

Boundary objects, trading zones, and stigmergy: the social and the cognitive in science

Ric Sims

Independent researcher

Abstract

The main proposal of this paper is that boundary objects and the trading zones in which they occur are the analogue of pheromone trails in the foraging of a termite colony. The colony can be construed as a *stigmergic* system where the traces of the actions of individual termites coordinate their further actions without the existence of any central control or planning structures. The coordinated systems approach proposed by this paper lends support to the idea that such a system is minimally cognitive in the sense that it is responsible for goal-directed behaviour. Boundary objects and trading zones in scientific practice play a similar functional role to pheromone traces because they are responsible for the same kind of coordination. This approach therefore provides a cognitive account of the social notions of boundary object and trading zone without making representationalist or computationalist assumptions. Moreover, it is scale-invariant – the same analytical technique can be applied at multiple scales simultaneously. It therefore provides a framework for an understanding of the complementarity of cognitive and social approaches to scientific investigation and points to areas for further ethnographic research.

1 Introduction

The starting point of this paper is what could be called cybernetic enactivism. This is an approach that takes a cognitive system to be a set of semi-autonomous processes coordinated by environmental variables and responsible for goal-directed behaviour. By applying such a coordinated systems approach to scientific practice, the paper seeks to understand some issues in the philosophy of science that bridge social and cognitive perspectives. The main focus will be on the notions of *trading zones* and *boundary objects* and how they can be seen as the analogues of pheromone traces in a termite colony. Specifically, this paper seeks to establish their cognitive function and hence provide a framework for an understanding of the complementarity of the cognitive and the social approaches to scientific investigation.

Peter Galison coined the term *trading zone* in his *Image and Logic* (1997, Chapter 9) for the space in which different groups of scientific practitioners with different sets of concepts and practices need to coordinate with each other. He supposed that this was essentially a linguistic problem and, analogous to the situation in anthropology where there may be literal trading between peoples of quite different cultures, suggested that a kind of pidgin evolves that facilitates the coordination. Susan Leigh Star and James Griesemer introduced *boundary objects* in their study of The Museum of Vertebrate Zoology at the University of California at Berkeley (Star and Griesemer, 1989; see also Bowker and Star, 1999). The boundary object was key to understanding coordination between different groups in a scientific project. It was a relatively stable material entity interpretable in a variety of ways by the different groups interacting with it. These objects facilitated coordination of complex actions across and between groups by virtue of their relative stability and interpretive richness despite the absence of a common understanding. The suggestion in this paper is that one way in which trading zones can function is through the use of boundary objects and that both play a crucial cognitive role.

The coordinated systems approach (CSA) was originally introduced to understand the central significance of the environment in minimal cognition (Sims, 2022, 2023; Sims and Yilmaz, 2023). The key feature of the proposal is that these processes are coordinated not by some central controller but rather by environmental variables themselves the result of actions of the system or its constituent processes – a characteristic of *stigmergy*. For example, the pheromone traces of termites are the result of previous building actions which coordinate future actions of the colony. Complementary to the coordination of system processes is a second set of processes working over a longer timescale by which the individual system elements become attuned to the appropriate environmental variable(s).

The main proposal in this paper is that the CSA provides a useful framework for understanding the cognitive roles of boundary objects within trading zones and hence knits together a cognitive and social understanding of scientific investigation.

The structure of the paper reflects this project of bringing together cognitive and social understandings. Section 2 takes a cognitive standpoint and outlines the coordinated systems approach. This is necessarily a skeletal outline, and the details can be found elsewhere (Sims, 2022, 2023; Sims and Yilmaz, 2023). Section 3 reviews the social notions of trading zones and boundary objects. Section 4 brings these two perspectives together identifying aspects of the stigmergic coordination account with features of trading zones and boundary objects. Section 5 evaluates the project and considers further implications.

Inevitably this paper offers an incomplete sketch of an approach that will need to be worked out in more detail if it is to account for key scientific concepts such as validation, evidence, explanation and understanding.

2 The Cognitive side

2.1 The Coordinated Systems Approach (CSA)

The starting point for a coordinated systems account of cognition is that a cognitive system is a loose collection of processes running over different timescales. What makes the system cognitive is that these processes are coordinated in such a manner as to produce goal-directed behaviour. That brings us to two tasks for this section. We need to describe what we mean by goal-directed behaviour and then find an appropriate mode of description to capture the essence of coordination. It should be emphasised that this is only a sketch that omits a great deal of technical detail (for a detailed account see Sims, 2022, 2023). The idea here is just to give a flavour of the account.

The activity of science plausibly fits such a framework – goal-directed action is the kind of thing that is produced by whole organisations. In the next section, for example, we shall meet the Museum of Vertebrate Zoology at the University of California which acts to establish and maintain itself as a museum devoted, amongst other things, to the ecology of California. This is an institutional goal which arguably emerges from the whole system and/or is imposed on it. Some individuals in the organisation may well act from exactly this motivation, but it is entirely plausible that some have other motivations, what Michael Bratman calls *alienated participants* (2022, pp. 66–75); it is not necessary for our purposes that all individuals share a group goal or in fact that any do. Individuals might primarily act from motivations that are local, such as, having a job or engaging in their interest in natural history in an amateur capacity and so on. However, the system is goal-directed not only because its behaviour is end-directed but also because the organisation of its constituent processes is set-up in a manner that delivers a goal-directed result possibly through allocation of the appropriate social roles that are created and policed by the relevant social norms (see Ritchie, 2020).

Goal directedness is a feature of the organisation of a complex system (see Christensen, 1996). A goal-directed system delivers goal-directed behaviour. In rough and ready terms, goal directed behaviour amounts to a relative insensitivity of the system to initial conditions – it is a kind of behavioural robustness. A dog and a thrown stick may both move across a field, but they are functionally quite different systems. The path of the stick is determined by the velocity of the throw. The path of the dog, on the other hand, is less sensitive to the initial conditions. It will, all things being equal, chase the stick

independently of its initial position. The behaviour of the dog is goal directed, while that of the stick is not (see also Walsh, 2015, p. 193). Both systems are characterised by their causal structure, but these causal structures are differently organised; in the case of the dog, the causal organisation ensures goal directedness while that of the stick does not.

At this point it is important to emphasise two points to avoid confusion. Firstly, saying that a system is goal directed is making a statement about the causal or functional structure of the system and that it is set up such that its behaviour is (relatively) independent of initial conditions as in the case of the dog and the stick. It is not required that the goal is represented anywhere or that the system is conscious or contains classical intentional states. Secondly, the task relativity of the system should be taken seriously. A system may be goal directed towards some tasks and not others. The dog may well be goal directed (and therefore cognitive) towards fetching the stick but not, I take it, towards solving the Times crossword. If we accept this notion, then a system is cognitive just in case it is set up to produce goal-directed behaviour with respect to a set of tasks. To do this, it will need to possess a certain kind of organisation with respect to these tasks.

Elsewhere it is argued that the system organisation that is characteristic of cognitive systems is best described using functional language applied to processes rather than in terms of mechanisms (Sims, 2022). I shall not rehearse these arguments here except to point out that there are advantages to be had adopting a process ontology for systems rather than an ontology of entities and activities as is customary in the new mechanisms literature. One clear advantage is that a process ontology avoids the problems of making systems too rigid, possessing fixed properties for example (see Kirchhoff, 2012, 2015a, 2015b), or assuming them to be clearly separable from their environments – thus avoiding the machine or container metaphors that are apt to mislead debates in the field of cognitive systems (see Nicholson, 2013, 2014), and making the notion available for the looser collections of processes found perhaps in social systems.

Having said this, there may be situations where a mechanistic description *is* appropriate, for example when the evolution of the structure of a system occurs on a long enough timescale relative to the that of the behaviour of the system for it to be taken to be more or less fixed. The functional stance taken here is, in principle, compatible both with new mechanistic thinking and more dynamic or enactive approaches to cognition (see for example Craver and Kaplan, 2020; Hutto and Myin, 2013, 2017; Kaplan, 2012; Kaplan and Bechtel, 2011). However, there are other elements of this account that depart from these approaches as we shall see.

I should emphasise that the coordinated systems approach described below is scale independent. It can operate at the level of individual organisms; for example, in connection with the kinds of functional

organisation required to string a set of simpler processes together to account for the coherent behavioural arc of the dog chasing the stick. But the beauty of the coordinated systems approach is that it also works at a social level. The same functional description covers the operations of a complex social organisation like a museum or a research institute. The scale independence also allows that the same cognitive system possesses components that are themselves cognitive (see Godfrey-Smith, 2016; Goldstone and Theiner, 2017; Theiner, 2014).

The idea that cognition is to do with coordinated processes is nothing new. Randall Beer and Tim van Gelder emphasised that cognition is an essentially diachronic and situated process (van Gelder and Port, 1995) expressed in terms of a set of coupled differential equations that link the dynamics of one part of the system with the rest of the system (Beer, 1995, 2000; van Gelder, 1998, 1999a, 1999b). This is a non-reductive theory since the behaviour of one part of the system is irreducibly coupled to the behaviour of the rest. It does not make any sense to examine the properties of one part in isolation. One of the reasons for this irreducibility is that system dynamics emerge from the interaction of all the system components.

I think that this is a correct description of the dynamics of a cognitive system but as an explanatory strategy it does not go far enough. It does not address the normative question *why* the system has the dynamics that it does. In the case of an organism there must be an appeal to the system as a self-organising whole – the *telos* of the system is to survive. In the case of scientific investigation, one expects an appeal to the social norms under which the practices of science are constituted and maintained and under which they produce knowledge. In both cases the normative framework translates a distal goal, say survival or successful production of knowledge, via local constraints into a set of local tasks which carry prescriptive weight and can be performed well or badly. In the case of an organism such as the *E.coli* bacterium the local task is to move to a position of greater glucose concentration than the current position. For a scientist in an investigation the local task might be to label a specimen brought in by a field operative and place it in the museum storage room. In each case what we need is that coordination is defined in terms of its relation to the task faced by the system in the light of environmental conditions and constraints, the system's capabilities, and a distal goal.

I shall return to tasks in a moment, but first I want to propose a high-level functional description of coordination that is general enough to cover all the examples we have discussed so far. Coordination is characterised in terms of two functional conditions placed on system processes. The first is the fulfilment of a function that tracks the performance of tasks and triggers other processes in a timely manner when a suitable stage is reached. Consider traffic flow on the approach to a roundabout. From the point of view of a driver of car A waiting for car B on the roundabout to complete its action, the

task of negotiating the roundabout is triggered by the exit of car B and is tracked by how far car A is from the desired exit. The task is complete when A leaves the roundabout. I call this the tracking and triggering condition. Tracking and triggering are functions performed by causal processes – they need not involve fully-fledged representations (although they are not ruled out) – nor do coordination processes need to be necessarily different from processes involved in task performance. In the case of urban traffic, it is the performance of the task of negotiating the roundabout, the behaviour of the system, that also plays a functional role as a tracker and trigger. The second condition is that coordination processes must be sensitive to changes in the task faced by the system. If car B collides with car C on the roundabout, then the task facing the driver of car A changes to ‘call the emergency services’. The new task emerges in real time from the ongoing state of play involving the system elements including the environment. The existence of the collision on the roundabout partially coordinates the performance of the new task in conjunction with the appropriate social norms. This second ‘task-sensitivity’ condition is more demanding than the first condition and ensures that the behaviour of the system is adaptive – that the system employs different strategies in response to changing environmental conditions to perform the required tasks.

I promised to return to the nature of tasks. Some readers will no doubt be worried about the role of teleological or normative elements in the theory. This is an issue that faces the theorist of cognition once they have opted for a task-driven approach since tasks possess norms for successful or adequate completion. On the other hand, there is a built-in tendency for science to eschew teleology. This tension has been noted in the literature; the German physiologist Ernst Wilhelm von Brücke inappropriately observed: “Teleology is a lady to a biologist: he cannot live without her but he’s unwilling to be seen with her in public”¹ (Rothschuh, 1953, p. 140). But there is nothing mysterious with teleology from a cybernetic standpoint, it is simply an organisational feature of certain kinds of complex system and emerges in systems that are not only self-organising and self-creating but also actively need to maintain their immediate environment in a state best suited to them (Bickhard, 2009a, 2009b). For example, organisms possess the goal of survival achieved through evolutionary processes. To survive in a precarious environment that threatens to disrupt their self-maintenance they need to perform certain tasks to maintain the right relation to this environment. These tasks have success conditions in the sense that certain kinds of failure spell the end of the self-maintaining structure. For example, organisms will need to find nutrients and avoid predators. In the long term they will need to reproduce. These goals do

¹ “Die Teleologie ist eine Dame, ohne die kein Biologe leben kann. Er scheut sich jedoch, sich mit ihr in der Öffentlichkeit zu zeigen”

not need to be represented anywhere, indeed they might not even be goals held by any given element of the system. Instead, they are the result of the self-maintaining processes at work in the system.

In social systems such distal goals may be the result of various kinds of social interactions and the establishment of social practices and norms. For example, from the point of view of the social system in which it is embedded, the goal of the roundabout might be to enable smooth traffic flow across a junction. This is not a goal held by any driver in the system who is likely to want to get to her desired exit in the shortest time. Rather the goal emerges from the totality of social interactions over a long time-interval and the establishment of norms that police traffic flow. These interactions may even be competitive and not entirely cooperative. The system consists in the roundabout as a physical structure together with a set of social norms for driving on public roads. Its beauty derives from the transformation it imposes on a situation of self-interested drivers wanting to get to their destinations in the least time to the vision of the city planners of smooth traffic flow across a complicated road intersection. This is an everyday illustration of the deeper idea that a characteristic of self-organising systems is that distal goals may emerge from the interactions of system elements including the environment.

I take tasks to be local versions of goals. Goals are far away in time and space and do not specify the exact means of their fulfilment. Tasks on the other hand are local and highly constrained, and a succession of them eventually leads to goal fulfilment. Tasks are specified by constraints in the environment (including the social environment) and constraints within the system itself that canalise the distal goal of the system. For an individual driver, the topology of the road system and the norms associated with driving constrain the actions available in approaching the roundabout. These constraints channel the distal goal of reaching the destination into a set of immediate tasks, such as: approach the roundabout, give way to traffic and then, when the roundabout is clear, proceed to the required exit following the course suggested in the Highway Code for the position of the exit. At the same time these norms canalise the whole system towards the task of making possible the local traffic manoeuvres that go towards the distal goal of efficient and fair traffic flow. The point here is that constraints both physical and social are profoundly enabling for the individual driver reducing the number of actions available and setting an upper bound to the complexity of task structure (see Anderson, 2015; Raja and Anderson, 2020 on different kinds of enabling constraints). What is interesting about these examples is that future actions are coordinated by the physical result of previous actions, that is, the position of a car relative to the road and other cars - coordination is *stigmergic*. It is to this that we now turn.

2.2 Stigmergic coordination and its benefits

In the example of traffic flow, the results of the actions of individual drivers, namely their position in the road system, coordinate their future actions. Locally, the actions of the drivers on the roundabout are mutually coordinative in the sense that they coordinate the speed and course of those drivers approaching it. The same kind of organisation can be found in human house building. The result of the previous stage of building coordinates the building of the next stage. The experienced builder knows when the foundations are finished, and work can start on the walls. The state of the building both tracks and triggers the actions of its builders.

It was a building project of a rather different kind that caught the attention of the entomologist Pierre Grassé (1959). How could it be that termites, with their miniscule cognitive capabilities, could build such elaborate and intricately functional structures such as a termite mound, which performs many functions for the whole colony such as a lung for breathing, temperature control, humidity control and a signalling system for structural problems? The solution is that the mound itself plays a key role in the coordination of the termites. In the construction process termites lay down pellets of sand containing a pheromone to which they are in turn sensitive. Encountering concentrations of the pheromone elicits pellet-dropping behaviour from pellet-carrying termites leading to a positive feedback loop that constitutes coherent and coordinated building behaviour. The physiologist J. Scott Turner in his book *The Extended Organism* (2000) describes other processes by which the building and maintenance of the mound is achieved in changing environmental conditions (for details see Turner, 2004, 2009, 2011, 2013, 2016; Turner and Soar, 2008). In each of these processes the results of the action of the termites, namely the mound, and the relevant environmental variable, such as the strength of the pheromone, not only trigger the tasks for other termites, but track the progress of the whole structure. Moreover, if part of the structure breaks, the pattern of the pheromone together with other indicators such as the frequency distribution of resulting air currents causes the termites to switch task to repairing the damaged area. These variables then satisfy the coordination conditions with respect to the tasks involved in building and maintaining the mound. In other words, variables such as the pheromone distribution and frequency of air currents, themselves a result of the building process, coordinate the building process. Grassé coined the term *stigmergy* for this kind of coordination.

Let us look more carefully at those features of the environment playing a coordinating role in the building of the mound – called *trace* variables in the stigmergy literature (Heylighen, 2016; Theraulaz and Bonabeau, 1999). In the case of the termites the pheromone trail laid down in depositing sand pellets is a trace – it tracks and triggers the stages of the building tasks. In the case of a human building project, some aspect of the building-in-progress will act as a trace with respect to a set of tasks. In

preparing the foundations for a house, only when the poured concrete reaches a certain level should the pouring action stop. Only when the concrete is dry can other tasks such as building the floors and walls get started. In the first task the relevant variable is the level of the poured concrete in the second it is the dryness of the concrete. If the concrete develops a crack, a new task is generated, and the system must adapt and initiate actions to repair it. In other words, the interactions of experienced builders with the relevant environmental variable, itself dependent on their previous actions, satisfies both the tracking/triggering condition and the task sensitivity condition.

While dynamical systems explanations present the dynamics of the system as being an explanation of its behaviour, a stigmergic coordination explanation involves a historical component that supplies the required normativity. The explanation must address the question of how the system acquired its sensitivity to a trace variable in the first place. In recent work I describe a second set of processes acting on longer timescales that attune elements of the system to the appropriate environmental variable(s)². In the case of the termite mound, this is the result of biological evolution over a long period of time. In the case of human building projects, trace sensitivity is acquired through processes of training and apprenticeship in which the trainee builder acquires an ‘eye’ for the kind of tasks that need doing and their state of completion. The concrete pourer has learned to monitor the depth of the concrete in order that her actions are tracked and triggered by it. Norms also play a role in regulating how tasks should be performed and what counts as a successful (or at least adequate) performance. The experienced builder is encultured into the correct way of pouring concrete and that it should proceed slowly so that no air is trapped making holes when it dries and so on. But norms work at a higher level too. They provide a framework within which tasks are generated and made intelligible. The task of pouring concrete only makes sense within a set of social practices and norms involving houses with concrete foundations, but these very practices and norms create that task for the concrete pourer. The experienced builder is encultured into these higher-level norms and practices again through a learning process, perhaps as an apprentice. When a set of system elements becomes sensitive to a trace relative to a set of tasks, the system is coordinated with respect to these tasks and, if these arguments are correct, the system is minimally cognitive. Social norms and their enculturation play a role not only in cultivating sensitivity to the appropriate trace variables and regulating the performance of tasks but also in the very generation of these tasks in the first place.

² In two articles under review I discuss attunement, roles and show how role distribution may also be coordinated through environmental variables.

In short then, a stigmergic system satisfies the coordination conditions with respect to a suitable set of tasks and hence is cognitive according to our theory. Moreover, stigmergy plays a role in explaining how goal-directed behaviour and significant self-organisation may emerge from systems possessing many simple elements that are made suitably responsive to the appropriate trace variable by long term attunement processes. In the human building example, the functioning of an environmental variable as a trace is possible because of the role of social processes such as enculturation of social norms and practices. These social processes produce norms that define the task structure of the situation, define the conditions for satisfactory performance of tasks, and define the ways in which an environmental variable can function as a stigmergic trace.

The astute reader may suppose this is all well and good but what are the benefits of this account over and above the dynamical systems approach. I have already argued that a coordinated systems account provides the required teleo-normative element and that it can apply recursively to the same system on different scales. Another benefit is that the account emphasises the distributed and non-centred nature of the system. Each driver approaching the roundabout possesses only a local view of the whole traffic system. However, the interactions with the road system and other drivers according to the norms of the Highway Code ensure that the whole system functions smoothly without any element of the system possessing an overview of the whole or without any single element in the system holding or even being aware of the overall system goal.

These are advantages, not only for the investigator of cognitive systems who wants her conceptual tools to apply to a wide range of systems, but also to the systems themselves in which information bottlenecks produced by central controllers are problematic for operational reasons. Hence, we can construct arguments for a coordinated systems view of scientific investigation not only from the standpoint of philosophy of science – that this is a good way of understanding science from the point of view of the philosopher of science – but also that it is a good way of doing or organising science – from the point of view of scientists themselves.

Avoiding information bottlenecks is a particular case of a more general class of optimisation benefits of stigmergic systems. The ideal house-building scenario where the trace variables are doing their thing delivers a nice just-in-time efficiency. Exactly when the concrete is dry, the wall builders can get to work. Exactly at the appropriate time the electrician gets busy installing the ring main and so on. Stigmergic systems are capable of solving optimisation problems that evade analytic methods (see for example Dorigo *et al.*, 1996; Parunak, 2006; Resnick, M., 1997) such as the travelling salesman

problem. Stigmergic ‘ant algorithms’ are a feature of computational methods in the optimisation literature as well as calving new research fields such as multi-agent system research and swarm robotics.

A second group of advantages concerns complexity. As systems grow and can perform more tasks, they typically become more complicated in the sense that there are more states available to them. In standard computational accounts this leads to an exponential increase in complexity and subsequent problems of limited bandwidth especially when a central controller is involved. In a stigmergic system, the local interaction of elements ensures that an increase in the task space is always accompanied by an increase in the number of system elements; the local complexity of the system remains roughly constant (see Parunak, 2006). A bigger termite mound involves more complex coordination but also more termites to do the ‘processing’. A bigger building site usually means proportionally more builders each coordinated by local trace variables. A larger roundabout, such as the Magic Roundabout in Swindon UK consisting of six smaller roundabouts, accommodates more drivers but involves more agents to do the local ‘processing’ and occupies a larger area and therefore keeps a lid on local complexity³.

A major intrinsic benefit of stigmergy follows from the prominent role of the environment - the fact that such systems use the world as its own model (Brooks, 1991, p. 140). Instead of spending energy and organisational resources on generating and maintaining representations for task book-keeping, builders only need attend to the state of the building-in-progress, just as termites do. To the experienced builder, the state of the house, at least in its larger aspects, speaks for itself. This is even clearer in the case of traffic flow. Individuals do not need to create internal representations of the whole provided they attend to stigmergic traces such as following the car in front and operating according to the encultured norms and habits of good driving.

Finally, stigmergic coordination knits together local and global aspects and thereby solves a major problem facing goal-directed systems, namely, how to bridge the gap between the local proximate environment or Umwelt and a distal goal (Keijzer, 2001; Tolman and Brunswik, 1935; von Uexkull and O’Neil, 2010). Both termites and humans exploit environmental regularities in their coordination of local action in the service of a distal goal. The pheromone trail is both a local trigger to action for the individual termite and a global structure that produces a coherent mound. The road layout ensures that the local actions of individual drivers will result in a large-scale goal-oriented displacement. A tree is an environmental regularity or enabling constraint that allows termites to reach a distal goal, food,

³ I use ‘processing’ as a metaphor – this account is not committed to a standard information processing reading of complexity.

through a sequence of local tasks. Building practices and norms ensure that local actions by builders result in a house.

Some readers may have been worrying about the house building example throughout this section. Surely, builders use maps and plans? This was supposed to be a non-representationalist story so what is the role of representations which in some systems play an important coordinative role? The answer should now be clear. Representations such as maps and plans just are a kind of environmental variable to which the system or parts of it are suitably sensitised. There is nowhere here the assumption that stigmergic traces possess representational content (and it is not at all clear whether iconic sign systems like plans possess propositional content anyway). They may do of course – but that is not required by the theory. Together the relation of the partially-built house to the representation on paper acts as a stigmergic trace with respect to the multiple building tasks. Moreover, roofers and electricians may read different traces from the map/building relation, with respect to their quite different tasks. This suggests a link with the social side and boundary objects in trading zones which is where we now turn.

3 The social side

3.1 Trading Zones

The notion of ‘trading zone’ was introduced by Peter Galison in his book *Image and Logic* (Galison, 1997, Chapter 9). The concept was introduced as part of Galison’s bigger project in the book of seeking a new approach to what he called the unification of science. In pushing back against the Kuhnian notion of incommensurability, Galison proposed that theory change was not the earth-shattering event that Kuhn supposed it to be but rather more piecemeal. To support this view, he supposed that the fabric of scientific practice was multilayered – a series of strata consisting of practices around theory, instrumentation, and experiment. Discontinuities in practice in one of these layers did not necessarily line up with discontinuities in the others. For example, Galison maintained that it is unlikely that radical changes in experimentation necessarily gave rise to immediate changes in instrumentation or theory. Indeed, he presented evidence for the autonomy of each layer and held that each of the layers were not reducible to the others. This autonomy, says Galison, is the source of the strength and stability of science (1997, p. 781).

Galison’s solution to the incommensurability problem relies on the operational autonomy of these different groups; the instrument makers get on with their own thing, and so for that matter do the theoreticians and the experimentalists. This raises a new challenge. How can life in one area, say experiment, continue unchanged given that, say, the theoretical foundations for these experiments

change? Are the theoretical entities identified by the experiment now radically different? How can these groups manage their coordination across such changes? Indeed, given that cooperation is the nature of the scientific project, how can these autonomous groups coordinate at all?

Galison saw this problem to be one of language. Each group involves its own set of practices, criteria for success in those practices, concepts, values and so on. When they need to work with other groups a kind of pidgin is developed. This is a set of linguistic expressions allowing coordination between the groups but possessing different meanings across them. The only condition is that the expressions used suffice for practical collaborative purposes. In Galison's example of quantum chromodynamics, terms like "quarks", "jets" and "hadronization", meant different things to different theoreticians and to experimenters (1997, p. 814). Similarly, "mass" meant something quite different depending upon whether one was talking to Lorentz or to Einstein (1997, p. 811). Nonetheless, these theoreticians could talk to each other and to the experimentalists and the instrument makers in these terms, the different meanings aligned in the different domains sufficiently for groups working together to coordinate their actions. That way, major changes in meaning in one area had little effect in another. Coordination could go on very much as before. Galison appropriates a term from anthropology to describe this. When trading occurred, for example, between different peoples of quite distinct cultures, a special interlanguage evolved with enough specificity to coordinate the trade. The location of this interaction was called a *trading zone*.

The geographical nature of the trading zone metaphor should be taken seriously. For Galison, they are literally spaces that can be located, for example, in buildings. Indeed, some buildings were designed specifically with such zones in mind. Striking examples are to be had here: the National Accelerator Laboratory near Chicago (NAL - later to become the Fermilab) doing high energy physics, or the MIT Radiation Laboratory (Rad Lab) located in Cambridge devoted to the development of radar during the Second World War. In both these cases, there are physical spaces created to bring together theoreticians, experimenters, engineers, and instrument specialists. These are created through an explicit administrative policy to facilitate such exchanges and to provide coordination (Galison, 1997, p. 819).

What is being coordinated? According to Galison, the interactions that take place in the trading zone coordinate action and belief (1997, p. 827). The present paper does not follow Galison all the way here. Granted, actions are being coordinated. Engineers at Fermilab are tasked with producing physical structures in which the experiments determined by the experimental community can take place. Likewise, experiments are devised based on current debates in theory, so there is a need for the various groups: theoreticians and experiments, to coordinate actions amongst and between themselves. But does

the coordination of action specifically require coordination of belief? This paper takes the perhaps unfashionable view that scientific knowledge, in the main, is ‘active knowledge’ in that it is closely related to goal-directed action in the world, and that it is founded more on know-how than propositional knowledge. I shall not defend this view here but refer the reader to the convincing arguments of Hasok Chang in his recent book (2022 especially chapter 1). The proposal is that the primary kind of coordination at stake here then is coordination of action. This will be discussed in the next section.

A second point of difference with Galison concerns the medium of coordination. In keeping with his emphasis on belief he emphasised that it is the evolution of a trading *language* - a pidgin - that lies at the heart of the coordinative processes. This pidgin should be understood in very broad terms, however. It is helpful to consider an example. In an extended case study, Galison examined the work done by radio engineers and theoretical physicists at Rad Lab in the development of a radar system suitable for military use in the 1940’s. The problems centre on the structures that channel microwave radiation along complex waveguides - long hollow metal boxes involving discontinuities such as protrusions and dividers. To treat the discontinuous geometries of these structures according to the methods of pre-war theoretical physics “would have involved intractable calculations of the currents around these protrusions and of the immensely complicated fields generated around them” (1997, p. 821). This is where the evolution of a trading zone comes in. An old technique among electrical engineers when faced with a complex electromechanical device such as a loudspeaker with its multiple feedback loops from the electrical domain to the mechanical and back again was to present a practical representation of this complexity as a simple ‘equivalent circuit’. This is a representation of the loudspeaker not in terms of its real electromechanical components but rather *as if* it were a circuit of purely electrical components. “They (symbolically) put the complicated physics of the loudspeaker’s electromechanically generated sound into a ‘black box’ and replaced it in their calculations with ‘equivalent’ electrical components” (1997, p. 821). One of the leading physicists, Julian Schwinger, wondered whether such a technique could be developed for the complex microwave circuits. He then went on to develop equivalent simple circuits involving notional inductors, resistors, and capacitors for the components of the waveguide system. That way, the radio engineers could perform calculations that modelled the basic characteristics of the microwave circuits while the physicists could manipulate the very same circuits without the intractable quantities. Neither side needed to know the specialist interpretations of the other, yet their interactions were coordinated.

Galison casted this situation in linguistic terms as a creation of a pidgin⁴. But is it not better expressed as the construction of some kind of virtual object? Are the equivalent circuits not just models of the microwave guides that capture their primary features and bracket the physical complexities in higher order correction terms? Would it not be better to think of these equivalent circuits in the same way as say the Hodgkin-Huxley equivalent circuit model of the action potential across a neuronal membrane, also presented in the form a circuit diagram using capacitors, variable resistors, and sources of electromotive force? (Craver, 2007, p. 115). Surely these equivalent circuit models can be manipulated both by the engineers and the physicists and it is this *manipulation* that coordinates their actions. This is the key point concerning trading zones - that rather than just being a zone in which a pidgin is spoken, they are zones in which objects are manipulated from both sides with different significance to each. They are *boundary objects*.

3.2 Boundary objects

The term boundary object enters the vocabulary of philosophy and sociology of science from a highly influential paper by Susan Leigh Star and James Griesemer (1989). Star and Griesemer were interested in the puzzle of how scientific cooperation could take place across quite disparate groups without much consensus or with the bare minimum of shared understanding. Actors trying to solve scientific problems come from different social worlds with their own sets of tasks, *modus operandi*, and, to an extent, criteria for success. Nonetheless they can be eminently successful in their cooperation despite these differences through being able to coordinate their *activities*. The key to this coordination is the boundary object.

The focus of their work was a detailed case study of the Museum of Vertebrate Zoology at the University of California in Berkeley. They described the participation of different groups from distinct social worlds in the scientific life of the museum: professional scientists, amateur naturalists, patrons, hired hands, and administrators. They observed that a university administrator in charge of grants and contracts was answerable to a different set of audiences and is engaged in a different set of tasks than an amateur field naturalist collecting specimens for the museum (1989, p. 388). The objects that they work with, grants, contracts, specimens, meant different things to these different groups of actors yet they somehow reconciled these meanings to coordinate their actions. Latour (1987) saw this process in terms of a flow of objects and concepts through the network of participating groups and their social

⁴ I want to thank an anonymous reviewer for pointing out that a pidgin may also be thought of as coordinating actions without reference to its traditional linguistic role.

worlds. Star and Griesemer adopted the same metaphor in their talk of the translation of objects and concepts between the different worlds. Both sets of authors, Latour on the one hand and Star and Griesemer on the other, saw the whole as being an ecology in which coordinative interactions are happening locally, leading to global behaviour of the right kind. This analysis is holistic in that the main unit of analysis is the whole enterprise itself, and no one viewpoint has epistemic primacy. In the case of the museum the “viewpoint of the amateurs is not inherently better or worse than that of the professionals” (Star and Griesemer, 1989, p. 389). The museum therefore provides a life-sciences parallel to Galison’s groups in the world of physics – the museum is literally one or more trading zones.

Star and Griesemer’s problem parallels that of Galison: how can the different actors in their different forms of life achieve sufficient coherence of understanding to enable cooperation between them? Boundary objects provide the key to the puzzle. I shall quote their definition in full since it will be used in our subsequent analysis and link to cognition.

[A boundary object] is an analytic concept of those scientific objects which both inhabit several intersecting social worlds (...) and satisfy the informational requirements of each of them. Boundary objects are objects which are both plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. They are weakly structured in common use, and become strongly structured in individual use. These objects may be abstract or concrete. They have different meanings in different social worlds, but their structure is common enough to more than one world to make them recognisable, a means of translation. The creation and management of boundary objects is a key process in developing and maintaining coherence across intersecting social worlds. (1989, p. 393).

Star and Griesemer listed the kinds of boundary object of interest to the various actors in the Museum network to be the following:

- (1) **Repositories** - ordered ‘piles’ of objects which are indexed in a standard fashion such as objects in the museum. These possess a modularity that allows people from different worlds to use them for their own purposes without others sharing those purposes.
- (2) **Ideal type** - this is a diagram or other description which does not accurately describe any particular thing but is abstracted from all domains. It adaptable to local domains precisely because it is vague and abstract. It can be used as a means for communicating and cooperating symbolically. In the museum situation an example of an ideal type is ‘species’. It accurately

describes no specimen but allows cooperation between the theoretical and practical worlds of the professional zoologist and the amateur collector.

(3) **Coincident boundaries** - objects with shared boundaries but different contents. An outline stencil of the state of California is such an example. The amateur collector uses the outline filled in with roads and campsites to coordinate her activities while the professional zoologist fills the interior of the boundary with quite different content, for example habitat information and so on.

(4) **Standardised forms** - these are boundary objects devised as methods for common communication between work groups given that they were working at diverse and scattered sites across the state, for example, the standard paper form that the amateur collector must fill in on handing over a specimen. These are objects that can be transported over great distances and convey unchanging information. Star and Griesemer point out that Latour calls these 'immutable mobiles'.

These objects are boundary objects in the sense that they are marginal - not that they bound anything physical or conceptual - but rather that they are accessible from different social worlds and mean different things in those worlds. They have multiple directions of access from different groups with different tasks, different immediate goals, and different ways of doing things. Schwinger's equivalent circuits are likewise boundary objects in this sense. They allow coordination between different groups.

The conceptions of boundary objects in a trading zone suggested here raise three important questions.

- (i) What is the role of boundary objects and the trading zones they inhabit in the coordination of scientific endeavour?
- (ii) Do boundary objects perform a purely social function, or can their role be construed as cognitive - at least in part?
- (iii) Are boundary objects primarily linguistic entities?

This paper proposes that (i) and (ii) can be answered by suggesting a cognitive role for boundary objects and that the answer to (iii) is 'no' at least according to a traditional conception of language. Aspects of boundary objects turn out to precisely the environmental variables identified in section 1 that coordinate cognitive systems. This connection is made in the next section.

4 Bringing the social and cognitive together

4.1 Boundary objects and multiple traces

Having introduced the notion of the stigmergic coordination of distributed cognitive systems it is time to bring all these elements together. How do boundary objects fit the coordinated systems approach to distributed cognition? The reader will doubtless have picked up on the hints in the last section that there are objects involving a stigmergic trace that help to coordinate distributed cognitive systems and that these are boundary objects. What makes them important in science as well as in other areas of goal-directed action is that the same object may possess *multiple* stigmergic traces. This is key to intergroup coordination. Before we explore this idea in connection with scientific practice, it may be useful to try these ideas on an example closer to home due to Marc Slors (2019; see also Sims, 2022).

Imagine shopping in a supermarket. One is confronted with a bewildering array of tasks as part of the experience. These tasks are generated, structured and their performance regulated and monitored through a web of interconnected social norms, conventions and practices. A shopping list needs to be consulted, the top item chosen, the correct shelf in the shop located, the item of the right kind identified, picked from the shelf, and placed in the shopping trolley. The item needs to be crossed off the list and so on. I argue that these activities are stigmergically coordinated through interactions with the shopping list. When every item is crossed off the list, the trolley is taken to the checkout and a new activity begins jointly with the cashier governed by a specific set of social norms. Slors argues that the cashier and the customer are not tightly functionally integrated because they possess quite different local goals and are relatively autonomous in the long term. However, they are bound together in the sense that their individual tasks are co-dependent. The act of choosing items from the shelves only makes sense in terms of the eventually checking-out and the checking-out procedure only makes sense after the customer has decided upon the goods she wishes to buy. Slors calls this situation *symbiotic cognition* in recognition of this mutual dependence (Slors, 2019). But I prefer to label it stigmergic cognition.

Let us look more closely at what happens when a customer, say Barbara who is on meal duty for her family that evening, interacts with the supermarket cashier Alan. Barbara places an item on the belt. Alan advances the belt and scans the item while Barbara places the next item and so on. Eventually Barbara switches task to placing the items in her bag. These actions are coordinated by their joint trace: the position of items on the belt relative to Alan. Yet they have quite different distal goals. Barbara's goal is to provide supper for her family, Alan's goal is to perform his supermarket duties well enough to continue to be employed by the supermarket. How can it be that the trace of their joint actions, the position of items on the belt, coordinates to two quite distinct goals?

The key to the puzzle is that the item on the belt is a boundary object. It is an object or environmental structure that gives rise to two different trace variables coordinating to different distal goals in two different ‘worlds’. Yet it is the same object thus its position physically coordinates Alan and Barbara’s coordinated joint action. In the case of Barbara, the salient trace variable consists of specific objects on the conveyor belt that bear the right relation to the supper project (via the intermediate coordinator of the shopping list). If the supper project changed in some substantial manner, the identity of these objects would also change. Alan, on the other hand is not sensitive to the identity of the objects but only interested in the location of a barcode that can be read by his scanner. What is of concern to him is that the belt keeps pace with the speed at which he is scanning the items. That fulfils his side of the joint task. Relevant to his goal is the throughput of supermarket goods and orderly management of the checkout. Despite being coordinated by different trace variable associated with the boundary object, the interaction produces a short-lived cognitive system delivering a goal-directed behavioural outcome. The small piece of spacetime in which the joint action occurs is a trading zone (in this case literally).

Marc Slors is right that social conventions play an important role in joint action coordination in these trading zones (Slors, 2021). The reason is that they provide a way in which the boundary object can possess multiple traces coordinating the different tasks on each side of the boundary. The participants in the joint task are equipped, by familiarity with the convention, to ‘read’ the relevant trace variable from the boundary object relative to their task and therefore coordinate their individual action towards distal and possibly quite distinct goals while at the same time coordinating their joint action. The main point here is that it is not true that the supermarket is one giant distributed cognitive system as some would have us believe (see Gallagher and Crisafi, 2009), but rather it should be seen as a kind of cognitive ecology that contains smaller cognitive systems centred on boundary objects involved in their stigmergic coordination. Can the same be said for science?

4.2 Boundary objects and stigmergic coordination

Let us see if the same kind of analysis helps elucidate the examples of section 2. What is needed is boundary objects that possess multiple ‘aspects’ that correspond to different trace variables for different groups of collaborators. Because different groups ‘read’ different tasks from the different aspects of the boundary object, yet the boundary object retains some stability and identity through the process, coordination between the groups is achieved. But this is precisely what Star and Griesemer specify in their definition of a boundary object (from section 3.2):

- (1) “They are plastic enough to adapt to local needs of the parties employing them” – they are flexible enough to register and track task completion.
- (2) “They are robust enough to maintain a common identity across sites” – so they can be recognised and ‘read’ by different groups.
- (3) “They are weakly structured in common use” – interpreted as possessing enough richness to exhibit multiple aspects that can coordinate multiple collaborative groups.
- (4) “They become strongly structured in individual use and have different meanings in different social worlds” – individuals or collaborative groups read a particular aspect in a structured way. Moreover, the use of the word ‘become’ hints at the attunement processes by which groups or individuals become sensitive to a particular trace aspect of a boundary object.

These properties seem to describe Schwinger’s representations of equivalent circuits of waveguides in the Rad Lab. The radio engineers respond to the equivalent circuit as describing an electronic circuit while the physicists see it as a description of a complex wavefunction. Both sides of the trade manipulate these representations as part of coordinating their actions by, for example, assembling basic waveguides like Lego to form more complex systems analogously aggregating the equivalent circuits.

In the case of the museum, recall there were 4 different kinds of boundary objects suggested by Star and Griesemer. There were listed as repositories, ideal type, coincident boundaries, standardised forms.

Repositories are ordered piles of objects indexed in standard fashion. Typically, these would be collections of objects in the museum. Thus, there are two kinds of object at stake here: the collections with their variable extension and changing organisation, and the objects themselves. It is plausible that the objects themselves have aspects that can be understood as trace variables in stigmergic coordination. For example, there are interventions required to prepare the object for display such as stabilising its state against decay and disintegration, presenting it in a visually appealing or ‘typical’ form, and connecting it to an index thought of perhaps as a translation of disciplining of a piece of the world into an element of a classification system (Latour, 1999, Chapter 2). A competent participant in each of the interest groups interacting with the object can read the progress of her task and discern the next task from the appropriate trace. For example, a person concerned with curation will be aware when the taxidermy process is complete in order that the object can be admitted to a collection. The collection itself is an object whose content is variable. Those responsible for the overall curation of the collection should be sensitive to emergent trace variables tracking its formation process to be able to judge whether

the collection fulfils its requirements. There may be criteria of completeness or representativeness that must be satisfied here, states of the trace to which the curators of the museum will be sensitive.

Ideal type objects are abstract and usefully vague. The example given by Star and Griesemer is ‘species’. It seems that something like a specific species is manipulated in the practices of the museum. What is striking here though is that it is not so much the object that tracks and triggers goal-directed action but more the relation of the object to other objects. A specimen is identified with an ideal type – its species – that is a relational property. We can imagine that this identification may be the subject of disagreement or discussion and may track and trigger further actions on behalf of the different groups in the museum. For example, in a situation in which there is uncertainty as to whether a specific specimen is indeed an example of a given species, the specimen may be rejected, or the ideal type modified. Although the amateur collector and the professional zoologist alike manipulate ideal types, generally only the professional, as part of a professional community, would have the authority to modify an ideal type⁵.

Coincident boundaries are objects providing an outline within which different groups place different kinds of content. The example here is an outline stencil of the state of California used by the collector to mark useful practical information such as the position of car parks and campsites but also used by the zoologist for marking the extent of habitats and so on, hence there are two trace aspects available for coordination here.

Bowker and Star discuss standardisation as a way in which coordination can be achieved between different groups (Bowker and Star, 1999). Standardised paper forms used in conjunction with specimens track the specimen collection process and doubtless contain some fields that are filled in by the zoologists in the museum on receipt of form and specimen. The forms possess multiple fields which may correspond to multiple trace aspects that track and trigger further manipulations of both form and specimen.

Although more detailed ethnographical research is required here, there is an outline of a cognitive function for boundary objects in coordinating the actions of different groups into a single minimally

⁵ As an afterthought here, keeping in mind the idea of the relation between a specimen and its ideal type, perhaps the insights of Star and Griesemer (and Bowker and Star) should be modified to include the possibility that there are boundary relations, that play the same kind of role as boundary objects (Bowker and Star, 1999; Star and Griesemer, 1989). This is supported by the fact that many examples of stigmergic trace variables in the wild are actually relational (see Sims and Yilmaz, 2023).

cognitive system. It would be interesting to see if fieldwork provided further evidence for the multiple trace account of boundary objects required to perform this coordination function.

Emerging from this admittedly sketchy analysis is the idea that environmental trace variables may be relational rather than simple properties of objects. This should not be a surprise given that, for example, the coordinating traces in the roundabout example are almost entirely relations between cars and infrastructure. The relation of a specimen and an ideal type then is legitimately thought of as a trace variable. Indeed, Bowker and Star acknowledge that boundary structures may be more complex than single objects in their discussion of ‘boundary infrastructure’ (cf Bowker and Star, 1999, p. 9).

5. Implications

Let us briefly recap what we have done so far. The proposal in this paper is that a cognitive system is distributed across different processes that are self-organised or coordinated in such a way as to be responsible for goal-directed behaviour. The account pinpoints a set of functional conditions that characterise coordination that may be realised in different ways by different processes. Because coordination is multiply realisable, it can be applied to cognition at different scales and levels including the social level. Many such systems are stigmergic in the sense that traces of the action of the system in the environment coordinate further actions like the termite nest. If science is such a system, then we would expect to find evidence of such a coordinating trace akin to the pheromone trace of the termite system. The argument here is that is precisely what we do find in the form of boundary objects occurring in trading zones within a scientific form of life.

This framework suggests a way of understanding the interaction of cognitive and social aspects of scientific practice. Social features of scientific practice such as the establishment of boundary objects play crucial cognitive role. Moreover since this proposal is essentially non-representational, it adds to existing distributed cognition accounts such as that of Ronald Giere (Giere, 1988, 2002a, 2002b, 2006, 2013; Giere and Moffatt, 2003; Vaesen, 2011). It may well go some way to answering some difficult questions that arise from such representationalist distributed cognition accounts (Toon, 2014; see also Magnus and McClamrock, 2015; Magnus, 2007).

For example, an objection that Adam Toon levels against a representationalist account of science as distributed cognition is that it may explain how science adopts the kinds of representation that it does but not *why* that particular representation rather than another. To be clear the CSA account does not require scientific representations to have propositional content but rather to possess aspects that are coordinative. Nonetheless, a similar objection may be levelled at the CSA. Could it be that the CSA

addresses the ‘how’ of the use of processes involving boundary objects for coordinative purposes but not the ‘why’?

This is where the historical dimension of the account pays dividends over more immediate accounts such as that of dynamical systems where the system is just given and neither history nor the evolution of normativity play a central role. Stigmergic coordination requires that a system becomes sensitive to an appropriate environmental trace variable that is correlated to the relevant task in for it to play the relevant coordinative role. This can be described as *sense-making*. A trace variable is salient, or makes sense, with respect to a task, or set of tasks – the system gives it value with respect to its coordinative role. But in section 2.2 it was seen that part of the explanation of minimal cognition as the production of goal-directedness includes an account of the evolution of sense-making, a process called *attunement* (Sims, 2023). It was stressed that for termites this is an evolutionary story, but in the case of social systems such as science, attunement happens through the development of social norms through cultural evolution. A boundary object, however idiosyncratic, that functions to coordinate scientific activity successfully will play the requisite role in the system. A particular way of drawing equivalent circuits will do provided it performs the relevant coordinative function which requires that those using it can read off the relevant trace variable. A paper specimen form, however idiosyncratic, will do provided it possesses enough structure to accommodate trace variables that coordinate the actions of the specimen collector and curator respectively. The answer to the ‘why’ question then involves a contingent history but also a necessary functional role. Yes, another document may have functioned just as well; there is some historical contingency – but not just *any* document will do. It must satisfy the coordination conditions to be able to do the job of coordinating the distributed cognitive system. In explanatory terms, cognition places constraints on these historical contingencies.

Another Toon-inspired worry could be whether minimal cognition in the sense used here is a thick enough notion to do the work of understanding scientific practice. Contentful representations were shown not to be necessary in explaining goal-directed action but doesn’t science output representations that have content and are truth-evaluable? This is a big question, and a proper detailed answer lies beyond the scope of the present paper. Taking cognition to be what is responsible for goal-directed action accounts for what is described as minimal cognition in the literature (see Lyon, 2020; van Duijn *et al.*, 2006), and I have suggested elsewhere that minimal cognition is at least continuous with more fully-fledged notions (Sims, 2023). The CSA is not alone in facing this ‘scaling up’ problem. The problem is common to all enactive theories of cognition and some progress has been made in addressing it (see Barrett, 2019, pp. 108–110; Hutto and Satne, 2015, 2017; Satne, 2016). Entirely within a CSA framework we note that folk psychological practices such as the attribution of beliefs and desires have

important coordinative functions in everyday life *whether or not* they play an explanatory role in cognitive science.

I anticipate that some readers might still object to the characterisation of such systems as cognitive, in the sense that Giere or Chang understands them, and want to hold on to the idea of science as a producer of true propositions. Putting this worry crudely: it seems at least awkward that the analysis in this paper applies perfectly well to pseudoscientific practices such as astrology or homeopathy that do not generate knowledge at all. I think this is a serious objection and one that should be met. While it may lie outside the scope of this paper, I have two brief comments. The first is that representationalist distributed cognition accounts cannot deal with this problem easily either. Astrology makes computations and outputs representations entirely in keeping with Edwin Hutchins' conception of distributed cognition (1991, 1995a, 1995b, 2011). It was never the intention of Hutchins, Giere or the present author to present an ambitious demarcation criterion for the scientific versus the pseudoscientific. Science may have cognitive features and be other things too, while other areas of human activity, like supermarket shopping, may also involve group cognition, while not outputting scientific knowledge. Secondly, if we are to use a coordinated systems approach to distinguish science from pseudoscience, we will need to show how truth-bearing representations emerge from the kind of system we have been studying, in which actions of individuals are coordinated through boundary objects carrying multiple stigmergic traces. The account here shows how scientific action may be coordinated but there is a further step required to link successful coordination of scientific action to truth or to other key scientific concepts such as understanding and validity as mentioned in the introduction. One speculative strategy would be to try to argue that in the long run truth is a good coordinator. This seems highly unlikely as the evidence for the reinforcement of positivity bias and other heuristics in Behavioural Economics attests (see for example Gilovich *et al.*, 2002; Kahneman, 2011; Kahneman and Tversky, 2000; Thaler and Sundstein, 2008). The cognitive properties of science as understood in this paper are, at least at a first approximation, distinct from its truth-generating and truth-transmitting properties⁶.

Giere also suggests that scientific models and diagrams may play key roles in these distributed systems. Quite so. In this paper I have tried to make plausible the idea that such models and diagrams may be boundary objects coordinating scientific activity. More targeted ethnographic fieldwork is needed to make these conclusions more secure and to establish, in the terms of this paper, how such models might function as stigmergic coordinators of such systems. Latour's comment that a lot of science seems to

⁶ The work of Joseph Henrich suggests that falsehoods make rather good coordinators too! (2016)

take place around symbols scribbled on a two-metre square surface makes a lot of sense in this context (1999). This is an area where there can be useful work carried out in the intersection of the social and cognitive in the context of stigmergically coordinated distributed cognitive systems.

Finally I should mention that these ideas, though perhaps not in this form, are familiar in some of the literature in technology studies which is perhaps another natural habitat for stigmergic systems (see for example Lepola *et al.*, 2020; Rambusch *et al.*, 2004; Sellberg and Susi, 2014; Sismondo, 2012; Susi, 2006; Wickman, 2015)⁷.

6. Conclusion

The paper has sketched cognitive roles for trading zones and boundary objects via a coordinated systems approach to scientific practice. This gives a framework for understanding the interaction between cognitive and social views and perhaps suggests a strategy to investigate whether science is irreducibly social.

This paper has only scratched the surface in terms of the possibilities of this approach and suggests a set of detailed ethnographic studies that closely examine the role of boundary objects in coordinating the processes constituting distributed cognitive systems in science. As one anonymous reviewer of an earlier version of this paper puts it – much important work lies ahead.

Acknowledgments

The author would like to thank two anonymous reviewers whose comments helped make this a much better paper.

References

Barrett, L. (2019) Supercharged apes versus super-sized minds: embracing continuity while accepting difference. In Colombo, M., Irvine, E., and Stapleton, M. (Eds.), *Andy Clark and His Critics*. Oxford: Oxford University Press.

⁷ My thanks to an anonymous reviewer for pointing this out.

Beer, R. D. (1995) A Dynamical-Systems Perspective on Agent Environment Interaction. *Artificial Intelligence* 72(1–2): 173–215.

Beer, R. D. (2000) Dynamical approaches to cognitive science. *Trends in Cognitive Sciences* 4(3): 91–99.

Bickhard, M. H. (2009a) The interactivist model. *Synthese* 166(3): 547–591.

Bickhard, M. H. (2009b) Interactivism: A manifesto. *New Ideas in Psychology* 27(1): 85–95.

Bowker, G. C., and Star, S. L. (1999) *Sorting things out: classification and its consequences*. Cambridge, Mass: MIT Press.

Bratman, M. (2022) *Shared and institutional agency: toward a planning theory of human practical organization*. New York, New York: Oxford University Press.

Brooks, R. (1991) Intelligence Without Representation. *Artificial Intelligence* 47(1–3): 139–159.

Chang, H. (2022) *Realism for Realistic People: A New Pragmatist Philosophy of Science*. Cambridge: Cambridge University Press doi:10.1017/9781108635738.

Christensen, W. (1996) A complex systems theory of teleology. *Biology & Philosophy* 11(3): 301–320.

Craver, C. F. (2007) *Explaining the brain: mechanisms and the mosaic unity of neuroscience*. Oxford: Oxford University Press.

Craver, C. F., and Kaplan, D. M. (2020) Are more details better? On the norms of completeness for mechanistic explanations. *British Journal for the Philosophy of Science* 71(1): 287–319.

Dorigo, M., Maniezzo, V., and Coloni, A. (1996) Ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems Man and Cybernetics Part B-Cybernetics* 26(1): 29–41.

Galison, P. (1997) *Image and logic: a material culture of microphysics*. Chicago, Ill. ; London: University of Chicago Press.

Gallagher, S., and Crisafi, A. (2009) Mental Institutions. *Topoi-an International Review of Philosophy* 28(1): 45–51.

Giere, R. N. (1988) *Explaining science: a cognitive approach*. Chicago, Ill. ; London: The University of Chicago Press.

Giere, R. N. (2002a) Scientific cognition as distributed cognition. In Carruthers, P., Stich, S. P., and Siegal, M. (Eds.), *The Cognitive Basis of Science*. Cambridge: Cambridge University Press.

Giere, R. N. (2002b) Discussion note: Distributed cognition in epistemic cultures. *Philosophy of Science* 69(4): 637–644.

Giere, R. N. (2006) The role of agency in distributed cognitive systems. *Philosophy of Science* 73(5): 710–719.

Giere, R. N. (2013) Distributed cognition without distributed knowing.

Giere, R. N., and Moffatt, B. (2003) Distributed cognition: Where the cognitive and the social merge. *Social Studies of Science* 33(2): 301–310.

Gilovich, T., Griffin, D., and Kahneman, D. eds. (2002) *Heuristics and biases: the psychology of intuitive judgment*. Cambridge: Cambridge University Press.

Godfrey-Smith, P. (2016) Individuality, subjectivity, and minimal cognition. *Biology & Philosophy* 31(6): 775–796.

Goldstone, R. L., and Theiner, G. (2017) The multiple, interacting levels of cognitive systems (MILCS) perspective on group cognition. *Philosophical Psychology* 30(3): 334–368.

Grassé, P.-P. (1959) La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux* 6(1): 41–80.

Henrich, J. P. (2016) *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton: Princeton University Press.

Heylighen, F. (2016) Stigmergy as a universal coordination mechanism I: Definition and components. *Cognitive Systems Research* 38: 4–13.

Hutchins, E. (1991) The social organisation of distributed cognition. In Resnick, L. B., Levine, J. M., and Teasley, S. D. (Eds.), *Perspectives on Socially Shared Cognition*. Washington DC: American Psychological Association.

Hutchins, E. (1995a) *Cognition in the wild*. Cambridge, Mass. ; London: MIT Press.

Hutchins, E. (1995b) How a Cockpit Remembers Its Speeds. *Cognitive Science* 19(3): 265–288.

Hutchins, E. (2011) Enculturating the Supersized Mind. *Philosophical Studies* 152(3): 437–446.

Hutto, D. D., and Myin, E. (2013) *Radicalizing enactivism: basic minds without content*. Cambridge, Mass: MIT Press.

Hutto, D. D., and Myin, E. (2017) *Evolving enactivism: basic minds meet content*. Cambridge, Mass: M.I.T. Press.

Hutto, D. D., and Satne, G. (2015) The Natural Origins of Content. *Philosophia* 43(3): 521–536.

Hutto, D. D., and Satne, G. (2017) Continuity skepticism in doubt. In Durt, C., Fuchs, T., and Tewes, C. (Eds.), *Embodiment, Enaction, and Culture*. Cambridge, Mass: M.I.T. Press.

Kahneman, D. (2011) *Thinking fast and slow*. New York: Farrar, Straus and Giroux.

Kahneman, D., and Tversky, A. eds. (2000) *Choices, values, and frames*. Cambridge: Cambridge University Press.

Kaplan, D. M. (2012) How to demarcate the boundaries of cognition. *Biology & Philosophy* 27(4): 545–570.

Kaplan, D. M., and Bechtel, W. (2011) Dynamical Models: An Alternative or Complement to Mechanistic Explanations? *Topics in Cognitive Science* 3(2): 438–444.

Keijzer, F. (2001) *Representation and behaviour*. Cambridge, Mass.: M.I.T. Press.

Kirchhoff, M. (2012) Extended cognition and fixed properties: steps to a third-wave version of extended cognition. *Phenomenology and the Cognitive Sciences* 11(2): 287–308.

Kirchhoff, M. (2015a) Cognitive assembly: towards a diachronic conception of composition. *Phenomenology and the Cognitive Sciences* 14(1): 33–53.

Kirchhoff, M. (2015b) Extended Cognition & the Causal-Constitutive Fallacy: In Search for a Diachronic and Dynamical Conception of Constitution. *Philosophy and Phenomenological Research* 90(2): 320–360.

Latour, B. (1987) *Science in action: how to follow scientists and engineers through society*. Cambridge, MA: Harvard University Press.

Latour, B. (1999) *Pandora's hope; essays on the reality of science studies*. Cambridge, Mass.: Harvard University Press.

Lepola, A., Karkkainen, H., Jalo, H., and Torro, O. (2020) Collaborative Virtual Reality as an Adaptable Boundary Object in the Design Phase of Facility Life Cycle. In Salgado, A., Bernardino, J., and Filipe, J. (Eds.), *Proceedings of the 12th International Joint Conference on Knowledge Discovery, Knowledge*

Engineering and Knowledge Management (Kmis), Vol 3. Setubal: Scitepress
doi:10.5220/0010020200630075.

Lyon, P. (2020) Of what is ‘minimal cognition’ the half-baked version? *Adaptive Behavior* 28(6): 407–424.

Magnus, P. D. (2007) Distributed cognition and the task of science. *Social Studies of Science* 37(2): 297–310.

Magnus, P. D., and McClamrock, R. (2015) Friends with benefits! Distributed cognition hooks up cognitive and social conceptions of science. *Philosophical Psychology* 28(8): 1114–1127.

Nicholson, D. J. (2013) Organisms \leftrightarrow machines. *Studies in history and philosophy of biological and Biomedical Sciences* 44: 669–678.

Nicholson, D. J. (2014) The machine conception of the organism in development and evolution: A critical analysis. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 48(PB): 162–174.

Parunak, H. V. D. (2006) A survey of environments and mechanisms for human-human stigmergy. In Weyns, D., Parunak, H. V. D., and Michel, F. (Eds.), *Environments for Multi-Agent Systems II* (Vol. 3830). Berlin: Springer-Verlag Berlin.

Rambusch, J., Susi, T., and Ziemke, T. (2004) Artefacts as Mediators of Distributed Social Cognition: A Case Study. In Forbus, K., Gentner, D., and Regier, T. (Eds.), *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*. Mahwah: Lawrence Erlbaum Assoc Publ.

Resnick, M. (1997) *Turtles, termites, and traffic jams: explorations in massively parallel microworlds*. Cambridge, Mass. ; London: MIT.

Ritchie, K. (2020) Minimal cooperation and group roles. In Fiebich, A. (Ed.), *Minimal Cooperation and Shared Agency* (1st ed. 2020.). Cham: Springer Nature Switzerland AG.

Rothschuh, K. E. (1953) *Gesichte der Physiologie*. Berlin: Springer Verlag.

Satne, G. (2016) A two-step theory of the evolution of human thinking: Joint and (various) collective forms of intentionality. *Journal of Social Ontology* 2(1): 105–116.

Sellberg, C., and Susi, T. (2014) Technostress in the office: a distributed cognition perspective on human-technology interaction. *Cognition Technology & Work* 16(2): 187–201.

Sims, R. (2022) *Getting their acts together: a coordinated systems approach to extended cognition*, PhD Thesis. University of Exeter.

Sims, R. (2023) Minimal cognition and stigmergic coordination: an everyday tale of building and bacteria. *Cognitive Systems Research* 79: 156–164.

Sims, R., and Yilmaz, O. (2023) Stigmergic coordination and minimal cognition in plants. *Adaptive Behavior* 31(3): 265–279.

Sismondo, S. (2012) Trading Zones and Interactional Expertise: Creating New Kinds of Collaboration. *Technology and Culture* 53(3): 696–697.

Slors, M. V. P. (2019) Symbiotic cognition as an alternative for socially extended cognition. *Philosophical Psychology*.

Slors, M. V. P. (2021) A cognitive explanation of the perceived normativity of cultural conventions. *Mind & Language*. doi:10.1111/mila.12265.

Star, S. L., and Griesemer, J. (1989) Institutional Ecology, Translations and Boundary Objects - Amateurs and Professionals in Berkeleys-Museum-of-Vertebrate-Zoology, 1907-39. *Social Studies of Science* 19(3): 387–420.

Susi, T. (2006) *The puzzle of social activity: The significance of tools in cognition and cooperation*, PhD. Linköping, Linköping.

Thaler, R., and Sundstein, C. (2008) *Nudge: improving decisions about health, wealth, and happiness*. New Haven, London: Yale University Press.

Theiner, G. (2014) A beginner's guide to group minds. In Sprevak, M. and Kallestrup, J. (Eds.), *New Waves in Philosophy of Mind*. Basingstoke, England: Palgrave Macmillan.

Theraulaz, G., and Bonabeau, E. (1999) A brief history of stigmergy. *Artificial Life* 5(2): 97–116.

Tolman, E. C., and Brunswik, E. (1935) The organism and the causal texture of the environment. *Psychological Review* 42: 43–77.

Toon, A. (2014) Friends at last? Distributed cognition and the cognitive/social divide. *Philosophical Psychology* 27(1): 112–125.

Turner, J. S. (2000) *The extended organism: the physiology of animal-built structures*. Cambridge, Mass.: Harvard University Press.

Turner, J. S. (2004) Extended phenotypes and extended organisms. *Biology & Philosophy* 19(3): 327–352.

Turner, J. S. (2009) *The Tinkerer's Accomplice: How Design Emerges from Life Itself*. De Gruyter: Harvard University Press.

Turner, J. S. (2011) Termites as models of swarm cognition. *Swarm Intelligence* 5(1): 19–43.

Turner, J. S. (2013) Superorganisms and superindividuality: The emergence of individuality in a social insect assemblage. In Bouchard, F. and Huneman, P. (Eds.), *From Groups to Individuals: Evolution and Emerging Individuality*. Cambridge, MA: The MIT Press.

Turner, J. S. (2016) Semiotics of a Superorganism. *Biosemiotics* 9(1): 85–102.

Turner, J. S., and Soar, R. C. (2008) Beyond biomimicry: What termites can tell us about realising the living building. In Wallis, I., Bilan, L., Smith, M., and Kazi, A. S. (Eds.), *I3CON BSRIA*. Presented at the First international conference on industrialised intelligent construction Loughborough University.

Vaesen, K. (2011) Giere's (In)Appropriation of Distributed Cognition. *Social Epistemology* 25(4): 379–391.

van Duijn, M., Keijzer, F., and Franken, D. (2006) Principles of Minimal Cognition: Casting Cognition as Sensorimotor Coordination. *Adaptive Behavior* 14(2): 157–170.

van Gelder, T. (1998) The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* 21(5): 615–+.

van Gelder, T. (1999a) What might cognition be, if not computation. In Lycan, W. (Ed.), *Mind and Cognition* (2nd ed.). Oxford: Blackwell.

van Gelder, T. (1999b) Dynamical approaches to cognition. In Wilson, R. A. and Keil, F. C. (Eds.), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, Mass.: M.I.T. Press.

van Gelder, T., and Port, R. F. (1995) It's about time: An overview of the dynamical approach to cognition. In Port, R. F. and van Gelder, T. (Eds.), *Mind as Motion: Explorations in the Dynamics of Cognition*. Cambridge, Mass: M.I.T. Press.

von Uexkull, J., and O'Neil, J. D. (2010) *A Foray into the Worlds of Animals and Humans: With A Theory of Meaning*. Jackson: University of Minnesota Press.

Walsh, D. M. (2015) *Organisms, agency, and evolution*. Cambridge: Cambridge University Press.

Wickman, C. (2015) Trading Zones in Technical and Scientific Communication. In *2015 Ieee International Professional Communication Conference (Ipc)*. New York: Ieee.

