

Alma Mater Studiorum Università di Bologna  
Archivio istituzionale della ricerca

Artificial Collective Intelligence Engineering: A Survey of Concepts and Perspectives

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

*Published Version:*

Artificial Collective Intelligence Engineering: A Survey of Concepts and Perspectives / Casadei Roberto. -  
In: ARTIFICIAL LIFE. - ISSN 1064-5462. - ELETTRONICO. - 29:4(2023), pp. 433-467. [10.1162/artl\_a\_00408]

*Availability:*

This version is available at: <https://hdl.handle.net/11585/955371> since: 2024-02-15

*Published:*

DOI: [http://doi.org/10.1162/artl\\_a\\_00408](http://doi.org/10.1162/artl_a_00408)

*Terms of use:*

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).  
When citing, please refer to the published version.

(Article begins on next page)

# Artificial Collective Intelligence Engineering: a Survey of Concepts and Perspectives

Roberto Casadei

Alma Mater Studiorum–Università di Bologna  
robby.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 behaviour 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 techno-scientific trends like the Internet of Things, swarm robotics, and crowd computing, just to name 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 recognised 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 paper 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.

## 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 the information and communication technology (ICT) is promoting a future where computation is deeply integrated in 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 towards 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 at 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 endeavour 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 (Stein et al., 2021; Bellman 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 just at making individual devices smarter, but also at 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, e.g., multi-agent systems (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 has also motivated the emergence of a general research field specifically aimed at studying how to build CI in artificial systems, also known as under terms such as *artificial collective intelligence (ACI)* (Zheng et al., 2018; Tumer and Wolpert, 2004) 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 multi-agent systems (Tumer and 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 environment and possibly involving “humans in the loop” (Schirner et al., 2013), which are to be thought as “programmable platforms” for services and applications benefitting by the CI emerging from their activity. The idea is to provide a research map on 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 and 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 of existing methods.

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

The article is organised as follows. First, a set of broad scoping questions are elicited, to provide a structure for the paper and its discussions. After this, a survey of existing reviews relevant to CI is presented, to also motivate the perspective of this very article. Then, the preliminary concepts of a “collective” and “(individual) intelligence” are briefly reviewed. Upon this basis, to understand what CI is, some reference definitions, examples, models, and classifications are reviewed from the literature. 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.

## Method

Our goal is to scope the large and fragmented area of (artificial) collective intelligence, in order to identify its key concepts, relevant perspectives, research problems, and gaps—with an emphasis on its engineering and its computational/artificial intelligence (AI) side. Accordingly, we perform a scoping review (Petticrew and 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, 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 focussed 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 collective intelligence?

SQ3) What behaviours can be termed “collectively intelligent”?  
Are there paradigmatic examples?

SQ4) What are the requirements for collective intelligence?

SQ5) What relationships exist between individual and collective intelligence?

SQ6) How does collective intelligence unfold/emerge?

SQ7) How can collective intelligence be measured?

SQ8) How can collective intelligence be built artificially/computationally?

SQ9) What is the state of the art of (computational) collective intelligence?

SQ10) How is the research community on collective intelligence 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.

## Tertiary Study

In order to motivate the need for a survey on CI, we performed a tertiary study where secondary studies (e.g., surveys, systematic reviews) and collections are reviewed. We organise these according to whether they consider CI in its generality (i.e., abstracting from its applications and areas) 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.

**Reviews on CI as a general topic** Two main surveys to date aim at addressing CI as a general topic. He et al. (2019) analyse 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 consider, integrate, and somewhat subsume previous, more limited, or less general CI models and reviews (Yu et al., 2018; Lykourantzou et al., 2009; Salminen, 2012; Krause et al., 2010), and so they will be discussed more extensively in later sections. However, to the best of our knowledge, there are no comprehensive mapping studies providing a broad overview of the field for computer scientists.

**Multi-disciplinary collections** In (Malone and Bernstein, 2015), essays on CI are collected from different fields including economics, biology, human-computer interaction (HCI), AI, organisational behaviour, and sociology. In (Millhouse et al., 2021), a collection of contributions from a workshop gathering scientists in different areas is 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.*”

**Reviews on ACI/CCI** The paper by Tumer and Wolpert (2004), published in 1999, surveys CCI systems across the categories of (i) AI and machine learning (ML), including multi-agent systems (MASs); (ii) social science-inspired systems, such as those found in economics and game theory; (iii) evolutionary game-theoretical approaches; (iv) biologically-inspired systems, like swarm intelligence, artificial life, and population approaches; (v) physics-based systems; and (vi) other research fields ranging from network theory and self-organisation. This is a very rich survey but covers research published before the year 2000 and is slightly focused towards automatic and utility-based approaches.

The editorial by Jung (2017) reviews special issue papers on the integration of CCI and big data, where it is considered how data-driven CI can help in (i) collecting data, (ii) analysing data, and (iii) using data e.g. to support decision making.

The review by Rossi et al. (2018) provides a survey and taxonomy of multi-agent algorithms for collective behaviour, classified into: consensus, artificial potential functions, distributed feedback control, geometric algorithms, state machines and behaviour composition, bio-inspired

algorithms, density-based control, and optimisation algorithms. What emerges is a rather sharp distinction between low-level (e.g., bio-inspired self-organisation) and high-level coordination.

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

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

Some surveys consider swarm intelligence applied to specific problems such as self-organising pattern formation (Oh et al., 2017), feature selection (Nguyen et al., 2020), clustering (Figueiredo et al., 2019), green logistics (Zhang et al., 2015), collective movement (Navarro and Matía, 2013). Other surveys consider 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).

**Reviews on CI for specific systems and settings** Reviews from specific viewpoints include collections and surveys on human CI (Salminen, 2012), deep learning (Ha and 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 human context, grouping contributions into (i) micro level, emphasising enabling factors; (ii) emergence (or meso) level, emphasising how global patterns arise from local activity; and (iii) macro level, emphasising the kinds of system output. A review on human CI by a crowd science perspective is provided by Yu et al. (2018). Krause et al. (2010) review and compare swarm intelligence in animals and humans.

Ha and Tang (2021) performed a survey of recent developments on the embedding of CI principles into deep learn-

ing methods. They discuss e.g. how CI can help in devising novel architectures and training algorithms, and recent works on multi-agent (reinforcement) learning. Studies like this one are important since they elicit and strengthen trans-disciplinary relationships which are key for complex inter-disciplinary fields like CI.

Narayanan et al. (2022) provide a survey of the CI emerging in human-machine socio-technical systems supported by 5G communications. The discussed applications include road traffic control, unmanned aerial vehicles, smart grid management, and augmented democracy. The point is that to realise their full potential, these kinds of decentralised socio-technical systems often require proper connectivity properties and capabilities to support and foster the emergence of CI. For instance, from the analysis, the authors foresee that the 5G communication technology can promote CI by enhancing aspects like connectivity with neighbour nodes, interaction protocols, knowledge exchange, and the exploration-exploitation tradeoff via improved speed, latency, and reliability. On the other hand, there are significant challenges addressed by current and hopefully by future research in terms of security, privacy, and radio resource management.

### 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.

#### 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 characterisation 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 characterise collectives and collective phenomena (Brodaric and Neuhaus, 2020; Galton and Wood, 2016; Wood and Galton, 2009; Bottazzi et al., 2006).

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), *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

**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)
Organisation	Composite / Collective	Structure, roles
System	Composite / Collective	Interacting elements; Boundary
Multi-agent system	System	Nature (agency)

depth, it is worth noticing that, unlike the *componenthood* relation in composites, *membership* in collectives is generally not transitive (Brodaric and Neuhaus, 2020). Composites can be defined as structured pluralities or groups of parts, called *components*, playing specific functions (Brodaric and Neuhaus, 2020). In the literature, it is generally assumed that composites are heterogeneous, while collectives are homogeneous (Brodaric and Neuhaus, 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 and Galton, 2009). Defining a general, comprehensive, and precise characterisation or taxonomy of collectives is not trivial (Wood and 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 e.g. by perspective, the key relation between items, or the fundamental property of the group. Some of these group-like notions are summarised in Table 1, with a

proposed classification—following Brodaric and Neuhaus (2020), though different meronomies are possible. A collective is a particular kind of plurality or group. Crowds, swarms, herds, flocks, schools can generally be considered specific kinds of collectives. Organisations and systems might be modelled as constructs based on the structural arrangement and heterogeneity of composites, but are also amenable to be characterised as collectives.

Like for the notion of intentional stance (Dennett, 1989), it may make sense to adopt a *collective stance* in which e.g. “*the human species [a group] is viewed as a single organism*” (Gaines, 1994), though the idea of *collective intentionality* is problematic and subject of intense philosophical debate (Schweikard and 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 kind of collectives are addressed, as this would help to clarify the assumptions and generality of a specific contribution.

## Intelligence

Intelligence is a controversial and elusive concept subject to philosophical debate (Legg et al., 2007), best understood as a nomological network of constructs (Reeve et al., 2011). Etymologically, 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), or the property that “*measures an agent’s ability to achieve goals in a wide range of environments*” (Legg et al., 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 *crystallised* and *fluid* abilities, respectively.

Problems about intelligence include, for instance, its definition and modelling, 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 (Hernández-Orallo, 2017; Chollet, 2019) is of course related to what representation or model of intelligence is considered, and is complicated by the need of distinguishing between causality and correlation, selecting a representative set of environments for evaluation, etc. Carrol defines an *ability* (i.e. an

intelligence factor) as a source of variance in performance for a certain class of tasks (Carroll, 1993). Measuring intelligence is based on *factor analysis*, i.e., 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 and Norvig, 2020). Both can be considered under the unifying notion of *abstract intelligence* (Wang, 2009).

## Understanding Collective Intelligence

On the basis of the preliminary concepts introduced in the previous section, this section focusses on *what* CI is, according to literature, discussing definitions, examples, models, and the main classifications of CI (namely ACI and CCI) we are interested into, 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.

### Definitions and characterisations of collective intelligence

Collective intelligence is the intelligence that can be ascribed to a collective—where a collective is a multiplicity of entities (commonly characterised 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 characterisations 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 characterisations.

**Reuse of (individual) intelligence definitions.** Some definitions do not attempt to re-define “intelligence” but merely bring existing characterisations of intelligence, common-sense 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 trans-disciplinarity.

**General vs. 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 of a general CI statistical factor *c*

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

**Collectives of different natures.** Some definitions largely abstract from the nature of collectives (cf. “collections” or “groups of individuals”, “artificial and/or natural”), some assume a minimal set of characteristics for individuals (cf. agency, ability to interact, etc.), some require that the individuals are connected in some way (cf. interaction, or existence of social structures).

**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.

**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 and Holvoet, 2004).

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

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

## Examples

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

**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 collective intelligence 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 and Bernstein, 2015). Accordingly, market-based abstractions have been considered in computer science to promote globally efficient systems (Mainland et al., 2004).

**Example 2 (Wisdom of crowds).** Crowds – groups of people – can be of different kinds (cf. 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) popularised the term “wisdom of crowds”, showing that groups are able 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 *decentralisation* (of individual knowledge acquisition)—whose importance has been confirmed by later studies such as, e.g., those by Woolley et al. (2010).

**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 sub-field that deals with very large groups and individuals behaving according to simple rules. Since the criteria of cardinality and simplicity are degrees, the boundaries of the field is also fuzzy.

**Example 4 (Learning multi-agent systems).** Another notable example of CI is given by MASs (Wooldridge, 2009). Unlike swarms, MASs usually comprise rational agents, possibly structured into organisations, and possibly exhibiting properties of strong agency (Wooldridge, 2009), i.e., 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 behaviour that they should follow to maximise some local or global notion of utility (Tumer and Wolpert, 2004).

**Example 5 (Human-machine collective intelligence).** A powerful example of CI is the so-called *human-machine collective intelligence (HMCI)* (Smirnov and Ponomarev, 2019) or *hybrid CI* (Peeters et al., 2021; Moradi et al., 2019), 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, which 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.

## Models

Here, we briefly review two main general models of CI from the literature, which comprehensively summarise and integrate previous models.

**Table 2**

Some definitions of CI from the literature.

Ref.	Definition	Remarks
Malone and Bernstein (2015)	<i>Groups of individuals acting collectively in ways that seem intelligent</i>	<ul style="list-style-type: none"> <li>• “Reuse” of the notion of intelligence</li> <li>• Collective action</li> </ul>
Nguyen et al. (2009)	<i>The form of intelligence that emerges from the collaboration and competition of many individuals (artificial and/or natural)</i>	<ul style="list-style-type: none"> <li>• Emergence</li> <li>• Mechanisms (collaboration, competition)</li> <li>• Members of different nature</li> </ul>
He et al. (2019)	<i>Collective intelligence (CI) refers to the intelligence that emerges at the macro-level of a collection and transcends that of the individuals.</i>	<ul style="list-style-type: none"> <li>• Emergence (transcendence)</li> <li>• Levels (macro, micro)</li> </ul>
Tumer and Wolpert (2004)	<i>[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.</i>	<ul style="list-style-type: none"> <li>• Requirements (interaction, decentralisation)</li> <li>• Embedded metric</li> </ul>
Szuba (2001)	<i>We can say that the phenomenon of CI has emerged in a social structure of interacting agents or beings, over a certain period, iff 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</i>	<ul style="list-style-type: none"> <li>• Requirements (social structure, interaction)</li> <li>• Embedded metric</li> <li>• Dynamic property</li> <li>• Negative and positive CI</li> </ul>
Lykourantzou et al. (2009)	<i>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.</i>	<ul style="list-style-type: none"> <li>• Hybrid or human-machine CI</li> </ul>

**The isolation-, collaboration-, and feedback-based CI paradigms (He et al., 2019)** He et al. (2019) propose a taxonomy of CI 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 behaviour results. Then:

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 behaviours. Isolation studies use statistical and mining tools.
2. *Collaboration paradigm.* There is direct or indirect *interaction* between the individuals. Indirect interaction can be modelled through a notion of *environment*. Aggregation operates on individual behaviour results and the environment state. The aggregation result does affect neither the individual behaviours nor the environment.
3. *Feedback paradigm.* This paradigm adds to the interaction paradigm a “downward causation” of the aggrega-

tion result on the individual behaviours and/or the environment.

**CI framework by Suran et al. (2020)** Suran et al. (2020) analyse 12 studies on CI and devise a *generic* model based on 24 CI attributes split into 3 CI components: individuals, coordination/collaboration activities, and communication means. The generic model is based on:

- Characterisation 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.
- Characterisation of *motivation* of CI actors: intrinsic or extrinsic.
- Characterisation of CI goals: individual and community objectives.
- Characterisation of CI processes: in terms of types of ac-



tivities (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 e.g. on state, data, aggregation, decentralisation, task allocation, and robustness.

### Factors and quantification of CI

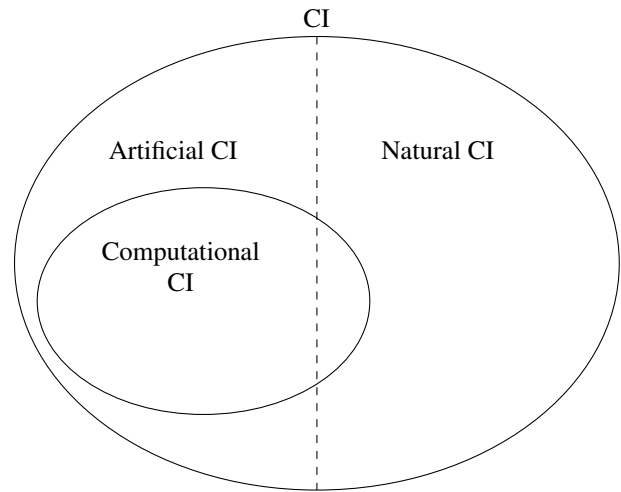
Key scientific questions, fundamental for both understanding and engineering CI, include what *factors* promote or inhibit CI, and, specifically, what is the relationship between individual and collective intelligence. We already mentioned the seminal work by Woolley et al. (2010) sustaining the idea 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 characterised 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 behaviour.

Works have been carried out to investigate these relationships. For instance, in a later study, Woolley et al. (2015) focus on (i) group composition, e.g., in terms of skills and diversity of the members of a group; and (ii) group interaction, e.g., 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) study the influence of the following factors: (i) concerning individuals: individual intelligence, individual reasoning/learning speed; (ii) concerning cooperation: cardinality of the collective, time to interact, communication protocol; and (iii) 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 and 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 non-trivial ways. These factors are also related to the components of CI models—a nice overview is provided in (Suran et al., 2020).

### Main kinds of CI

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

**Natural CI** Natural CI is the CI exhibited by collectives found in nature, such as swarms of insects, packs, herds, or



**Figure 1**

*The relationships between collective intelligence (CI), artificial collective intelligence (ACI), and computational collective intelligence (CCI). The dashed line is used to denote the false dichotomy between artificial and natural CI.*

groups of animals, crowds of people, flocks of birds, schools of fishes, etc. In all these systems there exist non-trivial collective phenomena and societal aspects that deserve deep investigation. Insect societies are analysed e.g. in the seminal book by (Bonabeau et al., 1999). For collective animal behaviour, one of the main references is the book by Sumpter (2010), which describe 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, 2002; Surowiecki, 2005); moreover, there are contributions for specific human settings like crowds of pedestrians (Sieben et al., 2017) and problems like e.g. the relationship between language and collective action (Smith, 2010).

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

**Artificial (ACI) and Computational Collective Intelligence (CCI)** ACI is the CI exhibited by human-made machines. Notice that, strictly speaking, natural and artificial CI constitute a false dichotomy since 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, i.e., “*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). *Soft computing* methods are those that help to address complex problems by overcoming approximation and uncertainty, using tech-

niques such as fuzzy logic, expert systems, machine learning, genetic algorithms, artificial neural networks (Ibrahim, 2016). In other words, the 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 is 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, since also mechanical machines can exhibit intelligence (Wang, 2009; Stradner et al., 2013)—cf. Braitenberg vehicles (Dvoretzskii et al., 2022), some of which are purely mechanical vehicles with hard-wired connections between sensors and actuators. Common sub-fields of ACI include e.g. semantic web, social networks, and multi-agent systems (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 or multi-agent intelligence may refer to natural, to artificial systems, or to a general model comprising both.

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

## Perspectives of Artificial Collective Intelligence Engineering

Building on the previous discussion of *what* CI is and its main models, this section focusses on *how CI can be engineered*, according to 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 (cf. the previous section), various *perspectives* and approaches to *CI engineering* can be devised, each one leveraging and providing peculiar sets of *CI mechanisms*.

### Knowledge-oriented vs. behaviour-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 (Segaran, 2007; Alag, 2008). There are several ways in which humans using web applications can provide data through their interaction: e.g., by what content they search, what paths they follow, what feedback they provide, what content they add, etc. Then, techniques like data mining, text mining, and machine learning can be used to classify information, cluster information, predict trends, recommend content, filter information, aggregate

information etc. We may call this *knowledge-oriented* CI since the collective intelligence 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 systems may not seem 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 is 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, socio-technical 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, etc.)—cf. the isolation, collaboration, and feedback-based paradigms.

Conversely, we may call *behaviour-oriented* CI the collective intelligence that drives the global behaviour of a system. This includes what Surowiecki (2005) refers to as *coordination* and *cooperation* problems. Examples include the form of intelligence driving the way in which robotic swarms move (Navarro and Matía, 2013), a computational ecosystem that self-organises into activity and communication structures (Pianini et al., 2021), and a market that self-regulates itself (Lo, 2015). However, as the latter example shows, since collective action is connected with collective decision making, which in turn is connected to collective knowledge, the border between knowledge-oriented and behaviour-oriented CI is fuzzy and so these types of CI should not be thought 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), e.g., based on the properties of autonomy, belonging, connectivity, diversity, and emergence (Boardman and Sauser, 2006), is also relevant in this discussion (Peeters et al., 2021). Research under the knowledge-oriented CI umbrella, typically involving socio-technical 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)).

### Manual vs. 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*.

**Manual design of ACI** In the manual approach, a designer specifies the behaviour of the computational agents making up the collective directly by providing *behavioural rules* (or *policies*).

Here, the key issue is determining what individual behaviour, when replicated or combined with other behaviours or phenomena, can give rise to the desired emergent behaviour. 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 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, e.g., for expressing the behaviour of robot swarms (Pinciroli and Beltrame, 2016) and distributed, IoT systems (Noor et al., 2019; Mizzi et al., 2018)—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 behaviour 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 (i) lifting standard values to field values; (ii) abstracting field computations through functions; (iii) stateful evolution of fields; and (iv) handling bi-directional communication through so-called *neighbouring fields*. Using such constructs, and library functions e.g. handling information flows through gradients or supporting higher-level patterns, a programmer can write an aggregate program that expresses the global behaviour of a possibly dynamic network of agents. The agents, by repeatedly evaluating the program in asynchronous sense-compute-interact rounds, and interacting with neighbours by exchanging data as dictated by the program, could steer self-organising behaviour hopefully fulfilling the intent of the program. Casadei et al. (2021) argue that multiple concurrent and dynamic aggregate computations could pave a path to CI engineering.

While AC adopts a swarm-like self-organisation model, another class of approaches for ACI is given by *multi-agent programming*, as supported e.g. by the JaCaMo platform (Boissier et al., 2020b), which comprises *Jason* for programming cognitive autonomous agents, *CARtAgO* for programming the distributed artifacts-based environment of the MAS, and *MOise* for programming agent organisations. 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 focusses on composites rather than collectives, i.e., well-structured organisations of heterogeneous intelligent agents rather than self-organising swarms of largely homogeneous and cognitively simple agents.

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

One of the early models and notable example is *Collective Intelligence*, or *COIN* (Wolpert, 2003). Essentially, COIN considers a *collective* as a system of self-interested agents, trying to maximise their *private utility* function, sharing an associated *world utility* 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 (Zhang et al., 2019b; Canese et al., 2021; Gronauer and Diepold, 2022). Learning of collective behaviour 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 behaviour 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 optimise models of robot controllers (e.g. mapping inputs from sensors to outputs to actuators) with respect to desired behavioural goals. Various techniques have been proposed in the literature to improve traditional evolutionary approaches, e.g., 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 behavioural 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 (Watson et al., 2002; Bredèche et al., 2018), 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, co-evolution, the role of environmental niches, the relationship between optimisation and selection pressure, locality of interaction, etc. 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. Bastani et al. (2020); Aguzzi et al. (2022). However, in the context of CI, this direction has not yet been investigated, representing an opportunity for future research (cf. 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).

## 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 behaviour-oriented CI. Other classifications can be made by considering what kind of entity plays the role of *controller* and *executor*:

1. tasking crowds of humans (Ganti et al., 2011; Guo et al., 2014; Zhen et al., 2021; Sari et al., 2019)—cf. crowdsourcing (Zhen et al., 2021) and crowdsensing (Guo et al., 2014);
2. using humans to guide machine operations, e.g., interactively (Yu et al., 2021);

or considering what entity plays the role of *input* and *output*

1. using AI to extract or mine intelligence from human contributions (Segaran, 2007; Alag, 2008);
2. 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 and Bederson, 2011).

or, finally, considering humans and machines as peers and hence the so-called *human-machine networks* (Tsvetkova et al., 2017) or *social machines* (Burégio et al., 2013; Berners-Lee, 1999).

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) propose three principles: (i) goals from the collective, technological, and human perspectives should be considered simultaneously; (ii) development effort should continuously embrace all the product’s lifecycle; and (iii) the requirements of observability, predictability, explainability, and directability should be considered at all abstractions levels (AI, team, and society).

## 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 (cf. 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) classify collective behaviours (of swarm robotics systems) into (i) spatially-organising behaviours, (ii) navigation behaviours, (iii) collective decision making, and (iv) others. Other reviews of swarm robotics tasks include (Bayindir, 2016; Nedjah and Junior, 2019). Moreover, collective tasks can be classified also according the three paradigms discussed in (He et al., 2019) and reviewed in previous sections: isolation, collaboration, and feedback.

In the following, 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.

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

Decision making and its collective counterpart can be classified according to the nature of the decision to be made. Reaching consensus and multi-agent task allocation are two common kinds of collective decision-making behaviours, typical in swarm robotics (Brambilla et al., 2013). Guttmann (2009) classifies MAS decision making by four dimensions: (i) *use of models of self vs. models of others*; (ii) *individual inputs vs. group input*; (iii) *learning vs. non-learning*, depending on whether decision making spans multiple rounds or just one round; and (iv) *collaboration vs. competition*. Surowiecki (2005) distinguishes three kinds of problems or

tasks of distributed decision making: (i) *cognition*, (ii) *co-operation*, and (iii) *coordination*.

Collective decision making is often supported by self-organisation mechanisms based on, e.g., collective perception (Schmickl et al., 2006), voter models (Valentini et al., 2014), opinion formation models (de Oca et al., 2011), and self-stabilising leader election (Pianini et al., 2022a).

Recent surveys on collective decision making include the following. Valentini et al. (2017) focus on discrete consensus achievement, and propose a formal definition of the *best-of- $n$*  problem (choice of the best alternative among  $n$  available options); then, they define a taxonomy based on different classes of the problem, and classify the literature on discrete consensus agreement accordingly. Zhang et al. (2019a) provide a review of consensus models in collective decision making, and compare them based on multiple criteria for measuring consensus efficiency. They also argue that two interesting research directions include (i) *large-scale* collective decision making and (ii) addressing social relationships and opinion evolution. Tang and Liao (2021) provide a review of literature around five challenges in large-scale collective decision making with big data: dimension reduction, weighting and aggregation of decision information, behaviour management, cost management, and knowledge distribution and increase. Rizk et al. (2018) provide a survey of decision making in MASs. The survey focusses 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. Particularly challenging is also decision making in dynamic environments (Rizk et al., 2018; Prasetyo et al., 2019). Other challenges include security, privacy, and trust; approaches to address these include, e.g., blockchain consensus (Pournaras, 2020b).

**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 both in areas like sociology and organisational theory (Garavan and Carbery, 2012; Fadul, 2009), and in AI research (Bock, 1993). Collective learning spans both the knowledge-oriented and behaviour-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 and Carbery, 2012). As a wide concept, collective learning can be interpreted along multiple perspectives (Garavan and Carbery, 2012): e.g., as the independent aggregation of individual learning, or as a collaborative activity. So, collective

learning is related 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 machine learning (Verbraeken et al., 2020): examples include centralised, federated, and decentralised machine learning systems. In *centralised 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 includes data generation. In *federated learning* (Kairouz et al., 2021), the idea is that individual independent workers perform machine learning tasks on local data sets, producing models that are then aggregated by a master into a global model without the need of sharing data samples. It enables to address data privacy issues. The combination of multiple models is also called *ensemble learning* (Dong et al., 2020). Hegedüs et al. (2021) propose *gossip-based learning* as a decentralised 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 decentralised 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*, and *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 (Zhang et al., 2019b): (i) *centralised structures*, involving a central controller aggregating information from the agents; (ii) *decentralised structures*, with no central entities and neighbourhood interaction; and (iii) *fully decentralised*, namely independent learning, with no information exchanged between the agents. Various formal frameworks have been proposed to address MARL problems, including *COIN* (Wolpert, 2003) and *Decentralised Markov Decision Processes (Dec-MDP)* (Oliehoek and Amato, 2016). The reader interested to MARL algorithms and frameworks can check out multiple comprehensive surveys on the topic (Busoniu et al., 2008; Zhang et al., 2019b; Hernandez-Leal et al., 2019).

There exist surveys on collective learning. D’Angelo et al. (2019) perform a systematic literature review on learning-based collective self-adaptive systems. Their analysis extracts, as the main characteristics of such systems, the application domains involving groups of agents with the abil-

ity 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 behaviours involved (e.g., selfish vs. collaborative). Accordingly, the authors provide a framework for learning collective self-adaptive systems, based on three dimensions: autonomy, knowledge access, and behaviour. The learning goals are analysed w.r.t. the target emergent behaviour; from the analysis, two clusters of works emerge: those where the emergent behaviour is associated to the anticipated learning task, and those where it is not. Among the learning techniques, they report 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) provides a review of 10-years research on human-centred collective learning for coordinated multi-objective decision making in socio-technical systems, within the context of the *Economic Planning and Optimized Selections (EPOS)* project. Collective learning is motivated as a way to address the long-standing *tragedy of the commons* problem, and argued to be a promising paradigm of artificial intelligence. As research opportunities and challenges, the author identifies: explainability and trust, resilience to plan violations and adversaries, collective learning in organic computing systems, co-evolution of collective human and machine learning, 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 (Mataric, 2007; Anderson et al., 2013). However, these techniques can also be combined (Nolfi and Floreano, 1999): learning can guide evolution (Hinton and Nowlan, 1987) and evolution can improve learning (cf. evolutionary learning—Telikani et al. (2022)), where different architectures for the combination are possible (Sigaud, 2022).

**Other collective tasks** *Collective action* (Oliver, 1993) commonly refers to the situation where 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 etc.), such as the *SOSIEL (Self-Organising Social & Inductive Evolutionary Learning)* simulation platform (Sotnik, 2018), for studying the problem and investigating solutions for human societies as well as for socio-technical and artificial systems. Collective actions may be supported by collective and self-organised decision-making processes, and leveraging abstractions like *electronic institutions* and *social capital* (Petruzzi et al., 2015).

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

*Distributed optimisation* (Yang et al., 2019) refers to the problem of minimising a global objective function, which is the sum of the local objective functions of the members of a collective, in a distributed manner. Distributed optimisation 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 such as collaboration (Hmelo-Silver, 2003), socio-technical infrastructures (Gruber, 2008), knowledge transfer (Huang and 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).

## A view of CI-related fields

Being CI a multi-disciplinary 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 (Casadei, 2020; Nicola et al., 2020; Bucchiarone et al., 2020; Ferscha, 2015). The idea is that CI engineering should be supported through inter-disciplinary research and a systems science perspective (Mobus et al., 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 such as, for instance, CPS, namely systems that combine discrete and continuous dynamics (Alur, 2015). Then, a set of inter-related fields can promote the study of peculiar CI phenomena such as emergence, self-organisation, ensemble formation, etc. Such fields include but are not limited to the field of coordination (Malone and Crowston, 1994), multi-agent systems (Wooldridge, 2009), autonomic/self-\* computing (Kephart and Chess, 2003), collective adaptive systems (Casadei, 2020; Nicola et al., 2020; Bucchiarone et al., 2020; Ferscha, 2015), 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.

We noticed multiple times in previous sections how interaction is a key element of CI. *Coordination* is the interdisciplinary study of interaction (Malone and Crowston, 1994). In computer science, interaction was early recognised as a concern related but clearly distinguished from computation (Gelernter and Carriero, 1992), hence amenable to separate modelling 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 behaviour 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 and Hankin, 1996; ter Beek and Sirjani, 2022).

*Collective adaptive systems (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 (Nicola et al., 2020; Bucchiarone et al., 2020). CASs are sometimes considered to be heterogeneous (Loreti and Hillston, 2016; Andrikopoulos et al., 2013), 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 characterised by (i) *large numbers* of (ii) *relatively simple (or not particularly intelligent) individuals* (Beni, 2004). Collectives are generally *quite homogeneous, at least at some level of abstraction* (Pianini et al., 2022b), though research works aim to address heterogeneous collective adaptive systems (Scekic et al., 2020) as well as heterogeneous swarms (Dorigo et al., 2013; Kengyel et al., 2015), e.g., with systems involving humans and robots (Hasbach and Bennewitz, 2022), or groups of robots with different morphology or behaviour. Swarm robotics is the combination of swarm intelligence and robotics (Brambilla et al., 2013; Beni, 2004).

Coordination, CASs, and swarm robotics can also be seen as sub-fields of the larger field of MASs (Wooldridge, 2009; Faliszewski et al., 2022), which itself stemmed from the field of distributed artificial intelligence (Ferber, 1999). In MASs 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 and Graesser, 1996; van der Hoek and Wooldridge, 2003), from which other properties like proactiveness, interactivity, and sociality stem. By a software programming and engineering point of view, agents can be considered as an abstraction following active objects and actors (Odell, 2002) that, together other first-class abstractions

like artifacts (Omicini et al., 2008), environments (Weyns and Michel, 2014), and organisations (Horling and Lesser, 2004), provide a support for the so-called *(multi-)agent-oriented programming* paradigm (Shoham, 1993; Boissier et al., 2020a). The MAS field/perspective is clearly intimately related to CI.

Like for MASs, the key property of autonomy is at the centre of *autonomic computing* (Kephart and Chess, 2003), namely the field devoted to the construction of 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 and Tahvildari, 2009), self-healing/repairing (Psaier and Dustdar, 2011), self-improving/optimising (Bellman et al., 2018), self-organising (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 and Chess, 2003). In so-called *architecture-based self-adaptation* (Garlan et al., 2004), architectural models of the managed systems are leveraged at runtime to organise the self-managing logic. The managing system could also be distributed and decentralised (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-organising*, characterised by processes that autonomously and resiliently increase/maintain order or structure (Wolf and Holvoet, 2004); it typically emerges from the interaction of several entities. Self-organisation can be considered as a key promoter or element of CI (Rodríguez et al., 2007).

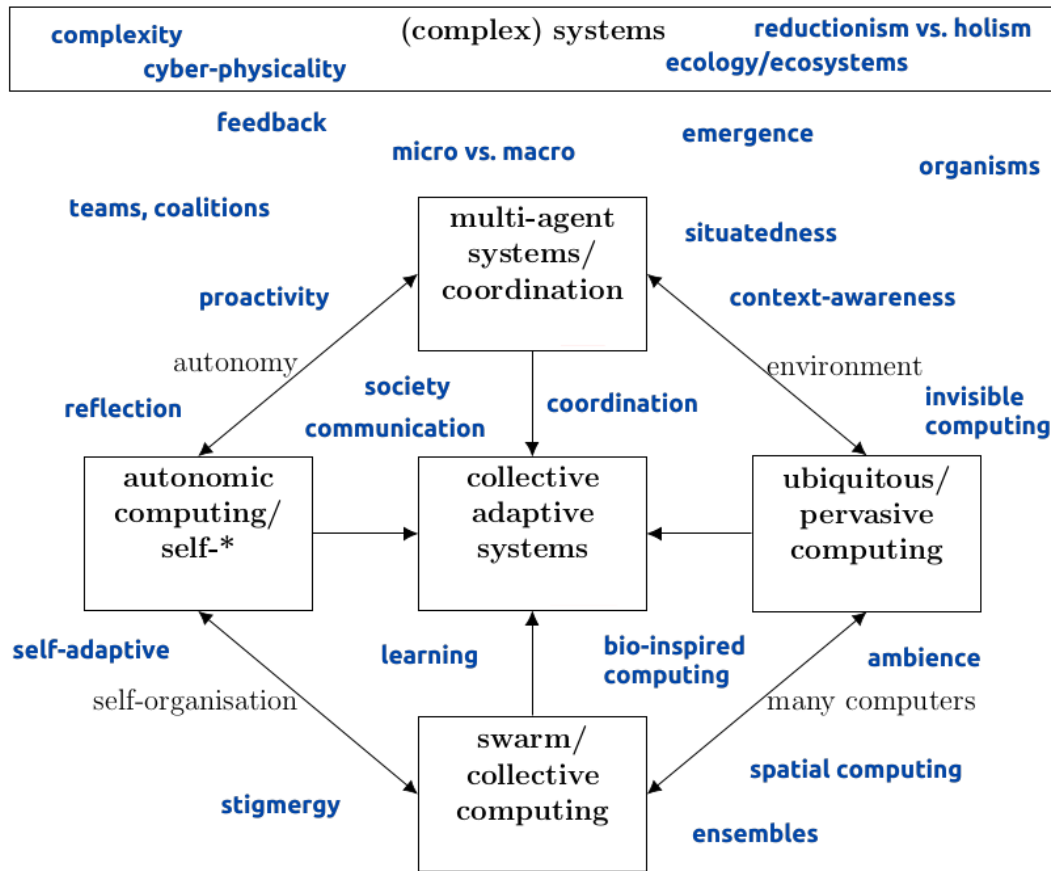
As a last remark, we stress that the aforementioned fields are highly inter- and trans-disciplinary. For instance, MASs can be considered by economical, sociological, organisational, and computational perspectives (Wooldridge, 2009). Same goes for coordination (Malone and Crowston, 1994). Moreover, a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

## Research Opportunities and Challenges

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

### Programming emergence and macro-programming

The problem of programming emergent and self-organising behaviours is an open research challenge (Gershenson et al.,



**Figure 2**

A research map of fields and concepts contributing to (research on) CI engineering.

2020; Varenne et al., 2015) intimately related to CI engineering. Term *macro-programming* emerged in early 2000s to identify programming approaches with the goal of defining the global behaviour of WSNs (Newton et al., 2005); currently, it generally denotes paradigms aiming at supporting the programming of system- or macro-level properties and behaviours. A recent survey by Casadei (2023) shows that, beside the first wave of research in the context of WSNs, we are witnessing a renewed interest in macro-programming fuelled by scenarios like the Internet of Things, robot swarms, and collective adaptive systems 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 behavioural rules of the individuals can promote the desired collective behaviour. In particular, we can distinguish two problems (Tumer and Wolpert, 2004). Given a set of individuals and the corresponding local behavioural 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 lo-

cal behaviours 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 behavioural rules (micro-programs) (Casadei, 2023).

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

### 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 multi-agent 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 synthesise macro-programs expressed in a macro-programming language (Casadei, 2023). This could be coupled with simulation to verify how systems executing synthesised programs operate in various environments. On one hand, since simulations may also be computationally-intensive, it might be necessary to limit simulation to 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, there exist some recent attempts at combining program synthesis and reinforcement learning (Verma et al., 2018; Bastani et al., 2020; Qiu and Zhu, 2022). For instance, Bastani et al. (2020) discuss 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 multi-agent oriented or a macro-programming language. An early attempt has been carried out e.g. in (Aguzzi et al., 2022), where 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 a collective adaptive behaviour that improves over a simple, manually encoded collective behaviour.

### 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 behaviour. Indeed, we are considering *feedback* CI paradigm (He et al., 2019), where the aggregation of contributions from the individuals and the environment in turn affects the individuals and the environment. This is also what Lykourantzou et al. (2009) call *active* CI systems, where collective behaviour is supported by the system level, which are contrasted from *passive* CI systems where 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 and Chess, 2003), with its *MAPE-K* (*Monitor-Analyse-Plan-Execute with Knowledge*) loops, and *organic computing* (Müller-Schloer and 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 and Sick, 2006). Attempts to defining 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 e.g. 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, on the other way, a deeper understanding on 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 e.g. formal studies carried out on *self-stabilisation* of aggregate computations (Pianini et al., 2022a), which guarantees that stable outputs are eventually achieved from stable inputs.

### Integrating humanity and technology: social machines

A key subfield of CI that is still at its early days is human-machine collective intelligence (HMCI) (Smirnov and Ponomarev, 2019), also known as *hybrid CI* (Peeters et al., 2021; Moradi et al., 2019), or *hybrid CASs* (Scekic et al., 2020). In the systems we consider in this article, we can identify two main domains (Beal et al., 2013): (i) the domain of *space-time*, which corresponds to physical environments and their evolution; and (ii) the domain of *information*, which evolves through computation. Of course, these two domains interact, e.g., 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 realisation 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 and Bederson, 2011)), and (computing) machines can operate in the physical world (cf. the notion of *cyber-physical system* (Alur, 2015)). Indeed, various terms or buzzwords are emerging to denote systems where such integration of humans, computation, and physical systems is present—cf., human CPSs (Liu and 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 characterisation of evolution of computing from mainframe computing to personal computing to ubiquitous computing (Weiser, 1991).

In order 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 organisational aspects (Bucchiarone et al., 2020). In other words, a key challenge and opportunity revolves around the design of social machines (Burégio et al., 2013; Berners-Lee, 1999), hybrid societies (Hamann et al., 2016), and socio-technical systems (Baxter and Sommerville, 2011). A social machine can be described as “*a computational entity that blends computational and social processes*” and that is at the intersection of social software, people as computational units, and software as sociable entities (Burégio et al., 2013; Berners-Lee, 1999). In this respect, elements whose formalisation 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 and Crowston, 1994). However, several challenges remain, related to proper modelling of human computation, achieving effective communication and coordination between humans and machines, achieving self-improving system integration (Bellman et al., 2021).

### 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 summarise, we recommend the following topics to be further investigated:

- 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);
- approaches and mechanisms for controlling or steering emergence and self-organisation (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));
- the role of CI across the various level 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, e.g., security, resilience, and resource efficiency;
- 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 (Bastani et al., 2020; Aguzzi et al., 2022) of macro-programs;
- integration of human intelligence with machine intelligence into hybrid, collectively intelligent systems (Smirnov and Ponomarev, 2019; Peeters et al., 2021), e.g., leveraging wearable computing (Ferscha

et al., 2014), ways for combining methods for human teamwork with AI, and self-organisation protocols considering both humans and artificial agents (Smirnov and Ponomarev, 2019; Scekic et al., 2020).

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

## Conclusion

Collective intelligence (CI) is a rich theme that builds on multi-, inter-, and trans-disciplinary collective endeavours. However, research is largely fragmented across several specific research problems (cf. types of collective tasks), research methods (cf. manual vs. automatic CI design), and even entire computer science research areas (cf. 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 and even a sense of CI-related work in their sub-field. This scoping review aimed at providing a comprehensive view on CI for computer scientists and engineers, with emphasis on concepts and perspectives, and also providing some research highlights on the forms of CI that most interest them, namely artificial collective intelligence (ACI), computational collective intelligence (CCI), and human-machine collective intelligence (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.
- Aguzzi, G., Casadei, R., and Viroli, M. (2022). Towards reinforcement learning-based aggregate computing. In ter Beek, M. H. and Sirjani, M., editors, *Coordination Models and Languages - 24th 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*, Lecture Notes in Computer Science. Accepted for publication.
- 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., and van Steen, M. (2013). Adaptive collective systems: herding black sheep.

- Andrikopoulos, V., Bucchiarone, A., Saez, S. G., Karastoyanova, D., and Mezzina, C. A. (2013). Towards modeling and execution of collective adaptive systems. In Lomuscio, A., Nepal, S., Patrizi, F., Benatallah, B., and Brandic, I., editors, *Service-Oriented Computing - ICSOC 2013 Workshops - CCSA, CSB, PASCEB, SWESE, WESOA, and PhD Symposium, Berlin, Germany, December 2-5, 2013. Revised Selected Papers*, volume 8377 of *Lecture Notes in Computer Science*, pages 69–81. Springer.
- Audrito, G., Casadei, R., Damiani, F., Salvaneschi, G., and Viroli, M. (2022). Functional programming for distributed systems with XC. In Ali, K. and Vitek, J., editors, *36th European Conference on Object-Oriented Programming, ECOOP 2022, June 6-10, 2022, Berlin, Germany*, volume 222 of *LIPIcs*, pages 20:1–20:28. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Badica, A., Trawinski, B., and Nguyen, N. T., editors (2014). *Recent Developments in Computational Collective Intelligence*, volume 513 of *Studies in Computational Intelligence*. Springer.
- Bastani, O., Inala, J. P., and Solar-Lezama, A. (2020). Interpretable, verifiable, and robust reinforcement learning via program synthesis. In Holzinger, A., Goebel, R., Fong, R., Moon, T., Müller, K., and Samek, W., editors, *xxAI - Beyond Explainable AI - International Workshop, Held in Conjunction with ICML 2020, July 18, 2020, Vienna, Austria, Revised and Extended Papers*, volume 13200 of *Lecture Notes in Computer Science*, pages 207–228. Springer.
- Baxter, G. D. and Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interact. Comput.*, 23(1):4–17.
- Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing*, 172:292–321.
- Beal, J., Dulman, S., Usbeck, K., Viroli, M., and Correll, N. (2013). Organizing the aggregate: Languages for spatial computing. In *Formal and Practical Aspects of Domain-Specific Languages: Recent Developments*, pages 436–501. IGI Global.
- Beaver, L. E. and Malikopoulos, A. A. (2021). An overview on optimal flocking. *Annu. Rev. Control.*, 51:88–99.
- Behjat, A., Manjunatha, H., Kumar, P. K., Jani, A., Collins, L., Ghassemi, P., Distefano, J. P., Doermann, D. S., Dantu, K., Esfahani, E. T., and Chowdhury, S. (2021). Learning robot swarm tactics over complex adversarial environments. In *International Symposium on Multi-Robot and Multi-Agent Systems, MRS 2021, Cambridge, United Kingdom, November 4-5, 2021*, pages 83–91. IEEE.
- Bellman, K. L., Botev, J., Diaconescu, A., Esterle, L., Gruhl, C., Landauer, C., Lewis, P. R., Nelson, P. R., Pournaras, E., Stein, A., and Tomforde, S. (2021). Self-improving system integration: Mastering continuous change. *Future Gener. Comput. Syst.*, 117:29–46.
- Bellman, K. L., Botev, J., Diaconescu, A., Esterle, L., Gruhl, C., Landauer, C., Lewis, P. R., Stein, A., Tomforde, S., and 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*, pages 160–167. IEEE.
- Bellman, K. L., Botev, J., Hildmann, H., Lewis, P. R., Marsh, S., Pitt, J., Scholtes, I., and Tomforde, S. (2017). Socially-sensitive systems design: Exploring social potential. *IEEE Technol. Soc. Mag.*, 36(3):72–80.
- Beni, G. (2004). From swarm intelligence to swarm robotics. In Sahin, E. and Spears, W. M., editors, *Swarm Robotics, SAB 2004 International Workshop, Santa Monica, CA, USA, July 17, 2004, Revised Selected Papers*, volume 3342 of *Lecture Notes in Computer Science*, pages 1–9. Springer.
- Berners-Lee, T. (1999). *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*. Harper San Francisco.
- Boardman, J. T. and 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*, pages 1–6. IEEE Computer Society.
- Bock, P. (1993). *The Emergence of Artificial Cognition: An Introduction to Collective Learning*. The Emergence of Artificial Cognition: An Introduction to Collective Learning. World Scientific.
- Boissier, O., Bordini, R., Hubner, J., and Ricci, A. (2020a). *Multi-Agent Oriented Programming: Programming Multi-Agent Systems Using JaCaMo*. Intelligent Robotics and Autonomous Agents series. MIT Press.
- Boissier, O., Bordini, R. H., Hubner, J., and Ricci, A. (2020b). *Multi-agent oriented programming: programming multi-agent systems using JaCaMo*. MIT Press.
- Bonabeau, E., Marco, D. d. R. D. F., Dorigo, M., Théraulaz, G., Theraulaz, G., et al. (1999). *Swarm intelligence: from natural to artificial systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford university press.
- Bose, T., Reina, A., and Marshall, J. A. (2017). Collective decision-making. *Current Opinion in Behavioral Sciences*, 16:30–34.
- Bottazzi, E., Catenacci, C., Gangemi, A., and Lehmann, J. (2006). From collective intentionality to intentional collectives: An ontological perspective. *Cognitive Systems Research*, 7(2-3):192–208.
- Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.*, 7(1):1–41.
- Bredèche, N., Haasdijk, E., and Prieto, A. (2018). Embodied evolution in collective robotics: A review. *Frontiers Robotics AI*, 5:12.
- Brodaric, B. and Neuhaus, F. (2020). Pluralities, collectives, and composites. In *Formal Ontology in Information Systems: Proceedings of the 11th International Conference (FOIS 2020)*, volume 330, page 186. IOS Press.
- Bucchiarone, A., D’Angelo, M., Pianini, D., Cabri, G., De Sanctis, M., Viroli, M., Casadei, R., and Dobson, S. (2020). On the social implications of collective adaptive systems. *IEEE Technol. Soc. Mag.*, 39(3):36–46.

- Burégio, V. A., de Lemos Meira, S. R., and 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*, pages 885–890. ACM.
- Busoniu, L., Babuska, R., and Schutter, B. D. (2008). A comprehensive survey of multiagent reinforcement learning. *IEEE Trans. Syst. Man Cybern. Part C*, 38(2):156–172.
- Canese, L., Cardarilli, G. C., Di Nunzio, L., Fazzolari, R., Giardinò, D., Re, M., and Spanò, S. (2021). Multi-agent reinforcement learning: A review of challenges and applications. *Applied Sciences*, 11(11):4948.
- Carroll, J. B. (1993). *Human cognitive abilities: A survey of factor-analytic studies*. Cambridge University Press.
- Casadei, R. (2020). *Engineering Self-Adaptive Collective Processes for Cyber-Physical Ecosystems*. PhD thesis, alma.
- Casadei, R. (2023). Macroprogramming: Concepts, state of the art, and opportunities of macroscopic behaviour modelling. *ACM Computing Surveys*.
- Casadei, R., Viroli, M., Audrito, G., Pianini, D., and Damiani, F. (2021). Engineering collective intelligence at the edge with aggregate processes. *Eng. Appl. Artif. Intell.*, 97:104081.
- Chakraborty, A. and Kar, A. K. (2017). Swarm intelligence: A review of algorithms. In Patnaik, S., Yang, X.-S., and Nakamatsu, K., editors, *Nature-Inspired Computing and Optimization: Theory and Applications*, pages 475–494. Springer International Publishing, Cham.
- Chmait, N., Dowe, D. L., Li, Y., Green, D. G., and Insa-Cabrera, J. (2016). Factors of collective intelligence: How smart are agent collectives? In *ECAI 2016 - 22nd European Conference on Artificial Intelligence*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 542–550. IOS Press.
- Chollet, F. (2019). On the measure of intelligence. *arXiv preprint arXiv:1911.01547*.
- Ciancarini, P. (1996). Coordination models and languages as software integrators. *ACM Comput. Surv.*, 28(2):300–302.
- Ciancarini, P. and Hankin, C., editors (1996). *Coordination Languages and Models, First International Conference, COORDINATION '96, Cesena, Italy, April 15-17, 1996, Proceedings*, volume 1061 of *Lecture Notes in Computer Science*. Springer.
- D'Angelo, M., Gerasimou, S., Ghahremani, S., Grohmann, J., Nunes, I., Pournaras, E., and Tomforde, S. (2019). On learning in collective self-adaptive systems: state of practice and a 3d framework. In Litoiu, M., Clarke, S., and Tei, K., editors, *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*, pages 13–24. ACM.
- Daniels, B. C., Ellison, C. J., Krakauer, D. C., and Flack, J. C. (2016). Quantifying collectivity. *Current Opinion in Neurobiology*, 37:106–113.
- 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., Geihs, K., Göschka, K. M., Gorla, A., Grassi, V., Inverardi, P., Karsai, G., Kramer, J., Lopes, A., Magee, J., Malek, S., Mankovski, S., Mirandola, R., Mylopoulos, J., Nierstrasz, O., Pezzè, M., Prehofer, C., Schäfer, W., Schlichting, R. D., Smith, D. B., Sousa, J. P., Tahvildari, L., Wong, K., and Wutke, J. (2010). Software engineering for self-adaptive systems: A second research roadmap. In de Lemos, R., Giese, H., Müller, H. A., and Shaw, M., editors, *Software Engineering for Self-Adaptive Systems II - International Seminar, Dagstuhl Castle, Germany, October 24-29, 2010 Revised Selected and Invited Papers*, volume 7475 of *Lecture Notes in Computer Science*, pages 1–32. Springer.
- de Oca, M. A. M., Ferrante, E., Scheidler, A., Pincioli, C., Birattari, M., and Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intell.*, 5(3-4):305–327.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Dong, X., Yu, Z., Cao, W., Shi, Y., and Ma, Q. (2020). A survey on ensemble learning. *Frontiers Comput. Sci.*, 14(2):241–258.
- 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., Guzzi, J., Longchamp, V., Magnenat, S., Mathews, N., de Oca, M. A. M., O'Grady, R., Pincioli, C., Pini, G., Rétonnaz, P., Roberts, J. F., Sperati, V., Stirling, T. S., Stranieri, A., Stützle, T., Trianni, V., Tuci, E., Turgut, A. E., and Vaussard, F. (2013). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics Autom. Mag.*, 20(4):60–71.
- Dorigo, M. and Stützle, T. (2019). Ant colony optimization: overview and recent advances. *Handbook of metaheuristics*, pages 311–351.
- Dvoretzskii, S., Gong, Z., Gupta, A., Parent, J., and Alicea, B. (2022). Braitenberg vehicles as developmental neurosimulation. *Artificial Life*, 28(3):369–395.
- Engelbrecht, A. P. (2007). *Computational intelligence: an introduction*. John Wiley & Sons.
- Fadul, J. A. (2009). Collective learning: Applying distributed cognition for collective intelligence. *The International Journal of Learning: Annual Review*, 16(4):211–220.
- Faliszewski, P., Mascardi, V., Pelachaud, C., and Taylor, M. E., editors (2022). *21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2022, Auckland, New Zealand, May 9-13, 2022*. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).
- Fenwick, T. (2008). Understanding relations of individual—collective learning in work: A review of research. *Management Learning*, 39(3):227–243.
- Ferber, J. (1999). *Multi-agent systems - an introduction to distributed artificial intelligence*. Addison-Wesley-Longman.

- Ferscha, A. (2015). Collective adaptive systems. In Mase, K., Langheinrich, M., Gatica-Perez, D., Gellersen, H., Choudhury, T., and Yatani, K., editors, *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*, pages 893–895. ACM.
- Ferscha, A., Lukowicz, P., and Zambonelli, F. (2014). The super-organism of massive collective wearables. In Brush, A. J., Friday, A., Kientz, J. A., Scott, J., and Song, J., editors, *The 2014 ACM Conference on Ubiquitous Computing, UbiComp '14 Adjunct, Seattle, WA, USA - September 13 - 17, 2014*, pages 1077–1084. ACM.
- Figueiredo, E. M. N., Macedo, M., Siqueira, H. V., Jr., C. J. S., Gokhale, A., and Filho, C. J. A. B. (2019). Swarm intelligence for clustering - A systematic review with new perspectives on data mining. *Eng. Appl. Artif. Intell.*, 82:313–329.
- Fister, I., Jr., I. F., Yang, X., and Brest, J. (2013). A comprehensive review of firefly algorithms. *Swarm Evol. Comput.*, 13:34–46.
- Francesca, G., Brambilla, M., Brutschy, A., Trianni, V., and Birattari, M. (2014). Automode: A novel approach to the automatic design of control software for robot swarms. *Swarm Intell.*, 8(2):89–112.
- Franklin, S. and Graesser, A. C. (1996). Is it an agent, or just a program?: A taxonomy for autonomous agents. In Müller, J. P., Wooldridge, M. J., and Jennings, N. R., editors, *Intelligent Agents III, Agent Theories, Architectures, and Languages, ECAI '96 Workshop (ATAL), Budapest, Hungary, August 12-13, 1996, Proceedings*, volume 1193 of *Lecture Notes in Computer Science*, pages 21–35. Springer.
- 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. and Wood, Z. (2016). Extensional and intensional collectives and the de re/de dicto distinction. *Applied Ontology*, 11(3):205–226.
- Ganti, R. K., Ye, F., and Lei, H. (2011). Mobile crowdsensing: current state and future challenges. *IEEE Commun. Mag.*, 49(11):32–39.
- Garavan, T. N. and Carbery, R. (2012). Collective learning. In *Encyclopedia of the Sciences of Learning*, pages 646–649. Springer US.
- Garlan, D., Cheng, S., Huang, A., Schmerl, B. R., and Steenkiste, P. (2004). Rainbow: Architecture-based self-adaptation with reusable infrastructure. *Computer*, 37(10):46–54.
- Gelernter, D. and Carriero, N. (1992). Coordination languages and their significance. *Commun. ACM*, 35(2):96–107.
- Gershenson, C., Trianni, V., Werfel, J., and Sayama, H. (2020). Self-organization and artificial life. *Artif. Life*, 26(3):391–408.
- Gignoux, J., Chérel, G., Davies, I. D., Flint, S. R., and Lateltin, E. (2017). Emergence and complex systems: The contribution of dynamic graph theory. *Ecological Complexity*, 31:34–49.
- Gomes, J. C., Urbano, P., and Christensen, A. L. (2013). Evolution of swarm robotics systems with novelty search. *Swarm Intell.*, 7(2-3):115–144.
- Gronauer, S. and Diepold, K. (2022). Multi-agent deep reinforcement learning: a survey. *Artif. Intell. Rev.*, 55(2):895–943.
- Gruber, T. (2008). Collective knowledge systems: Where the social web meets the semantic web. *J. Web Semant.*, 6(1):4–13.
- Gulwani, S., Polozov, O., and Singh, R. (2017). Program synthesis. *Found. Trends Program. Lang.*, 4(1-2):1–119.
- Guo, B., Yu, Z., Zhou, X., and 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*, pages 593–598. IEEE Computer Society.
- Gupta, A., Savarese, S., Ganguli, S., and Fei-Fei, L. (2021). Embodied intelligence via learning and evolution. *Nature Communications*, 12(1).
- Guttmann, C. (2009). Towards a taxonomy of decision making problems in multi-agent systems. In Braubach, L., van der Hoek, W., Petta, P., and Pokahr, A., editors, *Multiagent System Technologies, 7th German Conference, MATES 2009, Hamburg, Germany, September 9-11, 2009. Proceedings*, volume 5774 of *Lecture Notes in Computer Science*, pages 195–201. Springer.
- Ha, D. and Tang, Y. (2021). Collective intelligence for deep learning: A survey of recent developments. *CoRR*, 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., and Zamuda, A. (2016). Hybrid societies: Challenges and perspectives in the design of collective behavior in self-organizing systems. *Frontiers Robotics AI*, 3:14.
- Hasbach, J. D. and Bennewitz, M. (2022). The design of self-organizing human-swarm intelligence. *Adapt. Behav.*, 30(4):361–386.
- He, F., Pan, Y., Lin, Q., Miao, X., and Chen, Z. (2019). Collective intelligence: A taxonomy and survey. *IEEE Access*, 7:170213–170225.
- Hecker, A. (2012). Knowledge beyond the individual? making sense of a notion of collective knowledge in organization theory. *Organization Studies*, 33(3):423–445.
- Hegedüs, I., Danner, G., and Jelasity, M. (2021). Decentralized learning works: An empirical comparison of gossip learning and federated learning. *J. Parallel Distributed Comput.*, 148:109–124.
- Hernandez-Leal, P., Kartal, B., and Taylor, M. E. (2019). A survey and critique of multiagent deep reinforcement learning. *Auton. Agents Multi Agent Syst.*, 33(6):750–797.
- Hernández-Orallo, J. (2017). Evaluation in artificial intelligence: from task-oriented to ability-oriented measurement. *Artificial Intelligence Review*, 48(3):397–447.

- Hernández-Orallo, J. and Dowe, D. L. (2010). Measuring universal intelligence: Towards an anytime intelligence test. *Artificial Intelligence*, 174(18):1508–1539.
- Heylighen, F. (2013). From human computation to the global brain: The self-organization of distributed intelligence. In Michelucci, P., editor, *Handbook of Human Computation*, pages 897–909. Springer.
- Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Syst.*, 1(3).
- Hmelo-Silver, C. E. (2003). Analyzing collaborative knowledge construction: multiple methods for integrated understanding. *Comput. Educ.*, 41(4):397–420.
- Horling, B. and Lesser, V. R. (2004). A survey of multi-agent organizational paradigms. *Knowl. Eng. Rev.*, 19(4):281–316.
- Huang, Y.-C. and Chin, Y.-C. (2018). Transforming collective knowledge into team intelligence: the role of collective teaching. *Journal of Knowledge Management*, 22(6):1243–1263.
- Ibrahim, D. (2016). An overview of soft computing. *Procedia Computer Science*, 102:34–38.
- Iivari, J. (2016). Endogenously emergent information systems. In Goluchowski, J., Pankowska, M., Barry, C., Lang, M., Linger, H., and Schneider, C., editors, *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.
- Jensen, A. R. (1989). The relationship between learning and intelligence. *Learning and Individual Differences*, 1(1):37–62.
- Jung, J. J. (2017). Computational collective intelligence with big data: Challenges and opportunities. *Future Gener. Comput. Syst.*, 66:87–88.
- 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., Gruteser, M., Harchaoui, Z., He, C., He, L., Huo, Z., Hutchinson, B., Hsu, J., Jaggi, M., Javidi, T., Joshi, G., Khodak, M., Konečný, J., Korolova, A., Koushanfar, F., Koyejo, S., Lepoint, T., Liu, Y., Mittal, P., Mohri, M., Nock, R., Özgür, A., Pagh, R., Qi, H., Ramage, D., Raskar, R., Raykova, M., Song, D., Song, W., Stich, S. U., Sun, Z., Suresh, A. T., Tramèr, F., Vepakomma, P., Wang, J., Xiong, L., Xu, Z., Yang, Q., Yu, F. X., Yu, H., and Zhao, S. (2021). Advances and open problems in federated learning. *Found. Trends Mach. Learn.*, 14(1-2):1–210.
- Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., and Schmickl, T. (2015). Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. In Chen, Q., Torroni, P., Villata, S., Hsu, J. Y., and Omicini, A., editors, *PRIMA 2015: Principles and Practice of Multi-Agent Systems - 18th International Conference, Bertinoro, Italy, October 26-30, 2015, Proceedings*, volume 9387 of *Lecture Notes in Computer Science*, pages 201–217. Springer.
- Kephart, J. O. and Chess, D. M. (2003). The vision of autonomic computing. *Computer*, 36(1):41–50.
- Kimmerle, J., Cress, U., and Held, C. (2010). The interplay between individual and collective knowledge: technologies for organisational learning and knowledge building. *Knowledge Management Research & Practice*, 8(1):33–44.
- Kolling, A., Walker, P. M., Chakraborty, N., Sycara, K. P., and Lewis, M. (2016). Human interaction with robot swarms: A survey. *IEEE Trans. Hum. Mach. Syst.*, 46(1):9–26.
- Krajna, A., Brcic, M., Lipic, T., and Doncevic, J. (2022). Explainability in reinforcement learning: perspective and position. *CoRR*, abs/2203.11547.
- Krause, J., Ruxton, G. D., and Krause, S. (2010). Swarm intelligence in animals and humans. *Trends in Ecology & Evolution*, 25(1):28–34.
- Laan, A., Madirolas, G., and de Polavieja, G. G. (2017). Rescuing collective wisdom when the average group opinion is wrong. *Frontiers Robotics AI*, 4:56.
- Le Bon, G. (2002). *The crowd: A study of the popular mind*. Courier Corporation. (Original work published 1895).
- Legg, S., Hutter, M., et al. (2007). A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and applications*, 157:17.
- Liu, Z. and Wang, J. (2020). Human-cyber-physical systems: concepts, challenges, and research opportunities. *Frontiers Inf. Technol. Electron. Eng.*, 21(11):1535–1553.
- Lo, A. W. (2015). The wisdom of crowds vs. the madness of mobs. In *Handbook of Collective Intelligence*. MIT Press.
- Loreti, M. and Hillston, J. (2016). Modelling and analysis of collective adaptive systems with CARMA and its tools. In Bernardo, M., Nicola, R. D., and Hillston, J., editors, *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*, volume 9700 of *Lecture Notes in Computer Science*, pages 83–119. Springer.
- Lykourantzou, I., Vergados, D. J., and Loumos, V. (2009). Collective intelligence system engineering. In Chbeir, R., Badr, Y., Kapetanios, E., and Traina, A. J. M., editors, *MEDES '09: International ACM Conference on Management of Emergent Digital EcoSystems, Lyon, France, October 27-30, 2009*, pages 134–140. ACM.
- Mainland, G., Kang, L., Lahaie, S., Parkes, D. C., and Welsh, M. (2004). Using virtual markets to program global behavior in sensor networks. In Berbers, Y. and Castro, M., editors, *Proceedings of the 11st ACM SIGOPS European Workshop, Leuven, Belgium, September 19-22, 2004*, page 1. ACM.
- Maleszka, M. (2019). Application of collective knowledge diffusion in a social network environment. *Enterp. Inf. Syst.*, 13(7-8):1120–1142.
- Malone, T. W. and Bernstein, M. S. (2015). *Handbook of collective intelligence*. MIT Press.
- Malone, T. W. and Crowston, K. (1994). The interdisciplinary study of coordination. *ACM Comput. Surv.*, 26(1):87–119.

- Mamei, M., Zambonelli, F., and Leonardi, L. (2004). Co-fields: A physically inspired approach to motion coordination. *IEEE Pervasive Comput.*, 3(2):52–61.
- Mataric, M. J. (2007). *The Robotics Primer*. MIT Press.
- Mavrovouniotis, M., Li, C., and Yang, S. (2017). A survey of swarm intelligence for dynamic optimization: Algorithms and applications. *Swarm Evol. Comput.*, 33:1–17.
- Millhouse, T., Moses, M., and Mitchell, M. (2021). Frontiers in collective intelligence: A workshop report. *CoRR*, 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*, pages 1–39. Springer International Publishing.
- Mitchell, M. (2011). Ubiquity symposium: Biological computation. *Ubiquity*, 2011(February):3.
- Mizzi, A., Ellul, J., and 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*, pages 2:1–2:9. ACM.
- Mobus, G. E., Kalton, M. C., et al. (2015). *Principles of systems science*. Springer.
- Moradi, M., Moradi, M., Bayat, F., and Toosi, A. N. (2019). Collective hybrid intelligence: towards a conceptual framework. *International Journal of Crowd Science*.
- Müller-Schloer, C. and Sick, B. (2006). Emergence in organic computing systems: Discussion of a controversial concept. In Yang, L. T., Jin, H., Ma, J., and Ungerer, T., editors, *Autonomic and Trusted Computing, Third International Conference, ATC 2006, Wuhan, China, September 3-6, 2006, Proceedings*, volume 4158 of *Lecture Notes in Computer Science*, pages 1–16. Springer.
- Müller-Schloer, C. and Tomforde, S. (2017). *Organic Computing - Technical Systems for Survival in the Real World*. Birkhäuser.
- Munn, Z., Peters, M. D. J., Stern, C., Tufanaru, C., McArthur, A., and 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).
- Murray, D. G., Yoneki, E., Crowcroft, J., and Hand, S. (2010). The case for crowd computing. In *Proceedings of the second ACM SIGCOMM workshop on Networking, systems, and applications on mobile handhelds*, pages 39–44.
- 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., and Nardelli, P. H. J. (2022). Collective intelligence using 5g: Concepts, applications, and challenges in sociotechnical environments. *IEEE Access*, 10:70394–70417.
- Navarro, I. and Matía, F. (2013). A survey of collective movement of mobile robots. *International Journal of Advanced Robotic Systems*, 10(1):73.
- Nedjah, N. and Junior, L. S. (2019). Review of methodologies and tasks in swarm robotics towards standardization. *Swarm Evol. Comput.*, 50.
- Newton, R., Arvind, and Welsh, M. (2005). Building up to macro-programming: 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*, pages 37–44. IEEE.
- Nguyen, B. H., Xue, B., and Zhang, M. (2020). A survey on swarm intelligence approaches to feature selection in data mining. *Swarm Evol. Comput.*, 54:100663.
- Nguyen, N. T. (2008). Inconsistency of knowledge and collective intelligence. *Cybern. Syst.*, 39(6):542–562.
- Nguyen, N. T., Hwang, D., and Szerbicki, E. (2019). Computational collective intelligence for enterprise information systems. *Enterp. Inf. Syst.*, 13(7-8):933–934.
- Nguyen, N. T., Kowalczyk, R., and Chen, S., editors (2009). *Computational Collective Intelligence. Semantic Web, Social Networks and Multiagent Systems, First International Conference, ICCCI 2009, Wroclaw, Poland, October 5-7, 2009. Proceedings*, volume 5796 of *Lecture Notes in Computer Science*. Springer.
- Nicola, R. D., Jähnichen, S., and Wirsing, M. (2020). Rigorous engineering of collective adaptive systems: special section. *Int. J. Softw. Tools Technol. Transf.*, 22(4):389–397.
- Nielsen, C. B., Larsen, P. G., Fitzgerald, J. S., Woodcock, J., and Peleska, J. (2015). Systems of systems engineering: Basic concepts, model-based techniques, and research directions. *ACM Comput. Surv.*, 48(2):18:1–18:41.
- Nolfi, S. and Floreano, D. (1999). Learning and evolution. *Auton. Robots*, 7(1):89–113.
- Noor, J., Tseng, H., Garcia, L., and Srivastava, M. B. (2019). Ddflow: visualized declarative programming for heterogeneous iot networks. In Landsiedel, O. and Nahrstedt, K., editors, *Proceedings of the International Conference on Internet of Things Design and Implementation, IoTDI 2019, Montreal, QC, Canada, April 15-18, 2019*, pages 172–177. ACM.
- Odell, J. (2002). Objects and agents compared. *J. Object Technol.*, 1(1):41–53.
- Oh, H., Shirazi, A. R., Sun, C., and Jin, Y. (2017). Bio-inspired self-organising multi-robot pattern formation: A review. *Robotics Auton. Syst.*, 91:83–100.
- Oliehoek, F. A. and Amato, C. (2016). *A Concise Introduction to Decentralized POMDPs*. Springer Publishing Company, Incorporated, 1st edition.
- Oliver, P. E. (1993). Formal models of collective action. *Annual Review of Sociology*, 19(1):271–300.
- Omicini, A., Ricci, A., and Viroli, M. (2008). Artifacts in the a&a meta-model for multi-agent systems. *Auton. Agents Multi Agent Syst.*, 17(3):432–456.
- O’Neill, M. and Spector, L. (2020). Automatic programming: The open issue? *Genet. Program. Evolvable Mach.*, 21(1-2):251–262.

- Peeters, M. M. M., van Diggelen, J., van den Bosch, K., Bronkhorst, A. W., Neerinx, M. A., Schraagen, J. M., and Raaijmakers, S. (2021). Hybrid collective intelligence in a human-ai society. *AI Soc.*, 36(1):217–238.
- Petruzzi, P. E., Busquets, D., and 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*, pages 21–30. IEEE Computer Society.
- Petticrew, M. and Roberts, H. (2008). *Systematic reviews in the social sciences: A practical guide*. John Wiley & Sons.
- Pianini, D., Casadei, R., and Viroli, M. (2022a). Self-stabilising priority-based multi-leader election and network partitioning. In *2022 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*. IEEE.
- Pianini, D., Casadei, R., Viroli, M., and Natali, A. (2021). Partitioned integration and coordination via the self-organising coordination regions pattern. *Future Gener. Comput. Syst.*, 114:44–68.
- Pianini, D., Pettinari, F., Casadei, R., and Esterle, L. (2022b). A collective adaptive approach to decentralised k-coverage in multi-robot systems. *ACM Trans. Auton. Adapt. Syst.*, 17:4:1–4:39.
- Pincirol, C. and Beltrame, G. (2016). Buzz: A programming language for robot swarms. *IEEE Softw.*, 33(4):97–100.
- 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*, pages 205–214. IEEE.
- Pournaras, E. (2020b). Proof of witness presence: Blockchain consensus for augmented democracy in smart cities. *J. Parallel Distributed Comput.*, 145:160–175.
- Pournaras, E., Pilgerstorfer, P., and Asikis, T. (2018). Decentralized collective learning for self-managed sharing economies. *ACM Trans. Auton. Adapt. Syst.*, 13(2):10:1–10:33.
- Prasetyo, J., Masi, G. D., and Ferrante, E. (2019). Collective decision making in dynamic environments. *Swarm Intell.*, 13(3-4):217–243.
- Psaier, H. and Dustdar, S. (2011). A survey on self-healing systems: approaches and systems. *Computing*, 91(1):43–73.
- Qiu, W. and Zhu, H. (2022). Programmatic reinforcement learning without oracles. In *International Conference on Learning Representations*.
- Quinn, A. J. and Bederson, B. B. (2011). Human computation: a survey and taxonomy of a growing field. In Tan, D. S., Amershi, S., Begole, B., Kellogg, W. A., and Tungare, M., editors, *Proceedings of the International Conference on Human Factors in Computing Systems, CHI 2011, Vancouver, BC, Canada, May 7-12, 2011*, pages 1403–1412. ACM.
- Rainey, L. B. and Jamshidi, M., editors (2018). *Engineering Emergence*. CRC Press.
- Rajasekhar, A., Lynn, N., Das, S., and Suganthan, P. N. (2017). Computing with the collective intelligence of honey bees - A survey. *Swarm Evol. Comput.*, 32:25–48.
- Reeve, C. L., Bonaccio, S., Chamorro-Premuzic, T., Furnham, A., and von Stumm, S. (2011). The nature and structure of “intelligence”. *The Wiley-Blackwell handbook of individual differences*, pages 187–216.
- Reina, A., Miletitch, R., Dorigo, M., and Trianni, V. (2015). A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intell.*, 9(2-3):75–102.
- Ritz, F., Phan, T., Müller, R., Gabor, T., Sedlmeier, A., Zeller, M., Wiegardt, J., Schmid, R. N., Sauer, H., Klein, C., and Linnhoff-Popien, C. (2021). Specification aware multi-agent reinforcement learning. In Rocha, A. P., Steels, L., and van den Herik, H. J., editors, *Agents and Artificial Intelligence - 13th International Conference, ICAART 2021, Virtual Event, February 4-6, 2021, Revised Selected Papers*, volume 13251 of *Lecture Notes in Computer Science*, pages 3–21. Springer.
- Rizk, Y., Awad, M., and Tunstel, E. W. (2018). Decision making in multiagent systems: A survey. *IEEE Trans. Cogn. Dev. Syst.*, 10(3):514–529.
- Rodríguez, A., Grushin, A., and Reggia, J. A. (2007). Swarm intelligence systems using guided self-organization for collective problem solving. *Adv. Complex Syst.*, 10(supp01):5–34.
- Rossi, F., Bandyopadhyay, S., Wolf, M. T., and Pavone, M. (2018). Review of multi-agent algorithms for collective behavior: a structural taxonomy. *CoRR*, abs/1803.05464.
- Rostami, M., Kolouri, S., Kim, K., and Eaton, E. (2018). Multi-agent distributed lifelong learning for collective knowledge acquisition. In André, E., Koenig, S., Dastani, M., and Sukthankar, G., editors, *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, pages 712–720. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM.
- Russell, S. and Norvig, P. (2020). *Artificial Intelligence: A Modern Approach (4th Edition)*. Pearson.
- Salehie, M. and Tahvildari, L. (2009). Self-adaptive software: Landscape and research challenges. *ACM Trans. Auton. Adapt. Syst.*, 4(2):14:1–14:42.
- Salminen, J. (2012). Collective intelligence in humans: A literature review. *CoRR*, abs/1204.3401.
- Sari, A., Tosun, A., and Alptekin, G. I. (2019). A systematic literature review on crowdsourcing in software engineering. *J. Syst. Softw.*, 153:200–219.
- Scekic, O., Schiavinotto, T., Videnov, S., Rovatsos, M., Truong, H. L., Miorandi, D., and Dustdar, S. (2020). A programming model for hybrid collaborative adaptive systems. *IEEE Trans. Emerg. Top. Comput.*, 8(1):6–19.
- Schirner, G., Erdogmus, D., Chowdhury, K. R., and Padir, T. (2013). The future of human-in-the-loop cyber-physical systems. *Computer*, 46(1):36–45.



- Schmickl, T., Möslinger, C., and Crailsheim, K. (2006). Collective perception in a robot swarm. In Sahin, E., Spears, W. M., and Winfield, A. F. T., editors, *Swarm Robotics, Second International Workshop, SAB 2006, Rome, Italy, September 30-October 1, 2006, Revised Selected Papers*, volume 4433 of *Lecture Notes in Computer Science*, pages 144–157. Springer.
- Schranz, M., Caro, G. A. D., Schmickl, T., Elmenreich, W., Arvin, F., Sekercioglu, Y. A., and Sende, M. (2021). Swarm intelligence and cyber-physical systems: Concepts, challenges and future trends. *Swarm Evol. Comput.*, 60:100762.
- Schweikard, D. P. and Schmid, H. B. (2013). Collective Intentionality. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2021 edition.
- 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., and Porter, B. (2022). The state of the art of macroprogramming in iot: An update. *J. Internet Serv. Appl.*, 13(1):54–65.
- Shoham, Y. (1993). Agent-oriented programming. *Artif. Intell.*, 60(1):51–92.
- Siddique, N. H. and Adeli, H. (2015). Nature inspired computing: An overview and some future directions. *Cogn. Comput.*, 7(6):706–714.
- Sieben, A., Schumann, J., and Seyfried, A. (2017). Collective phenomena in crowds—where pedestrian dynamics need social psychology. *PLOS ONE*, 12(6):1–19.
- Sigaud, O. (2022). Combining evolution and deep reinforcement learning for policy search: a survey. *ACM Transactions on Evolutionary Learning and Optimization*.
- Smirnov, A. and 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*, pages 113–124, Cham. Springer International Publishing.
- Smith, E. A. (2010). Communication and collective action: language and the evolution of human cooperation. *Evolution and human behavior*, 31(4):231–245.
- Solar-Lezama, A. (2009). The sketching approach to program synthesis. In Hu, Z., editor, *Programming Languages and Systems, 7th Asian Symposium, APLAS 2009, Seoul, Korea, December 14-16, 2009. Proceedings*, volume 5904 of *Lecture Notes in Computer Science*, pages 4–13. Springer.
- Sotnik, G. (2018). The SOSIEL platform: Knowledge-based, cognitive, and multi-agent. *Biologically Inspired Cognitive Architectures*, 26:103–117.
- Stein, A., Tomforde, S., Botev, J., and Lewis, P. R. (2021). Lifelike computing systems. In *Proceedings of the Lifelike Computing Systems Workshop (LIFELIKE)*.
- Stradner, J., Thenius, R., Zahadat, P., Hamann, H., Crailsheim, K., and Schmickl, T. (2013). Algorithmic requirements for swarm intelligence in differently coupled collective systems. *Chaos, Solitons & Fractals*, 50:100–114.
- Sumpter, D. J. (2010). *Collective animal behavior*. Princeton University Press.
- Suran, S., Pattanaik, V., and Draheim, D. (2020). Frameworks for collective intelligence: A systematic literature review. *ACM Comput. Surv.*, 53(1).
- Surowiecki, J. (2005). *The wisdom of crowds*. Anchor.
- Szuba, T. M. (2001). *Computational collective intelligence*. John Wiley & Sons, Inc.
- Tang, M. and 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.
- Telikani, A., Tahmassebi, A., Banzhaf, W., and Gandomi, A. H. (2022). Evolutionary machine learning: A survey. *ACM Comput. Surv.*, 54(8):161:1–161:35.
- ter Beek, M. H. and Sirjani, M., editors (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*, volume 13271 of *Lecture Notes in Computer Science*. Springer.
- Trianni, V. (2008). *Evolutionary Swarm Robotics - Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, volume 108 of *Studies in Computational Intelligence*. Springer.
- Tsvetkova, M., Yasseri, T., Meyer, E. T., Pickering, J. B., Engen, V., Walland, P., Lüders, M., Følstad, A., and Bravos, G. (2017). Understanding human-machine networks: A cross-disciplinary survey. *ACM Comput. Surv.*, 50(1):12:1–12:35.
- Tumer, K. and Wolpert, D. H. (2004). A survey of collectives. In Tumer, K. and Wolpert, D. H., editors, *Collectives and the design of complex systems*, chapter 1. Springer Science & Business Media.
- Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers Robotics AI*, 4:9.
- Valentini, G., Hamann, H., and Dorigo, M. (2014). Self-organized collective decision making: the weighted voter model. In Bazzan, A. L. C., Huhns, M. N., Lomuscio, A., and Scerri, P., editors, *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS ’14, Paris, France, May 5-9, 2014*, pages 45–52. IFAAMAS/ACM.
- van der Hoek, W. and Wooldridge, M. J. (2003). Towards a logic of rational agency. *Log. J. IGPL*, 11(2):135–159.
- van Gerven, M. (2017). Computational foundations of natural intelligence. *Frontiers Comput. Neurosci.*, 11:112.
- Varenne, F., Chaigneau, P., Petitot, J., and Doursat, R. (2015). Programming the emergence in morphogenetically architected complex systems. *Acta biotheoretica*, 63(3):295–308.
- Verbraeken, J., Wolting, M., Katzy, J., Kloppenburg, J., Verbelen, T., and Rellermeyer, J. S. (2020). A survey on distributed machine learning. *ACM Comput. Surv.*, 53(2):30:1–30:33.

- Verma, A., Murali, V., Singh, R., Kohli, P., and Chaudhuri, S. (2018). Programmatically interpretable reinforcement learning. In Dy, J. G. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 5052–5061. PMLR.
- Viroli, M., Beal, J., Damiani, F., Audrito, G., Casadei, R., and Pianini, D. (2019). From distributed coordination to field calculus and aggregate computing. *J. Log. Algebraic Methods Program.*, 109.
- Wang, Y. (2009). On abstract intelligence: Toward a unifying theory of natural, artificial, machinable, and computational intelligence. *Int. J. Softw. Sci. Comput. Intell.*, 1(1):1–17.
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics Auton. Syst.*, 39(1):1–18.
- Wechsler, D. (1946). *The measurement of adult intelligence* (3rd ed.). Williams & Wilkins Co.
- Weiser, M. (1991). The computer for the 21 st century. *Scientific american*, 265(3):94–105.
- Weyns, D. and Michel, F. (2014). Agent environments for multi-agent systems - A research roadmap. In Weyns, D. and Michel, F., editors, *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*, volume 9068 of *Lecture Notes in Computer Science*, pages 3–21. Springer.
- Weyns, D., Schmerl, B. R., Grassi, V., Malek, S., Mirandola, R., Prehofer, C., Wuttke, J., Andersson, J., Giese, H., and Göschka, K. M. (2010). On patterns for decentralized control in self-adaptive systems. In de Lemos, R., Giese, H., Müller, H. A., and Shaw, M., editors, *Software Engineering for Self-Adaptive Systems II - International Seminar, Dagstuhl Castle, Germany, October 24-29, 2010 Revised Selected and Invited Papers*, volume 7475 of *Lecture Notes in Computer Science*, pages 76–107. Springer.
- Wolf, T. D. and Holvoet, T. (2004). Emergence versus self-organisation: Different concepts but promising when combined. In *Engineering Self-Organising Systems, Methodologies and Applications*, volume 3464 of *Lecture Notes in Computer Science*, pages 1–15. Springer.
- Wolpert, D. H. (2003). *Collective Intelligence*, chapter 17, pages 245–. Wiley.
- Wood, Z. (2016). Considering collectives: Roles, members and goals. In *FOIS*, pages 359–372.
- Wood, Z. and Galton, A. (2009). A taxonomy of collective phenomena. *Applied Ontology*, 4(3-4):267–292.
- Wooldridge, M. J. (2009). *An Introduction to MultiAgent Systems, Second Edition*. Wiley.
- Woolley, A. W., Aggarwal, I., and Malone, T. W. (2015). Collective intelligence and group performance. *Current Directions in Psychological Science*, 24(6):420–424.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., and Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688.
- Yang, T., Yi, X., Wu, J., Yuan, Y., Wu, D., Meng, Z., Hong, Y., Wang, H., Lin, Z., and Johansson, K. H. (2019). A survey of distributed optimization. *Annu. Rev. Control.*, 47:278–305.
- Yang, X. and He, X. (2013). Bat algorithm: literature review and applications. *Int. J. Bio Inspired Comput.*, 5(3):141–149.
- Yang, Y., Xiao, Y., and Li, T. (2021). A survey of autonomous underwater vehicle formation: Performance, formation control, and communication capability. *IEEE Commun. Surv. Tutorials*, 23(2):815–841.
- Yu, C., Chai, Y., and Liu, Y. (2018). Literature review on collective intelligence: a crowd science perspective. *International Journal of Crowd Science*.
- Yu, Z., Li, Q., Yang, F., and Guo, B. (2021). Human-machine computing. *CCF Trans. Pervasive Comput. Interact.*, 3(1):1–12.
- 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., and Ye, J. (2015). Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive Mob. Comput.*, 17:236–252.
- Zedadra, O., Guerrieri, A., Jouandeau, N., Spezzano, G., Seridi, H., and Fortino, G. (2018). Swarm intelligence-based algorithms within iot-based systems: A review. *J. Parallel Distributed Comput.*, 122:173–187.
- Zhang, H., Dong, Y., Chiclana, F., and Yu, S. (2019a). Consensus efficiency in group decision making: A comprehensive comparative study and its optimal design. *Eur. J. Oper. Res.*, 275(2):580–598.
- Zhang, K., Yang, Z., and Basar, T. (2019b). Multi-agent reinforcement learning: A selective overview of theories and algorithms. *CoRR*, abs/1911.10635.
- Zhang, S., Lee, C. K. M., Chan, H. K., Choy, K. L., and Wu, Z. (2015). Swarm intelligence applied in green logistics: A literature review. *Eng. Appl. Artif. Intell.*, 37:154–169.
- Zhang, T., Lowmanstone, L., Wang, X., and Glassman, E. L. (2020). Interactive program synthesis by augmented examples. In Iqbal, S. T., MacLean, K. E., Chevalier, F., and Mueller, S., editors, *UIST '20: The 33rd Annual ACM Symposium on User Interface Software and Technology, Virtual Event, USA, October 20-23, 2020*, pages 627–648. ACM.
- Zhen, Y., Khan, A., Nazir, S., Zhao, H., Alharbi, A., and Khan, S. (2021). Crowdsourcing usage, task assignment methods, and crowdsourcing platforms: A systematic literature review. *J. Softw. Evol. Process.*, 33(8).
- Zheng, L., Yang, J., Cai, H., Zhou, M., Zhang, W., Wang, J., and Yu, Y. (2018). Magent: A many-agent reinforcement learning platform for artificial collective intelligence. In McIlraith, S. A. and Weinberger, K. Q., editors, *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*,

*(AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New*

*Orleans, Louisiana, USA, February 2-7, 2018, pages 8222–8223. AAAI Press.*