



# Holarchic structures for decentralized deep learning: a performance analysis

Evangelos Pournaras<sup>1</sup> · Srivatsan Yadhunathan<sup>1</sup> · Ada Diaconescu<sup>2</sup>

Received: 31 July 2018 / Revised: 20 September 2018 / Accepted: 5 January 2019 / Published online: 13 February 2019  
© The Author(s) 2019

## Abstract

Structure plays a key role in learning performance. In centralized computational systems, hyperparameter optimization and regularization techniques such as dropout are computational means to enhance learning performance by adjusting the deep hierarchical structure. However, in decentralized deep learning by the Internet of Things, the structure is an actual network of autonomous interconnected devices such as smart phones that interact via complex network protocols. Self-adaptation of the learning structure is a challenge. Uncertainties such as network latency, node and link failures or even bottlenecks by limited processing capacity and energy availability can significantly downgrade learning performance. Network self-organization and self-management is complex, while it requires additional computational and network resources that hinder the feasibility of decentralized deep learning. In contrast, this paper introduces a self-adaptive learning approach based on holarchic learning structures for exploring, mitigating and boosting learning performance in distributed environments with uncertainties. A large-scale performance analysis with 864,000 experiments fed with synthetic and real-world data from smart grid and smart city pilot projects confirm the cost-effectiveness of holarchic structures for decentralized deep learning.

**Keywords** Deep learning · Optimization · Holarchy · Multi-agent system · Resilience · Smart city

## 1 Introduction

Smart citizens' devices with increasing processing power and high energy autonomy are becoming pervasive and ubiquitous in everyday life. The Internet of Things empowers a high level of interconnectivity between smart phones, sensors and wearable devices. These technological developments provide unprecedented opportunities to rethink about the future of machine learning and artificial intelligence: Centralized computational intelligence can be often used for privacy-intrusive and discriminatory services

that create 'filter bubbles' and undermine citizens' autonomy by nudging [12, 17, 31]. In contrast, this paper envisions a more socially responsible design for digital society based on decentralized learning and collective intelligence formed by bottom-up planetary-scale networks run by citizens [18, 40].

In this context, the structural elements of decentralized deep learning processes play a key role. The effectiveness of several classification and prediction operations often relies heavily on hyperparameter optimization [27, 53] and on the learning structure, for instance, the number of layers in a neural network, the interconnectivity of the neurons, the activation or deactivation of certain pathways i.e. dropout regularization [51], can enhance learning performance. Controlling and adjusting a deep hierarchical structure in a centralized computing system is straightforward in the sense that all meta information for model generation is locally available. However, in decentralized learning the challenge of optimizing the learning structure is not anymore exclusively a computational problem. Other challenges such as network latency, node and link failures

---

✉ Evangelos Pournaras  
epournaras@ethz.ch

Srivatsan Yadhunathan  
ysrivatsan@ethz.ch

Ada Diaconescu  
ada.diaconescu@telecom-paristech.fr

<sup>1</sup> Professorship of Computational Social Science, ETH Zurich, Zurich, Switzerland

<sup>2</sup> Telecom ParisTech, Paris-Saclay University, Paris, France

as well as the overall complexity of building and maintaining an overlay network [36] in a distributed environment perplex the feasibility of decentralized learning.

This paper introduces the concept of *holarchy* in deep hierarchical structures as the means to adapt to the aforementioned uncertainties of distributed environments. A learning process can be localized and performed over a holarchic structure in a recursive way without changing the core learning logic and without employing additional mechanisms to reconfigure the network. This is the proposed *self-adaptation approach* to decentralized learning that is by design highly reactive and cost-effective as it maximizes the utilization of the available communication and computational resources, in contrast to a complementary and more proactive *self-organization approach* that requires additional interactions between agents and therefore can increase communication and computational cost. By using holarchies for learning, forward propagation and backpropagation become recursive in nested levels over the deep hierarchical structure as the means to (i) *explore* improving solutions, (ii) *mitigate* learning performance in case part of the network is disconnected or even (iii) *boost* learning performance after the default learning process completes. These three scenarios are formalized by three holarchic schemes applied to a multi-agent system for decentralized deep learning in combinatorial optimization problems: I-EPOS, the *Iterative Economic Planning and Optimized Selections* [40].

A large-scale performance analysis with 864,000 experiments is performed using synthetic and real-world data from pilot projects such as bike sharing, energy demand and electric vehicles. Several dimensions are studied, for instance, topological properties of the deep hierarchical structure, constraints by the agents' preferences and the scale of the holarchic structures, i.e. number of nested layers. Results confirm the cost-effectiveness of the holarchic learning schemes for exploration, mitigation and boosting of learning performance in dynamic distributed environments. Nevertheless, in stable environments the localization of the learning process within holarchic structures may result in lower performance compared to a learning applied system-wide.

In summary, the contributions of this paper are the following:

- A novel self-adaptation approach to decentralized deep learning as the means to decrease or avoid the communication and computational cost of self-organization in distributed environments with uncertainties.
- The concept of holarchy as a self-adaptation approach for decentralized deep learning.

- The introduction of three holarchic schemes as self-adaptation means to explore, mitigate and boost deep learning performance.
- The applicability and extension of I-EPOS with the three holarchic schemes to perform collective decision-making in decentralized combinatorial optimization problems.
- An empirical performance analysis of 864,000 benchmark experiments generated with synthetic and real-world data from Smart Grid and Smart City pilot projects.

This paper is organized as follows: Sect. 2 introduces the concept of holarchy in decentralized learning as well as three holarchic schemes to explore, mitigate and boost deep learning performance under uncertainties of distributed environments. Sect. 3 illustrates a case study of a decentralized deep learning system to which the three holarchic schemes are applied: I-EPOS. Sect. 4 shows the experimental methodology followed to conduct a large-scale performance analysis of the three holarchic schemes. Section 5 illustrates the results of the experimental evaluation. Section 6 summarizes and discusses the main findings. Section 7 positions and compares this paper with related work. Finally, Sect. 8 concludes this paper and outlines future work.

## 2 Holarchic structures for decentralized deep learning

This paper studies *decentralized deep learning* processes in which the deep hierarchical structure is fully distributed and self-organized by remote autonomous (software) agents that interact over a communication network. In other words, the learning process is crowd-sourced to citizens, who participate by contributing their computational resources, for instance their personal computers or smart phones. The agents reside on these devices and collectively run the decentralized learning process. For example, in contrast to a conventional model of a centralized neural network, which performs training by locally accessing all citizens' data, a collective neural network consists of neurons that are remote citizens' devices interacting over the Internet to form an overlay peer-to-peer network [36]. It is this overlay network that represents the hierarchical neural network structure.

Methods for hyperparameter optimization of the hierarchical structure to improve the learning performance of a model are usually designed for centralized systems in which all information, including input data, the network structure and the learning model itself are locally available. This provides a large spectrum of flexibility to change the

deep hierarchical structure offline or even online [16, 42] and determine via hyperparameter optimization the settings that maximize the learning performance. In contrast to learning based on centralized computational systems, in decentralized deep learning the structure cannot arbitrary change without paying for some computational and communication cost. Network uncertainties such as node and link failures, latency as well as limited resources in terms of bandwidth, computational capacity or even energy in case of sensors and smart phones can limit performance, interrupt the learning process and increase the design complexity of decentralized hyperparameter optimization.

The aforementioned uncertainties of distributed environments introduce endogenous constraints in parts of the deep hierarchical structure: the learning process is interrupted and becomes *localized* within branches of the hierarchical structure. For instance, node and link failures or suspension of the learning processes due to conservation of resources [2, 52] in nodes are scenarios under which learning can be localized. This paper poses here the following question: How to preserve the learning capacity of a decentralized system, whose learning process is interrupted and localized by aforementioned uncertainties in distributed environments? On the one hand, it is known that localization and greedy optimization can underperform with search becoming trapped to locally optimum solutions that have a significant divergence from global optimality [3]. On the other hand, limiting the learning units of hierarchical structures can also increase learning performance by preventing overfitting as known by the dropout concept in neural networks [51]. This paper sheds light on the role of localization in decentralized learning.

In this context, the management of the learning performance of an algorithm is no longer entirely a computational challenge but rather a multifaceted self-adaptation process: Other aspects such as performance *exploration*, *mitigation* and *boosting* come to the analysis foreground. Exploration adapts the learning process and the search space to improve learning performance under localization. Mitigation is the maintenance of a high learning performance under a localization of the learning process. Finally, the feasibility of boosting the learning performance under localization is subject of this study as well.

This paper introduces the concept of *holarchy* in deep learning hierarchical structures to study the performance exploration, mitigation and boosting potential under the aforementioned uncertainties of distributed environments. A holarchy is a recursive hierarchical network of *holons* that represent part of the deep hierarchical structure as well as the whole structure. In the case of a tree topology, every possible branch (part) in the whole tree topology is also a tree topology (whole). When an agent (parent) connects two branches, it forms another nested holon that is the next

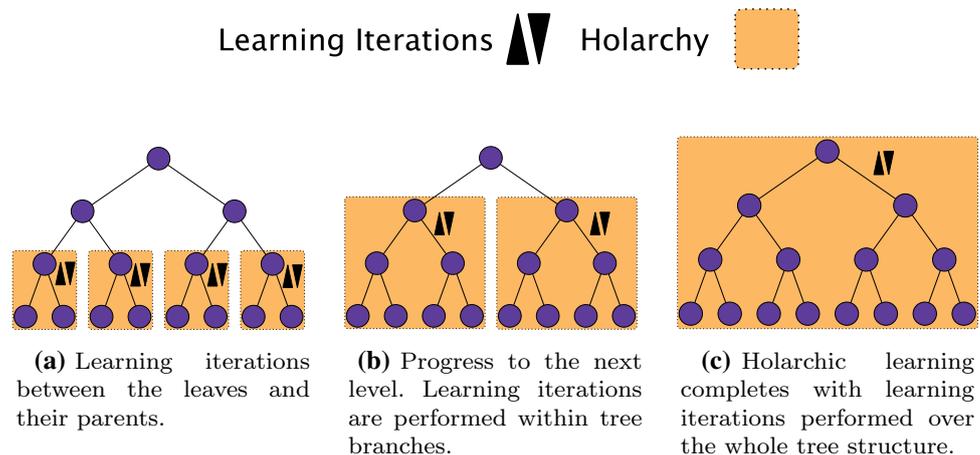
level of the holarchic structure. This recursive process starts from the parents of the leaves in the tree and progresses up to the root as shown in Fig. 1. Learning iterations can be independently executed in every holonic branch before the process progresses to the next level of the holarchic structure, in which a new series of learning iterations are executed. The top holonic branch is actually the whole tree topology and therefore the execution of learning iterations at this top level corresponds to the learning iterations without a holarchic structure. In other words, the concept of holarchy introduces multiple localized, nested and incremental learning processes.

The nested learning processes of a holarchic system enable the management of each level independently from the higher and lower levels, which further allows the management separation between the agents at each level. This is a divide-and-conquer approach to complex system management and hence improves localization, parallelism, reusability and diversity with a potential in improving system performance and robustness while decreasing costs [7]. From a software engineering perspective, the following key properties of a holonic design are identified [7, 9]: (i) *Bottom-up abstraction* that represents the aggregation of information from lower levels. (ii) *Partial isolation* within and among levels—ensuring an agent’s decision results from the aggregation of agents’ decisions at a lower level and results in the agent’s decision at a higher level, hence other agents do not influence. (iii) *Inter-level time tuning*—ensuring that the relative execution times at different holarchic levels are set in order to avoid cross-level oscillations and divergence.

This paper studies whether learning processes structured within holarchies are cost-effective countermeasures to adapt to the uncertainties of distributed environments. Holarchies can provide the following operational flexibility and adaptation: (i) The learning process can be limited to a targeted part of the network to prevent a network-wide use of computational and communication resources, i.e. agents can participate in the learning process on-demand. Any failure to serve the participation does not disrupt the learning process that can continue within part of the network. (ii) The mapping and deployment of holarchies on the network can be designed according to the network heterogeneity, i.e. varying latency, computational and battery capacities. For instance, higher performing nodes can be placed at the bottom of a holarchy to serve their more frequent recurrent use in bottom-up learning interactions.

This paper studies three self-adaptation scenarios of the holarchic concept in decentralized deep learning each designed for performance exploration, mitigation and boosting respectively: (i) *holarchic initialization*, (ii) *holarchic runtime* and (iii) *holarchic termination*. Assume a

**Fig. 1** The concept of holarchic learning. Learning iterations are performed in nested branches, the holons. **c** Actually depicts the default baseline learning strategy, while **a** and **b** show the earlier learning iterations performed within the holons



*baseline* scheme that involves a tree structure with agents interacting in a (i) *bottom-up phase* and a (ii) *top-down phase* that both complete a *learning iteration*. The former phase may represent a fit forward learning process starting from the leaves and completing to the root while the latter phase a backpropagation starting from the root and reaching back the leaves. Without loss of generality, an exact decentralized learning algorithm realizing these concepts is presented in Sect. 3. Learning iterations repeat to decrease a *cost function*. Learning *converges* when a certain number of iterations is performed or when the cost function cannot be decreased further. Figure 2a illustrates the baseline scheme.

Figure 1c depicts one baseline learning iteration, while within each nested holon formed during the bottom-up phase several learning iterations are performed. This process is common in all holarchic schemes. Holarchic initialization is applied before baseline to perform an exploration of the search space. Several learning iterations can be performed before switching to baseline as illustrated in Fig. 2b. In contrast, holarchic runtime applies holarchic learning throughout runtime without switching to baseline as shown in Fig. 2c. This scheme is applicable in self-adaptation scenarios of failures or conservation of resources to mitigate losses of learning performance. Finally, holarchic termination is applied after the baseline convergence as the means to further boost the baseline performance. This self-adaptation scheme is shown in Fig. 2d.

### 3 Applicability of holarchic schemes

This section illustrates a case study for the applicability of holarchic structures in decentralized deep learning for combinatorial optimization problems: I-EPOS,<sup>1</sup> the *Iterative economic planning and optimized selections* [35, 40].

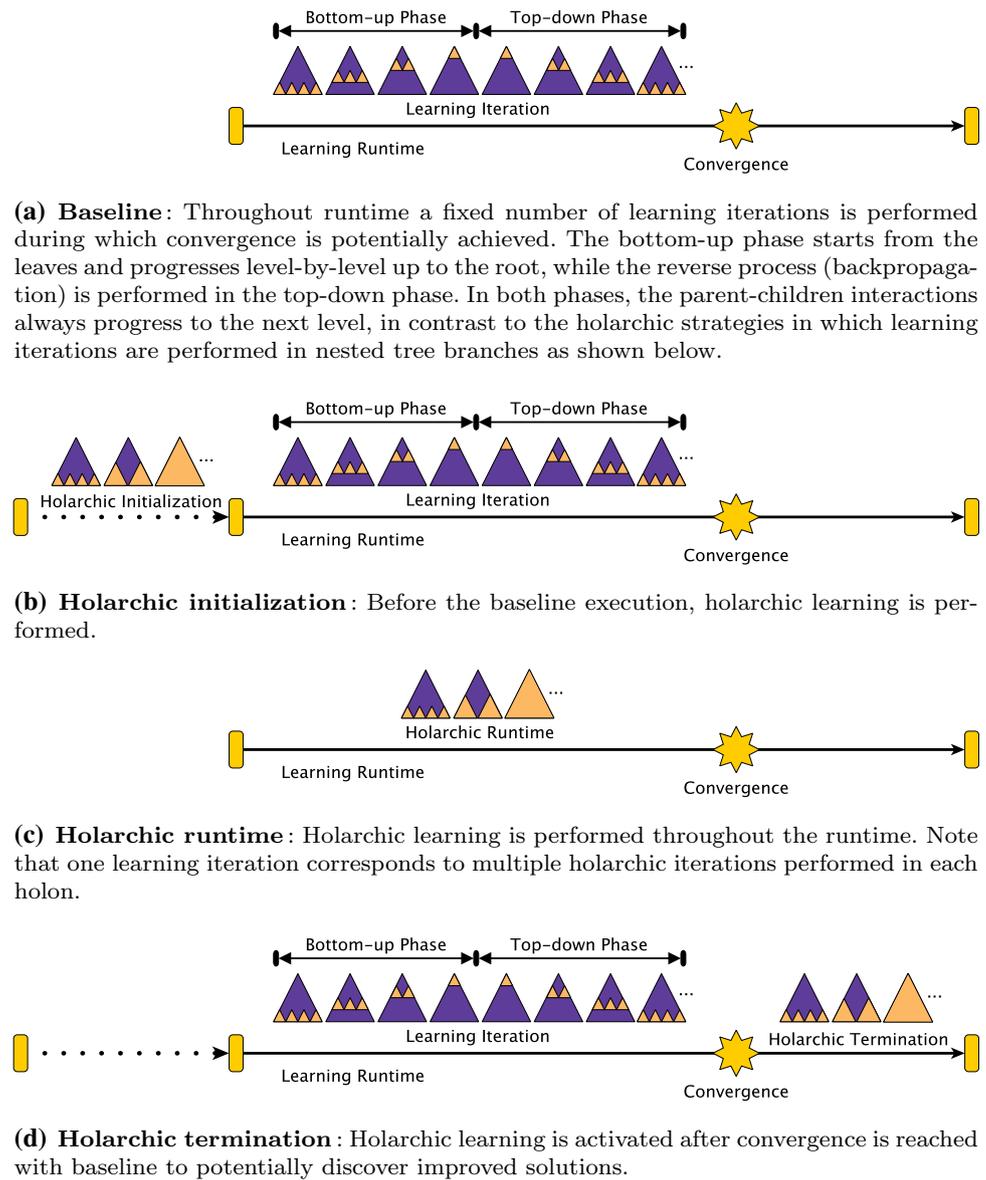
<sup>1</sup> Available at [www.epos-net.org](http://www.epos-net.org) (last accessed: September 2018)

I-EPOS consists of agents that autonomously plan resources they consume and produce. Planning is a process of *resource scheduling* or *resource allocation*. For instance agents may represent smart phone apps (personal assistants), cyber-physical controllers or smart home information systems with the capability to plan the energy consumption of residential appliances, the charging of electric vehicles or the choices of bike sharing stations. Planning serves the local resource requirements of users as well as system-wide objectives, for instance, the decrease of demand oscillations in the whole power grid to prevent blackouts [39], the synchronization of power demand with the availability of renewables [1] or the load-balancing of bike sharing stations to decrease the operational costs of manual bike relocations [40, 50].

I-EPOS introduces the computational model of locally generated *possible plans* that represent users' operational flexibility on how resources can be consumed or produced. For instance, a user may turn on a laundry machine earlier or later in time, can choose among two or more stations to return a shared bike, etc. Computationally, plans are vectors of real values and agents need to collectively choose one and only one of these plans to execute so that the summation of all these selected plans satisfies a system-wide objective measured by a *global cost* function. This paper focuses on the computational problem of the variance minimization that is a quadratic cost function, which cannot be locally minimized: coordination and collective decision-making is required [43]. Choosing the optimum combination of plans, whose summation minimizes a quadratic cost function is a non-convex combinatorial optimization problem known to be NP-hard [41].

Among the global cost that captures system-wide objectives, agents can also assess a *local cost* of their possible plans that may represent a notion of *discomfort* or *inconvenience* [37]. For instance, the longer the time period a user shifts the energy consumption earlier or later in time

**Fig. 2** The four learning schemes studied and compared in this paper



to prevent a power peak (global cost reduction) the higher the level of disruption is in the regular residential activities of a user (local cost increase). The agents’ choices can be autonomously parameterized to increase or decrease the priority of minimizing the local cost over the global cost. This trade-off is regulated by the  $\lambda$  parameter for each agent. A  $\lambda = 0$  results in choices that exclusively minimize the global cost, i.e. variance, and ignores the local cost. In contrast, a  $\lambda > 0$  biases the agents’ choices to favor plans with a lower local cost.

Learning is performed as follows: agents self-organize [36] in a tree network topology over which collective decision-making is performed—a plan is chosen by taking into account the following aggregate information (element-wise summation of vectors): (i) the aggregate plan choices of the agents in the branch underneath made available

during the bottom-up phase and (ii) the aggregate plan choices of all agents at the previous learning iteration made available during the top-down phase. Note that decision-making remains highly localized and decentralized as the planning information of the other agents is always at an aggregate level, i.e. the possible plans of other agents are not explicitly required.

Several of the following factors influence the learning performance defined by the level of the global cost reduction and the number of learning iterations required to minimize variance: (i) the positioning of the agents in the tree that determines the order of the collective choices made, (ii) the  $\lambda$  parameter that regulates the trade-off of global versus local cost and (iii) the overall topological structure and specifically in this paper the number of children  $c$  in balanced trees is studied.

Improving the learning performance by repositioning the agents in the tree or adapting the topology is complex and costly in distributed environments as the aforementioned self-organization methodologies are based on supplementary distributed protocols that consume resources, i.e. exchange of messages, computational power, energy, etc. Moreover, any node or link failure limits the computation of the aggregate plans at a branch level and therefore the collective decision-making cannot be anymore performed over all participating agents. The applicability of the holarchic schemes does not require any change in the logic of I-EPOS. The algorithm is localized and applied within multiple nested and connected branches even when the network becomes disconnected due to node and link failures: a disconnected agent triggers adaptation by children that turn to roots of holons and initiate the top-down phase of I-EPOS. The fact that a number of agents is isolated and does not participate in the learning process is the self-adaptation means to traverse the optimization space via alternative pathways. This has the potential to explore improving solutions, mitigate performance loss compared to a total interruption of I-EPOS, or even boost the reduction of the global cost that cannot be decreased anymore with I-EPOS.

Note that I-EPOS is selected to evaluate the applicability of holarchic schemes as it is a fully decentralized and hierarchical learning algorithm. The goal of learning in terms of whether it is designed for classification, prediction, pattern recognition, etc. or whether textual or image data are used, do not influence the applicability of holarchic schemes. The rest of this paper illustrates an empirical performance analysis of the three holarchic schemes and their applicability to I-EPOS.

## 4 Experimental methodology

This section illustrates the system prototyping, the varying dimensions in the experiments and the experimental settings. It also illustrates the evaluation metrics used to assess the holarchic learning schemes.

### 4.1 Prototyping and test runs

An open-source implementation of I-EPOS<sup>2</sup> is used for the prototyping of the holarchic schemes.<sup>3</sup> The software is implemented using the Protopeer distributed prototyping toolkit [13] and is designed to run in a simulation and live

mode. A software artifact<sup>4</sup> of EPOS for the broader scientific community is available for further research and evaluations [35, 40]. The actual experiments are deployed and parallelized in the Euler<sup>5</sup> cluster computing infrastructure of ETH Zurich.

### 4.2 Varying dimensions and performed experiments

The following system dimensions are studied: (i) *application scenarios*, (ii) *holarchic schemes*, (iii) *scale of holarchy*, (iv) *number of children  $c$  in the tree network*, (v) *different agent preferences  $\lambda$* .

Synthetic and empirical plans are generated for 1000 agents using data from real-world pilot projects. These four application scenarios are referred to as follows: (i) *synthetic*, (ii) *bike sharing*, (iii) *energy demand* and (iv) *electric vehicles*.

The synthetic dataset consists of 16 possible plans of size 100 generated from a standard normal distribution with a mean of 0 and a standard deviation of 1. A random local cost is assigned to the plans, i.e. the index of the plan represents its cost.

The bike sharing dataset<sup>6</sup> of the Hubway bike sharing system<sup>7</sup> in Paris is used to generate a varying number of plans of size 98 for each agent based on the unique historic trips performed by each user. Therefore, the plans represent the trip profiles of the users and they contain the number of incoming/outgoing bike changes made by each user in every station [38]. The local cost of each plan is defined by the likelihood of a user to not perform a trip instructed in the plan [40]. For instance, if three plans are chosen 4, 5 and 1 days during the measured time period respectively, the local cost for these plans is 0.6, 0.5 and 0.9 respectively.

The energy demand dataset<sup>8</sup> is generated via disaggregation of aggregate load obtained from the Pacific

<sup>2</sup> Available at <https://github.com/epoumaras/EPOS> (last accessed: September 2018).

<sup>3</sup> Available at <https://github.com/ysrivatsan/EPOS/tree/Srivatsan> (last accessed: September 2018).

<sup>4</sup> Available at <http://epos-net.org/software/exemplar/> (last accessed: September 2018).

<sup>5</sup> Available at <https://scicomp.ethz.ch/wiki/Euler> (last accessed: September 2018).

<sup>6</sup> Available at <http://epos-net.org/shared/datasets/EPOS-BICYCLES.zip> (last accessed: September 2018).

<sup>7</sup> The plans are generated using the dataset made available in the context of the Hubway Data Visualization Challenge: <http://hubwaydatachallenge.org/> (last accessed: September 2018). Although this dataset does not contain personalized records, user trips are extracted from user information: zip-code, year of birth and gender. All trips that have common values in these fields are assumed to be made by the same user. The timeslot is chosen from 8:00 a.m. to 10:00 a.m. All historic unique trips a user did in the defined timeslot of a week day are considered as the possible plans for that day.

<sup>8</sup> Available at <http://epos-net.org/shared/datasets/EPOS-ENERGY-SUBSET.zip> (last accessed: September 2018).

Northwest Smart Grid Demonstration Project (PNW) by Battelle.<sup>9</sup> The disaggregation algorithm and the raw data are illustrated in earlier work [39]. The generated dataset contains 5600 agents representing residential consumers. Every agent has 10 possible plans, each with length 144 containing electricity consumption records for every 5 min. The first plan corresponds to the raw data. The next three possible plans are obtained via the SHUFFLE generation scheme [39] that randomly permutes the values of the first plan. The next three plans are generated with the SWAP-15 generation scheme [39] that randomly picks a pair of values from the first plan and swaps their values. The process repeats 15 times. Respectively, the last three plans are generated by SWAP-30 [39] that applies the same process 30 times. The local cost of each plan represents the level of perturbation introduced on the first plan, i.e. on the disaggregated data, by the plan generation scheme. It is measured by the standard deviation of the difference between the raw and the perturbed plan values, element-wise:  $\sigma(x_1 - y_1, \dots, x_{144} - y_{144})$ .

The plans<sup>10</sup> for the electric vehicles are generated using data<sup>11</sup> from the Household Travel Survey of the California Department of Transportation during 2010–2012. The plans concern the energy consumption of the electric vehicles by charging from the power grid. Four plans per agent with size 1440 are generated by extracting the vehicle utilization using the historical data and then computing the state of charge by redistributing the charging times over different time slots. The methodology is outlined in detail in earlier work [38]. The local cost of each plan is measured by the likelihood of the vehicle utilization during the selected charging times.

The learning schemes studied are the *baseline* that is default I-EPOS and the three holarchic schemes: *Holarchic initialization* is used as an *exploration strategy* to evaluate its likelihood to improve the learning capacity. *Holarchic runtime* is used as a *mitigation strategy* to evaluate the maintenance of learning capacity in distributed environments under uncertainties. Finally, *holarchic termination* is used as a *boosting strategy* to evaluate the likelihood of improving the learning capacity. In each holon of the holarchic schemes within one main learning iteration,  $\tau = 5$  holarchic iterations are executed.

Two holarchic scales are evaluated: (i) *full* and (ii) *partial*. The full scale uses all levels of the baseline tree network to apply a holarchic scheme as also shown in

Fig. 1. In contrast, partial scale is applied in one branch under the root.

The influence of the topological tree structure and agent preferences on the learning capacity is studied by varying respectively the number of children as  $c = 2, \dots, 5$  and the  $\lambda$  parameter as  $\lambda = 0, 0.25, 0.5, 0.75$ . Table 1 summarizes the dimensions and their variations in the performed experiments.

The total experiments performed are calculated as follows: For the dimension of application scenarios, 4 variations are counted and 3 variations for the learning schemes given that the baseline and the holarchic termination can be generated within one experiment. The partial scale uses all possible branch combinations:  $2 + 3 + 4 + 5 = 14$  variations plus 4 variations for the full scale result in 18 total variations for the two dimensions of holarchic scale and number of children. Finally, 4 variations are counted for the dimension of agent preferences. The total number of  $4 * 3 * 18 * 4 = 864$  variation combinations is the total number of experiments performed. Each experimental combination is repeated 1000 times by (i) 1000 samples of possible plans in the synthetic scenario and (ii) 1000 random assignments of the agents in the tree network in the other three application scenarios. Therefore, the total number of experiments performed is  $864 * 1000 = 864,000$ .

### 4.3 Evaluation metrics

Performance is evaluated with the following metrics: (i) *standardized global cost*, (ii) *improvement index* and (iii) *communication cost*.

The standardized global cost is the variance of the global plan at convergence time. The minimization of the variance is the optimization objective and therefore the variance is used as the criterion of the learning capacity. Standardization<sup>12</sup> is applied on the variance so that the learning capacity among different datasets can be compared.

The improvement index  $I$  measures the reduction or increase of the global cost at convergence time for the holarchic schemes compared to the baseline. Positive values indicate an improvement of the baseline, while negative values show a deterioration. The improvement index is measured as follows:

$$I = \frac{C_b - C_h}{C_b + C_h} \quad (1)$$

where  $C_b$  is the global cost for the baseline and  $C_h$  is the global cost for a holarchic scheme, both at convergence

<sup>9</sup> Available upon request at <http://www.pnwsmartgrid.org/participants.asp> (last accessed: September 2018)

<sup>10</sup> Available at <http://epos-net.org/shared/datasets/EPOS-ELECTRIC-VEHICLES.zip> (last accessed: September 2018).

<sup>11</sup> Available at [www.nrel.gov/tsdc](http://www.nrel.gov/tsdc) (last accessed: September 2018). Electric vehicles equipped with GPS are selected.

<sup>12</sup> Standardization transforms the global cost values to have zero mean and units of variance as follows:  $\frac{x-\mu}{\sigma}$

**Table 1** Dimensions and their variations in the total of 864,000 experiments

Dimension	Variation 1	Variation 2	Variation 3	Variation 4
Application scenario	Synthetic	Bike sharing	Energy demand	Electric vehicles
Learning scheme	Baseline	Holarchic initialization	Holarchic runtime	Holarchic termination
Holarchic scale	Full	Partial	–	–
Number of children ( $c$ )	$c = 2$	$c = 3$	$c = 4$	$c = 5$
Agent preferences ( $\lambda$ )	$\lambda = 0$	$\lambda = 0.25$	$\lambda = 0.5$	$\lambda = 0.75$
Total experiments:	864,000			

time. The improvement index is calculated based on the principle of the symmetric mean absolute error, which compared to the mean absolute error can handle single zero values, it is bound to  $[-1, 1]$  and eliminates very large values originated by low denominators [21].

The communication cost measures the number of messages exchanged to complete a learning iteration and can be distinguished to *total* and *synchronized*. The total communication cost counts all exchanged messages between the agents during a learning iteration and it is calculated for the baseline  $M_b$  as follows:

$$M_b = 2(c^0 + \dots + c^l - 1) \quad (2)$$

where  $c$  is the number of children in a balanced tree. Equation 2 sums up the number of agents in each level of the tree and subtracts 1 to count the number of links. Multiplication by 2 counts both bottom-up and top-down phases. The total communication cost of a holarchic scheme can be measured as follows:

$$M_t = 2\tau \sum_{j=0}^l c^{l-j}(c^0 + \dots + c^j - 1) \quad (3)$$

where  $c$  is here again the number of children in a balanced tree/holarchy and  $\tau$  is the number of holarchic iterations. The summation starts from leaves (level  $j = 0$ ), and progresses to the root of the holarchy (at level  $j = l$ ). Equation 3 multiplies the number of agents  $c^{l-j}$  at each level  $l - j$  with  $2\tau$  times ( $\tau$  bottom-up and  $\tau$  top-down holarchic iterations) the number of agents in the branches underneath:  $c^0 + \dots + c^j - 1$ . For example,  $j = 0$  corresponds to Fig. 1a with a communication cost of  $2\tau 2^2(2^0 + 2^1 - 1) = 16\tau$ .  $j = 1$  corresponds to Fig. 1b with a communication cost of  $2\tau 2^1(2^0 + 2^1 + 2^2 - 1) = 24\tau$  and respectively, for  $j = 2$  and Fig. 1c communication cost is calculated as  $2\tau 2^0(2^0 + 2^1 + 2^2 + 2^3 - 1) = 28\tau$ . These nested calculations for the full holarchy sum up to  $M_t = (16 + 24 + 28)\tau = 68\tau$  messages.

The synchronized communication cost counts the number of messages exchanged within holons, while counting this number only once for holons at the same level

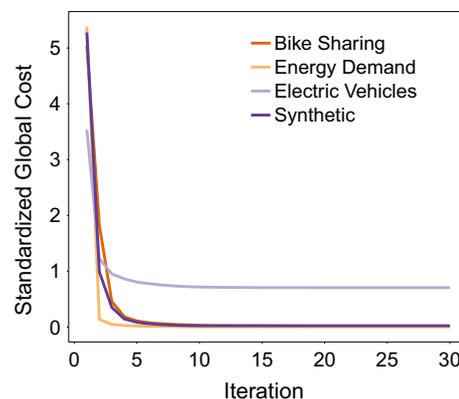
which can exchange messages in parallel. For instance, Fig. 1a illustrates four parallel holons with a total communication cost of  $2\tau 2^2(2^0 + 2^1 - 1) = 16\tau$  messages. Instead, the synchronized communication cost counts for  $2\tau(2^0 + 2^1 - 1) = 4\tau$  messages. The synchronized communication cost of a full holarchy can be measured as follows:

$$M_s = 2\tau \sum_{j=0}^l (c^0 + \dots + c^j - 1) \quad (4)$$

where  $c$  is the number of children and  $\tau$  the number of holarchic iterations. Note that the synchronized communication cost considers holons performing in parallel and not individual agents, that is why within each holon all messages exchanged are counted. For the same reason, the synchronized communication cost for baseline is not defined for a fairer comparison with holarchic schemes.

#### 4.4 Computational challenge

To better understand the computational challenge and context in which the performance of the holarchic schemes is studied, the performance characteristics of the baseline are illustrated in this section. Figure 3 illustrates the



**Fig. 3** Learning curves of I-EPOS for the four benchmark application scenarios with 1000 agents generating 4 plans and choosing a plan with  $\lambda = 0$

learning curves of the baseline in the four application scenarios.

The learning performance of I-EPOS shows the following behavior: Global cost decreases dramatically in very few iterations in all application scenarios, while the decrease is monotonous. Therefore, I-EPOS has a superior efficiency to learn fast combinations of plans that minimize the variance by executing 10–15 iterations. Convergence does not necessarily mean that the globally optimum solution is found. The evaluation of the global optimality in a system with 1000 agents and more than 4 plans per agents is computationally infeasible given the exponential complexity of the combinatorial space:  $4^{1000}$ . For this reason, the global optimality is evaluated using brute force search in a small-scale system of 10 agents with 4 plans per agent and  $\lambda = 0$ . Therefore the total number of solutions is  $4^{10}$ . Each experiment is repeated 10 times by shuffling the agents over a binary tree. Fig. 4 illustrates the global optimality results for each application scenario.

I-EPOS finds the 0.007%, 0%, 0.017% and 0.153% top solution in each of the application scenarios of Fig. 4. Note that such a significant optimality may not be achieved for systems with thousands of agents and several plans. Nevertheless, designing a new learning scheme to overpass this performance level, without introducing additional complexity and resources is a challenge and potentially not an endeavor worth pursuing. Instead this paper studies holarchic structures as learning strategies to explore, mitigate and boost the cost-effectiveness of I-EPOS in distributed environments and in this sense the notion of holarchy is the means for decentralized learning to tolerate their uncertainties.

## 5 Experimental evaluation

This section illustrates the learning capacity of the three holarchic schemes followed by the trade-offs and cost-effectiveness of the holarchic runtime. All main experimental findings are covered in this section and an appendix

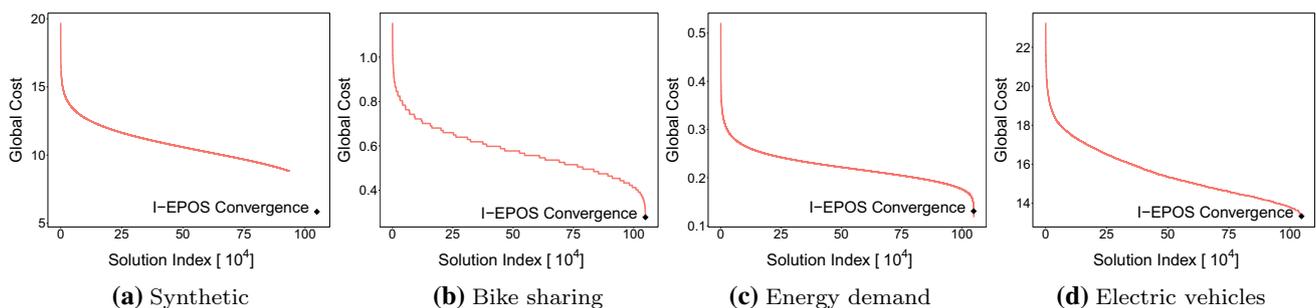
provides supplementary results that cover the broader range of the varying parameters. The results in the appendix are included for the sake of completeness of the paper and future reference.

### 5.1 Learning capacity

Figure 5 illustrates the learning curves for the four application scenarios and learning schemes. The partial scale with  $\lambda = 0$  and  $c = 2$  is illustrated. The learning curves for the full scale,  $\lambda = 0.5$  and  $c = 5$  are illustrated in Fig. 12 of Appendix 1.

The following observations can be made in Fig. 5: Holarchic runtime achieves the fastest convergence speed given the several multi-level holarchic iterations performed within a main learning iteration. However, a performance sacrifice in global cost is observed, which is though low and observable for the scenario of electric vehicles in which the global cost is 10% higher than the baseline. Moreover, within 7–8 iterations all holarchic schemes achieve the global cost reduction of the baseline. The convergence speed of the holarchic initialization is 1–2 iterations slower than the baseline, though this insignificant difference is not anymore observable in the scenario of electric vehicles.

Figure 6 illustrates the improvement index of the holarchic schemes for different application scenarios and different  $\lambda$  values. The exploration potential of the holarchic initialization is indicated by the error bars above the mean. Moreover, holarchic initialization does not influence significantly the global cost reduction, however, for higher  $\lambda$  values, i.e.  $\lambda = 0.75$ , an average increase of 0.5% is observed in the improvement index. This means that in more constrained optimization settings, the exploration strategy of the holarchic initialization contributes a low performance improvement. In contrast, the mitigation strategy of the holarchic runtime manages to preserve the baseline performance with an improvement index of values close to 0. An average performance boosting of 1.65% is observed via the holarchic termination in the bike sharing



**Fig. 4** Global optimality of I-EPOS for the four benchmark application scenarios with 10 agents generating 4 plans and choosing a plan with  $\lambda = 0$

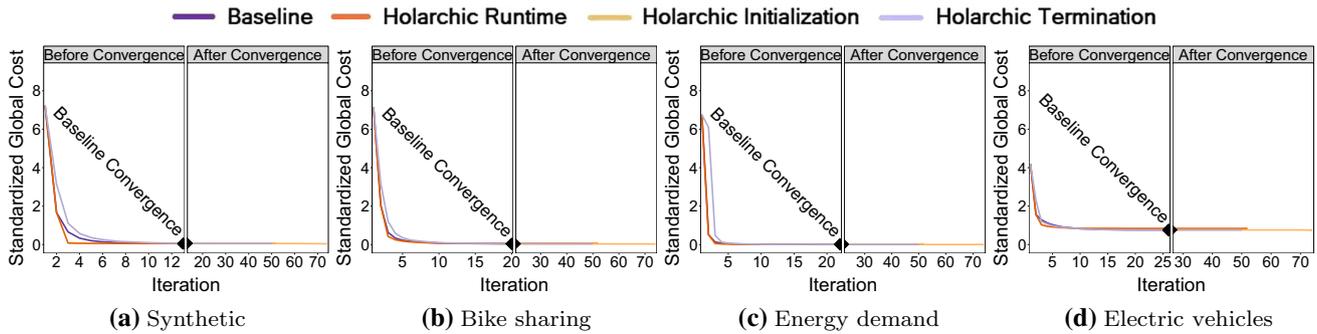


Fig. 5 Learning curves. *Dimensions* learning schemes, application scenarios. *Settings* partial scale,  $\lambda = 0$ ,  $c = 2$

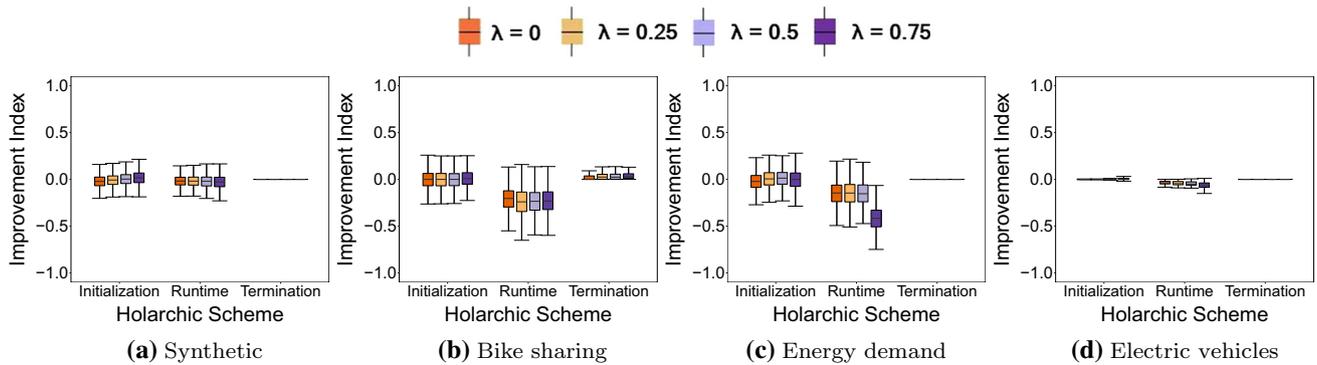


Fig. 6 Improvement index. *Dimensions* holarchic schemes, application scenarios, different  $\lambda$  values. *Settings* partial scale,  $c = 2$

scenario that has sparse data, while in the other scenarios no significant improvement is observed.

Figure 7 illustrates the improvement index of the holarchic schemes for different application scenarios and different  $c$  values. Holarchic initialization retains an average improvement index of  $-0.008$  while it can scale up the improvement index to values of  $0.225$  on average. The number of children does not influence the performance of this holarchic scheme. In contrast, the holarchic runtime shows an average increase of  $4.1\%$  in the improvement index by increasing  $c$  from  $2$  to  $5$ , while this holarchic scheme serves well its performance mitigation role: an average improvement index of  $-0.041$ . Finally, holarchic

termination boosts performance by  $1.1\%$  in the bike sharing scenario.

Figure 8 demonstrates the higher performance that the partial scale shows compared to full scale:  $0.25\%$  higher improvement index for holarchic initialization and  $6.5\%$  for holarchic runtime.

Figure 13, 14 and 15 of Appendix 1 show the probability density of the improvement index for different  $\lambda$ ,  $c$  values and holarchic scales respectively. In these figures one can study in more detail the density of the improvement index values behind the error bars of the respective Fig. 6, 7 and 8.

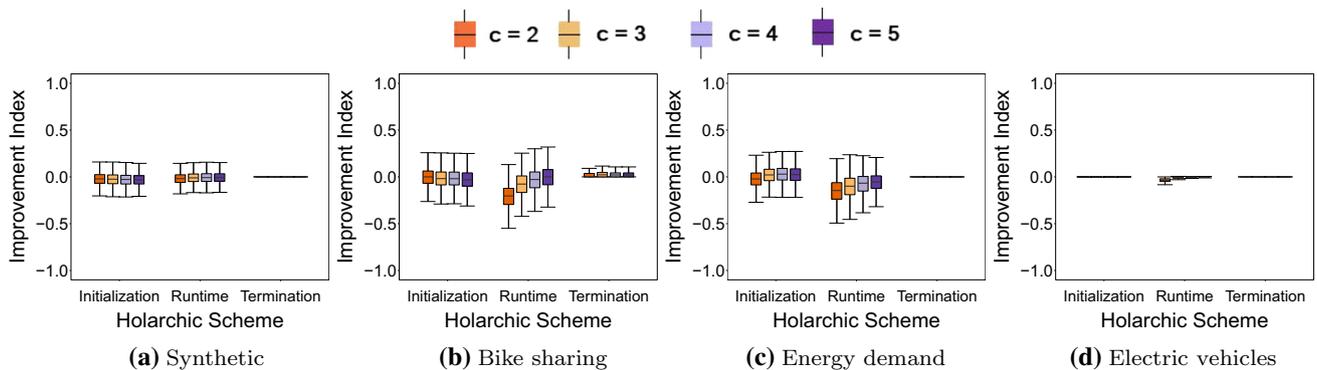
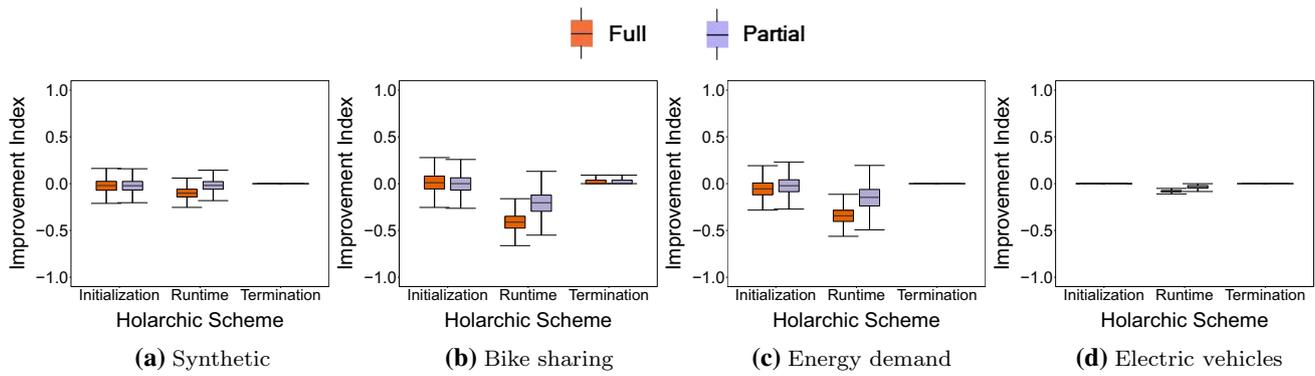
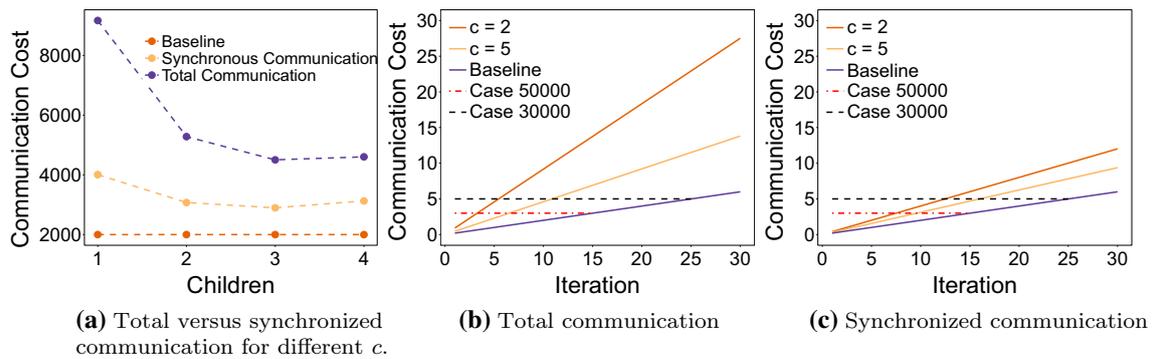


Fig. 7 Improvement index. *Dimensions* holarchic schemes, application scenarios, varying number of children. *Settings* partial scale,  $\lambda = 0$



**Fig. 8** Improvement index. *Dimensions* holarchic schemes, application scenarios, partial versus full scale. *Settings*  $\lambda = 0, c = 2$



**Fig. 9** Communication cost. *Settings* partial scale,  $\lambda = 0$  **a** *Dimensions* varying number of children, total versus synchronized communication cost, baseline versus holarchic runtime. **b** and **c** Sampling

case 30,000 and case 50,000 under total versus synchronized communication cost. *Dimensions*  $c = 2$  versus  $c = 5$ , number of iterations, baseline versus holarchic runtime

The rest of this section studies the trade-offs and the cost-effectiveness of the holarchic runtime designed for the mitigation of the learning performance.

### 5.2 Trade-offs and cost-effectiveness

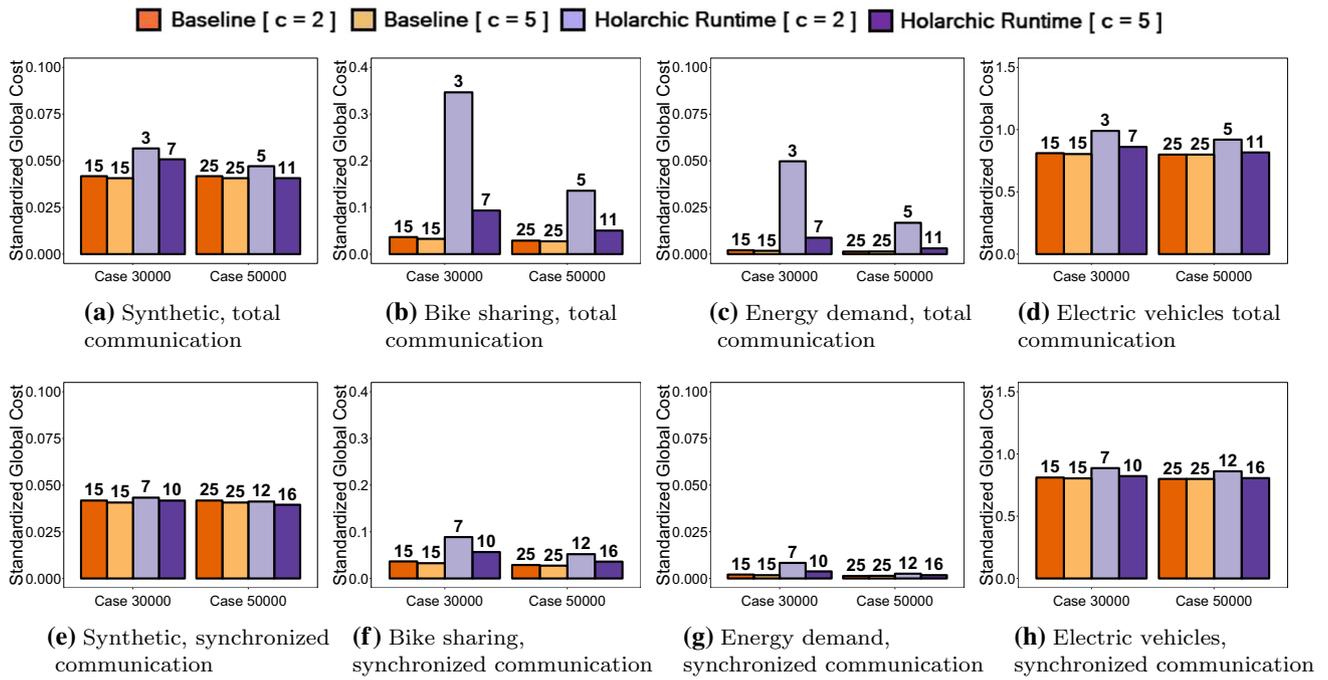
Figure 9a illustrates the communication cost per iteration of the baseline versus the total and synchronized communication cost of the holarchic runtime. This is a worse case scenario as applying the holarchy to a smaller branch or for a fewer than 5 holarchic iterations can make the communication cost equivalent<sup>13</sup> to the one of the baseline. In Fig. 9a, the communication cost of the holarchic runtime decreases as the number of children increases given that the recursion of the holarchy is limited to a lower number of levels in the tree, i.e. fewer holons are formed. The synchronized communication cost is on average 45% lower than the total communication cost.

The cost-effectiveness of the holarchic runtime is studied by fixing the communication cost for both baseline and holarching runtime and looking into the global cost reduction achieved for the same number of messages exchanged. Figure 9b, c illustrate this process for total and synchronous communication cost. Two cases are determined: (i) *case 30,000* and (ii) *case 50,000*. Each case runs for a given number of iterations that is determined by the intersection with the horizontal dashed lines for each of the  $c = 2$  and  $c = 5$ . Then, the global costs can be compared for the same number of exchanged messages as shown in Fig. 10.

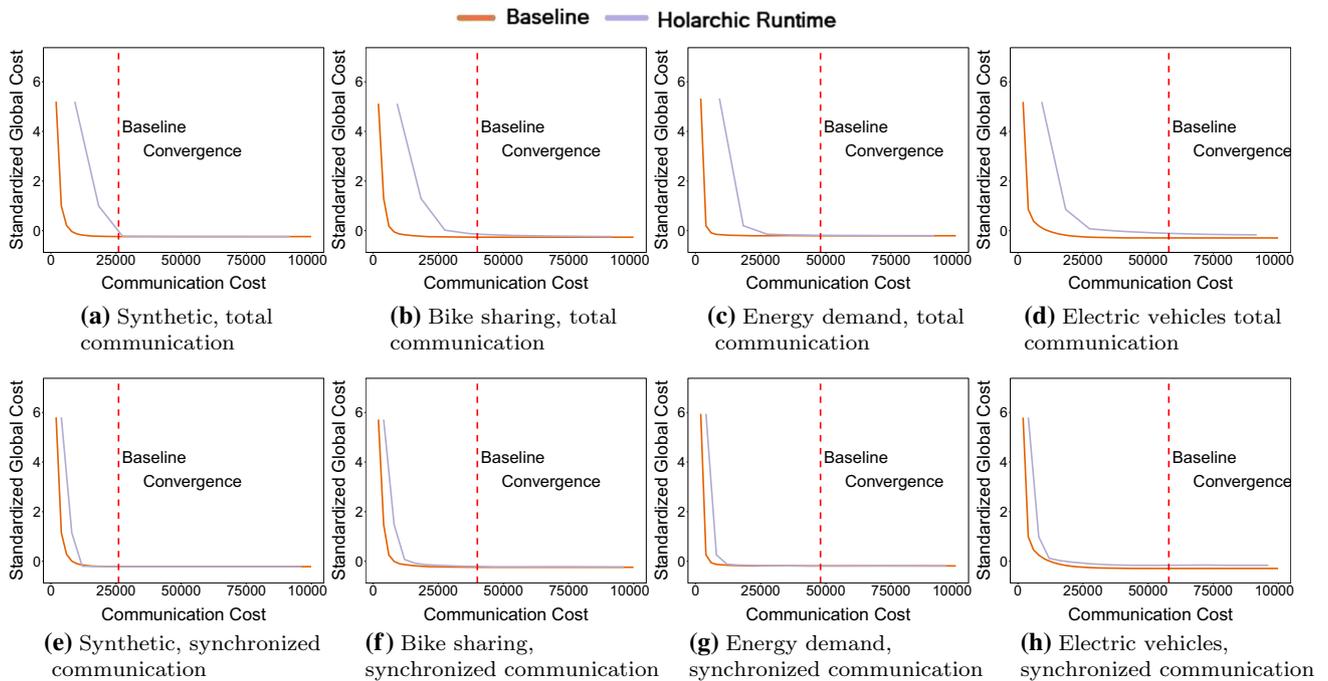
The following observations can be made in Fig. 10a–d for the total communication cost: When  $c = 5$ , the holarchic runtime achieves a highly equivalent performance with the baseline. This also holds for  $c = 2$  in the scenarios of synthetic and electric vehicles. The performance mitigation becomes 13.64% more significant when the synchronized communication cost is counted in Fig. 10e–h. This is also shown by the shifted probability densities of the improvement index in Fig. 16 of Appendix 1.

These findings can be generalized further for the broader range of communication cost as shown in Fig. 11 for  $c = 2$ .

<sup>13</sup> The communication cost of the holarchic runtime can even become lower than baseline assuming a holarchy at partial scale with the agents that do not belong to the holarchy being disconnected.



**Fig. 10** Global cost. *Dimensions* baseline versus holarchic runtime, total versus synchronized communication cost, application scenarios,  $c = 2$  versus  $c = 5$ , number of iterations given a communication cost, case 30,000 versus case 50,000. *Settings* partial scale,  $\lambda = 0$



**Fig. 11** Cost-effectiveness. *Dimensions* baseline versus holarchic runtime, application scenarios, total versus synchronized communication cost. *Settings* partial scale,  $\lambda = 0$ ,  $c = 2$

The key observation that confirms the mitigation capability of the holarchic runtime is the comparable global cost achieved using the communication cost required for baseline to converge. This mitigation potential is also demonstrated by the probability density of the relative global cost

between baseline and holarchic runtime in Fig. 17 of Appendix 1.

Figure 11 can also be compared with Fig. 18 of Appendix 1 that shows the cost-effectiveness under  $c = 5$ .

The performance mitigation is even higher with this setting.

## 6 Summary of results and discussion

The following key observations can be made about the learning capacity of holarchic structures: (i) The limitation to contribute improvements in terms of global cost as motivated in Figs. 3 and 4 is confirmed in the performed experiments. (ii) The performance exploration, mitigation and boosting for which the holarchic schemes are designed are confirmed: a low and sporadic performance improvement is observed by holarchic initialization, especially in constrained environments of high  $\lambda$  values that justifies the exploration potential of this scheme. Similarly, the mitigation potential of the holarchic runtime is also confirmed by the fast convergence and preservation of the performance, especially for trees with a higher number of children  $c$ . Performance boosting by holarchic termination is scarce but observable in the bike sharing scenario that has sparser data. (iii) Strikingly, holarchies applied at a partial scale demonstrate higher cost-effectiveness than full scale in all holarchic schemes, i.e. lower cost due to the higher localization that limits communication cost and higher effectiveness in terms of higher improvement index.

In the trade-offs of cost-effectiveness, the synchronized communication cost is almost half of the total communication cost via a parallel execution of the learning process using holarchic structures. Moreover, for the same communication cost in baseline and holarchic runtime, the global costs become equivalent, especially in terms of the synchronized communication cost and for trees with a higher number of children  $c$ .

Overall the results suggest that several parallel and small-scale holarchic structures decrease the probability of trapping to local optima, require a lower parallelizable communication cost, while they better serve the purpose they are designed for: make decentralized deep learning more resilient to the uncertainties of distributed environments. Because of the high efficiency of I-EPOS to which holarchic structures are applied, conclusions cannot be reached whether the learning capacity of other mechanisms can be enhanced using holarchic structures. Nevertheless, another promising use case of holarchies is the multi-level optimization of complex techno-socio-economic systems formed by agents with different goals. In other words, the localization of the learning process in part of the deep hierarchical structure can represent the collective effort of a community to meet its goal that can differ or even oppose the goal of another community, i.e. another part of the deep hierarchical structure.

## 7 Positioning and comparison with related work

Earlier work [48, 49] identifies the key role that hierarchical design plays in dealing with the complexity [8] of large-scale and highly dynamic systems. More specifically, it is studied how such hierarchies can address the challenge of *limited rationality*, which stems from the combinatorial explosion of alternative system compositions. In the context of this paper, the combinatorial explosion results from the multitude of resource planning alternatives. Such hierarchies are later on introduced as ‘holarchies’ by emphasizing their recursive and self-encapsulated nature [22] that represent both an entire system (whole) and a mere subsystem (part).

General-purpose holonic designs that integrate the above principles are proposed via holonic multi-agent system platforms [44, 45]; multi-level modeling and simulation approaches [14]; and hierarchical problem solving [24]. Customised holonic designs are been applied to various domains, including hierarchical planning [33], traffic control [10], manufacturing [6, 15] and smart grids [11, 25, 47]. The success of these applications motivate the adoption of similar holonic principles for designing deep learning systems for large-scale distributed environments.

The design of I-EPOS [35, 40], used in this paper as case study and performance baseline, adopts a hierarchical approach, featuring: (i) bottom-up abstraction—via the aggregation of plans from the lower level; (ii) partial isolation—as different tree branches operate in parallel; and (iii) time tuning—by synchronizing the execution of hierarchical levels. However, the learning process in I-EPOS executes over the entire hierarchy—sequentially, level-by-level—rather than being nested within multiple holarchic levels.

Various deep learning techniques employ a hierarchical structure for different purposes. This depends on the nature of the learning problem, e.g. classification or prediction, and on the context within which the learning process executes, i.e. amount of input data and availability of computational resources. Hierarchical models are employed to deal with complex processing of data, by performing learning tasks progressively within incremental modelling levels, e.g. image classification [20, 46] and text categorisation [23]. Hierarchical structures are employed to process online large amounts of distributed input data, hence scaling up machine learning techniques [28, 34]. Such learning approaches rely on data partitions—pre-existing or artificially created—over which they distribute the learning process. The partial results from each partition are then collected and aggregated into an overall learning

outcome. The work of this paper is positioned and compared below with some of these approaches.

An earlier survey [34] compares distributed learning approaches with respect to their capability to (i) combine learning outcomes among heterogeneous representations of data partitions and (ii) deal with privacy constraints. The studied approaches feature a two-layer hierarchy: one layer for distributed learning across data partitions and a second layer for collecting and aggregating the results. In contrast, the holarchic learning design proposed in this paper introduces multiple levels to further to enhance the learning capacity in distributed environments under uncertainties. The localized holons operating in parallel reuse their learning outcomes even when the topology is clustered by node or link failures.

MLNet [28] introduces a special-purpose communication layer for distributed machine learning. It uses tree-based overlay networks to aggregate progressively partial learning results in order to reduce network traffic. This approach draws parallels with the design proposed here, while it is more applicable at the lower communication layers.

In dynamic environments, rescheduling, or reoptimization [29] becomes a critical function for dealing with unpredictable disturbances. It usually raises the additional constraints of minimizing reoptimization time as well as the distance between the initial optimization solution and the one of reoptimization, e.g. dynamic rescheduling in manufacturing systems [26] or shift rescheduling [29]. For instance, a two-level holarchy is earlier adopted to combine global optimization scheduling with fast rescheduling when dynamic system disturbances occur [29]. In contrast, the holarchic schemes of this paper do not undermine the learning performance, while adding resilience by localizing the learning process without reinitiating it.

Another line of relevant related work concerns the initial selection and maintenance of the topology over which the distributed learning process operates. Multi-agent approaches often rely on self-organization by changing their interactions and system structure at runtime. A holonic multi-agent approach for optimizing facility location problems is earlier introduced [32]. Facilities include distribution of bus stops, hospitals or schools within a geographical area. The agents react to mutual attraction and repulsion forces to self-organize into a holarchy. Stable solutions represent optimal facility localization distributions. This self-organization process has a considerable computational and communication cost, in case of remote agents' interactions. In contrast, the holarchic schemes studied in this paper preserve the agents' organizational structure, while the decentralized learning process used for system optimization is self-adapted by localizing the learning span within part of the tree network.

The experimental work of this paper shows how different topological configurations, i.e. agents' positioning and number of children, influence learning performance. The key role that such hyperparameters play in ensuring the effectiveness of learning approaches is also confirmed by related work on the optimization of the hierarchical structure and its configuration variables. Some examples include exhaustive search of all possibilities, (which, however, suffers from exponential combinatorics), random search, Bayesian optimization, e.g. in neural networks and deep belief systems [4, 5] as well as gradient-based optimization [27]. This is also relevant for deep learning applications via unsupervised pre-training [53] and evolutionary algorithms, e.g. in deep learning neural networks [30, 54]. In the context of this work, such approaches can be used to determine the most effective holarchic structures and hyperparameter configurations.

## 8 Conclusion and future work

This paper concludes that holarchic structures for decentralized deep learning can be a highly cost-effective organizational artifact for managing learning performance under uncertainties of distributed environments. The communication cost of self-organization can be eliminated by self-adapting the span of the learning process at a more localized level within the hierarchical structure as the means to cope with failures, latency and constrained computational resources. An extensive experimental evaluation with more than 864,000 experiments fed with synthetic and real-world data from pilot projects confirm the potential to explore, mitigate and boost the learning performance using three respective holarchic schemes applied to the I-EPOS decentralized deep learning system for solving combinatorial optimization problems.

Results show that the exploration of improving solutions is feasible and more likely to happen under stricter agents' constraints, while performance mitigation is more effective in balanced tree topologies with higher number of children. Boosting the learning performance via holarchic structures is challenging yet consistently observed under computational problems with sparse data, i.e. the bike sharing application scenario. The partial scale of the holarchic structures is more cost-effective than the full scale. Nevertheless, when the uncertainties of distributed environments are not anymore a constraint, holarchic schemes cannot outperform the cost-effectiveness of learning systems that make use of the whole hierarchical structure, which is a finding consistent with earlier work on greedy optimization and suboptimum heuristics trapped in local optima [3].

Mechanisms for an automated activation and deactivation of holarchic schemes as well as the applicability of these schemes in other more complex hierarchical structures than tree topologies are subject of future work. Applying the concept of holarchy in non-hierarchical structures such as the unstructured learning network of COHDA [19] can provide new means to control high communication costs, while preserving a high learning performance.

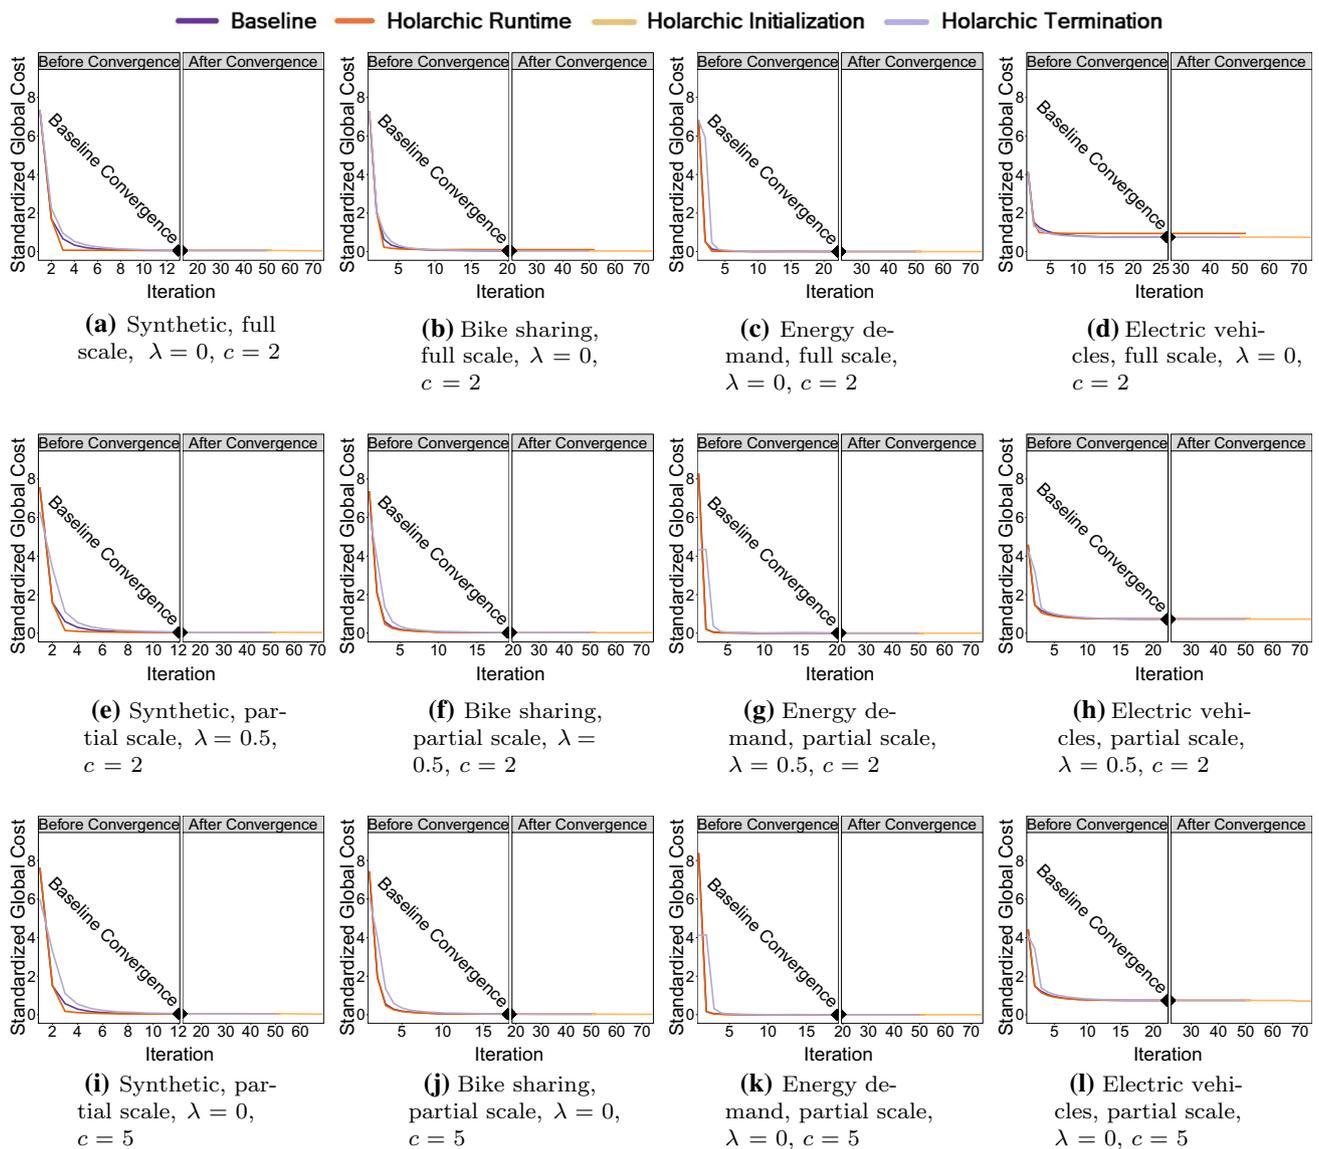
**Acknowledgements** This work is supported by the European Communitys H2020 Program under the scheme ‘INFRAIA-1-2014-2015: Research Infrastructures, Grant Agreement #654024 ‘SoBigData: Social Mining & Big Data Ecosystem (<http://www.sobigdata.eu>) and the European Communitys H2020 Program under the scheme ‘ICT-

10-2015 RIA’, Grant Agreement #688364 ‘ASSET: Instant Gratification for Collective Awareness and Sustainable Consumerism’.

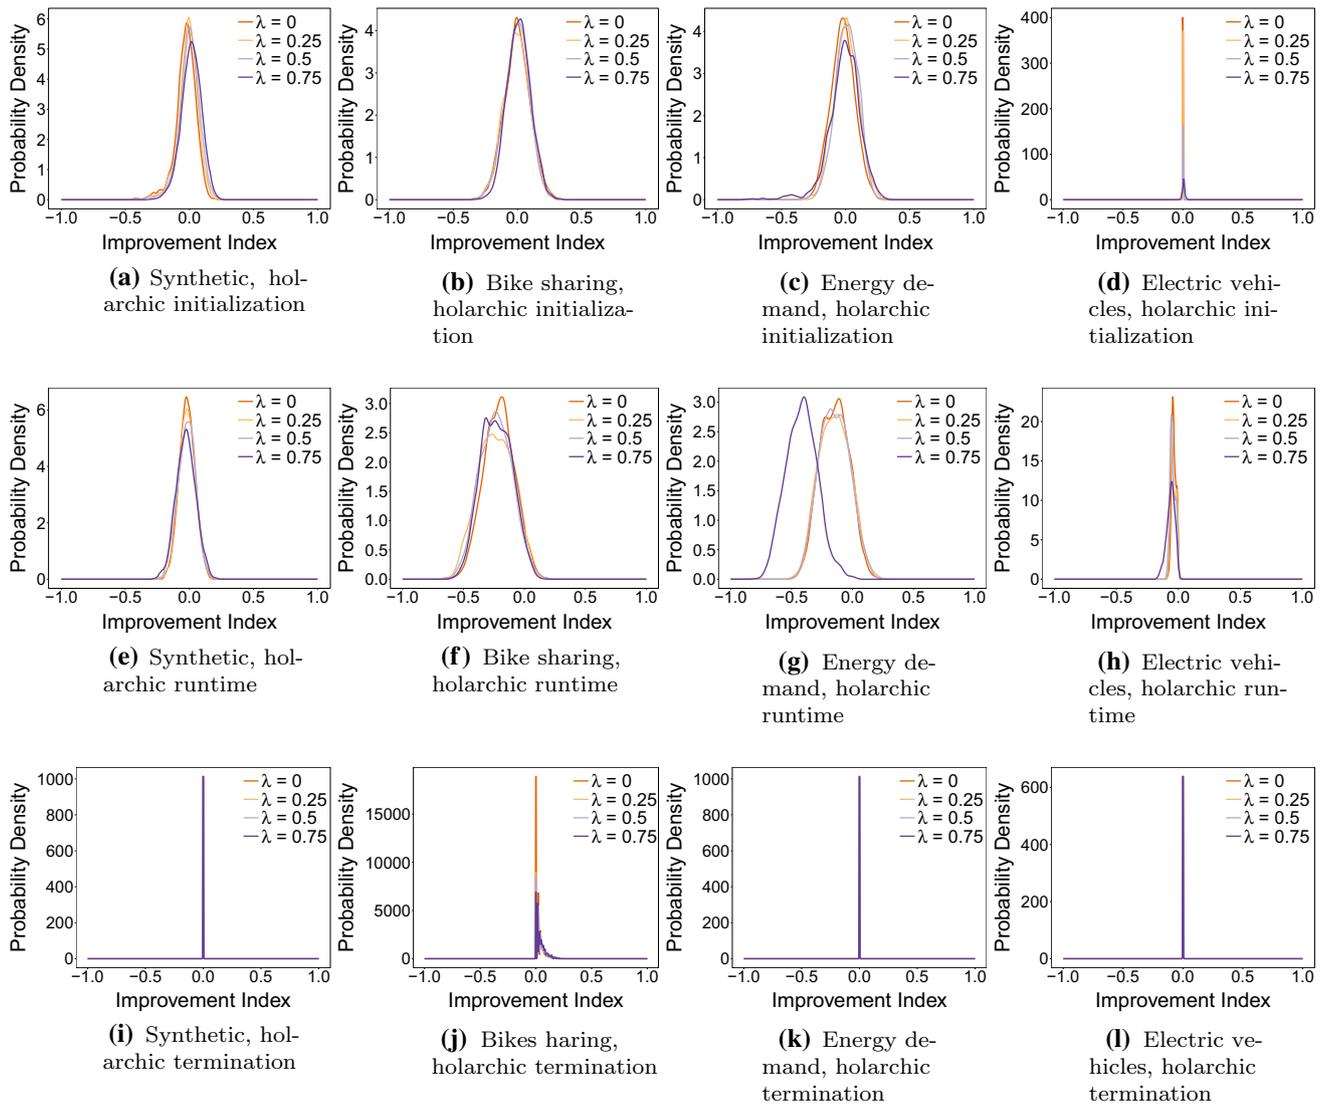
**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

### Appendix 1: Detailed experimental results

Figure 12 compares with Fig. 5 by varying partial scale to full (Fig. 12a–d),  $\lambda = 0$  to  $\lambda = 0.5$  (Fig. 12e–h) and  $c = 0$  to  $c = 5$  (Fig. 12i–l).



**Fig. 12** Learning curves for comparison with Fig. 5. Dimensions holarchic schemes, application scenarios, full versus partial scale, different  $\lambda$  values, varying number of children



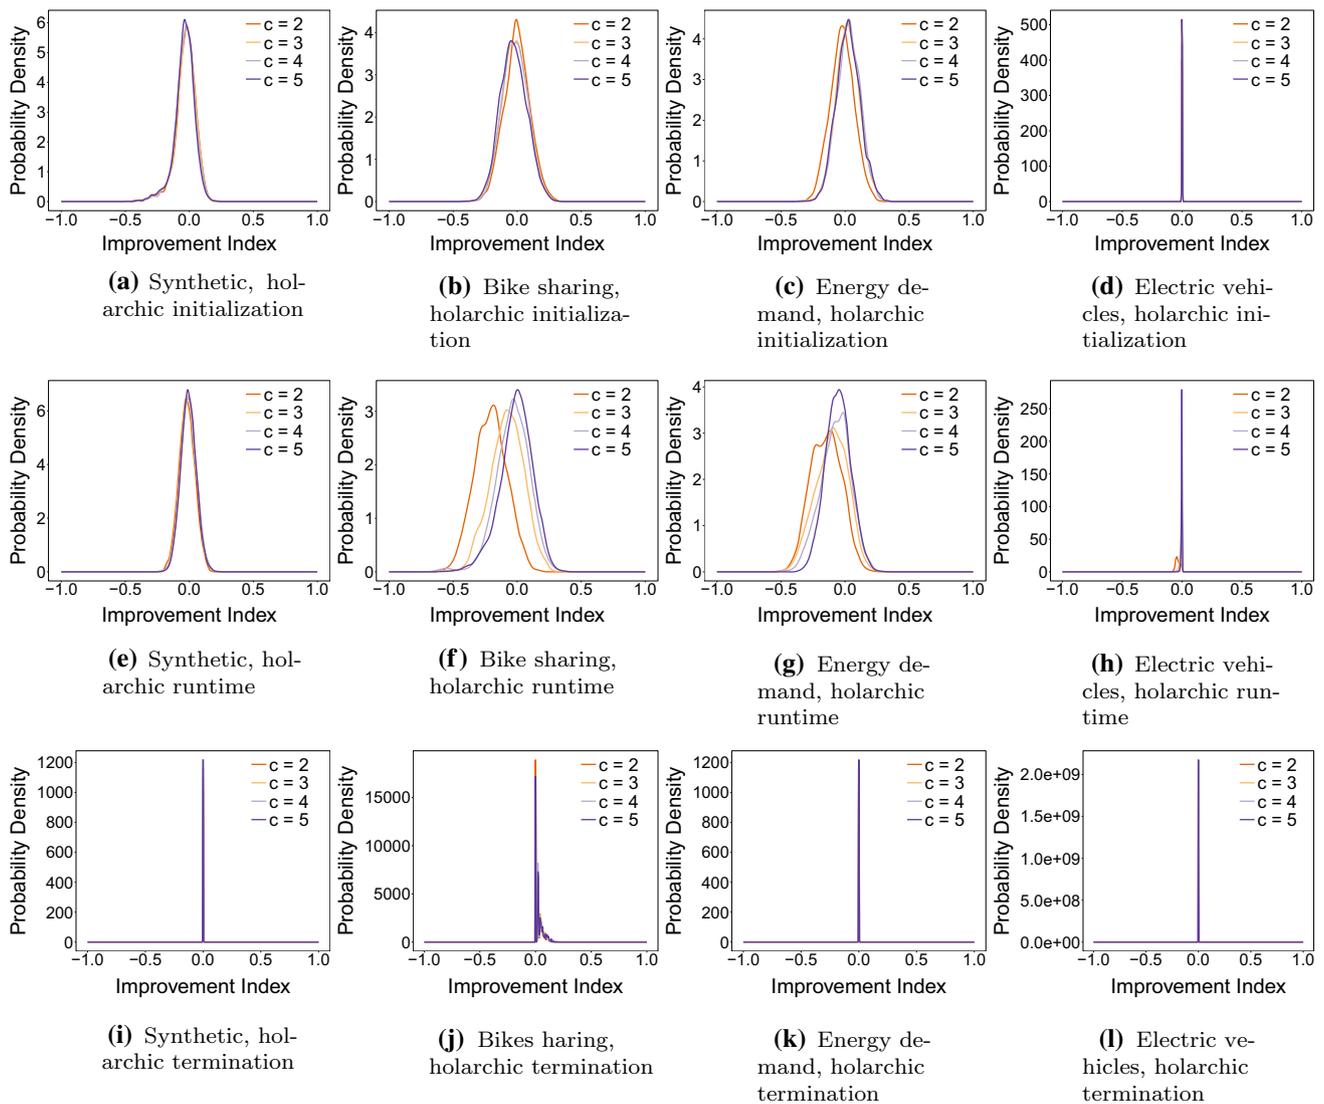
**Fig. 13** Probability density of the improvement index that elaborates on Fig. 6. *Dimensions* holarchic schemes, application scenarios, different  $\lambda$  values. *Settings* partial scale,  $c = 2$

Figure 13 elaborates on the Fig. 6. It illustrates the probability density function of the improvement index by fixing the holarchic scale to partial,  $c = 2$  and varying all other dimensions.

Figure 14 elaborates on the Fig. 7. It illustrates the probability density function of the improvement index by fixing the holarchic scale to partial,  $\lambda = 0$  and varying all other dimensions.

Figure 15 elaborates on the Fig. 8. It illustrates the probability density function of the improvement index by fixing  $\lambda = 0$ ,  $c = 2$  and varying all other dimensions.

Figure 16 contrasts the improvement index of the holarchic runtime on the basis of total versus synchronized communication cost. Results show an average increase of the improvement index by 9.5%, indicated by the shift to the right for synchronized communication.

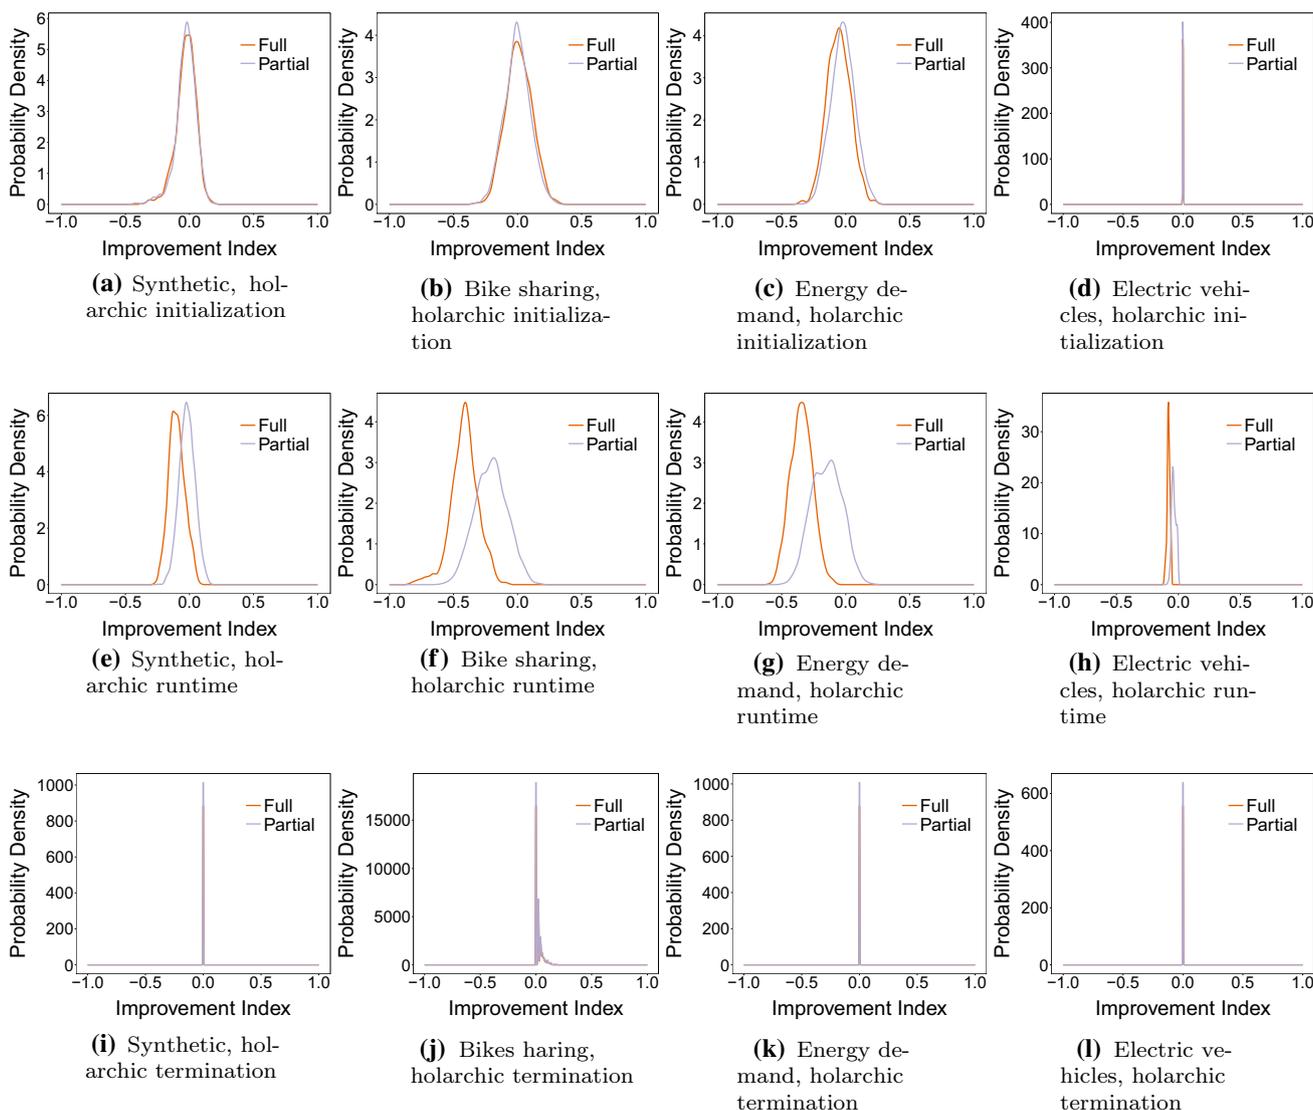


**Fig. 14** Probability density of the improvement index that elaborates on Fig. 7. *Dimensions* holarchic schemes, application scenarios, varying number of children. *Settings* partial scale,  $\lambda = 0$

The contrast of total versus synchronized communication cost is illustrated in Fig. 17 via the *relative performance* that is defined as follows:

$$P = \frac{C_h^{(1)} - C_h^{(T)}}{C_b^{(1)} - C_b^{(T)}} \quad (5)$$

where  $C_h^{(1)}, C_b^{(1)}$  is the global cost at the first iteration  $t = 1$  for holarchic runtime and baseline respectively, while  $C_h^{(T)}, C_b^{(T)}$  is the global cost at convergence  $t = T$ . This metric encodes the additional information of the improvement extent during the learning iterations. For instance, while Fig. 16f shows a density with improvement index values

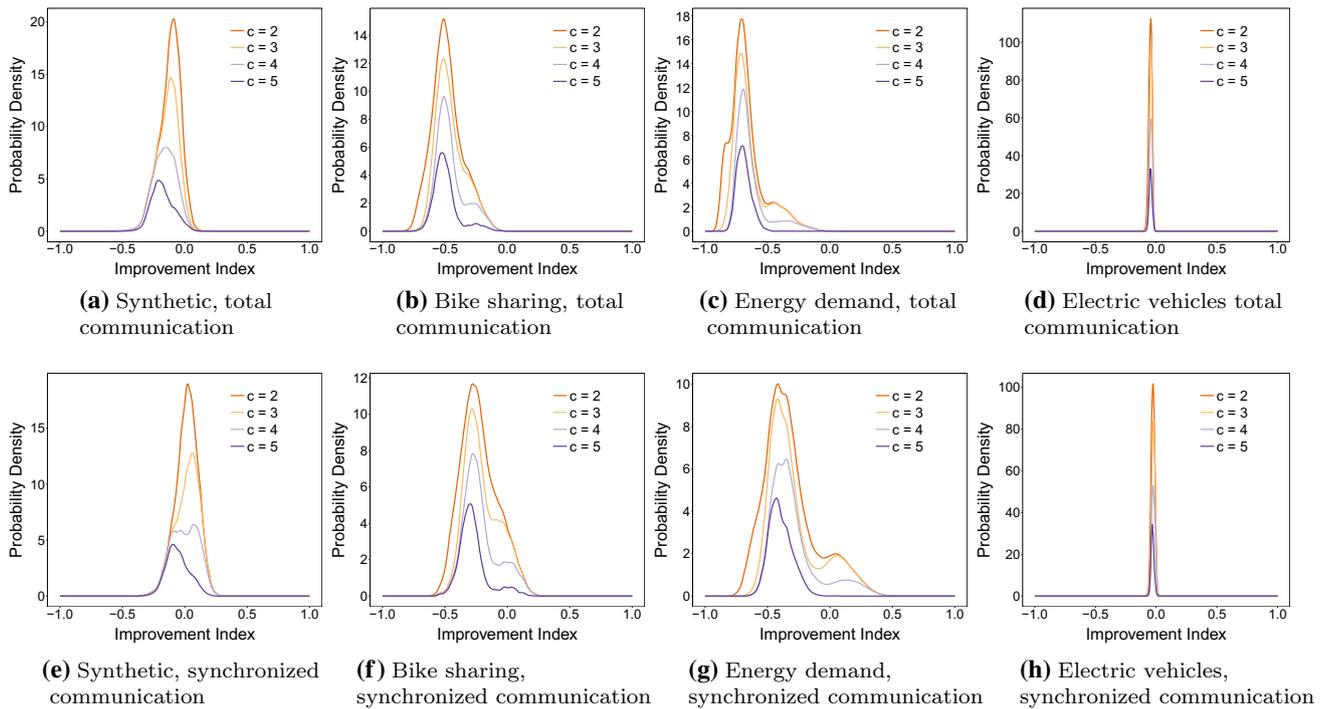


**Fig. 15** Probability density of the improvement index that elaborates on Fig. 8. *Dimensions* holarchic schemes, application scenarios, partial versus full scale. *Settings*  $\lambda = 0, c = 2$

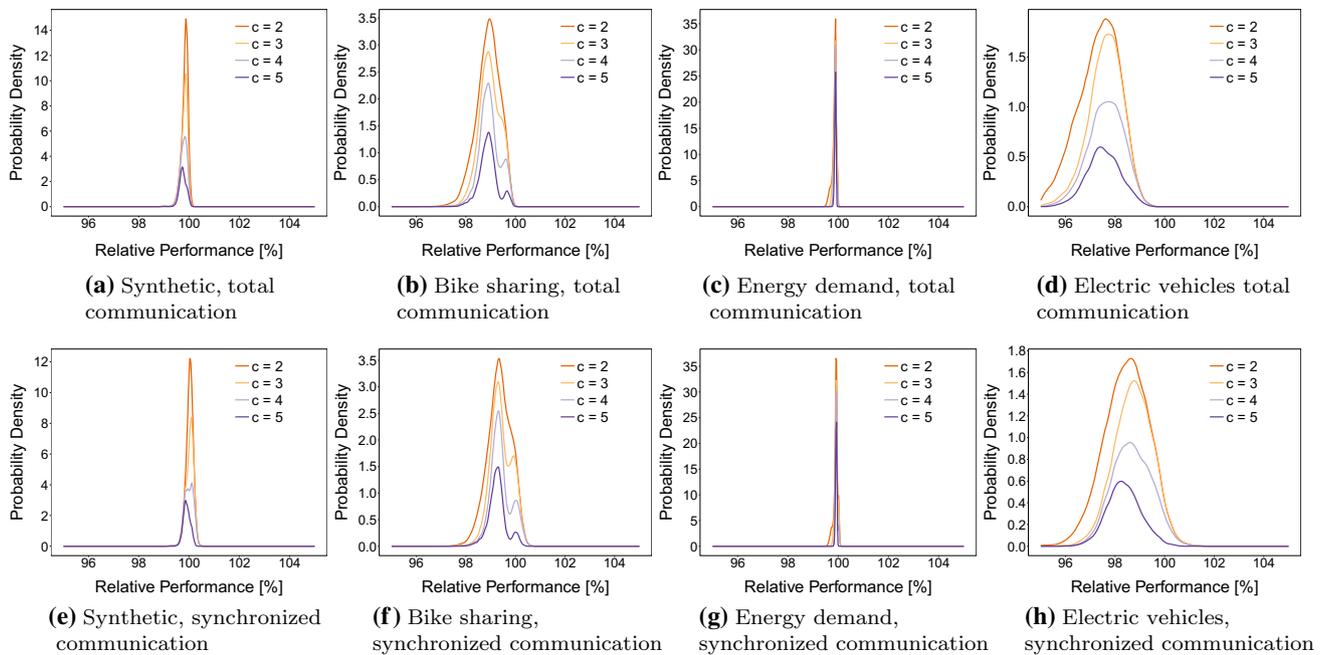
around  $-0.7$ , the relative performance values in Fig. 17f spread very close to 100%, which means that a performance peak is achieved.

Figure 18 compares with Fig. 11 by varying  $c = 2$  to  $c = 5$ . In this case, the global cost of the baseline and

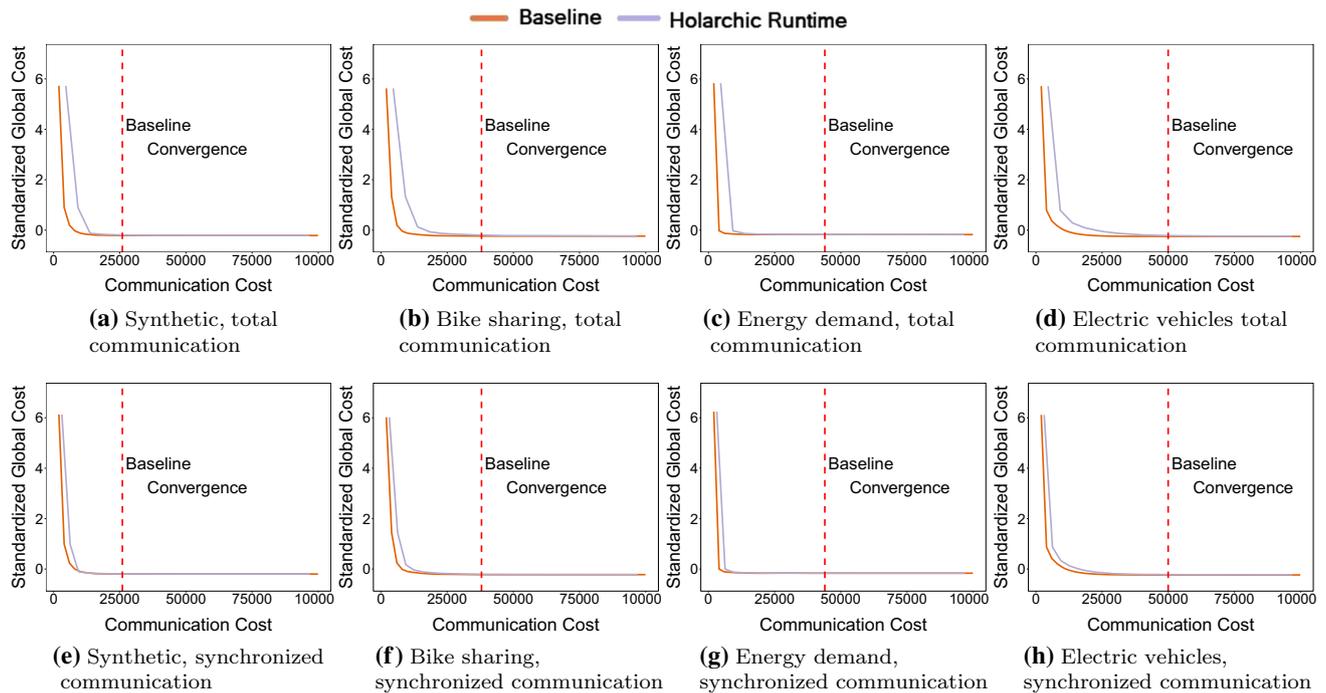
holarchic runtime is equivalent for the required communication cost to converge. The global cost of the holarchic runtime requires fewer messages to drop for  $c = 5$  compared to  $c = 2$  as indicated by the respective line shifted to the left.



**Fig. 16** Probability density of the improvement index. *Dimensions* total versus synchronized communication cost, application scenarios, varying number of children. *Settings* holarchic runtime, partial scale,  $\lambda = 0$



**Fig. 17** Probability density of the relative global cost reduction. *Dimensions* total versus synchronized communication cost, application scenarios, varying number of children. *Settings* holarchic runtime, partial scale,  $\lambda = 0$



**Fig. 18** Cost effectiveness for comparison with Fig. 11. *Dimensions* baseline versus holarchic runtime, application scenarios, total versus synchronized communication cost. *Settings* partial scale,  $\lambda = 0$ ,  $c = 5$

## References

- Aghaei, J., Alizadeh, M.I.: Demand response in smart electricity grids equipped with renewable energy sources: a review. *Renew. Sust. Energy Rev.* **18**, 64–72 (2013)
- Aziz, A.A., Sekercioglu, Y.A., Fitzpatrick, P., Ivanovich, M.: A survey on distributed topology control techniques for extending the lifetime of battery powered wireless sensor networks. *IEEE Commun. Surv. Tutor.* **15**(1), 121–144 (2013)
- Bang-Jensen, J., Gutin, G., Yeo, A.: When the greedy algorithm fails. *Discret. Optim.* **1**(2), 121–127 (2004)
- Bergstra, J., Bengio, Y.: Algorithms for hyper-parameter optimization. In: NIPS, pp. 2546–2554 (2011)
- Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **13**, 281–305 (2012)
- Christensen, J.H.: HMS/FB Architecture and Its Implementation, pp. 53–87. Springer, Berlin, Heidelberg (2003)
- Diaconescu, A.: Goal-oriented holonic systems. In: C. Müller-Schloer, S. Tomforde (eds.) *Organic Computing—Technical Systems for Survival in the Real World*, pp. 209–258. Springer, New York (2017)
- Diaconescu, A.: Organising complexity: hierarchies and holarchies. In: Müller-Schloer, C., Tomforde, S. (eds.) *Organic Computing—Technical Systems for Survival in the Real World*, pp. 89–105. Springer, New York (2017)
- Diaconescu, S.T.A., Müller-Schloer, C.: Holonic cellular automata: modelling multi-level self-organisation of structure and behaviour. In: *Proceedings of the 2018 conference on artificial life* (2018)
- Fischer, K.: Holonic multiagent systems theory and applications. In: Barahona, P., Alferes, J. (eds.) *Progress in Artificial Intelligence, LNCS*, vol. 1695, pp. 34–48. Springer, Berlin (1999)
- Frey, S., Diaconescu, A., Menga, D., Demeure, I.: A generic holonic control architecture for heterogeneous multiscale and multiobjective smart microgrids. *ACM Trans. Auton. Adapt. Syst.* **10**(2), 9:1–9:21 (2015)
- Friedman, A., Knijnenburg, B.P., Vanhecke, K., Martens, L., Berkovsky, S.: Privacy aspects of recommender systems. In: *Recommender Systems Handbook*, pp. 649–688. Springer, New York (2015)
- Galuba, W., Aberer, K., Despotovic, Z., Kellerer, W.: Protopeer: a p2p toolkit bridging the gap between simulation and live deployment. In: *Proceedings of the 2nd International Conference on Simulation Tools and Techniques*, p. 60. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering) (2009)
- Gaud, N.: A.: Towards a multilevel simulation approach based on holonic multiagent systems. In: UKSIM 08: *Proceedings of the Tenth International Conference on Computer Modeling and Simulation*, pp. 180–185 (2008)
- Giret, A., Botti, V.: Engineering holonic manufacturing systems. *Comput. Ind.* **60**(6), 428–440 (2009)
- Grande, R.C., Chowdhary, G., How, J.P.: Nonparametric adaptive control using Gaussian processes with online hyperparameter estimation. In: *IEEE 52nd Annual Conference on Decision and Control (CDC)*, pp. 861–867 (2013)
- Heidelberger, N., Karpinnen, K., D’Acunto, L.: Exposure diversity as a design principle for recommender systems. *Inf. Commun. Soc.* **2**(21), 191–207 (2018)
- Helbing, D., Frey, B.S., Gigerenzer, G., Hafen, E., Hagner, M., Hofstetter, Y., van den Hoven, J., Zicari, R.V., Zwitter, A.: Will democracy survive big data and artificial intelligence. *Scientific American* **25** (2017)
- Hinrichs, C., Lehnhoff, S., Sonnenschein, M.: A decentralized heuristic for multiple-choice combinatorial optimization problems. In: *Operations Research Proceedings 2012*, pp. 297–302. Springer, New York (2014)
- Huang, G.B., Lee, H., Learned-Miller, E.: Learning hierarchical representations for face verification with convolutional deep

- belief networks. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 2518–2525 (2012)
21. Hyndman, R.J., Koehler, A.B.: Another look at measures of forecast accuracy. *Int. J. Forecast.* **22**(4), 679–688 (2006)
  22. Koestler, A.: *The Ghost in the Machine*, 1st edn. Gateway Editions, Henry Regnery Co (1967)
  23. Kowsari, K., Brown, D., Heidarysafa, M., Jafari Meimandi, K., Gerber, M., Barnes, L.: Hdltext: hierarchical deep learning for text classification. In: IEEE International Conference on Machine Learning and Applications (2017)
  24. Landauer, C., Bellman, K.L.: New architectures for constructed complex systems. *Appl. Math. Comput.* **120**(1–3), 149–163 (2001)
  25. Lässig, J., Satzger, B., Kramer, O.: Hierarchically structured energy markets as novel smart grid control approach. In: Bach, J., Edelkamp, S. (eds.) *KI 2011: Advances in Artificial Intelligence*, pp. 179–190. Springer, Berlin (2011)
  26. Leitão, P., Restivo, F.: A holonic approach to dynamic manufacturing scheduling. *Robot. Comput. Integr. Manuf.* **24**(5), 625–634 (2008)
  27. Maclaurin, D., Duvenaud, D., Adams, R.P.: Gradient-based hyperparameter optimization through reversible learning. In: *Proceedings of the 32nd International Conference on International Conference on Machine Learning (ICML'15)*, vol. 37, pp. 2113–2122 (2015)
  28. Mai, L., Hong, C., Costa, P.: Optimizing network performance in distributed machine learning. In: *Proceedings of the 7th USENIX conference on hot topics in cloud computing, HotCloud'15*. USENIX Association, Berkeley (2015)
  29. Meignan, D.: A heuristic approach to schedule reoptimization in the context of interactive optimization. In: *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM (2014)
  30. Miiikkulainen, R., Liang, J.Z., Meyerson, E., Rawal, A., Fink, D., Francon, O., Raju, B., Shahrzad, H., Navruzyan, A., Duffy, N., Hodjat, B.: Evolving deep neural networks. *CoRR arXiv:1703.00548* (2017)
  31. Mik, E.: The erosion of autonomy in online consumer transactions. *Law Innov. Technol.* **8**(1), 1–38 (2016)
  32. Moujahed, S., Gaud, N., Meignan, D.: A self-organizing and holonic model for optimization in multi-level location problems. In *5th IEEE International Conference on Industrial Informatics*, vol. 2, pp. 1053–1058 (2007)
  33. Nolle, L., Wong, K., Hopgood, A.: Darbs: a distributed blackboard system. In: Bramer, M., Coenen, F., Preece, A. (eds.) *Research and Development in Intelligent Systems XVIII*, pp. 161–170. Springer, London (2002)
  34. Peteiro-Barral, D., Guijarro-Berdiñas, B.: A survey of methods for distributed machine learning. *Prog. Artif. Intell.* **2**(1), 1–11 (2013)
  35. Pilgerstorfer, P., Pournaras, E.: Self-adaptive learning in decentralized combinatorial optimization: a design paradigm for sharing economies. In: *Proceedings of the 12th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, pp. 54–64. IEEE Press (2017)
  36. Pournaras, E.: Multi-level reconfigurable self-organization in overlay services. Ph.D. thesis, TU Delft, Delft University of Technology (2013)
  37. Pournaras, E., Vasirani, M., Kooij, R.E., Aberer, K.: Measuring and controlling unfairness in decentralized planning of energy demand. In: *IEEE Energy Conference (ENERGYCON)*, pp. 1255–1262. IEEE International (2014)
  38. Pournaras, E., Jung, S., Yadhunathan, S., Zhang, H., Fang, X.: Socio-technical smart grid optimization via decentralized charge control of electric vehicles. *arXiv:1701.06811* (2017)
  39. Pournaras, E., Yao, M., Helbing, D.: Self-regulating supply-demand systems. *Future Gener. Comput. Syst.* **76**, 73–91 (2017)
  40. Pournaras, E., Pilgerstorfer, P., Asikis, T.: Decentralized collective learning for self-managed sharing economies. *ACM Trans. Auton. Adapt. Syst. (TAAS)* **13**(2), 10 (2018)
  41. Puchinger, J., Raidl, G.R.: Combining metaheuristics and exact algorithms in combinatorial optimization: a survey and classification. In: *International Work-Conference on the Interplay Between Natural and Artificial Computation*, pp. 41–53. Springer, New York (2005)
  42. Reymann, C., Renzaglia, A., Lamraoui, F., Bronz, M., Lacroix, S.: Adaptive sampling of cumulus clouds with uavs. *Autonomous robots* **42**(2), 491–512 (2018)
  43. Rockafellar, R.T., Uryasev, S.: Optimization of conditional value-at-risk. *J. Risk* **2**, 21–42 (2000)
  44. Rodriguez, S., Hilaire, V., Gaud, N., Galland, S., Koukam, A.: *Holonic Multi-Agent Systems*, 1st edn., chap. 11, pp. 238–263. *Self-Organising Software From Natural to Artificial Adaptation—Natural Computing*. Springer, New York (2011)
  45. Rodriguez, S., Gaud, N., Galland, S.: SARL: a general-purpose agent-oriented programming language. In: *The 2014 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. IEEE Computer Society Press, Warsaw (2014)
  46. Salakhutdinov, R., Tenenbaum, J.B., Torralba, A.: Learning with hierarchical-deep models. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1958–1971 (2013)
  47. Schiendorfer, A., Steghöfer, J.P., Reif, W.: Synthesis and abstraction of constraint models for hierarchical resource allocation problems. In: *Proceedings of the 6th International Conference on Agents and Artificial Intelligence (ICAART)*, vol. 2 (2014)
  48. Simon, H.A.: *The Architecture of Complexity*, vol. 106. American Philosophical Society, Philadelphia (1962)
  49. Simon, H.A.: *The Sciences of the Artificial*. MIT Press, Cambridge (1996)
  50. Singla, A., Santoni, M., Bartók, G., Mukerji, P., Meenen, M., Krause, A.: Incentivizing users for balancing bike sharing systems. In: *AAAI*, pp. 723–729 (2015)
  51. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**(1), 1929–1958 (2014)
  52. Sterbenz, J.P., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and survivability in communication networks: strategies, principles, and survey of disciplines. *Comput. Netw.* **54**(8), 1245–1265 (2010)
  53. Yao, C., Cai, D., Bu, J., Chen, G.: Pre-training the deep generative models with adaptive hyperparameter optimization. *Neurocomputing* **247**, 144–155 (2017)
  54. Young, S.R., Rose, D.C., Johnston, T., Heller, W.T., Karnowski, T.P., Potok, T.E., Patton, R.M., Perdue, G., Miller, J.: Evolving deep networks using HPC. In: *Proceedings of the Machine Learning on HPC Environments, MLHPC'17*, pp. 7:1–7:7. ACM, New York (2017)



**Evangelos Pournaras** is a senior scientist in the Professorship of Computational Social Science, at ETH Zurich, Zurich, Switzerland. He was earlier at Delft University of Technology and VU University Amsterdam in the Netherlands, where he completed his Ph.D. studies in 2013 with the thesis “Multi-level Reconfigurable Self-organization in Overlay Services”. Since 2007, he holds a M.Sc. with distinction in Internet Computing from University of

Surrey, UK and since 2006 a B.Sc. on Technology Education and Digital Systems from University of Piraeus, Greece. Evangelos has also been a visiting researcher at EPFL in Switzerland and has industry experience at IBM T.J. Watson Research Center in the USA. He serves the editorial board and the program committees of several international conferences and journals. He has several publications in high-impact journals and conferences in the area of distributed systems, including a best journal paper award. Smart Grids and social sensing/mining are some application domains of his expertise. Evangelos is currently working on projects focusing on fully decentralized and privacy-preserving services for social sensing and mining.



**Srivatsan Yadhunathan** is a student in the M.Sc. Statistics programme of ETH Zurich. His interests include Data Analytics, Machine Learning and Decentralized Optimization. He previously worked as a technical data analyst for Goldman Sachs for 2.5 years.



**Ada Diaconescu** is an assistant professor at Telecom ParisTech and member of the CNRS LTCI research laboratory, in France. Her research interests include autonomic and organic computing, software engineering for self-adaptive and self-organising systems, component and service-oriented architectures and interdisciplinary solutions for managing complex computing systems. She received a Ph.D. in computer science and electronic engineering from

Dublin City University (2006). Before joining Telecom ParisTech (in 2009), she has been an active researcher at University of Grenoble, Orange Labs, INRIA Rhône Alpes and Dublin City University. Her main publications concentrate on the software engineering of self-adaptive and self-organising systems, including a Springer book on *Autonomic Computing* (2013). In parallel, she has started investigating the application of her research to adjacent areas such as artificial life, unconventional computing and socio-technical systems.