# Holonic Active Distillation for Scalable Multi-Agent Learning in Multi-Sensor Systems

Dani Manjah<sup>1</sup><sup>(0)</sup>, Tim Bary<sup>1</sup><sup>(0)</sup>, Benoit Macq<sup>1</sup><sup>(0)</sup>, and Stéphane Galland<sup>2</sup><sup>(0)</sup>

<sup>1</sup> Institute of Information and Communication Technologies, Electronics and Applied Mathematics (ICTEAM), UCLouvain, 1348 Louvain-la-Neuve, Belgium {dani.manjah,benoit.macq}@uclouvain.be

<sup>2</sup> Université de Technologie de Belfort Montbéliard, UTBM, CIAD UR 7533, F-90010 Belfort cedex, France

stephane.gallandQutbm.fr

Abstract. The rapid expansion of sensor-based networks introduces major challenges in scalability, adaptability, and knowledge transfer, especially in open environments where new subsystems can dynamically join or leave. In this work, we propose a Holonic Active Distillation architecture within a Holonic Multi-Agent System (HMAS) to address these issues. Our approach integrates Clustered Stream-Based Active Distillation (CSBAD), a framework in which specialized student models collect local data, query pseudo-labels from teacher models, and cluster into groups of similar sensors.

Results show that the holonic organization balances local specialization with global generalization, while efficiently adapting to sensor departures and re-integrations. We also analyzed trade-offs among incremental model updates, system reorganization, and scalability limits.

Our findings highlight the advantages of holonic learning for multi-sensor systems while identifying key challenges related to model drift and longterm adaptation.

**Keywords:** Holonic Multi-Agent Systems · Distributed Learning · Collaborative Learning · Scalable Model Adaptation

# 1 Introduction

In recent years, sensor systems have evolved from isolated and manageable units to expansive and interconnected networks [38]. This transformation, while enabling broader coverage, challenges the traditional approach of deploying a single, universal Deep Neural Network (DNN) across all sensors [31]. The dynamic nature of real-world deployments, characterized by stochastic changes and the continuous addition of new sensors, introduces diverse contexts that demand ever larger training datasets and models [26]. This upscaling not only incurs significant costs, but also accumulates hidden technical debt, complicating the maintenance required to adapt to distribution shifts in sensor data [2, 42].

Traditional methods have relied on centralized, monolithic DNNs that struggle to scale with the increasing complexity and diversity of sensor networks.



Fig. 1: A large-scale distributed training system where DNN nodes are trained on datasets built from a specificity-diversity trade-off for effective learning. Lower nodes are typically tailored to their task (*i.e.*, analytics on a sensor) and operationally more efficient. Higher nodes, trained over vaster and more diverse data, provide generalization ability. This system can scale up or down, causing challenges in integration and adaptability.

These approaches often lead to inefficiencies and increased maintenance challenges. Recent advances, such as the Holonic Learning (HoL) framework, offer a promising alternative by embracing the agent paradigm to improve scalability and flexibility [13]. HoL leverages self-nested structures of agents, known as holons, to integrate local and global perspectives, facilitating easier subsystem integration and preventing the propagation of disturbances [9, 27, 41]. This hierarchical learning approach improves the efficiency of data and algorithm handling, particularly for large distributed datasets [11]. A question remains related to the design of a scalable distributed learning framework that supports continuous DNN refinement with minimal refactoring, while allowing each unit to update itself online, self-organize with its peers, and transfer knowledge as the system scales.

This work builds upon the HoL framework to introduce strategies for aggregation, communication, and commitment between learning holons. The contributions are twofold:

- We augment HoL with standardized organizations and roles, allowing newcomers to integrate knowing only their role and the associated protocol, inspired by active learning and distillation of knowledge [1, 4, 24, 41].
- We propose holonification mechanisms in which agents cluster horizontally and vertically based on the similarity of their sensor streams, preserving

confidentiality and balancing specificity and diversity. This recursive process improves the accuracy and resilience of the system, allowing seamless expansion or contraction without disruption [7, 14, 15, 18, 28, 32].

The remainder of this paper is structured as follows. Section 2 discusses the current landscape of learning in Multi-Agent System (MAS). Section 3 presents the high-level formulation of the organizations for the Holonic Active Distillation (HAD) architecture and the relationships between super- and subholons. Section 4 introduces a distributed, multi-tiered, self-nested structure for DNNs and the self-organization mechanisms that realize the specificity-diversity trade-off. Section 5 describes the materials used in the experiment, while Section 6 reports the results on sensor addition and removal, comparing partial reorganizations with complete retraining and measuring knowledge-transfer speed. Finally, Section 7 discusses the insights, limitations, and perspectives, followed by the conclusion in Section 8. This research advances the adaptability of multisensor holonic systems, materializing continuous learning in dynamic environments [23, 33, 45].

# 2 Related Works

Adaptive MAS networks leverage online learning strategies to dynamically respond to environmental changes, highlighting the importance of distributed and collaborative learning [35].

Wolpert and Macready [48] introduce a system utilizing reinforcement learning to align agent actions with collective goals, minimizing human oversight. Agents are organized into "sub-worlds" for focused collaboration, yet the application of reinforcement learning in complex scenarios with varied sensors and methods encounters obstacles such as unclear rewards and limited exploration, which hinders the required diversity of learning [39, 48, 51].

Organizational learning considers agents evolving through both personal and collective learning efforts, enhancing agents' abilities in MAS through management mediated interactions and task alignment to boost system efficiency [19, 44]. This model emphasizes the role of knowledge sharing in improving workflows and establishing structural knowledge, crucial for system resilience. Social science research [10] reflects on the applicability of this framework to understand the impact of staff turnover on management, analogous to agent dynamics in open MAS.

Hierarchical learning [12] uses hierarchical MAS to streamline Machine Learning (ML) training in various geographical locations. By modeling challenges as a hypergraph, the system organizes agents, each with unique skills and knowledge, into a structured multitiered network. This design not only facilitates the decentralized handling of ML algorithms and data, but also significantly improves the efficiency and scalability of processing large distributed datasets.

In the context of distributed ML, Gupta and Raskar [22] pioneered Federated Learning (FL) to train neural networks in distributed datasets, prioritizing data privacy and computational efficiency. However, FL faces hurdles in communication and training reliability. Hierarchical FL addresses these by grouping users to improve FL security and efficiency through group-specific updates [50]. Personalized FL [30] methods aim to produce personalized models for different users or groups of users [20] to keep track of their individualized requirements. Hierarchy has also been instrumental in Fog Learning. Unlike FL, which is based on a star topology of device-server interactions, Fog Learning explicitly considers the network and topology structures among devices and enables intelligent device collaborations through data and parameter offloading [25].

Esmaeili et al. [13] abstract FL with HoL, applying holonic principles to a collaborative learning framework. In that sense, FL can be seen as a first-order HoL. HoL enhances model cooperation with specific strategies for aggregation, communication, and commitment within holons, facilitating complex yet intuitive collaboration of nodes compared to Fog Learning. In this balance between local autonomy and coordinated decision making, holonic systems are better equipped to tackle challenges such as adaptability, and scalability.

HoL does not specify how learning agents should (re)organize, nor how a system can seamlessly expand to new domains or safely unlearn obsolete ones—shortcomings that become acute in applications requiring auto-scaling and auto-tuning [13, 47].

We address these limitations by enriching the holonic-agent paradigm with organizations and roles. These concepts provide an abstract interaction pattern that improves the architecture's robustness and flexibility [1]. Furthermore, our *Teacher–Student* distillation scheme supplies pseudo-labels that swiftly calibrate incoming *Students*, thereby enabling auto-tuning. Finally, self-organization emerges from a specificity–diversity trade-off among *Students*, while integration and deletion protocols dynamically scale sensor subsystems. Collectively, these mechanisms yield the first HoL variant that supports self-organization and auto-scaling.

# 3 Holonic Active Distillation (HAD) Architecture

We seek a design that minimizes refactoring and supports isolated updates, simplifying the integration of a new sensor, model, or more generally a subsystem of sensors and models [17, 42]. From the literature, we derived five main recommendations to **design scalable**, **multimethod learning systems**:

- 1. Establish standardized interaction protocols, aggregation strategies, commitment, and communication patterns within components. This facilitates the integration of new units, as they only need to understand their role and communication methods within the system, regardless of their operating mode [1, 41].
- 2. Render a method, a sensor, or by construction a subsystem as independent and self-contained as possible to limit the complexity between units. This aims to simplify a local update or maintenance [1, 21, 46].
- 3. Recursively divide a system into subsystems based on a key criterion. This prevents the propagation of disturbances [41]. Furthermore, integration of a

new component requires less communication as it requires only coordination with the upper layers of the system instead of with each subsystem [9].

- 4. Place units at certain levels of the hierarchy and provide representations of how other levels can contribute "information" or "models". This division simplifies the complexity of programming, allowing designers to focus on each module and facilitating reuse between different systems [9].
- 5. Exploit Active Distillation, where each Student unit collects data on the fly from its streams to train on them. Training is performed by querying a model Teacher [31].

Many agent-oriented organizational methodologies have been developed such as ADELFE [3, 21], ASPECS [8], Gaia [49], INGENIAS [37], SODA [36], or Tropos [6]. Each has its own focus: ADELFE on adaptive systems and cooperative agents, ASPECS on holonic multi-agent systems, Gaia on static organization and roles, and SODA on the environment.

Given our choice of the organizational holonic paradigm, we adopt the AS-PECS methodology [8]. The latter starts by defining an *Organization*, which denotes a subsystem in which components play a role and interact to achieve a shared goal in the context of this organization. Next, the *Roles* which are both expected behaviors to fulfill (part of) requirements, and status to the role's agent in the organization (Section 3.1). The subsequent activity (Section 3.2) is the definition of relationships between superholons (higher-level entities) and subholons (lower-level entities). As a reference later, a *holarchy* denotes the hierarchy of self-regulating holons.

#### 3.1 Teacher-Specialized Student (TSS)

Building on the Active Distillation framework and the specificity-diversity tradeoff from [32], we developed an organizational model that incorporates the roles of *Specialized Student* and *Teacher*, as illustrated in Fig. 2.

The **Specialized Student** role is designed to continuously collect data on subparts of the system's deployment environment. Under the oversight of a higher-order **Teacher** entity, these *Students* learn from these data, adapting their models' weights accordingly.

#### 3.2 Holarchy

The section begins with introducing a new notation. Then we present an example of a three-tiered holarchy structure. Each level of this holarchy is a possible instance of an organization defined in Section 3.1. To provide a more holistic perspective, we depict the Cyber-Physical Platform (CPP) data processing organization (see our previous work [33]) alongside the TSS but CPP is not the main focus of this paper<sup>3</sup>.

<sup>&</sup>lt;sup>3</sup> As a more detailed context, *CPP* is designed to respond to external requests with perceptions and to manage its finite resources to ensure fair access across multiple



Fig. 2: Organizational model of the *Teacher-Student*, using the ASPECS notation [8]. The **Specialized Student** role involves a component tasked with building expertise over a delineated sub-domain in the system, *i.e.*, a regional distribution. The *Teacher* role supervises the learning processes of the *Students*.

**Notations** A holarchy  $\mathcal{H}^L_O$  includes up to L vertical layers instantiating an organization O. A holon *i* in layer *l*, where *l* ranges from 0 to *L*, is denoted by  $\hbar_i^l$  and comprises:

- $\begin{array}{l} \ \mathbf{X}_{i}^{l} : \text{Set of operating data streams of a holon } \hbar_{i}^{l} . \\ \ \mathcal{T}_{i}^{l} : \text{Training set of a holon } \hbar_{i}^{l} . \\ \ \mathcal{V}_{i}^{l} : \text{Validation set of a holon } \hbar_{i}^{l} . \end{array}$

- $\theta^{\hbar_i^l}$ : Processing model of a holon  $\hbar_i^l$ .
- $\mathsf{SUB}_i^l$ : Inner members corresponding to layer l-1 of a holon  $\hbar_i^l$ .
- $SUP_i^l$ : Superior holon of a holon  $\hbar_i^l$ .

Multi-Scale Hierarchical Architecture The system architecture, shown in Fig. 3, includes two holarchies:  $\mathcal{H}^3_{CPP}$  that processes data on three levels and  $\mathcal{H}^2_{TSS}$  managing knowledge on two levels.

- At level 0: agents are the primary functional layer. They employ models designed for specific data streams. Proximity to other agents, geographically or related to the task, allows them to merge outputs and reduce errors. For example,  $\hbar_1^0$  and  $\hbar_2^0$  form Group G1 to fuse their outputs to feed the data request of a higher-order holon  $\hbar_1^1$ .

However, agents monitoring the same area may employ different models if their functions require learning different features. Consequently,  $\hbar_2^0$  and  $\hbar_3^0$ , as *Specialized Students*, form Group G4 to learn a shared model under the supervision of a *Teacher* holon via Group G5.

- At levels 1 and 2: higher levels above the agents integrate and synthesize data from specific areas of the system (e.g., data streams that share

surveillance operations. The **Resource Provider** role ensures a fair distribution of the resources among all parties. The **Observer** role has the ability to produce perceptions thanks to the data acquired by the *Sensor* role. The data acquisition could be based on another CPP.



Fig. 3: Holonic architecture inspired by the "cheese board" notation [8, 16]. Each level represents a different hierarchical position, defining both the semantic level of data and the degree of knowledge specialization. On the left, the  $\mathcal{H}^3_{CPP}$  instantiates the *CPP* organization, and on the right, the  $\mathcal{H}^2_{TS}$  is responsible for active learning. Agents may assume multiple roles and participate in multiple holarchies simultaneously.

attributes). The holons in the role **Observer**, such as  $\hbar_1^1$  and  $\hbar_1^2$ , elevate the collected data to a new semantic level.

Holon  $\hbar_2^1$ , a higher order **Specialized Student**, aggregates validation sets from  $\hbar_2^0$  and  $\hbar_3^0$  (*i.e.*,  $\mathcal{T}_2^1 = \mathcal{V}_2^0 \cup \mathcal{V}_3^0$ ), to create a broader and generalized model.

Generally, each semantic level consolidates knowledge across broader areas of the system, fostering a holistic view, such as city-scale tracking. Meanwhile, intermediate layers consisting of *Specialized Students* synthesize knowledge from lower levels, to deepen collective task understanding and increase holons' universality.

# 4 Holonic Learning Framework

This section presents mechanisms to create a multilevel learning framework. Next, it introduces a mechanism to incorporate new nodes by coordinating with the top layers and assigning each new node to the group whose DNN model is most accurate in its data stream.

#### 4.1 Holonification

In the holonic terminology, *holonification* is the process of grouping agents into a holarchy, resembling complex clustering based on criteria like capabilities and resource access [14, 15].

In this Section, we propose a multi-tiered learning structure (illustrated in Fig. 1), comprising a portfolio of models that range from sensor-specific to universal, deployable across the entire network. Specifically, upper-layer models are trained on larger datasets for broader coverage, while lower-layer models use smaller, more similar datasets for increased specificity. Having intermediate models at various levels of granularity not only ensures adaptability, **but also supports robust knowledge organization.** For example, city-wide vehicle detection may require multiple models specializing in certain domain representations [32]. However, these domain-specific models **benefit from interactions with peer models or a more fundamental model** that develops a fundamental understanding of object detection tasks [52]. Our agent-based modeling offers this flexibility to develop these vertical and horizontal interactions.

Formally, each holon in a layer l > 0 is allocated a budget  $B^l = B_0 \cdot 10^l$ , where  $B_0$  represents the number of images used for model fine-tuning. This budget limits the training of each layer to at most  $10^l$  from the preceding levels, ensuring that the size of the data set of any holon  $\hbar_i^l$  does not exceed  $B^l$ , *i.e.*,  $|\mathcal{T}_i^l| \leq B^l$ .

To merge holons, we adopt the premise from Manjah et al. [32] that models with similar performance have learned from comparable data.

The remainder of this section describes the holonification process.

STEP 1 – Cross-Performance Vector. Assuming holons can transfer their model weights to each other within the same layer. Each holon  $\hbar_i^l$  computes a performance vector  $P_i$  by evaluating the effectiveness of models from other holons and itself in the same layer on its own validation data  $\mathcal{V}_i^l$ , according to Equation 1.

$$P_i^l := \left[ f\left(\theta_1^l, \mathcal{V}_i^l\right); \cdots; f\left(\theta_{N^l}^l, \mathcal{V}_i^l\right) \right]^T \tag{1}$$

Where  $\theta_j$  denotes the model parameters of the *j*-th holon  $\hbar_j^l$ ,  $j \in 1, \dots, N^l$  and  $f(\theta, \mathcal{V})$  the score of a model performance  $\theta$  against a validation set  $\mathcal{V}$ .

STEP 2 – Pair-Wise Distance Computation. To quantify the differences between models trained in different domains, holons broadcast their cross-performance vectors  $P_i^l$  defined in Equation 1 and compute a pairwise distance between their performance in the datasets and the performance of the other holons. Generally, for a holon  $\hbar_i^l$ , the distance to a holon  $\hbar_j^l$  is given by Equation 2.

$$D_i(\hbar_j^l) = \sqrt{\sum_{k=1}^{N^l} (P_{ik}^l - P_{jk}^l)^2}$$
(2)

STEP 3 – Agglomerative Merging using Single Linkage. The merging of the holons is an iterative process. The set of holons  $\{\hbar_1^l, \dots, \hbar_{N^l}^l\}$  creates a higher-order holarchy  $\mathcal{H}^{l+1}$  to which they belong.

At each iteration, the set of holons  $\hbar^l$  transmits their smallest linkage distance. This is defined as the minimal distance between the inner members of the holons. Formally, for two holons  $\hbar_A^l$  and  $\hbar_B^l$ , the single link distance  $L(\hbar_A^l, \hbar_B^l)$  is given by Equation 3.

$$L(\hbar_A^l, \hbar_B^l) = \min\{D_{ij} : \hbar_i^{l-1} \in \mathsf{SUB}(\hbar_A^l), \ \hbar_j^{l-1} \in \mathsf{SUB}(\hbar_B^l)\}$$
(3)

After all linkages are evaluated, the pair with the smallest distance merges, involving a combination of their datasets. After merging, the set of holons has decreased,  $\{\hbar_1^l, \dots, \hbar_{N^l-1}^l\}$ , and the linkage distances are updated for all agents.

The process ends if there remains only one holon or if the previous merge leads to a holon with a dataset size that exceeds  $B^{l+1}$ . In the second scenario, the process goes back to STEP 1 for the set  $\{\hbar_1^{l+1}, \dots, \hbar_{N^{l+1}}^l\}$ .

*STEP 4* – *Model Training.* The final steps consist in training the cluster models on the aggregated data sets.

# 4.2 Domain Integration Process

A new holon  $\hbar_+$  joins a holarchy  $\mathcal{H}^L$  of L levels. Its integration starts at the highest hierarchical level, L, and progresses downward to the level 1. At each level,  $\hbar_+$  is associated with the holon  $\hbar_*^1$  that shows the highest performance in the new set of unit validations,  $\mathcal{V}_+$ , subject to meeting budget constraints, *i.e.*,  $|\mathcal{T}^l_* \cup \mathcal{T}_+| \leq B^l$ . Once integrated,  $\hbar_+$ 's dataset merges with that of the selected holon,  $\hbar_*^l$ , necessitating a retraining of the aggregated dataset. If no appropriate holons are available at a required level, the system can initiate reholonification, integrating  $\hbar_+$  with the set  $\{\hbar_1^l, \cdots, \hbar_{N^l}^l\}$ . A pseudocode is provided in Algorithm 1.

#### Algorithm 1 Integration of a New Holon into a Holarchy

```
Require: \mathcal{H}^L: L-level learning holarchy
Require: \hbar_+: A holon
 1: for l = L downto 1 do
                    > Identify sub-holons whose training sets do not exceed budget
      constraints
           FreeHolons \leftarrow \hbar^l : |\mathcal{T}^l \cup \mathcal{T}_+| \le B^l
 2:
           \mathbf{if} \ \mathsf{FreeHolons} = \emptyset \ \mathbf{then}
 3:
                Reholonification with the set \{\hbar_1^l, \cdots, \hbar_{N^l}^l\} \cup \hbar_+.
 4:
                break
 5:
           end if
 6:
                                                   > Select the optimal sub-holon for integration
 7:
           \mathsf{FreeModels} \leftarrow \mathsf{FreeHolons's} \ \mathrm{models}
 8:
           \theta_*^l \leftarrow \arg \max_{\theta^l \in \mathsf{FreeModels}} f\left(\theta^l, \mathcal{V}_+\right)
           Update model parameters \theta_*^l using \mathcal{T}_*^l \cup \mathcal{T}_+.
 9:
                                                                               \triangleright Integrate \hbar_+ into holon of \theta^l_*
           \mathsf{SUB}(\hbar^l_*) \leftarrow \mathsf{SUB}(\hbar^l_*) \cup \hbar_+
10:
11: end for
```

Remark 1. The cost of holonification is compared to integration on-the-fly on the basis of the amount of communication between the holons. It is built on a single linkage-Hierarchical Clustering, with a complexity of  $\mathcal{O}(N^2)$  [43]. On the other hand, the on-the-fly mechanism has a  $\mathcal{O}(N + L)$  complexity. This corresponds to the worst-case scenario in which the new agent is compared to all holons from the upper layer L to layer 0. This mechanism thus offers a cost-effective integration in comparison with a reholonification.

#### 4.3 Research Questions

From the setup and challenges described above, we formulate the following research questions:

- 1. Given a new data stream, how can we determine the most suitable existing model for fine-tuning?
- 2. Assume an effective integration of new sensors based on similarity with a group of sensors from the system:
  - (a) What are the consequences on model accuracy upon the integration of a new agent?
  - (b) How does the accuracy of the model scale when incrementally integrating  $N_+$  new agents versus performing a full system reorganization?
- 3. What are the long-term accuracy trade-offs between retaining versus discarding data from removed sensors?

# 5 Materials and Methods

The datasets, the training procedure, and the evaluation protocol are presented in this Section.

### 5.1 Datasets

We used two city-focused video datasets for a total of 16 cameras.

WALT [40] Features footage from nine static cameras over 1–4 weeks. Sampling rates vary (5,000-40,000 frames/week), with temporal bursts and diverse weather conditions (snow, rain, day/night).

AI-City [34]. Seven annotated videos, each approximately five minutes at 10 FPS. Camera angles and sensor types vary (vertical, dome, PTZ), ensuring coverage of multiple representation contexts.

# 5.2 Model Training

We follow Stream-Based Active Distillation (SBAD) [32] sampling 256 images per camera. A large YOLOv8x6 Teacher (261.1 GFLOPs) pseudolabels these samples. Each Student model is a YOLOv8n (8.7 GFLOPs), initialized with COCO weights [29], then fine-tuned at a learning rate of 0.01 (unless otherwise indicated).

### 5.3 Evaluation

We report the mean Average Precision (mAP) across various Intersection over Union (IoU) thresholds, spanning from 0.50 to 0.95 in increments of 0.05 (mAP50-95). We evaluated the holon's performance on its associated datasets.

# 6 Results

We begin by evaluating our holonification approach under different budgets and then proceed with incremental integration, departure handling, and knowledgetransfer experiments.

#### 6.1 Holonification Baseline Performance

We conducted a holonification, as proposed in Section 4.1, on a dataset comprising sixteen cameras. We set a multilayer budget framework  $B^l = 256 \cdot 10^l$ , implying that layer 0 holons do not exceed 256 training samples, and successive layers cannot exceed  $10^l$  sub-holons for a holon  $\hbar^l$ . Table 1 confirms that the

Table 1: Holonification with varying budgets. Shown are the final groupings and average mAP50-95 for the 16-camera dataset.

Layer	$\mathbf{B}^{1}$	mAP50-95	Holonic Structure
2	25600	0.65	$\mathcal{H}^2:\{\hbar^1_1,\hbar^1_2,\hbar^1_3\}$
			$\mathcal{H}^1_1: \{\hbar^0_1, \hbar^0_2, \hbar^0_3, \hbar^0_9\}$
1	2560	0.66	$\mathcal{H}_{2}^{1}:\{\hbar_{4}^{0},\hbar_{5}^{0},\hbar_{6}^{0},\hbar_{7}^{0},\hbar_{8}^{0}\}$
			$\mathcal{H}_{3}^{1}:\{\hbar_{16}^{0},\hbar_{17}^{0},\hbar_{18}^{0},\hbar_{19}^{0},\hbar_{20}^{0},\hbar_{22}^{0},\hbar_{24}^{0}\}$
0	256	0.67	$\mathcal{H}^1:\{\hbar^0_1,\ldots,\hbar^0_{16}\}$
YOLOv8n <sup>COCO</sup>	N.A	0.498	N.A

models require specificity to achieve maximum performance.

#### 6.2 Transfer, Integration and Departure

**Model Transfer Upon Increment** The transferability of holons across new, although similar, domains is investigated. Table 2 details the performance results for models trained in an all-but-one combination of domains as well as across all domains.

The results indicate that models struggle to transfer, even across similar camera domains, reinforcing the need to integrate the newcomer in a cluster, and the local retraining the cluster.

Table 2: mAP50-95 scores for models trained under an all-but-one camera to assess the transferability of those models on the remaining camera. A baseline is also provided where the model is trained across all cameras. Each model is trained for 10000 iterations.

Cluster	$\mathcal{V}_1^0$	$\mathcal{V}_2^0$	$\mathcal{V}_3^0$	$\overline{\mathcal{V}}_9^0$
$ heta^{\mathcal{H}_1^1\setminus\{\hbar_1^0\}}$	0.38	0.65	0.65	0.47
$ heta^{\mathcal{H}_1^1  ackslash \{\hbar_2^0\}}$	0.46	0.57	0.66	0.46
$ heta^{\mathcal{H}_1^1  ackslash \{\hbar_3^0\}}$	0.49	0.66	0.65	0.48
$ heta^{\mathcal{H}_1^1\setminus\{\hbar_9^0\}}$	0.47	0.64	0.66	0.42
$ heta^{\mathcal{H}_1^1}$	0.46	0.65	0.66	0.47

**Incremental Integration** We evaluated the impact of integrating  $N_+$  new units into a holonified system, structured with budget limits of  $B^{l} = 256 \cdot 10^{l}$ . Using our integration mechanism described in Algorithm 1, we evaluated two scenarios: integrating one  $(N_{+} = 1)$  and three  $(N_{+} = 3)$  additional agents.

In 16 agent configurations, the incremental integration maintained an average mAP50-95 of  $0.66 \pm 0.003$  ( $N_{+} = 1$ ) and  $0.66 \pm 0.006$  ( $N_{+} = 3$ ), showing no degradation compared to the baseline in Table 1.

Agent Departure When a sensor  $\hbar_i^0$  leaves, its data  $\mathcal{T}_i$  may be retained or discarded. We successively simulate the exit of each agent and track the accuracy of the global model on (i) remaining and (ii) left sensors. As Fig. 4 shows, removing a sensor's data yields small gains for the remaining sensors, but severely reduces performance if that sensor later re-enters the system. Note that, upon the departure of an agent, their data set  $\mathcal{T}_i^0$  is removed from the collective data set  $\mathcal{T}^2$ , and the model is re-trained for 10,000 iterations at a learning rate of 0.005.

#### 6.3 Inter-Holonic Knowledge Transfer

We test how effectively a holon trained on existing cameras can accelerate training and improve the peak accuracy of a newcomer domain. Specifically, we conducted 16 trials, each excluding one camera from the data set to simulate a "newcomer". The following pre-trained models serve as initial weights:

- $-\theta^2$ : Global holon (trained on 15 cameras),
- $-\theta_*^1$ : Group-specific holon, YOLOv8n<sup>COCO</sup>: General-purpose off-the-shelf model.

Fig. 5 shows that  $\theta^2$  or  $\theta^1_*$  consistently outperform the generic COCO baseline when fine-tuning the newcomer camera. Training spanned 5 epochs with a learning rate of 0.005.



Fig. 4: Difference in  $\hbar^2$  model performance between retaining and discarding each departed sensor's data. Blue: remaining sensors; yellow: departed sensors. Results show marginal gains for remaining sensors but a marked degradation on departed sensors.



Fig. 5: mAP50-95 per epoch for a new model starting from universal model  $\theta^2$ , group-specific  $\theta_*^1$ , and general-purpose  $\theta^{\text{COCO}}$ . The superiority of  $\theta^2$  highlights the efficiency of selecting a pretrained model closer to the source.

# 7 Discussion

#### 7.1 Insights

Our experiments confirm that a certain level of domain specificity improves accuracy (Section 6.1), though it also increases the number of models to maintain. Budget constraints help contain this growth but can reduce performance gains from specialized holons. Meanwhile, leveraging broader universal models accelerates learning for new domains (Section 6.3).

In the context of open systems (Section 6.2), our sanity check shows that a straightforward model transfer performs under, even when the model comes from similar domains. The observed performance gap motivated the development of integration mechanisms, which proved effective, but the experiments do not provide conclusive evidence regarding the maximum number of agents that can be integrated without a performance decline. Finally, agent departure highlights a trade-off between short-term gains and relearning costs if the environment reappears. In other words, discarding data should be considered in terms of agents' turnover rate.

## 7.2 Limitations

Machine Learning Lifecycle Machine learning based models are also subject to feedback loops, where data and interactions with the external world influence their behavior in unintentional ways [42]. The subsequent design of machine learning systems should account for the fact that their behavior evolves with environmental data and user interactions. This includes providing control mechanisms to avoid the accumulation of errors due to the self-supervised nature of the system.

Stress-tests We need to further stress test the system; that is, starting with a system of size N, stress tests can evaluate how many new components  $(N_+)$  can be integrated without compromising the quality of service.

#### 7.3 Perspectives

Modern systems integrate heterogeneous approaches (*e.g.*, physics-based modeling vs. deep learning) and diverse sensing modalities (*e.g.*, cameras, radar), producing richer analytics [5, 33]. Our architecture abstracts the holonic paradigm sufficiently to accommodate such heterogeneity. However, specialized coordination modules could further optimize collaborative performance among different modalities.

# 8 Conclusions

We present an organizational holonic learning design coupled with active learning to address the challenges of scaling learning in multisensor networks. Our self-organization mechanism, grounded in the specificity-diversity trade-off, allows for the establishment of various granularity levels and handles sensor addition and removal while maintaining strong predictive performance. Experimental results highlight the benefits of vertical and horizontal knowledge transfer, although more stress testing is needed to refine the upper limits on system growth. We also note that self-supervised processes risk model drift without robust monitoring, which may cause issues in autoscaling and autotuning. Future work aims to design colearning mechanisms for heterogeneous methods.

# Bibliography

- [1] Abbas, H.A.: Organization of multi-agent systems: An overview. International Journal of Intelligent Information Systems 4(3), 46 (2015), https://doi.org/10.11648/j.ijiis.20150403.11
- Beal, J., Viroli, M., Pianini, D., Damiani, F.: Self-adaptation to device distribution changes. In: 2016 IEEE 10th International Conference on Self-Adaptive and Self-Organizing Systems (SASO), pp. 60–69 (2016), https://doi.org/10.1109/SASO.2016.12
- [3] Bernon, C., Gleizes, M.P., Peyruqueou, S., Picard, G.: Adelfe: A methodology for adaptive multi-agent systems engineering. In: Petta, P., Tolksdorf, R., Zambonelli, F. (eds.) Engineering Societies in the Agents World III, pp. 156–169, Springer Berlin Heidelberg, Berlin, Heidelberg (2003), ISBN 978-3-540-39173-9, https://doi.org/10.1007/3-540-39173-8\_12
- [4] Cacciarelli, D., Kulahci, M.: Active learning for data streams: a survey. Machine Learning (2023), https://doi.org/10.1007/s10994-023-06454-2
- [5] Campagner, A., Ciucci, D., Cabitza, F.: Aggregation models in ensemble learning: A large-scale comparison. Information Fusion 90, 241-252 (2023), ISSN 1566-2535, https://doi.org/https://doi.org/10.1016/j.inffus.2022.09.015
- [6] Castro, J., Kolp, M., Mylopoulos, J.: Towards requirements-driven information systems engineering: the tropos project. Information Systems 27(6), 365–389 (2002), ISSN 0306-4379, https://doi.org/10.1016/ S0306-4379(02)00012-1
- [7] Chen, J., Gao, C., Sun, L., Sang, N.: Ccsd: cross-camera self-distillation for unsupervised person re-identification. Visual Intelligence 1(1), 27 (2023)
- [8] Cossentino, M., Gaud, N., Hilaire, V., Galland, S., Koukam, A.: Aspecs: an agent-oriented software process for engineering complex systems. Autonomous Agents and Multi-Agent Systems 20(2), 260–304 (2010), https: //doi.org/10.1007/s10458-009-9099-4
- [9] Diaconescu, A., Frey, S., Müller-Schloer, C., Pitt, J., Tomforde, S.: Goaloriented holonics for complex system (self-)integration: Concepts and case studies. In: 2016 IEEE 10th International Conference on Self-Adaptive and Self-Organizing Systems (SASO), pp. 100–109 (2016), https://doi.org/ 10.1109/SASO.2016.16
- [10] Dong, J., Liu, R., Qiu, Y., Crossan, M.: Should knowledge be distorted? managers' knowledge distortion strategies and organizational learning in different environments. The Leadership Quarterly 32(3), 101477 (2021), ISSN 1048-9843
- [11] Esmaeili, A., Gallagher, J.C., Springer, J.A., Matson, E.T.: Hamlet: A hierarchical agent-based machine learning platform. ACM Trans. Auton. Adapt. Syst. 16(3–4) (jul 2022), ISSN 1556-4665
- [12] Esmaeili, A., Gallagher, J.C., Springer, J.A., Matson, E.T.: Hamlet: A hierarchical agent-based machine learning platform. ACM Trans. Auton.

Adapt. Syst. 16(3-4) (Jul 2022), ISSN 1556-4665, https://doi.org/10. 1145/3530191

- [13] Esmaeili, A., Ghorrati, Z., Matson, E.T.: Holonic learning: A flexible agentbased distributed machine learning framework. In: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, p. 525–533, AAMAS '24, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2024), ISBN 9798400704864
- [14] Esmaeili, A., Mozayani, N., Jahed-Motlagh, M.R., Matson, E.T.: Towards topological analysis of networked holonic multi-agent systems. In: Advances in Practical Applications of Survivable Agents and Multi-Agent Systems: The PAAMS Collection, pp. 42–54, Springer International Publishing, Cham (2019), ISBN 978-3-030-24209-1
- [15] Esmaeili, A., Mozayani, N., Motlagh, M.R.J., Matson, E.T.: The impact of diversity on performance of holonic multi-agent systems. Engineering Applications of Artificial Intelligence 55, 186–201 (2016), ISSN 0952-1976
- [16] Feraud, M., Galland, S.: First comparison of sarl to other agentprogramming languages and frameworks. Procedia Computer Science 109, 1080-1085 (2017), ISSN 1877-0509, https://doi.org/https://doi.org/ 10.1016/j.procs.2017.05.389
- [17] Fowler, M.: Refactoring: improving the design of existing code. Addison-Wesley Professional (2018)
- [18] French, R.M.: Catastrophic forgetting in connectionist networks. Trends in Cognitive Sciences 3(4), 128–135 (1999), ISSN 1364-6613, https://doi. org/10.1016/S1364-6613(99)01294-2
- [19] Gherardi, S.: Learning: Organizational. In: International Encyclopedia of the Social & Behavioral Sciences (Second Edition), pp. 695–698, Elsevier, Oxford, second edition edn. (2015), ISBN 978-0-08-097087-5
- [20] Ghosh, A., Chung, J., Yin, D., Ramchandran, K.: An efficient framework for clustered federated learning. In: Advances in Neural Information Processing Systems, vol. 33, pp. 19586–19597, Curran Associates, Inc. (2020)
- [21] Gleizes, M.P.: Self-adaptive complex systems. In: Multi-Agent Systems, pp. 114–128, Springer, Berlin, Heidelberg (2012), ISBN 978-3-642-34799-3
- [22] Gupta, O., Raskar, R.: Distributed learning of deep neural network over multiple agents. Journal of Network and Computer Applications 116, 1–8 (2018), ISSN 1084-8045, https://doi.org/10.1016/j.jnca.2018.05.003
- [23] Hilal, A.R., Basir, O.A.: A scalable sensor management architecture using BDI model for pervasive surveillance. IEEE Systems Journal 9(2), 529–541 (2015), ISSN 19379234, https://doi.org/10.1109/JSYST.2014.2334071
- [24] Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. NIPS 2014 Deep Learning Workshop (2015)
- [25] Hosseinalipour, S., Brinton, C.G., Aggarwal, V., Dai, H., Chiang, M.: From federated to fog learning: Distributed machine learning over heterogeneous wireless networks. IEEE Communications Magazine 58(12), 41–47 (2020)
- [26] Jocher, G., Chaurasia, A., Qiu, J.: Ultralytics YOLO (2023), URL https: //github.com/ultralytics/ultralytics
- [27] Koestler, A.: The ghost in the machine. Hutchinson, London, UK (1967)

- 18 D. Manjah et al.
- [28] Le, J., Lei, X., Mu, N., Zhang, H., Zeng, K., Liao, X.: Federated continuous learning with broad network architecture. IEEE Transactions on Cybernetics 51(8), 3874–3888 (2021), ISSN 2168-2275, https://doi.org/10.1109/ TCYB.2021.3090260
- [29] Lin, T., Maire, M., Belongie, S.J., Bourdev, L.D., Girshick, R.B., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. CoRR abs/1405.0312 (2014)
- [30] Ma, Z., Xu, Y., Xu, H., Liu, J., Xue, Y.: Like attracts like: Personalized federated learning in decentralized edge computing. IEEE Transactions on Mobile Computing 23(2), 1080–1096 (2024), ISSN 1558-0660
- [31] Manjah, D., Cacciarelli, D., Standaert, B., Benkedadra, M., de Hertaing, G.R., Macq, B., Galland, S., De Vleeschouwer, C.: Stream-based active distillation for scalable model deployment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pp. 4998–5006 (2023)
- [32] Manjah, D., Cacciarelli, D., Vleeschouwer, C.D., Macq, B.: Camera clustering for scalable stream-based active distillation (2024)
- [33] Manjah, D., Galland, S., Vleeschouwer, C.D., Macq, B.: Autonomous methods in multisensor architecture for smart surveillance. In: Proceedings of the 16th International Conference on Agents and Artificial Intelligence, vol. 3, pp. 824–832 (2024), ISBN 978-989-758-680-4, ISSN 2184-433X
- [34] Naphade, M., Wang, S., Anastasiu, D.C., Tang, Z., Chang, M.C., Yang, X., Yao, Y., Zheng, L., Chakraborty, P., Lopez, C.E., Sharma, A., Feng, Q., Ablavsky, V., Sclaroff, S.: The 5th ai city challenge. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2021)
- [35] Nezamoddini, N., Gholami, A.: A survey of adaptive multi-agent networks and their applications in smart cities. Smart Cities 5(1), 318–347 (2022), ISSN 2624-6511
- [36] Omicini, A.: Soda: Societies and infrastructures in the analysis and design of agent-based systems. In: Agent-Oriented Software Engineering, pp. 185– 193, Springer Berlin Heidelberg (2001), ISBN 978-3-540-44564-7
- [37] Pavón, J., Gómez-Sanz, J.: Agent oriented software engineering with ingenias. In: Multi-Agent Systems and Applications III, pp. 394–403, Springer, Berlin, Heidelberg (2003), ISBN 978-3-540-45023-8
- [38] Perera, C., Zaslavsky, A., Christen, P., Georgakopoulos, D.: Sensing as a service model for smart cities supported by Internet of Things. Transactions on Emerging Telecommunications Technologies 25(1), 81–93 (2014), ISSN 2161-3915
- [39] Porter, B., Rodrigues Filho, R.: Distributed emergent software: Assembling, perceiving and learning systems at scale. In: 2019 IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems (SASO), pp. 127– 136 (2019)
- [40] Reddy, N.D., Tamburo, R., Narasimhan, S.G.: Walt: Watch and learn 2d amodal representation from time-lapse imagery. In: Proceedings of

19

the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9356–9366 (2022)

- [41] Rodriguez, S., Hilaire, V., Gaud, N., Galland, S., Koukam, A.: Holonic Multi-Agent Systems. Natural Computing Series 37, 251–279 (2011), https://doi.org/10.1007/978-3-642-17348-6\_11
- [42] Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.F., Dennison, D.: Hidden technical debt in machine learning systems. In: Advances in Neural Information Processing Systems, vol. 28, Curran Associates, Inc. (2015)
- [43] Sibson, R.: SLINK: An optimally efficient algorithm for the single-link cluster method. The Computer Journal 16(1), 30–34 (1973), ISSN 0010-4620
- [44] Terabe, M., Washio, T., Katai, O., Sawaragi, T.: A study of organizational learning in multiagents systems. In: Distributed Artificial Intelligence Meets Machine Learning Learning in Multi-Agent Environments, pp. 168– 179, Springer Berlin Heidelberg, Berlin, Heidelberg (1997), ISBN 978-3-540-69050-4
- [45] Valencia-Jiménez, J.J., Fernández-Caballero, A.: Holonic multi-agent systems to integrate independent multi-sensor platforms in complex surveillance. Proceedings - IEEE International Conference on Video and Signal Based Surveillance 2006, AVSS 2006 pp. 49–54 (2006), https://doi.org/ 10.1109/AVSS.2006.58
- [46] Wautelet, Y., Schinckus, C., Kolp, M.: Agent-based software engineering, paradigm shift, or research program evolution. In: Research Anthology on Recent Trends, Tools, and Implications of Computer Programming, pp. 1642–1654, IGI Global (2021)
- [47] Weyns, D., Gerostathopoulos, I., Abbas, N., Andersson, J., Biffl, S., Brada, P., Bures, T., Di Salle, A., Galster, M., Lago, P., Lewis, G., Litoiu, M., Musil, A., Musil, J., Patros, P., Pelliccione, P.: Self-adaptation in industry: A survey. ACM Trans. Auton. Adapt. Syst. 18(2) (2023), ISSN 1556-4665
- [48] Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation 1(1), 67–82 (1997), ISSN 1089778X, https://doi.org/10.1109/4235.585893
- [49] Wooldridge, M., Jennings, N.R., Kinny, D.: The gaia methodology for agentoriented analysis and design. Autonomous Agents and Multi-Agent Systems 3(3), 285–312 (2000), https://doi.org/10.1023/A:1010071910869
- [50] Xu, B., Xia, W., Wen, W., Liu, P., Zhao, H., Zhu, H.: Adaptive hierarchical federated learning over wireless networks. IEEE Transactions on Vehicular Technology 71(2), 2070–2083 (2021)
- [51] Yang, Y., Luo, J., Wen, Y., Slumbers, O., Graves, D., Bou Ammar, H., Wang, J., Taylor, M.E.: Diverse auto-curriculum is critical for successful real-world multiagent learning systems. In: Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, p. 51–56, AAMAS '21, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2021), ISBN 9781450383073
- [52] Yin, D., Pananjady, A., Lam, M., Papailiopoulos, D., Ramchandran, K., Bartlett, P.: Gradient diversity: a key ingredient for scalable distributed

learning. In: Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research, vol. 84, pp. 1998–2007, PMLR (2018)