Artificial Collective Intelligence Engineering: A Survey of Concepts and Perspectives

Roberto Casadei
Universitá di Bologna
Alma Mater Studiorum
roby.casadei@unibo.it

Abstract Collectiveness is an important property of many systems—both natural and artificial. By exploiting a large number of individuals, it is often possible to produce effects that go far beyond the capabilities of the smartest individuals or even to produce intelligent collective behavior out of not-so-intelligent individuals. Indeed, collective intelligence, namely, the capability of a group to act collectively in a seemingly intelligent way, is increasingly often a design goal of engineered computational systems—motivated by recent technoscientific trends like the Internet of Things, swarm robotics, and crowd computing, to name only a few. For several years, the collective intelligence observed in natural and artificial systems has served as a source of inspiration for engineering ideas, models, and mechanisms. Today, artificial and computational collective intelligence are recognized research topics, spanning various techniques, kinds of target systems, and application domains. However, there is still a lot of fragmentation in the research panorama of the topic within computer science, and the verticality of most communities and contributions makes it difficult to extract the core underlying ideas and frames of reference. The challenge is to identify, place in a common structure, and ultimately connect the different areas and methods addressing intelligent collectives. To address this gap, this article considers a set of broad scoping questions providing a map of collective intelligence research, mostly by the point of view of computer scientists and engineers. Accordingly, it covers preliminary notions, fundamental concepts, and the main research perspectives, identifying opportunities and challenges for researchers on artificial and computational collective intelligence engineering.

Keywords
Collective intelligence, artificial intelligence, swarm intelligence, distributed intelligence, collective behavior, multiagent systems

1 Introduction

Nowadays, technical systems are evolving in complexity: they are increasingly large scale, heterogeneous, and dynamic, posing several challenges to engineers and operators. For instance, progress in information and communication technology (ICT) is promoting a future where computation is deeply integrated into a large variety of environments: our bodies, homes, buildings, cities, planet, and universe. In other words, the vision of pervasive and ubiquitous computing is stronger than ever, with an increasing trend toward the mass deployment of a large number of heterogeneous devices nearly everywhere, to improve existing applications and create new ones. However, we still seem quite far from exploiting the full potential of the interconnected networks of devices at our disposal.

Nevertheless, there is some progress. New paradigms and solutions have been proposed, often drawing from that powerful source of mechanisms and solutions that is nature. Indeed, we are witnessing a long-term research endeavor aiming at bringing powerful properties and capabilities of living systems into technical systems (Stein et al., 2021). Intelligence, evolution, emergence of novel capabilities, resilience, and social integration (Bellman et al., 2021; Stein et al., 2021) are often observed in natural, living systems and considered important features of artificial, engineered systems as well. Indeed, computer scientists and engineers are increasingly often interested not only in making individual devices smarter but also in making whole ecosystems of devices (and people) more collectively intelligent. Creating collective intelligence (CI) in artificial systems, however, is challenging. Indeed, various computer science and engineering fields, such as multiagent systems (MASs) (Wooldridge, 2009) and swarm robotics (Brambilla et al., 2013), have often encountered problems related to this "CI challenge." Moreover, the generality of the problem and the possibility of transferring ideas and techniques across fields have motivated the emergence of a general research field specifically aimed at studying how to build CI in artificial systems, also known under terms such as *artificial collective intelligence* (ACI) (Tumer & Wolpert, 2004; Zheng et al., 2018) and *computational collective intelligence* (CCI) (Badica et al., 2014).

There exist some surveys on CI/ACI, but they tend to adopt specific viewpoints, limiting the overall scope of the study, such as models for social computing systems (Suran et al., 2020), interaction modality (He et al., 2019), or large-scale cooperative MASs (Tumer & Wolpert, 2004). The main goal of this article is to review the concepts, models, and perspectives needed for the *engineering of ACI*. We can say that the article mainly considers *cyber-physical collectives* as target systems, namely, groups of interconnected computing devices, possibly situated in physical environments and possibly involving "humans in the loop" (Schirner et al., 2013), which are to be thought of as "programmable platforms" for services and applications benefiting from the CI emerging from their activity. The idea is to provide a research map of CI for computer scientists and engineers, generally useful for the broad technoscientific community.

In summary, we provide the following contributions:

- We perform a *scoping review* (Petticrew & Roberts, 2008), different from existing surveys in scope and focus, covering concepts, models, and perspectives related to CI, ACI, and their (software) engineering, which can also be seen as a foundation for more systematic reviews.
- We provide a map and taxonomy of the ACI field by connecting it with related fields and providing categories to frame research works on ACI.
- We outline opportunities and challenges for further research, in terms of target domains and interesting developments in existing methods.

In other words, we provide a broad overview of the field of CI/ACI, larger in scope and more oriented toward software systems engineering with respect to He et al. (2019), Malone and Bernstein (2015), Suran et al. (2020), and Tumer and Wolpert (2004).

The article is organized as follows. First, a set of broad scoping questions is elicited, to provide a structure for the article and its discussions. After this, a survey of existing reviews relevant to CI is presented, also to motivate the perspective of this article. Then, the preliminary concepts of a "collective" and "(individual) intelligence" are briefly reviewed. On this basis, to understand what CI is, some reference definitions, examples, models, and classifications from the literature are reviewed. Then, to discuss how CI can be engineered, a number of perspectives are considered, under which some main approaches for CI engineering are pointed out. Building on such a presentation of approaches, a discussion of opportunities and challenges related to CI engineering is developed, providing directions for further research. Finally, a wrap-up is provided, with some conclusive thoughts.

2 Method

Our goal is to scope the large and fragmented area of (artificial) CI to identify its key concepts, relevant perspectives, research problems, and gaps—with an emphasis on its engineering and computational/artificial intelligence (AI) side. Accordingly, we perform a scoping review (Petticrew & Roberts, 2008). This tool can be preferred over systematic reviews whenever specific research questions are hard to identify or ineffective or when the goal is to identify the *types* of available evidence or to clarify notions and key characteristics/factors related to a concept (Munn et al., 2018). Indeed, we seek to provide a map of the field, supporting more focused and systematic reviews in the future.

We use a question-based method to drive the investigation and selection of the bibliography of this manuscript. In particular, we consider the following *scoping questions* (SQs):

SQ0) What is a collective?

SQ1) What is (individual) intelligence?

SQ2) What is CI?

SQ3) What behaviors can be termed "collectively intelligent"? Are there paradigmatic examples?

SQ4) What are the requirements for CI?

SQ5) What relationships exist between individual intelligence and CI?

SQ6) How does CI unfold/emerge?

SQ7) How can CI be measured?

SQ8) How can CI be built artificially/computationally?

SQ9) What is the state of the art of (computational) CI?

SQ10) How is the research community on CI structured?

SQ0–SQ1 cover the *preliminary concepts* underlying the notion of CI, setting the necessary background for addressing SQ2–SQ3 (which are about *what* CI is) and SQ4–SQ7 (which are about the *factors*, *characteristics*, and *mechanics* of CI). Then, SQ8 is about the problem of *engineering* CI, and SQ9–SQ10 are meta-questions concerning research in the field. Notice that these are broad, scoping questions aimed mainly at providing directions for the search and identification of the research works included in the survey.

3 Tertiary Study

To motivate the need for a survey on CI, we performed a tertiary study in which secondary studies (e.g., surveys, systematic reviews) and collections were reviewed. We organize these according to whether they consider CI in its generality (i.e., abstracting from its applications and areas) or focus on its artificial/computational form (ACI/CCI), its swarm-like form, or specific kinds of collectives or goals. Therefore this section also provides a partial answer to SQ10.

In general, we can observe a lack of comprehensive reviews and maps of the CI field. From this situation, we draw a motivation for this article: providing a map of the topic, especially aimed at computer scientists and engineers, showing different perspectives and providing some highlights from the state of the art in ACI.

3.1 Reviews on CI as a General Topic

Two main surveys to date aim at addressing CI as a general topic. He et al. (2019) analyzed CI across different fields based on a taxonomy that distinguishes between isolation, collaboration, and feedback-based CI paradigms. Suran et al. (2020) performed a systematic literature review to elicit a general model of CI and its attributes, with a focus on social computing. These two contributions considered, integrated, and somewhat subsumed previous, more limited, or less general CI models and reviews (Krause et al., 2010; Lykourentzou et al., 2009; Salminen, 2012; Yu et al., 2018), and so they will be discussed more extensively in later sections. However, to the best of our knowledge, no comprehensive mapping studies provide a broad overview of the field for computer scientists.

3.2 Multidisciplinary Collections

In Malone and Bernstein (2015), essays on CI were collected from different fields, including economics, biology, human–computer interaction, AI, organizational behavior, and sociology. In Millhouse et al. (2021, p. 1), a collection of contributions from a workshop gathering scientists in different areas was provided, with the goal of sharing "insights about how intelligence can emerge from interactions among multiple agents—whether those agents be machines, animals, or human beings."

3.3 Reviews on ACI/CCI

The article by Tumer and Wolpert (2004) surveys CCI systems across the categories of (a) AI and machine learning (ML), including MASs; (b) social science–inspired systems, such as those found in economics and game theory; (c) evolutionary game-theoretical approaches; (d) biologically inspired systems, such as swarm intelligence, Artificial Life, and population approaches; (e) physics-based systems; and (f) other research fields ranging from network theory to self-organization. This is a very rich survey but covers research published before the year 2000 and is slightly focused toward automatic and utility-based approaches.

The editorial by Jung (2017) reviewed special issue papers on the integration of CCI and big data, where it is considered how data-driven CI can help in (a) collecting data, (b) analyzing data, and (c) using data, for example, to support decision-making.

The review by Rossi et al. (2018) provided a survey and taxonomy of multiagent algorithms for collective behavior, classified into consensus, artificial potential functions, distributed feedback control, geometric algorithms, state machines and behavior composition, bio-inspired algorithms, density-based control, and optimization algorithms. What emerges is a rather sharp distinction between low-level (e.g., bio-inspired self-organization) and high-level coordination.

3.4 Reviews on Swarm Intelligence

Several reviews on swarm intelligence have been published (Brambilla et al., 2013; Chakraborty & Kar, 2017; Dorigo & Stützle, 2019; Figueiredo et al., 2019; Fister et al., 2013; Kolling et al., 2016; Mavrovouniotis et al., 2017; Navarro & Matía, 2013; Nguyen et al., 2020; Rajasekhar et al., 2017; Schranz et al., 2021; Yang & He, 2013; Zedadra et al., 2018; Zhang et al., 2015). In the swarm intelligence field, a large part of research is devoted to devising (meta-)heuristics and algorithms for solving complex optimization problems. Mavrovouniotis et al. (2017) focused on swarm algorithms for dynamic optimization, namely, in settings where the environment changes over time.

Moreover, reviews in this context have often adopted an angle based on what natural system inspired swarm intelligence mechanisms. For instance, Rajasekhar et al. (2017) provided a survey on algorithms inspired by honey bees, for example, based on mating, foraging, and swarming behaviors of honey bees; similar surveys exist for bat algorithms (Yang & He, 2013), firefly algorithms (Fister et al., 2013), ant colony optimization (Dorigo & Stützle, 2019).

Some surveys have considered swarm intelligence applied to specific problems, such as self-organizing pattern formation (Oh et al., 2017), feature selection (Nguyen et al., 2020), clustering (Figueiredo et al., 2019), green logistics (Zhang et al., 2015), and collective movement (Navarro & Matía, 2013). Other surveys have considered swarm intelligence in particular contexts or as exhibited by particular kinds of systems, such as Internet of Things (IoT) systems (Zedadra et al., 2018), cyber-physical systems (CPSs; Schranz et al., 2021), and robot swarms (Brambilla et al., 2013; Kolling et al., 2016).

3.4.1 Reviews of CI for Specific Systems and Settings

Reviews from specific viewpoints have included collections and surveys on human CI (Salminen, 2012), deep learning (Ha & Tang, 2021), enterprise information systems (Nguyen et al., 2019), and sociotechnical systems supported by 5G communications (Narayanan et al., 2022).

Salminen (2012) performed a literature review of CI in the human context, grouping contributions into (a) micro-level, emphasizing enabling factors; (b) emergence (or meso-) level, emphasizing how global patterns arise from local activity; and (c) macro-level, emphasizing the kinds of system output. A review of human CI by a crowd science perspective has been provided by Yu et al. (2018). Krause et al. (2010) reviewed and compared swarm intelligence in animals and humans.

Ha and Tang (2021) performed a survey of recent developments on the embedding of CI principles into deep learning methods. They discussed, for example, how CI can help in devising novel architectures and training algorithms and recent works on multiagent (reinforcement) learning. Studies like this one are important because they elicit and strengthen transdisciplinary relationships, which are key to complex interdisciplinary fields like CI.

Narayanan et al. (2022) provided a survey of the CI emerging in human–machine sociotechnical systems supported by 5G communications. The discussed applications included road traffic control, unmanned aerial vehicles, smart grid management, and augmented democracy. The point is that to realize their full potential, these kinds of decentralized sociotechnical systems often require proper connectivity properties and capabilities to support and foster the emergence of CI. For instance, from the analysis, the authors foresaw that the 5G communication technology can promote CI by enhancing aspects like connectivity with neighbor nodes, interaction protocols, knowledge exchange, and the exploration–exploitation trade-off via improved speed, latency, and reliability. On the other hand, significant challenges will be addressed by current and, it is hoped, by future research in terms of security, privacy, and radio resource management.

4 Preliminary Concepts

This section provides an introduction to the notions of collectives and individual intelligence, hence addressing SQ0 and SQ1 and providing preliminary concepts for introducing and discussing CI in the next section.

4.1 Collectives

Informally, a *collective* is a (possibly dynamic) group of largely *homogeneous* individuals, which are also called the *members* of the collective. Different works may use different or more specific definitions for a collective. Different fields often target different kinds of collectives, often resulting in implicit assumptions.

Devising a general and comprehensive characterization of collectives is an open research problem, addressed in the context of *mereology*, namely, the study of *parthood* relations, and *ontology*, namely, the study of "what there is." In the literature, a few formal theories attempt to deeply characterize collectives and collective phenomena (Bottazzi et al., 2006; Masolo et al., 2020; Galton & Wood, 2016; Wood & Galton, 2009).

For instance, in Wood and Galton (2009), a taxonomy of collective phenomena is provided, along the classification criteria of *membership* (concerned with the identity and cardinality of the members of a collective), *location* (of the collective as well as of its members), *coherence* (the source of "collectiveness"), *roles* (if members are distinguished by roles), and *depth* (concerning levels of collectives). In particular, two main sources for collectiveness can be devised: internal or external *causes*, and *shared purposes* or *goals*. Regarding depth, it is worth noticing that, unlike the *componenthood* relation in composites, *membership* in collectives is generally not transitive (Masolo et al., 2020). Composites can be defined as structured pluralities or groups of parts, called *components*, playing specific functions (Masolo et al., 2020). In the literature, it is generally assumed that composites are heterogeneous, whereas collectives are homogeneous (Masolo et al., 2020).

Moreover, a collective is often intended to be a "concrete particular" (i.e., not an abstraction like a mathematical set) and a "continuant" (i.e., a particular existing and possibly changing over a time span) (Wood & Galton, 2009). Defining a general, comprehensive, and precise characterization or taxonomy of collectives is not trivial (Wood & Galton, 2009). For instance, certain collectives may require a certain number of members or roles to be filled to exist (Wood, 2016) or may change identity following certain changes in their composition. Sometimes collectives may be abstracted by specific collective properties or collective knowledge (Nguyen, 2008). Collectiveness may also be considered as a degree, and hence a quantifiable property (Daniels et al., 2016) of phenomena and groups of individuals.

There exist several related group-like notions, which differ, for example, by perspective, the key relation between items, or the fundamental property of the group. Some of these group-like notions are summarized in Table 1, with a proposed classification—following Masolo et al. (2020)—though different meronomies are possible. A collective is a particular kind of plurality or group. Crowds, swarms, herds, flocks, and schools can generally be considered specific kinds of collectives. Organizations and systems might be modeled as constructs based on the structural arrangement and heterogeneity of composites but are also amenable to being characterized as collectives.

For example, for the notion of intentional stance (Dennett, 1989), it may make sense to adopt a *collective stance* in which "the human species [a group] is viewed as a single organism" (Gaines, 1994, p. 19), though the idea of *collective intentionality* is problematic and the subject of intense philosophical debate (Schweikard & Schmid, 2013). Indeed, we believe that the perspective of collectiveness can provide a complementary point of view to that of an individual for understanding and engineering various sorts of systems involving groups of individuals. However, when addressing themes involving collectives (such as CI), it is important to clarify what kinds of collectives are addressed, as this would help to clarify the assumptions and generality of a specific contribution.

4.2 Intelligence

Intelligence is a controversial and elusive concept subject to philosophical debate (Legg & Hunter, 2007), best understood as a nomological network of constructs (Reeve et al., 2011). Etymologically,

Table 1. Common group-like notions addressed in computer science and engineering.

Concept	(Typical) parent concept	(Typical) defining properties
Plurality; collection; group; set		Set-inclusion
Composite	Plurality	Componenthood; heterogeneity
Collective	Plurality	Membership; homogeneity
Crowd	Collective	Nature (humans)
Swarm	Collective	Nature (insect-like)
Robot swarm	Swarm	Nature (simple robots); structure(high numbers)
Herd; flock; school	Collective	Nature (animals)
Organization	Composite/collective	Structure; roles
System	Composite/collective	Interacting elements; boundary
Multiagent system	System	Nature (agency)

intelligence comes from Latin *intelligere*, which means "to understand." It can be defined as "the global capacity of the individual to act purposefully, to think rationally, and to deal effectively with the environment" (Wechsler, 1946, p. 7) or as the property that "measures an agent's ability to achieve goals in a wide range of environments" (Legg & Hunter, 2007). In general, there are two different interpretations, intelligence as either a collection of task-specific skills or a general learning ability (Chollet, 2019), which reflect the distinction between *crystallized* and *fluid* abilities, respectively.

Problems about intelligence include, for instance, its definition and modeling, such as devising the structure of intelligence (Reeve et al., 2011), its relation with action, its measurement and evaluation, its analysis, and its construction and development.

Concerning the theories of intelligence, there are two main traditions (Reeve et al., 2011): the *psychometric tradition*, based on the number and nature of basic cognitive abilities or *factors*, and the developmental or holistic perspective, based on acquired intellect.

The problem of the *measure* of intelligence (Chollet, 2019; Hernández-Orallo, 2017) is of course related to what representation or model of intelligence is considered and is complicated by the need to distinguish between causality and correlation, select a representative set of environments for evaluation, and so on. Carroll (1993) defines an *ability* (i.e., an intelligence factor) as a source of variance in performance for a certain class of tasks. Measuring intelligence is based on *factor analysis*, that is, it works by running specific tests (*observables*) and using factors (*unobservables*) as possible explanations for correlations among the observables, describing their variability. It is expected that the nature of the entity whose intelligence we are considering would drive and require the definition of suitable factor models.

Various taxonomies of intelligence have been proposed over time. A common distinction is between *natural* (van Gerven, 2017) and *artificial intelligence* (Russell & Norvig, 2020). Both can be considered under the unifying notion of *abstract intelligence* (Wang, 2009).

5 Understanding Collective Intelligence

On the basis of the preliminary concepts introduced in the previous section, this section focuses on *what* CI is, according to literature, discussing definitions, examples, models, and the main classifications of CI (namely, ACI and CCI) in which we are interested hence addressing SQ2 to SQ7. Understanding the goals, characteristics, and main frames of reference of CI is important before turning to the problem of CCI engineering in the next section.

5.1 Definitions and Characterizations of Collective Intelligence

CI is the intelligence that can be ascribed to a collective—where a collective is a multiplicity of entities (commonly characterized as discussed in the previous section). Indeed, by abstracting a collective as a *whole*, namely, as a *higher-order individual* in turn (consisting of other individuals, which are its *members*), it should be possible to transfer characterizations of individual intelligence to it.

Table 2 reports some definitions of CI taken from the literature. From them, it is possible to see recurrent as well as peculiar aspects of CI characterizations.

5.1.1 Reuse of (Individual) Intelligence Definitions

Some definitions do not attempt to redefine "intelligence" but merely bring existing characterizations of intelligence, commonsense acceptations, or its general meaning as a nomological network of concepts (Reeve et al., 2011) to the collective realm. This has the advantages of simplicity, generality, and *openness*, which may promote multi-, inter-, and transdisciplinarity.

5.1.2 General Versus Task-Specific

If we reuse existing notions of intelligence, it means that we may consider how different definitions in turn apply to collective entities. For instance, similarly to individual intelligence, CI may be considered as a general problem-solving ability or as a set of specific skills. Evidence for the existence

Table 2. Some definitions of CI from the literature.

Reference	Definition	Remarks
Malone and Bernstein (2015, p. 1)	"*Groups of individuals acting collectively in ways that seem intelligent*"	• "Reuse" of the notion of intelligence • Collective action
Nguyen et al. (2009, p. v)	"*The form of intelligence that emerges from the collaboration and competition of many individuals (artificial and/or natural)*"	• Emergence • Mechanisms (collaboration, competition) • Members of different nature
He et al. (2019, P. 170213)	"*Collective intelligence (CI) refers to the intelligence that emerges at the macro-level of a collection and transcends that of the individuals.*"	• Emergence (transcendence) • Levels (macro, micro)
Tumer and Wolpert (1999, p. 3)	"*[COllective INtelligence (COIN)] Any pair of a large, distributed collection of interacting computational processes among which there is little to no centralized communication or control, together with a 'world utility' function that rates the possible dynamic histories of the collection.*"	• Requirements (interaction, decentralization) • Embedded metric
Szuba (2001, p. 65)	"*We can say that the phenomenon of CI has emerged in a social structure of interacting agents or beings, over a certain period, if the weighted sum of problems they can solve together as a social structure is higher during the whole period than the sum of problems weighted in the same way that can be solved by the agents or beings when not interacting*"	• Requirements (social structure, interaction) • Embedded metric • Dynamic property • Negative and positive CI
Lykourentzou et al. (2009, p. 134)	"*Collective intelligence (CI) is an emerging research field which aims at combining human and machine intelligence, to improve community processes usually performed by large groups.*"	• Hybrid or human–machine CI

of a general CI statistical factor *c* in human groups has been provided by Woolley et al. (2010), where such a factor is shown to be more correlated with average social sensitivity and diversity, rather than with average or maximum individual intelligence of the members.

5.1.3 Collectives of Different Natures

Some definitions largely abstract from the nature of collectives (e.g., "collections" or "groups of individuals," "artificial and/or natural"), some assume a minimal set of characteristics for individuals (agency, ability to interact, etc.), and some require that the individuals be connected in some way (e.g., interaction, existence of social structures).

5.1.4 Different Sources for Collectiveness and Mechanisms for CI

Terms like *interaction*, *collaboration*, *competition*, and *social structure* might be used to further constrain the scope of CI to particular kinds of collectives or to different mechanisms thereof that are possible for supporting CI.

5.1.5 Connection to Emergence

Various definitions build on the notion of *emergence*, which relates to the production, in a system, of radically novel, coherent macro-level patterns from micro-level activity (Wolf & Holvoet, 2004).

5.1.6 Phenomenological Approach

Similarly to emergence, which is often studied phenomenologically (Minati, 2018; Rainey & Jamshidi, 2018), some CI definitions adopt a phenomological standpoint where the focus is not on what CI actually is but on the phenomena that may be associated to it.

5.1.7 Positive Versus Negative CI

It is common to consider CI as a *quantifiable* property and specifically as a *signed* quantity, that is, positive or negative. Indeed, various authors talk about *negative* collective intelligence (Laan et al., 2017; Szuba, 2001) to characterize the cases in which a collective would perform worse than one of its individual members. In such cases, the social constraints effectively hinder individual abilities with no benefit.

5.2 Examples

In the following, notable examples of CI are briefly reviewed.

5.2.1 Example 1: (Markets)

Markets are economic systems that consist of a large number of rational self-interested agents, buyers and sellers, that engage in transactions regarding assets. The prices of assets change to reflect supply and demand, as well as the larger context, and can be seen as a reification of the CI of the entire market (Lo, 2015). So, markets can be seen as a mechanism for sharing information and making decisions about how to allocate resources in a collectively intelligent way (Malone & Bernstein, 2015). Accordingly, market-based abstractions have been considered in computer science to promote globally efficient systems (Mainland et al., 2004).

5.2.2 Example 2: (Wisdom of Crowds)

Crowds—groups of people—can be of different kinds (compare physical vs. psychological crowds) and can exhibit different degrees of CI. A crowd can exhibit intelligent (Surowiecki, 2005) or unintelligent behavior (Laan et al., 2017). Surowiecki (2005) popularized the term *wisdom of crowds*, showing that groups are capable of good performance under certain circumstances, providing aggregate responses that incorporate and exploit the collective knowledge of the participants. Among the conditions required for a crowd to be wise, Surowiecki (2005) identified *diversity* (of individuals), *independence* (of individual opinions), and *decentralization* (of individual knowledge acquisition)—whose importance has been confirmed by later studies, such as that by Woolley et al. (2010).

5.2.3 Example 3: (Swarm Intelligence)

Swarm intelligence is the CI that emerges in groups of simple agents (Bonabeau et al., 1999). Swarm intelligence was first observed in natural systems, such as insect societies (e.g., ant colonies, beehives), which inspired mechanisms and strategies for improving the flexibility, robustness, and efficiency of artificial systems. With respect to the general field of CI, swarm intelligence may be considered as a subfield that deals with very large groups and individuals behaving according to simple rules. Because the criteria of cardinality and simplicity are degrees, the boundaries of the field are fuzzy.

5.2.4 Example 4: (Learning Multiagent Systems)

Another notable example of CI is given by MASs (Wooldridge, 2009). Unlike swarms, MASs usually comprise rational agents, possibly structured into organizations, and possibly exhibiting properties of strong agency (Wooldridge, 2009), which may in turn be individually intelligent. The agents as well as the MAS may be able to *learn* about the environment, themselves, or the behavior that they should follow to maximize some local or global notion of utility (Tumer & Wolpert, 2004).

5.2.5 Example 5: (Human–Machine Collective Intelligence)

A powerful example of CI is the so-called *human–machine collective intelligence* (HMCI) (Smirnov & Ponomarev, 2019) or *hybrid CI* (Moradi et al., 2019; Peeters et al., 2021), which is the one that applies to heterogeneous systems involving both machines and humans. The idea is to promote the synergy between artificial/machine intelligence and human intelligence, which are often seen as complementary forms of intelligence. An exemplar of HMCI is Wikipedia, a hypermedia system of interconnected collective knowledge that is created and revised by humans through the mediation of web technologies. Wikipedia data can also be autonomously processed by agents to build other kinds of applications leveraging its collective knowledge.

5.3 Models

Here we briefly review two main general models of CI from the literature that comprehensively summarize and integrate previous models.

5.3.1 The Isolation-, Collaboration-, and Feedback-Based CI Paradigms

He et al. (2019) proposed a taxonomy of CI that divides it into three paradigms of increasing power, based on the absence or presence of *interaction* and *feedback* mechanisms. In their view, CI can be generally regarded as an aggregation of individual behavior results, thus the following:

1. *Isolation paradigm.* The individuals are isolated and behave independently, producing results that are aggregated in some way. The aggregation result does not affect the individual behaviors. Isolation studies use statistical and mining tools.
2. *Collaboration paradigm.* There is direct or indirect *interaction* between the individuals. Indirect interaction can be modeled through a notion of *environment*. Aggregation operates on individual behavior results and the environment state. The aggregation result affects neither the individual behaviors nor the environment.
3. *Feedback paradigm.* This paradigm adds to the interaction paradigm a "downward causation" of the aggregation result on the individual behaviors and/or the environment.

5.3.2 CI Framework

Suran et al. (2020) analyze 12 studies on CI and devise a *generic* model based on 24 CI attributes split into three CI components: individuals, coordination/collaboration activities, and communication means. The generic model is based on (a) characterization of *who* is involved in a CI system, in terms of passive actors (users); active actors (CI contributors), which may be crowds or hierarchies; properties of actors in terms of diversity, independence, and critical mass; and interactions; (b) characterization of *motivation* of CI actors, intrinsic or extrinsic; (c) characterization of CI goals, that is, individual and community objectives; and (d) characterization of CI processes, in terms of types of activities (decide, contest, and voluntary) and interactions (dependent or independent).

Moreover, CI systems can be considered as complex adaptive systems and often are subject to requirements for proper functioning, for example, on state, data, aggregation, decentralization, task allocation, and robustness.

5.4 Factors and Quantification of CI

Key scientific questions, fundamental for both understanding and engineering CI, include what *factors* promote or inhibit CI and, specifically, the relationship between individual intelligence and

CI. We already mentioned the seminal work by Woolley et al. (2010) sustaining the idea of a general CI factor *c*, shown to be more correlated with the level of sociality than with the levels of intelligence of individuals. We also pointed out the example of swarm intelligence as a kind of CI emerging from a multitude of simple agents characterized by limited individual intelligence. In this example, clearly, it is the aspect of *interaction*—with other agents and/or the environment—that fosters the production of effective patterns of behavior.

Works have been carried out to investigate these relationships. For instance, in a later study, Woolley et al. (2015) focused on (a) group composition, for example, in terms of skills and diversity of the members of a group, and (b) group interaction, for example, in terms of structures and norms constraining and ruling the interaction. They found that the individual skills that contribute the most to CI are those that bring sufficient diversity and effectiveness in collaboration, whereas group-level psychological elements like satisfaction and cohesiveness are not influential. Considering different kinds of interactive cognitive systems, Chmait et al. (2016) studied the influence of the following factors: (a) concerning individuals, individual intelligence and individual reasoning/learning speed; (b) concerning cooperation, cardinality of the collective, time to interact, and communication protocol; and (c) concerning agent–environment interaction, search space complexity (through uncertainty) and algorithmic complexity of the environment. They quantify the CI of a group of agents as the mean accumulated reward in a set of test environments, hence extending the Anytime Universal Intelligence Test (Hernández-Orallo & Dowe, 2010) to collectives. What is observed is that such factors—considered independently and/or in joint configurations with other factors—do shape the CI of groups in nontrivial ways. These factors are also related to the components of CI models—a nice overview is provided in Suran et al. (2020).

5.5 Main Kinds of CI

A typical classification of CI is by the nature of the entities involved.

5.5.1 Natural CI

Natural CI is the CI exhibited by collectives found in nature, such as swarms of insects; packs, herds, or groups of animals; crowds of people; flocks of birds; and schools of fishes. In all these systems, there exist nontrivial collective phenomena and societal aspects that deserve deep investigation. Insect societies are analyzed, for example, in the seminal book by Bonabeau et al. (1999). For collective animal behavior, one of the main references is the book by Sumpter (2010, p. 1), which describes collective phenomena as those in which "repeated interactions among many individuals produce patterns on a scale larger than themselves." For CI in humans, some historical references include Le Bon (1895/2002) and Surowiecki (2005); moreover, there are contributions for specific human settings like crowds of pedestrians (Sieben et al., 2017) and problems like the relationship between language and collective action (Smith, 2010).

The study of natural CI is important because it is a powerful source of inspiration for CI mechanisms to be applied to artificial systems (Bonabeau et al., 1999).

5.5.2 ACI and CCI

ACI is the CI exhibited by human-made machines. Notice that, strictly speaking, natural and ACI constitute a false dichotomy because there is inherent subjectivity regarding where the line between the two is drawn, and these could also be considered as a gradation. ACI and CCI are mostly considered as synonyms in the literature. However, some authors refer to CCI as a particular kind of ACI, that is, "as an AI sub-field dealing with soft computing methods which enable making group decisions or processing knowledge among autonomous units acting in distributed environments" (Badica et al., 2014, p. v). *Soft computing* methods are those that help to address complex problems by overcoming approximation and uncertainty, using techniques such as fuzzy logic, expert systems, ML, genetic algorithms, and artificial neural networks (Ibrahim, 2016). In other words, the

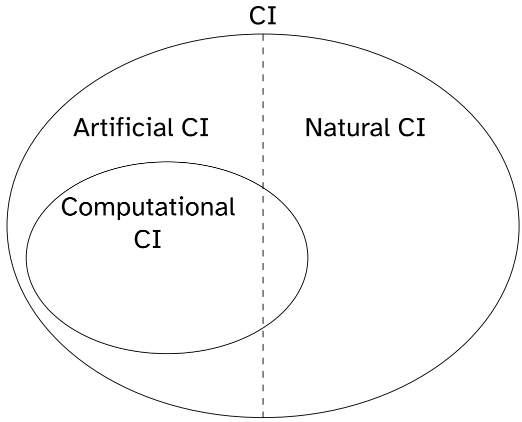

Figure 1. Relationships between collective intelligence (CI), artificial collective intelligence, and computational collective intelligence. The dashed line is used to denote the false dichotomy between artificial and natural CI.

distinction between ACI and CCI may follow the common way to distinguish between artificial and computational intelligence (Engelbrecht, 2007), where the former tends to prefer hard, symbolic approaches, while the latter tends to prefer soft and bio-inspired computing techniques. The relationships between CI, ACI, and CCI are shown in Figure 1. CCI might also be intended as a part of natural CI to account for the notion of *biological computation* (Mitchell, 2011), whereby biological systems are considered as computing devices (van Gerven, 2017). However, not all of ACI is necessarily computational, because mechanical machines also can exhibit intelligence (Stradner et al., 2013; Wang, 2009)—compare. Some Braitenberg vehicles (Dvoretskii et al., 2022) are purely mechanical vehicles with hard-wired connections between sensors and actuators. Common subfields of ACI include the semantic web, social networks, and MASs (Nguyen et al., 2009). Often, ACI and CCI include systems comprising both machines and humans. Possible taxonomies for ACI are proposed in the next section.

Notice that terms like *swarm intelligence* and *multiagent intelligence* may refer to natural systems, to artificial systems, or to a general model comprising both.

Other kinds of CI are described in the following section, as they are very much related to peculiar CI engineering methods and techniques.

6 Perspectives of Artificial Collective Intelligence Engineering

Building on the previous discussion of *what* CI is and its main models, this section focuses on *how CI can be engineered*, according to the literature, hence addressing SQ8 to SQ10. In doing this, we will picture a map of the state of the art in CI engineering, setting the stage for a discussion of research opportunities and challenges in the next section.

Depending on what kind of CI has to be achieved (see the previous section), various *perspectives* and approaches to *CI engineering* can be devised, each one leveraging and providing peculiar sets of *CI mechanisms*.

6.1 Knowledge-Oriented Versus Behavior-Oriented CI

From an industrial point of view, the engineering of CI often revolves around engineering the ICT platforms and algorithms for collecting data from human activity and extracting knowledge from collected data (Alag, 2008; Segaran, 2007). Humans using web applications can provide data through their interaction in several ways, for example, by what content they search, what paths they follow, what feedback they provide, and what content they add. Then, techniques like data mining, text mining, and ML can be used to classify information, cluster information, predict trends, recommend content, filter information, aggregate information, and so on. We may call this *knowledge-oriented* CI because the CI lies in the data produced and processed by a collective, and ICT has a role in supporting such information creation, ultimately promoting the emergence of *latent* collective knowledge. This is essentially what Surowiecki (2005) calls *cognition* problems.

This kind of system may not seem like a form of CI. Indeed, one might be tempted to completely abstract over the collective of agents providing the data and merely consider a conceptually single source of data and how data are processed and aggregated by a conceptually single process. However, the CI nature of all this starts to emerge once one considers the overall process by a larger, sociotechnical perspective. By this perspective, several agents produce information through their activity and reasoning, possibly interacting with other agents and with supporting tools (e.g., a network to share information, tools to make sense of others' contributions)—compare the isolation, collaboration, and feedback-based paradigms.

Conversely, we may call *behavior-oriented* CI the CI that drives the global behavior of a system. This includes what Surowiecki (2005) referred to as *coordination* and *cooperation* problems. Examples include the form of intelligence driving the way in which robotic swarms move (Navarro & Matía, 2013), a computational ecosystem that self-organizes into activity and communication structures (Pianini et al., 2021), and a market that self-regulates (Lo, 2015). However, as the latter example shows, because collective action is connected with collective decision-making, which in turn is connected to collective knowledge, the border between knowledge-oriented and behavior-oriented CI is fuzzy and so these types of CI should not be thought of as containers for mechanisms but rather as containers for typical CI goals.

The distinction between "plain" systems and systems of systems (SoS) (Nielsen et al., 2015), for example, based on the properties of autonomy, belonging, connectivity, diversity, and emergence (Boardman & Sauser, 2006), is also relevant to this discussion (Peeters et al., 2021). Research under the knowledge-oriented CI umbrella, typically involving sociotechnical systems, seems mostly related to the SoS framework. Research under the behavior-oriented CI umbrella, instead, seems more uniformly distributed along both the system (cf. swarm intelligence and aggregate computing; Bonabeau et al., 1999; Viroli et al., 2019) and SoS viewpoints (cf. hybrid human–machine systems; Peeters et al., 2021; Scekic et al., 2020).

6.2 Manual Versus Automatic ACI Development

Regarding ACI, it is possible to distinguish two main kinds of approaches: those based on *manual design* and those based on *automatic design*.

6.2.1 Manual Design of ACI

In the manual approach, a designer directly specifies the behavior of the computational agents making up the collective by providing *behavioral rules* (or *policies*).

Here the key issue is determining what individual behavior, when replicated or combined with other behaviors or phenomena, can give rise to the desired emergent behavior. Programming approaches that are thought to somehow address this goal are often known as *macro-programming* in the literature (Casadei, 2023), a term that recurred especially in the early 2000s in the context of wireless sensor networks (WSNs) (Newton et al., 2005). A foundational contribution to macro-programming is given by Reina et al. (2015), where a methodology is proposed for passing from macroscopic descriptions to a microscopic implementation through a design pattern, obtaining a quantitative correspondence between micro- and macro-dynamics. Research has produced different macro-programming frameworks, for example, for expressing the behavior of robot swarms (Pinciroli & Beltrame, 2016) and distributed IoT systems (Mizzi et al., 2018; Noor et al., 2019)—though most of them lack formal foundations.

A notable macro-programming approach that has recently been subject to intense research is *aggregate computing* (AC) (Viroli et al., 2019). AC consists of a functional macro-programming language expressing collective behavior in terms of computations over distributed data structures called *computational fields* (Mamei et al., 2004; Viroli et al., 2019). The basic language constructs provide support for dealing with (a) lifting standard values to field values, (b) abstracting field computations through functions, (c) stateful evolution of fields, and (d) handling bidirectional communication through so-called *neighboring fields*. Using such constructs and library functions, for example, handling information flows through gradients or supporting higher-level patterns, a programmer can write an aggregate program that expresses the global behavior of a possibly dynamic network of agents. The agents, by repeatedly evaluating the program in asynchronous sense-compute-interact rounds, and interacting with neighbors by exchanging data as dictated by the program, could steer self-organizing behavior, fulfilling, it is hoped, the intent of the program. Casadei et al. (2021) argued that multiple concurrent and dynamic aggregate computations could pave a path to CI engineering.

Whereas AC adopts a swarm-like self-organization model, another class of approaches for ACI is given by *multiagent programming*, as supported, for example, by the JaCaMo platform (Boissier et al., 2020), which comprises Jason for programming cognitive autonomous agents, CArtAgO for programming the distributed artifacts-based environment of the MAS, and MOise for programming agent organizations. However, it is worth noticing that the relationships between MAS research and CI research are often hindered by different terminologies and separate communities. A reason might be that a large part of MAS research properly focuses on composites rather than collectives, that is, well-structured organizations of heterogeneous intelligent agents rather than self-organizing swarms of largely homogeneous and cognitively simple agents.

6.2.2 Automatic Design of ACI

Because manually crafting control and behavioral rules of computational agents might be difficult, especially for complex tasks in nonstationary environments, a different approach consists in devising strategies for automatically designing behaviors. The idea is to provide hints about the intended behavior or the results to be attained by it (e.g., in terms of *specifications* or *data*) and to leverage mechanisms to generate or find behaviors that satisfy the specification. This can be addressed through *automatic programming* (O'Neill & Spector, 2020), *(machine) learning* (Behjat et al., 2021), and *search* (Russell & Norvig, 2020). For CI systems, these are essentially the approaches followed by prominent methods like multiagent reinforcement learning (MARL) (Busoniu et al., 2008) and evolutionary swarm robotics (Trianni, 2008).

One of the early models and a notable example is *collective intelligence*, or COIN (Wolpert, 2003). Essentially, COIN considers a *collective* as a system of self-interested agents trying to maximize their *private utility* function, sharing an associated *world utility* and giving a measure of the CI of the overall system. MARL is clearly a powerful technique for building CI, and it is currently a hot research area, with several surveys emerging (Canese et al., 2021; Gronauer & Diepold, 2022; Zhang et al., 2021). Learning of collective behavior may be related but should not be confused with collective learning, which is learning carried out by multiple agents that does not necessarily yield collective behavior models.

In evolutionary robotics (Trianni, 2008), the idea is to use evolutionary algorithms (i.e., algorithms that use mechanisms inspired by biological evolution for evolving populations of solutions) to optimize models of robot controllers (e.g., mapping inputs from sensors to outputs to actuators) with respect to desired behavioral goals. Various techniques have been proposed in the literature to improve traditional evolutionary approaches, for example, novelty search (Gomes et al., 2013). An interesting approach for the automatic design of the control logic of swarms is given by AutoMoDe (Francesca et al., 2014). AutoMoDe generates modular control software as a probabilistic finite state machine by selecting, composing, and configuring behavioral modules (bias). The idea is to leverage the bias to make the automatic design approach robust to differences between simulation and reality. Another relevant methodology for evolutionary robotics is the so-called *embodied evolution* approach (Bredèche et al., 2018; Watson et al., 2002), which is based on evolutionary processes that are *distributed* in a population of robots situated in an environment to support online and long-term adaptivity. Embodied evolution is an interesting setting for studying aspects like embodied intelligence, coevolution, the role of environmental niches, the relationship between optimization and selection pressure, and locality of interaction. The combination of learning and evolution is also a very interesting research direction (Gupta et al., 2021).

A second possibility for automatic design comes from *program synthesis* (Gulwani et al., 2017), which is the field studying the task of automatically crafting programs (in some given programming language) that satisfy a specified intent. Particularly interesting are the recent attempts of combining program synthesis and reinforcement learning (cf. Aguzzi et al., 2022; Bastani et al., 2020). However, in the context of CI, this direction has not yet been investigated, representing an opportunity for future research (see the next sections).

As a final remark, we stress that manual and automatic design can be seen as the extremes of a continuum and that hybrid approaches can be used (cf. interactive program synthesis Zhang et al., 2020).

6.3 Relationships Between Humans and Machines in HMCI

In HMCI, it is possible to distinguish multiple threads of research. A first classification could be based on the aforementioned distinction between knowledge-oriented and behavior-oriented CI. Other classifications can be made by considering what kind of entity plays the role of *controller* and *executor*, that is, (a) tasking crowds of humans (Ganti et al., 2011; Guo et al., 2014; Sari et al., 2019; Zhen et al., 2021) (compare crowdsourcing [Zhen et al., 2021] and crowdsensing [Guo et al., 2014]) or (b) using humans to guide machine operations, for example, interactively (Yu et al., 2021), or considering what entity plays the role of *input* and *output*, that is, (a) using AI to extract or mine intelligence from human contributions (Alag, 2008; Segaran, 2007) or (b) using humans (or *human computation*) to extract value from machine contributions, especially in tasks where machines cannot (yet) generally perform well, such as visual recognition and language understanding (Quinn & Bederson, 2011), or, finally, considering humans and machines as peers, hence the so-called *human–machine networks* (Tsvetkova et al., 2017) or *social machines* (Berners-Lee, 1999; Burégio et al., 2013).

Regarding the engineering of social machines, a notable macro-programming approach is given by the SmartSociety platform (Scekic et al., 2020), which is based on abstractions like persistent and transient teams of human–machine peers and collective-based tasks. The approach can be used for human orchestration and human tasking activities like those found in crowdsourcing and hybrid collectives.

Concerning the general design of ACI in social machines, Peeters et al. (2021) proposed three principles: (a) Goals from the collective, technological, and human perspectives should be considered simultaneously; (b) development effort should continously embrace all the product's life cycle; and (c) the requirements of observability, predictability, explainability, and directability should be considered at all abstractions levels (AI, team, and society).

6.4 Collective Tasks

Another main classification of CI engineering research is by the kind of collective task that is addressed. A *collective task* can be defined as a task that *requires* more than one individual to be carried out. Notice that CI may be seen as a requirement or mechanism for solving collective tasks (the general CI interpretation), or conversely, CI might be defined (and measured) in terms of the ability to solve a set of collective tasks in a variety of environments.

Multiple taxonomies of collective tasks have been proposed in the literature. For instance, Brambilla et al. (2013) classified collective behaviors (of swarm robotics systems) into (a) spatially organizing behaviors, (b) navigation behaviors, (c) collective decision-making, and (d) others. Other reviews of swarm robotics tasks include Bayindir (2016) and Nedjah and Junior (2019). Moreover, collective tasks can be classified according to the three paradigms discussed in He et al. (2019) and reviewed in previous sections: isolation, collaboration, and feedback.

In the following pages, we review material for two general, main kinds of collective tasks—collective decision-making and collective learning—and then point out references to other kinds of tasks.

6.4.1 Collective Decision-Making

Collective decision-making is the problem of how groups reach decisions without any centralized leadership (Bose et al., 2017; Prasetyo et al., 2019). This is also known as *group decision-making* (Tang & Liao, 2021; Zhang et al., 2019).

Decision-making and its collective counterpart can be classified according to the nature of the decision to be made. Reaching consensus and multiagent task allocation are two common kinds of collective decision-making behaviors, typical in swarm robotics (Brambilla et al., 2013). Guttmann (2009) classified MAS decision-making by four dimensions: (a) use of models of self versus models of others; (b) individual inputs versus group input; (c) learning versus nonlearning, depending on whether decision-making spans multiple rounds or just one round; and (d) collaboration versus competition. Surowiecki (2005) distinguished three kinds of problems or tasks of distributed decision-making: (a) cognition, (b) cooperation, and (c) coordination.

Collective decision-making is often supported by self-organization mechanisms based on, for example, collective perception (Schmickl et al., 2006), voter models (Valentini et al., 2014), opinion formation models (Montes de Oca et al., 2011), and self-stabilizing leader election (Pianini, Casadei, & Viroli, 2022).

Recent surveys on collective decision-making include the following. Valentini et al. (2017) focused on discrete consensus achievement and proposed a formal definition of the *best-of-n* problem (choice of the best alternative among *n* available options); then, they defined a taxonomy based on different classes of the problem and classified the literature on discrete consensus agreement accordingly. Zhang et al. (2019) provided a review of consensus models in collective decision-making and compared them based on multiple criteria for measuring consensus efficiency. They also argued that two interesting research directions include (a) *large-scale* collective decision-making and (b) addressing social relationships and opinion evolution. Tang and Liao (2021) provided a review of literature around five challenges in large-scale collective decision-making with big data: dimension reduction, weighting and aggregation of decision information, behavior management, cost management, and knowledge distribution and increase. Rizk et al. (2018) provided a survey of decision-making in MASs. The survey focuses on five cooperative decision-making models: Markov decision processes (and variants), control theory, game theory, graph theory, and swarm intelligence. These models are discussed along the dimensions of heterogeneity, scalability, and communication bandwidth—which are also crucial research challenges. Also particularly challenging is decision-making in dynamic environments (Prasetyo et al., 2019; Rizk et al., 2018). Other challenges include security, privacy, and trust; approaches to address these include, for example, blockchain consensus (Pournaras, 2020b).

6.4.2 Collective Learning

Learning is intimately related to intelligence (Jensen, 1989). Collective learning is learning backed by a collective process, with coordination and exchange of information between individuals and artifacts (Fadul, 2009). As a multi-disciplinary theme, it is studied in areas like sociology and organizational theory (Fadul, 2009; Garavan & Carbery, 2012) and in AI research (Bock, 1993). Collective learning spans both the knowledge-oriented and behavior-oriented perspectives of CI and is the main technique for automatic design of ACI. Goals of collective learning include supporting individual learning (Fenwick, 2008), producing collective knowledge, and promoting collective decision-making (Garavan & Carbery, 2012). As a wide concept, collective learning can be interpreted along multiple perspectives (Garavan & Carbery, 2012), for example, as the independent aggregation of individual learning or as a collaborative activity. So, collective learning is related to but not necessarily the same as cooperative and collaborative learning (Fadul, 2009). These different views can also be found in AI and ACI research.

Artificial collective learning includes distributed ML (Verbraeken et al., 2020); examples include centralized, federated, and decentralized ML systems. In *centralized learning*, the different individuals of the system provide data to a central entity that performs the actual learning process. So, in this case, the core learning process is not collective, though it would be collective if considered by a larger perspective that included data generation. In *federated learning* (Kairouz et al., 2021), the idea is that individual independent workers perform ML tasks on local data sets, producing models that are then aggregated by a master into a global model without the need for sharing data samples. It enables data privacy issues to be addressed. The combination of multiple models is also called *ensemble learning* (Dong et al., 2020). Hegedüs et al. (2021) proposed *gossip-based learning* as a decentralized alternative to *federated learning*, where no central entities are used and models are gossiped and merged throughout the nodes of the system. Collective learning might be supervised or unsupervised. An example of an unsupervised decentralized collective learning approach is provided by Pournaras et al. (2018).

Another important example of collective learning is MARL (Busoniu et al., 2008), which considers learning by collections of reinforcement-learning agents. MARL algorithms are commonly classified depending on whether they address *fully cooperative*, *fully competitive*, or *mixed cooperative/competitive* problems. In fully cooperative problems, the agents are given a *common reward signal* that evaluates the collective action of the MAS. Instead, in fully competitive problems, the agents have opposite goals. Mixed games are in between fully cooperative and fully competitive problems. Three common information structures in MARL are (a) *centralized structures*, involving a central controller aggregating information from the agents; (b) *decentralized structures*, with no central entities and neighborhood interaction; and (c) *fully decentralized structures*, namely, independent learning, with no information exchanged between the agents (Zhang et al., 2021). Various formal frameworks have been proposed to address MARL problems, including COIN (Wolpert, 2003) and *decentralised Markov decision processes* (Oliehoek & Amato, 2016). The reader interested in MARL algorithms and frameworks can consult multiple comprehensive surveys on the topic (e.g., Busoniu et al., 2008; Hernandez-Leal et al., 2019; Zhang et al., 2021).

There are surveys on collective learning. D'Angelo et al. (2019) performed a systematic literature review on learning-based collective self-adaptive systems. Their analysis extracted as the main characteristics of such systems, the application domains involving groups of agents with the ability to learn, the levels of autonomy of the agents, the levels of knowledge access (i.e., the way in which they explicitly share learning information), and the kinds of behaviors involved (e.g., selfish vs. collaborative). Accordingly, the authors provided a framework for learning collective self-adaptive systems based on three dimensions: autonomy, knowledge access, and behavior. The learning goals are analyzed with respect to the target emergent behavior; from the analysis, two clusters of works emerge: those for which the emergent behavior is associated with the anticipated learning task and those for which it is not. Among the learning techniques, the authors reported that the majority of research works leverage reinforcement learning, while game theory, supervised learning, probabilistic, and other approaches are less investigated in these settings. Resilience and security are deduced as the main open challenges in this research domain.

Pournaras (2020a) provided a review of 10 years of research on human-centered collective learning for coordinated multiobjective decision-making in sociotechnical systems, within the context of the Economic Planning and Optimized Selections project. Collective learning is motivated as a way to address the long-standing *tragedy of the commons* problem and is argued to be a promising paradigm of AI. As research opportunities and challenges, the author identified explainability and trust, resilience to plan violations and adversaries, collective learning in organic computing systems, coevolution of collective human learning and ML, and digital democracy.

Learning is also very related to evolution (Bredèche et al., 2018). Learning and evolution are generally considered as different mechanisms for adaptation working on different time and spatial scales (Anderson et al., 2013; Mataric, 2007). However, these techniques can also be combined

(Nolfi & Floreano, 1999): Learning can guide evolution (Hinton & Nowlan, 1987), and evolution can improve learning (cf. evolutionary learning; Telikani et al., 2022), where different architectures for the combination are possible (Sigaud, 2022).

6.4.3 Other Collective Tasks

Collective action (Oliver, 1993) commonly refers to the situation in which multiple individuals with conflicting goals as well as common goals would benefit from coordinated action. Clearly the ability to act collectively in an effective manner can be seen as an expression of CI. The problem is addressed mainly in sociology, but computer science also provides tools (e.g., simulations, models), such as the SOSIEL (Self-Organising Social and Inductive Evolutionary Learning) simulation platform (Sotnik, 2018), for studying the problem and investigating solutions for human societies as well as for sociotechnical and artificial systems. Collective actions may be supported by collective and self-organized decision-making processes and by leveraging abstractions like *electronic institutions* and *social capital* (Petruzzi et al., 2015).

Collective movement (Navarro & Matía, 2013) is the problem of making a group of agents (e.g., robots, drones, vehicles) move toward a common direction in a cohesive manner. Notice that this is not just about movement per se but rather moving in conjunction or to support other tasks as well—for example, distributed sensing, exploration, and rescue tasks. Two main subproblems can be identified (Navarro & Matía, 2013): (a) *formation control* (Yang et al., 2021), when the shape of the group and/or the individual positions are important, and (b) *flocking* (Beaver & Malikopoulos, 2021), where such aspects are less important.

Distributed optimization (Yang et al., 2019) refers to the problem of minimizing a global objective function, which is the sum of the local objective functions of the members of a collective, in a distributed manner. Distributed optimization can be a technique for collective decision-making.

Collective knowledge construction refers to the creation of new, distributed, and shared knowledge by a collective (Hecker, 2012). This topic is generally studied by considering aspects like collaboration (Hmelo-Silver, 2003), sociotechnical infrastructures (Gruber, 2008), knowledge transfer (Huang & Chin, 2018), the interplay between individual and collective knowledge (Kimmerle et al., 2010), models of information diffusion dynamics (Maleszka, 2019), and lifelong learning (Rostami et al., 2018).

6.5 A View of CI-Related Fields

CI being a multidisciplinary field, the engineering of CI and ACI can benefit from ideas and research results from a variety of fields. It would be useful to have a comprehensive map of research fields contributing to CI.

Though we consider providing a comprehensive research map of CI engineering as a future work, we provide a research map (see Figure 2) from the perspective of *collective adaptive system* (CAS) research (Bucchiarone et al., 2020; Casadei, 2020; Ferscha, 2015; Nicola et al., 2020). The idea is that CI engineering should be supported through interdisciplinary research and a systems science perspective (Mobus & Kalton, 2015), also providing a rigorous treatment of system-level properties that could be sustained by CI processes. This includes leveraging studies of abstract and fundamental kinds of systems, for instance, CPSs namely, systems that combine discrete and continuous dynamics (Alur, 2015). Then, a set of interrelated fields can promote the study of peculiar CI phenomena, such as emergence, self-organization, and ensemble formation. Such fields include but are not limited to coordination (Malone & Crowston, 1994), MASs (Wooldridge, 2009), autonomic/self-* computing (Kephart & Chess, 2003), collective adaptive systems (Bucchiarone et al., 2020; Casadei, 2020; Ferscha, 2015; Nicola et al., 2020), ubiquitous/pervasive computing (Weiser, 1991), swarm intelligence (Bonabeau et al., 1999), and collective computing (Abowd, 2016). Some of these are briefly overviewed in the following paragraphs.

We noticed multiple times in previous sections how interaction is a key element of CI. *Coordination* is the interdisciplinary study of interaction (Malone & Crowston, 1994). In computer science,

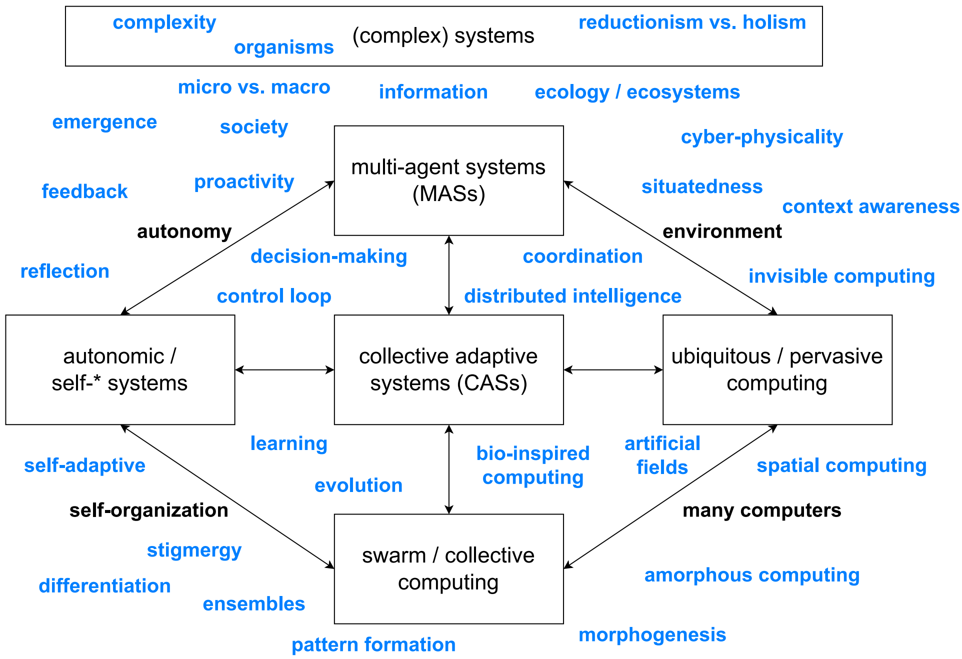

Figure 2. Research map of fields and concepts contributing to (research on) CI engineering.

interaction was early recognized as a concern related to but clearly distinguished from computation (Gelernter & Carriero, 1992), hence amenable to separate modeling by so-called coordination languages. A general meta-model of coordination (Ciancarini, 1996) consists of *coordinables* (the interacting entities), *coordination media* (the abstractions supporting and constraining interactions), and *coordination laws* (describing the behavior of a coordination medium). Languages, abstractions, and patterns can be used to define the way in which computational components coordinate across aspects like control, information, space, and time. This has motivated the birth of whole communities and long-standing research threads (Ciancarini & Hankin, 1996; ter Beek & Sirjani, 2022).

CASs are collectives of agents that can adapt to changing environments with no central controller. Their engineering poses several challenges, tackled in corresponding research communities (Bucchiarone et al., 2020; Nicola et al., 2020). CASs are sometimes considered to be heterogeneous (Andrikopoulos et al., 2014; Loreti & Hillston, 2016), contrasted to more homogeneous intelligent swarms, though we tend to disagree with this view. In our view, CASs are a superset of intelligent swarms, which are characterized by (a) large numbers of (b) relatively simple (or not particularly intelligent) individuals (Beni, 2004). Collectives are generally quite homogeneous, at least at some level of abstraction (Pianini, Pettinari, et al., 2022), though research aims to address heterogeneous CASs (Scekic et al., 2020) as well as heterogeneous swarms (Dorigo et al., 2013; Kengyel et al., 2015), for example, with systems involving humans and robots (Hasbach & Bennewitz, 2022) or groups of robots with different morphology or behavior. Swarm robotics is the combination of swarm intelligence and robotics (Beni, 2004; Brambilla et al., 2013).

Coordination, CASs, and swarm robotics can also be seen as subfields of the larger field of MASs (Faliszewski et al., 2022; Wooldridge, 2009), which itself stemmed from the field of distributed AI (Ferber, 1999). In MAS engineering, two main levels and corresponding problems are considered: the *micro-level* of agent design, and the *macro-level* of agent society design. *Autonomy* (encapsulation of control) and *agency* (the ability to act) are generally considered the two fundamental properties of agents (Franklin & Graesser, 1996; van der Hoek & Wooldridge, 2003) from which other properties, such as proactiveness, interactivity, and sociality, stem. From a software programming and engineering point of view, agents can be considered as an abstraction following active objects and actors (Odell, 2002) that, together with other first-class abstractions like artifacts (Omicini et al., 2008), environments (Weyns & Michel, 2014), and organizations (Horling & Lesser, 2004), provide support for the so-called *(multi-)agent-oriented programming* paradigm (Boissier et al., 2020; Shoham, 1993). The MAS field/perspective is clearly intimately related to CI.

Like for MASs, the key property of autonomy is at the center of *autonomic computing* (Kephart & Chess, 2003), namely, the field devoted to constructing computational systems that are able to manage/adapt themselves with limited or no human intervention. Following this vision, research has been carried out to find approaches and techniques to endow artificial systems with different *self-* properties*: self-adaptive (de Lemos et al., 2010; Salehie & Tahvildari, 2009), self-healing/repairing (Psaier & Dustdar, 2011), self-improving/optimizing (Bellman et al., 2018), self-organizing (Heylighen, 2013), and so on. To build autonomic systems, approaches typically distinguish between the *managed system* and the *managing system*, structuring the latter in terms of *monitoring, adaptation, planning, execution, and knowledge* (MAPE-K) components (Kephart & Chess, 2003). In so-called *architecture-based self-adaptation* (Garlan et al., 2004), architectural models of the managed systems are leveraged at run time to organize the self-managing logic. The managing system could also be distributed and decentralized (Weyns et al., 2010). If the managed system is a collective, then its self-* properties could be put in relation to its CI. Consider the property of being *self-organizing*, characterized by processes that autonomously and resiliently increase/maintain order or structure (Wolf & Holvoet, 2004); it typically emerges from the interaction of several entities. Self-organization can be considered as a key promoter or element of CI (Rodríguez et al., 2007).

As a last remark, we stress that the aforementioned fields are highly inter- and transdisciplinary. For instance, MASs can be considered by economical, sociological, organizational, and computational perspectives (Wooldridge, 2009). The same goes for coordination (Malone & Crowston, 1994). Moreover, a great source of inspiration is given by natural (e.g., physical and biological) systems, as recognized by a wealth of *nature-inspired coordination* (Zambonelli et al., 2015) and *nature-inspired computing* (Siddique & Adeli, 2015) contributions.

7 Research Opportunities and Challenges

With an understanding of the nature of CI and its engineering perspectives, in the following sections we discuss a few related research directions that include interesting opportunities and challenges for researchers in CI engineering.

7.1 Programming Emergence and Macro-programming

The problem of programming emergent and self-organizing behaviors is an open research challenge (Gershenson et al., 2020; Varenne et al., 2015) intimately related to CI engineering. The term *macro-programming* emerged in the early 2000s to identify programming approaches with the goal of defining the global behavior of WSNs (Newton et al., 2005); currently, it generally denotes paradigms aiming at supporting the programming of system- or macro-level properties and behaviors. A recent survey by Casadei (2023) showed that, beside the first wave of research in the context of WSNs, we are witnessing a renewed interest in macro-programming fueled by scenarios like the IoT, robot swarms, and CASs in general. This is also very much related to spatial computing (Beal et al., 2013), as space is often a constraint, a means, or a goal in systems.

The key challenge here is determining what local behavioral rules of the individuals can promote the desired collective behavior. In particular, we can distinguish two problems (Tumer & Wolpert, 2004): Given a set of individuals and the corresponding local behavioral rules, the *local-to-global mapping problem* (or *forward problem*) is the problem of determining what global outcomes will be produced; conversely, the *global-to-local mapping problem* (or *inverse problem*) is the problem of determining what local behaviors have produced the observed global outcomes. In macro-programming, the latter problem turns into how to map a description of a global intent (macro-program) into local behavioral rules (micro-programs) (Casadei, 2023).

It has been shown that approaches like aggregate computing (Viroli et al., 2019) can support forms of self-organization and CI with macro-programs that can be encoded as compositions of functions of reusable collective behaviors (Audrito et al., 2022). This is promising, but still, little research has as yet been devoted to investigating, systematizing, and formalizing the principles, concepts, and mechanisms of macro-programming in general or specific settings (Casadei, 2023).

7.2 Integration of Manual and Automatic CI Engineering Methods

In previous sections, we have discussed how CI can be programmed manually (e.g., through macro-programming languages, or using traditional techniques to connect and extract knowledge from human activity) or automatically (e.g., via multiagent reinforcement learning techniques or program synthesis). Arguably, the two approaches could be combined to overcome their individual issues. This is a still an unexplored research direction, but early works and ideas are emerging.

A first idea could be to use program synthesis (Gulwani et al., 2017) to synthesize macro-programs expressed in a macro-programming language (Casadei, 2023). This could be coupled with simulation to verify how systems executing synthesized programs operate in various environments. On one hand, because simulations may also be computationally intensive, it might be necessary to limit simulation to a few program candidates. On the other hand, the problem of generating macro-programs might be hard, especially if the space of possible programs is very large. Therefore macro-programming languages admitting few primitives or combinators may be more suitable for this.

Additionally, some recent attempts have been made at combining program synthesis and reinforcement learning (Bastani et al., 2020; Qiu & Zhu, 2022; Verma et al., 2018). For instance, Bastani et al. (2020) discussed approaches to reinforcement learning based on learning programmatic policies (i.e., policies in the form of a program), which can provide benefits in terms of interpretability, formal verification, and robustness. Therefore it would be interesting to consider the application of MARL where policies are expressed in a multiagent-oriented or macro-programming language. An early attempt has been carried out (e.g., Aguzzi et al., 2022) in which MARL has been used to fill a hole in a sketched aggregate computing program (cf. the *sketching* technique in program synthesis; Solar-Lezama, 2009), resulting in collective adaptive behavior that improves over a simple, manually encoded collective behavior.

7.3 Integration of Bottom-Up and Top-Down Processes

Another interesting research challenge and opportunity for our ability of engineering CI lies in achieving a better understanding of how bottom-up and top-down processes can be integrated—or, in other words, how emergence and downward causation/feedback can be exploited altogether to provide both flexibility and robustness in collective behavior. Indeed, we are considering the *feedback* CI paradigm (He et al., 2019), whereby the aggregation of contributions from the individuals and the environment in turn affects the individuals and the environment. This is also what Lykourentzou et al. (2009) called *active* CI systems, in which collective behavior is supported by the system level, which are contrasted to *passive* CI systems, in which no collective awareness or intentionality is present.

The problem of integrating top-down and bottom-up processes is indeed connected with the problem of *controlling emergence*, addressed in research fields like autonomic computing (Kephart & Chess, 2003), with its MAPE-K loops, and *organic computing* (Müller-Schloer & Tomforde, 2017), with *observer-controller* architectures. One issue is that emergence itself is a controversial concept, subject to philosophical and scientific investigation and often presented with definitions that hardly apply to systems engineering (Müller-Schloer & Sick, 2006). Attempts to define emergence based on hierarchical system models and ontological approaches (Gignoux et al., 2017) may prove useful. Initial, working classifications of emergence for reasoning in systems engineering may be based, for example, on whether it is *anticipated* or *not anticipated*, and whether it is *desirable* or *undesirable* (Iivari, 2016).

Some engineering techniques discussed in this section, such as macro-programming and MARL, could support the design of "controlled emergence," and in another direction, a deeper understanding of emergence and its relationship with feedback could provide insights for mechanisms or the implementation of such techniques. A macro-program, indeed, could be seen as a top-down structure for emergent processes. Also interesting in this respect are, for example, formal studies carried out on *self-stabilization* of aggregate computations (Pianini, Casadei, & Viroli, 2022), which guarantees that stable outputs are eventually achieved from stable inputs.

7.4 Integrating Humanity and Technology: Social Machines

A key subfield of CI that is still in its early days is HMCI (Smirnov & Ponomarev, 2019), also known as *hybrid CI* (Moradi et al., 2019; Peeters et al., 2021) or *hybrid CASs* (Scekic et al., 2020). In the systems we consider in this article, we can identify two main domains: (a) the domain of *space-time*, which corresponds to physical environments and their evolution, and (b) the domain of *information*, which evolves through computation (Beal et al., 2013). Of course, these two domains interact, for example, by measuring space-time to get associated information, and using information to manipulate space-time, through actuations. Now, addressing the integration of humans and machines passes through the realization that both kinds of individuals can fully operate on those two domains. That is, humans can be thought as computing machines (cf. the concept of *human computation*; Quinn & Bederson, 2011), and (computing) machines can operate in the physical world (cf. the notion of a CPS; Alur, 2015). Indeed, various terms or buzzwords are emerging to denote systems in which such integration of humans, computation, and physical systems is present—compare human CPSs (Liu & Wang, 2020), human-in-the-loop CPSs (Schirner et al., 2013), and crowd computing (Murray et al., 2010). From the perspective of computing, it is worth noting that *collective computing* based on heterogeneous human–machine collectives was identified by Abowd (2016) as the fourth generation in computing following Weiser's characterization of the evolution of computing from mainframe computing to personal computing to ubiquitous computing (Weiser, 1991).

To address the complexity of systems and unleash the potential of humans and technology, it is increasingly important to consider technical aspects together with human, social, and organizational aspects (Bucchiarone et al., 2020). In other words, a key challenge and opportunity revolves around the design of social machines (Berners-Lee, 1999; Burégio et al., 2013), hybrid societies (Hamann et al., 2016), and sociotechnical systems (Baxter & Sommerville, 2011). A social machine can be described as "a computational entity that blends computational and social processes" (Burégio et al., 2013, p. 886) and that is at the intersection of social software, people as computational units, and software as sociable entities (Berners-Lee, 1999; Burégio et al., 2013). In this respect, elements whose formalization and use might promote the engineering of CI into social machines may include macro-level and collective abstractions (Scekic et al., 2020), social concepts (Bellman et al., 2017), and coordination models (Malone & Crowston, 1994). However, several challenges remain, related to proper modeling of human computation, achieving effective communication and coordination between humans and machines, and achieving self-improving system integration (Bellman et al., 2021).

7.5 Summary of Recommendations for Future Research on ACI Engineering

This section has discussed multiple issues and directions providing for plenty of research opportunities and challenges. To summarize, we recommend that the following topics be further investigated: (a) language-based solutions to CI programming, as also fostered by recent research on macro-programming (Casadei, 2023; Sene Júnior et al., 2022), possibly also working as a foundation for explainability (Krajna et al., 2022); (b) approaches and mechanisms for controlling or steering emergence and self-organization (Gershenson et al., 2020; Varenne et al., 2015), together with efforts for building a deeper understanding of these very concepts (cf. Gignoux et al., 2017); (c) the role of CI across the various levels of modern computing systems (e.g., the application level, the middleware level, and the physical system level) (Sene Júnior et al., 2022), to address functional as well as non-functional aspects, including, for example, security, resilience, and resource efficiency; (d) designs for integrating manual and automatic approaches to CI engineering, for instance, along the lines of MARL with specifications (Ritz et al., 2021) or program synthesis (Aguzzi et al., 2022; Bastani et al., 2020) of macro-programs; and (e) integration of human intelligence with machine intelligence into hybrid, collectively intelligent systems (Peeters et al., 2021; Smirnov & Ponomarev, 2019), for example, leveraging wearable computing (Ferscha et al., 2014), ways for combining methods for human teamwork with AI, and self-organization protocols considering both humans and artificial agents (Scekic et al., 2020; Smirnov & Ponomarev, 2019).

Last, but not least, we strongly believe that the collective viewpoint has yet to find its place within software engineering practice. Recent efforts focused on formal models and languages for CASs (Nicola et al., 2020; Scekic et al., 2020; Viroli et al., 2019) might highlight a path in that direction.

8 Conclusion

CI is a rich theme that builds on multi-, inter-, and transdisciplinary collective endeavors. However, research is largely fragmented across several specific research problems (e.g., types of collective tasks), research methods (manual vs. automatic CI design), and even entire computer science research areas (hybrid systems, CASs, MASs, etc.), and comprehensive mapping studies are currently missing, making it difficult for people of diverse backgrounds to get a sense of the overall field or even of CI-related work in their subfield. This scoping review aimed at providing a comprehensive view on CI for computer scientists and engineers, with emphasis on concepts and perspectives, and also at providing some research highlights on the forms of CI that most interest them, namely, ACI, CCI, and HMCI. The final part reviews some interesting opportunities and challenges for researchers in computer science and engineering. These point at directions that, despite visionary and preliminary work, are yet to develop: CI programming, integration of manual and automatic techniques for CI engineering, integration of collectiveness and emergence, and hybrid human–machine systems.

References

Abowd, G. D. (2016). Beyond Weiser: From ubiquitous to collective computing. *Computer*, *49*(1), 17–23. https://doi.org/10.1109/MC.2016.22

Aguzzi, G., Casadei, R., & Viroli, M. (2022). Towards reinforcement learning-based aggregate computing. In M. H. ter Beek & M. Sirjani (Eds.), *Coordination models and languages—24th IFIP WG 6.1 international conference, COORDINATION 2022, held as part of the 17th international federated conference on distributed computing techniques, DISCOTEC 2022, Lucca, Italy, June 13–17, 2022, proceedings* (pp. 72–91). Springer. https://doi.org/10.1007/978-3-031-08143-9_5

Alag, S. (2008). *Collective intelligence in action*. Simon and Schuster.

Alur, R. (2015). *Principles of cyber-physical systems*. MIT Press.

Anderson, S., Bredeche, N., Eiben, A., Kampis, G., & van Steen, M. (2013). *Adaptive collective systems: Herding black sheep*. Bookprints.

Andrikopoulos, V., Bucchiarone, A., Saez, S. G., Karastoyanova, D., & Mezzina, C. A. (2014). Towards modeling and execution of collective adaptive systems. In A. Lomuscio, S. Nepal, F. Patrizi, B. Benatallah, & I. Brandic (Eds.), *Service-oriented computing—ICSOC 2013 workshops—CCSA, CSB, PASCEB, SWESE, WESEA, and PhD symposium, Berlin, Germany, December 2–5, 2013, revised selected papers* (pp. 69–81). Springer. https://doi.org/10.1007/978-3-319-06859-6_7

Audrito, G., Casadei, R., Damiani, F., Salvaneschi, G., & Viroli, M. (2022). Functional programming for distributed systems with XC. In K. Ali, & J. Vitek (Eds.), *36th European conference on object-oriented programming, ECOOP 2022, June 6–10, 2022, Berlin, Germany* (20:1–20:28). Schloss Dagstuhl—Leibniz-Zentrum für Informatik. https://doi.org/10.4230/LIPIcs.ECOOP.2022.20

Badica, A., Trawinski, B., & Nguyen, N. T. (Eds.). (2014). *Recent developments in computational collective intelligence* (Vol. 513). Springer. https://doi.org/10.1007/978-3-319-01787-7

Bastani, O., Inala, J. P., & Solar-Lezama, A. (2020). Interpretable, verifiable, and robust reinforcement learning via program synthesis. In A. Holzinger, R. Goebel, R. Fong, T. Moon, K. Müller, & W. Samek (Eds.), *XXAI—Beyond explainable AI—International workshop, held in conjunction with ICML 2020, July 18, 2020, Vienna, Austria, revised and extended papers* (pp. 207–228). Springer. https://doi.org/10.1007/978-3-031-04083-2_11

Baxter, G. D., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interactive Computing*, *23*(1), 4–17. https://doi.org/10.1016/j.intcom.2010.07.003

Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing*, *172*, 292–321. https://doi.org/10.1016/j.neucom.2015.05.116

Beal, J., Dulman, S., Usbeck, K., Viroli, M., & Correll, N. (2013). Organizing the aggregate: Languages for spatial computing. In *Formal and practical aspects of domain-specific languages: Recent developments* (pp. 436–501). IGI Global.

Beaver, L. E., & Malikopoulos, A. A. (2021). An overview on optimal flocking. *Annual Review of Control, Robotics, and Autonomous Systems*, *51*, 88–99. https://doi.org/10.1016/j.arcontrol.2021.03.004

Behjat, A., Manjunatha, H., Kumar, P. K., Jani, A., Collins, L, Ghassemi, P., Distefano, J. P., Doermann, D. S., Dantu, K., Esfahani, E. T., & Chowdhury, S. (2021). Learning robot swarm tactics over complex adversarial environments. In *International symposium on Multi-Robot and Multi-Agent Systems, MRS 2021* (pp. 83–91). IEEE. https://doi.org/10.1109/MRS50823.2021.9620707

Bellman, K. L., Botev, J., Diaconescu, A., Esterle, L., Gruhl, C., Landauer, C., Lewis, P. R., Nelson, P. R., Pournaras, E., Stein, A., & Tomforde, S. (2021). Self-improving system integration: Mastering continuous change. *Future Generation Computer Systems*, *117*, 29–46. https://doi.org/10.1016/j.future.2020.11.019

Bellman, K. L., Botev, J., Diaconescu, A., Esterle, L., Gruhl, C., Landauer, C., Lewis, P. R., Stein, A., Tomforde, S., & Würtz, R. P. (2018). Self-improving system integration—status and challenges after five years of SISSY. In *2018 IEEE 3rd International Workshops on Foundations and Applications of Self* Systems (FAS*W), Trento, Italy, September 3–7, 2018* (pp. 160–167). IEEE. https://doi.org/10.1109/FAS-W.2018.00042

Bellman, K. L., Botev, J., Hildmann, H., Lewis, P. R., Marsh, S., Pitt, J., Scholtes, I., & Tomforde, S. (2017). Socially-sensitive systems design: Exploring social potential. *IEEE Technology and Society Magazine*, *36*(3), 72–80. https://doi.org/10.1109/MTS.2017.2728727

Beni, G. (2004). From swarm intelligence to swarm robotics. In E. Sahin & W. M. Spears (Eds.), *Swarm robotics, SAB 2004 international workshop, Santa Monica, CA, USA, July 17, 2004, revised selected papers* (pp. 1–9). Springer. https://doi.org/10.1007/978-3-540-30552-1_1

Berners-Lee, Tim. (1999). *Weaving the web: The original design and ultimate destiny of the World Wide Web*. Harper.

Boardman, J. T., & Sauser, B. J. (2006). System of systems—the meaning of of. In *1st IEEE/SMC International Conference on System of Systems Engineering, SoSE 2006, Los Angeles, CA, USA, 24–26 April 2006* (pp. 1–6). IEEE. https://doi.org/10.1109/SYSOSE.2006.1652284

Bock, P. (1993). *The emergence of artificial cognition: An introduction to collective learning*. World Scientific. https://doi.org/10.1142/1831

Boissier, O., Bordini, R. H., Hubner, J., & Ricci, A. (2020). *Multi-agent oriented programming: Programming multi-agent systems using JaCaMo*. MIT Press.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford University Press. https://doi.org/10.1093/oso/9780195131581.001.0001

Bose, T., Reina, A., & Marshall, J. A. (2017). Collective decision-making. *Current Opinion in Behavioral Sciences*, *16*, 30–34. https://doi.org/10.1016/j.cobeha.2017.03.004

Bottazzi, E., Catenacci, C., Gangemi, A., & Lehmann, J. (2006). From collective intentionality to intentional collectives: An ontological perspective. *Cognitive Systems Research*, *7*(2–3), 192–208. https://doi.org/10.1016/j.cogsys.2005.11.009

Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, *7*(1), 1–41. https://doi.org/10.1007/s11721-012-0075-2

Bredèche, N., Haasdijk, E., & Prieto, A. (2018). Embodied evolution in collective robotics: A review. *Frontiers in Robotics and AI*, *5*, 12. https://doi.org/10.3389/frobt.2018.00012, PubMed: 33500899

Bucchiarone, A., D'Angelo, M., Pianini, D., Cabri, G., De Sanctis, M., Viroli, M., Casadei, R., & Dobson, S. (2020). On the social implications of collective adaptive systems. *IEEE Technology Society Magazine*, *39*(3), 36–46. https://doi.org/10.1109/MTS.2020.3012324

Burégio, V. A., de Lemos Meira, S. R., & Rosa, N. S. (2013). Social machines: A unified paradigm to describe social web-oriented systems. In *22nd international World Wide Web conference, WWW '13, Rio de Janeiro, Brazil, May 13–17, 2013, companion volume* (pp. 885–890). ACM. https://doi.org/10.1145/2487788.2488074

Busoniu, L., Babuska, R., & Schutter, B. D. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, *38*(*2*), 156–172. https://doi.org/10.1109/TSMCC.2007.913919

Canese, L., Cardarilli, G. C., Di Nunzio, L., Fazzolari, R., Giardino, D., Re, M., & Spanò, S. (2021). Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences*, *11*(11), 4948. https://doi.org/10.3390/app11114948

Carroll, J. B. (1993). *Human cognitive abilities: A survey of factor-analytic studies*. Cambridge University Press. https://doi.org/10.1017/cbo9780511571312

Casadei, R. (2020). *Engineering self-adaptive collective processes for cyber-physical ecosystems* [Unpublished doctoral dissertation]. Alma Mater Studiorum–Universitá di Bologna. http://amsdottorato.unibo.it/9380/

Casadei, R. (2023). Macroprogramming: Concepts, state of the art, and opportunities of macroscopic behavior modelling. *ACM Computing Surveys*, *55*(13s). https://doi.org/10.1145/3579353

Casadei, R., Viroli, M., Audrito, G., Pianini, D., & Damiani, F. (2021). Engineering collective intelligence at the edge with aggregate processes. *Engineering Applications of Artificial Intelligence*, *97*, 104081. https://doi.org/10.1016/j.engappai.2020.104081

Chakraborty, A., & Kar, A. K. (2017). Swarm intelligence: A review of algorithms. In S. Patnaik, X.-S. Yang, & K. Nakamatsu (Eds.), *Nature-inspired computing and optimization: Theory and applications* (pp. 475–494). Springer. https://doi.org/10.1007/978-3-319-50920-4_19

Chmait, N., Dowe, D. L., Li, Y., Green, D. G., & Insa-Cabrera, J. (2016). Factors of collective intelligence: How smart are agent collectives? In *ECAI 2016—22nd European Conference on Artificial Intelligence*, *285*, 542–550. https://doi.org/10.3233/978-1-61499-672-9-542

Chollet, F. (2019). *On the measure of intelligence*. ArXiv. https://doi.org/10.48550/arXiv.1911.01547

Ciancarini, P. (1996). Coordination models and languages as software integrators. *ACM Computing Surveys*, *28*(2), 300–302. https://doi.org/10.1145/234528.234732

Ciancarini, P., & Hankin, C. (Eds.). (2019). *Coordination languages and models, first international conference, COORDINATION'96, Cesena, Italy, April 15–17, 1996, proceedings* (Vol. 1061, p. 452). Springer. https://doi.org/10.1007/3-540-61052-9

D'Angelo, M., Gerasimou, S., Ghahremani, S., Grohmann, J., Nunes, I., Pournaras, E., & Tomforde, S. (2019). On learning in collective self-adaptive systems: State of practice and a 3d framework. In M. Litoiu, S. Clarke, & K. Tei (Eds.), *Proceedings of the 14th international symposium on software engineering for adaptive and self-managing systems, seams@icse 2019, Montreal, QC, Canada, May 25–31, 2019* (pp. 13–24). ACM. https://doi.org/10.1109/SEAMS.2019.00012

Daniels, B. C., Ellison, C. J., Krakauer, D. C., & Flack, J. C. (2016). Quantifying collectivity. *Current Opinion in Neurobiology*, *37*, 106–113. https://doi.org/10.1016/j.conb.2016.01.012, PubMed: 26874472

de Lemos, R., Giese, H., Müller, H. A., Shaw, M., Andersson, J., Litoiu, M., Schmerl, B. R., Tamura, G., Villegas, N. M., Vogel, T., Weyns, D., Baresi, L., Becker, B., Bencomo, N., Brun, Y., Cukic, B., Desmarais, R. J., Dustdar, S., Engels, G., ... Wuttke, J. (2010). Software engineering for self-adaptive systems: A second research roadmap. In R. de Lemos, H. Giese, H. A. Müller, & M. Shaw (Eds.), *Software engineering for self-adaptive systems II—International seminar, Dagstuhl Castle, Germany, October 24–29, 2010 revised selected and invited papers* (pp. 1–32). Springer. https://doi.org/10.1007/978-3-642-35813-5_1

Dennett, D. C. (1989). *The intentional stance*. MIT Press.

Dong, X., Yu, Z., Cao, W., Shi, Y., & Ma, Q. (2020). A survey on ensemble learning. *Frontiers in Computer Science*, *14*(2), 241–258. https://doi.org/10.1007/s11704-019-8208-z

Dorigo, M., Floreano, D., Gambardella, L. M., Mondada, F., Nolfi, S., Baaboura, T., Birattari, M., Bonani, M., Brambilla, M., Brutschy, A., Burnier, D., Campo, A., Christensen, A. L., Decugniere, A., Caro, G. D., Ducatelle, F., Ferrante, E., Förster, A., Gonzales, J. M., . . . Vaussard, F. (2013). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics and Automation Magazine*, *20*(4), 60–71. https://doi.org/10.1109/MRA.2013.2252996

Dorigo, M., & Stützle, T. (2019). Ant colony optimization: Overview and recent advances. In *Handbook of Metaheuristics* (pp. 311–351). Springer. https://doi.org/10.1007/978-3-319-91086-4_10

Dvoretskii, S., Gong, Z., Gupta, A., Parent, J., & Alicea, B. (2022). Braitenberg vehicles as developmental neurosimulation. *Artificial Life*, *28*(3), 369–395. https://doi.org/10.1162/artl_a_00384. PubMed: 35881679

Engelbrecht, A. P. (2007). *Computational intelligence: An introduction*. John Wiley. https://doi.org/10.1002/9780470512517

Fadul, J. A. (2009). Collective learning: Applying distributed cognition for collective intelligence. *International Journal of Learning: Annual Review*, *16*(4), 211–220. https://doi.org/10.18848/1447-9494/cgp/v16i04/46223

Faliszewski, P., Mascardi, V., Pelachaud, C., & Taylor, M. E. (Eds.) (2022). *21st international conference on autonomous agents and multiagent systems, AAMAS 2022, Auckland, New Zealand, May 9–13, 2022*. International Foundation for Autonomous Agents; Multiagent Systems. https://doi.org/10.5555/3535850

Fenwick, T. (2008). Understanding relations of individual—collective learning in work: A review of research. *Management Learning*, *39*(3), 227–243. https://doi.org/10.1177/1350507608090875

Ferber, J. (1999). *Multi-agent systems—an introduction to distributed artificial intelligence*. Addison-Wesley-Longman.

Ferscha, A. (2015). Collective adaptive systems. In K. Mase, M. Langheinrich, D. Gatica-Perez, H. Gellersen, T. Choudhury, & K. Yatani (Eds.), *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing and proceedings of the 2015 ACM international symposium on wearable computers, UBICOMP/ISWC adjunct 2015, Osaka, Japan, September 7–11, 2015* (pp. 893–895). ACM. https://doi.org/10.1145/2800835.2809508

Ferscha, A., Lukowicz, P., & Zambonelli, F. (2014). The superorganism of massive collective wearables. In A. J. Brush, A. Friday, J. A. Kientz, J. Scott, & J. Song (Eds.), *The 2014 ACM conference on ubiquitous computing, UBICOMP'14 adjunct, Seattle, WA, USA - September 13–17, 2014* (pp. 1077–1084). ACM. https://doi.org/10.1145/2638728.2659396

Figueiredo, E. M. N., Macedo, M., Siqueira, H. V., Santana, C. J., Jr., Gokhale, A., & Bastos-Filho, C. J. A. (2019). Swarm intelligence for clustering—A systematic review with new perspectives on data mining. *Engineering Applications of Artificial Intelligence*, *82*, 313–329. https://doi.org/10.1016/j.engappai.2019.04.007

Fister, I., Fister, I., Jr., Yang, X., & Brest, J. (2013). A comprehensive review of firefly algorithms. *Swarm and Evolutionary Computation*, *13*, 34–46. https://doi.org/10.1016/j.swevo.2013.06.001

Francesca, G., Brambilla, M., Brutschy, A., Trianni, V., & Birattari, M. (2014). AutoMoDe: A novel approach to the automatic design of control software for robot swarms. *Swarm Intelligence*, *8*(2), 89–112. https://doi.org/10.1007/s11721-014-0092-4

Franklin, S., & Graesser, A. C. (1996). Is it an agent, or just a program? A taxonomy for autonomous agents. In J. P. Müller, M. J. Wooldridge, & N. R. Jennings (Eds.), *Intelligent agents III, agent theories, architectures, and languages, ECAI '96 workshop 1238 (ATAL), Budapest, Hungary, August 12–13, 1996, proceedings* (pp. 21–35). Springer. https://doi.org/10.1007/BFb0013570

Gaines, B. R. (1994). The collective stance in modeling expertise in individuals and organizations. *International Journal of Expert Systems*, *7*(1), 19–49.

Galton, A., & Wood, Z. (2016). Extensional and intensional collectives and the de re/de dicto distinction. *Applied Ontology*, *11*(3), 205–226. https://doi.org/10.3233/AO-160168

Ganti, R. K., Ye, F., & Lei, H. (2011). Mobile crowdsensing: Current state and future challenges. *IEEE Communications Magazine*, *49*(11), 32–39. https://doi.org/10.1109/MCOM.2011.6069707

Garavan, T. N., & Carbery, R. (2012). Collective learning. In *Encyclopedia of the sciences of learning* (pp. 646–649). Springer. https://doi.org/10.1007/978-1-4419-1428-6_136

Garlan, D., Cheng, S., Huang, A., Schmerl, B. R., & Steenkiste, P. (2004). Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, *37*(10), 46–54. https://doi.org/10.1109/MC.2004.175

Gelernter, D., & Carriero, N. (1992). Coordination languages and their significance. *Communications of the ACM*, *35*(2), 96–107. https://doi.org/10.1145/129630.376083

Gershenson, C., Trianni, V., Werfel, J., & Sayama, H. (2020). Self-organization and artificial life. *Artificial Life*, *26*(3), 391–408. https://doi.org/10.1162/artl_a_00324, PubMed: 32697161

Gignoux, J., Chérel, G., Davies, I. D., Flint, S. R., & Lateltin, E. (2017). Emergence and complex systems: The contribution of dynamic graph theory. *Ecological Complexity*, *31*, 34–49. https://doi.org/10.1016/j.ecocom.2017.02.006

Gomes, J. C., Urbano, P., & Christensen, A. L. (2013). Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, *7*(2–3), 115–144. https://doi.org/10.1007/s11721-013-0081-z

Gronauer, S., & Diepold, K. (2022). Multi-agent deep reinforcement learning: A survey. *Artificial Intelligence Review*, *55*(2), 895–943. https://doi.org/10.1007/s10462-021-09996-w

Gruber, T. (2008). Collective knowledge systems: Where the social web meets the semantic web. *Journal of Web Semantics*, *6*(1), 4–13. https://doi.org/10.1016/j.websem.2007.11.011

Gulwani, S., Polozov, O., & Singh, R. (2017). Program synthesis. *Foundations and Trends in Programming Languages*, *4*(1–2), 1–119. https://doi.org/10.1561/2500000010

Guo, B., Yu, Z., Zhou, X., & Zhang, D. (2014). From participatory sensing to mobile crowd sensing. In *2014 IEEE international conference on pervasive computing and communication workshops, PerCom 2014 workshops, Budapest, Hungary, March 24–28, 2014* (pp. 593–598). IEEE. https://doi.org/10.1109/PerComW.2014.6815273

Gupta, A., Savarese, S., Ganguli, S., & Fei-Fei, L. (2021). Embodied intelligence via learning and evolution. *Nature Communications*, *12*(1), 5721. https://doi.org/10.1038/s41467-021-25874-z, PubMed: 34615862

Guttmann, C. (2009). Towards a taxonomy of decision making problems in multi-agent systems. In L. Braubach, W. van der Hoek, P. Petta, & A. Pokahr (Eds.), *Multiagent system technologies, 7th German conference, MATES 2009, Hamburg, Germany, September 9–11, 2009, Proceedings* (pp. 195–201), Springer. https://doi.org/10.1007/978-3-642-04143-3_19

Ha, D., & Tang, Y. (2021). *Collective intelligence for deep learning: A survey of recent developments*. ArXiv. https://arxiv.org/abs/2111.14377

Hamann, H., Khaluf, Y., Botev, J., Soorati, M. D., Ferrante, E., Kosak, O., Montanier, J., Mostaghim, S., Redpath, R., Timmis, J., Veenstra, F., Wahby, M., & Zamuda, A. (2016). Hybrid societies: Challenges and perspectives in the design of collective behavior in self-organizing systems. *Frontiers in Robotics and AI*, *3,*14. https://doi.org/10.3389/frobt.2016.00014

Hasbach, J. D., & Bennewitz, M. (2022). The design of self-organizing human-swarm intelligence. *Adaptive Behavior*, *30*(4), 361–386. https://doi.org/10.1177/10597123211017550

He, F., Pan, Y., Lin, Q., Miao, X., & Chen, Z. (2019). Collective intelligence: A taxonomy and survey. *IEEE Access*, *7*, 170213–170225. https://doi.org/10.1109/ACCESS.2019.2955677

Hecker, A. (2012). Knowledge beyond the individual? Making sense of a notion of collective knowledge in organization theory. *Organization Studies*, *33*(3), 423–445. https://doi.org/10.1177/0170840611433995

Hegedüs, I., Danner, G., & Jelasity, M. (2021). Decentralized learning works: An empirical comparison of gossip learning and federated learning. *Journal of Parallel Distributed Computing*, *148*, 109–124. https://doi.org/10.1016/j.jpdc.2020.10.006

Hernandez-Leal, P., Kartal, B., & Taylor, M. E. (2019). A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, *33*(6), 750–797. https://doi.org/10.1007/s10458-019-09421-1

Hernández-Orallo, J. (2017). Evaluation in artificial intelligence: From task-oriented to ability-oriented measurement. *Artificial Intelligence Review*, *48*(3), 397–447. https://doi.org/10.1007/s10462-016-9505-7

Hernández-Orallo, J., & Dowe, D. L. (2010). Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, *174*(18), 1508–1539. https://doi.org/10.1016/j.artint.2010.09.006

Heylighen, F. (2013). From human computation to the global brain: The self-organization of distributed intelligence. In P. Michelucci (Ed.), *Handbook of human computation* (pp. 897–909). Springer. https://doi.org/10.1007/978-1-4614-8806-4_73

Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, *1*(3). http://www.complex-systems.com/abstracts/v01_i03_a06/

Hmelo-Silver, C. E. (2003). Analyzing collaborative knowledge construction: Multiple methods for integrated understanding. *Computers and Education*, *41*(4), 397–420. https://doi.org/10.1016/j.compedu.2003.07.001

Horling, B., & Lesser, V. R. (2004). A survey of multi-agent organizational paradigms. *Knowledge Engineering Review*, *19*(4), 281–316. https://doi.org/10.1017/S0269888905000317

Huang, Y.-C., & Chin, Y.-C. (2018). Transforming collective knowledge into team intelligence: The role of collective teaching. *Journal of Knowledge Management*, *22*(6), 1243–1263. https://doi.org/10.1108/jkm-03-2017-0106

Ibrahim, D. (2016). An overview of soft computing. *Procedia Computer Science*, *102*, 34–38. https://doi.org/10.1016/j.procs.2016.09.366

Iivari, J. (2016). Endogenously emergent information systems. In J. Goluchowski, M. Pankowska, C. Barry, M. Lang, H. Linger, & C. Schneider (Eds.), *Information systems development: Complexity in information systems development—proceedings of the 25th international conference on information systems development, ISD 2016, Katowice, Poland, August 24–26, 2016*. University of Economics in Katowice/Association for Information Systems. http://aisel.aisnet.org/isd2014/proceedings2016/ISDMethodologies/5

Jensen, A. R. (1989). The relationship between learning and intelligence. *Learning and Individual Differences*, *1*(1), 37–62. https://doi.org/10.1016/1041-6080(89)90009-5

Jung, J. J. (2017). Computational collective intelligence with big data: Challenges and opportunities. *Future Generation Computer Systems*, *66*, 87–88. https://doi.org/10.1016/j.future.2016.08.021

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K. A., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., Rouayheb, S. E., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, *14*(1–2), 1–210. https://doi.org/10.1561/2200000083

Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., & Schmickl, T. (2015). Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. In Q. Chen, P. Torroni, S. Villata, J. Y. Hsu, & A. Omicini (Eds.), *PRIMA 2015: Principles and practice of multi-agent systems—18th international conference, Bertinoro, Italy, October 26–30, 2015, proceedings* (pp. 201–217). Springer. https://doi.org/10.1007/978-3-319-25524-8_13

Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. *Computer*, *36*(1), 41–50. https://doi.org/10.1109/MC.2003.1160055

Kimmerle, J., Cress, U., & Held, C. (2010). The interplay between individual and collective knowledge: Technologies for organisational learning and knowledge building. *Knowledge Management Research and Practice*, *8*(1), 33–44. https://doi.org/10.1057/kmrp.2009.36

Kolling, A., Walker, P. M., Chakraborty, N., Sycara, K. P., & Lewis, M. (2016). Human interaction with robot swarms: A survey. *IEEE Transactions on Human–Machine Systems*, *46*(1), 9–26. https://doi.org/10.1109/THMS.2015.2480801

Krajna, A., Brcic, M., Lipic, T., & Doncevic, J. (2022). *Explainability in reinforcement learning: Perspective and position*. ArXiv. https://doi.org/10.48550/arXiv.2203.11547

Krause, J., Ruxton, G. D., & Krause, S. (2010). Swarm intelligence in animals and humans. *Trends in Ecology and Evolution*, *25*(1), 28–34. https://doi.org/10.1016/j.tree.2009.06.016, PubMed: 19735961

Laan, A., Madirolas, G., & de Polavieja, G. G. (2017). Rescuing collective wisdom when the average group opinion is wrong. *Frontiers in Robotics and AI*, *4*, 56. https://doi.org/10.3389/frobt.2017.00056

Le Bon, G. (2002). *The crowd: A study of the popular mind*. Courier (Original work published 1895).

Legg, S., & Hutter, M. (2007). A collection of definitions of intelligence. In B. Goertzel & P. Wang (Eds.), *Proceedings of the 2007 conference on Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms: Proceedings of the AGI Workshop 2006* (pp. 17–24). IOS Press.

Liu, Z., & Wang, J. (2020). Human-cyber-physical systems: Concepts, challenges, and research opportunities. *Frontiers of Information Technology and Electronic Engineering*, *21*(11), 1535–1553. https://doi.org/10.1631/FITEE.2000537

Lo, A. W. (2015). The wisdom of crowds vs. the madness of mobs. In *Handbook of collective intelligence* (pp. 21–37). MIT Press.

Loreti, M., & Hillston, J. (2016). Modelling and analysis of collective adaptive systems with CARMA and its tools. In M. Bernardo, R. D. Nicola, & J. Hillston (Eds.), *Formal methods for the quantitative evaluation of collective adaptive systems—16th international school on formal methods for the design of computer, communication, and* *software systems, SFM 2016, Bertinoro, Italy, June 20–24, 2016, advanced lectures* (pp. 83–119). Springer. https://doi.org/10.1007/978-3-319-34096-8_4

Lykourentzou, I., Vergados, D. J., & Loumos, V. (2009). Collective intelligence system engineering. In R. Chbeir, Y. Badr, E. Kapetanios, & A. J. M. Traina (Eds.), *MEDES '09: International ACM conference on management of emergent digital ecosystems, Lyon, France, October 27–30, 2009* (pp. 134–140). ACM. https://doi.org/10.1145/1643823.1643848

Mainland, G., Kang, L., Lahaie, S., Parkes, D. C., & Welsh, M. (2004). Using virtual markets to program global behavior in sensor networks. In Y. Berbers & M. Castro (Eds.), *Proceedings of the 11st ACM SIGOPS european workshop, Leuven, Belgium, September 19–22, 2004* (p. 1), ACM. https://doi.org/10.1145/1133572.1133587

Maleszka, M. (2019). Application of collective knowledge diffusion in a social network environment. *Enterprise Information Systems*, *13*(7–8), 1120–1142. https://doi.org/10.1080/17517575.2018.1526325

Malone, T. W., & Bernstein, M. S. (2015). *Handbook of collective intelligence*. MIT Press.

Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Computing Surveys*, *26*(1), 87–119. https://doi.org/10.1145/174666.174668

Mamei, M., Zambonelli, F., & Leonardi, L. (2004). Co-fields: A physically inspired approach to motion coordination. *IEEE Pervasive Computing*, *3*(2), 52–61. https://doi.org/10.1109/MPRV.2004.1316820

Masolo, C., Vieu, L., Ferrario, R., Borgo, S., & Porello, D. (2020). Pluralities, collectives, and composites. In B. Brodaric & F. Neuhaus (Eds.), *Formal ontology in information systems: Proceedings of the 11th international conference (FOIS 2020)* (Vol. 330, pp. 186–200). IOS Press. https://doi.org/10.3233/FAIA330

Mataric, M. J. (2007). *The robotics primer*. MIT Press.

Mavrovouniotis, M., Li, C., & Yang, S. (2017). A survey of swarm intelligence for dynamic optimization: Algorithms and applications. *Swarm and Evolutionary Computation*, *33*, 1–17. https://doi.org/10.1016/j.swevo.2016.12.005

Millhouse, T., Moses, M., & Mitchell, M. (2021). *Frontiers in collective intelligence: A workshop report*. ArXiv. https://arxiv.org/abs/2112.06864

Minati, G. (2018). Phenomenological structural dynamics of emergence: An overview of how emergence emerges. In *The systemic turn in human and natural sciences* (pp. 1–39), Springer. https://doi.org/10.1007/978-3-030-00725-6_1

Mitchell, M. (2011). Ubiquity symposium: Biological computation. *Ubiquity*, *2011*(February), 3. http://doi.acm.org/10.1145/1940721.1944826

Mizzi, A., Ellul, J., & Pace, G. J. (2018). D'Artagnan: An embedded DSL framework for distributed embedded systems. In *Proceedings of the Real World Domain Specific Languages Workshop, RWDSL@CGO 2018, Vienna, Austria, February 24–24, 2018* (pp. 2:1–2:9). ACM. https://doi.org/10.1145/3183895.3183899

Mobus, G. E., & Kalton, M. C. (2015). *Principles of systems science*. Springer.

Montes de Oca, M. A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., & Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, *5*(3–4), 305–327. https://doi.org/10.1007/s11721-011-0062-z

Moradi, M., Moradi, M., Bayat, F., & Toosi, A. N. (2019). Collective hybrid intelligence: Towards a conceptual framework. *International Journal of Crowd Science*, *3*(2), 198–220.

Müller-Schloer, C., & Sick, B. (2006). Emergence in organic computing systems: Discussion of a controversial concept. In L. T. Yang, H. Jin, J. Ma, & T. Ungerer (Eds.), *Autonomic and trusted computing, third international conference, ATC 2006, Wuhan, China, September 3–6, 2006, proceedings* (pp. 1–16). Springer. https://doi.org/10.1007/11839569_1

Müller-Schloer, C., & Tomforde, S. (2017). *Organic computing—technical systems for survival in the real world*. Birkhäuser. https://doi.org/10.1007/978-3-319-68477-2

Munn, Z., Peters, M. D. J., Stern, C., Tufanaru, C., McArthur, A., & Aromataris, E. (2018). Systematic review or scoping review? Guidance for authors when choosing between a systematic or scoping review approach. *BMC Medical Research Methodology 18*(1), 143. https://doi.org/10.1186/s12874-018-0611-x, PubMed: 30453902

Murray, D. G., Yoneki, E., Crowcroft, J., & Hand, S. (2010). The case for crowd computing. In *Proceedings of the second ACM SIGCOMM workshop on networking, systems, and applications on mobile handhelds* (pp. 39–44). ACM. https://doi.org/10.1145/1851322.1851334

Narayanan, A., Korium, M. S., Melgarejo, D. C., Hussain, H. M., de Sena, A. S., Silva, P. E. G., Gutierrez-Rojas, D., Ullah, M., Nezhad, A. E., Rasti, M., Pournaras, E., & Nardelli, P. H. J. (2022). Collective intelligence using 5G: Concepts, applications, and challenges in sociotechnical environments. *IEEE Access*, *10*, 70394–70417. https://doi.org/10.1109/ACCESS.2022.3184035

Navarro, I, & Matía, F. (2013). A survey of collective movement of mobile robots. *International Journal of Advanced Robotic Systems*, *10*(1), 73. https://doi.org/10.5772/54600

Nedjah, N., & Junior, L. S. (2019). Review of methodologies and tasks in swarm robotics towards standardization. *Swarm and Evolution Computation*, *50*. https://doi.org/10.1016/j.swevo.2019.100565

Newton, R., A, & Welsh, M. (2005). Building up to macroprogramming: An intermediate language for sensor networks. In *Proceedings of the fourth international symposium on information processing in sensor networks, IPSN 2005, April 25–27, 2005, UCLA, Los Angeles, California, USA* (pp. 37–44). IEEE. https://doi.org/10.1109/IPSN.2005.1440891

Nguyen, B. H., Xue, B., & Zhang, M. (2020). A survey on swarm intelligence approaches to feature selection in data mining. *Swarm and Evolutionary Computing*, *54*, 100663. https://doi.org/10.1016/j.swevo.2020.100663

Nguyen, N. T. (2008). Inconsistency of knowledge and collective intelligence. *Cybernetics and Systems*, *39*(6), 542–562. https://doi.org/10.1080/01969720802188268

Nguyen, N. T., Hwang, D., & Szczerbicki, E. (2019). Computational collective intelligence for enterprise information systems. *Enterprise Information Systems*, *13*(7–8), 933–934. https://doi.org/10.1080/17517575.2019.1640394

Nguyen, N. T., Kowalczyk, R., & Chen, S. (Eds.) (2009). *Computational collective intelligence: Semantic web, social networks and multiagent systems, first international conference, ICCCI 2009, Wroclaw, Poland, October 5–7, 2009. proceedings* (Vol. 5796). Springer. https://doi.org/10.1007/978-3-642-04441-0

Nicola, R. D., Jähnichen, S., & Wirsing, M. (2020). Rigorous engineering of collective adaptive systems: Special section. *International Journal on Software Tools for Technology Transfer*, *22*(4), 389–397. https://doi.org/10.1007/s10009-020-00565-0

Nielsen, C. B., Larsen, P. G., Fitzgerald, J. S., Woodcock, J., & Peleska, J. (2015). Systems of systems engineering: Basic concepts, model-based techniques, and research directions. *ACM Computing Surveys*, *48*(2), 18:1–18:41. https://doi.org/10.1145/2794381

Nolfi, S., & Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, *7*(1), 89–113. https://doi.org/10.1023/A:1008973931182

Noor, J., Tseng, H., Garcia, L., & Srivastava, M. B. (2019). DDFlow: Visualized declarative programming for heterogeneous IoT networks. In O. Landsiedel, & K. Nahrstedt (Eds.), *Proceedings of the international conference on internet of things design and implementation, IoTDI 2019, Montreal, QC, Canada, April 15–18, 2019* (pp. 172–177). ACM. https://doi.org/10.1145/3302505.3310079

Odell, J. (2002). Objects and agents compared. *Journal of Object Technology*, *1*(1), 41–53. https://doi.org/10.5381/jot.2002.1.1.c4

Oh, H., Shirazi, A. R., Sun, C., & Jin, Y. (2017). Bio-inspired self-organising multi-robot pattern formation: A review. *Robotics and Autonomous Systems*, *91*, 83–100. https://doi.org/10.1016/j.robot.2016.12.006

Oliehoek, F. A., & Amato, C. (2016). *A concise introduction to decentralized POMDPs* (1st ed.). Springer.

Oliver, P. E. (1993). Formal models of collective action. *Annual Review of Sociology*, *19*(1), 271–300. https://doi.org/10.1146/annurev.so.19.080193.001415

Omicini, A., Ricci, A., & Viroli, M. (2008). Artifacts in the A&A meta-model for multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, *17*(13), 432–456. https://doi.org/10.1007/s10458-008-9053-x

O'Neill, M., & Spector, L. (2020). Automatic programming: The open issue? *Genetic Programming and Evolvable Machines*, *21*(1–2), 251–262. https://doi.org/10.1007/s10710-019-09364-2

Peeters, M. M. M., van Diggelen, J., van den Bosch, K., Bronkhorst, A. W., Neerincx, M. A., Schraagen, J. M., & Raaijmakers, S. (2021). Hybrid collective intelligence in a human-AI society. *AI Society*, *36*(1), 217–238. https://doi.org/10.1007/s00146-020-01005-y

Petruzzi, P. E., Busquets, D., & Pitt, J. V. (2015). A generic social capital framework for optimising self-organised collective action. In *2015 IEEE 9th international conference on self-adaptive and self-organizing systems, Cambridge, MA, USA, September 21–25, 2015* (pp. 21–30). IEEE. https://doi.org/10.1109/SASO.2015.10

Petticrew, M., & Roberts, H. (2008). *Systematic reviews in the social sciences: A practical guide*. John Wiley.

Pianini, D., Casadei, R., & Viroli, M. (2022). Self-stabilising priority-based multi-leader election and network partitioning. In *2022 IEEE international conference on autonomic computing and self-organizing systems (ACSOS)* (pp. 81–90). IEEE. https://doi.org/10.1109/acsos55765.2022.00026

Pianini, D., Casadei, R., Viroli, M., & Natali, A. (2021). Partitioned integration and coordination via the self-organising coordination regions pattern. *Future Generration Computing Systems*, *114*, 44–68. https://doi.org/10.1016/j.future.2020.07.032

Pianini, D., Pettinari, F., Casadei, R., & Esterle, L. (2022). A collective adaptive approach to decentralised k-coverage in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, *17*, 4:1–4:39. https://doi.org/10.1145/3547145

Pinciroli, C., & Beltrame, G. (2016). Buzz: A programming language for robot swarms. *IEEE Software*, *33*(4), 97–100. https://doi.org/10.1109/MS.2016.95

Pournaras, E. (2020a). Collective learning: A 10-year odyssey to human-centered distributed intelligence. In *IEEE international conference on autonomic computing and self-organizing systems, ACSOS 2020, Washington, DC, USA, August 17–21, 2020* (pp. 205–214). IEEE. https://doi.org/10.1109/ACSOS49614.2020.00043

Pournaras, E. (2020b). Proof of witness presence: Blockchain consensus for augmented democracy in smart cities. *Journal of Parallel Distributed Computing*, *145*, 160–175. https://doi.org/10.1016/j.jpdc.2020.06.015

Pournaras, E., Pilgerstorfer, P., & Asikis, T. (2018). Decentralized collective learning for self-managed sharing economies. *ACM Transactions on Autonomous and Adaptive Systems*, *13*(2), 10:1–10:33. https://doi.org/10.1145/3277668

Prasetyo, J., Masi, G. D., & Ferrante, E. (2019). Collective decision making in dynamic environments. *Swarm Intelligence*, *13*(3–4), 217–243. https://doi.org/10.1007/s11721-019-00169-8

Psaier, H., & Dustdar, S. (2011). A survey on self-healing systems: Approaches and systems. *Computing*, *91*(1), 43–73. https://doi.org/10.1007/s00607-010-0107-y

Qiu, W., & Zhu, H. (2022, April 25). *Programmatic reinforcement learning without oracles* [Paper presentation]. The tenth international conference on learning representations ICLR 2022, virtual. https://openreview.net/forum?id=6Tk2noBdvxt

Quinn, A. J., & Bederson, B. B. (2011). Human computation: A survey and taxonomy of a growing field. In D. S. Tan, S. Amershi, B. Begole, W. A. Kellogg, & M. Tungare (Eds.), *Proceedings of the international conference on human factors in computing systems, CHI 2011, Vancouver, BC, Canada, May 7–12, 2011* (pp. 1403–1412). ACM. https://doi.org/10.1145/1978942.1979148

Rainey, L. B., & Jamshidi, M. (Eds.). (2018). *Engineering emergence*. CRC Press. https://doi.org/10.1201/9781138046412

Rajasekhar, A., Lynn, N., Das, S., & Suganthan, P. N. (2017). Computing with the collective intelligence of honey bees—a survey. *Swarm and Evolutionary Computation*, *32*, 25–48. https://doi.org/10.1016/j.swevo.2016.06.001

Reeve, C. L., Bonaccio, S., Chamorro-Premuzic, T., Furnham, A., & von Stumm, S. (2011). The nature and structure of "intelligence." In *The Wiley-Blackwell handbook of individual differences* (pp. 187–216). Wiley-Blackwell.

Reina, A., Miletitch, R., Dorigo, M., & Trianni, V. (2015). A quantitative micro–macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, *9*(2–3), 75–102. https://doi.org/10.1007/s11721-015-0105-y

Ritz, F., Phan, T., Müller, R., Gabor, T., Sedlmeier, A., Zeller, M., Wieghardt, J., Schmid, R. N., Sauer, H., Klein, C., & Linnhoff-Popien, C. (2021). Specification aware multiagent reinforcement learning. In A. P. Rocha, L. Steels, & H. J. van den Herik (Eds.), *Agents and artificial intelligence—13th international conference, ICAART 2021, virtual event, February 4–6, 2021, revised selected papers* (pp. 3–21). Springer. https://doi.org/10.1007/978-3-031-10161-8_1

Rizk, Y., Awad, M., & Tunstel, E. W. (2018). Decision making in multiagent systems: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, *10*(3), 514–529. https://doi.org/10.1109/TCDS.2018.2840971

Rodríguez, A., Grushin, A., & Reggia, J. A. (2007). Swarm intelligence systems using guided self-organization for collective problem solving. *Advances in Complex Systems*, *10*(Suppl. 1), 5–34. https://doi.org/10.1142/S0219525907001069

Rossi, F., Bandyopadhyay, S., Wolf, M. T., & Pavone, M. (2018). *Review of multi-agent algorithms for collective behavior: A structural taxonomy*. ArXiv. http://arxiv.org/abs/1803.05464

Rostami, M., Kolouri, S., Kim, K., & Eaton, E. (2018). Multi-agent distributed lifelong learning for collective knowledge acquisition. In E. André, S. Koenig, M. Dastani, & G. Sukthankar (Eds.), *Proceedings of the 17th international conference on autonomous agents and multiagent systems, AAMAS 2018, Stockholm, Sweden, July 10–15, 2018* (pp. 712–720). International Foundation for Autonomous Agents/ACM. http://dl.acm.org/citation.cfm?id=3237489

Russell, S., & Norvig, P. (2020). *Artificial Intelligence: A modern approach* (4th ed.). Pearson. http://aima.cs.berkeley.edu/

Salehie, M., & Tahvildari, L. (2009). Self-adaptive software: Landscape and research challenges. *ACM Transactions on Autonomous and Adaptive Systems*, *4*(2), 14:1–14:42. https://doi.org/10.1145/1516533.1516538

Salminen, J. (2012). *Collective intelligence in humans: A literature review*. ArXiv. http://arxiv.org/abs/1204.3401

Sari, A., Tosun, A., & Alptekin, G. I. (2019). A systematic literature review on crowdsourcing in software engineering. *Journal of Systems and Software*, *153*, 200–219. https://doi.org/10.1016/j.jss.2019.04.027

Scekic, O., Schiavinotto, T., Videnov, S., Rovatsos, M., Truong, H. L., Miorandi, D., & Dustdar, S. (2020). A programming model for hybrid collaborative adaptive systems. *IEEE Transactions on Emerging Topics in Computing*, *8*(1), 6–19. https://doi.org/10.1109/TETC.2017.2702578

Schirner, G., Erdogmus, D., Chowdhury, K. R., & Padir, T. (2013). The future of human-in-the-loop cyber-physical systems. *Computer*, *46*(1), 36–45. https://doi.org/10.1109/MC.2013.31

Schmickl, T., Möslinger, C., & Crailsheim, K. (2006). Collective perception in a robot swarm. In E. Sahin, W. M. Spears, & A. F. T. Winfield (Eds.), *Swarm robotics, second international workshop, SAB 2006, Rome, Italy, September 30–October 1, 2006, revised selected papers* (pp. 144–157). Springer. https://doi.org/10.1007/978-3-540-71541-2_10

Schranz, M., Caro, G. A. D., Schmickl, T., Elmenreich, W., Arvin, F., Sekercioglu, Y. A., & Sende, M. (2021). Swarm intelligence and cyber-physical systems: Concepts, challenges and future trends. *Swarm and Evolutionary Computation*, *60*, 100762. https://doi.org/10.1016/j.swevo.2020.100762

Schweikard, D. P., & Schmid, H. B. (2013). Collective intentionality. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2021 ed.). https://plato.stanford.edu/entries/collective-intentionality/

Segaran, T. (2007). *Programming collective intelligence: Building smart Web 2.0 applications*. O'Reilly.

Sene Júnior, I. G., Santana, T. S., Bulcão-Neto, R. F., & Porter, B. (2022). The state of the art of macro programming in IoT: An update. *Journal of Internet Services and Applications*, *13*(1), 54–65. https://doi.org/10.5753/jisa.2022.2372

Shoham, Y. (1993). Agent-oriented programming. *Artificial Intelligence*, *60*(1), 51–92. https://doi.org/10.1016/0004-3702(93)90034-9

Siddique, N. H., & Adeli, H. (2015). Nature inspired computing: An overview and some future directions. *Cognitive Computation*, *7*(6), 706–714. https://doi.org/10.1007/s12559-015-9370-8, PubMed: 26693257

Sieben, A., Schumann, J., & Seyfried, A. (2017). Collective phenomena in crowds—where pedestrian dynamics need social psychology. *PLoS ONE*, *12*(6), 1–19. https://doi.org/10.1371/journal.pone.0177328, PubMed: 28591142

Sigaud, O. (2022). Combining evolution and deep reinforcement learning for policy search: A survey. *ACM Transactions on Evolutionary Learning and Optimization*. https://doi.org/10.1145/3569096

Smirnov, A., & Ponomarev, A. (2019). Decision support based on human-machine collective intelligence: Major challenges. In *Internet of things, smart spaces, and next generation networks and systems* (pp. 113–124), Springer. https://doi.org/10.1007/978-3-030-30859-9_10

Smith, E. A. (2010). Communication and collective action: Language and the evolution of human cooperation. *Evolution and Human Behavior*, *31*(4), 231–245. https://doi.org/10.1016/j.evolhumbehav.2010.03.001

Solar-Lezama, A. (2009). The sketching approach to program synthesis. In Z. Hu (Ed.), *Programming languages and systems, 7th Asian symposium, APLAS 2009, Seoul, Korea, December 14–16, 2009, proceedings* (pp. 4–13). Springer. https://doi.org/10.1007/978-3-642-10672-9_3

Sotnik, G. (2018). The SOSIEL platform: Knowledge-based, cognitive, and multi-agent. *Biologically Inspired Cognitive Architectures*, *26*, 103–117. https://doi.org/10.1016/j.bica.2018.09.001

Stein, A., Tomforde, S., Botev, J., & Lewis, P. R. (2021). Lifelike computing systems. In *Proceedings of the Lifelike Computing Systems Workshop (LIFELIKE)* ceur-ws.org. https://ceur-ws.org/Vol-3007/2021-paper-1.pdf

Stradner, J., Thenius, R., Zahadat, P., Hamann, H., Crailsheim, K., & Schmickl, T. (2013). Algorithmic requirements for swarm intelligence in differently coupled collective systems. *Chaos, Solitons, and Fractals*, *50*, 100–114. https://doi.org/10.1016/j.chaos.2013.01.011, PubMed: 23805030

Sumpter, D. J. (2010). *Collective animal behavior*. Princeton University Press. https://doi.org/10.1515/9781400837106

Suran, S., Pattanaik, V., & Draheim, D. (2020). Frameworks for collective intelligence: A systematic literature review. *ACM Computing Surveys*, *53*(1). https://doi.org/10.1145/3368986

Surowiecki, J. (2005). *The wisdom of crowds*. Anchor.

Szuba, T. M. (2001). *Computational collective intelligence*. John Wiley.

Tang, M., & Liao, H. (2021). From conventional group decision making to large-scale group decision making: What are the challenges and how to meet them in big data era? A state-of-the-art survey. *Omega*, 100, 102141. https://doi.org/10.1016/j.omega.2019.102141

Telikani, A., Tahmassebi, A., Banzhaf, W., & Gandomi, A. H. (2022). Evolutionary machine learning: A survey. *ACM Computing Surveys*, *54*(8), 161:1–161:35. https://doi.org/10.1145/3467477

ter Beek, M. H., & Sirjani, M. (Eds.). (2022). *Coordination models and languages—24th IFIP WG 6.1 international conference, COORDINATION 2022, held as part of the 17th international federated conference on distributed computing techniques, DISCOTEC 2022, Lucca, Italy, June 13–17, 2022, proceedings* (Vol. 13271). Springer. https://doi.org/10.1007/978-3-031-08143-9

Trianni, V. (2008). *Evolutionary swarm robotics—Evolving self-organising behaviours in groups of autonomous robots*. Springer. https://doi.org/10.1007/978-3-540-77612-3

Tsvetkova, M., Yasseri, T., Meyer, E. T., Pickering, J. B., Engen, V., Walland, P., Lüders, M., Følstad, A., & Bravos, G. (2017). Understanding human-machine networks: A cross-disciplinary survey. *ACM Computing Surveys*, *50*(1), 12:1–12:35. https://doi.org/10.1145/3039868

Tumer, K., & Wolpert, D. H. (1999). *Avoiding Braess' paradox through collective intelligence*. ArXiv. https://doi.org/10.48550/arXiv.cs/9912012

Tumer, K., & Wolpert, D. H. (2004). A survey of collectives. In K. Tumer & D. H. Wolpert (Eds.), *Collectives and the design of complex systems* (pp. 1–42). Springer. https://doi.org/10.1007/978-1-4419-8909-3_1

Valentini, G., Ferrante, E., & Dorigo, M. (2017). The best-of-*n* problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, *4*, 9. https://doi.org/10.3389/frobt.2017.00009

Valentini, G., Hamann, H., & Dorigo, M. (2014). Self-organized collective decision making: The weighted voter model. In A. L. C. Bazzan, M. N. Huhns, A. Lomuscio, & P. Scerri (Eds.), *International conference on autonomous agents and multi-agent systems, AAMAS '14, Paris, France, May 5–9, 2014* (pp. 45–52). IFAAMAS/ACM. http://dl.acm.org/citation.cfm?id=2615742

van der Hoek, W., & Wooldridge, M. J. (2003). Towards a logic of rational agency. *Logic Journal of the IGPL*, *11*(2), 135–159. https://doi.org/10.1093/jigpal/11.2.135

van Gerven, M. (2017). Computational foundations of natural intelligence. *Frontiers in Computational Neuroscience*, *11*, 112. https://doi.org/10.3389/fncom.2017.00112, PubMed: 29375355

Varenne, F., Chaigneau, P., Petitot, J., & Doursat, R. (2015). Programming the emergence in morphogenetically architected complex systems. *Acta Biotheoretica*, *63*(3), 295–308. https://doi.org/10.1007/s10441-015-9262-z, PubMed: 26024971

Verbraeken, J., Wolting, M., Katzy, J., Kloppenburg, J., Verbelen, T., & Rellermeyer, J. S. (2020). A survey on distributed machine learning. *ACM Computing Surveys*, *53*(2), 30:1–30:33. https://doi.org/10.1145/3377454

Verma, A., Murali, V., Singh, R., Kohli, P., & Chaudhuri, S. (2018). Programmatically interpretable reinforcement learning. In J. G. Dy & A. Krause (Eds.), *Proceedings of the 35th international conference on machine learning, ICML 2018, Stockholmsmäs-san, Stockholm, Sweden, July 10–15, 2018* (pp. 5052–5061). PMLR. http://proceedings.mlr.press/v80/verma18a.html

Viroli, M., Beal, J., Damiani, F., Audrito, G., Casadei, R., & Pianini, D. (2019). From distributed coordination to field calculus and aggregate computing. *Journal of Logical and Algebraic Methods in Programming*, 109, 100486. https://doi.org/10.1016/j.jlamp.2019.100486

Wang, Y. (2009). On abstract intelligence: Toward a unifying theory of natural, artificial, machinable, and computational intelligence. *International Journal of Software Science and Computational Intelligence*, *1*(1), 1–17. https://doi.org/10.4018/jssci.2009010101

Watson, R. A., Ficici, S. G., & Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, *39*(1), 1–18. https://doi.org/10.1016/S0921-8890(02)00170-7

Wechsler, D. (1946). *The measurement of adult intelligence* (3rd ed.). Williams & Wilkins. https://doi.org/10.1037/11329-000

Weiser, M. (1991). The computer for the 21st century. *Scientific American*, *265*(3), 94–105.

Weyns, D., & Michel, F. (2014). Agent environments for multi-agent systems—a research roadmap. In D. Weyns & F. Michel (Eds.), *Agent environments for multi-agent systems IV—4th international workshop, E4MAS 2014—10 years later, Paris, France, May 6, 2014, revised selected and invited papers* (pp. 3–21). Springer. https://doi.org/10.1007/978-3-319-23850-0_1

Weyns, D., Schmerl, B. R., Grassi, V., Malek, S., Mirandola, R., Prehofer, C., Wuttke, J., Andersson, J., Giese, H., & Göschka, K. M. (2010). On patterns for decentralized control in self-adaptive systems. In R. de Lemos, H. Giese, H. A. Müller, & M. Shaw (Eds.), *Software engineering for self-adaptive systems II—international seminar, Dagstuhl Castle, Germany, October 24–29, 2010 revised selected and invited papers* (pp. 76–107). Springer. https://doi.org/10.1007/978-3-642-35813-5_4

Wolf, T. D., & Holvoet, T. (2004). Emergence versus self-organisation: Different concepts but promising when combined. *Engineering Self-Organising Systems, Methodologies and Applications*, *3464*, 1–15. https://doi.org/10.1007/11494676_1

Wolpert, D. H. (2003). Collective intelligence. In *Computational intelligence: The experts speak* (pp. 245–260). John Wiley.

Wood, Z. (2016). *Considering collectives: Roles, members and goals*. IOS Press.

Wood, Z., & Galton, A. (2009). A taxonomy of collective phenomena. *Applied Ontology*, *4*(3–4), 267–292.

Wooldridge, M. J. (2009). *An introduction to multiagent systems* (2nd ed.). John Wiley.

Woolley, A. W., Aggarwal, I., & Malone, T. W. (2015). Collective intelligence and group performance. *Current Directions in Psychological Science*, *24*(6), 420–424. https://doi.org/10.1177/0963721415599543

Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, *330*(6004), 686–688. https://doi.org/10.1126/science.1193147, PubMed: 20929725

Yang, T., Yi, X., Wu, J., Yuan, Y., Wu, D., Meng, Z., Hong, Y., Wang, H., Lin, Z., & Johansson, K. H. (2019). A survey of distributed optimization. *Annual Reviews in Control*, *47*, 278–305. https://doi.org/10.1016/j.arcontrol.2019.05.006

Yang, X., & He, X. (2013). Bat algorithm: Literature review and applications. *International Journal of Bio-Inspired Computation*, *5*(3), 141–149. https://doi.org/10.1504/IJBIC.2013.055093

Yang, Y., Xiao, Y., & Li, T. (2021). A survey of autonomous underwater vehicle formation: Performance, formation control, and communication capability. *IEEE Communications Surveys and Tutorials*, *23*(2), 815–841. https://doi.org/10.1109/COMST.2021.3059998

Yu, C., Chai, Y., & Liu, Y. (2018). Literature review on collective intelligence: A crowd science perspective. *International Journal of Crowd Science*, *2*(1), 64–73. https://doi.org/10.1108/IJCS-08-2017-0013

Yu, Z., Li, Q., Yang, F., & Guo, B. (2021). Human–machine computing. *CCF Transactions on Pervasive Computing and Interaction*, *3*(1), 1–12. https://doi.org/10.1007/s42486-020-00051-1

Zambonelli, F., Omicini, A., Anzengruber, B., Castelli, G., Angelis, F. L. D., Serugendo, G. D. M., Dobson, S. A., Fernandez-Marquez, J. L., Ferscha, A., Mamei, M., Mariani, S., Molesini, A., Montagna, S., Nieminen, J., Pianini, D., Risoldi, M., Rosi, A., Stevenson, G., Viroli, M., & Ye, J. (2015). Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive and Mobile Computing*, *17*, 236–252. https://doi.org/10.1016/j.pmcj.2014.12.002

Zedadra, O., Guerrieri, A., Jouandeau, N., Spezzano, G., Seridi, H., & Fortino, G. (2018). Swarm intelligence-based algorithms within IoT-based systems: A review. *Journal of Parallel and Distributed Computing*, *122*, 173–187. https://doi.org/10.1016/j.jpdc.2018.08.007

Zhang, H., Dong, Y., Chiclana, F., & Yu, S. (2019). Consensus efficiency in group decision making: A comprehensive comparative study and its optimal design. *European Journal of Operational Research*, *275*(2), 580–598. https://doi.org/10.1016/j.ejor.2018.11.052

Zhang, K., Yang, Z., & Basar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. In K. G. Vamvoudakis, Y. Wan, F. L. Lewis, & D. Cansever (Eds.), *Handbook of reinforcement learning and control* (Vol. 325, pp. 321–384). Springer. https://doi.org/10.1007/978-3-030-60990-0_12

Zhang, S., Lee, C. K. M., Chan, H. K., Choy, K. L., & Wu, Z. (2015). Swarm intelligence applied in green logistics: A literature review. *Engineering Applications of Artificial Intelligence*, *37*, 154–169. https://doi.org/10.1016/j.engappai.2014.09.007

Zhang, T., Lowmanstone, L., Wang, X., & Glassman, E. L. (2020). Interactive program synthesis by augmented examples. In S. T. Iqbal, K. E. MacLean, F. Chevalier, & S. Mueller (Eds.), *UIST '20: The 33rd annual ACM symposium on user interface software and technology, virtual event, USA, October 20–23, 2020* (pp. 627–648). ACM. https://doi.org/10.1145/3379337.3415900

Zhen, Y., Khan, A., Nazir, S., Zhao, H., Alharbi, A., & Khan, S. (2021). Crowdsourcing usage, task assignment methods, and crowdsourcing platforms: A systematic literature review. *Journal of Software: Evolution and Process*, *33*(8), e2368. https://doi.org/10.1002/smr.2368

Zheng, L., Yang, J., Cai, H., Zhou, M., Zhang, W., Wang, J., & Yu, Y. (2018). MAgent: A many agent reinforcement learning platform for artificial collective intelligence. *Proceedings of the AAAI Conference on Artificial Intelligence*, *32*(1). https://doi.org/10.1609/aaai.v32i1.11371