Collective Learning: A 10-Year Odyssey to Human-centered Distributed Intelligence

Evangelos Pournaras

School of Computing, University of Leeds, Leeds, UK

e.pournaras@leeds.ac.uk

Abstract—This paper illustrates a 10-year research endeavor on collective learning, a paradigm for tackling tragedy of the commons problems in socio-technical systems using human-centered distributed intelligence. In contrast to mainstream centralized artificial intelligence (AI) allowing algorithmic discrimination and manipulative nudging, the decentralized approach of collective learning is by-design participatory and value-sensitive: it aligns with privacy, autonomy, fairness and democratic values. Engineering such values in a socio-technical system results in computational constraints that turn collective decision-making into complex combinatorial NP-hard problems. These are the problems that collective learning and the EPOS research project tackles. Collective learning finds striking applicability in energy, traffic, supply-chain and the self-management of sharing economies. This grand applicability and the social impact are demonstrated in this paper along with a future perspective of the collective learning paradigm.

Keywords-collective learning; artificial intelligence; human-centered AI; multi-agent system; combinatorial optimization; distributed computing; EPOS

I. INTRODUCTION

The pervasiveness of artificial intelligence (AI) in society has brought unprecedented technological opportunities, which though come along with risks on human, social and democratic values [1]. Alternative learning approaches are required that would pass new kind of Turing tests about human values such as privacy, non-discriminatory decisions that are fair and free from manipulative nudging, while remaining decentralized and participatory by-design to prevent power concentration and serve as public good. A human-centered distributed AI co-evolving along with human collective intelligence and human values is the vision behind ambitious research on the paradigm of collective learning that is the subject of this paper.

Collective learning is a human-centered and fully decentralized approach for coordinated multi-objective decision-making in socio-technical multi-agent systems. It is an approach that contributes solutions to *tragedy of the commons* problems [2] without top-down regulation using digital means. More specifically, it computationally solves a large class of combinatorial NP-hard optimization problems as a result of constraints introduced to preserve human values such as privacy, autonomy and fairness. EPOS [3], [4], [5], the *Economic Planning and Optimized Selections*, is the project under which research on collective learning has been conducted the last 10 years. Its main artifact is the I-EPOS algorithm, a highly cost-effective and general-purpose method to perform collective learning in a large spectrum of load-balancing applications scenarios of socio-technical systems: energy, traffic, bike sharing, electric vehicles, fog computing as well as a new application portfolio introduced in this paper.

Figure 1. Word cloud of 17 publication abstracts on collective learning with the following frequent keywords: decentralized, learning, smart, agents.

A plenary view of collective learning research conducted in the last 10 years (see Figure 1) is introduced in this paper to demonstrate the full potential and impact of this alternative approach. This paper also illustrates a future perspective and open research questions on collective learning to motivate further research on human-centered distributed intelligence. Moreover, this paper illustrates an MSc course on self-organizing multi-agent systems designed and taught through the prism of collective learning. Strikingly, by introducing a modeling guideline to foster creativity and innovation, students managed to incept and build a novel portfolio of application scenarios with systematic evaluations taught in the class. This confirms the feasibility of collective learning as a promising paradigm of human-centered artificial intelligence.

The contributions of this paper are summarized as follows: (i) A plenary view on collective learning research conducted the last 10 years with the key milestones. (ii) An educational approach on collective learning derived by hands-on experience in designing an MSc course. (iii) A novel application portfolio as a result of a modeling guideline [6] introduced to foster students' creativity, innovation, while assessing the feasibility of collective learning as an AI paradigm. (iv) A list of future research opportunities and open research questions.

This paper is summarized as follows: Section II discusses and compares related theories and paradigms. Section III

outlines the human-centered collective learning approach. Section IV reviews the research developments on collective learning the last 10 years and identifies the key milestones. Section V introduces the design of an MSc course via which the feasibility of collective learning as an AI paradigm is evaluated. Section VI illustrates a novel application portfolio, the modeling decisions of students, performance results and an overall evaluation of the designed MSc course. Section VII introduces a future perspective of open research challenges on collective learning and human-centered distributed intelligence. Finally, Section VIII concludes this paper.

II. Related Work

Collective learning [7] is a broadly defined, complex and highly inter-disciplinary concept [8]. It usually refers to the production of knowledge as an emergent result of dynamic and evolutionary interactive processes in which information is shared between the elements of a socio-technical system, i.e. humans and software agents.

Several manifestations of collective learning appear in artificial intelligence literature. The concept initially appears in the theory of *collective learning systems* [9]. This theory studies statistical reinforcement learning agents referred to as *collective learning automatons*. They learn appropriate responses for each environmental stimulus by choosing responses until one of them emerges as statistically optimal [10]. While this theory envisioned fully networked automatons, earlier research is limited to the learning of individual automatons [11]. The area of *autonomic computing* followed with a significant contributions on modeling further the complexity and interactions of such intelligent agents with their environment [12], [13].

More recently, *collaborative interactive learning* focuses on human and social aspects of learning processes [14]. It studies the interplay of machine-machine, human-machine and human-human collaboration, for instance, providing higher-quality and up-to-date labeling of data for training in supervised learning. In contrast, the collective learning paradigm illustrated in this paper is mainly unsupervised and focuses on artificial and human collective intelligence as a result of decentralized interactions. *Federated learning* [15] is a distributed machine learning paradigm for training models with local data across multiple edge devices. Distributed computing principles are applied to move learning models to data, instead of data to models as in mainstream centralized machine learning. This paradigm inherits desired properties such as privacy by-design and personalization. However, federated learning does not necessarily determine an alternative design for more human-centered or value-sensitive learning models.

In the last decade, a few collective learning approaches for coordination problems are studied. For instance, an incremental social learning framework is applied to decentralized decision-making problems in which agents' choices create interference. Coordination is achieved by controlling the population of agents and orchestrating the learning process as the population grows [16]. Collective learning is also studied in the context of norms and their emergence in different agents'

interaction scenarios [17], [18]. Reinforcement learning shows applicability to distributed coordination problems in robotics and traffic control [19], [20], yet most of such methods focus on supervised learning with limited scalability [21]. The theory of *probability collectives* involves agents that learn the probability distribution of selecting a particular action that maximizes both local and global utility [22]. It lies its foundations in game theory, statistical physics and optimization.

Recent methods for human-centered AI focus on the involvement of policy-makers to mitigate undesirable biases and regulate trade-offs between model accuracy and fairness [23]. Explainability models based on contrastive explanations without external assistance are also studied. They allow policy-makers to directly answer what-if questions [24]. Autonomous vehicles are also subject of active research on human-centered artificial intelligence. For instance, a value-sensitive design approach for secure, fair, legal and respectful automated vehicle speed control is recently introduced. Optimal policies are computed based on a partially observable Markov decision process and dynamic programming [25]. Despite this progress, the applicability of such human-centered AI approaches to more decentralized learning systems remains an open challenge.

III. Human-centered Collective Learning

The socio-technical computational problems of collective learning research conducted the last 10 years and outlined in this paper fall in the class of decentralized multi-agent combinatorial optimization as a result of a human-centered design. In these computational problems, humans autonomously decide about the consumption and production of some shared resources such as energy. It is known that such individuals' decisions driven by their own self-interests often result in *tragedy of the commons* [2], a situation in which independent decisions oppose common good and result in depletion of shared resources or, in general, an inefficient utilization of resources. Traffic jams, power blackouts and any undesirable load-imbalance in socio-technical systems is often a result of some tragedy of the commons problem.

To improve the outcome of collective human decisions, software agents are introduced for decision-support. In practice such agents can be a personal digital assistant in a smart phone, a home energy management system, a vehicle navigation system or any other digital twin of a human. The advantage of such agents is the automation, efficiency and facilitation of complex coordination actions required between humans to prevent tragedy of the commons. However, for humans to trust such software agents, their solutions should be superior to existing top-down regulation, market-based and economic mechanisms that often result in inequalities and concentration of power. Their design should also be human-centered and value-sensitive [26] to minimize risks such as sacrifice of autonomy over automation, privacy violation, algorithmic discrimination as well as manipulative nudging.

To make this possible, a number of critical system design constraints and choices are introduced. The multi-agent system

is designed as a fully decentralized one, no agent has full information or control over other agents. Coordination of human decisions requires flexibility, meaning a finite number of self-determined discrete options to choose from. These options are referred to as *possible plans*. Each plan is a sequence of real values that *schedule resources* over time or *allocate resources* over consumers and producers. Such plans can be generated with the assistance of software agents that can computationally generate such feasible plans for humans, e.g. reasoning based on past behavior and satisfying current human goals [27], [28]. A human may treat the plans as equivalent without preferences in the plan selection or may have preferences measured locally by a *discomfort cost*. Having multiple plans to choose is a form of operational flexibility crowdsourced by humans to agents for improving coordination efficiency and fairness measured system-wide by an *inefficiency and unfairness cost*. Inefficiency cost refers to the load-balancing of the scheduled or allocated total resources in the system. Unfairness cost is the dispersion of discomfort cost among the agents. For instance, consumers aim at having thermal comfort during cold winters (low discomfort cost), while their energy management system may require coordination of the power demand to prevent a blackout (low inefficiency cost) and make sure that consumers experience equal comfort (low unfairness cost). As a result of bringing together different human values in collective decision-making, plan selection is formalized as a socio-technical multi-objective optimization problem:

$$s = \operatorname*{arg\,min}_{j=1}^{k_i} (\gamma_i \cdot I_{i,j} + \beta_i \cdot D_{i,j} + \alpha_i \cdot U_{i,j}), \quad (1)$$

subject to $\alpha_i + \beta_i + \gamma_i = 1$ and $\alpha_i, \beta_i, \gamma_i \in [0, 1],$ (2)

$$I_{i,j} = f_\mathsf{I}(p_{i,j}, \sum_{u=1}^{n} p_{u,s}), \quad (3)$$

$$D_{i,j} = f_\mathsf{D}(p_{i,j}), \quad (4)$$

$$U_{i,j} = f_\mathsf{U}(p_{i,j}, \sum_{u=1}^{n} D_{u,s}, \sum_{u=1}^{n} (D_{u,s})^2). \quad (5)$$

The selected plan $p_{i,s}$ of agent i is among its k_i possible plans $s \in \{1, .., k_i\}$ and the one that minimizes the inefficiency $I_{i,j}$, discomfort $D_{i,j}$ and unfairness $U_{i,j}$ cost. The parameters α_i, β_i, γ_i model the individual behavior of the agent. For $\alpha_i = 0, \beta_i = 0, \gamma_i = 1$, the agent represents an *altruistic* individual, who prioritizes for system-wide efficiency, whereas, $\alpha_i = 0, \beta_i = 1, \gamma_i = 0$ represents a *selfish* individual, who prioritizes comfort. For $\alpha_i > 0$, the agent also employs fairness criteria in its plan selection.

Among the cost functions in Equation 3-5, the ones for inefficiency and unfairness are evaluated system-wide $\forall u \in \{1, ..., n\}$: using aggregate plan selections of all n agents summed up element-wise (Equation 3) as well as the sum of the discomfort cost and its sum of squares among all n agents (Equation 5) respectively. The selection of the optimum combination of plans is a combinatorial NP-hard optimization problem [5]. This is a result of a human-centered design: (i) Agent choices are made among discrete options

self-determined by humans. (ii) Human values represented in system-wide objectives are evaluated by non-linear cost functions. While linear cost functions can be optimized locally, e.g. choosing the plan with the minimum energy consumption minimizes the total energy consumption, quadratic and other more complex cost functions require coordination as each agent's choice depends on all other agents' choices [5], e.g. choosing the plan with the minimum variance does not minimize the variance of the total energy consumption (peak shaving). Such coordination is required for *balancing* (minimizing the variance of the aggregated plan selections) and *matching* (minimizing the root mean square error between the aggregated plan selections and a goal signal) objectives.

A practical, yet general-purpose, cost-effective, human-centered and decentralized approach to solve such challenging computational problems in a broad spectrum of applications motivates the research on collective learning in the last decade. The core artifact of this research is the I-EPOS learning algorithm, the *Iterative Economic Planning and Optimized Selections* [5]. I-EPOS solves the following problem in Equation 1-5: *Plan selection is augmented by making locally available to each agent the aggregated plan selections of other agents. Based on this collective information, a more informed plan selection is made that coordinates with the choices of other agents, while incrementally improving earlier plan choices.*

The collective learning algorithm works as follows: Agents are self-organized in a connected acyclic undirected graph, i.e. a tree, based on which they structure their learning interactions. Trees are known for their efficient aggregation of information and decision-making and as a result collective learning can be designed without redundant communication. Solid methods for building and maintaining tree structures in distributed environments exist [29], [30]. Each learning iteration consists of a *bottom-up* and a *top-down* phase as illustrated in Figure 2. The bottom-up phase aggregates the plan selections performed level-by-level in the tree that become coordination input to the agents of the next level above. Agents have a minimal memory of the plan selections made in the previous iteration. As a result, they can make new plan selections that lower further the cost of Equation 1. In this phase, the agents experience the information gap of not having available the plan selections of the agents above them. To guarantee monotonic improving solutions across the learning iterations and prevent falling into worse solutions when the bottom-up phase completes at the root, the agents initiate a top-down phase. During this phase each agent assesses whether itself and the agents below should switch back to the plan selections of the previous learning iteration letting the agents above contributing any potential improvement. With this back-propagation mechanism, agents achieve a continuous self-improvement until they find the optimal solution or they are trapped in a local minima. Strikingly earlier optimality results confirm the top 3% solution discovery under a monotonically decreasing inefficiency cost [5], [31]. I-EPOS also shows superior cost-effectiveness compared to related combinatorial optimization approaches [32]. A detailed illustration of the

algorithm and its complexity can be found in earlier work [5].

(a) Collective learning

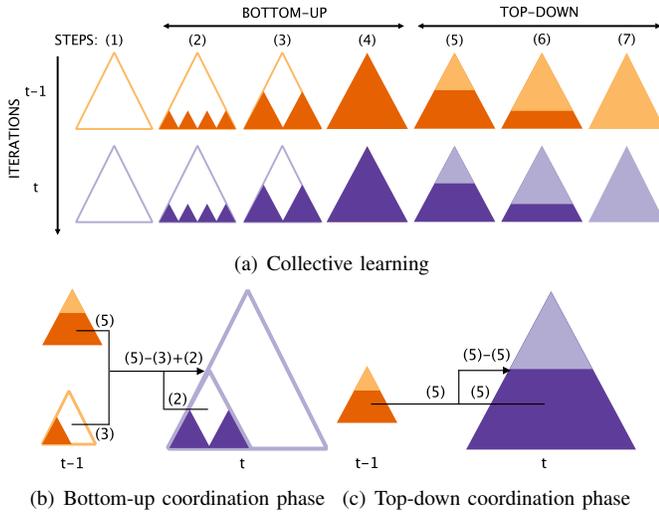(b) Bottom-up coordination phase   (c) Top-down coordination phase

Figure 2. Visual outline of collective learning and coordination in I-EPOS. (a) Each learning iteration consists of a peer-to-peer exchange of collective information between children-parent (bottom-up phase) and parent-children (top-down phase). This collective information augments the plan selection, see Equation 3 and 5. An agent coordinates based on: (b) The earlier aggregate plan selections of all agents (Step 5 at $t-1$) and the aggregate plan selection updates of the agents below (Step 2 at t and Step 3 at $t-1$). (c) The aggregate plan selection updates of all agents (Step 5 at t and $t-1$).

All learning interactions between agents rely on the exchange of aggregate plans, i.e. selected plans summed up element-wise. Agents do not share their possible plans to preserve privacy. They also do not delegate their decisions to other agents. Note also that the agent behavioral weights in Equation 1 are self-determined by the human to preserve autonomy. In other words, the design of the collective learning algorithm is also human-centered and satisfies the problem constraints that stem from the value-sensitive requirements of the multi-agent system. All these novel features together with the superior learning performance and versatile real-world applicability is what distinguishes and sets this work apart from other learning approaches in artificial intelligence.

IV. RESEARCH EVOLUTION AND MILESTONES

Figure 3 illustrates the research timeline of collective learning during the last decade. This body of work consists of 17 papers [6] with 28 coauthors from 12 different institutes (including an industry partner) as well as a PhD thesis that incepted the first concept [29]. Three related MSc theses are written and a novel MSc course is designed on the topic (see Section V). Other researchers have studied the concept of collective learning and the EPOS project beyond the author's involvement [32]. This research has also produced open-source code and artifacts [4], [33], [34], documentation and tutorials [35], [34], [6] as well as open data [36], [37].

The progress during these 10 years can be grouped into the following six milestones: (i) *Initial inception* (2010-2013). (ii) *Socio-technical human-centered approach* (2014-now). (iii) *From distributed optimization to collective learning* (2017-2018). (iv) *Generalized applicability* (2017-now). (v) *The role*

of structure in collective learning (2019-2020). (vi) *Collective learning as a robust live service* (2020).

The initial inception during the author's PhD studies involved a single objective (inefficiency cost), a bottom-up decentralized optimization approach without learning iterations and a single application domain (energy) [29].

The socio-technical human-centered approach involved the study of trade-offs between efficiency, comfort and fairness. The latter two were not yet part of the optimization process and they are simply measured. Different cost functions are incepted for the domain of energy to regulate the trade-offs at the design phase [28] rather endogenously via the optimization process. Different planning mechanisms are introduced and their role in optimization performance is established [27].

The simplicity of the initial distributed optimization concept and some early insights on observed emergent behavior (see the reverse deviations cost function [29]) motivated further experimentation and impactful design improvements that resulted in the I-EPOS algorithm [5]. The optimality, complexity and scalability of collective learning are profiled in depth [5].

Along with the introduction of collective learning, the applicability of this research expanded significantly beyond energy [5]. This expansion continues with new disruptive application scenarios: sharing economies (bike sharing), charging control of electric vehicles [38], traffic flow optimization [39], edge-to-cloud load-balancing and even an art project about the sonification of collective learning as an alternative means for general public to conceive such complex systems.

Several new insights are gained by studying the role of structure in collective learning: agents' placements in the tree structure that improve efficiency are determined as well as the most effective plan features based on which such placements can be made [31]. New large-scale optimality benchmarks are also introduced [31], [37]. Furthermore, the resilience of collective learning in distributed environments with uncertainties is studied, with findings showing that localizing collective learning in clustered branches as a result of agent failures can mitigate and even often boost the learning performance, i.e. exploration by structural adaptation [40].

Last but not least, recent software engineer developments result in turning I-EPOS to a live collective learning service of high Technology Readiness Level (TRL 6).

V. COLLECTIVE LEARNING AS AN AI PARADIGM?

Collective learning is mainly an incremental research effort and built experience of the author and his collaboration network. The question that arises is to what extent collective learning can be applied to a broader context and independently by diverse individuals as a standardized practice to solve computational learning problems in different application domains. In other words, this paper aspires to study the potential of collective learning to stand as a paradigm of artificial intelligence. The following hypothesis is formulated:

Hypothesis 1. *Collective learning is a feasible paradigm of artificial intelligence.*
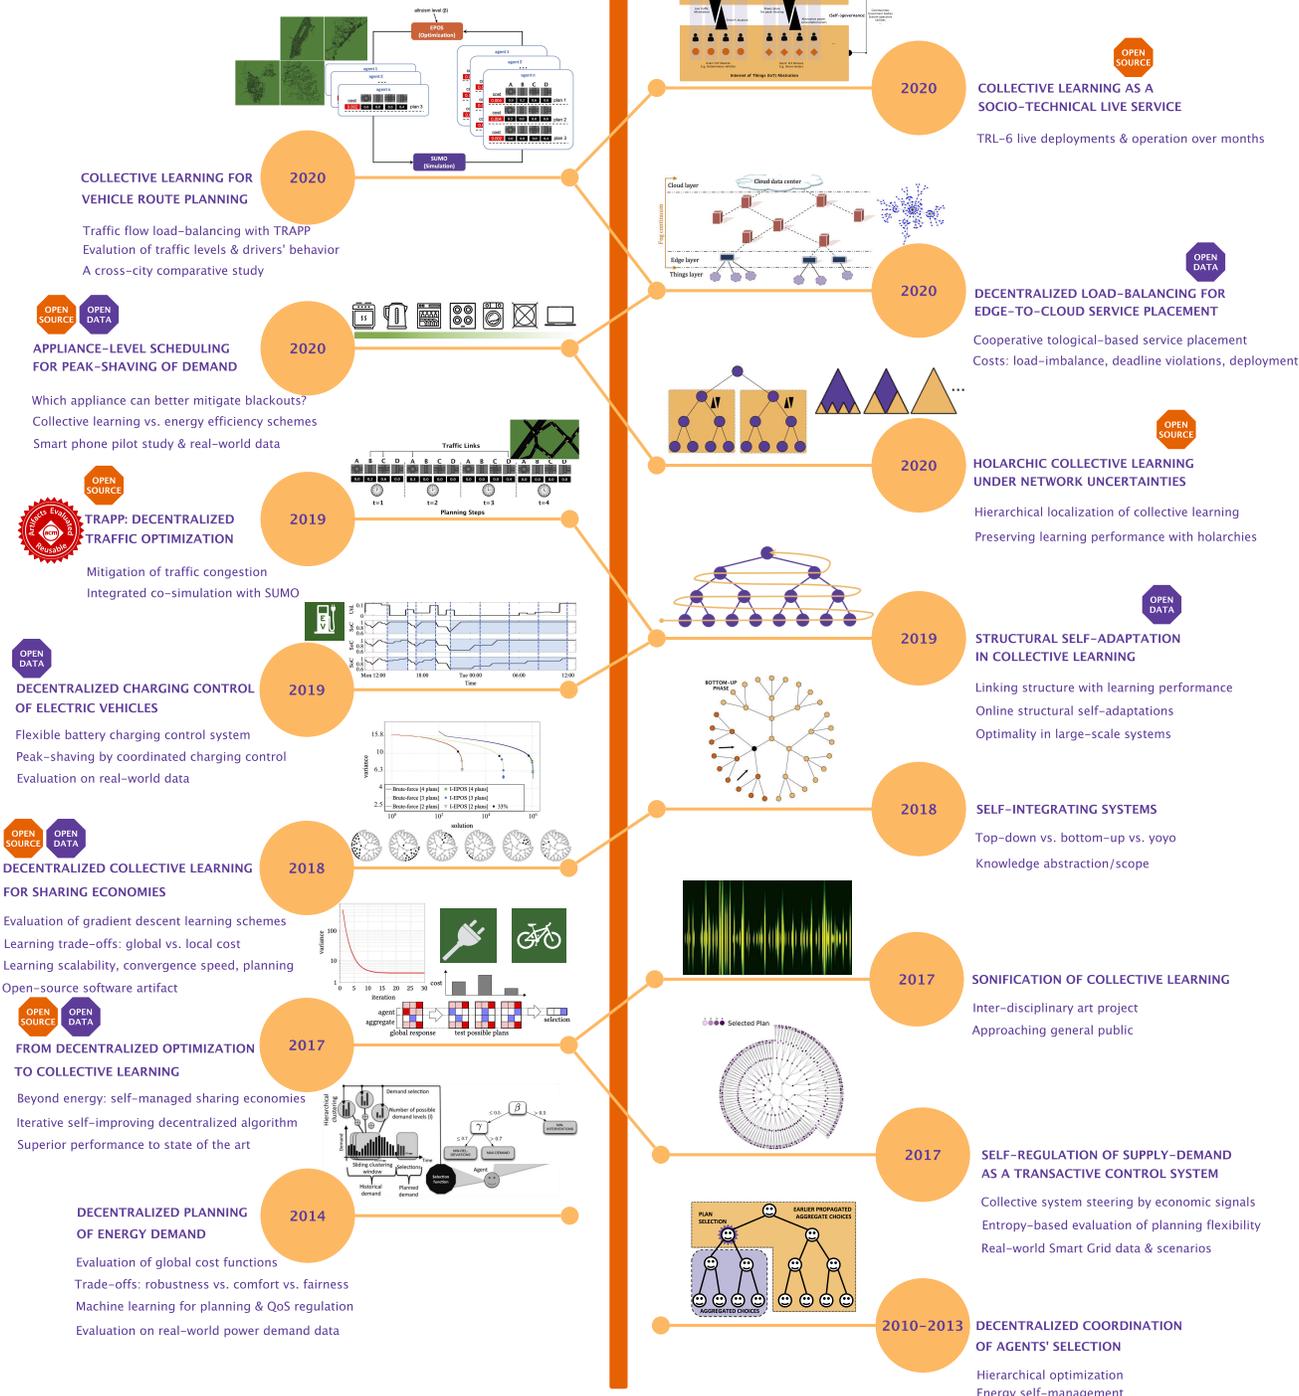
Figure 3. The research timeline of collective learning and the EPOS project in the last 10 years (2010-2020).

Applying theoretical models [41] that predict the emergence and adoption of such paradigms is not feasible without data that model several economic, political, social and business factors [42]. Instead, this paper introduces a simple, yet, pragmatic and empirical approach to study systematically Hypothesis 1: *The independent applicability of collective learning by students in a university class*. The limitation of this approach is that the necessary conditions observed in such a class and used for the acceptance or refute of the hypothesis cannot confirm a universal paradigm establishment but only its potential as illustrated below.

A new MSc course is designed and taught during the autumn semester of 2018 at ETH Zurich entitled "*Self-organizing Multi-agent Systems*" [43]. The course had a highly inter-disciplinary affinity in the university curriculum given that it ran by the Chair of Computational Social Science [44] in the department of Humanities, Social and Political Sciences. It hosted 29 students from different disciplines, including Robotics, System and Control, Informatics, Mechanical Engineer, Physics and other. Collective learning has been a central focus and subject of this course. More specifically, it was used as a educational artifact to teach and understand decentralized multi-agent systems according to the costructivism and transformative learning theory [45]. Topics such as agents' actions, goals, autonomy, altruistic vs. selfish behavior as well as the more complex concepts of coordination, combinatorial optimization and learning were introduced using problem formulations and modeling concepts of collective learning. The course though covered several other concepts beyond collective learning, for instance, autonomic computing, reinforcement learning, game theory and other. Moreover, the course relied extensively on use cases and application scenarios including traffic, energy, pedestrian dynamics, sharing economies and other. Finally, the course provided practical tutorials [35] on the EPOS software artifact [5], [4].

For the assessment for this course, students had to "*design, prototype, evaluate and present an application scenario of a distributed multi-agent system*" that performs collective learning. In other words, the students had to actually incept a new application scenario of a multi-agent system that solves a decentralized combinatorial problem according to the modeling approach of collective learning. Students worked in teams of maximum of 3 people. They were given full freedom to select and model themselves the application scenario, the dataset, the optimization problem, the cost functions, the plan generation techniques and all other design features of collective learning and its application scenario. There was neither a predefined list of application scenarios to choose from nor a list of datasets. Undoubtedly, coping with this degree of freedom at an educational-level poses a significant challenge for students, comparable to incepting and pursuing novel research such as the one conducted for the decentralized charging control of electric vehicles [38]. This approach though can provide the necessary conditions to accept or refute the Hypothesis 1: the level of the students' understanding on the material, the knowledge applicability during such challenging assignments

as well as their overall creativity and innovation on their projects can be used as indicators to assess the feasibility of collective learning as a paradigm of artificial intelligence. These conditions are observed in three incremental assignments that students had to deliver during the semester. Such assignments kept students focused and exposed with learning activities throughout the semester, while having the chance to receive feedback for improvements and track their progress.

For the first assignment (40% of the grade) delivered by the middle of the semester (28.10.2018), students are given a modeling guideline [6] with six concepts to realize in an application domain of their choice: (i) *Agents*: Determine which physical and/or virtual actor, i.e. end users or other stakeholders, are involved in the selected application scenario. (ii) *Resources*: Determine which resources are managed by the agents and the system as a whole. (iii) *Agents' plans*: Determine how agents can generate multiple options to manage their resources, meaning scheduling or allocating the consumption/production of resources. (iv) *Local agents' objectives*: Determine criteria with which agents have preferences over their generated plans (discomfort cost function). (v) *Global system-wide objectives*: Determine the system-wide objectives, which the agents need to collectively satisfy (inefficiency and unfairness cost functions). (vi) *Agents' behavior*: Determine the priority of the agent over the global vs. local objectives.

These concepts are covered in depth during the lectures along with three application scenarios as inspiring examples of how such a guideline could be realized: (i) *energy*, (ii) *bike sharing* and (iii) *electric vehicles*. These three scenarios together have an education value due to their modeling diversity: The application scenarios on energy and electric vehicles rely of schedules, while the bike sharing scenario relies on resource allocation. All scenarios have very different models to generate possible plans and their costs using historic data and algorithms, e.g. plan discomfort is modeled by shifting the power demand over time, whereas for the electric vehicles scenario by the likelihood of vehicle usage during charging.

For the second assignment (40% of the grade), three deliverables were expected by the end of the semester (5.12.2018): (i) *evaluation report*, (ii) *dataset of plans* and (iii) *plan generation code*. The evaluation report contains the following: (i) *Analysis of the generated plans* – number, values and cost of plans among the agents. (ii) The *System evaluation* – inefficiency cost reduction compared to baselines. (iii) *Socio-technical evaluation* – varying the agents' behavior. Such evaluations are introduced during the lectures and they cover the three taught application scenarios. The third assignment (20% of the grade) is the presentation and defense of the project to the lecturers and classmates (10.12.2020 and 17.12.2020).

VI. Application Portfolio: Creativity–Innovation

Students delivered a novel application portfolio of 11 projects. One project did not yield convincing results with questionable modeling choices. All other 10 projects completed successfully with an average grade of 5.4/6.0 ($\sigma = 0.46$) among the three assignments. The satisfaction level on the

lectures was on 72%, while a satisfaction level of 82% is measured on how interesting the learning material was according to official university evaluations (available upon request).

It is worth scrutinizing further the students' projects to demonstrate the significant modeling capacity of collective learning to solve real-world practical problems. Out of the 10 successful projects, 4 of them are variations of traffic and public transport optimization and as a result only one of these four is demonstrated here. In total, an application portfolio of collective learning with 6 new application scenarios is shown in Table I among with the three application scenarios taught to students during the semester for a comparison. The table outlines the modeling choices that students made as well as the prototyping and experimental choices for the evaluations.

The project on *libraries* studies how students can collectively choose the times and libraries to study to avoid overcrowded spaces and noise. This can be particularly critical in times such as pandemic lockdowns to minimize the risk of infections. The *evacuation* project studies how coordinated selections of evacuation points reduce traffic congestion and evacuation times (see Figure 4). The *parking* project studies the load-balancing of available parking spots in a city. The *food supply* project load-balances the times, the restaurants and the dishes that students select to reduce food waste. The *traffic project* introduces a load-balance of traffic flow that considers the travel time, the waiting time as well as the cost of fuel and tolls. In the *manufacturing* project, the different manufacturing elements coordinate the scheduling and allocation of the machines to minimize manufacturing duration. Finally, the *Airbnb* project introduces a price and occupancy coordination scheme for the allocation of guests to apartments to alleviate overcrowded urban spots (see Figure 5).
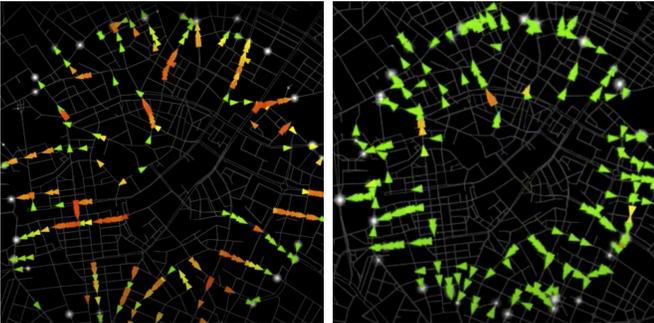
(a) Congested evacuation: 45 steps   (b) Optimized evacuation: 13 steps

Figure 4.   Snapshot [53] on the 5th step of the MATSim simulation [54].

Strikingly, all six applications scenarios are novel, highly diverse among each other and different than the ones taught in the class. There a few key observations about these projects: (i) High experimental variability: Number of agents–from 27 (manufacturing) to 4000 (traffic). Number of plans–from 16 (traffic) to ≤ 3600 (manufacturing). Plans size: from 36 (parking) to 480 (libraries). (ii) The traffic and the evacuation projects introduce a novel data-driven integration of EPOS with the SUMO [55] and MATSim [54] simulation tools. (iii) The food supply project involves a novel data collection about
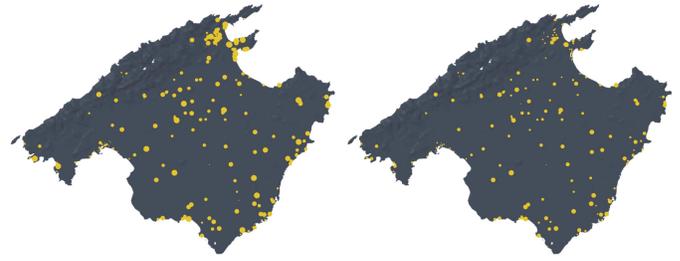
(a) Baseline overcrowded allocation   (b) Balanced allocation

Figure 5.   Allocation of guests to Airbnb apartments in Mallorca. Yellow dots size is proportional to the number of hosted guests.

the time and dish choices that students make at university restaurants. (iv) The manufacturing and Airbnb projects follow a different approach to modeling agents, i.e. the agents are modeled by the manufacturing elements and the apartments instead of machines and apartment guests respectively. An open research question is to what extent such modeling choices influence the learning performance.

The datasets and the software code of the students' projects are used here to run new experiments with the following settings: each experiment is the result of 40 learning iterations and each experiment is repeated 50 times with a random assignment of the agents over a balanced binary tree topology. For each application scenario, the designed experiments run for altruistic ($\alpha = 0, \beta = 0, \gamma = 1$) and selfish ($\alpha = 0, \beta = 1, \gamma = 0$) agents. The three performance indicators of Equation 1 are measured at each experiment: (i) *inefficiency*, (ii) *discomfort* and (iii) *unfairness*. Each one is contextualized as shown in Table I. Figure 6 illustrates the performance of collective learning at each application scenario.

The results demonstrate the following: (i) All application scenarios achieve reduction in inefficiency in exchange of a discomfort increase for altruistic agents. (ii) The application scenario of libraries and evacuation show the highest relative inefficiency reduction for altruistic agents (but also a significant increase in discomfort cost). (iii) Fairness increases by altruistic agents compared to selfish agents, however, a common lower bound in the discomfort cost among the agents results is the minimum unfairness of 0 for energy and libraries.

In conclusion, there are strong indicators to accept Hypothesis 1: Students with diverse background managed to independently incept and model effectively novel application scenarios of collective learning. Students' success stem from following the modeling and evaluation guideline as well as their continuous learning exposure during the semester. As a result they achieved very high grades although they found the course demanding, yet interesting and valuable. Several students' projects have (publishable) scientific value and could also be the basis for new business model and entrepreneurship endeavors. Such an outcome exceeds by far the initial expectations: It sets foundations for a more universal impact by repeating the course with new experience, running hackathons as well as reaching out entrepreneurs and experienced practitioners in industry. This effort settles a blueprint towards a more universal adoption and paradigm shift.

Load-balancing Scenario	Modeling Collective Learning			
	Agent Representation	Plans	Inefficiency Cost Function	Discomfort Cost Function
Energy [5]	Household (energy management system)	Energy demand over time	Residual of sum squares	Shifting discomfort
Bike Sharing [5]	Cyclist (app)	Use of bike stations	Variance	Likelihood of station usage
Electric Vehicles [38]	Driver (charge controller)	Energy charging demand over time	Root mean square error	Likelihood of vehicle usage
Libraries	Student (app)	Utilization of study places over time & libraries	Variance	Shifting discomfort & library distance
Evacuation	Citizen (app)	Use of evacuation nodes	Root mean square error	Dijkstra's distance
Parking	Driver (navigation system/app)	Use of parking spots	Variance	Shortest distance (Bee line)
Food Supply	Student (app)	Dish, restaurant & eating time	Residual of sum squares	Dish price, dish desire & waiting time
Traffic	Driver (navigation system/app)	Use of road	Variance	Travel & waiting time, fuel & toll cost
Manufacturing	Manufacturing element	Use of machines over time	Residual of sum squares	Completion time / deadline
Airbnb	Airbnb apartment	Guests, price & occupancy	Root mean square error	Price, occupancy & occupancy type

| Load-balancing Scenario	Prototyping & Experimental Choices					
	Num. of Agents	Num. of Plans	Plans Size	Plans Generation	Dataset	Goal Signal
Energy [5]	1000	10	144	Shifting	PNW [36]	✓– Ramp down generation [27]
Bike Sharing [5]	1000	23	98	Past choices	Hubway Data Visualization Challenge Paris [36]	✗
Electric Vehicles [38]	2779	4	1440	Multiple charging slots	California Household Travel Survey [36]	✓– Price signals
Libraries	500	30	480	Shifting & multiple libraries	Library Navigator [46]	✗
Evacuation	961	65	65	Roads crossing evacuation boundary	OpenBerlin for MATSim [47]	✓– Capacity of exit nodes (vehicles/hour)
Parking	166	36	36	Parking options	Parking Guidance System of Zurich [48]	✗
Food Supply	926	≤ 288	288	Arrive times, dishes & restaurants	Collected & ETH Zurich [49], [50]	✓– Num. of available dishes (scaled)
Traffic	4000	16	360	Top-k routes with min discomfort cost	Grid of 10x10 in SUMO [51]	✗
Manufacturing	27	≤ 3600	450	Possible manufacturing times & machines	Synthetic factory & machine catalogue	✓– Synthetic
AirBnb	100	55	300	Guests & price	Inside Airbnb Malloca [52],	✓– Density & type of available apartments
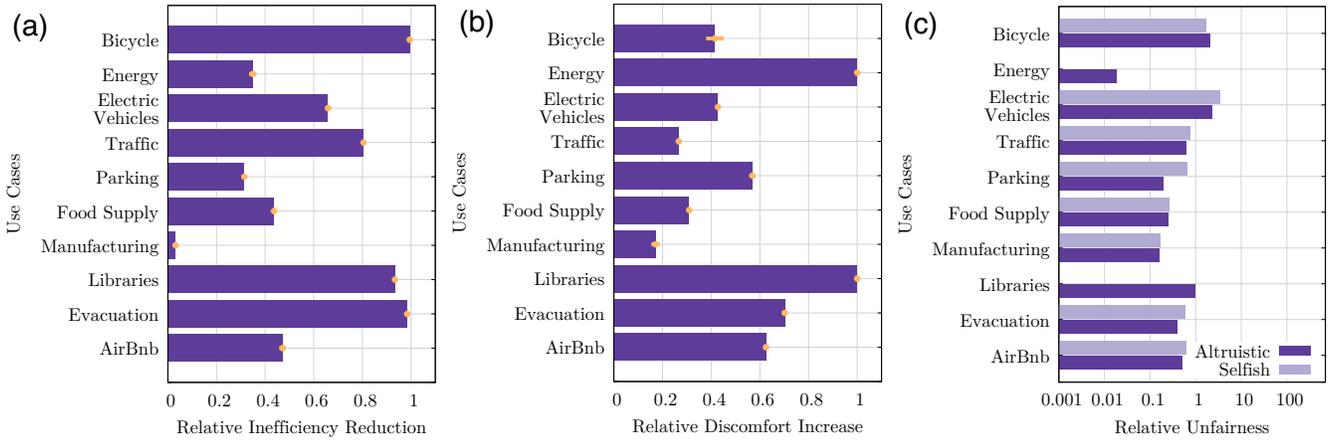
Figure 6. Performance of altruistic agents that coordinate their choices via collective learning to optimize system-wide goals of different application scenarios, as shown in Table I. The relative inefficiency cost reduction and discomfort cost increase are calculated between the altruistic and selfish agents (baseline).

VII. FUTURE PERSPECTIVE & RESEARCH

This section outlines a future perspective and open research questions on collective learning.

A. Explainability and trust

Explaining and interpreting the outcomes of collective learning is challenging [56]. In particular, explainability requires dissecting the complexity of emergent behavior that decentralized systems exhibit. It may not always be intuitive to humans how a certain selected plan serves the global and local objectives due to the combinatorial nature of the computational problem [57]. Moreover, different levels of explainability are required: (i) *System-level* explainability can support system operators, policy-makers and designers to understand, tune and steer collective decisions. (ii) *Agent-level* explainability that can empower trust of humans to collective learning. The former may depend on the latter and vice versa. Distributed ledgers designed to empower trust without central control are a promising approach in this context [58].

B. Learning resilience against plan violations & adversaries

Planning uncertainties may result in violations during plan execution, i.e. planned actions may fail. Such violations invalidate the solutions of collective learning in combinatorial optimization problems. Initiating a new learning process for each plan violation can hinder performance by adding up significant communication and processing overhead as well as latency. If such latency expands over the planning horizon, the learning process becomes impractical. Having the option to seamlessly roll forward to another comparable solution by influencing a minimum number of other agents is a self-adaptation that can make collective learning more resilient [59]. Pre-computing such backup solutions before violations happen or acquiring them efficiently on-demand after they occur are possible approaches to build up such resilience [60].

Given the sensitivity of the learning solution to every agent's choice, adversaries may heavily disdord the coordination process. For instance, assume free rider agents that persist to selfish behavior. Or agents selecting plans that maximize

the inefficiency cost or even agents that arbitrary violate the execution of their selected plans. Making collective learning tolerant to Byzantine faults may require methods to identify and isolate such agents or novel collective actions by other agents to remedy the effect of adversary behavior [61].

C. Organic collective learning in systems-of-systems

Complex systems such as power grids can perform collective learning at multiple nested encapsulated levels corresponding to consumers, aggregators, power utilities, producers, etc. The solution of a subsystem can form the goal signal of another subsystem above or below creating a holarchic and federated system-of-systems [62]. Via such an approach, agents can co-evolve their objectives in an organic way [63], [64] by self-adapting their cost functions. Moreover, preserving decentralization and the cost-effectiveness of collective learning in unstructured networks is a challenge to tackle [32].

D. Co-evolving human and artificial collective learning

A co-evolving and augmented collective intelligence can emerge by mutual decision-support of collective actions between human and software agent populations [65]. For instance, to what extent is it possible to optimize via collective learning the working hours of employees, who travel to work with electric vehicles and also coordinate their charging control via collective learning? Can new norms and culture for flexible working hours and energy sustainability emerge as a result of such coupled and augmented learning processes?

E. An opportunity for digital democracy?

Existing majority voting mechanisms often fail to achieve fair outcomes, consensus and social cohesion. They relate to low citizens' participation, poor inclusion and legitimacy, 'tyranny of the majority' effects and rise of populism [66], [67]. Multi-option preferential voting is an alternative to better promote consensus and inclusion. Such model shares common features with the operational flexibility model of collective learning. What if intelligent and transparent methods such as collective learning are used for governance, participatory and direct democracy scenarios in which citizens collectively decide for a broad spectrum of complex topics without delegating decisions to representatives? As radical as it sounds in the current status quo, scientific progress in these emerging areas makes more plausible such alternatives.

VIII. CONCLUSION

This paper concludes that collective learning is a feasible and promising paradigm of a human-centered distributed intelligence. This is demonstrated by several significant milestones reached during the last 10 years as well as the novel application portfolio incepted by students in an MSc course designed to foster modeling creativity and innovation. Nevertheless, a further involvement of industry and public sector as well as a community building promise a more universal establishment of the paradigm for tackling some of the grand future challenges outlined in this paper.

REFERENCES

[1] D. Helbing, B. S. Frey, G. Gigerenzer, E. Hafen, M. Hagner, Y. Hofstetter, J. Van Den Hoven, R. V. Zicari, and A. Zwitter, "Will democracy survive big data and artificial intelligence?" in *Towards digital enlightenment*. Springer, 2019, pp. 73–98.

[2] E. Ostrom, *Governing the commons: The evolution of institutions for collective action*. Cambridge university press, 1990.

[3] "EPOS Website," July 2020. [Online]. Available: http://epos-net.org

[4] "EPOS Open Source Code," July 2020. [Online]. Available: https://github.com/epournaras/epos

[5] E. Pournaras, P. Pilgerstorfer, and T. Asikis, "Decentralized collective learning for self-managed sharing economies," *ACM Transactions on Autonomous and Adaptive Systems*, vol. 13, no. 2, pp. 1–33, 2018.

[6] "EPOS Research Guideline," Jun 2020. [Online]. Available: https://figshare.com/articles/EPOS_Research_Guideline/12443507/1

[7] T. N. Garavan and R. Carbery, *Collective Learning*. Boston, MA: Springer US, 2012, pp. 646–649. [Online]. Available: https://doi.org/10.1007/978-1-4419-1428-6_136

[8] G. Salomon, *Distributed cognitions: Psychological and educational considerations*. Cambridge University Press, 1997.

[9] P. Bock, *The emergence of artificial cognition: An introduction to collective learning*. World Scientific Publishing Company, 1993.

[10] ——, "Observation of the propertes of a collective stochastic learning automaton," in *Proceedings of the International Information Sciences Symposium*, 1976.

[11] A. Stefanovski and P. Bock, "Team-based learning and leadership effect in multi-agent system of mixed personalities," *Procedia Computer Science*, vol. 20, pp. 177–182, 2013.

[12] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.

[13] M. C. Huebscher and J. A. McCann, "A survey of autonomic computing—degrees, models, and applications," *ACM Computing Surveys (CSUR)*, vol. 40, no. 3, pp. 1–28, 2008.

[14] A. Calma, S. Oeste-Reiß, B. Sick, and J. M. Leimeister, "Leveraging the potentials of dedicated collaborative interactive learning: Conceptual foundations to overcome uncertainty by human-machine collaboration," in *Proceedings of the 51st Hawaii International Conference on System Sciences*, 2018.

[15] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.

[16] M. A. M. De Oca, T. Stützle, M. Birattari, and M. Dorigo, "Incremental social learning applied to a decentralized decision-making mechanism: Collective learning made faster," in *4th International Conference on Self-Adaptive and Self-Organizing Systems*. IEEE, 2010, pp. 243–252.

[17] J. Hao, J. Sun, G. Chen, Z. Wang, C. Yu, and Z. Ming, "Efficient and robust emergence of norms through heuristic collective learning," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 12, no. 4, pp. 1–20, 2017.

[18] C. Yu, M. Zhang, and F. Ren, "Collective learning for the emergence of social norms in networked multiagent systems," *IEEE Transactions on Cybernetics*, vol. 44, no. 12, pp. 2342–2355, 2014.

[19] A. Yahya, A. Li, M. Kalakrishnan, Y. Chebotar, and S. Levine, "Collective robot reinforcement learning with distributed asynchronous guided policy search," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 79–86.

[20] J. Jin and X. Ma, "Hierarchical multi-agent control of traffic lights based on collective learning," *Engineering applications of artificial intelligence*, vol. 68, pp. 236–248, 2018.

[21] L. Bu, R. Babu, B. De Schutter *et al.*, "A comprehensive survey of multiagent reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 2, pp. 156–172, 2008.

[22] A. Kulkarni, A. Abraham, and K. Tai, *Probability Collectives*. Springer, 2015.

[23] Y. Zhang, R. Bellamy, and K. Varshney, "Joint optimization of ai fairness and utility: A human-centered approach," in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2020, pp. 400–406.

[24] B. Mittelstadt, C. Russell, and S. Wachter, "Explaining explanations in ai," in *Proceedings of the conference on fairness, accountability, and transparency*, 2019, pp. 279–288.

[25] S. M. Thornton, B. Limonchik, F. E. Lewis, M. J. Kochenderfer, and J. C. Gerdes, "Toward closing the loop on human values," *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 3, pp. 437–446, 2019.

[26] K. Bellman, J. Botev, H. Hildmann, P. R. Lewis, S. Marsh, J. Pitt, I. Scholtes, and S. Tomforde, "Socially-sensitive systems design: exploring social potential," *IEEE Technology and Society Magazine*, vol. 36, no. 3, pp. 72–80, 2017.

[27] E. Pournaras, M. Yao, and D. Helbing, "Self-regulating supply–demand systems," *Future Generation Computer Systems*, vol. 76, pp. 73–91, 2017.

[28] E. Pournaras, M. Vasirani, R. E. Kooij, and K. Aberer, "Decentralized planning of energy demand for the management of robustness and discomfort," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, pp. 2280–2289, 2014.

[29] E. Pournaras, "Multi-level reconfigurable self-organization in overlay services," Ph.D. dissertation, Delft University of Technology, 2013.

[30] T. Qiu, X. Liu, L. Feng, Y. Zhou, and K. Zheng, "An efficient tree-based self-organizing protocol for internet of things," *IEEE Access*, vol. 4, pp. 3535–3546, 2016.

[31] J. Nikolic and E. Pournaras, "Structural self-adaptation for decentralized pervasive intelligence," in *2019 22nd Euromicro Conference on Digital System Design (DSD)*. IEEE, 2019, pp. 562–571.

[32] C. Hinrichs and M. Sonnenschein, "A distributed combinatorial optimisation heuristic for the scheduling of energy resources represented by self-interested agents." *IJBIC*, vol. 10, no. 2, pp. 69–78, 2017.

[33] "TRAPP: Traffic Reconfigurations via Adaptive Participatory Planning," July 2020. [Online]. Available: https://github.com/iliasger/TRAPP

[34] "EPOS Software Artifact," Mar 2019. [Online]. Available: https://figshare.com/articles/EPOS_Software_Artifact/7800701/1

[35] "EPOS Tutorial and Documentation," July 2020. [Online]. Available: http://epos-net.org/software_doc/documentation/index.html

[36] "Agent-based Planning Portfolio," Apr 2019. [Online]. Available: https://figshare.com/articles/Agent-based_Planning_Portfolio/7806548/4

[37] "Optimality Benchmark for Combinatorial Optimization," Mar 2019. [Online]. Available: https://figshare.com/articles/Optimality_Benchmark_for_Combinatorial_Optimization/7803005/3

[38] E. Pournaras, S. Jung, S. Yadhunathan, H. Zhang, and X. Fang, "Socio-technical smart grid optimization via decentralized charge control of electric vehicles," *Applied Soft Computing*, vol. 82, p. 105573, 2019.

[39] I. Gerostathopoulos and E. Pournaras, "TRAPPed in Traffic? A Self-Adaptive Framework for Decentralized Traffic Optimization," in *2019 IEEE/ACM 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. IEEE, 2019, pp. 32–38.

[40] E. Pournaras, S. Yadhunathan, and A. Diaconescu, "Holarchic structures for decentralized deep learning: a performance analysis," *Cluster Computing*, vol. 23, no. 1, pp. 219–240, 2020.

[41] P. Lai, "The literature review of technology adoption models and theories for the novelty technology," *JISTEM-Journal of Information Systems and Technology Management*, vol. 14, no. 1, pp. 21–38, 2017.

[42] S. Alsheibani, Y. Cheung, and C. Messom, "Artificial Intelligence Adoption: AI-readiness at Firm-Level," *Artificial Intelligence*, vol. 6, pp. 26–2018, 2018.

[43] "Self-Organizing Multi-Agent Systems," July 2020. [Online]. Available: https://coss.ethz.ch/education/past/FS20180/multi-agent.html

[44] "Chair of Computational Social Science at ETH Zurich," July 2020. [Online]. Available: https://coss.ethz.ch

[45] M. I. Romadhon, O. Muhammad, and N. Qomariyah, "Instructional design learning: A framework of a blended discovery-project model in electronic engineering," in *Proceedings of the 2020 11th International Conference on E-Education, E-Business, E-Management, and E-Learning*, 2020, pp. 1–4.

[46] "Library Navigator," July 2020. [Online]. Available: https://www.uzh.ch/cmsssl/en/outreach/libraries/search.html

[47] D. Ziemke and K. Nagel, "Development of a fully synthetic and open scenario for agent-based transport simulations–the matsim open berlin scenario," *Transport Systems Planning and Transport Telematics–Technische Universität Berlin, Tech. Rep*, 2017.

[48] "Parking Guidance System of Zurich," July 2020. [Online]. Available: https://www.pls-zh.ch

[49] "ETH Zurich Course Catalogue," July 2020. [Online]. Available: http://www.vvz.ethz.ch/Vorlesungsverzeichnis/sucheLehrangebotPre.view?lang=en

[50] "ETH Zurich Annual Report 2018," July 2020. [Online]. Available: https://ethz.ch/en/the-eth-zurich/portrait/eth-zurich-in-figures.html

[51] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. WieBner, "Microscopic traffic simulation using sumo," in *21st International Conference on Intelligent Transportation Systems*. IEEE, 2018, pp. 2575–2582.

[52] "Inside Airbnb Malloca," July 2020. [Online]. Available: http://insideairbnb.com/mallorca/

[53] "Congested vs. optimized evacuation scenario in MATSim," July 2020. [Online]. Available: https://www.youtube.com/watch?v=c36tq5XoThI,https://www.youtube.com/watch?v=OwKodo8r9BU

[54] A. Horni, K. Nagel, and K. W. Axhausen, *The multi-agent transport simulation MATSim*. Ubiquity Press London, 2016.

[55] "SUMO – Simulation of Urban MObility," July 2020. [Online]. Available: https://www.dlr.de/ts/en/desktopdefault.aspx/tabid-9883/16931%20read-41000/16931_read-41000/

[56] S. Anjomshoae, A. Najjar, D. Calvaresi, and K. Främling, "Explainable agents and robots: Results from a systematic literature review," in *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 1078–1088.

[57] T. Miller, "Explanation in artificial intelligence: Insights from the social sciences," *Artificial Intelligence*, vol. 267, pp. 1–38, 2019.

[58] D. Calvaresi, Y. Mualla, A. Najjar, S. Galland, and M. Schumacher, "Explainable multi-agent systems through blockchain technology," in *International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems*. Springer, 2019, pp. 41–58.

[59] T. Bartold and E. Durfee, "Limiting disruption in multiagent replanning," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 2003, pp. 49–56.

[60] E. H. Durfee, C. L. Ortiz Jr, M. J. Wolverton *et al.*, "A survey of research in distributed, continual planning," *AI magazine*, vol. 20, no. 4, pp. 13–13, 1999.

[61] Y. Dong, J. Cheng, M. J. Hossain, and V. C. Leung, "Secure distributed on-device learning networks with byzantine adversaries," *IEEE Network*, vol. 33, no. 6, pp. 180–187, 2019.

[62] M. Abdoos, N. Mozayani, and A. L. Bazzan, "Holonic multi-agent system for traffic signals control," *Engineering Applications of Artificial Intelligence*, vol. 26, no. 5-6, pp. 1575–1587, 2013.

[63] C. Müller-Schloer and S. Tomforde, "Building organic computing systems," in *Organic Computing–Technical Systems for Survival in the Real World*. Springer, 2017, pp. 171–258.

[64] A. Diaconescu and J. Pitt, "Holonic institutions for multi-scale polycentric self-governance," in *International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems*. Springer, 2014, pp. 19–35.

[65] B. van Opheusden and W. J. Ma, "Tasks for aligning human and machine planning," *Current Opinion in Behavioral Sciences*, vol. 29, pp. 127–133, 2019.

[66] P. Emerson, *Majority voting as a catalyst of populism*. Springer, 2020.

[67] C. C. Wagenaar, "Beyond for or against? multi-option alternatives to a corrective referendum," *Electoral Studies*, vol. 62, p. 102091, 2019.