Traces of thinking: a stigmergic approach to 4E cognition

Ric Sims
University of Gävle
ricsis@hig.se
sharpcomposer@hotmail.com

Abstract:

This paper outlines an approach to analysing minimal cognition that brings out its social and historical dimensions. It proposes a model, the coordinated systems approach (CSA), which understands cognition as a coordinated coalition of loosely autonomous processes responsible for goal-directedness in a system. On this view, even individual cognition has something of a social flavour to it. The central concept of the paper is *stigmergy*: a process where the material[1] trace of actions of system elements in their environment is a *sign* that coordinates a group of semi-autonomous processes in future actions – this is the social dimension. The historical dimension refers to longer term processes which establish the coordinative power of the sign and endow it with normative force. According to this proposal, a full explanation of cognitive capabilities should reference both dimensions. In the second half of the paper the CSA is let loose on some puzzles in 4E cognition. Can the model deal with old problems such as that of cognitive bloat, or new problems such as the supposed external memory of the slime mould *Physarum polycephalum*? Potentially, the approach could be used to analyse minimal cognitive phenomena over a range of scales from bacteria to human beings.

1. Introduction

The current paper seeks to make sense of two important insights. The first is that there is an historical dimension to cognition. This is not a new insight – it forms the central hypothesis of work by teleosemanticists such as Ruth Millikan, Fred Dretske, Karen Neander, and David Papineau – however the approach taken here is slightly different. While these authors seek to naturalise representational content, the goal of the current paper is to understand how such historical processes make possible a task-driven analysis of cognition. Cummins and de Jesus (2016) may be a better jumping off point here where they describe the ahistorical bio-cognitive analysis of the single cell as a myth. They theorise a social dimension to cognition even at the cellular level and point out that this is largely undiscussed in the enactive literature, hence their expression 'the loneliness of the enactive cell'. In the same vein, Levin and Dennett (2020) regard a cognitive system as being a collection of agent-like elements and

---

[1] The term material here is used loosely for simplicity. We shall see that the trace can also be relational.

processes possessing some kind of autonomy. This paper attempts to put some detail on these insights and suggests that ‘stigmergy’ is a process that unites them.

The background for the paper is previous work done on what I call a Coordinated Systems Approach (Sims, R., 2022, 2023a, 2023b; Sims, R. and Yilmaz, 2023). This is a framework that could be described as a kind of cybernetic enactivism. Enactive in the sense that cognition is constituted through active engagement of the system with the world, and cybernetic because the system can be thought of in terms of individual elements acting in a semi-autonomous manner, but coordinated to be able to perform tasks that benefit the whole system. The reader is no doubt puzzled by the term ‘semi-autonomous’. As we shall see later, this means that some of the individual elements of the system can be thought of as agents in their own right, but their actions are brought together through the coordinative powers of certain aspects of the environment. Like a single termite in a termite nest such agents may not be viable outside the system itself, but neither are they subject to strict central control. Coordination is achieved through interaction with environmental features - natural signs - that are themselves results of the task-performance of the system of agents. Historicity is necessary to present the system with tasks and goals in the first place and to make the agents in the system appropriately sensitive to signs that are the result of previous task performance. It is through historical processes that tasks acquire their normativity, and that signs acquire their coordinative powers.

The proposal, then, is to understand cognition in terms of two sets of interlocking processes operating on different timescales. Longer-term processes reinforce the operation of short term processes. For example, they are responsible for the attunement of agents in the system to the relevant environmental signs to coordinate the performance of tasks. They are also responsible for the constraints that define tasks in the first place. The attunement process operates not on individual systems, but on collections or populations of them - this is where historicity comes in. Let us call the processes of coordination type-1 processes, and the processes of attunement and task generation type-2 processes; the proposal, then, is for a dual-process model of cognition[2]. Type-1 processes take care of coordination leading to performance of tasks conducing to distal goals, while type-2 processes can broadly be described as ‘evolutionary’ or ‘ratcheted’ and are responsible for the normativity of these tasks and the salience of the environmental variables from the point of view of the system. Explaining cognition involves both types of process. I call the proposal a coordinated systems approach (CSA).

The central concept in this paper is *stigmergy.* This is the process by which the traces of actions of system agents in the environment coordinate further actions. I mentioned termites above because they

---

[2] The terms type-1 and type-2 have nothing to do with the dual process model of Kahneman and Tversky (Kahneman, 2011).

are the classic example for this process; the term stigmergy was coined by the entomologist Pierre-Paul Grassé in 1959 as a solution to the coordination paradox (Grassé, 1959). How could it be that individual termites could coordinate in such a refined way as to produce the intricate structure of the termite mound of which they surely have little or no conception? The answer according to Grassé was that termites build and maintain their nest through coordinated dropping of pellets of sand in response to a pheromone trace. Each termite deposits the pheromone on the pellet that it carries, but it is also sensitive to concentrations of the pheromone in the environment and will tend to drop the pellet when this concentration reaches a certain threshold. Therefore, a positive feedback loop leads to the construction of the mound – the pheromone trace is a result of building actions, so tracks them, but also in turn triggers them. Although individual termites, with their limited capacities, do not have an overview of the whole building project, the stigmergic process involving the pheromone coordinates the building process to such an exquisite degree as to result in the complex and beautiful structures that are termite mounds. In a sense then, it is the termite mound that coordinates the building processes. The concentration of the pheromone is the relevant environmental variable that tracks and triggers these building actions and in stigmergy theory it is called a *trace* variable. It is an elegant way to produce order by coordinating loosely autonomous agents. As Ted Lewis and Leslie Marsh put it, stigmergy is a theoretical framework for understanding the evolution of higher-order low-entropy systems through cooperation of lower-order high-entropy systems (Lewis and Marsh, 2016) – it is a mechanism for self-organisation of a complex system. In the case of the termite colony, biological evolution, a type-2 process, has equipped the individual termites with both the machinery for producing the marker pheromone but also the ability to detect it and respond appropriately to it. But I argue that what makes the pheromone significant in the explanation is that its concentration is correlated in a one-one fashion with the task of building the termite mound. It is the existence of long-term type-2 reinforcement processes such as biological evolution that can exploit this correlation and make it salient to the individual termites and therefore it is these processes that generate tasks for them.

I want to suggest that the CSA is a good framework for understanding a broad range of cognitive phenomena and that many, if not all, cognitive systems involve stigmergies. This includes human beings, or groups of human beings, as well as higher animals, but also bacteria, plants and protists. It has impeccable 4E credentials because of the key role played by the environment but given that stigmergy can also explain brain processes, has potential for internalist cognitive explanations[3]. The enactive component of the CSA means that cognition is constituted by processes of interaction between system elements and the environment without the mediation of anything like contentful representations. The cybernetic character is due to the ubiquitous sociality and multilevel agency in the theory. Hence

---

[3] My thanks to an anonymous reviewer for pointing this out. One way to cash this out at a neural level, might be to think of the reinforcement of synpatic connections in Hebbian learning as being a stigmergic process.

the claim that the CSA is a variety of cybernetic enactivism.

The approach brings together ideas across a wide range of fields. For example, there are those who are interested in the application of stigmergy theory to cognitive systems (see Consiglio, 2020; Lewis, 2013; Marsh and Onof, 2008; Ricci *et al.*, 2007; Sims, M. and Kiverstein, 2022; Susi, 2016; Susi and Ziemke, 2001; Turner, 2011). Sometimes these approaches have involved swarm intelligence or swarm robotics (Beckers, R. *et al.*, 1990; Beckers, Ralph *et al.*, 1994; Sharkey, 2007; Singh *et al.*, 2012; Spezzano, 2019). Moreover, there is a broader group exploring the distribution of cognition over a collection of agents, whether within a representationalist context or not (Cummins and De Jesus, 2016; Giere, 2002a, 2002b, 2011, 2013; Giere and Moffatt, 2003; Goldstone *et al.*, 2006, 2008; Goldstone and Roberts, 2006; Goldstone and Theiner, 2017; Hollan *et al.*, 2000, 2000; Hutchins, 1991, 1995a, 1995b, 2011; Kirsh, 1999, 2006; Levin and Dennett, 2020; Roberts and Goldstone, 2011; Theiner *et al.*, 2010; Vaesen, 2011). Furthermore if we consider the stigmergic trace as being a Peircean sign - see section 4 - then there is a huge literature to draw on from Peirce onwards (Bacigalupi, 2022; Barbieri, 2009; Bickhard, 2017; De Jesus, 2016a, 2016b, 2018; Deacon, 1997, 2021; Favareau, 2021; Heras-Escribano and De Jesus, 2018; Hoffmeyer, 2007, 2020; Menary, 2007, 2018; Peirce *et al.*, 1960; Turner, 2016). Furthermore, this is a diachronic process so owes a lot to process ontologists such as Dupré (2012, 2020; Dupré and Nicholson, 2018) and those who emphasise the diachronic in cognition such as Michael Kirchhoff (2012, 2015a, 2015b; Kirchhoff and Kiverstein, 2020, 2023).

As a side note, I want to resist applying the label ‘distributed cognition’ to the CSA, despite its obviously ‘distributed’ character, because this term traditionally connotes positions with representationalist or computationalist tendencies. Sure, the paper will discuss environmental traces as ‘signs’ in a Peircean sense, but I do not want to suggest that we are committed to fully-fledged symbolic representations with propositional content. In this sense, the paper moves in a different direction to that of Kirsh (1999) and Hutchins (1995a). I also want to add that I shall refer to these semi-autonomous system components as elements or agents. For the purposes of this paper I shall not dwell too long on what makes them agentive, save to point out in section 4 that they are individual consumers of signs. The nature of agenthood of these elements is not a central concern of the current paper and I shall leave this question largely unanswered.

The paper is structured as follows. Section 2 describes the process of stigmergy and how this mechanism plays a key role in producing goal-directed action in systems: type-1 processes. Taking cognitive processes to be those that generate goal-directed action implies that the presence of stigmergy signals the presence of cognition, though we shall leave to one side the question whether this works in reverse, i.e. whether stigmergy is a necessary condition for cognition. Section 3 discusses the crucial role of type-2 reinforcement processes. These processes not only ensure that the agents in the system are attuned to the relevant environmental variables, but they generate the normativity of type-1 processes through a process of *teleogenesis*. Stigmergic traces and their associated tasks come into being together. These reinforcement processes may be biological or social in nature thus creating the possibility that group cognition is appropriate for stigmergic analysis. This section also briefly discusses the distance between the CSA and classical teleosemantic arguments. Section 4 links the trace variable to Peirce's theory of signs. Although this move is not required for the central arguments of CSA, it pays off handsomely when thinking about agency. An agent is simply the consumer of signs. This paves the way for a theory of multiple simultaneous levels of agency. Finally, section 5 applies this theoretical structure to some old and new puzzles in 4E cognition to give a sense of how it can be used.

2 Coordination (type-1) processes

The first kind of process in the coordinated systems account are coordination processes. Coordination is the key to any kind of distributed or multi-agent approach to cognition. If the whole system is to work together then the main concern is how to align the contribution of the individual elements or agents temporally to produce a coherent outcome. Indeed coordination has been the central element of a variety of approaches (for example De Jaegher and Di Paolo, 2007; Kirsh, 1999; van Duijn *et al.*, 2006). David Kirsh writes:

> Our effective environment is a shifting coalition of resources and constraints, some physical, some social, some cultural, some computational (involving internal and external resources). When this shifting coalition of resources is appropriately coordinated, the tasks we set out to achieve are accomplished" (Kirsh, 1999, p. 2).

I depart from Kirsh with respect to the role of representations and computations but agree with the thrust of his statement regarding the coordination of a coalition of elements, or agential processes as I call them, and environmental features, in pursuit of a task. Andy Clark also reminds us that even the most representational environmental structure, language, may derive its significance from its coordinative power: "(…) language works its magic not (or not solely) by means of translation into appropriate expressions of neuralese of the language of thought but also by something more like coordination dynamics" (2011, p. 53). In this section I shall argue that coordination is achieved without content-carrying representations at all, but rather through deceptively simple agent-environmental coupling involved in stigmergy.

In earlier work, I have described in some detail how the theory of stigmergy can be used to understand the coordination of many semi-autonomous processes to goal-directed action (Sims, R., 2022, 2023a). Recall, in the case of the termite nest each termite has limited capacities, but when their actions are coordinated the whole system is able to self-organise to produce the remarkable termite mounds that can be found in the wild. In his wonderful book *The Extended Organism* J. Scott Turner shows a picture of a mound of *Macrotermes michaelsemi* in Namibia that is over 6m high (2000, p. 195). Even more remarkable are the functions that such mounds perform. They are not just accommodation complexes for the termites, but they are lungs for the colony, they are temperature and humidity regulators and stores of environmental resources and perform many other ‘physiological’ functions. It is Turner’s contention that they are in some sense the exoskeleton for the superorganism that is the whole colony. We saw in section 1 how the termite mound is constructed by the individual termites according to the principle of stigmergy. This is the positive feedback loop between the pheromone-imbued pellets deposited in the partially-built nest structure, and the tendency of an individual termite to drop its pellet on sensing a sufficient concentration of the pheromone. As the mound takes shape, the geometry of the pheromone gradient changes producing a global coordination structure. For example, two mounds being built close to each other will converge eventually producing an arch.

Through the coordination of many simple actions by the pheromone feedback loop, the system self-organises in such a way as to cause the emergence of the performance of the building task. Turner charts many other stigmergic processes in the mound such as processes for repairing the mound and for regulating moisture. Regarding the repair function, the mound is, in effect, a huge low-pass filter that attenuates high frequency oscillations in the surrounding air. Termites are sensitive to the frequency of air pressure variations within the mound, so that if the mound were damaged, the resultant unusual high frequency oscillations would trigger them to repair the breach. Stigmergy explains all these processes of self-organisation.

Although the original context was social insects, theorists have explored stigmergies in situations where there is only one agent present. All that is needed for stigmergy is an agent acting in the world and being coordinated by the results of that action. For example, together with Özlem Yilmaz I have argued that a bean plant can be thought of as a collection of growth tips with associated stigmergies (Sims, R. and Yilmaz, 2023). Each of these growth tips is stigmergic because by growing it changes its position relative to its environment, for example, relative to the goal of reaching a pole which to which it can attach itself (Calvo *et al.*, 2017). A second example is the nest-building activity of the solitary Euminid wasp *Abispa ephippium*, which is a favourite in the stigmergy literature. Theraulaz and Bonabeau (1999, p. 102) describe how each stage of the building process of the nest is triggered by the previous stage - the geometry of the nest itself acting as a coordinating trace. If, in the final stage of the project, a hole is made in the side of the clay nest structure, the wasp will start again from the beginning and build a completely new nest on top of the old one rather than repairing the hole. Nest building is stigmergic - tracked and triggered by the incomplete structure. Since the damaged nest does not resemble any stage of the sequence coordinating construction, it starts the whole process anew. Traditionally, such behaviour is seen as tropistic in the literature, but I have argued elsewhere that such ‘tropism’ is what happens when stigmergic coordination goes wrong, when the normally reliable one-one relation between task and trace variable breaks down (Sims, R., 2022, p. 278). Stigmergy therefore can apply to single agents and its breakdown can explain certain kinds of individual pathology such as tropistic behaviour.

This broad notion of stigmergy can be found in the literature. Francis Heylighen (2016a, p. 4) defines stigmergy as an “indirect mediated mechanism of coordination between actions in which the *trace* of an action in a *medium* stimulates the performance of a subsequent action”. Again there is no requirement for multiple agents in stigmergic interaction. For Heylighen the medium is the part of the environment that is accessible to the organism. The trace is an environmental variable that is correlated with the state of a task[4], and is itself the result of previous performance of the task. In the case of the termites, the trace is the pheromone concentration distribution, while in the case of the *Abispa* wasp it is the geometry (and topology) of the funnel. The stigmergy literature is extensive (see Doyle and Marsh, 2013; Goldstone and Roberts, 2006; Heylighen, 2016b; Holland and Melhuish, 1999; Lewis, 2013; Lewis and Marsh, 2016; Marsh and Onof, 2008; Nieto-Gomez, 2016; Parunak, 2006; Ricci *et al.*, 2007; Secretan, 2013; Sharkey, 2007; Susi, 2016; Susi and Ziemke, 2001; Theraulaz and Bonabeau, 1999; Tsankova and Georgieva, 2004; Turner, 2011; Zamfirescu *et al.*, 2014) and some stigmergy theorists have argued that stigmergic coordination is ubiquitous (for example Parunak, 2006). We are engaged in stigmergic coordination whenever we are navigating around the environment (here the trace is the relation to environment features), building something (trace is incomplete state of building), writing something (trace is incomplete state of text), dancing (trace is relative bodily attitude of dancers), playing a sport (trace is relative bodily attitude of players relative to position of ball say), playing jazz (trace is current musical production plus bodily attitude of players), engaging in online cooperative software development (trace is incomplete code), engaging in market activity (trace is price), or conversing (trace is the current conversational fragment)[5]. These examples give a sense of how broad the notion of stigmergy is and how ubiquitous it is in the world.

How does stigmergy help in understanding cognition? The problem is to work out how a distributed network of processes or agents can self-organise to produce action that is directed towards a goal without having a central controller. Stigmergy is one mechanism by which these constituent processes can be coordinated to produce such emergent organisation. It is not the intention to suggest that stigmergy is the only mechanism that can do this, but it is certainly important. Coordination processes perform the key functions of (1) tracking and triggering task performances by system processes in real

---

[4] Again a reminder that references to tasks does not imply any commitment to representationalism. Tasks are not represented anywhere.

[5] My thanks to Bart Geurts for discussions about the coordination of conversation.

time and (2) switching tasks when necessary – elsewhere I have called these the coordination conditions (Sims, R., 2022, pp. 136–137). These conditions are automatically satisfied by stigmergic systems. The environmental variable that coordinates the system is responsible for task tracking and triggering and is sensitive to task switching. For example, the vector between the growth tip of the bean plant and the pole is a stigmergic variable for the task of reaching the pole, but when it is reached, it is responsible for switching the task to climbing up the pole (Sims, R. and Yilmaz, 2023, p. 271).[6]

The beauty of a stigmergic approach from the point of view of the system is that it is metabolically cheap. There is no need for an elaborate monitoring mechanism or the creation of internal representations to stand in for task completion. It is the performance of the task itself coupled with a sensitivity to the appropriate environmental variable that does the job. The only requirement is that the relevant environmental variable is connected to the state of the task in a one-one relation. This is the stand-in relation, the 'sign' function, that nature obligingly provides. The *Abispa* wasp does not need an internal plan of the nest-to-be together with a representation of its current state to decide when to cease building. It just needs a sensitivity to its geometry. The termites do not need metabolically expensive machinery for monitoring the state of the mound. They just need a simple mechanism for detecting the pheromone and dropping their load.

The stigmergic picture painted here has phenomenological consequences for the system. Stigmergic traces are salient for the organism and stand out from other features of the environment - they have significance. I argue, in the next section, that this significance is established by the longer-term type-2 processes. Termites care about the pheromone because of a long history of biological evolution. Jazz players are sensitive to, and care about the current chord changes because of a longish cultural history of jazz performance. These traces can play their coordinative function because they matter to the components of the system. Phenomenology is nature's (and culture's) way of connecting the system to its past[7]. Stigmergy therefore provides a framework for understanding value. For reasons of space, I shall put this idea aside for now and come back to it in section 5.

The reader should not think that stigmergic systems are destined to be somehow primitive or simple. Through compounding stigmergic traces, systems can act in highly complex and conditional ways; for

---

[6] Richard Menary (2018) has examined the related question of how Epistemological Tracking Tools (ETTs) can keep track of environmental variables. I am interested here in how environmental variables can keep track of tasks performed by the system.

[7] This links with the enactivist concept of sensemaking and the autopoietic idea that the Umwelt is 'brought forth' by the organism (De Jesus, 2015; see Heras-Escribano and De Jesus, 2018; von Uexkull and O'Neil, 2010). I do not have space to take this idea further here but I want to acknowledge that although the origins of this paper lie in research done prior to my PhD thesis in 2022, I have since been inspired by the work of Paolo de Jesus especially his (2016b, 2016a, 2018).

example, stigmergic traces can be arranged in such a way as to act as logic gates. Indeed, in other work I have shown how boundary objects in scientific discourse function as coordination tools because they have the capacity to carry multiple traces that are significant to different groups of people at the same time (Sims, R., 2023b). These can be called multi-trace objects, and they are central to coordination regimes.

To illustrate these ideas in a human social context, consider the example of human beings working towards a joint task in a norm-governed cultural environment such as bartenders working on a large order in a bar. As has been pointed out in the literature (Beach, 1993; Menary, 2007), bartenders rely on material cues in the environment to aid the production of the order such as placing the appropriate kind of glass in the appropriate place on the bar top to keep track of the order. For example, to remember the order of a half pint of IPA the bartender may put a half pint glass near the IPA pump. Moreover, it seems plausible that, in common with many other kinds of coordination between human beings, the bodily attitude and position of bartender colleagues aids the coordination of the joint task. When King Beach forced experienced bartenders to use identically proportioned opaque black glasses in a drinks-pouring experiment, errors increased 17-fold. Richard Menary interprets the use of specialised glasses and position on the bar top as a kind of external memory store for the drinks order (2007, p. 94). The CSA takes a slightly different approach regarding the whole set up, bar top, glasses, pumps and optics, and bartenders as a distributed cognitive system in which certain environmental variables such as type of glass and position act as traces for stigmergic coordination. Stigmergic principles then can explain processes in human social life as well as that of insects.

For those who worry about the boundaries of cognitive systems, the CSA identifies the core of the system - the part that is responsible for the production of goal-directedness - with coordination processes. In stigmergic systems these processes may well be inseparable from ‘plant’ processes such as the termite building operations therefore making the system holistic and nicely identifiable. Ironically, systems that may have more fuzzy edges are those which are more obviously modular. I shall come back to this point in section 5.

The CSA is a task-based account in the sense that it analyses the system as being a performer of tasks. So far I have not said much about tasks and do not have space to develop the idea fully in this paper. Suffice it to say that tasks are by their nature normative - they have some kind of success or adequacy conditions. The reader may rightly wonder where these conditions come from. This is where type-2 reinforcement processes come in.

3 Reinforcement (type-2) processes

In the previous section we saw how trace variables in the environment, like the termite pheromone, can coordinate task performance through stigmergy. The variable being correlated with the result of task performance can both track task completion and trigger further task performance in a dynamic manner. If we accept that such coordination to goal-directed action in a complex system is the hallmark of a cognitive process, then stigmergic coordination is one way in which nature produces cognition.

However, this is not the whole story. We may well ask how it is that system agents become attuned to trace variables associated with a task in the first place, and for that matter where tasks come from. The answer to these questions requires us to look at the evolution of the system on a much longer timescale than that typically associated with coordination processes. It requires us to understand how it is that termite species become appropriately sensitive to the pheromone that coordinates their building actions or to variations in air pressure that triggers and tracks their repair actions. I shall argue that these processes are reinforcement or 'ratchet' processes; they possess a kind of directedness that itself confers normativity upon the tasks faced by agents in the system. I shall say more about this in a moment and warn the reader that we shall move freely between different examples of sociality such as termites and human social systems.

I call these processes 'reinforcement processes' in an attempt to highlight that they are, in some sense, irreversible trajectories in the space of the evolution of a population of similar systems. The term is designed to describe both familiar processes of biological evolution but also the cultural evolutionary 'ratchet' processes discussed by Michael Tomasello and colleagues - processes such as the cultural evolution of tools or language (Tennie *et al.*, 2009). Social processes such as the construction of norms that police the coordination of joint action are also included (see Heras-Escribano *et al.*, 2015; Kiverstein and Rietveld, 2021). In the bartending example, the term describes the evolution of norm-driven patterns of behaviour which allow coordination of joint action through attunement to the appropriate stigmergic traces in the environment. But familiar biological evolutionary processes are also covered such as the attunement of the termite to the pheromone or the sensitivity of the *Abispa* wasp to the topology of the partially built nest and the sensitivity of a bean plant to the position of its growth apices in relation to a pole. Type-2 processes therefore endow our coordination processes with normative force. They explain how there is a need for coordination to produce behaviour that 'gets it right'.

First let us look at two examples to understand the nature of these type-2 processes. In the case of the termites, the key to understanding the role of the pheromone is to note that there is a one-one correspondence between the distribution of pheromone intensity and the building project. It is this relation that allows the pheromone to function as a stand-in for the building task specifically as a stigmergic trace - a trigger and tracker of the task of building. We saw in the previous section how a possible evolutionary argument could proceed. Since the metabolic cost of endowing each termite with a sensitivity to the pheromone vastly undercuts the cost of supplying something like an internal map of the work required, it is conceivable that the stigmergic solution is more often, and more easily, reached in evolutionary history. Moreover, once reached, it is unlikely that such a solution will be lost from the repertoire of the species, at least while nest building remains a critical task. Instead, it may well be refined by adding task-specific detection modules for associated tasks. The conjunction of such modules will serve as a kind of conditional trigger for tasks with more complex structure. This is the basis for arguments for stigmergy in plants where the relation of the organism to its environment is complex and possesses multiple growth points (Sims, R. and Yilmaz, 2023). Once a stigmergic trace sensitivity is set up by type-2 processes, it is likely to be reinforced in future generations given similar environmental conditions. This is what is meant by convergence or the ratchet effect. Nature does not easily give up stigmergic correspondences easily simply because the net metabolic benefits are so large.

The bartender example in the last section illustrates how type-2 reinforcement processes might work in human social systems and how stigmergic processes might emerge, for example, how the use of the position of glasses on the bar top, or the bodily attitude and position of fellow bar staff may come to function as stigmergic trace variables. The argument is that tasks and their stigmergic signalling co-emerge in the evolution of the practices of bar management. Hand-in-hand with the requirement to serve drinks comes a set of physical performances which themselves can act as stigmergic traces to coordinate the relevant tasks. Not only can individual bartenders learn to ‘read’ aspects of these physical performances as coordinative signals, but the interlocking roles associated with bartending may evolve in a convergent ‘ratcheted’ manner through conscious teaching, learning and imitation (for more on roles see Ritchie, 2015, 2020a, 2020b). At one level, then, what gives the bartender system its distinctive unity and systematicity, is the interlocking set of roles (sets of tasks together with characteristic trace variables) that constitute the system. This is how systems such as companies and bars (and termite nests) can retain their identity even when individual bartenders and termites change. And lest we forget, bars operate under conditions where there are economic selection pressures.

Skipping a lot of detail, let us just concentrate on the conclusion of these complex arguments. In a situation involving coordinated joint action, practice and repeated action reinforces both what constitutes the relevant tasks and the set of stigmergic traces that perform the coordination for these tasks. Practice induces normativity in the sense that tasks are generated with well-defined constraints and success conditions - I call this *teleogenesis*. Roles are generated in terms of sets of tasks and stigmergic signs become salient to the players of these roles in order to coordinate the joint action - *semiogenesis.* I argue that the whole situation constitutes a cognitive system: bartenders, material objects, stigmergic traces and the normative tasks and roles. Moreover teleogenesis and semiogenesis go hand-in-hand.

Before we move on, I want to sound a word of warning about normativity. In some parts of the literature, normativity is described as explicit or implicit rule-following by an agent (Brennan *et al.*, 2013, p. 2; De Caro and Macarthur, 2010; Wright, 1963, Chapter 1). In this paper a somewhat weaker notion of norm is adopted. Normativity is tied to the prescriptive nature of tasks, the fact that they can be performed well or badly, and that in the long run there are consequences for the system if too many tasks are performed badly. This is expressed nicely by Manuel Heras-Escribano and colleagues, in the context of an autopoietic account of system viability. Tasks are normative because they are adaptive; there is an ever-present requirement to move the system, in some abstract state space, away from the boundary of viability (Heras-Escribano *et al.*, 2015, p. 23). In the termite example there will be actions on behalf of the whole colony that move it away, in state space, from the boundary where continued existence ceases. In the bartending example, social norms concerning the tasks and roles associated with the situation give tasks their normative force which presumably are linked to the bar continuing to function as an economic entity. Of course, the naturalisation of normativity is a huge field, and it is not my intention to settle the question once and for all here. Suffice it to say that the argument in this paper assumes the possibility of a complex-systems account of normativity and that the stigmergic approach may show how the kind of self-organisation needed for such an account can get off the ground[8].

In short then, type-2 processes are responsible, over a long period of time, for establishing tasks that the system ought to perform. Because of biological evolution the termites collectively ought to perform the task of building a nest. Evolution has put paid to variants that could not do this. Individual termites, elements of the system, ought to be sensitive to the pheromone, and ought to drop their pellets in response to it. Cultural evolution is a type-2 process that is responsible for the establishing the practices of serving drinks in a bar. The bar system ought to serve drinks and get the order right. Individual bartenders ought to recall the order correctly possibly with the help of reading the sign that is the glass placed in the appropriate place on the bar-top. Tough economic competition has put paid to bars that could not perform these tasks consistently. Reinforcement is responsible for producing tasks with success or adequacy conditions, and for endowing signs with salience.

Suppose the reader accepts, for the sake of the argument, that type-2 processes do what I say they do. Still, there is the question what part these processes play in the explanation of goal-directedness or cognition. The answer is that reinforcement processes play an indirect role. The immediate explanation for goal-directedness, say, of termite nest building, is, as Grassé originally proposed, the coordinative

---

[8] Of course one of the problems with the account of Heras-Escribano et al, is that when the system reaches the boundary of viability it is too late. The system is already disrupted. This is why type-2 processes are needed to guide the system away from the boundary before this state is reached on the basis of past successful actions of the system. As mentioned, the way history ‘communicates’ this to the system in the here-and-now is through phenomenology. I do not have space to develop this idea here but it will form the basis for a future paper.

processes involving the pheromone and termite pellet-dropping behaviour. However, the whole explanation, why we consider the termite nest and its pheromonal environment as the relevant system, rather than say that of other biochemical exudates of the termites, is to do with the role played by these elements in type-2 processes. Through type-2 processes the pheromone has become salient to the termites as a stigmergic trace whereas other exudates have not (even if they bear the right relationship to the task – there is a degree of contingency). Type-2 processes bestow coordinative significance on environmental variables.

It is this secondary explanatory role attached to type-2 processes that blocks cases of accidental stigmergic coordination as being genuinely cognitive. For example, Nick Brancazio and Russell Meyer consider the case of active materials which are self-propelled, non-living entities which exhibit a number of cognitively interesting behaviours such as gradient-following and obstacle-avoidance (Brancazio and Meyer, 2023). They consider, for example, a self-catalyzing oil droplet (Hanczyc, 2011; Hanczyc and Ikegami, 2010; Holler and Hanczyc, 2019). Such a droplet can travel through a liquid environment autonomously. In certain circumstances it can follow gradients and exhibit behaviour that is analogous to that of chemotaxis in bacteria – it can solve mazes for example. To cut a long story short, Brancazio and Meyer do not find that these active materials exhibit essential features of cognition but ask the interesting question whether they can model cognition in an explanatorily fruitful manner. For its part, the CSA gives a simple reason why these active materials are not cognitive systems despite imitating what looks like type-1 stigmergically coordinated chemotaxis: the lack of requisite type-2 processes to generate normativity of task and salience of environmental variables. The movement through the environment is incidental – it is not task driven. It does not benefit the oil drop to do so – the system is not moved away from the boundary of viability as a consequence of the movement. To put it another way, the oil droplet is not a consumer of these environmental variables as signs. We shall look at signs in the next section.

Before we do that, I want to outline briefly the differences between the CSA and classical teleosemantic theories such as those of Millikan (1984, 1989), Papineau (MacDonald and Papineau, 2006; 1984), Neander (1995, 2018), and Dretske (1981, 1988, 2006)[9]. I shall start with the similarities. Both approaches emphasise the significance of historical processes and both do so because of the realisation that cognitive processes are normative. The teleosemanticist emphasise that proper functions, the functions performed by traits that are selected, are normative – they have success conditions. For the CSA, tasks are normative and contribute to the well-being of the system. They do not do so randomly, they are canalised by type-2 processes to be established patterns of behaviour that are goal-directed, and these goals are normative. Both approaches emphasise the importance of the consumer – for

---

[9] My thanks to an anonymous reviewer for urging me to make this difference clearer.

teleosemantics the consumer of the representation – for the CSA the consumer of the trace variable, the environmental sign. The principal differences occur when examining misrespresentation or task failure. Because the teleosemanticists want to provide a strong platform for a theory of semantic content they need a notion of representation that is robust enough to do the job. Misrepresentation occurs when a subsystem whose proper function it is to recognise a representation as meaning X does not do so. X is fully-fledged representational content. However in the CSA 'misprepresentation' occurs when the system produces an inappropriate action given the state of an environmental variable. This could be due to the breakdown in the correlation between the variable and the task. For example an experimenter could mess up the foraging activities of the ants by laying a false trail of the pheromone. This breaks the stigmergy of the system. The difference comes down to different kinds of content. For Millikan and friends, representations carry fully-fledged propositional content about the world. For the CSA stigmergic traces carry no content at all in this sense. If they have content at all, it is something like 'appropriate action' – a kind of irreducible mixture of indicative and prescriptive content (perhaps analogous to Millikan's pushmi-pullyu representations (1996)). The disruption of the pheromone by the experimenter is not really misrepresentation in a robust sense. The sign is not saying that the world is other than it is. What goes wrong is that the action that results from the sign does not necessarily lead to the benefits to the well-being of the system that issued in the past[10]. The approach taken here, then, borrows a little from the teleosemanticists in order to show how tasks can be normative in the here and now and how signs can be salient. This line of argument is in keeping with recent enactivist work such as that of Hutto and Myin (see 2017, Chapter 5) especially their point about "stripping teleosemantics of its semantic ambitions and putting it towards a new and different theoretical use" (2017, p. 104). This seems to be exactly what the CSA attempts to do.

4. Signs and semiotics

The previous sections have referred, in a rather cavalier fashion, to environmental trace variables as stigmergic *signs* and have done so without justification. This section seeks to justify a semiotic interpretation of stigmergy through Peirce's theory of signs - after all the term has a Greco-Latin origin based on the word stigma meaning mark. I am aware that other authors have made this point (De Jesus, 2016b; Dipple *et al.*, 2014; Huang *et al.*, 2008; Tummolini and Castelfranchi, 2007), although I differ in varying respects from these accounts.

Before we start, a word or two of warning. I described the approach taken here as being a kind of cybernetic enactivism. As with many varieties of enactivism, it is based on a rejection of standard cognitivist accounts of cognition as constituted by the manipulation of contentful representations. It

---

[10] This is discussed in more detail elsewhere (Sims, R., 2023a) and in a paper in preparation.

may therefore be surprising that this section seems bent on portraying the stigmergic trace as a sign. However, as we saw in the previous section, the line taken here is that a commitment to a Peircean sign is not thereby a commitment to fully fledged representational content. I shall explain this further below.

In his theory Peirce understands a sign to be a tripartite relation: a sign is “something which stands to somebody for something in some respect or capacity” (Peirce *et al.*, 1960, p. 2:228). A sign signifies an ‘object’ (its meaning) to an ‘observer’ that interprets it. An object can be, for example, an actual object like an animal walking in the snow, but it can also be an action to be performed, and the observer or consumer of the sign could be a subsystem or system agent that performs an action in response to the sign. Peirce uses the term ‘interpretant’ to label the disposition of the observer to associate the sign with the object. For example, take an animal as the object. The signifier is the track of the animal in the snow, the observer may be someone in the forest and the interpretant the disposition to associate danger with the observed track because of the causal link between the dangerous animal and the track. In the case of the termites the object could be the collection of pellets, the observer the individual termite, the signifier is the pheromone, and the interpretant is the disposition of the termite to drop the pellet that it is carrying. There need not be any internal representation corresponding to the sign such as a Fodorian meaning in a language of thought (Fodor, 2008). For example, the pheromone concentration signifies dropping a pellet as interpreted by an individual termite – but this meaning need not be represented anywhere in the termite. Peirce posits three categories of signs: icon, index and symbol. The symbol is the standard bearer of conventional meaning thoroughly familiar to human beings and beloved of the cognitivist, the icon is the sign that signifies by virtue of some kind of structural similarity, and the index, which is what I propose stigmergic traces to be, signifies because of a lawlike connection between sign and interpretant based on regularities in nature. It is because the pheromone reliably accompanies previously built structures that it can serve to trigger further building. The advantage of Peirce’s framework is that it is well disposed to deal with situations which do not fit classical communication where there is a sender, a message, and a receiver. In stigmergic systems the environment does the ‘sending’ and the environment does not have a ‘message’ to send.

Signs are only useful if they can be misinterpreted or if they can, with the caveats already discussed, ‘misrepresent’. As discussed in the previous section, ‘misrepresentation’ occurs if the trace signal does not correspond in the right way to a task. For example, we discussed the case where a wicked entomologist deposits a decilitre of pheromone inside a perfectly built termite nest to trigger a spate of unnecessary construction. Similarly, the experimenter poking her finger through the surface of the *Abispa* nest funnel disrupts the topology of the nest and creating a sign for a re-initiation of the whole building process from the point of view of the wasp as interpreter. In both these cases, the stigmergic trace does not trigger the appropriate task. The termites engage in unnecessary construction and the *Abispa* starts her whole nest-building project anew. These are not appropriate tasks in the sense that they do not contribute to moving the position of the system away from the limits of viability in the appropriate state space according to Heras-Escribano's definition. Misrepresentation of a stigmergic trace sign results in inappropriate action.

At this point the reader could be forgiven for wondering why bother introducing signs given the potential for misunderstanding, (and the potential for alienating potential supporters of the theory such as ecological psychologists and radical enactivists), and given that we can do all the analysis in the previous sections without them. One strong reason is that we can use some of the well-developed machinery of biosemiotics to refine the theory. For example, we can introduce the concepts of semiotic depth - the degrees of freedom available for which the trace variable can coordinate a variety of complex actions, and semiotic freedom - the degree to which the system is free to interpret the trace variable rather than being purely determined in its response (Hoffmeyer, 2007)[11]. This last notion helps defend the CSA from the objection that it only seems to deal with stimulus-response interactions and does not entirely capture the 'creative engagement' characteristic of cognitive systems.

But by far the biggest advantage of thinking of the stigmergic trace as a sign is that it gives us a way of dealing with agency. We have already spoken of semi-autonomous system elements as agents. A semiotic interpretation helps bolster this move to supply us with a relatively thin notion of agency. An agent is simply a consumer of signs. Thus, individual termites are agents in that they consume or interpret the pheromone trace as part of the coordination of their building activities. But the whole colony is also an agent as the consumer of the trace signal given by the amount of food left in a food source (and as mentioned earlier the whole colony may change its behaviour depending on the distance richness and amount remaining of a food source). This gives us the possibility of explaining how it can be that there are agents at multiple levels of organisation within a system which solves puzzles raised by colonies of single celled organisms such as slime moulds (cf Godfrey-Smith, 2016; Goldstone and Theiner, 2017). Signs provide a simple way of thinking about agency.

5. Putting the CSA to work

The CSA originated as an attempt to provide theoretical grip on some classical problems in the 4E literature (Sims, R., 2022). Of course, there is not enough space here to do justice either to the problems or their putative solutions, but the aim of this section is just to give a flavour of the coordinated systems approach in action. The three problems are selected to illustrate how the approach may cast light upon

---

[11] Note that in his definition of stigmergy Heylighen carefully refers to the trace as *stimulating* task performance not determining it (Heylighen, 2016a, p. 4). Thus, there is room here for 'degrees' of compliance in an interpretation. A sign is not the same as a cause.

a diverse range of issues. The CSA analysis is sketched in the broadest outline, and it is taken for granted that there are omissions and the need for refinement.

5.1 The problem of cognitive bloat

This is one of the classical problems raised in response to the original Clark and Chalmers paper mainly by Fred Adams and Ken Aizawa (Adams and Aizawa, 2001, 2009, 2010; Aizawa and Adams, 2005). The danger with any kind of active externalist account of cognition, especially those accounts that invoke functional similarity with inner cognitive processes, is that they become too inclusive and that any process that is causally involved with the system acquires cognitive status. One example here is the use of the internet to aid my biological memory in looking up an address of a museum I want to visit. It appears that the use of a website containing the address of the museum is functionally similar to biological memory. Clark and Chalmers' parity principle (1998, p. 8) permits the inference that such a website is part of my mental apparatus. The worry is now that this argument can be extended to the whole internet and that the concept of a cognitive system has become too bloated to do any useful work.

I have argued elsewhere that the coordination conditions set up by the CSA, tracking, triggering and task sensitivity, can be taken to be a 'mark of the cognitive' (Sims, R., 2022). These conditions are stronger than the functional similarity of the parity principle. For example, while the use of the internet may trigger tasks such as walking to the museum it does not satisfy the tracking condition. The causal relation is just one way. But reciprocal causal coupling on its own is still not enough according to the CSA *pace* some theorists such as Palermos (2014). For part of the internet to be a stigmergic trace it requires that the trace updates with the performance of the task in the right way, as an index sign. It must preserve the lawlike relation that allows the trace to be interpreted as, say, progress of the task towards the museum. From a semiotic point of view a sign plays a very specific role with respect to an interpreting system that cannot be captured merely by saying that it is causally coupled.

An example here would be the use of GPS in guiding the task of walking to the museum. Now there is stronger case for a limited part of the internet to be included in the cognitive system. The little blue dot on the GPS map corresponds in the right way to task of walking to the museum. Thanks to a complex network of geostationary satellites, that little blue dot is a stigmergic trace for the task. The satellite network ensures the right lawlike relation between task and sign. Performing the walking task moves the position of my phone which then interacts with the satellite network and its ground infrastructure in the right way to move the blue dot accordingly. The blue dot position is therefore a consequence of previous walking actions and is a stigmergic trace.

The lesson here is that cognitive bloat is a result of analysing the system at too low an explanatory level. The existence of causal relations, even reciprocal ones, is not sufficient for cognitive systemhood. What is needed is that these causal relations satisfy the coordination conditions which are pitched at the functional level, and, if the system is stigmergic, at the semiotic level.

5.2 Sims and Kiverstein's problem of external memory

The second application is suggested by a very interesting paper by Matt Sims (no relation) and Julian Kiverstein (2022) in which they take up the question of external memory in slime moulds (2022) – for brevity I shall refer to this paper as SK. The CSA may offer further lines of development to their approach and possibly solve a few outstanding issues. I should say at the outset that I am deeply sympathetic to this paper, not least because of its mobilization of stigmergy theory. My criticism is perhaps that they can take the stigmergic approach further and by so doing avoid some problems that are raised by confusing stigmergy with classical communication of messages. I shall only sketch part of their argument here and suggest a possible way forward.

Slime moulds have been much studied in relation to topics of interest to 4E theorists (Castro Garcia, 2011; Reid, 2023; Reid *et al.*, 2012; Reid and Latty, 2016; Smith-Ferguson *et al.*, 2017; Smith-Ferguson and Beekman, 2020; Walmsley, 2020). In particular, the slime mould *Physarum polycephalum* has being touted as an example of an organism with an external memory. *P. polycephalum* is a unicellular multi-nucleate slime mould that in the vegetative (plasmodium) stage of its life cycle has shown itself to be quite capable of performing cognitive tasks such as navigating mazes and making decisions (Nakagaki *et al.*, 2000, 2004). Moreover, when a *Physarum plasmodium* moves[12] through its environment it leaves behind a clear residue of non-living extracellular slime (Sims, M. and Kiverstein, 2022, p. 29). When it encounters this slime in the environment it tends to avoid it, thus producing a dynamic landscape in which depleted food sources are repellers - the opposite of attractors. SK claim that the slime trail acts as a kind of external memory for the organism.

They justify this claim by appealing to the functional similarity with internal generalised biological memory defined by Baluška and Levin in the following terms: "memory is defined as experience-dependent modification of internal structure, in a stimulus-specific manner that alters the way the system will respond to a stimulus in the future as a function of *its* past" (Baluška and Levin, 2016, p. 902 emphasis added). But traditional internal memory is located inside an organism which bestows privileged access and a special 'ownership' relation between memory and organism.

---

[12] The locomotion of the plasmodium is a complex affair more akin to growth.

I mention this feature of memory - that it has an owner – “its past” in the definition - because the ownership relation may be problematic in cases where memory is external to the organism. What fact about the relation between the external structure and the organism fixes ownership? In the human case ownership of memories is fixed, at least on a first analysis, by phenomenology. Only my memories are phenomenologically present to me. But what happens with external memory which in principle can be available to other individuals? Further, do we want a criterion for fixing ownership that is available to the organism itself - what I shall call the *transparency condition*? Perhaps we are satisfied with a condition that will allow attribution of ownership of memory by an omniscient observer but not necessarily available to the organism itself.

Alternatively, ownership of external memory can be fixed by some kind of physical connection. For example, in thinking of a notebook as an external memory module, Clark and Chalmers seek to avoid this problem by introducing the glue and trust conditions (Clark and Chalmers, 1998, p. 17); namely that he notebook is constantly accessible, and the information it contains is directly available and so on. Ownership of this kind will be a problem for organism-centred extended cognition where external processes are taken to be extensions of an organism’s cognitive capabilities. Distributed cognitive systems which include organisms and parts of the environment but that are not organism-centred are less prone to this problem. The CSA belongs to the latter category.

Fixing ownership of external memory is not a problem in the somewhat artificial situation in which there is just one *P. polycephalum* plasmodium X in a petri dish. By definition, any slime trail that the organism encounters must belong to X and claiming that the slime trail is the external memory of X seems unproblematic. But given two organisms X and Y in the same dish, a trail produced by X will modulate the behaviour of Y in exactly the same way that it modulates the behaviour of X, and in each case satisfying an externalized version of Baluska and Levin’s definition of generalised biological memory even though it is not produced by Y. In other words, the problem is that there is no *transparent* fact of the matter what makes the trail the memory of X rather than Y given that X and Y cannot distinguish ownership of such a trail.

SK are aware of this problem. Their answer is interesting but unconvincing:

> “If the cue/signal is used by the same organism that produced it then that external memory belongs to the producer/user; it is tied to that organism’s past interactions with its environment. On the other hand, if the cue/signal is used by an organism other than the one which produced it, then the external memory belongs to the *joint producer/user system*.” (Sims, M. and Kiverstein, 2022, p. 31 emphasis original).

The problem is that the transparency condition is violated; the organism itself cannot tell whether a given slime trail is its own production or not. It would require an observer with perfect knowledge of the history of the organism to make that connection. The attribution of external memory becomes parasitic upon the memory of the external observer. One solution is to drop the transparency condition. However, it seems that SK are committed to it because of their use of an extended parity principle. The systems which we commonly take to be a benchmark for the parity principle are human cognitive capacities, or those of higher mammals, that satisfy the transparency condition. Therefore, one might expect that the transparency condition is necessary for extended cognitive capacities if one is committed to these capacities being functionally similar to traditional internal capacities.

When there is more than one organism involved, SK switch tack. Instead of an organism-centred system they now argue for a distributed one. They claim that X and Y is now a collective entity that possesses a joint external memory. But this raises the question: on what grounds can X and Y be considered a collective? It is true that they are conspecifics and happen to be interacting with each other’s slime trails, but does this mean that whether X and Y are a collective entity depends on the contingent fact of whether they encounter each other’s trials? SK admit as much: “a cue or signal that is never used does not count as part of external memory”. This seems to be a matter of chance – and it is difficult to see how it can do any useful metaphysical work. When they do encounter each other’s trails, we could claim that there is some coordination. However, there is no normativity applying to the group over and above that applying to the individuals. The coordination is by chance; the system is not set up to be like this to conduce to a group goal.

One option is to jettison the psychological category ‘memory’ and with it the extended parity principle. By so doing they also sidestep the ownership and transparency issues. But this is a move SK are reluctant to make since they are committed to the biogenic approach of Pamela Lyon and Fred Keijzer and others, where memory is one of many tools in the cognitive toolbox and as such is owned by the organism (Lyon, 2006, 2020; Lyon *et al.*, 2021; Lyon and Keijzer, 2007). As it stands SK have the resources within their grasp to resolve this issue but in my view do not go far enough. They recognise the importance of stigmergy in their paper without realising that a stigmergic approach can solve the puzzle.

So how does the CSA solve the puzzle? For a start it simplifies things. The slime trail is a sign which is coordinative of the foraging task whether produced by the organism or not. However, it is not helpful to think of the slime trail as being external memory belonging to the organism, it is not helpful to take an organism-centred view because of the difficulty in resolving the ownership question. Rather they should take the cybernetic enactivist approach of this paper and accept that it is the coordination of the whole system which is at stake not just a question of memory tasks being ‘outsourced’ into the environment by the organism. The coordinative slime is part of the system for coordinating the foraging task. The CSA also denies that two *Physarum* plasmodia X and Y together constitute a group, the foraging task is not a group task, so the CSA would not endorse SK’s solution in this case.

Why don’t two *Physarum* plasmodia X and Y constitute a system, and why is the foraging task not a group task? It will be helpful to draw a distinction with the termite colony. In the case of the termites the building task is a joint task, and the system can be thought of as a collective. Why? Because type-2 processes acting over evolutionary timescales have produced an emergent group goal which is linked to the collective building task coordinated by the pheromone. A failure in the building task threatens the viability of all termites, however a failure by *Physarum* plasmodium Y to forage does not threaten X’s viability. In fact, the opposite is the case, the failure of Y to complete the foraging task adequately benefits X since, in this simple situation, it is a zero-sum game and Y’s loss is X’s gain. In the case of X and Y the combined slime trail coordinates tasks conducing to individual goals and there is no sign for which X and Y is a consumer, while in the case of the termites the combined pheromone trace coordinates to a collective goal and there are stigmergic signs in the environment such as richness/distance of food source which are consumed by the colony as a whole. This is because the requisite reinforcement processes are not in place to create a group task for X and Y and therefore a cognitive system[13].

As a final point, offered in a spirit of wanting to be helpful, there are some distinctions in the paper that are not really needed. Signals are distinguished from cues, stigmergy from patch-marking, and pilotage from stigmergy. Perhaps this could be simplified - all you need is that the stigmergic trace is a sign that coordinates the action of one or more agents into a whole system. The requirement for cognition here is that stigmergic coordination is normatively set up by type-2 processes. I shall deal briefly with these issues to show how the CSA can offer a way of cutting through these distinctions.

The worry whether stigmergic traces are signals - involving a sender - or cues - without a sender seems to be misplaced *pace* Tummolini and Castelfranchi (2007). Taking a Peircean view removes these problems and considerably simplifies the picture. Signs can be produced by organisms, but they need not be. They can be features of the natural world such as light intensities or degrees of humidity. All they need to be signs is that they are interpreted as standing for something and, in the case of an index, this is established by a regularity of nature. Light intensity may be interpreted by a plant as a sign to start growing leaves again in the spring. A lack of humidity may be interpreted likewise as a sign to

---

[13] Of course, there could be a mechanism through which X and Y do cooperate, say against a species predator, and by so doing are sensitive to a group-level trace variable, in which case there may be a possibility for asserting degrees of group agency – but the simple analysis here assumes that this is not the case.

grow a deeper root system. Moreover, these signs do not have propositional content in the standard sense.

SK unnecessarily distinguish stigmergy from patch-marking - the process of marking some parts of the environment as having been visited in the foraging task. On encountering such a mark, the organism is averse to it and takes a different route. This fits the definition of stigmergy in Heylighen (2016a, p. 4) that SK themselves cite. As mentioned in section 1, there is no requirement for multiple agents in stigmergic interaction. Therefore, there is no need for the distinction.

Finally, SK write that pilotage is not stigmergic (2022, p. 31) and that a bacterium navigating up a sucrose gradient is not engaged in stigmergy. I want to bluntly refute this. As I have written elsewhere (Sims, R., 2023a; Sims, R. and Yilmaz, 2023) such a bacterium *is* engaging in stigmergy - *positional stigmergy*. The position of the bacterium relative to its environment is a result of its previous locomotive actions and triggers future locomotive actions. The stigmergic trace is therefore the relation of the organism to features of the environment. It has been a commonplace for some time in some quarters that relational environmental variables are coordinative of action - for example see the literature on tau theory (Delafield-Butt, J. *et al.*, 2010; Delafield-Butt, J. T. *et al.*, 2012; Frazier *et al.*, 2020; Lee, 1998; Lee *et al.*, 2001) and of course the whole field of ecological psychology (Gibson, 2014). It seems that to fail to recognise stigmergy in signs from the environment that are the result of previous motor activity is perhaps to ignore one of the most powerful kinds of stigmergy available, that of sensorimotor coupling.

In my view, this otherwise excellent paper could benefit by taking stigmergy further along the lines of a CSA approach.

5.3 The problem of group cognition and multiple levels of agency

This problem is perhaps related to the previous one when considering what kinds of arguments can be made to support the claim that groups can usefully be considered themselves cognitive agents. Again returning to termites, it makes sense to claim that there is a goal-directed group agent for two reasons:
(1) The actions of the termites in a colony are coordinated to a group goal stigmergically[14].
(2) The group goal and the attunement to the stigmergic trace are the result of type-2 reinforcement processes.

---

[14] Remember that a group goal does not require that it is ‘held’ individually or that individuals are even aware of it.

It is not only that there is an environmental variable coordinating behaviour. Rather it is that this coordination and the attendant sensitivity to the variable is set up by longer-term type-2 processes in a way that makes them task-relative and normative. The termites are tasked with building the nest. Failure to do so, or failure to do in an appropriate manner, will push them towards, and even over, the boundary of viability. The teleogenesis of the task and the associated norms in terms of attunement to the stigmergic trace and the tracking and triggering of appropriate actions are all to be found in type-2 processes. Recall that it is due to the absence of these processes that the collection of *P. polycephalum* plasmodia is not a group agent.

A similar argument can be put together in the case of human joint action. The bartenders in a bar acting jointly on a drinks order also satisfy (1) and (2). They are coordinated by a stigmergic trace consisting of the position and state of the glasses and the position and bodily attitude of the other bartenders. This coordination can be thought of in terms of roles which are, roughly speaking, sets of tasks and associated stigmergic traces. But there are also type-2 processes associated with the evolution of such tasks and roles and the normativity that attends them. These processes are predominantly cultural and are themselves embedded in a complex network of roles and tasks – that elsewhere I have described as *role-interlocking*[15]. It is sufficient to point to the existence of coordinative processes and suggest type-2 process mechanisms to demonstrate the existence of a group agent.

One such mechanism is the process by which conventions become normative. That way conventions that define, separate and coordinate roles take on type-1 status (Brandl *et al.*, 2015; Csibra and Gergely, 2011; Geurts, 2018; Keupp *et al.*, 2013; Rakoczy *et al.*, 2008; Schmidt *et al.*, 2011; Slors, 2021; Wexler, 2006). For example, the simple repetition of a set of behaviours can set up a coordinative convention, that over time, becomes a norm that is penalised if transgressed. Whatever processes are responsible, and I shall leave it up to the anthropologists to supply the details, these are type-2 processes responsible for the establishment of stigmergies that create group cognition. A good example here might be the phenomenon of the Mexican wave at a sports stadium, where the position of the wave is the stigmergic trace. Reinforcement processes that endow repeated behaviours with normativity are responsible for transforming such a behaviour into a task.

These examples naturally lead to a discussion of cognitive agency. In previous sections I have been rather gung-ho about using the term ‘agent’ to apply to a set of processes that are coordinated stigmergically to perform a task. Levin and Dennett (2020) are happy to use the term ‘agent’ for such system components even though they are not capable of thought in a traditional sense:

---

[15] Talk given at the workshop: *Coordination, Norms, and Group Action,* at the Center for Cognition, Culture and Language, Radboud University 11th July, 2023.

> “(…) individual cells are not just building blocks, like the basic parts of a ratchet or pump; they have extra competences that turn them into (unthinking) *agents* that, thanks to information they have on board, can assist in their own assembly into larger structures, and in other large-scale projects that they needn’t understand” (2020, pp. 7–8 emphasis added).

Quite so. Often such an agent is constituted by components that are themselves agents in some sense. For example, the termite colony is such a group agent that is constituted by termites (and other material items such as pheromone traces). It is reasonable to consider each termite as an individual agent given that it possesses some minimal autonomy, perhaps they have more claim to agency than Levin and Dennett’s cells. Therefore, there seems to be a situation in which agency occurs at different levels of organisation of a system simultaneously.

Although counterintuitive perhaps, multiple interacting levels of agency may well be a feature of the CSA account rather than a bug. Indeed, other theorists have not shied away from multiple levels of agency: Hutchins explores the dynamics of agents consisting of agents in distributed cognitive systems (1995a), Goldstone and Theiner (2017) offer a computational systems account of group cognition in such terms, and Peter Godfrey-Smith (2016) explores the transition between single and multicelled organisms in terms of agency and subjectivity. It is entirely possible to conceive of a multicellular slime mould both as a population of single celled organisms and a single compound organism. The CSA is completely sympathetic to these approaches and indeed may be able to contribute an understanding of an agent as a biosemiotic entity; an agent is an interpreter of the stigmergic trace sign. In this sense as interpreters of signs, single-celled organisms are agents such as those linked to the production of fruiting bodies in slime moulds or the individual termites interpreting the pheromone trace. But the whole system is also simultaneously an agent because it interprets and responds to other stigmergic traces. As mentioned earlier, the termite colony stigmergically responds to the quality and distance of food sources. This is not a semiotic activity that is taking place at the individual termite level but at the group level. A stigmergic account has no problem with the multiplicity of salience of environmental signs at different levels of the system or that there is interaction between agency at these different levels. Similarly, there is no problem analysing the whole bar system as responding to a drinks order or a whole company responding to a change in economic conditions while at the same time asserting that individual employees are also agents responding to a different set of signs. If we can overcome the idea that signs need to possess content or need to be understood somehow as messages, then they can do useful work in understanding the kind of thing that a system is and in what sense(s) it is agential.

Conclusions

There is more work to do, but the paper has set out the CSA framework and shows how it helps think of cognition in both social and historical terms. The systems is composed of many agents, and these agents face tasks and react to signs that are made salient through type-2 processes. It shows how complex goal-directed behaviour by a complex system can be the result of a special kind of stigmergic coordination process of much simpler subsystems, and that these may be thought of as agents. It links historical processes to the production both of normative tasks and to phenomenology and the way in which environmental signs become salient. It therefore offers a solution to the transparency problem by which the system can be aware of features of its own past.

Of course it raises new puzzles. If agency is a matter of consuming signs in a way that promotes coordinated action in the world, and if sign consumption comes with its own phenomenology – that signs become salient to agents, then do we need to worry about the phenomenology of multiple levels of agency? Maybe we need to think about the characteristics of type-2 processes that do a lot of work in the account and hold the key to both teleogenesis, semiogenesis, and agentogenesis. In this connection, what should we make of Putnam’s warning that we cannot get out of evolutionary processes more intentionality than we put in (1995, p. 33)?

But the approach presented here seems to have possibilities for furthering the debate on some questions regarding 4E cognition. It provides a framework for dealing with stubborn problems such as cognitive bloat by adverting to a description of cognitive systems that is at a higher and more abstract level than causality. It suggests a way of dealing with minimal cognition without needing cognitive categories such as memory that cause unnecessary complications of ownership when extended outside the organism. It suggests a way forward for dealing with group cognition by treating sub-personal and super-personal processes as being part of the same conceptual framework. It has the potential to solve problems to do with individuation of collective systems through a description of a system as a set of interlocking roles. It suggests that a Peircean approach to coordination via natural signs supplies a good mechanism for self-organisation of certain complex systems including biological ones.

Perhaps most of all the framework promises to bring together cybernetic, enactive, dynamical systems, distributed cognition, ecological psychology, and perhaps even computationalist approaches to cognition.

References:

Adams, F., and Aizawa, K. (2001) The bounds of cognition. *Philosophical Psychology* 14(1): 43–64.

Adams, F., and Aizawa, K. (2009) Why the mind is still in the head. In Robbins, P. and Aydede, M. (Eds.), *The Cambridge Handbook of Situated Cognition*. Cambridge University Press.

Adams, F., and Aizawa, K. (2010) Defending the bounds of cognition. In Menary, R. (Ed.), *The Extended Mind*. Cambridge, Mass.: M.I.T.Press.

Aizawa, K., and Adams, F. (2005) Defending non-derived content. *Philosophical Psychology* 18(6): 661–669.

Bacigalupi, J. A. (2022) Semiogenesis: A Dynamic System Approach to Agency and Structure. *Biosemiotics* 15(2): 261–284.

Baluška, F., and Levin, M. (2016) On Having No Head: Cognition throughout Biological Systems. *Frontiers in Psychology* 7.

Barbieri, M. (2009) A Short History of Biosemiotics. *Biosemiotics* 2(2): 221–245.

Beach, K. (1993) Becoming a Bartender - the Role of External Memory Cues in a Work-Directed Educational Activity. *Applied Cognitive Psychology* 7(3): 191–204.

Beckers, R., Deneubourg, J., Goss, S., and Pasteels, J. (1990) Collective Decision-Making Through Food Recruitment. *Insectes Sociaux* 37(3): 258–267.

Beckers, Ralph, Holland, O., and Deneubourg, J.-L. (1994) Fom Local Actions to Global Tasks: Stigmergy and Collective Robotics. *Proceedings of the Workshop on Artificial Life* 4.

Bickhard, M. H. (2017) Information, Representation, Biology. *Biosemiotics* 10(2): 179–193.

Brancazio, N., and Meyer, R. (2023) Minimal model explanations of cognition. *EUROPEAN JOURNAL FOR PHILOSOPHY OF SCIENCE* 13(3).

Brandl, J. L., Esken, F., Priewasser, B., and Rafetseder, E. (2015) Young children's protest: what it can (not) tell us about early normative understanding. *Phenomenology and the Cognitive Sciences* 14(4): 719–740.

Brennan, G., Eriksson, L., Goodin, R., and Southwood, N. (2013) *Explaining norms*. Oxford: Oxford University Press.

Calvo, P., Raja, V., and Lee, D. (2017) Guidance of circumnutation of climbing bean stems: an ecological exploration. *bioRxiv* 122358.

Castro Garcia, O. (2011) Principles of Minimal Cognition in Smart Slime Molds and Social Bacteria. *Pensamiento* 67(254): 787–797.

Clark, A. (2011) *Supersizing the mind: embodiment, action and cognitive extension*. Oxford: Oxford University Press.

Clark, A., and Chalmers, D. (1998) The extended mind (Active externalism). *Analysis* 58(1): 7–19.

Consiglio, F. (2020) *Culture as a collective mind: the emergent properties of a community*. University of Granada, Granada.

Csibra, G., and Gergely, G. (2011) Natural pedagogy as evolutionary adaptation. *Philosophical Transactions of the Royal Society B-Biological Sciences* 366(1567): 1149–1157.

Cummins, F., and De Jesus, P. (2016) The loneliness of the enactive cell: Towards a bio-enactive framework. *Adaptive Behavior* 24(3): 149.

De Caro, M., and Macarthur, D. eds. (2010) *Naturalism and normativity*. New York ; Chichester: Columbia University Press.

De Jaegher, H., and Di Paolo, E. A. (2007) Participatory sense-making: An enactive approach to social cognition. *Phenomenology and the Cognitive Sciences* 6(4): 485–507.

De Jesus, P. (2015) Autopoietic enactivism, phenomenology and the deep continuity between life and mind. *Phenomenology and the Cognitive Sciences* 15(2): 265.

De Jesus, P. (2016a) Making Sense of (Autopoietic) Enactive Embodiment: A Gentle Appraisal. *Phainomena* 25(98/99): 33–56.

De Jesus, P. (2016b) From enactive phenomenology to biosemiotic enactivism. *Adaptive Behavior* 24(2): 130.

De Jesus, P. (2018) Thinking through enactive agency: sense-making, bio-semiosis and the ontologies of organismic worlds. *Phenomenology and the Cognitive Sciences* 17(5): 861.

Deacon, T. W. (1997) *The Symbolic Species*. New York, London: Norton.

Deacon, T. W. (2021) How Molecules Became Signs. *Biosemiotics* 14(3): 537–559.

Delafield-Butt, J., Galler, A., Schogler, B., and Lee, D. N. (2010) A perception-action strategy for hummingbirds. *Perception* 39(9): 1172–1174.

Delafield-Butt, J. T., Pepping, G.-J., McCaig, C. D., and Lee, D. N. (2012) Prospective guidance in a free-swimming cell. *Biological Cybernetics* 106(4–5): 283–293.

Dipple, A., Raymond, K., and Docherty, M. (2014) General theory of stigmergy: Modelling stigma semantics. *Cognitive Systems Research* 31–32: 61–92.

Doyle, M. J., and Marsh, L. (2013) Stigmergy 3.0: From ants to economies. *Cognitive Systems Research* 21: 1–6.

Dretske, F. (1981) *Knowledge and the flow of information*. Cambridge, Mass.: M.I.T. Press.

Dretske, F. (1988) *Explaining behavior: reasons in a world of causes*. Cambridge, Mass.: M.I.T. Press.

Dretske, F. (2006) Representation, teleosemantics, and the problem of self-knowledge. In MacDonald, G. and Papineau, D. (Eds.), *Teleosemantics*. Oxford: Oxford University Press.

Dupré, J. (2012) *Processes of life*. Oxford: Oxford University Press.

Dupré, J. (2020) Life as a process. *Epistemology & Philosophy of Science* 57(2): 96–113.

Dupré, J., and Nicholson, D. J. (2018) A manifesto for a processual philosophy of biology. In Dupré, J. and Nicholson, D. J. (Eds.), *Everything Flows*. Oxford: Oxford University Press.

Favareau, D. (2021) Facing Up to the Hard Problem of Biosemiotics A commentary on Terrence Deacon's 'How molecules became signs'. *Biosemiotics* 14(3): 603–615.

Fodor, J. (2008) *LOT2: The language of thought revisited*. Oxford: Oxford University Press.

Frazier, P. A., Jamone, L., Althoefer, K., and Calvo, P. (2020) Plant Bioinspired Ecological Robotics. *Frontiers in Robotics and Ai* 7: 79.

Geurts, B. (2018) Convention and common ground. *Mind & Language* 33(2): 115–129.

Gibson, J. J. (2014) *The ecological approach to visual perception*. New York: Psychology 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. (2011) Distributed Cognition in the Lab. *Science* 333(6039): 159–160.

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.

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

Goldstone, R. L., Jones, A., and Roberts, M. E. (2006) Group path formation. *IEEE Transactions on Systems, Man, and Cybernetics Part A:Systems and Humans* 36(3): 611–620.

Goldstone, R. L., and Roberts, M. E. (2006) Self-organized trail systems in groups of humans. *Complexity* 11(6): 43–50.

Goldstone, R. L., Roberts, M. E., and Gureckis, T. M. (2008) Emergent processes in group behavior. *Current Directions in Psychological Science* 17(1): 10–15.

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.

Hanczyc, M. M. (2011) Metabolism and motility in prebiotic structures. *PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B-BIOLOGICAL SCIENCES* 366(1580): 2885–2893.

Hanczyc, M. M., and Ikegami, T. (2010) Chemical Basis for Minimal Cognition. *ARTIFICIAL LIFE* 16(3): 233–243.

Heras-Escribano, M., and De Jesus, P. (2018) Biosemiotics, the Extended Synthesis, and Ecological Information: Making Sense of the Organism-Environment Relation at the Cognitive Level. *Biosemiotics* 11(2): 245–262.

Heras-Escribano, M., Noble, J., and de Pinedo, M. (2015) Enactivism, action and normativity: a Wittgensteinian analysis. *Adaptive Behavior* 23(1): 20–33.

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

Heylighen, F. (2016b) Stigmergy as a universal coordination mechanism II: Varieties and evolution. *Cognitive Systems Research* 38: 50–59.

Hoffmeyer, J. (2007) *Biosemiotics: An examination into the signs of life and the life of signs*. Chicago: University of Scranton Press.

Hoffmeyer, J. (2020) The swarming body. In Rauch, I. and Carr, F. (Eds.), *Semiotics around the World: Synthesis in Diversity: Proceedings of the Fifth Congress of the International Association for Semiotic Studies, Berkeley 1994* (Reprint 2020., Vol. 2). De Gruyter: De Gruyter Mouton.

Hollan, J., Hutchins, E., and Kirsh, D. (2000) Distributed cognition: Towards a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7(2): 174.

Holland, O., and Melhuish, C. (1999) Stigmergy, Self-Organization, and Sorting in Collective Robotics. *Artificial Life* 5(2): 173–202.

Holler, S., and Hanczyc, M. M. (2019) Droplet-based synthetic biology: chemotaxis and interface with biology. In Fellermann, H., Bacardit, J., GoniMoreno, A., and Fuchslin, R. (Eds.), *ALIFE 2019: THE 2019 CONFERENCE ON ARTIFICIAL LIFE*. Presented at the Conference on Artificial Life (ALIFE) - How Can Artificial Life Help Solve Societal Challenges? Cambridge: Mit Press.

Huang, H., Ren, C., and Jin, S. (2008) ‘Signs’ is the Sign: Towards a Unified View of Stigmergy. In *2008 International Conference on Information and Automation, Vols 1-4*. New York: Ieee.

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. (2017) *Evolving enactivism: basic minds meet content*. Cambridge, Mass: M.I.T. Press.

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

Keupp, S., Behne, T., and Rakoczy, H. (2013) Why do children overimitate? Normativity is crucial. *Journal of Experimental Child Psychology* 116(2): 392–406.

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.

Kirchhoff, M., and Kiverstein, J. (2020) Attuning to the World: The Diachronic Constitution of the Extended Conscious Mind. *Frontiers in Psychology* 11: 1966.

Kirchhoff, M., and Kiverstein, J. (2023) Dissolving the Causal-Constitution Fallacy: Diachronic Constitution and the Metaphysics of Extended Cognition. In Caspar, M. O. and Artese, G. F. (Eds.), *Situated Cognition Research: Methodological Questions*. Cham, Switzerland: Springer Nature Switzerland AG.

Kirsh, D. (1999) Distributed cognition, coordination and environment design. *Proceedings of the European congerence on cognitive sience* : 1–11.

Kirsh, D. (2006) Distributed cognition: A methodological note. *Pragmatics and Cognition* 14(2): 249–262.

Kiverstein, J., and Rietveld, E. (2021) Skilled we-intentionality: Situating joint action in the living environment. *Open Research Europe*.

Lee, D. N. (1998) Guiding movement by coupling taus. *Ecological Psychology* 10(3–4): 221–250.

Lee, D. N., Georgopoulos, A. P., Clark, M. J. O., Craig, C. M., and Port, N. L. (2001) Guiding contact by coupling the taus of gaps. *Experimental Brain Research* 139(2): 151–159.

Levin, M., and Dennett, D. C. (2020) Cognition all the way down. *Aeon* (13 October 2020).

Lewis, T. G. (2013) Cognitive stigmergy: A study of emergence in small-group social networks. *Cognitive Systems Research* 21: 7–21.

Lewis, T. G., and Marsh, L. (2016) Human stigmergy: Theoretical developments and new applications. *Cognitive Systems Research* 38: 1–3.

Lyon, P. (2006) The biogenic approach to cognition. *Cognitive Processing* 7(1): 11–29.

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

Lyon, P., and Keijzer, F. A. (2007) The human stain. In Wallace, B., Ross, A., Davies, J., and Anderson, T. (Eds.), *The Mind, the Body, and the World: Psychology after Cognitivism*. Exeter: Imprint Academic.

Lyon, P., Keijzer, F., Arendt, D., and Levin, M. (2021) Reframing cognition: getting down to biological basics. *Philosophical Transactions of the Royal Society B-Biological Sciences* 376(1820): 20190750.

MacDonald, G., and Papineau, D. eds. (2006) *Teleosemantics*. Oxford: Oxford University Press.

Marsh, L., and Onof, C. (2008) Stigmergic epistemology, stigmergic cognition. *Cognitive Systems Research* 9(1–2): 136–149.

Menary, R. (2007) *Cognitive integration: mind and cognition unbounded*. Australia, Australia/Oceania: Research Online.

Menary, R. (2018) Cognitive integration: how culture transforms us and extends our cognitive capabilities. In Newen, A., de Bruin, L., and Gallagher, S. (Eds.), *The Oxford Handbook of 4E Cognition*. Oxford: Oxford University Press.

Millikan, R. G. (1984) *Language, thought and other biological categories: new foundations for realism*. Cambridge, Mass. ; London: MIT Press.

Millikan, R. G. (1989) In Defense of Proper Functions. *Philosophy of Science* 56(2): 288–302.

Millikan, R. G. (1996) Pushmi-pullyu representations. In May, L., Friedman, M., and Clark, A. (Eds.), *Mind and Morals: Essays on Cognitive Science and Ethics*. Cambridge, Mass ; London: MIT Press.

Nakagaki, T., Kobayashi, R., Nishiura, Y., and Ueda, T. (2004) Obtaining multiple separate food sources:: behavioural intelligence in the *Physarum plasmodium*. *PROCEEDINGS OF THE ROYAL SOCIETY B-BIOLOGICAL SCIENCES* 271(1554): 2305–2310.

Nakagaki, T., Yamada, H., and Tóth, A. (2000) Maze-solving by an amoeboid organism. *NATURE* 407(6803): 470–470.

Neander, K. (1995) Misrepresenting and Malfunctioning. *Philosophical Studies* 79(2): 109–141.

Neander, K. (2018) *A mark of the mental : in defense of informational teleosemantics*. The MIT Press.

Nieto-Gomez, R. (2016) Stigmergy at the edge: Adversarial stigmergy in the war on drugs. *Cognitive Systems Research* 38: 31–40.

Palermos, S. O. (2014) Loops, constitution, and cognitive extension. *Cognitive Systems Research* 27: 25–41.

Papineau, D. (1984) Representation and Explanation. *Philosophy of Science* 51(4): 550–572.

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.

Peirce, C. S., Hartshorne, C., and Weiss, P. (1960) *Collected papers of Charles Sanders Peirce. Vols.1 and 2: Principles of philosophy, and Elements of logic*. Cambridge [Mass.]: Belknap Press of Harvard University Press.

Putnam, H. (1995) *Renewing Philosophy*. (1st ed.). Cambridge: Harvard University Press doi:10.2307/j.ctv260718z.

Rakoczy, H., Wameken, F., and Tomasello, M. (2008) The sources of normativity: Young children’s awareness of the normative structure of games. *Developmental Psychology* 44(3): 875–881.

Reid, C. R. (2023) Thoughts from the forest floor: a review of cognition in the slime mould *Physarum polycephalum*. *ANIMAL COGNITION*. doi:10.1007/s10071-023-01782-1.

Reid, C. R., and Latty, T. (2016) Collective behaviour and swarm intelligence in slime moulds. *Fems Microbiology Reviews* 40(6): 798–806.

Reid, C. R., Latty, T., Dussutour, A., and Beekman, M. (2012) Slime mold uses an externalized spatial ‘memory’ to navigate in complex environments. *Proceedings of the National Academy of Sciences of the United States of America* 109(43): 17490–17494.

Ricci, A., Omicini, A., Viroli, M., Gardelli, L., and Oliva, E. (2007) Cognitive stigmergy: Towards a framework based on agents and artifacts. In Weyns, D., Parunak, H. V. D., and Michel, F. (Eds.), *Environments for Multi-Agent Systems Iii* (Vol. 4389). Berlin: Springer-Verlag Berlin.

Ritchie, K. (2015) The Metaphysics of Social Groups. *Philosophy Compass* 10(5): 310–321.

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

Ritchie, K. (2020b) Social Structures and the Ontology of Social Groups. *Philosophy and Phenomenological Research* 100(2): 402–424.

Roberts, M. E., and Goldstone, R. L. (2011) Adaptive Group Coordination and Role Differentiation. *Plos One* 6(7): e22377.

Schmidt, M. F. H., Rakoczy, H., and Tomasello, M. (2011) Young children attribute normativity to novel actions without pedagogy or normative language. *Developmental Science* 14(3): 530–539.

Secretan, J. (2013) Stigmergic dimensions of online creative interaction. *Cognitive Systems Research* 21: 65–74.

Sharkey, A. J. C. (2007) Swarm robotics and minimalism. *Connection Science* 19(3): 245–260.

Sims, M., and Kiverstein, J. (2022) Externalized memory in slime mould and the extended (non-neuronal) mind. *Cognitive Systems Research* 73: 26–35.

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

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

Sims, R. (2023b) Boundary objects, trading zones, and stigmergy: the social and the cognitive in science. *Synthese* 202(4): 117.

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

Singh, G., Kumar, N., and Verma, A. K. (2012) Ant colony algorithms in MANETs: A review. *Journal of Network and Computer Applications* 35(6): 1964–1972.

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

Smith-Ferguson, J., and Beekman, M. (2020) Who needs a brain? Slime moulds, behavioural ecology and minimal cognition. *Adaptive Behavior* 28(6): 465–478.

Smith-Ferguson, J., Reid, C. R., Latty, T., and Beekman, M. (2017) Hansel, Gretel and the slime mould-how an external spatial memory aids navigation in complex environments. *Journal of Physics D-Applied Physics* 50(41): 414003.

Spezzano, G. ed. (2019) *Swarm robotics*. Basel: MDPI.

Susi, T. (2016) Social cognition, artefacts, and stigmergy revisited: Concepts of coordination. *Cognitive Systems Research* 38: 41–49.

Susi, T., and Ziemke, T. (2001) Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. *Cognitive Systems Research* 2(4): 273–290.

Tennie, C., Call, J., and Tomasello, M. (2009) Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B-Biological Sciences* 364(1528): 2405–2415.

Theiner, G., Allen, C., and Goldstone, R. L. (2010) Recognizing group cognition. *Cognitive Systems Research* 11(4): 378–395.

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

Tsankova, D. D., and Georgieva, V. S. (2004) From local actions to global tasks: simulation of stigmergy based foraging behavior. In *2004 2nd International IEEE Conference on 'Intelligent Systems'. Proceedings (IEEE Cat. No.04EX791)* (Vol. 1). Presented at the 2004 2nd International IEEE Conference on 'Intelligent Systems'. Proceedings (IEEE Cat. No.04EX791).

Tummolini, L., and Castelfranchi, C. (2007) Trace signals: The meanings of stigmergy. In Weyns, D., Parunak, H. V. D., and Michel, F. (Eds.), *Environments for Multi-Agent Systems III* (Vol. 4389). Berlin: Springer-Verlag Berlin.

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

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

Turner, J. S. (2016) Swarm Cognition and Swarm Construction: Lessons from a Social Insect Master Builder. In Portugali, J. and Stolk, E. (Eds.), *Complexity, Cognition, Urban Planning and Design*. Dordrecht: Springer doi:10.1007/978-3-319-32653-5_7.

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.

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.

Walmsley, L. D. (2020) Lessons from a virtual slime: marginal mechanisms, minimal cognition and radical enactivism. *Adaptive Behavior* 28(6): 453–464.

Wexler, B. E. (2006) *Brain and culture: neurobiology, ideology, and social change*. Cambridge, Mass: MIT Press.

Wright, G. von (1963) *Norm and action: a logical enquiry*. London: Routledge and Kegan Paul.

Zamfirescu, C.-B., Candea, C., and Radu, C. (2014) A stigmergic approach for social interaction design in collaboration engineering. *Neurocomputing* 146: 151–163.