

This is the peer reviewed version of the following article:

An Argumentation-based Perspective over the Social IoT / Lippi, Marco; Mamei, Marco; Mariani, Stefano; Zambonelli, Franco. - In: IEEE INTERNET OF THINGS JOURNAL. - ISSN 2327-4662. - 5:4(2018), pp. 2537-2547. [10.1109/JIOT.2017.2775047]

Terms of use:

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.

02/05/2024 12:51

An Argumentation-based Perspective over the Social IoT

Marco Lippi, Marco Mamei, Stefano Mariani, and Franco Zambonelli, *IEEE Fellow*

Abstract—The crucial role played by social interactions between smart objects in the Internet of Things is being rapidly recognized by the Social Internet of Things (SIoT) vision. In this paper, we build upon the recently introduced vision of Speaking Objects – “things” interacting through argumentation – to show how different forms of human dialogue naturally fit cooperation and coordination requirements of the SIoT. In particular, we show how speaking objects can exchange arguments in order to seek for information, negotiate over an issue, persuade others, deliberate actions, and so on, namely, striving to reach consensus about the state of affairs and their goals. In this context, we illustrate how argumentation naturally enables such a form of conversational coordination through practical examples and a case study scenario.

Index Terms—Speaking Objects, Argumentation, Internet of Things, Social IoT, Coordination

I. INTRODUCTION

In recent years, research efforts in the *Internet of Things* (IoT) landscape started focusing on methodologies and frameworks for the development of smart objects capable of establishing *social relationships* with each other and with humans, seamlessly. This novel paradigm, usually referred to as the *Social Internet of Things* (SIoT) [1], revisits the classic architecture of the IoT where sensors and actuators have each a single responsibility: gathering data, and reacting to commands, respectively. In the SIoT, devices have the additional capability of performing some form of social behaviour: leveraging knowledge sharing, service composition, trust, and goal-oriented computing [2]. The key idea behind the SIoT is in fact to imitate the dynamics of human interaction within social networks and turning *smart* devices into *social* devices, as clearly stated by Atzori et al. [1]:

“[...] these patterns of interaction among human beings are directly applicable to possible social behaviors of typical objects that implement pervasive applications.”

Several enabling technologies are increasingly contributing to realize the SIoT vision, such as machine learning [3], [4], [5], computer vision [6], [7], ontologies [8] and commonsense reasoning [9], human-computer interfaces [10], among the many. Devices in the IoT arena are thus becoming smarter, by becoming able to perform operations such as processing data on chip so as to produce higher-level information, interacting

with other devices and humans as well through natural language processing and generation, and autonomously deciding their own course of actions towards the achievement of their (individual or collectively shared) goals. Following this vision, in [11] we introduced “speaking” and “hearing” objects as smart devices featuring all of the aforementioned properties, whose distinguishing feature is that of *arguing* about their beliefs, goals, and decisions.

Argumentation, indeed, is one of the most natural expression of humans’ capability to interact through *dialogue*: people continuously engage in conversations to persuade, negotiate, share ideas, motivate their choices, and so on [12]. In a similar way, we expect speaking and hearing objects in the SIoT scenario to communicate and *argue* with each other, so as to find an agreement on the perceived state of the affairs and to share decisions on how to act.

Exploiting argumentation within this domain potentially presents several advantages (better discussed in Section IV), among which:

- interaction and communication protocols may arise in a decentralized way by *emergence* [13], thus overcoming the limitations of traditional rule-based systems that require *a priori* encoding of coordination laws—making it hard to adapt their behaviour in unexpected situations
- decision making based on argumentation, leveraging declarative approaches, make decisions amenable of *interpretation*, so that any action may be easily explained and justified by a chain of arguments both to human users or supervisor systems
- humans-in-the-loop may benefit of *natural language* interaction with the system, being natural language generation well-supported by argumentation graphs [14]
- *reliability* and *security* are naturally supported, despite uncertainty of perceptions and system openness. Argumentation in fact, on the one hand enables reaching consensus despite discrepancies in measured metrics by guaranteeing that only well-backed claims win a debate, on the other hand may be used to spot malicious behaviours by proof-checking false claims

Building on the work in [11], in this paper we move the discussion about conversational coordination in the Speaking Objects vision one step forward, by (i) framing the different kinds of *social relationships* that can be established between speaking and hearing objects following [1]; (ii) grounding social interactions in the *kinds of dialogue* categorized by Walton and Krabbe [15]; and (iii) exemplifying how dialogues between speaking and hearing objects expressed in natural

Authors are with the Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, via Amendola 2, Reggio Emilia, Italy, email: {name.surname}@unimore.it.

Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

language could be turned into formal argumentation graphs. Accordingly, conversations may arise, i.e., due to the need of obtaining some missing information required to take a decision, or to disambiguate a contradictory situation. The resulting dialogues may assume different specific forms, depending on a number of factors such as the *goal* of the initiating agent, that of the responder one(s), that of the dialogue itself, and the kind of information subject to argumentation. For instance, when an agreement over a state of the affairs is sought for, *persuasive* and *negotiation* dialogues are likely to occur. Instead, a *deliberation* dialogue is likely to emerge after consensus is reached among the participants, and then a shared goal (or plan to act) is to be agreed upon.

The remainder of the paper is organised as follows. Section II summarizes the Speaking Objects vision. Section III discusses in detail the forms of social interaction that this novel setting envisions, also with the help of several examples and of a case study scenario. Section IV highlights challenges and opportunities for the proposed vision, whereas Section V details related work. Finally, Section VI concludes the paper.

II. SPEAKING OBJECTS

In traditional IoT frameworks, *smart objects* are typically referred to as physical objects or places equipped with some ICT device that allows them to communicate with the environment. Two kinds of smart objects are typically distinguished: *sensors*, which collect information by observing the world to continuously produce data streams, and *actuators*, which are capable of performing specific actions in their environment when they are instructed with (often hardwired) commands.

Recent advancements in the development of ICT devices allow to embed in them more and more computational power with relatively cheap cost. In parallel, machine learning techniques are also steadily improving, so that incorporating intelligent components within smart objects, with the goal of processing data and producing novel information *on chip*, is now reality [16], [7]. For instance, cameras can now accurately distinguish different objects within a scene, and more general purpose devices such as smartphones are capable of natural language processing and generation.

Indeed, the time is ripe for a further evolution of smart objects, which will soon acquire the capabilities for recognizing complex situations and higher-level activities. This implies the ability of performing some reasoning tasks, i.e. to find causal relationships between activity patterns or elementary situations, and thus the skill for making *assertions* about what is happening around. Recent methodologies such as those coming from deep learning [6], [7] or statistical relational learning [17] will likely play a major role within this context. Also knowledge representation and *commonsense reasoning* represent essential aspects for the creation of such high-level assertions. For instance, a camera should be able to assemble an assertion such as “A bike just crossed the intersection with a red traffic light” even without contacting an external server.

We call this kind of smart sensors, capable of producing articulate, high-level granules of information *speaking objects* [11]. Complementarily, actuator objects will likely

become able to understand *goals*, both individual and collective ones, and consequently behave so as to achieve them. Actions should then be undertaken by adopting a *goal-oriented* attitude, which, in the case of collective goals, may lead to either competition or (opportunistic) cooperation of multiple objects that agree on a shared *state of the affairs*—either currently recognized or to be achieved. We name these novel kind of smart actuators *hearing objects*. Interaction among speaking and hearing objects will be enabled by the use of dedicated ontologies, the injection of domain knowledge, and the integration of context-aware mechanisms [8].

IoT scenarios are intrinsically highly dynamic, typically featuring many devices which enter and leave the system at will and interact in many different ways. Also, interactions occur at different levels and in multiple forms, ranging from data fusion to knowledge sharing and decentralized coordination. On purpose, the SIoT vision [1] is gaining momentum, since it is especially focussed with empowering devices with social capabilities. The issues involved in coordinating a set of distributed autonomous devices (i.e., a multi-agent system) has been studied extensively [18]. A variety of different algorithms have been developed for the synchronization and coordination of the activities of distributed processes (i.e. distributed *mutual exclusion* to gain access to shared resources).

The key idea underlying the Speaking Objects approach is to exploit *argumentation* as a common basis for interaction amongst speaking and hearing objects. Argumentation [19] is an interdisciplinary research field, that lies at the intersection between logic, philosophy, social and political sciences, psychology, computational linguistics, and artificial intelligence. The aim of argumentation is to study dialogues and debates, trying to understand their dynamics, and to analyze the reasoning process that typically drives their evolution and conclusion. In particular, *computational argumentation* proposes to exploit computational models and techniques to automatically describe, construct, and analyze arguments and their interactions. This research area has recently been receiving growing attention [12], [20], and it is seen as one of the most promising directions for the study of social interactions through natural language.

As already mentioned, exploiting argumentation for coordination and cooperation presents several advantages, especially regarding interpretability of decision making, reliability of computations, adaptiveness of coordination, robustness w.r.t. contingencies, and trustability by human users, which are deeply discussed, along with the many challenges implied, in Section IV.

III. ARGUMENTS & CONVERSATIONS

In a Speaking Objects setting, coordination between speaking and hearing objects naturally assumes the form of a distributed multi-party *conversation*, or *dialogue* [21]. We define a conversation as a session of interaction between an ensemble of distributed components aimed at supporting achievement of a goal [15], [21]—either shared, or pursued by an individual. Conversations take place by having speaking and hearing objects exchange assertions about the current or

desirable state of the affairs, which can be contradicted or strengthened by others engaging in the conversation.

The dynamics of interaction just described may be well-supported by *argumentation-based coordination* [22], that is, the practice of letting an ensemble of autonomous agents figure out themselves the interaction protocol to adopt in a particular situation, instead of super-imposing some prescribed coordination rules [23], by arguing in compliance with a shared argumentation framework [24]—defining the interaction protocol and coordination rules at the meta-level [11].

In the following subsections we expand our previous work [11] – which was focused on the whole architectural organization of the speaking objects vision – by analyzing more systematically how the kind of dialogues identified by Walton and Krabbe [15], [25] (described in Subsection III-B) may apply in a few real-world, IoT-related scenarios (Subsection III-C). As a necessary premise, we describe argumentation frameworks [24], [26] in Subsection III-A. Then, Subsection III-D will further detail in a more formal way a specific IoT case study among those described.

A. Argumentation Frameworks

Argumentation is the discipline of studying how humans converse, debate, and reason [26]. Generally speaking, an argument can be defined as an assertion, or a series of assertions, regarding a certain matter of discussion. The investigation of argumentation using computer science, and especially natural language processing, led to the development of the field of computational argumentation.

Dung [24] formally defined an argumentation framework (AF) as a pair $AF = \langle AR, Attacks \rangle$ where AR is a set of arguments, and $Attacks$ is a binary relation on arguments, that is $Attacks \subseteq AR \times AR$. The $Attacks$ relation is clearly not symmetric (if $a \in AR$ attacks $b \in AR$, not necessarily b attacks a), thus the resulting argumentation framework can be easily mapped to a directed graph, where nodes represent arguments, and edges represent attacking relations. Figure 1 shows an example of such a graph, including three arguments and four attack relations, inspired by a simple IoT setting.

Such a model represents an *abstract* argumentation framework, since it does not require any assumption on the internal structure of each argument, which is instead considered as an atomic entity. This way it is possible to study which arguments in a graph can be justified, according to the underlying semantics of attack relations.

Models that instead take into account what is *inside* an argument fall in the category of *structured* argumentation frameworks [27]. Such frameworks can exploit a variety of different argument models, that have been proposed throughout the years. An exhaustive analysis of these models goes beyond the scope of this paper, and we refer to [28] for a comprehensive survey. In our setting, we adopt the simple model in which an argument consists of a *claim* possibly linked to one or more *premises* through a relation called *support* [26]. A claim can be seen as the conclusion of an argument, or a thesis to be justified, whereas a premise is a grounded fact that is used as a supporting reason of a claim.




Fig. 1: A graphical depiction of Dung’s argumentation framework, grounded on an IoT setting. Nodes in the graph represent arguments, whereas edges represent attack relations.

Argumentation frameworks have also been extended by introducing the concept of *strength* of arguments: it is worth mentioning weighted argumentation systems [29] and value-based argumentation frameworks [30], which attach real-valued weights or discrete preference values to attack links, respectively. This kind of extension is meant to better deal with those *open* and *dynamic* scenarios in which there is the need to compare arguments despite their relevance be changing over time or when new information is available—as typically happens in IoT-like applications.

B. Dialogue Types

Arguments, their internal components, and their support and attack relationships are what constitute dialogues, or conversations. Walton and Krabbe [15] propose a classification of conversations, following a pragmatic approach based on (i) the actual *situation* motivating the need for initiating a conversation session, (ii) the *goal* of the initiating participant, and (iii) the goal of the dialogue *itself*—which may be in conflict with that of the participants (which, in turn, not always share the same goal) [15], [25]. The classification identifies 7 *dialogue types*¹:

- *Information Seeking*. In this kind of dialogue, initiators of the conversation session are in need of some information, and they believe to know *who* amongst the other participants could provide it, thus start to ask questions.
- *Inquiry*. In this case, initiators and responders collaborate in *collectively* answering questions meant to acquire information which *no one* believes some other participant individually knows. The information to seek can be interpreted as *evidence* to a proof (information supporting other information), thus participants are engaged to either prove (provide sound information) or disprove it (provide contradicting information).
- *Discovery*. Discovery dialogue is similar to inquiry but differs in a fundamental way: in the latter, the information sought for is agreed upon at the beginning of the dialogue, whereas in the former it *emerges* during the course of the dialogue itself. According to the proof-theoretic interpretation, whereas in inquiry participants seek for evidence, in discovery participants seek for the best *hypothesis* (and accordingly strive to defend it).

¹*Discovery* dialogue has been actually added by McBurney and Parsons [31] and later recognised by Walton himself [25].

- *Persuasion*. In this kind of dialogue the initiator wants to *convince* a participant to believe in an opinion, to adopt a standpoint, to accept a fact (s)he does not currently hold. The responder has a conflicting belief, and may not share the initiator's goal of persuading the other participant—thus simply refuting each persuasion attempt.
- *Negotiation*. Participants of a negotiation dialogue *bargain* over an issue, i.e., the division of some scarce resource, each pursuing a goal which is not necessarily aligned with the goal of the collective, nor of the dialogue session—i.e., maximizing own satisfaction. A negotiation session converges when a resolution of the issue which is acceptable by all participants is found, fails otherwise—where acceptability may be defined based upon each participant's individual goal, the goal of the collective (the system), the goal of the dialogue, or any combination of the three.
- *Deliberation*. In this kind of dialogue a decision should be *collectively* taken about what *course of actions* to adopt given a situation. The participants may either share a common set of goals, and decide how to collectively act to reach the shared goals – thus what course of actions each of them should commit to – or share the willingness to reach a consensus about whether they have shared intentions.
- *Eristic*. Participants of a eristic dialogue argue as a substitution to physical *fight*—this dialogue is mentioned for completeness solely, but will not be further discussed.

Apart from the eristic conversation, each of the others fits well some specific purpose within the general issue of coordinating distributed components in IoT scenarios.

It is worth emphasizing that the kind of dialogue-based interactions just described can be naturally placed within the categorization of social relationships proposed by Atzori et al. for the SIoT vision [1]:

- the *social* object relationship captures the opportunistic interactions between personal devices stemming from their owners' relationships (i.e., friends going out together or colleagues working in the same place). Thus, discovery, information seeking, or inquiry dialogues seems particularly relevant here
- the *co-work* object relationship captures sporadic interactions but between any kind of smart object and caused by application requirements (i.e., sensors and actuators collaborating to provide an application functionality). Here, any dialogue type mentioned above fits
- the *co-location* relationship captures stable relationships amongst objects placed in proximity, regardless of whether they share or not any purpose w.r.t. the application functionalities. Provided that, instead, within the Speaking Objects vision some shared application goal should exist, dialogues such as deliberation and discovery can easily be thought of as being based on some notion of locality. For instance, when heterogeneous sensors close to each other contribute to shape a partial view of the surrounding world, or when actuators nearby jointly plan a shared course of actions to affect a local portion of the

environment

- the *ownership* relationship covers all the objects belonging to the same user, thus, it can easily capture any of the dialogue types considered in the Speaking Objects vision

The next subsection provides practical examples of how natural it is to decline the interaction protocols typically found in IoT applications within the reach of our Speaking Objects vision—that is, as argumentation-based dialogues.

C. Arguments & Dialogues in the Real World

We take as a reference two prominent examples of IoT-enabled application scenarios:

- *Assisted living*. We imagine that a smart home, as well as its inhabitants, are equipped with an array of speaking and hearing objects (sensors and actuators) meant to support them in their everyday life activities to ensure safety and promote healthy ageing—i.e., cameras, smoke detectors, the A/C system, automatic windows and curtains, self-driving wheel-chairs [32], the inhabitants' smartphones, smart TVs, RFIDs [33], and wearable medical devices.
- *Traffic management*. We imagine a set of crossroads (i.e., those belonging to the same block) and the approaching self-driving cars to feature speaking and hearing objects meant to autonomously manage traffic safely and efficiently—i.e., cameras, smart traffic lights, the whole array of sensors and actuators of a self-driving car.

There, conversational coordination amongst speaking and hearing objects is likely to happen as a *combination* of dialogue types, in which, for instance, a dialogue pursues an application goal by embedding multiple kinds of dialogue for achieving specific sub-goals.

In this paper we describe all the dialogues involving speaking objects using natural language. On the one hand, this notably improves paper readability, as introducing a formal schema for these dialogues (i.e., a process algebraic notation) would be cumbersome and distract from the main point we are trying to deliver—that is, the practice of letting agents figure out themselves the interaction protocol to adopt in a particular situation by arguing in compliance with a shared argumentation framework. On the other hand, natural language could be actually used in situations where human users are involved, whereas other, more formal and machine-friendly dialogue representations (i.e., based on interoperable standards such as XML and XML Schema) are suitable for machine-to-machine argumentation.

1) *Traffic management*: In the traffic management case, consider vehicle *A* approaching traffic light *T* from the east, as depicted in Fig. 2. In order not to wait at the crossroad, *A* asks *T* how much time is left for the green light, and, in case time is running out, whether *T* could extend green light duration so as to allow *A* to cross.

The above dialogue may unfold as follows (*i* stands for *inquiry* dialogues, *n* stands for *negotiation* dialogues):




Fig. 2: The crossroad in the traffic management scenario, depicting approaching cars A, B , traffic light T , and cameras C_i , where $i \in 1, \dots, N$.

- i_1) A : Hi T , how long will the light remain green?
- i_2) T : Hello A , it would last 30 seconds.
- n_1) A : Could you keep the green on 30 seconds more? I'm a bit late to work.

Initially (utterances $i_{1,2}$), it is just an *inquiry* dialogue between A and T . Then (utterance n_1) A starts a *negotiation* phase to convince T to let it pass². In order to answer, T needs to know whether some other vehicles are approaching the crossroad soon, thus it asks traffic cameras C_i , where $i \in 1, \dots, N$ (denoted as a collective) if any of them is detecting approaching vehicles (*is* stands for *information seeking* dialogues):

- is_1) T : Is any vehicle reaching the crossroad within a minute?
- is_2) C_i : Yes, one vehicle approaching from south in 40 seconds.

This is a typical *information seeking* dialogue, since it is not known a priori who will respond—differently from the opening inquiry dialogue. Camera C_i delivers the information that vehicle B is approaching the crossroad in about 40 seconds (utterance is_2). Now then, T may attempt to convince B to let A traverse the crossroad in advance. Here, after an inquiry phase (utterances $i_{1,2}$), a *persuasion* dialogue is initiated by T (utterances $p_{1,4}$).

- i_1) T : Hi B , are you crossing straight?
- i_2) B : Hello T , yes I am.
- p_1) T : Do you mind waiting for one minute?
- p_2) B : Sorry, I need to cross now for reaching home soon.
- p_3) T : I see. Could you turn right and reach home anyway? It's just to help another vehicle.
- p_4) B : Sure, that will take about the same time.

²Notice that this dialogue cannot be defined as *persuasive* yet, since it is not known whether T has conflicting opinions with respect to A .

The conversation session terminates (successfully) when B agrees to turn right, so that A can go straight without any of the two being forced to wait.

The above conversation is a clear example of a near-future, real-world scenario envisioned by the Speaking Objects vision, where the coordination of the interacting objects through argumentation leads to the achievement of a state of the affairs in which all participants are satisfied. The overall conversational coordination process thus appears as a sort of complex negotiation protocol embedding information inquiry and seeking, and persuasion. It is worth noting that also a *deliberation* phase may be detected, because with utterance p_3 T actually suggests a plan to B to bring about its own goal, which is alternative to B 's own plan (utterance p_2), and is accepted by B with utterance p_4 .

2) *Assisted living*: In the assisted living scenario, consider that one of the inhabitants, Walter, is suffering of chronic obstructive pulmonary disease (COPD), thus he may undergo unexpected respiratory/asthma crisis with different degrees of severity, up to requiring immediate hospitalization. Assume now that one of those crisis happens when the elderly person is temporarily home alone.

First of all, the harmful situation has to be recognized by the speaking objects (the sensors). Although this may seem at first a trivial task, for instance simply involving the smartphone microphone hearing a coughing sound, if the smart home is to *pro-actively* and *autonomously* assist Walter, then more information is needed so as to take appropriate action. In fact, the coughing episode may be, for instance, a food choking event actually, not a respiratory crisis due to COPD, thus requiring a very different intervention.

To disambiguate the situation, an *information seeking* dialogue is required, which involves, for instance, the smartphone S , the wristband W monitoring heartbeat, and the indoor surveillance cameras C_i , with $i \in 1, \dots, N$ (denoted as a collective):

- is_1) S : Everybody listen, I'm detecting a clear coughing sound, should we worry?
- is_2) W : Probably yes, since the heart beat is accelerating steadily!
- is_3) C_i : Wait a minute, Walter was just eating peanuts a moment ago, could it be food choking instead?
- is_4) C_j : I can confirm, he is now punching his own breast and trying to puke.

Whichever the cause behind the coughing episode just detected is, an appropriate action has to be taken to help Walter. This naturally leads towards a *deliberation dialogue*, which the smartphone, the cameras (also external ones, collectively denoted as $C_{ek}, k \in 1, \dots, M$), Walter's smart watch Wa , and his self-driving wheel-chair Ch may participate to (*d* stands for *deliberation* dialogues):

- d_1) S : Status report, please.
 d_2) C_j : He is not puking yet.
 d_3) S : I text Linda [Walter's daughter] to get here.
 d_4) Wa : You may try, but she's at the dentist now thus she may not come in time!
 d_5) Ch : Go for it, I'll take Walter out in the meanwhile.
 d_6) Ce_k : Agreed, there's a neighbour mowing the lawn just across the road.
 d_7) S : I'll also call the ambulance.

Notice that we made every speaking and hearing object proactive in sharing information just to avoid unnecessary complications, but the same conversational coordination session may as well embed a *discovery* dialogue amongst speaking objects so as to find arguments supporting a claim (the actual shared plan, in this case). For instance, the information conveyed by utterances d_2 , d_4 and d_6 are either *supporting* (d_6) or *attacking* (d_4) each other, and jointly constitute a valid argument for applying the plan of taking Walter out while texting Linda—with the goal of getting help from either her or the neighbour.

D. Towards Conversational Coordination

We now focus on the *assisted living* scenario to better describe the way in which argumentation may support coordination in the Speaking Objects scenario, by formalizing (part of) the dialogues therein discussed according to a combination of abstract and structured argumentation frameworks³ described in Subsection III-A. In particular, we provide graphs depicting *attack* and *support* relationships amongst arguments, with a twofold aim: on the one hand, to show that complex coordination requirements motivated by real-world applications may be naturally expressed and dealt with by adopting an argumentation-based standpoint (which also further motivates our approach, ultimately); on the other hand, to discuss the feasibility and the benefits of the proposed approach with respect to the SIoT landscape and the currently available technologies (Section IV).

Figures 3–6 depict the *information seeking* dialogue where Walter's smartphone, his wristband, and some indoor cameras debate whether the coughing episode Walter is suffering right now is due to his COPD or to the peanuts he just ate. Despite the apparent simplicity of the dialogue, having just 4 utterances, it conveys 14 argument components and 11 relations (all of kind support)—at least⁴.

For instance, Fig. 3 and Fig. 4 both represent an argument having as a claim the fact that Walter is having a respiratory crisis. In both cases, the claim is supported by two distinct kinds of premises: observations of the *current situation*, such as the fact that Walter is coughing, or that his heartbeat is accelerating, and *background knowledge*, such as the evidence that cough and fast heartbeat are symptoms of COPD. Fig. 5 and Fig. 6 propose similar arguments for the thesis that Walter

³The aim is to balance the tradeoff between unambiguity of formalization and accessibility to a wide audience.

⁴Some arguments may be further expanded and a few relations have been omitted with the usual aim of balancing the tradeoff between unambiguity of formalisation and accessibility to a wide audience.




Fig. 3: Arguments (circles) and support relations (dashed arrows) in utterance is_1 from Subsection III-C2.




Fig. 4: Arguments and support relations in utterance is_2 from Subsection III-C2.




Fig. 5: Arguments and support relations in utterance is_3 from Subsection III-C2.

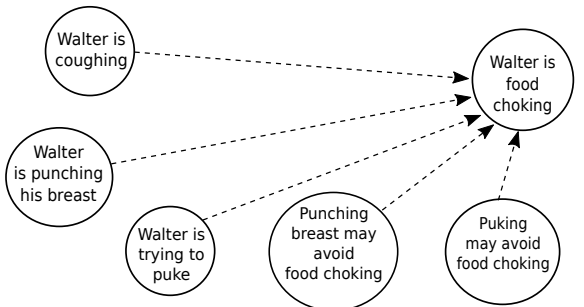


Fig. 6: Arguments and support relations in utterance is_4 from Subsection III-C2.

is instead food choking, this time supported by premises such as that he is trying to puke by punching his own breast, which




Fig. 7: Arguments and *attack* relations (solid arrows) in the *information seeking* dialogue ($is_{1,4}$) from Subsection III-C2.

is a known technique to avoid food choking.

We want to emphasize that similar representations are possible solely by adopting the *structured* standpoint over argumentation, because only there arguments can be further decomposed into *claims* and *premises*—linked by a support relation. From the perspective of *abstract* argumentation instead, a possible representation is that of Fig. 7, where the arguments of respiratory crisis and that of food choking are reciprocally attacking one another. In Fig. 8 we let the reader appreciate how structured argumentation easily enables to “zoom-in” arguments and navigate amongst supporting and conflicting claims and premises: there, in fact, it is apparent how different premises support conflicting claims.

Regardless of whether abstract or structured argumentation is adopted, *weights* or *preference values* could be attached to relations so as to unambiguously and automatically decide who is winning the debate. For instance, in Fig. 8 we naively decorated all support relations with the same weight (equal to 1), and by attaching a weight to attacks simply computed as the sum of the weights of supporting premises. We also added a third argument (asserting that Walter might just be healthy, since it has been running) to highlight that a dynamic evolution of the argument graph is possible, and likely. Thus, in this case it is clear that the winner argument is *a*, that is that Walter is choking due to peanuts ingestion. Alternatively, one could assign a different weight to premises based, i.e., on their nature (background knowledge vs. current situation), so as to, for instance, give more relevance to a priori knowledge, which can be considered well agreed upon and more stable, w.r.t. observations, which may be prone to errors and defeated by new perceptions.

Finally, it is worth emphasizing that weights (or preference values, or any other quantitative label) may be adjusted *dynamically*, as time flows or novel information is acquired, especially in those application domains featuring openness, uncertainty, and incomplete information—as the (S)IoT typically is.

IV. CHALLENGES AND BENEFITS

The Speaking Objects vision has yet to address several challenging tasks in order to be fully realized.

Consensus on Argumentation Framework: First of all, argumentation requires that all participants in a debate to be resolved automatically either *agree* on the *argumentation-based framework* deciding who wins the dispute, or abide by a well respected authority acting as *dispute arbiter*. For instance, in the traffic management scenario (Subsection III-C1) implementing the latter solution is straightforward since the traffic light(s) may naturally play the role of the arbiter

deciding over arguing cars, whereas in the assisted living one (Subsection III-C2) it appears less obvious who could be the arbiter, thus a more decentralised solution where sensors, actuators, and personal devices share consensus upon the argumentation rules may be preferred. Both solutions have some cons, such as being vulnerable to a single point of failure in the former case, or requiring additional communications to dynamically reach consensus in the latter, but they can effectively be dealt with by exploiting already established techniques or technologies, for instance, *replication* and *fail-over* for recovering from arbiter failures, and high-throughput low-bandwidth Message Oriented Middleware (MOM) for rapidly reaching consensus—such as *Akka* [34].

It is worth noting, here, that either the acknowledged arbiter, or the agreed upon argumentation rules themselves, should also act as the “guardian” authority enforcing mission-critical *safety constraints*, as well as *ethical compliance*. For instance, by monitoring whether the argumentation process itself and its results comply to well-defined boundaries [35] (the “do and don’t” of the system) it should guarantee that some properties of the system are preserved no matter what.

Feasibility & Performance: Recent hardware advancements allow to embed a fair amount of computational power in sensor and actuator devices as well as in single chips [36]. In parallel, machine learning techniques are also steadily improving, so that incorporating “intelligence” within smart objects is already a feasible reality [16], [7]. Then, a certain minimum performance is also required to ensure that short dialogues such as the ones presented in Subsection III-C are decided nearly in *real-time*, thus arguments exchange as well as inner argumentation-based reasoning must be extremely fast. This is a challenging task which partly requires low latency MOM and partly extremely efficient reasoning engines. Besides the aforementioned Akka framework, *Apache ActiveMq Real Time* [37] is a fork of Apache flagship IoT-focussed MOM explicitly aimed at real-time messaging, while *MiniMe* [38] is an extremely lightweight reasoning engine conceived especially for Semantic Web of Things (SWoT) applications. With the latter, for instance, carrying out simple reasoning tasks supporting argumentation, such as *subsumption*, *entailment*, *abduction*, and *satisfiability*, is particularly efficient, thus feasible on resource-constrained devices such as smartphones, RasPi, and directly on-chip, as already mentioned.

Complexity & Tractability: Another challenge is represented by the *state explosion* problem induced by expanding utterances and dialogues into arguments and relations, which are possibly further expanded in claims and premises. In fact, as shown in Subsection III-D, a simple dialogue made up of 4 utterances can be easily decomposed in 20+ between arguments and relations. In this respect, we would like to emphasize that not all the participants to the debate may argue at the same *level of abstraction*. For instance, in the assisted living scenario thoroughly discussed in Subsection III-D, the smartphone is not at all interested in the premises supporting claims, but only in knowing which argument between “Walter is food choking” and “Walter has respiratory crisis” is stronger. Similarly, the cameras may be totally unaware of which claim their arguments are supporting to actively cooperate to the




Fig. 8: Premises, claims, and support/attack relations obtained by combining Figures 3–7, and by adding a third argument—some arguments such as those stemming from a priori knowledge have been omitted to avoid cluttering the graph. Edges are attached to weights, enabling automatic decision of the winner argument, that is argument *a* (“Walter is food choking”).

dispute resolution process. Therefore, there is no need to represent and reason about the *whole* argumentation graph in participant devices, as they are usually only concerned with a portion of it.

Automatic Argumentation & Natural Language Generation: When interacting with the users, the automation of the process that produces these argumentation graphs, starting from natural language dialogues, is clearly a highly challenging task. There are many technologies that have to be combined and exploited to reach such an ambitious goal. First of all, *argumentation mining* techniques [20] have to be employed for the automatic detection of argument components within dialogue utterances, as well as for the synthesis and generation of new arguments. Within this context, it is worth remarking that *machine learning* tools are already widely employed in argumentation mining, both for the detection of argument components in text, and for the recognition of attack/support relationships between arguments, or part thereof (see [20] and references therein).

Knowledge representation: An open challenge related to the previous one is to combine the aforementioned instruments with *commonsense reasoning*, that has to be distilled within speaking and hearing objects, so as to exploit background knowledge stored in the form of ontologies, or logic facts and rules. Also, *ontology alignment* [39] is an issue which demands for effective and efficient solutions such as the one based on similarity-based clustering proposed in [40], [41]: there, the problem of heterogeneous agents adopting different ontologies in an open MAS is dealt with by enabling agents to learn new terms by letting other agents explain them, and negotiating about their intended semantics.

Besides the numerous challenges discussed above, as many benefits may be as well emphasized.

Interpretability: A self-evident benefit of argumentation is its *interpretability*. A network of speaking and hearing objects acts as a sort of *gray-box*, where low-level data is processed

by machine learning, computer vision, and signal processing tools, often in the form of *black-box* models (as it happens for deep neural networks), but the high-level granules of information that are generated are amenable of interpretation by both humans and other agents. Argumentation, in fact, encourages interacting agents to disclose information and the strategy behind their “negotiation moves” so as to persuade the other parties and reach a deal, which is something that negotiation framework not based on argumentation do not usually do. For instance, auction-based negotiation is just about exchanging subjective values attributed by agents to the object of the negotiation. With argumentation, smart devices will be capable of *explaining* their behaviour and *motivating* their choices and decisions [14], also improving their capability of interaction with *humans-in-the-loop*.

Dealing with Uncertainty: The scenario we described opens the doors also to other interesting research lines, where background knowledge could be continuously enriched and refined thanks to the *learnt experience* reported by speaking and hearing objects. Similarly, the weights attached to edges could be obtained with a *dynamic function* that changes over time and that is tuned according to the knowledge of previous situations. The opportunity here is to advance beyond the state of art in situation recognition, especially in those scenarios where *uncertainty* of information is the norm [42]—such as the IoT, indeed.

Adaptiveness: The *flexibility* of argumentation-based coordination, in particular the fact that coordination protocols and rules may arise by *emergence* from the argumentative interactions of the participants, straightforwardly enables the system as a whole as well as individual agents to *adapt* to the ever-changing goals, constraints, and unexpected contingencies arising during operation, by seamlessly adjusting the coordination protocols upon need [43], [22]. Furthermore, this adaptation capability is embedded in the coordination process enacted by agents through argumentation, thus it does not need

additional mechanisms to monitor the system behaviour, plan adaptation actions, and then execute them.

Robustness: Argumentation may also be exploited to enhance robustness of the system w.r.t. errors and malicious behaviour. For instance, regarding the former, conflicts in sensors' perceptions can be solved by arguing for achieving consensus, while in the latter case, false or harmful claims made by attacker agents may be effectively challenged by "good" ones, *fact-checking* suspected claims so as to detect anomalies and act accordingly.

Trust: Last but not least, when taking into account IoT scenarios with *humans-in-the-loop*, argumentation has the great advantage of promoting trustability of a system: if users can get *justifications* about the decision making undergoing "behind the scenes", about *why* a system is pursuing a given course of actions, and *how* it came up with a precise conclusion about the state of the world, they are likely to increase their *confidence* in relying on the autonomous capabilities of the system. We hereby remark that striving to provide trustability and interpretability is necessary to guarantee *accountability* of systems and decision making, an increasingly hot topic in many fields of AI – from big data [44] to algorithms in general [45] – as witnessed by the recent *transparency initiative* endorsed by many organisations worldwide (<http://www.transparency-initiative.org/>).

V. RELATED WORK

Argumentation, and in particular computational argumentation, has been applied to several application domains, yet, to the best of our knowledge, it has never been exploited in the IoT setting. Nevertheless, several existing applications represent interesting case studies that confirm the suitability of argumentation models and technologies for many contexts that share analogies with the scenario of Speaking Objects. For example, Parsons and Shaw [46] propose to integrate argumentation with traditional decision theory by reasoning about the expected value of actions under uncertain conditions as part of the deliberation process. Hulstijn and van der Torre [47] employ an argumentation framework built upon Dung's seminal work with the goal of performing planning operations and of modeling goals to be achieved. Natural language generation has also been addressed with argumentation-based techniques in [48]. In [49], Jung provides a *computational model* supporting argumentation, showing that it can greatly speed up the negotiation phase of a coordination process.

Although these advancements are certainly relevant for our framework, we believe that, in order to exploit argumentation within the Speaking Objects vision, there are still several steps that have to be carried out. The main open issue is given by the need of combining argumentation with situated reasoning [50], so as to model knowledge of the environment, to represent beliefs, and to exploit relationships with other agents [43]. Moreover, research on argumentation has largely been theoretical, and the applicability of argumentation models to real-world, large-scale problems has yet to be proven [49].

Closely related to argumentation is the research thread on *agreement technologies*, that is, those technologies concerned

with supporting an ensemble of heterogeneous agents, possibly with conflicting goals, in reaching an agreement about some state of affairs or deal. There, argumentation is one of the many ingredients needed, together with ontologies alignment, negotiation, and e-norms enforcement. With respect to the Speaking Objects vision, the literature on agreement technologies may help, for instance, in shedding light upon the benefits of integrating argumentation in a multi-agent setting, in the spirit of [51]. Despite early success regarding application of argumentation-based models for reaching agreement has been reported [52], work is still to be done to fully evaluate the approach.

In the IoT arena, plenty of approaches for the coordination and cooperation of components have been based on *rule-based* systems, and on middleware engines that check and enact them [53]. Such rules dictate how the components should be activated (and their services executed), depending both on the situations that are happening, and on those that – in reaction – should be achieved. However, in a scenario of interacting (goal-oriented) objects, such an approach falls short, due to the impossibility of foreseeing and defining all possible events and state of the affairs, and all the possible ways in which components can be activated. It is in fact unfeasible to design all the possible composition rules that orchestrate the behaviors of the components. Thus, while the possibility of defining rules and constraints for the "do" and the "don't" of the systems (e.g., *safety* and *liveness* properties that should be always guaranteed [54]), the actual way the components act and interact should be identified at run-time by the components themselves, still in respect of global system goals and constraints.

Clearly, many relations exist between our proposal and the literature on *agent-oriented* computing, there including coordination in multi-agent systems [23], and means-ends reasoning [55]. By taking as a reference one the latest AgentLink reports [56], for instance, back in 2005 (i) application domains strictly related to the IoT vision were considered as main drivers of agent-oriented technologies – i.e. web services and service-oriented computing in general, and ambient intelligence – (ii) industry sectors increasingly looking at IoT-based solutions – such as healthcare, manufacturing, and transportation – were already declaring interest in agent-oriented technologies adoption (Figure 6.1 in [56]), and (iii) agent-based technologies currently exploited in the IoT technological stack had already survived the infamous "through of disillusionment" in Gartner's hype cycle (Figure 6.12 in [56]), such as chatterbots, agent-based web services, and intelligent agents in general.

Furthermore, goal-oriented agent-based modeling and simulation were considered already mature enough to take over both academic and industrial practice, and since then have in fact been successfully exploited in the context of smart home security [57], unmanned vehicle governance [58], and surveillance [59], to mention a few. In this respect, it is worth emphasizing that simulation, either agent-based or not, is likely to play a key role in engineering a system adopting the Speaking Objects vision, to ensure that collective behaviours arising by emergence are within certain desired boundaries

[60], [61]. There, in fact, the flexibility of the argumentation process, the dynamism intrinsic in the envisioned deployment scenarios, and the autonomy of components in deciding the course of actions to undertake, strongly demands for an iterative development cycle featuring a “simulate-then-tune” loop at each new software release.

Nevertheless, despite agent-oriented models and technologies being quite mature from the perspective of agent-based engineering, the viability and effectiveness of such techniques in a highly dynamic, heterogeneous, resource-constrained, and scale-demanding domain such as the SIoT still remains to be fully assessed.

VI. CONCLUSIONS

The near future of IoT will see the development of smart objects capable of performing social interactions, thus leading to the so-called SIoT framework. Within this context, we believe that argumentation can play a major role, as it allows to study and analyze debates, reasoning and persuasion in dialogues, with a well-grounded theoretical framework. Building upon the Speaking Objects vision, we have shown how different dialogues can occur in the SIoT context, highlighting the advantages of employing argumentation. In particular, interpretability of decision making, tolerance to uncertainty, adaptiveness of coordination protocols, robustness of the system, and improved trust by end users and amongst interacting components, are the most notable benefits we envision for the Speaking Objects approach.

One major aspect that differentiates the scenario of Speaking Objects from many of the existing application domains of argumentation is that, in our framework, it is difficult to establish a priori an agreement among participants regarding a fixed set of pre-defined behaviours and moves that can be undertaken by the actors in the scene. In argumentation, a typical assumption is to have an external judge or authority that has to control the whole argumentation process, but this is very unlikely in a dynamic, distributed scenario like the one we propose. This aspect will certainly be the subject of future work. A possible solution is to require only temporary agreements among participants, exploiting a limited set of currently-valid rules, which could be further subject to change in the future. Another possibility would be to employ a number of external referees sharing the load of arbitrating argumentations amongst a limited number of participants.

ACKNOWLEDGMENT

This work was supported by the CONNECARE (Personalised Connected Care for Complex Chronic Patients) project (EU H2020-RIA, Contract No. 689802).

REFERENCES

- [1] L. Atzori, A. Iera, G. Morabito, and M. Nitti, “The social internet of things (siot)—when social networks meet the internet of things: Concept, architecture and network characterization,” *Computer networks*, vol. 56, no. 16, pp. 3594–3608, 2012.
- [2] M. Nitti, R. Girau, and L. Atzori, “Trustworthiness management in the social internet of things,” *IEEE Transactions on knowledge and data engineering*, vol. 26, no. 5, pp. 1253–1266, 2014.
- [3] J. Ye, S. Dobson, and S. McKeever, “Situation identification techniques in pervasive computing: A review,” *Pervasive and Mobile Computing*, vol. 8, no. 1, pp. 36–66, 2012.
- [4] R. Gravina, P. Alinia, H. Ghasemzadeh, and G. Fortino, “Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges,” *Information Fusion*, vol. 35, no. 1, pp. 68–80, 2017.
- [5] Z. R. Hesabi, Z. Tari, A. Goscinski, A. Fahad, I. Khalil, and C. Queiroz, *Data Summarization Techniques for Big Data—A Survey*. New York, NY: Springer New York, 2015, pp. 1109–1152.
- [6] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [7] V. Radu, N. D. Lane, S. Bhattacharya, C. Mascolo, M. K. Marina, and F. Kawsar, “Towards multimodal deep learning for activity recognition on mobile devices,” in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, 2016.
- [8] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos, “Context aware computing for the internet of things: A survey,” *Communications Surveys Tutorials, IEEE*, vol. 16, no. 1, pp. 414–454, First 2014.
- [9] E. Davis and G. Marcus, “Commonsense reasoning and commonsense knowledge in artificial intelligence,” *Communications of the ACM*, vol. 58, no. 9, pp. 92–105, 2015.
- [10] M. Kranz, P. Holleis, and A. Schmidt, “Embedded interaction: Interacting with the internet of things,” *IEEE internet computing*, vol. 14, no. 2, pp. 46–53, 2010.
- [11] M. Lippi, M. Mamei, S. Mariani, and F. Zambonelli, “Coordinating distributed speaking objects,” in *37th IEEE International Conference on Distributed Computing Systems, ICDCS 2017, Atlanta, USA, June 5-8, 2017*.
- [12] F. H. Van Eemeren and R. Grootendorst, *Argumentation, communication, and fallacies: A pragma-dialectical perspective*. Routledge, 2016.
- [13] F. Zambonelli, A. Omicini, B. Anzenberger, 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, Part B, pp. 236–252, 2015.
- [14] M. W. Caminada, R. Kutlak, N. Oren, and W. W. Vasconcelos, “Scrutable plan enactment via argumentation and natural language generation,” in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, ser. AAMAS ’14. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 1625–1626. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2615731.2616095>

- [15] D. Walton and E. Krabbe, *Commitment in Dialogue: Basic concept of interpersonal reasoning*. Albany NY: State University of New York Press, 1995.
- [16] K. Bourzac, “Millimeter-scale computers: Now with deep-learning neural networks on board,” Feb 2017. [Online]. Available: <https://goo.gl/sciVTC>
- [17] L. Getoor and B. Taskar, *Introduction to Statistical Relational Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2007.
- [18] Y. Cao, W. Yu, W. Ren, and G. Chen, “An overview of recent progress in the study of distributed multi-agent coordination,” *Industrial Informatics, IEEE Transactions on*, vol. 9, no. 1, pp. 427–438, 2013.
- [19] F. H. Van Eemeren, R. Grootendorst, R. H. Johnson, C. Plantin, and C. A. Willard, *Fundamentals of argumentation theory: A handbook of historical backgrounds and contemporary developments*. Routledge, 2013.
- [20] M. Lippi and P. Torroni, “Argumentation mining: State of the art and emerging trends,” *ACM Trans. Internet Technol.*, vol. 16, no. 2, pp. 10:1–10:25, Mar. 2016. [Online]. Available: <http://doi.acm.org/10.1145/2850417>
- [21] L. Amgoud and S. Parsons, *Agent Dialogues with Conflicting Preferences*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 190–205.
- [22] S. Parsons and P. McBurney, “Argumentation-based dialogues for agent co-ordination,” *Group Decision and Negotiation*, vol. 12, no. 5, pp. 415–439, 2003.
- [23] A. Omicini and M. Viroli, “Coordination models and languages: From parallel computing to self-organisation,” *The Knowledge Engineering Review*, vol. 26, no. 1, pp. 53–59, Mar. 2011.
- [24] P. M. Dung, “On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games,” *Artif. Intell.*, vol. 77, no. 2, pp. 321–357, Sep. 1995.
- [25] D. Walton, “Types of dialogue and burdens of proof,” in *Proceedings of the 2010 Conference on Computational Models of Argument: Proceedings of COMMA 2010*. Amsterdam, The Netherlands, The Netherlands: IOS Press, 2010, pp. 13–24.
- [26] —, “Argumentation theory: A very short introduction,” in *Argumentation in Artificial Intelligence*, G. Simari and I. Rahwan, Eds. Springer US, 2009, pp. 1–22. [Online]. Available: http://link.springer.com/chapter/10.1007/978-0-387-98197-0_1
- [27] P. Besnard, A. J. García, A. Hunter, S. Modgil, H. Prakken, G. R. Simari, and F. Toni, “Introduction to structured argumentation,” *Argument & Computation*, vol. 5, no. 1, pp. 1–4, 2014. [Online]. Available: <http://dx.doi.org/10.1080/19462166.2013.869764>
- [28] T. J. Bench-Capon and P. E. Dunne, “Argumentation in artificial intelligence,” *Artificial intelligence*, vol. 171, no. 10-15, pp. 619–641, 2007.
- [29] P. E. Dunne, A. Hunter, P. McBurney, S. Parsons, and M. Wooldridge, “Weighted argument systems: Basic definitions, algorithms, and complexity results,” *Artificial Intelligence*, vol. 175, no. 2, pp. 457–486, 2011.
- [30] T. J. Bench-Capon, “Persuasion in practical argument using value-based argumentation frameworks,” *Journal of Logic and Computation*, vol. 13, no. 3, pp. 429–448, 2003.
- [31] P. McBurney and S. Parsons, *Chance Discovery Using Dialectical Argumentation*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2001, pp. 414–424. [Online]. Available: http://dx.doi.org/10.1007/3-540-45548-5_57
- [32] M. Scudellari, “Self-driving wheelchairs debut in hospitals and airports [news],” *IEEE Spectrum*, vol. 54, no. 10, pp. 14–14, October 2017.
- [33] F. Rahman, M. Z. A. Bhuiyan, and S. I. Ahamed, “A privacy preserving framework for rfid based healthcare systems,” *Future Generation Computer Systems*, vol. 72, no. 1, pp. 339–352, 2017.
- [34] Lightbend Inc., “Akka,” 2010. [Online]. Available: <http://akka.io/>
- [35] L. Ly, F. Maggi, M. Montali, S. Rinderle-Ma, and W. van der Aalst, “Compliance monitoring in business processes: Functionalities, application, and tool-support,” *Information Systems*, vol. 54, no. 1, pp. 209–234, 2015.
- [36] N. D. Lane, S. Bhattacharya, A. Mathur, P. Georgiev, C. Forlivesi, and F. Kawsar, “Squeezing deep learning into mobile and embedded devices,” *IEEE Pervasive Computing*, vol. 16, no. 3, pp. 82–88, 2017.
- [37] Apache Software Foundation, “Apache ActiveMQ™ Real Time,” 2017. [Online]. Available: <http://activemq.apache.org/activemq-real-time.html>
- [38] F. Scioscia, M. Ruta, G. Loseto, F. Gramegna, S. Ieva, A. Pinto, and E. Di Sciascio, “A mobile matchmaker for the ubiquitous semantic web,” *Int. J. Semant. Web Inf. Syst.*, vol. 10, no. 4, pp. 77–100, Oct. 2014. [Online]. Available: <http://dx.doi.org/10.4018/ijswis.2014100104>
- [39] N. Keskes and A. Rahmoun, “Meaning negotiation based on merged individual context ontology and part of semantic web ontology,” *International Journal on Human Machine Interaction*, vol. 2, no. 1, pp. 1–12, 2016.
- [40] S. Garruzzo and D. Rosaci, “Agent clustering based on semantic negotiation,” *ACM Transactions on Autonomous and Adaptive Systems*, vol. 3, no. 2, 2008.
- [41] —, “Ontology enrichment in multi agent systems through semantic negotiation,” in *Proceedings of the International Conference on Cooperative Information Systems (COOPIS 2007) - LNCS 4803 - Springer Verlag*, Vilamoura, Portugal, 2007.
- [42] P. Krause, S. Ambler, M. Elvang-Goransson, and J. Fox, “A logic of argumentation for reasoning under uncertainty,” *Computational Intelligence*, vol. 11, no. 1, pp. 113–131, 1995.
- [43] I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. Mcburney, S. Parsons, and L. Sonenberg, “Argumentation-based negotiation,” *Knowl. Eng. Rev.*, vol. 18, no. 4, pp. 343–375, Dec. 2003.
- [44] “More accountability for big-data algorithms,” *Nature*, vol. 537, no. 7621, pp. 449–449, sep 2016. [Online]. Available: <http://www.nature.com/doi/10.1038/537449a>
- [45] L. Medsker, “Algorithmic Transparency and Accountability – AI Matters,” 2017. [Online].

Available: <https://sigai.acm.org/aimatters/blog/2017/06/01/algorithmic-transparency-and-accountability/>

- [46] S. Parsons and S. Green, *Argumentation and Qualitative Decision Making*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1999, pp. 328–339.
- [47] J. Hulstijn and L. W. N. van der Torre, “Combining goal generation and planning in an argumentation framework,” in *NMR*, J. P. Delgrande and T. Schaub, Eds., 2004, pp. 212–218.
- [48] M. Elhadad, “Using argumentation in text generation,” *Journal of Pragmatics*, vol. 24, no. 1–2, pp. 189 – 220, 1995, argumentation Within Language.
- [49] H. Jung, M. Tambe, and S. Kulkarni, “Argumentation as distributed constraint satisfaction: Applications and results,” in *Proceedings of the Fifth International Conference on Autonomous Agents*, ser. AGENTS ’01. New York, NY, USA: ACM, 2001, pp. 324–331.
- [50] N. Maudet, S. Parsons, and I. Rahwan, *Argumentation in Multi-Agent Systems: Context and Recent Developments*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 1–16.
- [51] E. Argente, O. Boissier, C. Carrascosa, N. Fornara, P. McBurney, P. Noriega, A. Ricci, J. Sabater-Mir, M. Schumacher, C. Tampitsikas, K. Taveter, G. Vizzari, and G. Vouros, “The role of the environment in agreement technologies,” *Artificial Intelligence Review*, vol. 39, no. 1, pp. 21–38, 2013.
- [52] N. C. Karunatilake, N. R. Jennings, I. Rahwan, and T. J. Norman, *Argument-Based Negotiation in a Social Context*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 104–121. [Online]. Available: https://doi.org/10.1007/11794578_7
- [53] M. A. Razzaque, M. Milojevic-Jevric, A. Palade, and S. Clarke, “Middleware for internet of things: A survey,” *IEEE Internet of Things Journal*, vol. 3, no. 1, pp. 70–95, Feb 2016.
- [54] F. Zambonelli, N. R. Jennings, and M. Wooldridge, “Developing multiagent systems: The gaia methodology,” *ACM Trans. Softw. Eng. Methodol.*, vol. 12, no. 3, pp. 317–370, 2003.
- [55] M. J. Wooldridge, *Reasoning about rational agents*. MIT press, 2000.
- [56] M. Luck, P. McBurney, O. Shehory, and S. Willmott, “Agent technology: computing as interaction (a roadmap for agent based computing),” 2005.
- [57] B. Bos, L. Chmielewski, J. H. Hoepman, and T. S. Nguyen, “Remote management and secure application development for pervasive home systems using jason,” in *Third International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing (SecPerU 2007)*, July 2007, pp. 7–12.
- [58] R. S. Barros, V. H. Heringer, C. E. Pantoja, N. M. Lazzarin, and L. M. de Moraes, “An agent-oriented ground vehicle’s automation using jason framework,” in *ICAART* (2), 2014, pp. 261–266.
- [59] F. Castanedo, M. A. Patricio, J. García, and J. M. Molina, “Extending surveillance systems capabilities using bdi cooperative sensor agents,” in *Proceedings of the 4th*

ACM International Workshop on Video Surveillance and Sensor Networks, ser. VSSN ’06. New York, NY, USA: ACM, 2006, pp. 131–138.


- [60] G. Vizzari, L. Manenti, K. Ohtsuka, and K. Shimura, “An agent-based pedestrian and group dynamics model applied to experimental and real-world scenarios,” *Journal of Intelligent Transportation Systems*, vol. 19, no. 1, pp. 32–45, 2015.
- [61] G. Fortino, R. Gravina, W. Russo, and C. Savaglio, “Modeling and simulating internet-of-things systems: A hybrid agent-oriented approach,” *Computing in Science Engineering*, vol. 19, no. 5, pp. 68–76, 2017.



Marco Lippi received the Ph.D. in Computer and Automation Engineering from the University of Florence in 2010. Currently, he is an Assistant Professor (tenure track) at the University of Modena and Reggio Emilia. He previously held positions at the Universities of Florence, Siena and Bologna, and he was a visiting scholar at Université Pierre et Marie Curie, Paris. His work focuses on machine learning and artificial intelligence, with applications to bioinformatics, time-series forecasting, computer vision, game playing, and argumentation mining.



Marco Mamei is associate professor in Computer Science at the University of Modena and Reggio Emilia, since 2014. He received the PhD in Computer Science from the same University in 2004. He has been visiting researcher at Telecom Italia Lab (IT), Nokia Research Center (USA), Harvard University (USA), Cypcorp Europe (SLO) and Yahoo Research (ES). His current research interests include: mobility data analysis and applications for pervasive and mobile computing.



Stefano Mariani is assistant professor of Computer Science at the University of Modena and Reggio Emilia. He got his PhD in Computer Science from the University of Bologna in 2016. He has been involved in the EU FP7 Project SAPERE, and currently is involved in EU H2020 Project CON-NECARE. His research interests include: coordination models and languages, agent-oriented technologies, pervasive computing, self-organisation mechanisms, socio-technical systems.



Franco Zambonelli is full professor of Computer Science at the University of Modena and Reggio Emilia. He got his PhD in Computer Science and Engineering from the University of Bologna in 1997. He has been scientific manager of the EU FP6 Project CASCADAS and coordinator of the EU FP7 Project SAPERE. He is ACM Distinguished Scientist, member of the Academia Europaea, and IEEE Fellow. His research interests include: pervasive computing, multi-agent systems, self-adaptive and self-organizing systems.