i^{(l)} = \mathcal{M}_j^{(l')} \quad (9)$$

Communication Weights. As seen before, the weights that each holon assigns to the communication with its neighbors, subordinates, and superior allow for more nuanced control over the influence of other holons and can be dynamically defined based on metrics such as reliability, trustworthiness, data quantity, etc. Presuming the locality of the holons belonging the same connected graph $\mathcal{G}_i^{(l)}$ and that they have similar characteristics in terms of performance, data quality, etc., HoAL defines the inter- and intra-holon communication weights across different levels base on the size of the training data. Formally, let $w_{ij}^{(l)}$ denote the weight that holon $h_i^{(l)}$ uses for a peer $h_j^{(l')} \in \mathcal{B}_i^{(l)} \cup \mathcal{N}_i^{(l)} \cup \mathcal{U}_i^{(l)}$. We define $w_{ij}^{(l)} = d_j^{(l)}$, where:

$$d_j = \begin{cases} |\mathbb{X}_j^{(l')}| & \text{if } h_j^{(l')} \text{ is terminal} \\ \sum_{h_k^{(l'+1)} \in \mathcal{B}_j^{(l')}} d_k^{(l'+1)} & \text{if } h_j^{(l')} \text{ is non-terminal} \end{cases} \quad (10)$$

Please note that despite its recursive definition, the weight does not need to be calculated recursively each time it is needed. Each non-terminal holon keeps track of its training size whenever it receives updates from the subordinates.

Updates Stopping Criteria. Each holon $h_i^{(l)}$ in HoAL employs a numerical threshold $k_i^{(l)}$ to control the number of its local updates. To put it formally, we have:

$$C_i^{(l)} := t_i^{(l)} > k_i^{(l)} \quad (11)$$

Aggregation. As stated before, HoAL employs weighted averaging in each holon to aggregate the values of model parameters collected in its contribution matrix. Strictly speaking, we have:

$$\mathcal{A}_i^{(l)}(\Theta_i^{(l)}, W_i^{(l)}) := \frac{1}{W_i^{(l)\top} \mathbf{1}_z} (\Theta_i^{(l)} W_i^{(l)}) \quad (12)$$

where z denotes the size of vector $W_i^{(l)}$ at the time of aggregation, and $\mathbf{1}_z$ represents a z -dimensional column vector of ones. In this paper, we utilize mini-batch SGD as the optimizer and assume each training iteration equal to a specific number of epochs at the terminal level.

4 EXISTING MODELS AS SPECIAL CASES

The strength of holonic learning, in its original definition as presented in section 3.1, is its flexibility to encapsulate an expansive spectrum of distributed learning scenarios. An example that showcases its versatility is when we further investigate the structure of the HoAL. While being a specific instantiation of the HoL with some limitations, HoAL still extends the flexibility of several existing distributed approaches, such as federated averaging, hierarchical federated averaging, and peer-to-peer learning to name a few. Figure 4 visualizes a high-level example HoAL model accompanied by its particular embodiment to support the aforementioned standard models. Assuming SGD as the optimization method employed

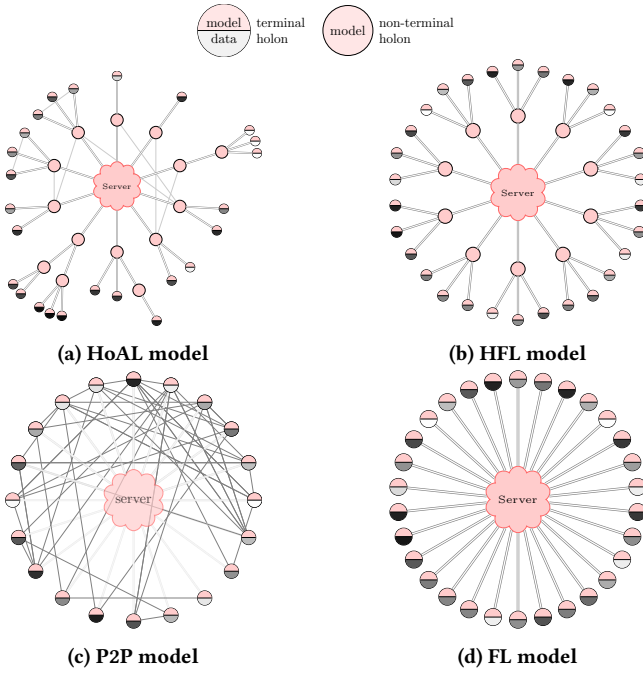


Figure 4: The holonic averaging learning model alongside different special cases implementing federated averaging, hierarchical federated averaging, and peer-to-peer models. Different shades in the lower half of the terminal holons represent the difference in their local data.

by the terminal holons, the depicted instantiations can further be discussed as follows:

Federated Averaging. : An FL system can be modeled by a one-level HoAL, where the server is modeled by the root holon $\mathbf{h}_1^{(0)}$ and each of the clients with terminal holon $\mathbf{h}_i^{(1)}$, $i = 1, \dots, n_1^{(0)}$. Additionally, $\mathcal{G}_1^{(0)} = \bar{K}_n$, where n and \bar{K}_n denote the number of clients and a null graph of size n . As depicted in Figure 4d, all clients use their own private data to learn the same ML model. At each update carried out by holons $\mathbf{h}_1^{(0)}$ and $\mathbf{h}_i^{(1)}$ the contribution matrix and weights vectors would respectively be:

$$\Theta_1^{(0)} = \Theta_{\mathcal{B}_1}^{(0)} = \left[\theta_1^{(1)} | \dots | \theta_{n_1^{(0)}}^{(1)} \right] \quad (13)$$

$$\mathbf{W}_1^{(0)} = \mathbf{w}_{\mathcal{B}_1}^{(0)} = \left[|\mathbb{X}_1^{(1)}|, \dots, |\mathbb{X}_{n_1^{(0)}}^{(1)}| \right]^\top \quad (14)$$

$$\Theta_i^{(1)} = \theta_i^{(1)} = \theta_i^{(1)} - \eta_i^{(1)} \left[\frac{1}{|\mathbb{X}_i^{(1)}|} \sum_{j=1}^{|\mathbb{X}_i^{(1)}|} \nabla \mathcal{F}_i^{(1)}(\theta; \xi^{(j)}) \right] \quad (15)$$

$$\mathbf{W}_i^{(1)} = \left[|\mathbb{X}_i^{(1)}| \right]^\top \quad (16)$$

Peer-to-Peer Learning. : Similar to FL, a P2P learning model can also be modeled by a one-level HoAL, as depicted in Figure 4c. A key distinction her is that the communication graph connecting all the terminal holons, i.e. $\mathcal{G}_1^{(0)}$ is connected. Additionally, as there

is no need for server-side aggregation of the learned models, the root $\mathbf{h}_1^{(0)}$, being connected to at least one client, only serves as the initiator of the learning process. Once initiated, the model updates of each holon $\mathbf{h}_i^{(1)}$ will be calculated using the following contribution matrix and weight vector configurations:

$$\Theta_i^{(1)} = \left[\theta_i^{(1)} | \Theta_{\mathcal{N}_i}^{(1)} \right] = \left[\theta_i^{(1)} - \frac{\eta_i^{(1)}}{|\mathbb{X}_i^{(1)}|} \sum_{j=1}^{|\mathbb{X}_i^{(1)}|} \nabla \mathcal{F}_i^{(1)}(\theta; \xi^{(j)}) \mid \theta_e^{(1)} - \frac{\eta_e^{(1)}}{|\mathbb{X}_e^{(1)}|} \sum_{j=1}^{|\mathbb{X}_e^{(1)}|} \nabla \mathcal{F}_e^{(1)}(\theta; \xi^{(j)}) \mid \dots \mid \theta_f^{(1)} - \frac{\eta_f^{(1)}}{|\mathbb{X}_f^{(1)}|} \sum_{j=1}^{|\mathbb{X}_f^{(1)}|} \nabla \mathcal{F}_f^{(1)}(\theta; \xi^{(j)}) \right] \quad (17)$$

$$\mathbf{W}_i^{(1)} = \left[|\mathbb{X}_i^{(1)}|, \mathbf{w}_{\mathcal{N}_i}^{(1)} \right]^\top = \left[|\mathbb{X}_i^{(1)}|, |\mathbb{X}_e^{(1)}|, \dots, |\mathbb{X}_f^{(1)}| \right]^\top \quad (18)$$

Hierarchical Federated Learning. : The client-edge-server hierarchical federated learning can be modeled by a two level balanced holarity, in which the root holon $\mathbf{h}_1^{(0)}$ serves as the cloud, the non-terminal holons $\mathbf{h}_i^{(1)}$, $i = 1, \dots, n_1^{(0)}$ serves as the edge model-aggregators, and finally the terminal holon $\mathbf{h}_j^{(2)}$, $j = 1, \dots, n_i^{(1)}$ represent the clients employed to learn optimize the models using their own data. Additionally, all the communication graphs of the holons are null, i.e. $\mathcal{G}_1^{(0)} = \bar{K}_{n_1^{(0)}}$ and $\mathcal{G}_i^{(1)} = \bar{K}_{n_i^{(1)}}$, $i = 1, \dots, n_1^{(0)}$. Similar to the above-mentioned federated averaging case, the contribution matrices and the weight vectors would be as follows:

$$\Theta_1^{(0)} = \Theta_{\mathcal{B}_1}^{(0)} = \left[\theta_1^{(1)} | \dots | \theta_{n_1^{(0)}}^{(1)} \right] = \left[\mathcal{A}_1^{(1)} \left(\Theta_1^{(1)}, \mathbf{W}_1^{(1)} \right), \dots, \mathcal{A}_{n_1^{(0)}}^{(1)} \left(\Theta_{n_1^{(0)}}^{(1)}, \mathbf{W}_{n_1^{(0)}}^{(1)} \right) \right], \quad (19)$$

$$\mathbf{W}_1^{(0)} = \mathbf{w}_{\mathcal{B}_1}^{(0)} = \left[\sum_{\mathbf{h}_j^{(2)} \in \mathcal{B}_1^{(1)}} |\mathbb{X}_j^{(2)}|, \dots, \sum_{\mathbf{h}_j^{(2)} \in \mathcal{B}_{n_1^{(0)}}^{(1)}} |\mathbb{X}_j^{(2)}| \right]^\top \quad (20)$$

$$\Theta_i^{(1)} = \Theta_{\mathcal{B}_i}^{(1)} = \left[\theta_1^{(2)} | \dots | \theta_{n_i^{(1)}}^{(2)} \right] \quad (21)$$

$$\mathbf{W}_i^{(1)} = \mathbf{w}_{\mathcal{B}_i}^{(1)} = \left[|\mathbb{X}_1^{(2)}|, \dots, |\mathbb{X}_{n_i^{(1)}}^{(2)}| \right]^\top \quad (22)$$

$$\Theta_i^{(2)} = \theta_i^{(2)} = \theta_i^{(2)} - \eta_i^{(2)} \left[\frac{1}{|\mathbb{X}_i^{(2)}|} \sum_{j=1}^{|\mathbb{X}_i^{(2)}|} \nabla \mathcal{F}_i^{(2)}(\theta; \xi^{(j)}) \right] \quad (23)$$

$$\mathbf{W}_i^{(2)} = \left[|\mathbb{X}_i^{(2)}| \right]^\top \quad (24)$$

5 EXPERIMENTS

This section details our experiments investigating the convergence and performance of HoAL under various design scenarios and data distributions. The HoAL framework was implemented using pytorch along with its multiprocessing and distributed libraries¹. We have utilized training loss and test accuracy as the measure of learning convergence and performance.

5.1 Settings

The experimental setup comprised 10 terminal holons (agents) organized into four distinct holarchical structures (see Figure 5). These structures were chosen based on their shared characteristic of utilizing a collaboration network among terminal holons, representing various facets of collaborations. For instance, HoAL1P resembles a P2P learning setting, primarily relying on inter-holon communications. In contrast, the remaining three structures differ in the number of holarchical levels while confining collaborations to peers within the holons at each level.

As previously discussed, each terminal holon possesses its own personalized machine learning models and dataset. While the model parameters are shared through inter-holon collaborations, ensuring uniform architecture, the datasets remain private. These datasets are utilized by their respective holons for local model training and parameter updates.

Throughout the experiments, personal data for terminal holons and the collaboration network at the terminal level remained constant. We utilized the standard MNIST handwritten digits classification dataset, with all terminal holons equipped with a Convolutional Neural Network Model as in [17] initialized with the same values for its 21,840 trainable parameters. Terminal holons used mini-batch Stochastic Gradient Descent (SGD) with a batch size of 32 and an initial learning rate of 0.01. Both terminal and non-terminal holons utilized separate CPU processes for training/aggregating their models in parallel. This deliberate implementation aims to simulate cross-device decentralization and the independence of holons. Furthermore, the holons in this implementation utilize threading to manage communications with their peers through the Gloo backend. It is important to emphasize that the proposed framework is not bound to any specific implementation technologies; hence, its behavior should remain consistent as long as there is a proper utilization of such technologies.

Data distribution, a pivotal factor in distributed machine learning, was evaluated in three settings: IID, EqNIID, and UEqNIID. In EqNIID, each terminal holon possessed 500 samples from 2 class labels, while in UEqNIID, terminals had varying numbers of samples from sorted groups.

Collaboration schemes varied across holarchy levels: each terminal holon underwent 250 rounds of a 5-epoch budget, collaborating with peers and superordinates at the end of each round. Non-terminal holons, on the other hand, had a 2-round local budget for aggregations and guiding subordinates before communicating with their own superordinates. Non-terminal holons budget is reset after each update from their superordinates. To make comparisons fair, we made the duration of experiments be controlled by terminal holons. That is, each terminal holon quits as soon as its budget of

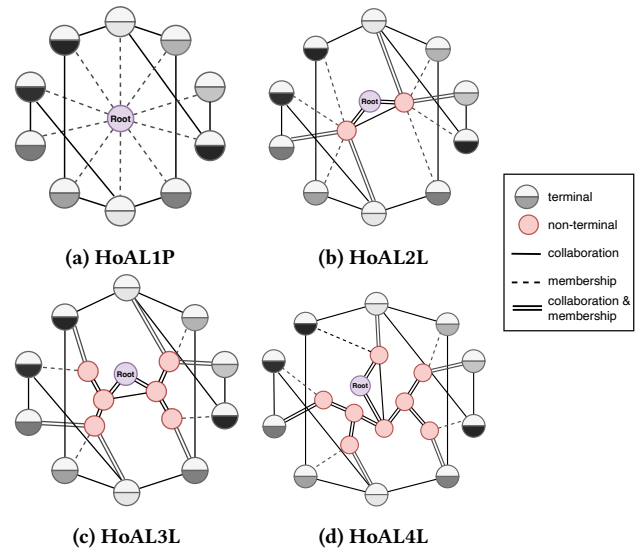


Figure 5: The different holarchical HoAL structures used for experiments.

200 rounds is exhausted, and the non-terminal holons quit once all their subordinates finished their job.

5.2 Results

In conducting an empirical analysis of HoL’s convergence, as shown in Figure 6, we observed consistent behavior across various holarchies and data distribution settings. Figures 6a and 6b present a comparative view of the training loss and test accuracy, respectively, for the experimented holarchies on IID data distribution. These metrics are averaged over all terminal holons. As seen, the holarchical structure has yielded relatively improved performance, in both training loss and test accuracy, compared to the plain usage of collaboration network among terminal holons as represented by HoAL1P model.

This trend persists under the UEqNIID data distribution, as demonstrated in Figures 6c and 6d. Here, the terminal holons exhibit a similar pattern of enhanced performance compared to the collaboration network model, though with comparatively smaller improvements than the IID case. The nuances observed in these results indicate the insensitivity of the models to unbalanced data distribution. Such resilience to uneven data distributions further underscores the potential robustness of the HoAL learning framework.

In the EqNIID setting, illustrated in Figures 6e and 6f, we persistently observe similar behavior, albeit with more noticeable distinctions in the performance of the examined structures. Unlike the IID setting, the introduction of more holarchical levels, as in HoAL4L and HoAL3L cases, has led to a degradation in the generalization capability of terminal holons on average. This observed phenomenon can be primarily attributed to the collaboration delays enforced by extra levels of the holarchy.

It is worth highlighting that the evolution of performance on higher-level holons is intentionally not included for two main reasons: (1) due to the guiding role of non-terminal holons at higher

¹Code available at: <https://github.com/aesmaeli/HoAL>

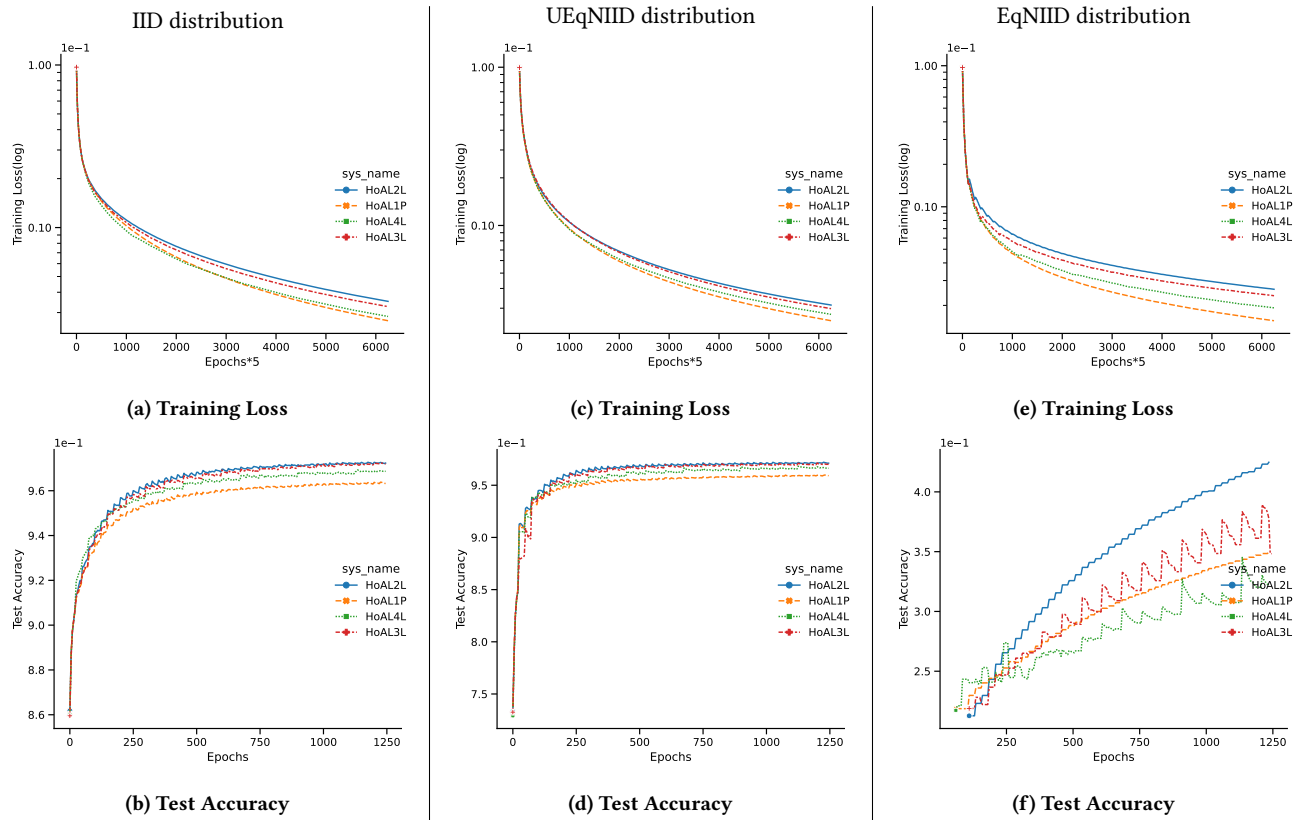


Figure 6: The average performance of all terminal holons of different holarchical structures trained on MNIST dataset. Each column corresponds to training and test performances on distinct data distributions, including IID, unequal Non-IID (UEqNIID), and equal Non-IID (EqNIID).

levels and their lack of a learning process and independent dataset, which would result in redundantly mirroring the convergence already reported for terminal holons; (2) the varied compositions of higher-level holons and the use of different holarchical structures, which would render their performance incomparable and limit meaningful insights. This is not meant to diminish the significance of such comparisons, but to align with our specific objectives of showcasing the convergence of inter-holon collaborations and highlighting the impact of holarchical structures on relative improvements.

While we acknowledge that the presented experimental settings are on a small scale, limiting our ability to fully characterize the extensive potential of the introduced HoL framework, the observed trends and insights serve as a foundational exploration. These initial findings provide a compelling basis for further, larger-scale investigations on the capabilities and robustness of the HoL framework in diverse and complex real-world scenarios.

6 CONCLUSION

This paper introduced Holonic Learning (HoL), an agent-based distributed machine learning framework inspired by the self-similar holonic paradigm. Within HoL, the self-contained holons are responsible for managing individual datasets and model aggregation

processes, with the unique characteristic that lies in its collaborative synergy across the holarchy, both horizontally and vertically. This holistic approach enables each holon to maintain autonomy over local data and aggregation strategies, facilitating the integration of diverse policies and learning strategies at different levels. We also presented HoAL, a homogeneous instantiation with a weighted parameter averaging aggregation strategy and demonstrated its extending flexibility compared to existing distributed approaches like federated learning. Empirical investigation revealed HoAL’s convergence and effectiveness in various design and data distribution settings, providing insights into the applicability of HoL.

As we envision the future of distributed machine learning, Holonic Learning stands out by harmonizing individual agency with collaborative intelligence, and leveraging holonic multi-agent systems, it offers promising potential for scalable, adaptable, privacy-focused, and heterogeneous learning environments. This paper’s empirical results scrutinize HoAL’s behavior, emphasizing the pivotal role of holarchical design in convergence rate and generalization, even in small-scale settings. In contemplating the next steps, conducting theoretical investigation of HoL’s convergence in both convex and non-convex optimization settings and integrating the gained insight into the design of large-scale holarchies are crucial for its real-world applications.

REFERENCES

- [1] Mehdi Salehi Heydar Abad, Emre Ozfatura, Deniz Gunduz, and Ozgur Ercetin. 2020. Hierarchical federated learning across heterogeneous cellular networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8866–8870.
- [2] José Barbosa. 2015. *Self-organized and evolvable holonic architecture for manufacturing control*. Ph.D. Dissertation. Université de Valenciennes et du Hainaut-Cambresis.
- [3] Abla Chaouni Benabdellah, Imane Bouhaddou, and Asmaa Benghabrit. 2021. SmartDFRelevance: a holonic agent based system for engineering industrial projects in concurrent engineering context. In *Artificial Intelligence and Industrial Applications: Smart Operation Management*. Springer, 103–123.
- [4] Fritjof Capra. 1984. *The turning point: Science, society, and the rising culture*. Bantam.
- [5] Ahmad Esmaeili, John C. Gallagher, John A. Springer, and Eric T. Matson. 2022. HAMLET: A Hierarchical Agent-Based Machine Learning Platform. *ACM Trans. Auton. Adapt. Syst.* 16, 3–4, Article 9 (jul 2022), 46 pages. <https://doi.org/10.1145/3530191>
- [6] Ahmad Esmaeili, Nasser Mozayani, Mohammad Reza Jahed Motlagh, and Eric T. Matson. 2017. A Socially-Based Distributed Self-Organizing Algorithm for Holonic Multi-Agent Systems. *Cogn. Syst. Res.* 43, C (2017), 21–44. <https://doi.org/10.1016/j.cogsys.2016.12.001>
- [7] Klaus Fischer, Michael Schillo, and Jörg Siekmann. 2003. Holonic multiagent systems: A foundation for the organisation of multiagent systems. In *International conference on industrial applications of holonic and multi-agent systems*. Springer, 71–80.
- [8] Christian Gerber, Jörg Siekmann, and Gero Vierke. 1999. *Holonic multi-agent systems*. Technical Report. German Research Center for Artificial Intelligence (DFKI).
- [9] Otkrist Gupta and Ramesh Raskar. 2018. Distributed learning of deep neural network over multiple agents. *Journal of Network and Computer Applications* 116 (2018), 1–8.
- [10] Vincent Hilaire, Abder Koukam, and Sebastian Rodriguez. 2008. An Adaptive Agent Architecture for Holonic Multi-Agent Systems. *ACM Trans. Auton. Adapt. Syst.* 3, 1, Article 2 (mar 2008), 24 pages. <https://doi.org/10.1145/1342171.1342173>
- [11] Seyedali Hosseinalipour, Sheikh Shams Azam, Christopher G Brinton, Nicolo Michelusi, Vaneet Aggarwal, David J Love, and Huaiyu Dai. 2022. Multi-stage hybrid federated learning over large-scale D2D-enabled fog networks. *IEEE/ACM transactions on networking* 30, 4 (2022), 1569–1584.
- [12] Zhanhong Jiang, Aditya Balu, Chinmay Hegde, and Soumik Sarkar. 2017. Collaborative Deep Learning in Fixed Topology Networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 5906–5916.
- [13] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.
- [14] Arthur Koestler. 1968. *The ghost in the machine*. Macmillan.
- [15] Arthur Koestler. 1978. *Janus: A Summing Up*. Random House.
- [16] Lumín Liu, Jun Zhang, SH Song, and Khaled B Letaief. 2020. Client-edge-cloud hierarchical federated learning. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*. IEEE, 1–6.
- [17] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [18] Omer Rana, Theodoros Spyridopoulos, Nathaniel Hudson, Matt Baughman, Kyle Chard, Ian Foster, and Aftab Khan. 2023. Hierarchical and Decentralised Federated Learning. *arXiv preprint arXiv:2304.14982* (2023).
- [19] Nguyen H Tran, Wei Bao, Albert Zomaya, Minh NH Nguyen, and Choong Seon Hong. 2019. Federated learning over wireless networks: Optimization model design and analysis. In *IEEE INFOCOM 2019-IEEE conference on computer communications*. IEEE, 1387–1395.
- [20] Mihaela Uliuru. 2002. Emergence of holonic enterprises from multi-agent systems: A fuzzy-evolutionary approach. *Soft Computing Agents: A new perspective on dynamic information systems* (2002), 187–215.
- [21] PAUL Valckenaers, Hendrik Van Brussel, Luc Bongaerts, and Jozef Wyns. 1997. Holonic manufacturing systems. *Integrated Computer-Aided Engineering* 4, 3 (1997), 191–201.
- [22] Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. 2017. Decentralized Collaborative Learning of Personalized Models over Networks. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, 509–517. <https://proceedings.mlr.press/v54/vanhaesebrouck17a.html>
- [23] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. 2018. Split learning for health: Distributed deep learning without sharing raw patient data. *arXiv preprint arXiv:1812.00564* (2018).
- [24] Shiqiang Wang, Tiffany Tuor, Theodoros Salonidis, Kin K Leung, Christian Makaya, Ting He, and Kevin Chan. 2019. Adaptive federated learning in resource constrained edge computing systems. *IEEE journal on selected areas in communications* 37, 6 (2019), 1205–1221.
- [25] Bo Xu, Wenchao Xia, Wanli Wen, Pei Liu, Haitao Zhao, and Hongbo Zhu. 2021. Adaptive hierarchical federated learning over wireless networks. *IEEE Transactions on Vehicular Technology* 71, 2 (2021), 2070–2083.
- [26] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 2 (2019), 1–19.