

Crowd-Based Socio-Cognitive Systems

Pablo Noriega¹ and Mark d’Inverno²

¹ IIIA-CSIC, Barcelona, Spain pablo@iia.csic.es

² Goldsmiths College, London, UK
dinverno@gold.ac.uk

Abstract. In recent years there have been several successful examples of crowd-based systems, used for very different purposes and built on a variety of technological artefacts, some ad-hoc, some generic. We presume it is possible to give a precise characterisation of what we call crowd-based sociocognitive systems and postulate that it is possible to formulate a framework to model and implement actual crowd-based sociocognitive systems in a principled way. In this paper we outline the research program, propose the main features of a metamodel for modelling crowd-based sociocognitive systems and make a call to arms for future development and collaboration.³

1 Introduction

Despite the explosion of crowd-based systems and our increasing desire to engage in social activity online, there has been surprisingly little interest from the multi-agent community to use their methods to provide an analysis of such systems. This is surprising, as at the heart of these systems is the notion of coordination, cooperation, emergence, regulation, trust and reputation. This is especially true when we consider that such systems must be modelled from a *socio-cognitive* perspective where individual agents (human or computational) not only behave with some degree of rationality but where that rationality is based on the model that agents have of the other agents in the system. In this paper we will refer to such systems as *Crowd-Based Socio-Cognitive Systems* (or CBSCS) which at their most fundamental level have large numbers of rational agents, each with the ability to model the other agents in the system, and that interact in order to achieve shared or individual goals.

Thinking about the nature of crowd-based socio-cognitive systems in this way might lead readers from a multi-agent system (MAS) background to consider them as wonderful ideal testbeds, with huge amounts of data, to see whether our theories of regulated

³ This paper reflects ideas from conversations with Harko Verhagen and Julian Padget. Both of them and Mark d’Inverno received support from the European Network for Social Intelligence, SINTELNET (FET Open Coordinated Action FP7-ICT-2009-C Project No. 286370) for short term visits to the IIIA. Work supported by projects PRAISE, ACE, COR, AT and by Generalitat of Catalunya grant 2009-SGR-1434. Some of the work was undertaken as part of the FP7 project in the Technology-Enhanced Learning Program called Practice and Performance Analysis Inspiring Social Education (PRAISE). The Principal Investigator for the project at Goldsmiths is Mark d’Inverno and the partners also include IIIA-CSIC, Barcelona

behaviour can be used to describe how people operate in the virtual worlds of crowd-based systems. It is our belief that there is a wonderful opportunity for our research community to make a concerted effort to try to understand these systems from an MAS perspective. Not only to test out own models and theories with lots of real examples of regulated multi-agent systems but to provide proper tools for the analysis of such systems that can provide us with the proper language for engaging with the huge body of work from sociologists, psychologists, anthropologists and cultural theorists investigating human behaviour in such systems. This paper is a first foray into trying to engage the MAS community in the specification, design, analysis and engineering of such systems as we believe that the community has much to offer.

Not only is there an opportunity to understand the social activity that happens online but also an opportunity to build a language to relate the virtual worlds of crowd-based systems to the real emotional and physical worlds we inhabit with our systems. We need tools and analysis techniques for exploring this relationship, for a technical language for describing and understanding the meaning and ramifications of various kinds of online social activity into the real world, and even to start map out the space to see if there are new kinds of opportunity for building new kinds of online systems supporting new kinds of social activity. Not surprisingly many systems are driven by technological possibility and financial gain rather than social good. If we, as researchers in this area, can provide a conceptual framework to map out what is happening within these systems from a multi-agent perspective might there be an opportunity to take part (as just one example) in the discussions about the social responsibility of such systems.

The current situation appears to us that we have no clear technical grounding to adequately describing such systems from an MAS socio-cognitive perspective. We do not have the models, language, theory or tools for the description, analysis and creation of such systems. In order to try to form a bridge between the work in MAS and regulated systems and the plethora of emerging systems this paper sets out to define a conceptual framework for such systems using the language of agents, norms and communication in order to do so. By doing so we hope to seed an emerging research area concerned with developing theories, tools, languages and methodologies for designing such systems. If we can demonstrate the applicability and usefulness of our modelling techniques then we may potentially provide a bridge to other subject areas, such as sociology, in order to have a more rounded understanding of the social and psychological responsibility that should be considered in the design of such systems.

1.1 Socio-Cognitive Systems

Socio-Cognitive Systems (SCS) are characterised across the following characteristics. The reader should consider this to be an indicative list of the qualities of the class of SCS rather than an exhaustive one.

- Dimension 1. The system contains *agents*. Agents are either computational or human and can exhibit purposeful behaviour.
- Dimension 2. The population with a system may be a *mix* of human and software agents.
- Dimension 3.. The agents have a *model* of the world in which they operate.

- Dimension 4. The agents within the system are *rational* in that they are capable of choosing different courses of action based on their own models (however simple or complex these may be).
- Dimension 5. The agents are *social* in that they interact with other agents.
- Dimension 6. The agents have *social models* (either complex or simple) of some of the other agents in the system,
- Dimension 7. At least some of the agents are *socio-cognitive* in the sense that they based their decisions on some decision-making process which takes into account the models of the social world in which they are situated. This includes the capability to plan for future desire states in the environment whilst taking into account the motivations and models of other agents. Such agents can reason about who to collaborate with other agents to achieve individual and joint goals.
- Dimension 8. The agents have *social capabilities* including potentially awareness and models of others, an ability to understand the norms of a system and adopt attitudes relating to norm-compliance and the ability to have altruistic goals
- Dimension 9. Any such system is defined by the system of interacting agents which means that the state of the system can never be known in full as there is no general access to the internal state of agents. This is often referred to as *opacity*.
- Dimension 10. Agents may enter and leave an SCS at any time. It cannot be known either by the designer of the system or by other agents which agents may join or leave. Agents may often be able to join or leave without it being known to other agents.
- Dimension 11. Such systems are *regulated* either intrinsically because of the way the system is designed, the way that some agents have been specified to operate or naturally through the agreement of agents within the system. The point about regulated systems is that not all actions are available to all agents at all times which enables more effective *social coordination* to be facilitated.
- Dimension 12. Agents are *autonomous* and so march to the beat of their own drum and so are not necessarily socially-considerate, benevolent, or honest and so may fail to act as expected or desired or promised.
- Dimension 13. All interactions are mediated by technological artefacts and may therefore be wrapped as communicative acts or messages. Systems that have this property are referred to as *dialogical*.

This list includes the characteristics of systems that have been investigated by the research community looking at *regulated multi-agent systems* and attempts to reflect the recent discussion on Socio-Cognitive Technical Systems that is arising from the Sintelnet project (see Position Papers in www.sintelnet.eu/wiki/index.php/Sourcebook and in particular [2, 1])

1.2 Towards a description of Crowd-based Socio-Cognitive Systems

According to the work of Surowiecki [9] crowd-based systems are concerned with connecting or collecting diverse collections of independently deciding individuals. The basic thesis is that diverse collections of independent autonomous agents with different models, perceptions, motivations and rationality can often analyse or predict scenarios

or data more effectively than individuals even when those individuals are specialists in their area of expertise.

He discusses three types of system advantages

1. *Cognition*. This is about how crowds can make judgements through thinking and information processing faster than individual experts
2. *Coordination*. This is whether social or physical coordination can emerge naturally in large communities of agents. It relates to how a shared view of the reactions of a community provide often accurate judgements about how the community will react to events.
3. *Cooperation*. Again this relates to the emergence of ways in which trust and reputation can emerge naturally without needing a top-down set of norms for social cooperation and co-ordination of activity.

There are then four criteria to distinguish wise crowds from unwise crowds that we summarise using our own MAS terminology as follows.

1. *Diversity of opinion*. Each agent has its own private information that cannot be known by others.
2. *Independence*. Agents' opinions are not completely determined by the opinions of those around them, agents also have a degree of autonomy in the way they form their opinions.
3. *Decentralization*. Agents have different local knowledge and different perceptions of their local environment.
4. *Aggregation*. Some mechanism exists for turning private judgments into a collective decision.

So the question becomes what further characteristics do we need to add to our descriptions of *Socio-Cognitive Systems*. In the four criteria above we already have items 1, 2 and 3 from our own definition. So we add two more

Crowd-based Socio-Cognitive Systems are Socio-Cognitive Systems which have the following three sometimes rather nebulous characteristics.

- Dimension 14. There is a significant population of agents.
- Dimension 15. The system allows for norms (for social cooperation and coordination), trust and reputation to arise natural.
- Dimension 16. The system provides mechanism for turning individual analysis, goals or work into collective analysis, goals and work.

Crowd-based Socio-Cognitive Systems are thus systems that exhibit some features of what is accepted as crowdsourcing or crowd-based behaviour systems but have the distinguishing characteristic that individuals need to reason about themselves and their social environment, because their behaviour is affected by that social environment and also because with their behaviour they may influence the social environment to some extent. A system which implements all of these characteristics also appears in this conference [12].

1.3 Motivation of our work

Such systems are a new phenomenon that involves thousands and sometimes millions of people. What is striking, and perhaps a little unnerving too, is that most such systems are being developed without any theoretical underpinning and in such a way that it is often not easy to see what is underneath the bonnet. Of course there are many kinds of definitions and descriptions but they are not necessarily conducive to a principled analysis or design of CBSCS. Our motivation is to want to understand them in principled ways and describe them in a systematic and formal way that can then be used for the design and implementation of such systems based on principles developed from work in regulated MAS. Our overall ambition is to be part of the design systems of these systems where we can more clearly articulate the social benefit for those participating within it. Our wide-ranging goals to support this ambitions can be described as at least containing the following enumerated below.

1. How can the MAS community take part in the design of CBSCS?
2. What could we offer to the design of such systems in general?
3. How should we present our work in such a way that any system of designers would ever care to notice it?
4. Could we imagine collaborative research projects with designers where research could be developed through the process of design about the nature of designing such systems and understanding how MAS techniques could be applied?
5. In this light, does it make sense to define the universe of SCS (in terms of normative systems and institutions?)
6. If this is not possible (look at the thousands of agent definitions that derailed many scientific and investigations because of lack of a common conceptual and definitional framework from which proper scientific enquiry and engineering systems integration could take place) would it be possible to identify and define the key concepts of CBSCS systems that is useful, engaging and relevant?
7. How might we turn this round and highlight to the MAS community the potential of CBSCS systems for investigating social systems from an MAS perspective?
8. What is a good way to map out the key research issues for a regulated MAS approach to analysing CBSCS systems?
9. Could we identify the potential influence into the design of CBSCS systems for new kinds of collective activity for communities such as ours?

We believe that with a combined effort we can produce answers to these questions through models, metamodels, tools, design methodologies for interaction and interfaces that would underpin the principled design, specification and analysis of such systems. The way to do this, we believe, is to undertake an empirical study of such systems, and attempt their characterisation using our conceptual framework (models, data structures and languages).

We want to consider the social needs of users and the kinds of actions users want, to coordinate within communities and use this as part of the design of new systems. If we can develop clear, useful and principled models of such systems that can be used by designers and communities then the emphasis can be much more focussed on the end use than the specific goals of the engineer who builds it.

Not only do we want principled design and tools and interfaces that make it easy to build these things, but also so that potential communities can understand the range of options that are available to them. So there is a political element here to develop models and design methodologies that can put the user and the community in control of the system they want to be a part of, rather than be part of a system that has been developed by sets of engineers from multi-national companies with less clear motives.

At heart is the question of not wanting to be left out of the design and investigation of crowd-based systems when the MAS community has put such effort into understanding them. Indeed many of us joined the research effort into understanding what social action is from a multi-agent system perspective because we wanted to understand and support cooperation and coordination. How could we coordinate the activities of agents with different personal goals coming together for a common need? Political and social activism, environmentalism, local community, learning, fun and games—we want to be able to build such systems to support a whole range of coordinated social activity as well as investigate the potential range of social activity that can be supported by such systems. If now is not the time that the work we, the MAS community, have developed over the last 20 years or so in understanding social systems from a computational perspective then, when will it ever be?

2 Background

The research programme we envision is to achieve an understanding of socio-cognitive systems, in general, so that we may eventually be able to design new systems with a principled approach. We propose to address the general problem, first by delimiting the universe to an explicit set of features that may allow us to decide whether a given system—existing or being designed—belongs to that universe, and second, developing an abstract understanding of what is common to these systems by separating two fundamental objects of concern: the actual agents (be they human or artificial) and the social space where these agents interact.

On modelling agents, we will assume that these agents may need to reason not only about themselves but also about that social space, because the social space influences and determines in some sense their actions, and also because agents do influence the social space. Thus agents would have to exhibit capabilities or cognitive dispositions to be aware of other agents, to interpret what is the state of the world, and to hold expectations of what possibilities of action are available (for itself or for other agents) and what the consequences of those actions may be. Likewise, the modelling of the social space determines what inputs and outputs will be accessible to the agents, and therefore one has to devise the means to model what the social space “affords” agents to act upon and to be aware of, and the means by which the space may influence the activity of agents. In other words, what objects exist in that space, how agents communicate, how can activities may be coordinated, what types of organisations can an agent belong to, and so on.⁴ Consequently, in abstract terms, we shall speak of meta-models of socio-cognitive agents and metamodels of social spaces. For each of these metamodels

⁴ We use the notion of *affordance* in the spirit of Norman [8].

we would then attempt to produce precise, even formal, descriptions that would allow the specification (and formal analysis) of actual models of agents and of social spaces. Metamodels that in turn need to be accompanied by technological artefacts that enable the actual implementation of socio-cognitive systems where artificial or natural agents pullulate in an artificial social space.

We realise that attempting to find a single metamodel for agents and for social spaces is at best impractical. However we glimpse the possibility of sketching some generic metamodels for families of socio-cognitive systems (SCS). For example, on-line marketplaces, massive on-line role playing games, mixed-level simulation, or policy-making support systems. We postulate that one of these families are crowd-based SCS.

Consequently, we may rephrase the next steps in our research programme, in terms of CBSCS, in the following four steps:

1. Compile a corpus of CBSCS
2. Understand what is "structural" of crowd-based SCS.
3. Map the universe from these exemplars
4. Outline a conceptual framework, identify adequate tools and methodological guidelines

2.1 Our bias

In two previous papers [10, 7] we discussed the basic tenants of our research programme: (i) a three-fold view of socio-cognitive systems, the notion of shared context and the relationships among the three views, its abstract (*platform-independent*) model, and the implementable (*platform-specific*) model. (ii) The relationship between metamodel, environment, computational architecture and platform. (iii) The separation between agents and social space and between design environment and enactment environment of the socio-cognitive systems.

We are confident that the approach we propose has some hope to succeed because we already have done a similar task for the abstract notion of "electronic institutions" ([3]). In this case, the universe of systems it may model and implement is close to the very general notion of socio-cognitive systems that we have as the ultimate goal of our research programme.

After a decade of development of the conceptual framework, associated tools and a considerable number of application cases, we were able to produce a formal metamodel, and the associated technological artefacts integrated in a working development platform [3]. The electronic institution framework is actually only applicable to the social interaction space, since one of its key assumptions is that participating agents are black boxes who are able and willing to comply with the space conventions, thus there is no commitment to any agent model. Thus, the space itself is a regulated multiagent system where agents interact through speech acts that are organised as interrelated conversations (or "scenes") where the illocutionary exchange is prescribed with regimented procedural rules.

Figure 1 summarises the actual metamodel. Briefly speaking, it includes two parts: the "dialogical framework" that provide participants those elements that are involved in action (these are "dialogical" because interactions are understood as conversations); and

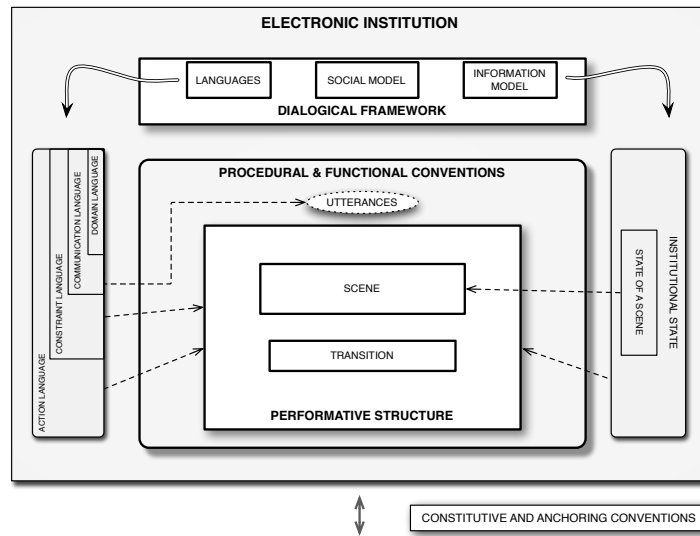


Fig. 1. The conceptual model of electronic institutions. After [6].

the regulations that govern those actions. It is beyond the scope of this paper to explain these components in detail but we shall make some quick remarks on the *affordances* of this metamodel, in order to give a flavour of what needs to be made explicit and properly formalised in order to provide *ad-hoc* metamodels for crowd-based SCS:

- Ontologies about the domain of interaction, of possible actions, its preconditions and effects and how these are related to the social model. These ontologies are captured and reflected through a collection of languages that are part of that dialogical framework.
- A social model that in this case is limited to a finite set of roles and relations among these.
- A means to specify and structure local contexts (scenes) where a subgroup of the participants may participate in a collective activity (conversation) and a structure that establishes how agents playing specific roles may move between scenes or participate simultaneously in different scenes (the performative structure)
- A normative model that allows to specify and regiment those conventions that regulate the conversational exchanges and the pragmatics of those exchanges.
- An information model that includes (i) a collection of variables and parameters that refer to the domain of interaction, participants in the interactions and interactions themselves, (ii) the “state of the scene” that keeps track of the evolution of the conversation and the “state of the institution” that includes those states and, in essence, keeps track of the values of those variables and parameters at every point in time.

3 Towards a framework for crowd-based SCS

3.1 Outlook

We aim to define a framework to model and eventually implement crowd-based SCS. Consequently, we distinguish between the conceptual framework that is used for modelling a crowd-based SCS and the artefacts that implement actual models. With the conceptual framework we intend to model social spaces that enable participating agents to perform a variety of social tasks, that are particular of crowd-based SCS. Namely, define and broadcast a collective challenge, accept a call, perform the individual tasks involved in the challenge, compile the responses to the challenge, assess whether the tasks are properly completed and the rewards entailed by accomplishing a task are being properly granted, etc. Along this conceptual framework we presume that there will be technological artefacts that serve to implement actual systems that are modelled with the conceptual framework.

In order to enable crowdsourcing tasks, the conceptual framework needs to “afford” the designer, and ultimately the participating agents, the means to specify and apprehend the social space where those tasks take place. Because we intend to use the conceptual framework to model crowd-based SCS, we will refer to it as a *metamodel*.

Although it is beyond the scope of this paper to give a formal definition, a metamodel is a collection of affordances each of these specified as a class of formalisms. Loosely speaking, the metamodel should *afford* agents those aspects of the social context that enable them to act proficiently. Hence, the metamodel should include means to formalise (i) the contents of the shared context, (ii) the features of the agents that participate in it, (iii) their interactions and (iv) the overall behaviour of the system.

These affordances should involve, for example, the means to establish and become aware of how much is revealed of the individual’s identity, and how individuals become aware of other participants in a given activity. Moreover, agents would need to have means to communicate in non-ambiguous terms, and that implies that, at the very least, the metamodel shall afford some shared ontology, communication language and interaction model. Coordination is usually achieved by organising the crowdsourced system into activities, usually collective, that may be executed in parallel or following certain rules that link them by time, causality or whatever (for instance, in Wikipedia, editing articles, handling disputes, quality review,...). In most SCS one may reify a social structure, however primitive, where individuals are meant to play roles (e.g., *editor*, *administrator*, *bureaucrat* in *Wikipedia*) that entail some capabilities and are frequently subject to rules that apply to individuals only while they play that role. More sophisticated social structures involving groups, hierarchies and organisations are used to better coordinate complex projects.⁵ Likewise there are usually means through which the designer may specify the “rules of the game” (see for example the “five pillars” of

⁵ For instance in the crowd-sourced drafting of the Constitution of Iceland, Parliament created a Constitutional Council, whose members were citizens voted by the population, this Council was then organised in three working groups who produced recommendations to the council, that received comments from the public through the Council’s webpage, and were submitted in turn to the Council and incorporated into a “process document” that was again open to public comments until a “draft proposition” was made by the Council. <http://stjornlagarad.is/english/>

Wikipedia and the associated guidelines and polices) and how to enforce and update them; and for participants to be informed of these rules, of their application and of any possibility of changing them.

3.2 A tentative list of affordances for Crowd-Based SCS

We presume it is possible to propose a single conceptual framework that allows the modelling of a large class of CBSCS. That is an empirical question that can only be answered with a systematic analysis of existing crowd-based sociocognitive systems and a serious attempt at the design of new ones. Nevertheless, after our experience with electronic institutions [3], a discussion of normative MAS [6] and a superficial inspection of three other paradigmatic SCS classes (gaming, simulation and policy-making [10, 7]) we propose a tentative list of affordances that we claim should be included and made precise in an abstract conceptual framework for crowd-based sociocognitive systems.

Agent types and agent socio-cognitive models Probably three types are enough: (A1) humans or real world organisations that commission or execute a task; (A2) software agents that commission or execute a task; (B) server agents that perform management or support functions (for instance, enact a collective process, search for potential executioners, remind executioners of pending tasks, evaluate task performance, perform police-like and time-keeping functions). For each of these types we may want to make explicit the socio-cognitive dispositions that agents have or should have.

Domain ontology. This will include the elements that are used to define the content of collective contexts and interactions. For example, in `change.org` the ontology would include petition, signature, motivation, proposer, number of signatures. in Wikipedia, articles, review, update, guideline, editor, bureaucrat, dispute, etc.

Languages. These are needed to define the behaviour of the system and the way it is regulated. These may be organised as a hierarchy of languages that starts with a *domain* language (to refer to the domain ontology); *communication* language, *action* languages (description of an action); *constraint* languages (preconditions and post-conditions of actions); *normative* languages (procedural, functional or operational directions; behavioural rules,...) and so on, depending on the complexity of the crowd-based system

Interaction contexts are needed to define separate collective activities and their interrelations. They are ideal locations or activities where several agents interact simultaneously, sharing the *same state* (they correspond to EI's "scenes"). For instance a Turk challenge, a Wikipedia dispute over an article, the trading process of a prediction market. When a CBSCS involves several activities, it should be possible to specify how several local context may be connected and how individuals may move between them.

Actions . Atomic actions like "speak", "move to another interaction context"; complex actions like "broadcast", "execute", e.g. in Wikipedia: create, edit, censor an article, introduce a guideline or a norm, participate in a dispute.

Information structures. The (shared) state of the system (the value of each and every variable that may change through the action of some agent or the passing of

time) and the shared state of local contexts (generally, subsets of the state of the system); profile of participants, performance indicators, data structures associated with composite actions, ...

Social constructs. Describe the way individuals are related among themselves and also serve as means to refer to individuals and groups of agents by the role they play, rather than by who they actually are. These may include: roles; relations among roles (n-ary relationships between individuals as well as higher-order relationships. i.e, groups, hierarchies of roles, power relationships and so on); organisations (groups plus coordination conventions).

Regulatory system. To allow top-down or bottom-up articulation of interactions: e.g., norms of different types (procedural, constitutional, rules of behaviour,...) with their associated features (relationships between norms, incentives, effects of compliance and non-compliance,...)

Inference. Assumptions about different ways of inferring intended or observed behaviour. Ways to model reasoning under *uncertainty* and alternatives to classical forms of inference like argumentation of coherence.

Social order mechanisms. To allow top-down or bottom-up governance. Among these: regimentation (rendering some actions impossible, strict application of sanctions,...); mechanisms for assessment, evaluation, prosecution and punishment of non-compliance; social devices (trust, reputation, prestige, status, gossip); policing devices (law enforcement),...

Performance indicators. To measure the behaviour of the system by the designer and qualified participants.

Evolution. Means by which the system may change over time (adaptation of agents, bottom-up, negotiated, external change of system regulations) and devices involved in producing that change happens: performance indicators, normative transition functions and such.

4 Concluding Remarks

4.1 The affordances challenge

The list of affordances we used in the previous section is biased by our previous attempts with games, simulation and policy-making. This list is inadequate for two main reasons: It is a mix of heterogenous notions : languages and ontology, for instance, are not like “inference” or “social order mechanisms”, which are easier to assimilate to a collection of formalisms each providing a different flavour to the same type of functionality. Second, it is not complete, or not explicit enough. For instance, What are the affordances that explain the main functionalities of the *Amazon Turk*? or Where does one capture the requirements of *crowdness* or the way one may filter the participation of given individuals?

There is another matter to ponder: Should the list of affordances be different in crowd-based SCS and other classes of SCS (say games or electronic markets)? What would the advances be for one answer over the other? How may this issue be settled?

4.2 The expressiveness trade-off.

As suggested above, an affordance should entail those conceptual elements that when properly specified allow for a precise specification of the way the CBSCS enables certain functionalities. Ideally, that precise incarnation ought to be made operational through technological artefacts that implement those functionalities. As discussed in [6] the metamodel may choose among several available formalisations of a given affordance, depending on the functionality desired, and each formalisation will be implemented possibly in different ways choosing among available artefacts. However, there is a “whorfian” expressiveness paradox: once a formalism is chosen, then the implementation of the affordance is conditioned by the chosen formalism and the corresponding choice of artefacts, and *vice-versa*.⁶

When there is a collection of artefacts that are coherent, interoperate and are integrated on a working computational architecture, they are usually called a *platform*. Ideally, the metamodels we foresee should allow a formal representation of *platform-independent models* whose implementation will eventually depend on available artefacts. If one is lucky enough to have a platform that integrates those artefacts, including a specification language, the transcription of those platform-independent into platform-dependent implementation is a relatively simple task, modulo the “expressiveness paradox.”

As with other representation of a class of problems (e.g. norms, planning, work-flows), there is a trade-off between the generality of the framework that is used to describe (and formalise) a sub-class of those problems and the ease with which a sub-class of those problems is represented in the framework.⁷

Currently, many working CBSCS are not platform-independent and some forms of crowd coordination have proliferated because practical platforms are available. Let's examine three paradigmatic cases, of platform dependent classes of systems: *Mechanical Turk*-enabled projects, prediction markets and *Ushahidi* crisis mapping .[4] Two obvious remarks apply to the three examples: First, the three examples are designed using platform-dependent models. That is, they are founded on particular technological artefacts (or platform) that restrict CBSCS design to involve only those features that are afforded and implemented by the particular artefacts (or platform). In other words, they can model only the systems that can be implemented with the corresponding artefacts. Second, these platforms are the nuclei of the corresponding crowd-based SCS but are not necessarily powerful enough to model and implement an actual crowd-based SCS and, as we shall show, not enough for the three examples at hand. Now, in particular,

- The Amazon *Mechanical Turk* (<https://www.mturk.com>) is a full-fledged *platform* that applies to different sorts of microwork, and its metamodel affords means for specifying, enacting and evaluating projects that have one single activity, performed

⁶ The usual, but not altogether false, misreading of B.L. Whorf (in [11]) as the postulate that the structure of anyone's native language strongly influences or fully determines the worldview he will acquire as he learns the language.

⁷ We have come across that trade-off in the case of electronic institutions where the rich language for describing transitions between scenes are unnecessarily cumbersome in work-flows that can be hard-wired(in e-commerce, for instance).

as a multitude of micro-tasks by populations whose members may be filtered into the project from an open pool. It is Amazon, and not the Turk platform, who provides the additional services that allow the management of projects through other artefacts.

- Prediction markets may be implemented in different platforms—for example the Iowa Electronic Market (tippie.uiowa.edu/iem/) or iPredict (www.ipredict.co.nz/)—each of them is a regimented implementation of a particular futures market with its own notion of contracts, its trading protocols, entry and compensation requirements, the definition of an event to predict. and so on They are open to traders that fulfil some requirements but there is a unique model for each platform, and no more affordances may be included in their CBSCS.
- The *Ushahidi* platform (<http://ushahidi.com/products/ushahidi-platform>) gravitates around the *Ushahidi* map, an artefact that consists of a graphical representation of geo-referenced data that belong to an explicit taxonomy of relevant events. This platform supports the collection, interactive mapping and visualisation of events but the implementation of an actual crisis follow-up SCS needs to be complemented with other ad-hoc artefacts for integrating, filtering and deciding on how to use incoming data. Thus, while the mapping activity has a single model, the metamodel for crisis management needs to afford other features depending on the organisational structure of the management organisation and the functions it assumes during the crisis.

We presume that it is possible and useful to strive for platform-independent models, and we believe that it is possible to meet this challenge. Evidently, a good analysis of currently available platforms is an immediate step to take towards the affordances challenge.

4.3 A call to arms

This is simply to say to our colleagues in the research field of normative multi-agent systems and in areas of the social sciences, that the theories, models and methodologies that we have developed in the last 20 years or so can be brought to the design and analysis of the increasing plethora of crowd-based socio-cognitive systems. There are several goals that we outlined above but in the immediate term there are several issues we need to address first that we consider here.

We propose to start with an empirical study of existing crowd-based sociocognitive systems that should enlighten the development of an abstract conceptual framework for modelling CBSCS, or support the convenience of developing several metamodels. We think that the achievement of a clear description of that conceptual framework should open the way to the assembly of technological artefacts and the development of crowd-based platforms that allow design and use of CBSCS whose properties can be ascertained formally and ideally proven to be correctly implemented.

The steps to follow, in our opinion, are:

1. Compile a set of CBSCS “examples” that map out the space of such systems clearly.
2. Produce a more rigorous description of the universe of CBSCS by developing the 16 dimensions we have posed above.

3. Identify precise descriptions of the entities (in each example) that permitted modelling and implementation of affordances (also to be described with precision) present in the models.
4. Develop this model and framework using formal methods in order to ensure clarity, rigour and portability of our ideas. (Using techniques for developing formal conceptual frameworks such as [5] for example.)
5. Contrast this analysis with the tentative list of characteristics we have developed above and refine these.
6. Based on that outcome, draft a single conceptual framework.
7. Classify existing and future CBSCS and formulate classification schemata to apply these ideas.
8. Identify useful technological artefacts to implement the components of the conceptual framework.
9. Develop proof-of-concept CBSCS, analyse and postulate preliminary methodological guidelines for those interested in designing such systems.

We hope that we have made it clear why our models and method are of interested to those participating in the workshop as well as a introducing a tentative road map of how an interdisciplinary community might emerge through contribution to the explicit components that we identify in the roadmap of research above.

References

1. Alexander Artikis Andrew Jones and Jeremy Pitt. The design of intelligent socio-technical systems. *Artificial Intelligence Review*, 39(1):5–20, May 2013.
2. Cristiano Castelfranchi. Minds as social institutions. *Phenomenology and the Cognitive Sciences*, pages 1–23, forthcoming.
3. Mark d’Inverno, Michael Luck, Pablo Noriega, Juan A. Rodriguez-Aguilar, and Carles Sierra. Communicating open systems. *Artificial Intelligence*, 186(0):38 – 94, 2012.
4. Jessica Heinzelman, Rachel Brown, and Patrick Meier. Mobile technology, crowdsourcing and peace mapping: New theory and applications for conflict management. In Marta Poblet, editor, *Mobile Technologies for Conflict Management*, volume 2 of *Law, Governance and Technology Series*, pages 39–53. Springer Netherlands, 2011.
5. M. Luck and M. d’Inverno. Structuring a Z specification to provide a formal framework for autonomous agent systems. In J. P. Bowen and M. G. Hinchey, editors, *ZUM’95: 9th International Conference of Z Users, Lecture Notes in Computer Science*, pages 47–62, Heidelberg, 1995. Springer-Verlag.
6. Pablo Noriega, Amit K. Chopra, Nicoletta Fornara, Henrique Lopes Cardoso, and Munindar P. Singh. Regulated MAS: Social Perspective. In Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leendert W. N. van der Torre, editors, *Normative Multi-Agent Systems*, volume 4 of *Dagstuhl Follow-Ups*, pages 93–133. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2013.
7. Pablo Noriega and Julian Padget. Approaching social coordination from a normative perspective. In Harko Verhagen, editor, *Analytical Sociology, Social Coordination and Informatics*, 06/06/2013 2013.
8. Donald A. Norman. Affordance, conventions, and design. *interactions*, 6(3):38–43, May 1999.

9. James Surowiecki. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Little, Brown, 2004. ISBN 0-316-86173-1.
10. Harko Verhagen, Pablo Noriega, and Mark d'Inverno. Towards a design framework for controlled hybrid social games. In Harko Verhagen, Pablo Noriega, Tina Balke, and Marina de Vos, editors, *Social Coordination: Principles, Artifacts and Theories (Social.PATH)*, pages 83–87, Exeter, UK, 03/04/2013 2013. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour.
11. Benjamin Lee Whorf. The relation of habitual thought and behavior to language. In J.B. Carroll, editor, *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*, pages 134–159. MIT Press, 1956. ISBN 0-262-73006-5.
12. Matthew Yee-King and Mark d'Inverno. Pedagogical agents for social music learning in crowd-based socio-cognitive systems. In *Crowd Intelligence: Foundations, Methods and Practices*. European Network for Social Intelligence, Barcelona, January 2104.