On the Social Implications of Collective Adaptive Systems

# On the Social Implications of Collective Adaptive Systems

Antonio Bucchiarone*, Giacomo Cabri†, Roberto Casadei‡, Mirko D'Angelo§, Martina De Sanctis¶,
Simon Dobson‖, Danilo Pianini‡, and Mirko Viroli‡

*Fondazione Bruno Kessler IT
bucchiarone@fbk.eu

†Università di Modena e Reggio Emilia IT
giacomo.cabri@unimore.it

‡Alma Mater Studiorum—Università di Bologna IT
{roby.casadei,danilo.pianini,mirko.viroli}@unibo.it

§Linnaeus University SE
mirko.dangelo@lnu.se

¶Gran Sasso Science Institute (GSSI) IT
martina.desanctis@gssi.it

‖School of Computer Science, University of St Andrews UK
simon.dobson@st-andrews.ac.uk

❖

**Abstract**—Many *Collective Adaptive Systems (CASs)* exist in nature: think of ant colonies, where large collectives of ants operate autonomously but interact with other ants and the environment to provide resilient global behaviours that sustain their colony. Following scientific studies that were aimed at understanding and predicting the evolution of these systems, and fueled by technological advances, research has started to investigate *CAS engineering*: the methods, tools, and techniques for building CASs. This naturally leads to a vision where collectives of humans and computational elements, situated both in the digital and physical worlds, collaborate to give rise to "intelligent" collective behaviour supporting novel kinds of applications and services. Humans can be involved in two ways: both as *users* and as *components* of the CAS, in the sense that human behaviours and limitations are often integral to the system description. This has significant social implications that need to be considered by CAS researchers: in this paper, we share a discussion that took place between some experts thinking about CAS engineering, focusing on the social implication of CASs and related open research challenges. We hope that this provides a useful context for future research projects, research grant proposals, and research directions.

## 1 INTRODUCTION

Current technological and societal trends foster the development of artificial systems composed of large-scale ensembles of widely distributed, largely autonomous and heterogeneous entities situated both in the digital and physical world. As business and societal demands upon services and applications increasingly exceed the capabilities of individual devices, it becomes crucial to exploit higher levels of *system effects* ("the whole is greater than the sum of the parts") and *integration* of human and artificial intelligence. This can be afforded by systems made of collectives of humans and computational components able to work in some target environment in a co-ordinated fashion.

We envision a future where humans interact with, and take part in, *Collective Adaptive Systems (CASs)* [1], [2], by which we mean computational systems exhibiting collective and decentralized adaptation mechanisms that must simultaneously address critical changes in run-time conditions while maintaining the consistent functioning of the ensemble. CASs consist of diverse heterogeneous entities or actors – computers, services, devices, sensors, humans, networks, robots, and so forth – that are individually autonomous but that must co-operate with each other to accomplish collective tasks [3]. As a rule, the environment in which such systems operate is continuously changing. These changes include those in the context in which entities live, the availability of new entities (or their exit from the system), and the system requirements and preferences. This means that CASs require *continuous adaptation* to operate in a continuously-changing environment. Similar concepts have been studied in various domains:

- *Swarm Intelligence*, where actors are essentially homogeneous and are able to adapt their behaviour considering only local knowledge [4], [5];
- *Autonomic and Self-Adaptive Computing*, the field devoted to the theory and practice of *self-\* systems* (i.e., systems able to self-repair, self-optimise, self-

configure, etc.) [6], [7];

- *Service-Based Systems*, where services are designed independently by different service providers and are composed to achieve a predefined goal [8];
- *Multi-Agent Based Systems*, where activities of different actors are regulated by certain collectively defined rules (norms) [9], [10];
- *Pervasive Computing* and the *Internet of Things*, where everyday objects and activities are mediated through software; and
- *Complex Adaptive Systems* [11], in which many independent elements or agents interact, leading to emergent outcomes that are often difficult (or impossible) to predict simply by looking at the individual interactions[1].

All the initiatives, in the context of CAS, have the objective of investigating the foundations of large-scale collaborative systems by developing concepts and mechanisms for the autonomous configuration of *ensembles* [12], collections of autonomous entities, and their evolution through self-adaptation. For the last decade the scientific community has invested considerable effort into making distributed computing systems adaptive across diverse scales. The grand vision of enabling all computing resources around us to join in and solve problems in a collective effort has been tackled in various projects [13], [14], [15], [16], and has also opened up many challenges that still need to be considered [17].

This article captures a vision of the social implications of CASs and the corresponding open research challenges. The idea comes from the 3rd eCAS Workshop on Engineering Collective Adaptive Systems[2] held in conjunction with the 12th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2018) in Trento, Italy[3]. We synthesize the discussions emerged during this workshop and build on the following three core topics:

1) what is the role of CASs for human society, their *benefits* and *risks*?,
2) what is the role of humans in CASs, what *values* can they provide?, and
3) what are the difficulties in *engineering* CASs *for people and the society*, especially with regard to their *limitations* and *reliability*?

In the rest of the article, we explore these issues with a view to sharing with the community the vision of CAS as a key approach to building future software systems for human society.

## 2 THE RISE OF COLLECTIVE ADAPTIVE SYSTEMS

In recent decades we have witnessed huge technological advances with tremendous impact on science, business, and life in general. Computers are everywhere, and pervasive networks connect the world and its inhabitants. More and more devices are being deployed into our environments, making ourselves, our homes and cities *smarter* and *more connected*.

1. https://code.org/educate/science/files/CS_in_Science_Background_papers.pdf
2. http://archive.ph/wip/QdnHm
3. https://saso2018.fbk.eu/

Two long-standing trends are at play in this evolution. The first is the *distribution* of computing across components and spaces, by which processing can be injected into multiple locations to parallelize activity, exploit locality, and tolerate (some) faults. The other endows computational devices with *autonomic* capabilities, in order to make them *adapt* and evolve in a dynamic environment with little or no human intervention.

These issues have especially been addressed in the broad area of *distributed artificial intelligence (DAI)* [18] and, in particular, in the context of *multi-agent systems (MASs)* research [19]. A MAS is a system of *autonomous* entities, *agents*, situated into some *environment* and interacting with other agents to achieve both individual and collective goals. Autonomy implies *reactivity* (responding to environmental change), *proactivity* (taking the initiative), and *sociality* (interacting and collaborating with other agents) [20], [21]. In MASs, there are two key perspectives: of individual agents (the *micro level*), which can be conceived according to a *strong* or *weak* notion of agency (depending on whether or not agents exhibit *intelligence* through human-like, *mental components*); and of the agent society (the *macro level*), where issues like *co-ordination* (the ordering of interaction) and *self-organization* (having the system itself maintain its own order) emerge. When agents in a MAS sense and respond directly to their environment they are often referred to as *cyber-physical systems*, and we shall use these terms largely interchangeably in this article.

The scientific community in ICT has invested considerable effort into constructing intelligent distributed computing systems, with an initial focus on small scale applications. Notable advances have been made in the areas of robotics [22], [23], [24], [25], [26] and in selected fields of distributed computing. We have seen new adaptive middleware systems [27], [28], [29], new programming paradigms [30], [31], [32], bio-inspired computing [33], [34], [35] and artificial intelligence [36].

As the density and scale of cyber-physical systems increases, more and more emphasis is put on their *collective dimension* and the ability of adapting *groups* of devices consistently. This has led to specific research efforts to address the problem of capturing the *global, system-level behaviour* developed by groups of interacting entities. For instance, swarm intelligence [37], – inspired by natural systems like ant colonies or bird flocks – builds sophisticated global behaviour (*e.g.*, resilient or ordered responses) from interactions based on very simple local rules. Work such as reported by Beal *et al.* [32] develops this idea to engineer robust distributed components of resilient behaviour, applicable at any scale.

More generally, and especially when taking into account the heterogeneity of modern cyber-physical systems which may involve components with various capabilities, autonomy, and "smartness", we stray into the field of *collective adaptive systems (CASs)*. A CAS can be defined as *a distributed cyber-physical system composed of entities that interact with each other and are able to modify their behaviour to adapt to different external or internal circumstances [2].*

# 3 CASs: Concepts and Perspectives

By definition, a CAS is a *system*: a set of interacting entities that is crucially:

- *collective* – the behaviours of individuals have coherence at the collective level; and
- *adaptive* – it is able to change its control rules as a consequence of internal or external factors.

An example of a *non-collective* (but still possibly adaptive) system is one in which all the individuals do interact but are driven only by their own goals. An example of collective but *non-adaptive* system is a cellular automaton, driven by fixed rules.

Key for a characterization of a collective system is the definition of at least two levels – the *micro* level of the units, and the *macro* level of the whole. In true system-building style a collective may consist of multiple (transient or permanent) sub-collectives and may be part of super-collectives. The boundaries and containment of collectives essentially depend on the perspective taken for both analysis and synthesis.

As for MASs, CASs often consist of autonomous entities. Autonomy, generally conceived as the ability to "self-govern", is actually a complex notion that can have multiple characterizations and is not simply a black-or-white feature. Importantly for systems designers, autonomy is often a source of *unpredictability* as behaviours and responses vary depending on an individual agent's self-determination or collective determination of the system. The challenge in CASs, as for MASs, is to promote collaboration between components, such that together they can carry out tasks that none of them, as an individual component, would be able to perform [9]. This is especially hard in competitive settings, where each component is fully self-interested, but has to interact with other components to achieve its goal [38]; this may involve *negotiation* and *trust*.

The other key source for unpredictability is the environment, because it generally has complex dynamics that are at most partially observed. CASs are essentially *situated*: they are made up of components that are immersed into some (logical and/or physical) environment and are engaged in non-trivial interactions with it. Indeed, such systems are adaptive just because they need to *evolve* in order to respond to changes in the environment, the input patterns, or the responses required.

From a structural point of view, a CAS may exhibit various dynamic structures, often built and sustained through self-organization: a robust, internal process by which order is continually sought, often in an emergent way [39]. The presence of several (autonomous) components requires an appropriate organization to be enforced so as to assign roles and responsibilities to the components themselves [40] and promote the ordered development of collective behaviour. This is closely related to the dual idea of co-ordination, where rules are enacted to constrain interaction.

Crucially, CASs often feature *emergence*, whereby macro-level properties and behaviours spring from decentralized, micro-level activity. The key point is that the global properties arising from the interactions within the system cannot easily be traced back to the properties and behaviours of the individual parts. This poses significant challenges in the understanding, design, and control of CASs, the extent that many engineering approaches attempt to eliminate the phenomenon. However, emergent behaviour must often be expected in CASs consisting of autonomous components and featuring decentralized control, since there is no single locus of control over behaviour. Moreover decentralization is fundamental in allowing systems to scale (adding more components to avoid overloading) and to achieve robustness (as functionality is not confined to few, critical components).

CASs exist in the world, in our minds as mental models, in theory as mathematical models, and in software [41]. Nature has been a great source of inspiration for mechanisms used by engineers to endow artificial systems with features like self-organization and resilience [42]. CAS-oriented features are especially useful in scenarios that are highly dynamic with complex dynamics, highly open with a changing component population, and operate at a very large scale.

## 3.1 CASs and humans

Humans are *prosumers* of CASs. In general they actively participate in the CAS by directly producing data (for example through handheld or wearable devices), by monitoring and controlling components, or by providing specific services in a larger workflow. At the same time, they are consumers of the information and services provided by the CASs, as the ultimate goal of the system is to provide some form of utility to human users. Of course, humans are autonomous, generally unpredictable, have goals evolving with time and context, and are definitely not directly controllable: humans are therefore at the same time one of the most valuable components, and one of the largest sources of design challenges, for CASs. The ultimate value of humans in CASs is simply that software, whatever its intelligence, cannot always substitute for human experience and skill in evaluating situations.

Among the most relevant issues a CAS needs to tackle when leveraging users' data is *privacy*. Although the work on privacy in computational systems is large and well-developed, little is found specifically devoted to CASs despite the fact that their peculiarities seem to call for tailored practices. Due to the very nature of a CAS being "collective" at its core, data may be processed by multiple agents in multiple locations, and it can therefore be challenging to provide guarantees on where and when data in CASs is being used and processed. This may be in conflict with specific privacy requirements. To the best of our knowledge, no CAS-specific best practices for privacy assurance have been proposed so far.

*Security* is related to privacy, but focuses on how to detect and counteract malicious intrusion. As with privacy, very little work currently specifically addresses CAS security [43], [44]; works found in literature do not generally tackle the issue at large, but typically propose designs inheriting security properties from the underlying platform [45], or investigate collective resilience to attacks in specific settings [46]. Fortunately, related fields such as autonomic computing have developed a body of literature that can be leveraged as a starting point for new, CAS-specific practices for both detection and response to attacks.

The situation is much brighter for *safety*, the ability to respond to accidental disruption or degradation. CASs are designed for resilience, and in many cases can promptly return a nominal state after disruptive events. A sub-set of such events may include damage by an attacker.

To what degree do users need to understand a CAS before they will trust it? Human behaviour is strongly influenced by knowledge about the services and perceived reputation of providers. If there is insufficient understanding and of over- or under-trust in its provider, humans can be reluctant to exploit services or – possibly worse – be prone to misuse them beyond their designed scope. There are several examples of these phenomena in adaptive automotive systems [47], [48]. Human trust in a system is influenced, among other thing, by its *usability* and *observability*. The former is related to the ease to access, use, and control a service: high usability of a well-functioning system leads to increased trust. The latter can be a double-edged sword: an "inscrutable" black box may serve people correctly and lower the learning curve, and positively influence usability, but poor or incorrect service may then remain imperceptible until an unacceptably late stage.

Finally we should consider *liability*. It is as yet unclear, both ethically and legally, who should be held accountable for unwanted, unwarranted, illegal, or dangerous behaviours in a CAS—as is indeed the case for many other classes of systems. The matter is becoming pressing [49] because of the technological advances leading to general availability of autonomous systems, such as in the automotive domain.

## 3.2 On human-oriented CAS design and engineering

While CASs can be extremely useful, with broad impact on humans and the society, they are not easy to engineer and often feature significant unavoidable complexity. Their design and construction is further complicated when human interactions enter the equation.

CASs usually rely on learning and emergence to cope with change and unpredictability [50]. Both techniques (often used in conjunction) need subsequent phases of design, validation, and tuning. Once the overall approach for the realization of a CAS has been devised and a first candidate prototype has been readied, a verification phase is necessary.

Unfortunately, nominal conditions of operation can often not be produced on demand: consequently, validation is performed using techniques that either operate on a model of the system, or emulate the environment in which the final system will operate. Techniques include *formal verification* (proofs of theorems on a mathematical model of the system), *model checking* (verification of properties through complete or probabilistic state space exploration [51]), and *simulation* (stochastic sampling of the execution of a system model). These approaches complement each other and verify a system at different levels of granularity. Comprehensive validation of a CAS should, when possible and applicable, include all of them. Very frequently, however, verification is performed without considering the human and societal aspects. As a consequence, a promising system may actually fail once it is deployed through the developers having missed relevant aspects within the modelling and/or simulation

environments. Some recent work aims to capture some of the cognitive, emotional, and interpersonal aspects into a computational model [52], but these approaches remain limited.

## 4 APPLICATION DOMAINS

In this section we present some applications of CASs in order to deepen their social implications.

### 4.1 Ambient Intelligence

CASs can be used to make environments "smart".

When collectives with adaptation capabilities are deployed pervasively into an environment, cyber-physical ecosystems emerge to realise Weiser's vision [53] of truly pervasive computing. By spreading collective computation and connectivity into the spaces attended by humans (including homes, offices, streets, green spaces, and other places), it becomes possible to provide contextual services and adapt their provisioning to user preferences and situations while also providing continuity and coherence. By co-ordinating their activities, CAS entities can adjust the environment in which they reside to make it more ergonomic or functional for the humans who inhabit that environment. In a conference room, for example, a CAS involving the lights, projectors, and thermostats could reactively modulate the environment to make presenters and their audience comfortable. In a different scenario, an ensemble of smart screens and speakers could dynamically provide content to a person based on her preferences, even tracking her while she was walking [16]. In an office, chairs, desks, and monitors may adapt their height and inclination to the preferences of a worker moving to a new workstation.

### 4.2 Smart Logistics

In the domain of smart mobility and transport, the general goal is to improve the efficiency and/or the quality of the movement of people and goods. In particular, since people and goods often have common needs (many people may need to be at some location at some point in time, or many goods may need to be transported to nearby destinations with the same packaging conditions), these needs can be satisfied in a co-ordinated, collective fashion so that both individual and global costs can be reduced by sharing strategies and efforts.

#### 4.2.1 Moving people

Smart mobility systems aim to support the efficient exploitation of the city transport facilities as well as providing sustainable mobility within an urban environment. To accomplish this, systems must take into account the dynamic variabilities in the environment such as users' mobility requirements and service availability. When we consider the common needs of users moving around, in addition to the emergence of the sharing economy such as ride-sharing and car-pooling, CASs may play a relevant role. They may be successfully exploited for forming and managing collectives of users sharing common goals, as well as dynamically adapting these collectives in the face of changes in user needs or changes in the environment.

A relevant example is given by the *urban mobility system* [15], where heterogeneous available services, including traditional transportation services and ride-share services, are integrated to support all the phases of users' journeys for the whole duration of travel. In case of context changes affecting users' journeys, the system may offer alternative solutions by exploiting the adaptation capabilities of the collective.

Users may also be called upon to actively participate in the system they use by playing different roles. For instance, users can act both as consumers (by requiring a ride-share) and as providers (by offering ride sharing or car-pooling). Moreover they may be actively and collectively involved in the dynamic evolution of the system, which may be regulated by complex collaboration patterns such as negotiation, competition and collaboration [9]. Leveraging CASs to induce users to prefer collective mobility solutions can also be considered as a means to support sustainable mobility and reduce the congestion in city centres. However, keeping people engaged with collectives, for example when common needs evolve or are overcome by individual requirements, pose challenges in terms of ensuring user satisfaction. The CAS needs to make collective solutions continually more attractive with respect to traditional ones.

### 4.2.2  Moving goods

Logistics scenarios refer to the transfer and delivery of goods from manufacturers to retailers and then to end-users. Logistics systems aim to support an efficient management of goods, for any kind of goods (with their specific requirements), while taking account of different-sized transportation networks and geographical areas covered by diverse vehicle types [54]. To accomplish this, logistics systems must deal with different dynamics and constraints (legal, environmental, economical) at different levels of granularity [55]. Given the large number of diverse actors involved (including importers, authorities, middlemen, and logistics service providers) that must organize and co-ordinate themselves, CASs may be seen as a natural solution for automation.

Despite working at different levels of the system and with different responsibilities, the various actors mostly share common goals – for example to deliver goods in time, in an acceptable condition, at minimal cost. For each flow of goods, the actors should behave as a collective able to execute efficient and decentralized decision making during the nominal execution of the system and promptly react to changes affecting the delivery process.

Although CASs are intuitively suitable for logistics scenarios, interesting challenges arise when looking at the impact of runtime adaptation on those actors who actually carry out the deliveries. For instance, truck drivers are directly affected by dynamic re-allocation of routes or goods. In these contexts, CASs must balance adaptation in pursuit of global optima against the local needs of drivers and other service providers. One advantage of a CAS solution is that these ethical issues can be encoded into the fabric of the automation.

### 4.3  Public Safety and Security

Public safety and security involve countermeasures for protection against threats to the safety and integrity of the general public, including natural disasters as well as crimes. CASs may be usefully leveraged for monitoring and for the co-ordination of responses to dangerous situations.

A motivating example is a distributed surveillance system, where a collection of cameras are employed for intrusion detection, target identification and tracking, or behavioural analysis. The key problems with multi-camera configurations are co-ordination and control of individual, heterogeneous cameras such that they perform collaboratively [56]. The set of cameras should in other words behave as a CAS in order to react to local and global contingencies, and to provide distributed global functions [57].

Regulation in emergency situations is another task that can be best performed in a collective-adaptive manner. For instance, coordinating the movement of people in dangerous situations such as in overcrowded spaces requires timely, co-ordinated adaptation, since it is easy for individual and small-scale deviations to escalate into global chaos, leading to disaster. This problem is tackled by crowd engineering approaches, which are often founded on CAS paradigms [32], [58].

Environmental monitoring (discussed in more detail in Section 4.4) can also be considered from the perspective of public safety, as it supports timely identification and intervention in natural disasters.

However, safety-critical applications pose challenges for CASs related to verification and validation, real-time guarantees, and to the satisfaction of "hard" constraints that are difficult to provably satisfy in highly variable and possibly unknown scenarios [59]. This remains an area requiring significant research.

### 4.4  Ecology

The ecology domain offers a wide range of scenarios that can benefit from the application of collective approaches and strategies to better handle common needs. These scenarios span from environmental monitoring (for alerts management in case of wildfires or floods) to rubbish collection (in the cities as well as the seas) and the protection for endangered animal species. Such ecology-oriented systems aim to support the management of medium-to-large areas (*e.g.*, city, sea, national park) while by guaranteeing efficient operation during the nominal behaviour of systems (such as forest monitoring), and effective adaptation in case of constraints violations (such as hunting of protected animals) or problems detection (such as burning forests). Moreover most of the adaptation needs of these systems strongly require timely decision-making and co-ordinated and co-operative intervention. Different actors are involved (including police, firemen, rangers, navy, national parks, protected reserves, municipalities, environmental managers, and citizens) with diverse physical distribution and diverse responsibilities in a hierarchy of managerial roles (for example where a fire brigade unit intervenes on the spot while headquarters remotely coordinates several units).

With these parameters, it is easy to see the advantages of using CASs for the overall operation and dynamic adaptation of systems where heterogeneous entities can efficiently operate only by collectively sharing strategies and efforts [60]. Furthermore, ecology-oriented systems are

also included in the Internet of Things domain, since smart devices are often exploited in ecology scenarios, which further increases system complexity. For instance, smoke sensors, motion sensors, GPS, and the like can be used for monitoring tasks, and even though the IoT might be an enabler of ecology-oriented systems it might also emphasize the problems of heterogeneity and interoperability among the involved agents.

Interesting challenges arise when thinking about the roles played by different kinds of actor. Smart devices can both act as passive actors (for example smoke sensors in a monitoring phase) and as active actors during an adaptation phase (for example drones flying over a fire to measure its extent and guide rescue efforts). The same idea applies to humans. Another interesting challenge comes from how and to what extent human actors and devices co-operate, support, or complement each other in specific situations. In these contexts, CASs can represent a proper base from which to investigate dynamic role assignment and communication protocols, taking account of a wide range of diverse actors, constraints, peculiarities and uncertainties.

### 4.5 Operations Support

One final scenario comes from using CASs to create resilient, adaptive structures that can aid, sustain, protect, or replace, human activity in dynamic, unknown, and risky environments. Collaboration between humans and robots is generally useful in situated problem solving scenarios [61] such as city maintenance or waste collection.

A particularly relevant case is that of humanitarian relief operations [62], where people need to be rescued after environmental catastrophes such as hurricanes, tsunamis, floods, or during conflict. In these situations, teams of robots (unmanned ground or aerial vehicles) and possibly humans going on missions need to self-organize in an insecure environment. This may involve setting up *ad hoc* communication networks to replace damaged infrastructure, partially repairing physical structures or building temporary ones, and organizing field activity [63].

Another interesting case is exploration in unknown or "human off-limits" environments such as ocean depths or planetary surfaces. There, the CAS need to be resilient to unexpected hazards while preserving collected information and leveraging collective knowledge and physical possibilities to safely overcome obstacles.

## 5 CHALLENGES AND FUTURE AGENDA

Inevitably, CASs influence and are influenced by human behaviour. The same "classic" software system may yield very different results if adopted to different human communities, and this effect is exacerbated in the context of self-adaptive systems that would try to tailor themselves to the specific context, with final effects being hard to predict.

In fact, adaptive software systems interacting with a society of humans inevitably undergo a *co-adaptation*, as depicted in Figure 1 borrowing notation from control theory: while the adaptive part tries to steer the system towards the desired goal, the co-adaptation introduced by human behaviour may change the global effect. The two adaptations



Fig. 1. (Top) Adaptive software is usually conceptualized as steering the behaviour of a controlled system by sensing its outputs and acting on its inputs, in a manner familiar from control theory. (Bottom) When the whole system (including controller and controlled) operates in the context of a human-populated society, the system becomes more similar to a double-loop control structure with humans co-evolving and responding themselves to the system adaptation. Concurrent adaptation by users and software may result in hard-to-predict behaviours.

together are similar to a double-control loop structure [64], which can be very difficult to design and analyse for conflicts.

### 5.1 Designing CASs with and for humans

Human behaviour is currently often considered in CASs in a very simplistic way – often treated as part of the "unpredictability" that may affect the system and tackled as any other perturbation, i.e., by relying on the resilience and adaptivity of the system to compensate for unpredicted behavior. This approach is hardly sufficient for systems meant to enrich the society.

An attempt at better modeling human behaviour can be found in literature devoted to creating models of cognitive agents [52], where beliefs, desires, and intentions are considered as part of the human model. Such approaches allow designers to model (and mimic) phenomena such as fear contagion, opening the door to more in-depth design and verification of these aspects of CASs. More broadly, scaling up realism in the verification phase has the potential to retroactively impact the entire design process: one might be able to verify correct and coherent (with respect to human expectations) system adaptations even in case of human behaviour considered (arguably correctly) irrational, and as such (arguably wrongly) usually discarded entirely.

This is only a side of a multifaceted issue. Besides being able to model human behaviour more realistically, CASs must be able to explain themselves to their users. It is somewhat a case of mutual understanding: the system needs to understand society better to serve it, and society understands the system so as to harmoniously co-adapt

with it. Humans tend to resist changes that require effort to be understood unless they solve a pressing and evident problem, especially if they are very accustomed to "the way things are". Examples of such behaviour can be found in relation to traditional software: a study conducted on the user acceptance of a new interface for a common word processor found that, the higher the experience with the traditional interface (and hence the higher the experience with the tool), the higher the resistance to change [65]. The same issues appear in safety critical contexts such as healthcare. Longitudinal studies have found that resistance to the adoption of novel technologies is found in all groups of stakeholders, mirroring the complexity of the target organization [66]. Further work suggests that time is a key variable, as different individuals accept and adapt changes at different rates [67].

Adaptation cannot be allowed to preclude explanation. In fact, not providing clear and understandable justifications for a behaviour may trigger counter-reactions on the part of users, with detrimental effect on the global (cyber-physical) system behaviour. Ongoing work in self-explanation of adaptive systems is trying to tackle such issue [68]. In order to improve our understanding and capacity to model human and collective responses to changes in the technological environment, we suggest deeper collaboration among experts in collective adaptive systems and the social and psychological studies communities.

From a safety perspective, CASs operating in a society must provide mechanisms for monitoring to detect undesirable behaviours or responses. A key contribution in this context may be provided by exploring a closer integration with control theory, which might provide bounds on adaptativity and thereby exclude by design a whole class of potentially dangerous behaviours.

## 5.2 Evaluating CASs

The pervasive integration of CASs into human society make research into CAS systems engineering of utmost importance, from both a quantitave and qualitative perspective. Feng *et al.* [69] consider the problem of defining useful performance measure for CASs. The authors classify performance measures derived from probability distributions that can be extracted from a CAS model. Innes and Booher [70] propose a framework for evaluating consensus-building in complex adaptive systems. Loreti *et al.* [71] build a language based on stochastic process algebra for the specification and analysis of CASs. Vandin and Tribastone [72] review a number of techniques dealing with the problem of efficiently analyzing large-scale CASs for quantitative properties. From the perspective of correctness-by-construction, Viroli *et al.* [73] address the resilience of large-scale systems by formally proving the self-stabilization properties of basic distributed building blocks and their functional composition.

CASs provide collective behaviour by making use of shared resources. Privacy and ethical issues therefore have to become first-class citizens when designing and evaluating these systems. Hartswood *et al.* [74] discuss these issues in the context of peer profiling within CASs. They propose to return the control of data to users, supporting the principles of privacy and transparency. Other privacy-preserving mechanisms based on decentralized data ownership are proposed in the literature as a way to enhance privacy [75], [76]. However, in a world dominated by centralized platforms making profits out of the users' personal data, designers face a major socio-technical challenge in executing such a paradigm shift. The centralized/decentralized dichotomy affects the scalability of the proposed CAS solutions too, of course, and so centralization may contain within itself the seeds of a move towards more decentralized solutions on technical grounds [77].

## 6 CONCLUSIONS

The social implications of CASs already lie deep in the technology of the sharing economy, crowd-sourcing, crowd-funding and other domains. This is an indication that social aspects have a potential to change the way adaptive software systems are and will be designed. A wide variety of stakeholders are interested in building *social CAS* applications, ranging from public administrations to businesses, informed by the "data for good" and other movements. But nobody yet knows how to build these systems effectively.

In this article we presented a number of challenges that can be considered as research directions for the CAS field. This is especially pertinent when humans are involved in:

- designing CASs that consider the often unpredictable human behaviour;
- designing CASs where humans are part of the system;
- providing a support to evaluate CASs, both quantitatively and qualitatively; and
- addressing privacy, reliability, usability, observability, and liability in CASs.

Only when these aspects are addressed will Collective Adaptive Systems be able to fulfill their potential as the foundation for the next generation of cyber-physical systems.

## REFERENCES

[1] A. Ferscha, "Collective adaptive systems," in *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, ser. UbiComp/ISWC'15 Adjunct. New York, NY, USA: ACM, 2015, pp. 893–895.

[2] "FoCAS: Fundamentals of Collective Adaptive Systems - FP7 Coordination Action," 2013-2016.

[3] S. Kernbach, T. Schmickl, and J. Timmis, "Collective adaptive systems: Challenges beyond evolvability," 2011.

[4] C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. D. Caro, F. Ducatelle, T. S. Stirling, Á. Gutiérrez, L. M. Gambardella, and M. Dorigo, "Argos: A modular, multi-engine simulator for heterogeneous swarm robotics," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2011, San Francisco, CA, USA, September 25-30, 2011*, 2011, pp. 5027–5034.

[5] P. Levi and S. Kernbach, *Symbiotic Multi-Robot Organisms - Reliability, Adaptability, Evolution*, ser. Cognitive Systems Monographs. Springer, 2010, vol. 7.

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

[7] S. Dobson, S. Denazis, A. Fernández, D. Gaïti, E. Gelenbe, F. Massacci, P. Nixon, F. Saffre, N. Schmidt, and F. Zambonelli, "A survey of autonomic communications," *ACM Transactions on Autonomous and Adaptive Systems*, vol. 1, no. 2, pp. 223–259, December 2006. [Online]. Available: doi://10.1145/1186778.1186782

[8] A. Bucchiarone, M. De Sanctis, and A. Marconi, "Decentralized dynamic adaptation for service-based collective adaptive systems," in *Service-Oriented Computing - ICSOC 2016 Workshops - ASOCA, ISyCC, BSCI, and Satellite Events, Banff, AB, Canada, October 10-13, 2016. Revised Selected Papers*, 2016, pp. 5–20.

[9] G. Cabri, M. Puviani, and F. Zambonelli, "Towards a taxonomy of adaptive agent-based collaboration patterns for autonomic service ensembles," in *2011 International Conference on Collaboration Technologies and Systems, CTS 2011, Philadelphia, Pennsylvania, USA, May 23-27, 2011*, 2011, pp. 508–515.

[10] E. Lavinal, T. Desprats, and Y. Raynaud, "A generic multi-agent conceptual framework towards self-management," in *Management of Integrated End-to-End Communications and Services, 10th IEEE/IFIP Network Operations and Management Symposium, NOMS 2006, Vancouver, Canada, April 3-7, 2006. Proceedings*, 2006, pp. 394–403.

[11] J. H. Holland, "Studying complex adaptive systems," *J. Systems Science & Complexity*, vol. 19, no. 1, pp. 1–8, 2006.

[12] M. M. Hölzl, A. Rauschmayer, and M. Wirsing, "Software engineering for ensembles," in *Software-Intensive Systems and New Computing Paradigms - Challenges and Visions*, 2008, pp. 45–63.

[13] M. Wirsing, M. M. Hölzl, N. Koch, and P. Mayer, Eds., *Software Engineering for Collective Autonomic Systems - The ASCENS Approach*, ser. Lecture Notes in Computer Science. Springer, 2015, vol. 8998.

[14] V. Andrikopoulos, A. Bucchiarone, S. G. Saez, D. Karastoyanova, and C. A. Mezzina, "Towards modeling and execution of collective adaptive systems," in *Service-Oriented Computing - ICSOC 2013 Workshops - CCSA, CSB, PASCEB, SWESE, WESOA, and PhD Symposium, Berlin, Germany, December 2-5, 2013. Revised Selected Papers*, 2013, pp. 69–81.

[15] "ALLOW Ensembles. FP7 ICT-2011.9.10 - FET Proactive: Fundamentals of Collective Adaptive Systems (FOCAS)," 2013-2016.

[16] F. Zambonelli, A. Omicini, B. Anzengruber, G. Castelli, F. L. D. Angelis, G. D. M. Serugendo, S. Dobson, J. L. Fernandez-Marquez, A. Ferscha, M. Mamei, S. Mariani, A. Molesini, S. Montagna, J. Nieminen, D. Pianini, M. Risoldi, A. Rosi, G. Stevenson, M. Viroli, and J. Ye, "Developing pervasive multi-agent systems with nature-inspired coordination," *Pervasive and Mobile Computing*, vol. 17, pp. 236–252, feb 2015.

[17] F. Partners, "FoCAS Manifesto - A roadmap to the future of Collective Adaptive Systems," 2016, http://www.focas.eu/focas-manifesto.pdf.

[18] A. H. Bond and L. Gasser, Eds., *Readings in Distributed Artificial Intelligence*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.

[19] M. Wooldridge, *An introduction to multiagent systems*. John Wiley & Sons, 2009.

[20] G. Valetto, A. Bucchiarone, K. Geihs, M. Büscher, K. Petersen, A. Nowak, A. Rychwalska, J. V. Pitt, J. Shalhoub, F. Rossi, P. Silingardi, and P. Bernardeschi, "All together now: Collective intelligence for computer-supported collective action," in *2015 IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, SASO Workshops 2015, Cambridge, MA, USA, September 21-25, 2015*, 2015, pp. 13–18.

[21] K. L. Bellman, J. Botev, H. Hildmann, P. R. Lewis, S. Marsh, J. Pitt, I. Scholtes, and S. Tomforde, "Socially-sensitive systems design: Exploring social potential," *IEEE Technol. Soc. Mag.*, vol. 36, no. 3, pp. 72–80, 2017.

[22] S. Keller, R. Hausmann, L. Kressner, and A. Koenig, "An approach of a computerized planning assistant to the system design of collaborative robot installations," in *21st IEEE International Conference on Emerging Technologies and Factory Automation, ETFA 2016, Berlin, Germany, September 6-9, 2016*, 2016, pp. 1–4.

[23] N. Bredèche, E. Haasdijk, and A. Prieto, "Embodied evolution in collective robotics: A review," *Front. Robotics and AI*, vol. 2018, 2018.

[24] H. Ding, M. Schipper, and B. Matthias, "Collaborative behavior design of industrial robots for multiple human-robot collaboration," in *Proceedings of the 44th Internationel Symposium on Robotics,*

[25] A. Amanatiadis, K. Charalampous, I. Kostavelis, B. Birkicht, B. Andel, V. Meiser, C. Henschel, S. Baugh, M. Paul, R. May, and A. Gasteratos, "Autonomous vehicle emergency recovery tool: A cooperative robotic system for car extraction," *J. Field Robotics*, vol. 33, no. 8, pp. 1058–1086, 2016.

[26] T. Schmickl, R. Thenius, C. Möslinger, J. Timmis, A. M. Tyrrell, M. Read, J. A. Hilder, J. Halloy, A. Campo, C. Stefanini, L. Manfredi, S. Orofino, S. Kernbach, T. Dipper, and D. K. Sutantyo, "Cocoro - the self-aware underwater swarm," in *Fifth IEEE Conference on Self-Adaptive and Self-Organizing Systems, SASOW 2011, Ann Arbor, MI, USA, October 3-7, 2011, Workshops Proceedings*, 2011, pp. 120–126.

[27] M. D'Angelo, A. Napolitano, and M. Caporuscio, "Cyphef: a model-driven engineering framework for self-adaptive cyber-physical systems," in *Proceedings of the 40th International Conference on Software Engineering: Companion Proceeedings, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, 2018, pp. 101–104.

[28] G. Castelli, M. Mamei, A. Rosi, and F. Zambonelli, "How to develop pervasive social applications with the SAPERE middleware," *Computing and Informatics*, vol. 34, no. 1, pp. 185–209, 2015.

[29] D. Ingram, "Reconfigurable middleware for high availability sensor systems," in *Proceedings of the Third ACM International Conference on Distributed Event-Based Systems, DEBS 2009, Nashville, Tennessee, USA, July 6-9, 2009*, 2009.

[30] G. Salvaneschi, C. Ghezzi, and M. Pradella, "Contexterlang: introducing context-oriented programming in the actor model," in *Proceedings of the 11th International Conference on Aspect-oriented Software Development, AOSD 2012, Potsdam, Germany, March 25-30, 2012*, 2012, pp. 191–202.

[31] C. Cetina and V. Pelechano, "Variability in autonomic computing," in *Systems and Software Variability Management, Concepts, Tools and Experiences*, 2013, pp. 261–267.

[32] J. Beal, D. Pianini, and M. Viroli, "Aggregate programming for the Internet of Things," *Computer*, vol. 48, no. 9, pp. 22–30, sep 2015.

[33] F. Dressler, "Bio-inspired feedback loops for self-organized event detection in sanets," in *Self-Organizing Systems, Third International Workshop, IWSOS 2008, Vienna, Austria, December 10-12, 2008. Proceedings*, 2008, pp. 256–261.

[34] G. Stevenson, J. L. Fernandez-Marquez, S. Montagna, A. Rosi, J. Ye, A. Tchao, S. Dobson, G. D. M. Serugendo, and M. Viroli, "Towards situated awareness in urban networks: A bio-inspired approach," in *Sixth IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, SASOW 2012, Lyon, France, September 10-14, 2012*, 2012, pp. 223–228.

[35] E. D. Nitto, D. J. Dubois, and R. Mirandola, "On exploiting decentralized bio-inspired self-organization algorithms to develop real systems," in *2009 ICSE Workshop on Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2009, Vancouver, BC, Canada, May 18-19, 2009*, 2009, pp. 68–75.

[36] M. D'Angelo, M. Caporuscio, V. Grassi, and R. Mirandola, "Decentralized learning for self-adaptive qos-aware service assembly," *Future Generation Computer Systems*, vol. 108, pp. 210 – 227, 2020.

[37] M. Dorigo, "Editorial: Ten years of swarm intelligence," *Swarm Intelligence*, vol. 10, no. 4, pp. 245–246, Dec 2016.

[38] J. O. Kephart, A. Diaconescu, H. Giese, A. Robertsson, T. Abdelzaher, P. Lewis, A. Filieri, L. Esterle, and S. Frey, "Self-adaptation in collective self-aware computing systems," in *Self-Aware Computing Systems*. Springer, 2017, pp. 401–435.

[39] T. De Wolf and T. Holvoet, "Emergence versus self-organisation: Different concepts but promising when combined," in *International workshop on engineering self-organising applications*. Springer, 2004, pp. 1–15.

[40] F. Zambonelli, N. R. Jennings, and M. Wooldridge, "Multi-agent systems as computational organizations: the gaia methodology," in *Agent-oriented methodologies*. IGI Global, 2005, pp. 136–171.

[41] G. E. Mobus, M. C. Kalton *et al.*, *Principles of systems science*. Springer, 2015.

[42] N. Suri and G. Cabri, *Adaptive, dynamic, and resilient systems*. CRC Press, 2014.

[43] D. Pianini, R. Casadei, and M. Viroli, "Security in collective adaptive systems: a roadmap," in *2019 IEEE 4th International Workshops on Foundations and Applications of Self* Systems (FAS*W)*, 2019, to appear.

[44] A. Aldini, "Design and verification of trusted collective adaptive

systems," *ACM Trans. Model. Comput. Simul.*, vol. 28, no. 2, pp. 9:1–9:27, Feb. 2018.

[45] D. Pianini, G. Ciatto, R. Casadei, S. Mariani, M. Viroli, and A. Omicini, "Transparent protection of aggregate computations from byzantine behaviours via blockchain," in *Proceedings of the 4th EAI International Conference on Smart Objects and Technologies for Social Good, GOODTECHS 2018, Bologna, Italy, November 28-30, 2018*, 2018, pp. 271–276. [Online]. Available: https://doi.org/10.1145/3284869.3284870

[46] R. Casadei, A. Aldini, and M. Viroli, "Towards attack-resistant aggregate computing using trust mechanisms," *Sci. Comput. Program.*, vol. 167, pp. 114–137, 2018.

[47] M. Dikmen and C. Burns, "Trust in autonomous vehicles: The case of tesla autopilot and summon," in *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, Oct. 2017.

[48] B. Brown and E. Laurier, "The trouble with autopilots," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI 2017*. ACM Press, 2017.

[49] A. Rouvroy, Ed., *Law, human agency and autonomic computing : the philosophy of law meets the philosophy of technology*. Routledge, 2011.

[50] M. D'Angelo, S. Gerasimou, S. Ghahremani, J. Grohmann, I. Nunes, E. Pournaras, and S. Tomforde, "On Learning in Collective Self-Adaptive Systems: State of Practice and a 3D Framework," in *Proceedings of the 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '19. IEEE Press, 2019, p. 13–24.

[51] C. Baier and J.-P. Katoen, *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.

[52] T. Bosse, M. Hoogendoorn, M. C. A. Klein, J. Treur, C. N. van der Wal, and A. van Wissen, "Modelling collective decision making in groups and crowds: Integrating social contagion and interacting emotions, beliefs and intentions," *Autonomous Agents and Multi-Agent Systems*, vol. 27, no. 1, pp. 52–84, Jun. 2012.

[53] M. Weiser, "The computer for the 21st century," *Scientific American*, September 1991.

[54] A. Farinelli, N. Boscolo, E. Zanotto, and E. Pagello, "Advanced approaches for multi-robot coordination in logistic scenarios," *Robotics and Autonomous Systems*, vol. 90, pp. 34–44, 2017.

[55] T. Démare, C. Bertelle, A. Dutot, and L. Lévêque, "Modeling logistic systems with an agent-based model and dynamic graphs," *Journal of Transport Geography*, vol. 62, pp. 51–65, 2017.

[56] P. Natarajan, P. K. Atrey, and M. Kankanhalli, "Multi-camera coordination and control in surveillance systems: A survey," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 11, no. 4, pp. 57:1–57:30, Jun. 2015.

[57] D. Pianini, S. Dobson, and M. Viroli, "Self-stabilising target counting in wireless sensor networks using Euler integration," in *11th IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2017, Tucson, AZ, USA, September 18-22, 2017*. IEEE Computer Society, 2017, pp. 11–20.

[58] C. Borean, R. Giannantonio, M. Mamei, D. Mana, A. Sassi, and F. Zambonelli, "Urban crowd steering: An overview," in *Internet and Distributed Computing Systems - 8th International Conference, IDCS 2015, Windsor, UK, September 2-4, 2015. Proceedings*, ser. Lecture Notes in Computer Science, G. D. Fatta, G. Fortino, W. Li, M. Pathan, F. T. Stahl, and A. Guerrieri, Eds., vol. 9258. Springer, 2015, pp. 143–154.

[59] B. Schmerl, J. Andersson, T. Vogel, M. B. Cohen, C. M. Rubira, Y. Brun, A. Gorla, F. Zambonelli, and L. Baresi, "Challenges in composing and decomposing assurances for self-adaptive systems," in *Software Engineering for Self-Adaptive Systems III. Assurances*. Springer, 2017, pp. 64–89.

[60] Ö. Bodin, "Collaborative environmental governance: Achieving collective action in social-ecological systems," *Science*, vol. 357, no. 6352, p. eaan1114, 2017.

[61] R. Casadei, C. Tsigkanos, M. Viroli, and S. Dustdar, "Engineering resilient collaborative edge-enabled iot," in *2019 IEEE International Conference on Services Computing, SCC 2019, Milan, Italy, July 8-13, 2019*, 2019, pp. 36–45. [Online]. Available: https://doi.org/10.1109/SCC.2019.00019

[62] R. Casadei and M. Viroli, "Collective abstractions and platforms for large-scale self-adaptive iot," in *2018 IEEE 3rd International Workshops on Foundations and Applications of Self* Systems (FAS*W), Trento, Italy, September 3-7, 2018*, 2018, pp. 106–111. [Online]. Available: https://doi.org/10.1109/FAS-W.2018.00033

[63] J. Beal, K. Usbeck, J. P. Loyall, M. Rowe, and J. M. Metzler, "Adaptive opportunistic airborne sensor sharing," *TAAS*, vol. 13, no. 1, pp. 6:1–6:29, 2018.

[64] F.-Y. Yang, C.-L. Chen, and M.-C. Tsai, "A double-loop control structure for tracking control and disturbance attenuation," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 204–209, 2014.

[65] M. Dostál, "User acceptance of the microsoft ribbon user interface," in *Proceedings of the 9th WSEAS International Conference on Data Networks, Communications, Computers*, ser. DNCOCO'10. Stevens Point, Wisconsin, USA: World Scientific and Engineering Academy and Society (WSEAS), 2010, pp. 143–149.

[66] E. R. Nilsen, J. Dugstad, H. Eide, M. K. Gullslett, and T. Eide, "Exploring resistance to implementation of welfare technology in municipal healthcare services – a longitudinal case study," *BMC Health Services Research*, vol. 16, no. 1, Nov. 2016.

[67] S. Safi, T. Thiessen, and K. J. Schmailzl, "Acceptance and resistance of new digital technologies in medicine: Qualitative study," *JMIR Research Protocols*, vol. 7, no. 12, p. e11072, Dec. 2018.

[68] K. Welsh, N. Bencomo, P. Sawyer, and J. Whittle, "Self-explanation in adaptive systems based on runtime goal-based models," in *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2014, pp. 122–145.

[69] C. Feng, M. Gribaudo, and J. Hillston, "Performance analysis of collective adaptive behaviour in time and space," *Electr. Notes Theor. Comput. Sci.*, vol. 318, pp. 53–68, 2015.

[70] J. E. Innes and D. E. Booher, "Consensus building and complex adaptive systems," *Journal of the American Planning Association*, vol. 65, no. 4, pp. 412–423, 1999.

[71] M. Loreti and J. Hillston, "Modelling and analysis of collective adaptive systems with carma and its tools," in *International School on Formal Methods for the Design of Computer, Communication and Software Systems*. Springer, 2016, pp. 83–119.

[72] A. Vandin and M. Tribastone, *Quantitative Abstractions for Collective Adaptive Systems*. Cham: Springer International Publishing, 2016, pp. 202–232.

[73] M. Viroli, G. Audrito, J. Beal, F. Damiani, and D. Pianini, "Engineering resilient collective adaptive systems by self-stabilisation," *ACM Transaction on Modelling and Computer Simulation*, vol. 28, no. 2, pp. 16:1–16:28, Mar. 2018.

[74] M. Hartswood, M. Jirotka, R. Chenu-Abente, A. Hume, F. Giunchiglia, L. A. Martucci, and S. Fischer-Hübner, "Privacy for peer profiling in collective adaptive systems," in *Privacy and Identity Management for the Future Internet in the Age of Globalisation*, J. Camenisch, S. Fischer-Hübner, and M. Hansen, Eds. Cham: Springer International Publishing, 2015, pp. 237–252.

[75] G. Zyskind, O. Nathan, and A. . Pentland, "Decentralizing privacy: Using blockchain to protect personal data," in *2015 IEEE Security and Privacy Workshops*, May 2015, pp. 180–184.

[76] E. Pournaras, J. Nikolic, P. Velásquez, M. Trovati, N. Bessis, and D. Helbing, "Self-regulatory information sharing in participatory social sensing," *EPJ Data Science*, vol. 5, no. 1, p. 14, Apr 2016.

[77] R. de Lemos *et al.*, *Software Engineering for Self-Adaptive Systems: A Second Research Roadmap*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 1–32.