Nesta...

COLLECTIVE INTELLIGENCE:
HOW DOES IT EMERGE?

Stefana Broadbent and **Mattia Gallotti***

June 2015

While collective intelligence is not a new phenomenon, the internet and digital media have greatly extended the possibilities of remote collective collaboration. The joint effect of a networked public, digital platform and open organisations are allowing knowledge to emerge from truly novel forms of social interaction. However, not all collective digital projects are successful in engaging people and generating significant outcomes. It is therefore important to understand what are the mechanisms that underlie co-operative processes of knowledge exchange and development.

This paper discusses some of the prerequisites for collective intelligence: the cognitive predisposition that allows humans to elaborate shared intentions, the presence of cultural artefacts that allow co-ordination across time and space, the interaction with digital tools that embed social practices, the existence of systems of governance that encourage the free transformation of knowledge.

These requirements describe a process of socially developed knowledge that relies on tools and environments designed to enable contributors to access aggregated information and freely transform the results. We can start envisaging how to translate these mechanisms in organisational and design principles.

*School of Advanced Studies, University of London.

1. INTRODUCTION

The term collective intelligence appears more and more frequently in discussions about the internet as a tool to foster new forms of public engagement. Generally, collective intelligence refers to the process by which large groups of individuals pool their knowledge, data and skills to contribute in solving societal issues. By inputting environmental or clinical data, mapping territories, discussing and voting, coding and writing, citizens can contribute their knowledge and ideas to data collection tasks, analysis and public debates.

Two converging phenomena have brought to the forefront the idea of leveraging collective intelligence to address problems of public policy. One is the pressing demand for innovation in models of public governance. The combined effect of loss of trust in elected officials and government institutions, coupled with the increased complexity of social and political issues, is pushing traditional public services to incorporate more participatory models of governance. The other is the increasing ease with which networked publics can carry out participatory activities online. The simplification of the interfaces, combined with growing digital literacy and an ever increasing sophistication of computational systems, has expanded the range and complexity of activities in which citizens can now be involved.

Spectacular examples of collective online action, such as Wikipedia, the Linux operating system, or numerous other results from the open source movement, have shown the potential of digital tools and a distributed networked public to generate forms of collective intelligence. Likewise, at the local level, there are many other attempts to involve citizens in reporting accidents, budgeting, writing political manifestos and creating registers of patients. All the examples showcase the potential of digital tools to mobilise the skills of myriads of people. These rising forms of collaboration and joint action seem to be qualitatively different from other and more traditional forms of social action. The number and distribution of people, the presence of digital platforms and specific structures of co–ordination, as well as unprecedented modes of participation and governance, are constitutive elements of these novel forms of collective intelligence. However, it is difficult to say exactly what is unique to accomplished cases of collective intelligence such as Wikipedia, Github, and Ushaidi, compared to classic organisations as well as to other forms of online participation such as Youtube and Facebook.

In this paper, we shall address the question of the nature and structure of collective intelligence in the context of present discussions about online tools of public engagement. The overarching goal is to identify a set of **preconditions** for collective intelligence to come about. In the past, it was common to argue that joint action might involve some capacity for mindreading based on proximity among the people involved in collaborative acts. However, the internet revolution has made it possible for geographically distant individuals to collaborate and come up with joint results that do not require physical proximity. This phenomenon makes the question about the mechanisms enabling dispersed individuals to produce shared outcomes all the more interesting and timely. Do concepts of distributed and

shared knowledge and expertise imply that collective intelligence require the agents' intention to act jointly? If so, what would the nature and mechanism of the 'meeting of minds' that sustains forms of collective intentionality be? And what would the impact be on the way we think about the organisation and governance of online collaborative systems?

In order to answer these questions, we will present in detail the **Missing Maps project** where thousands of people use an open source platform to map vulnerable regions of the world where humanitarian organisations are deployed. This case provides the background to outline a number of key questions and factors that are common to modern forms of collective intelligence.

We will then draw upon the existing literature on collective intentionality across the social and cognitive sciences to identify the **cognitive mechanisms** that enable and sustain the emergence of collective intelligence.

We will then examine the role of **cultural artefacts** in sustaining the process of co–ordination across space and time.

We will then discuss the crucial role of digital tools in the co–ordination process.

Finally, we will conclude by examining the type of **governance** that best enables collective intelligence to emerge.

2. A CASE OF COLLECTIVE INTELLIGENCE: THE MISSING MAPS PROJECT

The Missing Maps project is a fascinating case study of collective intelligence in the making. Borne out of a collaboration between the charities MSF (Doctors without Borders), the British and the American Red Cross, and the open source foundation Openstreetmap, the project aims to provide maps of vulnerable areas for humanitarian intervention and disaster relief, and has set a goal of two years to complete the task. Using satellite imagery and Openstreetmap, an open data–mapping platform led by a community of mappers and software developers, the project engages thousands of volunteers tracing areas where official maps are limited or do not exist at all. Missing Maps is a feat that would be inconceivable without the combination of a network of people, specialised digital tools and a process to co–ordinate the efforts. It is also an excellent example of online and offline hybridity, where part of the task is done remotely and digitally while part is completed on site by people on the ground. The whole Missing Map process unfolds in phases in order to build resilience in the project, by ensuring for instance that the first phase of mapping carried out on satellite images by 'lay' volunteers at home, is systematically double checked by 'validators' to spot inconsistencies. Later, once large areas have been mapped and validated, the maps are printed out by humanitarian agents on the ground, in Sudan or Nepal for instance, and enriched with fine–grained details by hand. People in the field visit the locations and add the paths, names, and geographical features that are missing in the first drafts of the maps. This phase provides an additional layer of information and serves as a supplementary validation of the work done by remote mappers. Yet another series of volunteers are eventually tasked to transcribe the handwritten notes into the digital platform. As it unfolds, the full mapping process builds upon layers of increasing detail and precision.

On top of this process, software developers are constantly improving the Openstreetmap platform to add tools and features. As any open source project there is a tension to improve and adapt what is being built because the systems are always incomplete. The participation and interaction of different people with vastly different levels of competence, engagement and expertise means that the tools themselves are permanently under pressure to be fixed and adapted. The software JOSM editor for Openstreetmap for instance, is being fixed and updated on a daily basis in reply to the requests of its users who notify the presence of bugs or inconsistencies. Similarly, there are continuous debates and readjustments of the terminology to describe spatial elements on the ground. Developing and improving mapping tools therefore is an integral part of the creative process of the collective, which is elaborating new notational and functional systems in response to the needs and requirements of the mapping activity itself.

While the majority of the initial mapping work on the satellite images, is carried out individually and separately by thousands of volunteers dispersed all over the world, there are occasions in which people come together to work side by side. There are, for instance,

mapping events in various cities across the world in which tens or hundreds of volunteers meet in a room to map together for a few hours. The structure of a mapping 'party' is extremely revealing and it gives insights into the mechanisms at play in an exercise of collective intelligence. Before the event, participants are recruited online, generally through social media, and are asked to sign up for the meeting and register with Openstreetmap. No prior knowledge of mapping tools or the Openstreetmap software is required.

The event usually takes place in a hall where the organisers (from the Red Cross, MSF and Openstreetmap) have simply organised tables and internet access as participants bring their own laptops to work with. The session begins with an organiser explaining the region that will be mapped, giving precise instructions as to which elements will be mapped (e.g. only main roads and inhabited areas) and giving some ground images to exemplify what roads or houses look like. While this initial phase provides participants with the objective, scope, and context of the exercise, in the next phase another organiser gives a short tutorial on how to use the mapping tools. The tutorial includes a set of recommendations on how to co-ordinate the joint work, as well as a step-by-step explanation of the software tools. Not all the participants are newcomers, some of them are active mappers or have attended previous events of the same kind. More experienced mappers are offered the possibility to work with a more complex mapping tool (JOSM) or to validate the work of others.

During the few hours of the event, participants join in a predefined task, learn how to use the online mapping tools, and start contributing by drawing roads, trails, buildings, etc., over a satellite image of the area. They scroll through satellite views and spot roads or villages that they then mark with lines or circles. It is no easy task especially for unskilled mappers, because satellite images are typically blurred, for example, it is difficult to distinguish trees from houses. All in all, the mapping session is an informative and social event in which people help each other, volunteers give support and a 'common' sense of purpose is constructed. When individuals leave the hall, they are asked to continue their work at home. One of the outcomes of the event is to increase the likelihood that participants will become regular mappers. However, only a fraction of volunteers of the Missing Maps project live close enough to a mapping party and will in fact ever attend a meeting of this sort.

2.1. The open source development model

The case of Missing Maps is not unique, as there is a now long tradition of similar collective endeavours. Open source communities, such as the Debian project, have operated for years on the premise that a dispersed community of coders can collaborate effectively to create and maintain complex software programs and packages of code. As stated in the history page of the Debian Foundation, Debian is both a completely distributed system of development and a tightly governed one. Like in the Missing Maps project, Debian co-ordinators help define some of the priorities and tasks, and they control the releases of code. Unlike Open Street Maps the criteria of acceptance in the community are quite stringent and demanding (although contributing to the Openstreetmap or JOSM code is also subject to more severe forms of screening). Debian is produced by almost a thousand active developers spread around the world, who like Missing Maps occasionally meet at conventions to code together, take organisational decisions and consolidate their relationships. While each developer codes individually, the whole project relies on the fact that code is layered on top of other code, and it can be integrated as a component of other pieces of code.

There are a number of prominent open source organisations involved in open source development such as the Linux Foundation, the Mozilla Foundation, the Apache Software Foundation and the original Free Software Foundation. These communities have different modes of governance, tools, repositories, etc. All of them have rather precise processes to contribute code to a project, processes that are technical but also social in order to guarantee

co–ordination, quality and relevance. Open source developers have different backgrounds and levels of expertise, but compared to Missing Maps mappers they are all knowledgeable coders. Openstreetmap was intentionally designed to allow anyone to map their neighbourhood and thus had to be made as accessible as possible. Such openness is not common and certainly not suitable for most developer communities because of the nature of the programming activities which requires each piece of code to be fully operational in order not to disrupt the code it is embedded in. However, we must again make a distinction between the mappers and the Openstreetmap developers involved in extending, improving and fixing the code that supports the mapping activities. Here we are in a configuration that is similar to any other open source development project, with a governance structure based on merit and specific coding expertise.

2.2. From division of labour to collective intelligence

The openness and inclusiveness to a lay community of contributors, makes the Missing Maps project an excellent case to discuss the specificity of collective intelligence. The project shows that single agents with minimal topical expertise can generate results through collaboration, the scale of which exceeds the capacity of individuals. We could argue however, that this is a generic effect of any form of collaboration in which individuals join efforts. According to several accounts of the evolutionary and cultural origins of sociality, early humans adapted to pool their mental attitudes and skills as they realised that they could not respond to increasing environmental pressure by relying on their forces alone (Tomasello 2014). Since knowledge is distributed among, and shared by, individual agents whose cognitive abilities and rational faculties are bounded in many respects (Kahneman 2013), the only way for individuals to overcome these limitations was to join forces and work together to construct cultural artefacts that can be shared across people, space and time. Unsurprisingly, this is also the principle behind the division of labour that sustains the functioning of complex social and economic systems: when humans and possibly other non–human animals work together, they achieve bigger results than individuals working on their own.

Indeed, the people involved in mapping sessions engage in collaboration on the assumption that they could achieve something bigger only if they join forces. They are willing therefore, to contribute their time and efforts to contribute in building an external information system that can be accessible to all. The outcome of a joint activity like Missing Maps is, not only bigger, but better than individual contributions. The term ‘better’ hints at the possibility that group activities might yield a qualitatively different, though not necessarily optimal, performance than the sum of individual actions. If the result of pooling resources together could be explained (away) in terms of the single agents’ individual contributions, nothing would be left out of the sheer ‘aggregation’ of what the agents can contribute in terms of their own bodies of knowledge and experience, ideas and motivations, know–hows and goals. In other words, if all we need for a network or group of people to produce an intelligent solution is for labour to be divided (in a classic scenario of division of labour) among multiple agents, it would be difficult to explain how the result could differ from the individual knowledge and expertise. But the most visible cases of collective intelligence show that what turns the sum of individual intelligences into a qualitatively distinct phenomenon is the fact that, by joining forces, sometimes people achieve results that could only emerge as a result of their interactions. The product of these interactions is what, we would argue, is qualitatively different from the simple aggregation because the interaction itself contributes to generating something radically new. So, back to Missing Maps project. The final map ‘emerges’ when the participants online and on the ground, acting jointly with the software developers who have created the tools and visualisations, the humanitarian organisations that have provided the object to be mapped, all pool their topical resources – ideas, experiences, and so forth – in such a way that the outcome of their interactions can no longer be partitioned over the single agents.

2.3. The effect of interaction

The fact that the joint outcome displays features that cannot be fully described in terms of how much, or what, each agent knows, shows that there is more to the intelligence of groups of geographically proximate and dispersed individuals than mere division of labour. The question thus arises as to **where the 'plus' of collective intelligence emerges from**. This is a particularly salient question in the case of digital communities where the 'plus' is not generated in a central overseeing organisation that assembles and organises the single contributions (as would be the case in traditional industrial organisations) but emerges from the network itself. In fact, talk of joint action is often associated with expressions such as 'smart' groups and the wisdom of crowds, which gesture in the direction of a flat networked organisation and more importantly, a qualitatively, rather than merely quantitatively, phenomenon.

We take this to be the problem of collective intelligence – the problem of explaining **what makes the intelligence of jointly interacting individuals more than the mere sum of individual intelligences prompted by division of labour**. This is the question about *"what sets a scientific paradigm apart from a jumble of scientists, Wikipedia from a collection of people with an internet connection and topical expertise, or a vibrant cluster such as Silicon Valley from a hodgepodge of companies that happen to be located in the same place"* (Nesta working paper, p. 5).

Why, and how, does the interaction of individuals 'enlarge' the potential of their actions so as to bring about collective results that exceed their own capacity? What conditions must be in place for individual thoughts and experiences to be shared, and for expressions of true group intelligence to emerge? Although much social–science research entails claims about collective intelligence, the questions about the nature and causal structure of this phenomenon remain unanswered. In the next section, we will tackle these questions by providing a conceptual framework for thinking about the preconditions of collective intelligence in light of recent theoretical and empirical research on collective intentionality in social ontology and social cognition.

3. THE COGNITIVE PRECONDITIONS FOR COLLECTIVE INTELLIGENCE

In the previous section, we put forward an intuitive 'indicator' of when groups are smarter than their individual components. In cases of genuinely collective intelligence, the activity and outcomes of a group appear as something which we would want to say that they, the people involved in the group, did it together. Although there is a sense in which the same happens when labour is distributed among individual agents in society, the point of collective intelligence is not so much that each person must contribute specific ideas and competences in order to achieve something that exceeds their own capacity, as the fact that people do it in interaction rather than in isolation. It is the integration of personal know–hows and goals that lead to collective outcomes which are irreducible to the sum of the parts as interaction stretches the boundaries of individual expertise. But, what makes groups smarter? What preconditions must be in place for collective intelligence to emerge from the interaction of jointly interacting individuals? The problem about the nature and origin of collective intelligence can be tackled at many different levels of analysis, which by and large fall into the scope and motives of 'social ontology'. How persons relate socially to one another and to the social facts they constitute is the problem of social ontology (Pettit 1993). In this section, we distinguish between two possible ways of articulating the foundations of collective intelligence, drawing on classic and modern debates about the relationship between mind and society.

3.1 The problem with emerging social phenomena

Questions about 'what there is' in the social world date back to the earliest stages of philosophical investigation, but it is only in the modern age that they have received systematic formulations in the work of, for example, early sociologists like Emile Durkheim (1893). The first difficulty that we encounter in analysing the foundations of collective intelligence concerns the meaning of emergence. In general terms, 'emergentism' is the view that complex phenomena depend for their existence and identity on the component parts and their interactions. Yet, an account of the organisation and outputs of these macro–phenomena is not fully exhausted by an explanation in terms of their micro–parts (Sawyer 2002). This contrast is epitomised in the old truism that there cannot be society without individuals, yet society is more than the sum of the individuals (Sawyer 2001). Notice that there are two meanings to this claim, ontological and epistemological. The former is about the conditions that are needed for the phenomenon to be what it is, as a matter of objective fact; the latter is about the concepts that we employ in understanding and explaining the phenomenon. To illustrate the difference, consider the case of corporate organisations. Companies are (ontologically) constituted by individuals, each with distinctive backgrounds and objectives, functional roles and expectations. As complex systems, however, companies exhibit features than none of their employees display. For instance, when we hold them legally, if not morally, responsible for their conduct, as if they were minded creatures like us, we imply that they can

'think' and 'act' in the pursuit of goals that might even diverge from individual preferences. This amounts to the (epistemic) claim that the behaviour of a corporate organisation, as well as of its underlying mechanisms and processes, is not fully exhausted by an account of the psychological attitudes of the people working in it.

3.2 From supra–individual cognition...

In light of this distinction, a theory of the cognitive preconditions of collective intelligence can be seen as a response to at least two different ways of conceiving of the nature (ontology) and explanation (epistemology) of collective intelligence. In the present discussion, we use the term 'cognitive' to refer to all sorts of attributes of the mental states of interacting agents, including cognitive (i.e. beliefs), conative (i.e. intentions) and affective (i.e. emotions) attitudes. The first response characterises collective intelligence as an attribute of the bearer holding the relevant attitude. In the case of corporate organisations, to continue with the same example, some authors argue that when we make truthful attributions of mental attitudes and ethical dispositions to corporates, as testified by evidence of their explanatory and predictive success in various social–science research programmes like rational choice theory (List and Pettit 2012), we hold a realist and not purely metaphorical or fictional, i.e. 'as if', stance (Tollefsen 2002). That is, to attribute aspects of mentality – thinking, reasoning, imagining, feeling, intending, and so forth – to complex organisations, follows from the fact these multi–person entities become autonomous and minded 'subjects' once they emerge from the single parts. On this approach, the notion of collective intelligence would refer to the capacity of plural subjects to think and act in ways that bring about different results from the thoughts and actions of their individual components. Although this view has been vulnerable to a number of critiques, much research has been done across the cognitive sciences, lately, to articulate a scientifically reasonable account of forms of 'macrocognition' or collective mentality (Huebner 2014; Tollefsen 2015). One general problem facing these explanatory attempts, however, is to make sense of the idea that plural subjects have their own attitudes emergent from those of their individual constituents. In fact, provided that plural subjects intend and do what their members individually intend and do, we could not explain their emergent attitudes unless we know what it takes for the single agents to understand one another in interaction. One consequence is that the view that collective intelligence is a feature of plural subjects, no matter how scientifically and philosophically robust, entails a response to the question about why, and how, individuals share mental states in the first instance when they engage in social interactions.

3.3 ... to individual social cognition

To shift the focus of social ontology towards the social cognition of individual agents, namely the cognitive mechanisms and processes whereby they each understand the others' minds as they co–ordinate in interaction, suggests an alternative way to articulate the concept of collective intelligence. On this approach, the key would be to investigate the type of psychological attitudes that interacting agents feature in the 'meeting of minds' characteristic of truly joint activities. On the assumption that these intentional attitudes are collective, how can single agents act collectively out of individual attitudes? Before we elaborate on the meaning and relevance of collective, or shared, intentionality for an account of collective intelligence, it is worth noticing that this approach differs from the previous one in at least one important respect. The analysis now takes the psychology of the individuals involved in the group activity, rather than the psychology of an allegedly supra–individual plural subject, as its explanatory unit. Nonetheless, the motivation of early proponents of analyses

of collective intentionality for endorsing what seems more of an individualistic approach is, in fact, similar to the Durkheimian lesson about the relationship between (individual) minds and complex social phenomena (Olen and Turner, manuscript). The problem, still, is to provide a social–scientific account of the fact that emergent entities are constitutively grounded in their individual components, whereas they acquire attributes that make them appear more than the sum of the parts. But the answer differs in that most discussions of collective intentionality discard (ontologically) controversial claims about entities such as group minds, i.e. collective consciousness, while they turn out to be constrained by the results of cognitive science, which has risen to the front stage of research in the (social) mind in the second half of the 20th century.

3.4 The concept of intentionality

One clarification about the concept of collective intentionality concerns the meaning of 'intentionality'. Intentionality – from the Latin verb 'intendere', which literally means to stretch (Crane 2001: 9) – is a technical term of philosophical jargon used to capture, not just the idea of intending (to do) something when we engage in action, but more generally the power of all sorts of psychological states to be about objects or states of affairs in the world. For example, if you desire to talk to someone after work, and wonder whether you have time for it, because you fear that the current meeting is going to last too long – your thoughts are 'directed toward' something (the person you want to meet, time, the current meeting). Another way to exemplify the concept of intentionality is to say that intentional mental states represent things in the world as thus–and–so. What, then, makes the intentional attitudes of individuals shared in the course of interaction? In ordinary language, the thoughts of any two people are shared when they represent things in the world in the same, or similar, ways. Indeed, parties to a joint action must share some understanding of aspects of the task that are essential for achieving the common goal, like the goal itself, their own contributions, how action unfolds in space and time, and so forth, for the interaction to occur smoothly. But if 'sameness' of understanding were a precondition of collective intelligence, then it would be impossible to make sense of the fact that geographically dispersed agents, each with very distinctive cultural and social perspectives on the world, could pool their ideas together and produce a collective outcome. Successful manifestations of mass collaboration and collective intelligence – think about Wikipedia, for instance – show that it is not necessary that there is an exact replica in each agent's mind of the meanings of the project. Jointly interacting individuals do not need to 'see', that is represent, things in the world in the same way for them to be able to come up with a macro result which may turn out to be (epistemologically) irreducible to features of individual psychology. According to pragmatic theories of language, communication is a process by which utterances and gestures produce relevant inferences that are supported by the context. On the inferential view, utterances are not signals but pieces of evidence about the speaker's meaning, and comprehension is achieved by inferring this meaning from evidence provided not only by the utterance but also by the context. (1995 : 2) In pragmatics, 'context' is not just the physical ambience of play; rather, it refers to all features of the perceptual and cognitive environment of the agent performing communicative acts, which allows the hearer to infer the speaker's meanings (Sperber and Wilson 1995).

Thus, collaborative projects need not expect all members to share the same understanding of aspects of the action 'scene' for them to be successful. Instead, communicative processes are framed in environments that elicit a sense of relevance and the possibility of generating inferences that are meaningful for the participants, whatever the content might be.

3.5 **We-intentionality**

In the late 1980s, early theorists of collective intentionality noticed that most social science research has been characterised by an implicit, though pervasive, assumption about what counts as relevant for people acting together in social settings. The assumption is that their behaviour is always given from the perspective of the first-person thinking and experiencing subject I. Thus, collective behaviour has long been seen as guided by representations in the minds of the agents that specify aspects of the action scene in an 'egocentric', or first-person singular, mode. On the contrary, most episodes of sociality, ranging from real-time co-ordination to planned and orchestrated co-operative behaviour, to strategic interaction and team-reasoning, require the single agents to adopt a 'we-perspective' instead. Namely, individuals can conceive of things in the world, including themselves and the other people, from the perspective of the group (the 'we') of which they happen to be members at any given time. So, instead of having a thought expressed by the words 'I intend to do this, or that, with you.' people's individual behaviour can be guided by a shared understanding of the world. The content of this understanding is articulated in terms of the idea of taking each other's perspective into account, and it is expressed in a way that makes prominent the use of the first-person plural pronoun ("We intend to do this together"). In his writings on the logic of morality, philosopher Wilfrid Sellars called this psychological attitude 'we-mode' and characterised it as the primitive state of consciousness entertained by one person thinking and acting 'as-one-of-us' (1963, 204–5; emphasis in original). Along similar lines, in *The Construction of Social Reality*, perhaps the most articulate and well-known attempt to provide a grand theory of social ontology in contemporary philosophy, John Searle defined any fact involving collective intentionality as a social fact (1995: 172) and construed the relevant psychological attitude as an explanatory 'building block' necessary to make the structure of collective-intelligence phenomena fully intelligible (Gallotti 2012).

According to the 'we-mode hypothesis', which has rapidly gained consensus in philosophy and science, when any two persons interact, they experience their thoughts and actions as being about something that is to be jointly achieved – a sense that 'what you and I do, we do it together' (Gallotti and Frith 2013). In broad strokes, this experience amounts to what psychologist Michael Tomasello presents as the "uniquely human sense of 'we', a sense of shared intentionality" (2009, 57). Yet, whilst current research on the we-mode has benefited from theoretical and empirical advances in social cognitive research (neuroscience), mostly the central claim of the hypothesis is hardly new. The agentive experience of shared thought and agency in the first-person plural was at the core of a well-established tradition of studies dating back to the work of phenomenologists, like Edmund Husserl, on issues of intersubjectivity, empathy and we-intentionality (Zahavi 2014). Furthermore, the phenomenon of group behaviour and thinking are at the focus of a rich experimental literature on social identity and categorisation, initiated by social psychologist Henri Tajfel at the end of the 1960s. There is evidence that, in many episodes of everyday life, human beings can express who they are in terms of 'we', as well as 'I'. When they represent aspects of the world from the perspective of their interacting partners, as a 'we', the individuals have access to a broader repertoire of information about things, which then result in different patterns of behaviour, than if they were acting outside of the group or, most strikingly, as part of another group. This effort of articulating the concept of we-mode beyond intuition and in a manner that would be consistent with the claims of cognitive science reveals something important about the foundations of collective intelligence. It might well be the case that the capacity of geographically proximate as well as remote agents to acquire intelligent traits, as a group, is in fact grounded in a cognitive specialisation of individual minds for sharing mental states, which has evolved and develops through life span in response to environmental pressure towards more social forms of thought and action.

3.6 **Joint attention**

The we-mode is a cognitive mechanism enabling individual agents to share attitudes and jointly produce forms of collective intelligence. As an illustration of the functioning and relevance of this mechanism, let us examine the ability of individuals, most notably human infants, to share attentional states with adults, i.e. their caretakers. 'Joint attention' designates a suite of phenomena of triangular interaction based on a perceptual relation between two people attending to a third object together (Moore and Dunham 1995; Tomasello 1999). As data from studies in developmental and comparative psychology and psycholinguistics show, at around nine months of age human infants start to experience the various components of the world differently. There is an increase of complexity in their ability to understand the set of causes and mechanisms lying behind the behaviour of their caretakers, which suggests a form of interpersonal engagement that outgrows ritualisation and manifests itself in the ability to engage in triadic episodes of interaction. Individuals attending to the same object together are now attributed a level of understanding of each other's intentional behaviour that, given the age, is supposed to play a 'scaffolding' role for later forms of social cognition. There have been many studies in psychology to capture the onset of children's capacity of social understanding in terms of preconditions such as mindreading, role-taking, imitation, and more generally the ability to take another person's point of view into account. The main constrain is that the cognitive and communicative, i.e. pre-linguistic, faculties of nine-month olds are still limited in many fundamental respects. It is at this point that the question about the mechanisms of joint attention, where the psychological understanding ascribed to infants originate from, becomes urgent. The question is what enables infants to engage with adults in ways that bring about a novel understanding of the world and, therefore, new possibilities for individual action, brought to light by engagement in a shared space of thoughts and experiences – one in which the agents ('we') each know that what they attend to, they do it together (Tomasello 1995).

Indeed, Moll and Tomasello (2007) argue that the notion of perspective depends on us first having a joint attentional focus, which we may then come to view differently (otherwise we just see completely different things). Coming back to the question about the cognitive preconditions of collective intelligence, as the quotation here above suggests, there seems to be a parallel between the capacity of individual agents to act in the pursuit of a joint goal, while they each play an individual role, and the capacity to take another's perspective into account while maintaining a joint attentional focus. To say it more simply, the fact that we can 'tune' into the others' minds and act in a shared attentional space allows us to envisage a collective goal while holding on to individual-level specific roles. This 'twofold' state of mind, influentially illustrated by Tomasello in *A Natural History of Human Thinking* (2014), seems to be a recurrent feature of most examples of collective intelligence in which distant participants exhibit the ability to focus and adhere onto a set of goals – such as mapping an area of Africa by focusing on a predefined set of tools and representations – as well as the ability to represent one's own task and meanings.

The role of joint attention in the development of full-blown cognitive functions, which are essential for the development of languages, sophisticated cultural artefacts and social norms, exemplifies the idea that phenomena of sociality rely on certain cognitive preconditions for their constitution and persistence, where 'cognitive' means that these preconditions consist in attributes of the mental states of interacting agents.

4. THE CULTURAL PRECONDITIONS FOR COLLECTIVE INTELLIGENCE

The discussion in the previous sections does not yet include fundamental components in the process of collective intelligence which are **cultural artefacts**. Artefacts which include everything from writing, to tools and the built environment, but also rituals, stories and science, are public expressions of mental representations. Humans transcend the boundaries of the individual mind and make their mental states available to others in a variety of forms. In his book *Explaining Culture* (1996) Dan Sperber explains the role of public representations in human cognition. Individual, mental representations become public when they are spoken, written, or inscribed in tools and gestures. Some artefacts, as Sperber says below, are produced in order to cause mental representations in other people.

> *"The mental things involved are mental representations and processes. These mental representations and processes may cause behaviours that alter the environment in ways that can be perceived and thus serve as stimuli to further cognitive processes. Some of these environmental changes are perceptible as processes, e.g. bodily movements, speech sounds; others are perceptible as stable states of the environment, e.g. the presence of paths, buildings, artefacts, or writings. I will call all such perceptible behaviours and effects of behaviour 'public productions.' Some public productions, for instance utterances, signals or pictures, are produced for the purpose of being perceived and causing mental representations. These 'public representations' form a particularly important subclass of public productions."*
>
> **Sperber 2001: 301**

Sperber repeatedly points out that public representations do not elicit exact replicas in different people's heads: an utterance does not generate a replica of my mental representation in someone else's mind. Utterances or any other public production trigger inferences which depend on context and prior knowledge (Grice 1958, Sperber and Wilson 1989). It is not necessary for successful communication to happen, to generate replicas in other people's minds. Both speaker and listener are aware that they can rely on context to generate relevant inferences. The context, in a certain sense, is doing part of the work and the capacity for joint attention and shared intentions supports the process. What is important for our discussion, is that publicly available productions can be perceived and used by different people to collaborate and co–ordinate, in a way that private mental states cannot. In fact we would argue that the **materialisation of mental states in public representations** (whichever form they take) **is the only way joint activities can come about**.

Another important aspect of public representations is that they can become objects of thought in themselves. A piece of writing, an utterance, a number, a drawing allow people to reflect upon the representations themselves. This form of abstraction has a huge impact on the mind and the way we think. The anthropologist Jack Goody in *The Domestication of the Savage Mind* (1978), describes how writing, historically, allowed human beings to abstract their thoughts, creating the basis for new forms of classificatory thinking which ultimately

led some societies to develop systems of administration. Lists of goods became the basis for clustering and classifying items, which in turn led to significant social transformations in trade and agriculture. Similarly Bruno Latour mentions how maps, charts, graphs, and tables enable the superimposition of representations on otherwise incommensurable items (Latour 1986) making them understandable, comparable and abstractable.

Once representations are public and in forms that are accessible and sharable, they can also become the building blocks for other artefacts. Lists can become categories, tools can be assembled. People are not only capable to metacognitively reflect on symbolic systems, they are brilliant in using them to build new representations and new artefacts. The materialisation of mental states, therefore, also make them modifiable and reusable.

4.1 The epidemiology of ideas

The availability of public representations in time and space, makes them more stable and less likely to be lost during transmission. While individual memory is quite short, the materialisation (in writing, tools, built environment, etc.) extends the duration of the representation in time as it can be revisited many years or centuries later. Public representations can also be potentially transferred to people who are not in co–presence via different media, as is happening through the internet. This transferability over time and space forms the basis for cultural transmission of ideas and practices.

However, not all public representations circulate successfully, some disappear immediately and some are around for centuries. As Sperber says we have particular cognitive predispositions which explain the epidemiology of ideas:

> *A human population is inhabited by a much wider population of mental representations. The common environment of that population is furnished with the public productions of its members, some long lasting, like buildings, other ephemeral, like the sounds of speech. Particularly important among these productions are (tokens of) public representations. Typically, productions have mental representations among their causes, and mental representations have productions (in particular public representations) among their causes. There are thus complex causal chains where mental representations and public productions alternate. In many cases, representations (mental or public), occurring in these causal chains inherit some of the semantic properties of the representations (mental or public) of which there are causal descendants. A variety of inter–individual processes bring about this match between causal and semantic relationships. Processes of imitation and communication can be described as having the function of bringing about such semantic similarity. However, neither communication nor imitation are replication mechanisms, so that, I argue, standard selectionist models don't properly apply. The psychological dispositions that make possible the individual formation and social distribution of representations, are, on the other hand, evolved mechanisms to be explained in straightforward Darwinian terms.*
>
> **Sperber: 1996**

This means that within collectives some ideas circulate more widely and better than others and that external supports to memory and cognition are an integral part in eliciting internal representations that are relevant to multiple people. In other words there are some ideas which minds are more predisposed to entertain and share. Culturally shared knowledge usually refers to those cases. Some religious beliefs which can be found across cultures (Boyer 1996), children's rhymes or games, memes, etc., have characteristics that make them contagious and easily transferrable. Other public representations, such as advanced algebra or genetic molecular biology, require expert training and considerably more effort to be acquired. Educational systems have been developed among other reasons, to ensure that more complex ideas are circulated in society. This is not trivial when thinking of collective intelligence because we can expect that some representational formats will be more likely to be quickly and widely adopted and distributed. One of the functions of computational systems is to transform the complexity of the conceptual frameworks being operated upon, thus increasing the likelihood of distribution. When complex spatial tasks involving measurements, ratios, etc., are transformed into a process of drawing a line, they become activities that are accessible to many.

5. THE DIGITAL PRECONDITIONS FOR COLLECTIVE INTELLIGENCE

Coming back to the Missing Map example, the role of the mapping platform Openstreetmap and the editing tools it provides (ID , JOSM and Potlatch2), is extremely significant in co-ordinating the efforts of the volunteers. The platform offers a visualisation of the territory, a set of bounded drawing tools, a graphical language to represent geographical elements (lines for roads, circles for homes, etc.), measuring devices that provide the unit to translate the graphical elements in geographical proportions, a bi-dimensional view of the territories, etc., etc. The platform also embeds some of the rules for co-ordinating the efforts of the users, such as colour coding areas that require work, areas that have to be evaluated, and those that are finished and should not be opened. The system stores the work centrally, it allocates a certain amount of time to tasks, it provides 'ownership' of an area for the period on which it is worked on by a volunteer by blocking access to other users.

By inscribing in the software these rules it relieves participants from the task of negotiating among each other issues of order, time, and control. By providing the code and tools which rely on established mapping conventions, it allows volunteers to embrace a professional system of notation without having to go through a lengthy professional training to acquire it. These along with many other elements inscribe into the system norms, rules, conventions and processes that do not need to be made explicit or negotiated between the collective that contributes to the mapping exercise. Furthermore, the platform ensures continuity across time and location. Tasks can be done at different times in an unrelated fashion by different people, because the system ensures the storage of the work, the constant update and the continuity between apparently unrelated events. The platform Openstreetmap is therefore designed to support distributed collaboration, the software introduces forcing functions and sequences that literally limit or enable certain actions, artificially imposing sequences and dependencies between the activities of users. This inscription of social norms in the system is nothing new and corresponds to what has been extensively analysed in literature on socio-technical analysis and in Actor Network Theories (Long 1986, Bejkers 1992 and Callon 1986 among others).

5.1 Distributed cognition

Beyond the entwining of the social and the technical, there is also a cognitive dimension that is worth discussing here. A very relevant approach to collective intelligence is Hutchin's (1992) theory of Distributed Cognition. This approach claims that cognitive processes are always distributed among people, artefacts and processes and that significant parts of the computational work involved in any activity is carried out in interaction with external systems. Artefacts can be supports for memory, embed partial solutions, represent previously computed solutions, but in all cases are an integral part of realising a cognitive task. For example, measuring the length of a table is done in interaction with a ruler which is an artefact that embeds some prior cultural knowledge on the conventional length of a meter

and thus simplifies the task (to one of matching two objects, adding or multiplying, and translating the result in a measure). Think of the complexity of the task if each time we were to devise a system of measurement that is transferrable and shareable. Artefacts therefore, like procedures, embed solutions found by others, and transform tasks into something manageable by an individual or a group at a given moment. Systems for collaboration such as a Debian repository, a mapping system, a document editing system, tend to embed solutions to recurrent co–ordination problems such as how to avoid multiple people working on the same area. Interestingly, open source projects such as these are constantly looking and finding new solutions to problems as they emerge (fixing bugs and adding functionalities at the requests of users) thus giving a 'live' demonstration of how culture works. Adjustments and improvements to tools and artefacts in the past could take decades but now are visible in hours.

Our cognitive system, as mentioned in previous sections, is not only geared for interacting with others, it is also capable of distributing knowledge and computations onto systems of public representations (Sperber 1996). In this perspective, to come back to our argument that collective intelligence emerges from the interaction with other entities, cognition results from the interaction of people processes and artefacts.

Realising that artefacts and processes are an integral part of bringing about collective intelligence is crucial. The internet and the services that rely upon it, are not simply means to bring people together to think, discuss, share and act. The tools with which collectives interact and work are a part of the creative process because they extend the mind (Clark 2002), both individually and collectively. We would argue that in the case of distant networked individuals, they constitute both a precondition for collective intelligence and a result of the process. These common systems materialise and make manifest the intentions of the collective and more importantly the solutions. By materialising the organisational principles, the objects to be computed and the solutions reached collectively they contribute to framing the collective endeavour.

Not all internet systems however, have the capacity to support collective intelligence. We have witnessed numerous failed attempts to generate 'collective wisdom' by bringing together citizens on loosely designed discussion platforms. Public forums to discuss policies, budgets, health or other crucial topics often fail to generate the collective creative output that we witness in Openstreetmaps or Wikipedia. Often these platforms simply offer text boxes where individuals can write opinions and reply to others. In practice they offer an open environment for debate with limited functions for storing and grouping the discussions. In these systems the burden of co–ordination and synthesis is on the participants and it comes as no surprise that participation is low or dwindles after few interactions.

These types of environments do not offer objects on which to jointly operate, they do not provide visualisations of the common results, they do not make tangible the 'whole' emergent results. What they offer, on the contrary, are exactly the contrary elements which display the individual contributions, the parts or fragments. By not providing a sense of the emerging collective output they inhibit the creation of the collective itself. From these experiences what we can learn is that for digital systems to become components of the collective extended mind and foster collective intelligence, they have to have some characteristics. These characteristics include the possibility to modify the tools, aggregate the contents and contributions, retrace the history, visualise the aggregations and the single elements, include analytic tools and allow the negotiation of terms. In other words the systems should support learning and metacognition through synthesis and modification mimicking the process of cultural evolution.

6. THE ORGANISATIONAL PRECONDITIONS FOR COLLECTIVE INTELLIGENCE

Collective intelligence emerges in social contexts which are never isolated but operate in interaction, opposition or relation to existing institutions that want to leverage collective action for an objective. Missing Maps, MSF and the Red Cross have an overarching objective of mapping areas of the world in which they operate. They are therefore mobilising volunteers to help them in achieving this goal in a much shorter period than would be possible without Openstreetmap and a crowd of networked people. These organisations have chosen a certain mode of collaboration, of knowledge sharing and management based on the integration of distant participants and local networks of volunteers. Openstreetmaps offers an environment that requires limited governance on the part of the humanitarian organisations. As an open source project, it follows some of the classic approaches of open source development (Raymond 1999): the data produced is free and can be reused and modified by others; the contributors are not remunerated and have a loose affiliation with the project; volunteers contribute with content but are also involved in the constant improvement of the tools themselves.

Chris Kelty (2008) calls this a *"recursive public,"* a concept to describe the unique phenomenon of a public which builds and modifies the tools and conditions that enable its existence.

> *"A recursive public is a public that is vitally concerned with the material and practical maintenance and modification of the technical, legal, practical, and conceptual means of its own existence as a public; it is a collective independent of other forms of constituted power and is capable of speaking to existing forms of power through the production of actually existing alternatives. ... Recursive publics, and publics generally, differ from interest groups, corporations, unions, professions, churches, and other forms of organization because of their focus on the radical technological modifiability of their own terms of existence."*
>
> **Kelty 2008: 15.**

The constant refashioning of tools and representations is an integral part of the collective intelligence of the project as it serves as a constant rekindling and redefinition of the joint objectives.

> *"The second reason I use the concept of a recursive public is that conventional publics have been described as 'self-grounding,' as constituted only through discourse in the conventional sense of speech, writing, and assembly. Recursive publics are 'recursive' not only because of the 'self-grounding' of commitments and identities but also because they are concerned with the depth or strata of this self-grounding: the layers of technical and legal infrastructure which are necessary for, say, the Internet to exist as the infrastructure of a public. Every act of self–grounding that constitutes a public relies in turn on the existence of a medium or ground through which communication is possible – whether face-to-face speech, epistolary communication, or net-based assembly – and recursive publics relentlessly question the status of these media, suggesting that they, too, must be independent for a public to be authentic."*
>
> **Kelty 2008 : 16**

6.1 Models of governance

The free software, open source models are not the only models that are being invoked to organise and elicit collective intelligence. There are various models emerging in the digital world, ranging from classical vertical structures where a central unit defines the objectives and tasks, categorises and analyses the contributions, synthesises the results, and eventually reaps the benefits of the value generated by the collective. Some commercial big data companies or some social media platforms follow this approach and use data generated from the analysis of participant behaviours for their own benefit. In this model, knowledge and information is extracted from myriad individuals but the assembling and analysing is done centrally by an organisation which maintains control over the whole. This approach is compatible with division of labour as it presupposes an overseeing entity defining the overall plan.

Other more horizontal models, suggest modes of functioning in which tasks and activities emerge from a process of deliberation and where iterations are an integral part of the collaborative framing of issues and objectives. These are mechanisms more closely associated with new forms of political engagement (M5S in Italy, Podemos in Spain etc.) where supporters are given online space to propose initiatives and actions. The challenge in this case is to ensure that the individual contributions are assembled and become coherent.

6.2 The open source model

The most successful cases of online collaborations tend to have a co–ordinating component that structures the tasks and analyses and assembles the contributions. Wikipedia, and the open source development foundations for instance, have progressively developed governance structures to ensure the quality of the content produced, avoid degradation and support their means of existence. Committees, boards and local chapters participate in the steering of the projects and foundations. They are responsible for funding but also liaising with other institutions, scoping the work and main tasks, defining the rules of participation and representing the organisation externally. Interestingly, most open source organisations have some form of governing body to carry out some of these functions. These co–ordinating structures usually work on principles of meritocracy with the best developers or most

committed writers having some form of control on the content and rules of contribution. We would want to argue that these bodies play not only a social and institutional role but also an epistemic one by providing the framework of activity and meaning ensuring the continuity and perennity of the projects.

As we have mentioned in the previous section, the challenge for collective intelligence to emerge in social groups is not only to create settings for sharing and communication, but also to provide the **means for knowledge to be made public, to be assembled, sedimented and reflected upon**. There is a fine balance between preservation and modification: ensuring the continuity of the knowledge produced, while allowing constant refashioning and appropriation. In order to move from division of labour to something better than the sum of individual contributions, the systems of knowledge that are being constructed have to be **collectively visible, modifiable, discardable, operational and reusable**. This is exactly the process we see in action in cultural transmission where ideas are made manifest, constantly refashioned, rejected, and reassembled. It is also what we are witnessing in an accelerated form in many open source development projects.

Any reflection on the governance models to sustain processes of collective intelligence, should therefore consider what are the legal, economic and administrative requirements for an open and reusable system of knowledge creation and sharing which recognises the multiple entanglements of knowledge systems (Coombe 2008, Cohen 2012). The ongoing battles around intellectual property law (Cohen 2006) are part of a crucial debate on the mechanisms to ensure that collective intelligence can in fact emerge and constitute a good example of why governance models are part of the prerequisites for collaboration.

If we follow the argument, as we have expressed it in this paper, that collective intelligence relies on a cognitive apparatus that allows us to join in other peoples' intentions (and actually entertain we–intentions), but also on a system of public representations, tools and artefacts that enables the reflection and transformation of knowledge, we set ourselves squarely in the wider discussion about cultural and scientific transmission. Within this realm, there is little doubt or controversy on the need for open and reusable information for cultural and scientific development, and the question of the organisational model must simply be addressed from this perspective. As we stated in the beginning of the paper, collective intelligence is an emergent property of interaction between people which goes beyond division of labour and as such requires a social model that is appropriate for sustaining interactions. When we add to this equation the need to ensure that the tools and public productions that enable the co–ordination and elaboration of knowledge, are available and transformable by the collective, we find ourselves inevitably with a different question. What forms of commons can best ensure the circulation and appropriation of ideas in order for collective intelligence to flourish ?

7. CONCLUSIONS

This paper has attempted to identify some of the preconditions for collective intelligence to emerge. We have challenged the idea that division of labour, in which each individual contributes a small part of work to a predefined whole, can explain the quality of the product emerging from the interaction between networked people. The most common and naive use of the term collective intelligence is often associated with the idea of an overarching entity 'pooling' people's brains or knowledge, extracting from a mass of individuals elements of information, expertise or computational capacity. This model does not, in our view, capture the full extent of the transformations that occur when networks of people interact around common goals as for instance in the case of the Missing Maps project we have been analysing. Our objective has not been so much to qualify the nature of this different output, as much as explain how it is possible for a group of geographically remote people, often volunteers with a loose organisation, to co-ordinate and produce complex solutions and fashion their own tools. In the case of the new digitally enabled forms of collective intelligence we have witnessed in recent years, division of labour is not an organisational or cognitive model that applies. These new forms have an emergent property that in our view is far more evocative of the processes underlying cultural transmission than industrial production.

We have argued that humans have a cognitive predisposition for collaboration which is manifested in the capacity to elaborate we-intentions and thus joint objectives. This process is sustained by the production of public representations such as utterances, writing, tools or other artefacts which not only make manifest individual mental states but become material forms that can be revisited, modified, reflected upon, transmitted by multiple people. Digital networks are obviously greatly extending and accelerating the distribution of such public productions and computational systems are allowing for an infinitely greater potential of transformation. Collectives are not only producing content but are actively working on the tools, notations and organisational models to bring these about. If we take the Missing Maps project, the intelligence of the collective is being expressed not only in the production of the maps, but also the software, the notation, the labeling, the usage modes, the organisation and learning processes. The full visibility, access and modifiability of the system are a source of motivation for the participants and the basis for learning and improving. The collective learning is materialised in the tools themselves which in turn gives rise to improved maps, which then require the tools to be improved. This process of collective learning and refashioning is, in our view the most compelling description of collective intelligence.

When thinking about the prerequisites for collective intelligence it seems therefore essential to stress the interdependence of the technology with a certain governance model in order to enable cognitive predispositions for cultural transmission to fully emerge.

REFERENCES

Bijker, W. and Law, J. (Eds.) (1992) 'Shaping Technology, Building Society: Studies in Sociotechnical Change.' Cambridge MA: MIT Press.

Boyer, P. (1996) 'The Naturalness of Religious Ideas : a cognitive theory of religion.' Oakland CA : University of California Press.

Callon, M., Law, J., et al. (Eds.) (1986) 'Mapping the Dynamics of Science and Technology: Sociology of Science in the Real World.' London: Macmillan.

Clark, A. and Chalmers, D. (1999) The Extended Mind. 'Analysis.' 58(1): 10–23.

Clark, H.H. (1992) 'Arenas of Language Use.' Chicago IL: University of Chicago Press.

Cohen, J. (2012) 'Configuring the Networked Self: Law, Code, and the Play of Everyday Practice.' New Haven CT: Yale University Press.

Cohen, J. (2006) Copyright, Commodification, and Culture: Locating the Public Domain. In Hugenholtz , P.B. and Guibault, L. (Eds.) 'The Future of the Public Domain.' Pp.121–66. The Hague: Kluwer Law International.

Coombe, R. (1998) 'The Cultural Life of Intellectual Properties: Authorship, Appropriation and the Law.' Durham: Duke University Press. (Reprinted 2008.)

Crane, T. (2001) 'The Mechanical Mind.' 2nd Edition. London: Routledge.

Durkheim, E. (1893) 'Sociology and Philosophy.' Glencoe IL: Free Press.

Gallotti, M. (2012) A Naturalistic Argument for the Irreducibility of Collective Intentionality. 'Philosophy of the Social Sciences.' 4: 3–30.

Gallotti, M. and Frith, C.D. (2013) Social cognition in the we-mode. 'Trends in Cognitive Sciences.' 17: 160–165.

Grice, H.P. (1957) Meaning. 'Philosophical Review.' 66(3).

Huebner, B. (2014) 'Macrocognition.' New York: Oxford University Press.

Hutchins, E. (1995) 'Cognition in the Wild.' Boston MA : MIT Press.

Kahneman, D. (2012) 'Thinking, Fast and Slow.' London: Penguin.

Kelty, C. M. (2008) 'Two Bits: The cultural significance of Free Software.' Durham NC: Duke University Press.

Latour, B. (1991) Technology is Society Made Durable. In Law, J. (Ed.) 'A Sociology of Monsters. Essays on Power, Technology and Domination.' Sociological Review Monograph. London: Routledge. 38: 103-131

List, C. and Pettit, P. (2011) 'Group Agency.' Oxford: Oxford University Press.

Moll, H. and Tomasello, M. (2007) How 14- and 18-month-olds know what others have experienced. 'Developmental Psychology.' 43(2), 309-317.

Moore, C. and Dunham, P. J. (Eds.) (1995) 'Joint Attention: Its Origins and Role in Development.' Hillsdale, NJ: Erlbaum.

Olen, P. and Turner, S. (under review) 'Durkheim, Sellars, and The Origins of Collective Intentionality.' Manuscript.

Pettit, P. (1993) 'The Common Mind.' Oxford: Oxford University Press.

Sawyer, R. K. (2002) Durkheim's Dilemma: Toward a Sociology of Emergence. 'Sociological Theory.' Vol.20, Issue 2, 227-247.

Sawyer, K. (2001) Emergence in Sociology: Contemporary Philosophy of Mind and Some Implications for Sociological Theory. 'American Journal of Sociology.'

Searle, J. (1995) 'The Construction of Social Reality.' New York NY: Free Press.

Sellars, W. (1963) Imperatives, Intentions, and the Logic of 'Ought'. In Castañeda, H.N. and Nakhnikian, G. (Eds.) 'Morality and the Language of Conduct.' Detroit MI: Wayne State University Press.

Sperber, D. and Wilson, D. (1995) 'Relevance. Communication and Cognition.' 2nd edition. Oxford: Blackwell.

Sperber, D. (1996) 'La Contagion des Idées.' Paris : Odile Jacob.

Tollefsen, D. (2015) 'Groups as Agents.' Polity Press.

Tollefsen, D. (2002) Organizations as True Believers. 'Journal of Social Philosophy.' 33: 395-411.

Tomasello, M. (2014) 'A Natural History of Human Thinking.' Cambridge MA: Harvard University Press.

Tomasello, M. (2009) 'Why We Cooperate.' Cambridge MA: MIT Press.

Tomasello, M. (1999) 'The Cultural Origins of Human Cognition.' Cambridge MA: Harvard University Press.

Tomasello, M. (1995) Joint Attention as Social Cognition. In Moore, C. and Dunham, P. (Eds.) 'Joint Attention: Its Origins and Role in Development.' Hillsdale NJ: Erlbaum.

Zahavi, D. (2014) 'Self and Other.' Oxford: Oxford University Press.

Nesta is an innovation charity with a mission to help people and organisations bring great ideas to life.

We are dedicated to supporting ideas that can help improve all our lives, with activities ranging from early-stage investment to in-depth research and practical programmes.

Nesta...

1 Plough Place
London EC4A 1DE

skills@nesta.org.uk
@nesta_uk
www.facebook.com/nesta.uk

www.nesta.org.uk

Nesta is a registered charity in England and Wales with company number 7706036 and charity number 1144091. Registered as a charity in Scotland number SCO42833. Registered office: 1 Plough Place, London, EC4A 1DE.

CC BY NC SA