

Computational theories of group behavior

Bryce Huebner and Joseph Jebari¹

Repeated interactions between animals can generate stable patterns of collective behavior (Couzin 2009). For example, colonies of eusocial insects can make collective decisions about where to build their nests, and where to forage (Reid et al. 2015; Sasaki et al, 2013; Seeley 1995; 2010). And people working together develop the skills that are required to fly commercial airplanes, and to navigate modern Naval vessels (Hutchins 1995a; 1995b). In each of these cases, and in many others as well, computational models have provided useful insights about the nature and structure of group-level behavior; and some philosophers have suggested that such models provide a plausible foundation for thinking about collective mentality. But while group structure often impacts individual computations, and shapes individual mentality, it is substantially less clear whether these group-level computations ever yield forms of group-mindedness. By looking to coordination dynamics, we can begin to understand how the strength of informational relations between the components of distributed systems can stabilize collective behavior (Anderson, Richardson, & Chemero 2012). But it is more difficult to say when such relations yield collective mentality.

In addressing the issue of collective mentality, some philosophers have adopted a dynamical approach. They have focused on patterns of self-organization, and argued that collective mentality requires neither collective computation nor collective mental representation (Palermos 2016). Others have focused on the ways that group members acquire, store, transmit, and manipulate information as they perform collaborative tasks; they hold that collective mentality emerges as a result of informational transactions, which can occur even in the absence of collective mental representations (Theiner 2013; 2014). And still others have suggested that we should only posit collective mentality “where no subsystem is capable of producing an authoritative representation and where the representations of multiple subsystems can be coordinated and integrated to yield flexible, goal-directed system-level behavior” (Huebner 2014, 14). Stepping back from these disagreements, we can begin to see substantial overlap between ‘pure’ cases of dynamic coordination and more complex forms of coordination that depend on mental representations and informational relations (cf., Dale et al 2013). And from this perspective, it becomes possible to discern cases that don’t fit nicely within any of these perspectives, but which seem to play a role in biologically significant forms of cooperative behavior. Our aim in this paper is to clarify the variety of different ways in which individual and collective computations unfold in cooperative

¹ The writing, research, and revisions of this paper were fully collaborative; both authors approved the manuscript for submission.

contexts. While we offer some reflections along the way, our aim is not to establish where collective mentality occurs; instead, we hope to clarify the variety of different ways in which a computational approach can be used to understand collective behavior, and the possibility of collective mentality.

Throughout the paper, we treat dynamic forms of coordination as physical computations, which are best modeled using differential equations (Piccinini & Scarantino 2011). This is consistent with the rapidly expanding empirical literature on collective behavior, which often appeals to dynamical computations without specifying what is represented—or whether anything is. We believe that many of these collective computations are likely to occur in the absence of group-level representations (cf., Piccinini 2008). And for this reason, we adopt a mechanistic account of computation that requires: 1) a functional mechanism that processes variables that can change states, 2) in accordance with rules that map inputs to outputs, 3) in ways that are sensitive to the properties of these variables and to differences between different portions of them (Piccinini & Bahar 2013, 458; Piccinini & Scarantino 2011, 10). By paying attention to the precise nature of the computations that have been posited to explain different kinds of group behavior, we hope to show that there is a continuum between minimal forms of coordination—such as synchronous flashing in fireflies—and maximal forms of coordination—such as collective decision-making on the basis of distributed sources of information.

1. Individual computations

Let's begin with the form of dynamic updating that allows some colonies of fireflies to synchronize their flashing behavior. Individual fireflies flash to attract mates. But in one colony in the Great Smoky Mountains, stable patterns of synchronization reliably emerge for two weeks each summer. Each insect functions as an intrinsic oscillator that tracks the distribution and frequency of nearby flashes, and adjusts its rate and frequency of flashing against the observed value (Ramírez-Ávila et al 2003, 255). Individuals flash within a fixed period (965 ± 90 ms); but if another firefly flashes within the first 840 ms of a cycle, this will inhibit the upcoming flash, and cause a one second delay before the next one; and if an insect flashes within the last 160 ms of this period, the next flash will occur normally, and subsequent flashes will advance to synchronize with observed flashes (Buck et al 1981). This pattern of individual updating yields coordinated collective behavior. At the beginning of the night, flashing is chaotic; but order emerges as each firefly tracks the locally observable firing patterns, and uses this information to update its own behavior (Mirollo & Strogatz 1990). As the night goes on, one or two individuals will produce one or two flashes, triggering synchronized flashing in every member of a colony. In most cases, a “burst of five to eight species-specific

flashes is produced, and then group flashing abruptly ends” (Moiseff and Copeland 1994, 403); but after a brief period of inactivity the process repeats.

We contend that this is a relatively pure case of dynamical updating, which occurs in the absence of mental representations. Individual fireflies act as functional mechanisms (intrinsic oscillators), which track nearby flashes, and mechanically adjust their behavior in ways that yield stable patterns of coordinated behavior. Since these insects are functioning as intrinsic oscillators, the behavior of nearby insects will become dynamically coupled. Finally, the emergence of synchrony yields a transition from a collection of independent and randomly flashing individuals, to a well-ordered system with stable collective dynamics.

Significantly, this means that individual behavior sustains group-level coordination because individuals update their behavior in light of the behavior of other group members. But while this allows groups of fireflies to function as coordinated systems, the resulting patterns of emergent behavior play no further computational role in the unfolding of individual or collective behavior. None of the individuals tracks the emergent state of the system; instead, each insect changes its behavior in response to flashes that occur nearby. Moreover, the group doesn’t track any changes in internal or external variables, and it doesn’t update its state in ways that are sensitive to the properties of such variables, or to differences in their portions. Nonetheless, coordinated behavior bubbles-up through patterns of individual behavior, yielding interesting group-level stabilities. We claim that this kind of dynamic, group-level behavior is sustained by individual computations that occur in the absence of representations, and that only the most minimal form of group-level computations is being carried out. Put differently, synchronously flashing fireflies reveal the power of simple tracking mechanisms for creating and sustaining stable patterns of collective behavior.

This point is more significant than it might initially seem, as humans and other animals often “‘dance’ like fireflies” (Schmidt & Richardson 2008, 287), taking perceptually available information as input to behavioral oscillatory systems that facilitate behavioral entrainment. A well-known dynamical model (The HKB-model) explains how such ordered states emerge, why they stabilize, and why shifts occur between different coordinative states (Kelso 2008). For example, people can wag their fingers in parallel at low frequencies; but increasing the speed of this movement triggers a switch to an anti-parallel pattern of movement. Something similar appears to hold across a wide range of interpersonal phenomena, where each person’s behavior constrains the activity of the other, and each person allows their activity to be so constrained (Walton, Richardson & Chemero 2014, 12). Through this process, people draw one another into robust patterns of coordination. For example, in the context of improvisational jazz, musicians are “constrained by the sonic and kinesthetic results of the activities of the other improvisers”, and their coordinated interactions yield a musical system that is constituted by the

improvisers and their ongoing patterns of interaction (Walton, Richardson & Chemero 2014, 19).

There is good reason to believe that dynamical updating of this sort is “a universal self-organizing strategy that occurs at multiple scales of nature” (Schmidt & Richardson 2008, 285). And we contend that it is driven by a minimal form of collective computation: rhythmic oscillations become synchronized through dynamic coupling between individuals, yielding group-level regularities that accord with the HKB-model. Across many cases, individuals can reasonably be seen as components in computational systems, which track changes in each other’s states, and update their behavior in accordance with behavioral rules that are sensitive to such changes. But many kinds of mutual alignment require more robust forms of information processing, which go beyond oscillatory resonance. Colonies of insects, schools of fish, flocks of birds, and herds of mammals often move together as groups, as individuals adjust their behavior in light of the observed behavior of nearby organisms. And behavioral alignment emerges as individuals adjust their behavior in light of “locally acquired cues such as the positions, motion, or change in motion, of others” (Couzin 2009, 36). And as we argue in the next section, this is true even where nothing is represented by the group.

2. Group-level information processing

Individuals within a school, a flock, or a herd typically adjust their direction of travel to avoid isolation, and these adjustments reliably scale to yield stable forms of collective behavior. But in some cases, animals also exploit the information about unobserved resources or predators that is embodied in the movement decisions of nearby neighbors, treating observed directional changes as evidence of a biological salient object or event (Kao et al. 2014). Where the location of resources and predators is known to a subset of a population, naïve individuals who track the movement of nearby neighbors, and treat their behavior as information about inaccessible features of the local environment, expand their ‘effective range’ of perception (Couzin 2007, 715). And where animals adjust their willingness to commit to a course of action in light of the number of committed animals, this can generate cohesive forms of group-level behavior that reduce the cost of seeking new information for all individuals (Hein et al 2015; Sumpter et al 2008, 1775).

Understanding these kinds of phenomena doesn’t require adding much to the form of dynamical updating that we discussed in the previous section. But particular forms of information processing are important to the kind of computational phenomena that occurs in these kinds of cases. Individuals adjust their direction of travel, using simple behavioral rules that map inputs to outputs, in ways that are sensitive to the movements of others and to differences in the

distribution of such movements. But just as importantly, they treat observed behavior as a source of natural information that provides evidence about the presence of resources and predators. Consequently, these behavioral adjustments often cascade through a group, allowing each animal to update its own decisions in ways that are sensitive to the group's overall informational state. As animals move toward a foraging location, or away from a predator, they provide a signal that can be interpreted by others as behaviorally relevant. And as individual animals change their state in light of this information, using simple interaction rules to map the distribution and value of observed movements onto behavioral outputs, their behavior constitutes a signal for further observation. To understand these kinds of collective behavior, it is thus necessary to examine both the behavioral rules that are employed by individuals, as well as the sensitivity to the information encoded in observed behavior.

Consider the way that colonies of Olive baboons make decisions about where and when to forage. These baboons are “most likely to follow when there are many initiators with high agreement. However, when agreement is low, having more concurrent initiators decreases the likelihood that a baboon will follow anyone” (Strandburg-Peshkin et al. 2015, 1361). Within colonies, group-level cascades of information help individuals make decisions that will guide them toward preferable foraging locations, while reducing the cost of leaving the group to seek new information. The decision an individual makes thus depends on the decisions that others are making; and since all individuals adjust their decisions in light of the changes in group-level properties, these shifts can be seen as inputs into a parallel processing algorithm that reliably moves the group toward consensus (Sumpter & Pratt 2009, 743). Here, as in many different species, a simple “quorum-response” rule is used to effectively integrate multiple sources of information about a course of action; uninformed and misinformed individuals then correct their behavior in light of group-wide behaviors; and since individuals raise their decision thresholds as group size increases, being in a larger group improves the accuracy of such decisions (Sumpter et al 2008, 1776). Where this pattern of tracking and updating occurs in the context of a relatively stable group, it yields a group-based computational process, where small changes in individual states yield large changes in group-level behavior, and where cascades of biologically significant information—embodied in the prevalence and distribution of movement decisions—lead individual preferences to converge on the best available option, in ways that would otherwise be impossible (Sumpter & Pratt 2009).

Similarly, consider the case of predator avoidance in Golden Shiners, which emerge as individual fish adjust their speed and direction to align with their nearest-neighbors (Berdahl et al. 2013). Golden Shiners are skittish, and they often display spontaneous avoidance behavior in the absence of predators—but crucially, the effect of this behavior is dampened in schools (Rosenthal et al. 2015). Within a

school, fish respond to the avoidance displays of any fish they can see. And since few fish will see the spontaneous avoidance behavior of a single fish, skittishness only triggers a limited response in the school. By contrast, the presence of a predator leads multiple fish on the same edge of a school to produce simultaneous avoidance signals. This yields a threat signal that cascades through the school, as each fish will observe several simultaneous avoidance displays. Where this occurs, the information that is embodied in predator avoidance triggers group-wide avoidance behavior, as the information cascades through the entire school. The important thing to notice, here, is that informational signals are amplified by parallel patterns of behavioral resonance. And information processing cascades occur when “the growing number of adherents to an option increases its attractiveness to undecided animals” (Sumpter & Pratt 2009, 745). Likewise, informational signals are dampened by behavioral dissonance, and negative feedback prevents uninformed individuals from guiding group-level behavior.

In each of these cases, adjustments to individual behavior unfold dynamically in real-time, and they are sensitive to the value of food resources, and to the presence or absence of predators. But as with the cases we addressed above, the most salient computations are those that are carried out by individuals. Each animal adjusts its behavior in light of its own observations, with patterns of mutual adjustment yielding a computational process that governs group-level behavior. As information cascades through a group, this generates successful forms of individual and collective behavior—without individuals needing to know that a collective computation is occurring. But here, we think it is substantially less clear whether these group-based computations yield collective mentality. Perhaps colonies of baboons and schools of fish should be treated as cognitive systems, as they are carrying out classically cognitive tasks (i.e., avoiding predators, and finding foraging patches), and they are doing so in a way that depends on integrated networks of computational processes. Or perhaps the individuals are carrying out these cognitive tasks, by tracking the distributed computation that is being carried out by the group; on this latter approach, each animal is looking for the best foraging option, and their success depends on using the information that other animals have broadcast. Of course, acting in the context of a group is a good way to succeed in this task—and it may be the only biologically feasible way for such organisms to track preferable food sources and dangerous predators.

In general, we prefer this second approach, which focuses on the ways that individuals use the information embodied in the structure of the group. There are likely to be many cases where aggregate success depends on the independence of decision-makers (Surowiecki 2004). And in such cases, tendencies toward local control should be preserved, even in species where collective action is common. More importantly, where consensus is reached too rapidly, informational cascades can often generate forms of groupthink, and “when individuals sense too much of

the group, the result is a filtering of the local influences and an averaged (compromised) collective response” (Leonard et al. 2012: 232). Consequently, it is biologically plausible that individual-level computations will tend to retain a high degree of salience in many cases, as they are necessary to prevent group-level processing from repeatedly leading to sub-optimal decisions (Torney et al. 2015). But there are cases where the properties of a network become just as important as individual computations, and in the remainder of this chapter, we turn to forms of collective behavior where this seems to be the case.

3. Network structure and informational processing

Simple responses to the position and motion of others produce many forms of group-level behavior. But in well-organized groups, network structure, as well as flows of information, and relations of independence and interdependence can play critical roles in individual and group-level computations (Derex & Boyd 2016). In some groups, individuals learn to track the network structure of the group they belong to; and this can open up novel individual strategies, which in turn transform the structure of the group. This kind of computational phenomena is complex, but it is pervasive. And to see what it amounts to, it will help to work through a particular case in detail. Here, we consider the way that pigtailed macaques track and maintain stable dominance hierarchies.

A smart animal should only compete for resources if their chance of winning is high—otherwise they should acquiesce. However, in competitive colonies, such decisions are often made under conditions of uncertainty: success is subject to numerous environmental factors, and immediate past successes or failures are often unreliable predictors of future outcomes. To solve this problem, pigtailed macaques have arrived upon an ingenious strategy for collectively determining the “temporally stable factors that predict who will be the winner on average” (Flack 2012, p.1804). Over the course of repeated dyadic and polyadic conflicts, macaques generate information about their relative fighting ability, as the less adept fighters are sure to lose more often over multiple bouts. After losing numerous fights, a monkey will begin signaling submission by baring its teeth toward the previous winner. This shift to subordination signaling consolidates the information generated by the fights, and communicates a general willingness to acquiesce, functioning as a kind of contract that allows the macaques to interact without fighting. So through repeated interactions, these monkeys are able to compute a measure of relative fighting ability and encode it in a subordination signal; and when these signals are expressed across multiple overlapping pairs of monkeys, they generate a subordination signaling network, whose overall structure encodes the dominance hierarchy of the group. Thus, while noisy signals are produced by the result of any single competition, the integration of these signals

into a larger network of mutually adjusted responses produces a stable hierarchy organized according to a robust measure of relative fighting ability.

What is the nature of the computation that occurs in this case? As with the golden shiners, the interactions between colony members have the function of producing a robust measure of a biologically salient value: relative fighting ability. Moreover, the distributed nature of this process allows for the emergence of a stable collective order. But unlike the cases we have examined thus far, this subordination signaling network allows each macaque to locate its position within the larger group structure, by determining its own rank within the dominance hierarchy (Flack & Krakauer 2006). Each macaque can estimate its social power, as well as its position in the hierarchy, by tracking the number of monkeys that signal subordination to them, and the frequency with which they do so. And by integrating over this information, each macaque can obtain a reliable estimate of how much power it has within the group. Thus, the subordination signaling network gives the individual macaques access to the output of multiple parallel competitions—they not only acquire a position in the hierarchy, they also obtain knowledge of that position.

To summarize, individual monkeys continually generate information about relative fighting ability, and encode that information through their patterns of signaling behavior within the subordination network. By tracking their history of agonistic interactions, they can determine the likelihood they will win in a fight against a given conspecific and then use the resulting representations to decide when to signal subordination, and to whom. Since multiple monkeys are doing this in parallel, this yields multiple interaction networks, with individual monkeys as nodes, and fights and subordination signals as edges (i.e. the connections between the nodes). Within these networks, altercations constitute physical computations, which collectively determine the dominance relations within the colony. Once stabilized, these dominance relations are encoded as constitutive features of the subordination signaling network; and in this context, individual monkeys can use the information encoded in the subordination signaling network to infer their own location within the hierarchy (Flack & Krakauer 2006, E93). This is possible because each macaque can average over multiple signals to accurately measure its position in the colony and track its social power.

In this way, these monkeys become able to “see” the output of the computations that are occurring across the colony and use this information to guide colony-relevant actions, making it possible to feed the output of the collective process back into the dynamics of the process. As Jessica Flack and her colleagues (2006) argue, the computations carried out within this signaling network guide forms of individual decision-making that make new and valuable forms of social decision-making possible. Specifically, the ability to determine one’s own position in the overall hierarchy increases the stability and cohesiveness of the group by allowing for a form of policing, where dominant macaques monitor the behavior of

their group-mates, and intervene in conflicts that would destabilize the group. Such interventions are generally risky, making this kind of policing biologically unlikely. However, since these monkeys can reliably evaluate their own power within a group, dominant monkeys can engage in policing behavior with minimal risk of being harmed; for after all, the dominance hierarchy ensures that other monkeys will acquiesce to them. This form of policing stabilizes the network structure that allows policing to occur, and this means that the stability of the group structure is both the cause and the effect of effective policing. As a result, colonies where policing occurs are larger, and they have higher rates of partner diversity, as well as increased possibilities for forms of socially contagion and cooperation; by contrast, colonies without policing have high rates of conflict, yielding less integrated groups, with less stable social interactions (Flack et al., 2006).

Similar forms of computational phenomena are ubiquitous within human groups; and the expressive power of language allows us to exchange signals about an indefinitely large class of topics, far beyond considerations of social dominance. Humans can identify useful environmental stabilities that can be signaled, and they can establish robust network structures that are capable of processing collective information about these stabilities (Barkoczi & Galesic 2016). This yields a kind of social learning that can contribute to individual- and group-level performance by providing a way to diffuse successful strategies through a communication network (e.g., Apicella et al 2012; Hill et al 2011; Rand et al 2011). However, in complex adaptive environments, communities of social learners risk settling on locally optimal strategies, while being unable to successfully explore superior strategies; and whether the most successful strategy can be identified and spread across a network largely depends on the structure of the network and the social learning rule used by the agents (Derex & Boyd 2016).

This brings us to an intriguing fact about human sociality: Where multiple agents execute a social learning routine in parallel, it is possible for a network of interacting agents to efficiently search for the best strategy, in ways that go beyond what would be possible for a lone individual (Derex & Boyd 2016). This happens as individual agents repeatedly sample the group-level process to identify better strategies, and use these strategies as the basis for further individual search. When superior strategies are found individually, these feed back into the group-level computational process, ratcheting up the overall performance of the network. This is possible because the agents are embedded in a signaling network that allows other agents to communicate both the strategy they are using and its value. A focal agent can then integrate over the signals they receive, to infer the relative value of their own strategy, and this in turn can guide decision-making. Although the integration algorithm is different from the one that undergirds macaque policing, it yields a similar type of phenomena: an agent uses their position in a signaling network to infer their own position with respect to that network, and uses this information to

guide their own individual decisions. The outcome of these individual decisions then feeds back into the signaling network, altering the dynamics of the collective computation. This process creates complex computational processes that are optimized by striking a balance between independence and interdependence; and this capacity provides a foundation for a more robust form of collective decision making, which we examine in the next section (List et al 2008).

4. Decisions

Identifying cases of group-level cognition is, at least in part, “a matter of determining how a cognitive system at a higher level can subsume cognitive systems at a lower level, and how the systems at multiple levels can strengthen rather than diminish one another” (Goldstone & Theiner 2017). Thus far, we’ve argued that dynamic forms of coupled processing allow animals to engage in flexible and adaptive forms of behavior; but do coupled systems ever “instantiate cognitive mechanisms in virtue of which a variety of systems perform important cognitive functions associated with flexible, adaptive, and intelligent behavior” (Goldstone & Theiner 2017)? In a recent paper, Rob Goldstone and Georg Theiner (2017) have argued that diffusion-based computations facilitate the accumulation of evidence, and generate rapid and accurate decisions; they also contend that the computational processes that allow individuals to make perceptual decisions under uncertainty are realized in the decision making of social insect colonies. In each case, interacting populations use competitive algorithms, arriving at decisions when one population exceeds an uncertainty threshold, which is adjusted in light of the salience of speed and accuracy (Marshall et al 2011). We think that Goldstone and Theiner are on the right track.

Colonies of ants and honeybees are often able to function as “parallel information-processing systems capable of intricate collective decision-making during essential tasks such as foraging, moving home or constructing a nest” (Couzin 2009: 39). In each case, collective decisions arise through a process that parallels the winner-take-all algorithms that are used to explain how the visual system categorizes objects (Riesenhuber & Poggio 1999), or by forms of diffusion-based processing that yield rapid and accurate perceptual decisions (Marshall et al 2011). By enhancing the activity of some computational units (here, ants or honeybees), while inhibiting or suppressing the activity of others, responses can be pooled in ways that will achieve an accurate representation of the biologically salient features of an object (here a foraging location or a nest site). Consequently, decisions are reached that are relevant to group-level behavior by way the friendly-competition between group members.

One of the clearest examples of this kind of phenomena occurs when a honeybee hive splits, scouts carry out a random search for a new nest site,

evaluating each alternative in terms of “cavity volume, entrance height, entrance area, entrance direction, entrance position relative to the cavity floor, and presence of combs from a previous colony” (Seeley & Buhrman 1999, 31). Few bees visit more than one site, but the colony settles on a decision by representing the variety of different options and selecting among them. Each option is evaluated in terms of its attributes (as listed above), and the group’s choice is typically optimal with respect to the value of each attribute for colony survival and reproduction (Seeley & Buhrman 1999). When scouts return to the hive, they dance in support of the site they visit, using a waggle dance that varies in intensity as a function of the quality of the site. Few scouts dance in support of more than one site; and “most bees that dance initially for a site other than the ultimately chosen site terminate their dancing for this site by ceasing their dancing altogether, not by switching their dancing to the chosen site” (Seeley & Buhrman 1999, 30). Some bees that dance for the ultimately chosen site stop before consensus is reached; but in general, scouts who find a preferable site tend to recruit others to inspect the same site; and the increasing levels of recruitment to a site further increase support for that site (List et al 2008; Seeley et al 2012). Decisions are thus made by “quorum”, with colonies settling on a preferred alternative as soon as there is sufficient support for it (Seeley & Visscher 2003). Importantly, their “independence in assessing the various sites’ quality and their interdependence through communication are both necessary and sufficient for the reliability of the bees’ decision process” (List et al 2008, 758).

A similar consensus-based algorithm is used by ants to compare multiple potential nest sites, which differ with respect to cavity volume, interior dimness, and entrance size (Pratt et al. 2002; Sasaki et al. 2013). But more intriguingly, when army ants build bridges to create a shortcut in a foraging trail, for example, they make adjustments to the length of the bridge in response to the flow of traffic across their bodies (Reid et al. 2015: 15114). As traffic decreases, ants abandon their position; and as traffic increases, ants are recruited to the bridge. But bridge expansion will often stop before the maximum foraging shortcut has been achieved. This phenomenon is interesting because no individual can represent the costs and benefits to the colony—yet the variations in recruitment underwrite a form of parallel information processing that is sensitive to “the diminishing returns of shortening the trail to avoid the cost of locking up an increasing number of workers in the structure” (Reid et al. 2015: 15116).

The important thing to notice about these two cases is that both group-level computations, and group-level representations, play a critical role in the production of collective behavior. In contrast to the cases that we have discussed above, the individuals are serving as nodes in an integrated computational network that solves a group-relevant task. Individual bees observe informationally-rich dances, and adjust their behavior in light of what they perceive; but *the colony* chooses the best nest site, for a large range of parameter conditions, by aggregating over individual

patterns of behavior. Likewise, individual ants join or abandon their position in a bridge, in a way that is sensitive to the number of ants currently using the bridge; but the colony determines the optimal bridge-length for successful hunting. While an account of the independent decisions of individual animals would be interesting, it is only by understanding the trade-off between independence and interdependence that we can understand what these colonies are doing, and why (cf., List et al. 2008)

There have been speculative extensions of these types of models to human decision-making. For example, John Dyer and colleagues (2008) has argued that the motion of crowds is likely to be driven by a computational process like his. But while there is a computation at the group-level according to his account, the relevant forms of motion-guidance aren't used by the group for any group-relevant ends. Bryce Huebner (2014, 69ff) has suggested that a stock market could display computational properties that were best modeled in terms of a competitive algorithm; but he worries that all of the relevant representations would be used for individual decision-making, not for any sort of collective decision-making. And Bernard Grofman and Scott Feld (1988) have argued that democratic group decisions could arise through a process of Condorcet aggregation (using a process of recruitment like the one suggested by List et al 2008). But here too, there is little evidence that successful democratic decision-making relies on these specific kinds of computations.

Against this backdrop, it's worth asking what other kinds of system-level computations might lead to the production of group-level behavior in human groups, and whether any of these might yield stable forms of collective mentality. Much of the existing research on distributed cognition has focused on the ways that representations are passed between small numbers of people to yield higher-level regularities. For example, research on collaborative retrieval (Harris et al 2011, 2013, 2014; Michaelian & Sutton 2013) and transactive memory (for reviews, see Ren & Argote 2011; Theiner 2013) has examined the ways that groups of two or three people broadcast and receive semantic information to one another. These projects often proceed at a relatively high level of abstraction; and as Mateo Colombo (2015) suggests, emerging research on hierarchical Bayesian algorithms may offer an interesting, biologically plausible model of collective learning and agency that can be used to flesh out these proposals. But at present, this suggestion remains quite speculative.

One way of moving forward on this speculative suggestion is to examine cases of top-top cognition. As Andreas Roepstorff and Chris Frith (2004) argue, we are able to communicate with one another (both linguistically and non-linguistically), in ways that allow us to draw one another into patterns of mutual alignment. Where this process is successful, it does seem to yield shared representations of a situation or an event, and there is some recent evidence that

these similarities are grounded in patterns of overlapping neural activity (Clark 2015, 286-287; Friston & Frith 2015; Gallotti, Fairhurst, & Frith 2017; Dikker et al 2017). Perhaps this is a way of linking the top-level structures of multiple interacting agents, and perhaps this allows for the possibility of treating people as computational nodes in an integrated computational network. At present, this remains a theoretically interesting possibility, which is only beginning to be demonstrated empirically—but it does recommend a fruitful path for future exploration, which we examine in the final section of this paper.

5. A Speculative Conclusion

It would be incredibly interesting if human forms of collective mentality were implemented by powerful forms of machine learning, such as hierarchical Bayesian algorithms. But which algorithms produce collective mentality, and whether they are the same kind across forms of biological cognition, are open empirical questions. We contend that mentality is likely to emerge wherever self-organizing systems achieve enough unity and stability to process information that is relevant to group-level ends, and to adjust group-level behavior in accordance with a group's representation of the world. This can be achieved in many different ways, and this is why we refrain from making claims about the class of algorithms and the kinds of representations that are necessary for collective mentality. We are unaware of cases where group-level cognitive capacities have been shown to depend on hierarchical Bayesian algorithms, but ants and honeybees do exhibit forms of collective mentality implemented by consensus algorithms and winner-take-all algorithms. And while it would be premature to attempt to pull all forms of collective mentality under the hierarchical Bayesian umbrella, there are interesting questions in this vicinity.

Collective behavior of various sorts is ubiquitous in human life. For as long as we have been human, we have been immersed in a integrated network of interaction, which is integral to the human phenotype, and to the possibility of human lifeways (cf., Jebari & Huebner in press); furthermore, there is good reason to believe that our ability to flexibly exploit the forms of collective computation that are carried out within this network structure are central to our capacity to sustain increasingly complex social systems and acquire increasingly complex understandings of the world (Dedeo 2013; Goldstone & Theiner 2017). One of the major attractions of the hierarchical Bayesian framework in computational neuroscience is that it promises to unify within a single framework the mosaic of algorithms that have been proposed to explain neural function. And one tantalizing possibility is that something similar could be achieved with respect to human social systems. As we noted in the previous section, the research investigating the nature of human collective computation has been largely exploratory, emphasizing the

vast array of different (and perhaps complementary) mechanisms that plausibly undergird the various dimensions of human sociality. We believe the question of unification should be taken up more seriously; it is possible that there is a universal algorithm characterizes the basic processing structure of all human social systems—and it may be a hierarchical Bayesian process, or something like it. If this were established, it would revolutionize sociology, political science, and public policy. And recent empirical results suggest a way of moving forward on this approach to collective computation.

We already know that human social systems preserve certain network properties across scales, and that this is a feature of human social organization (Derex & Boyd 2015, 2016; Salali et al 2016). Analyzing a range of data involving multiple different measures of social network structure, Wei-Xing Zhou and colleagues (2005) have found that social networks across the developed world observe a hierarchical fractal-like scaling ratio, with units at each level of analysis constituted by ~ 4 units from the level below. Specifically, they show that small groups composed of ~ 4 closely-associated individuals tend to associate with ~ 3 other groups of similar size, forming larger, more loosely associated groups. Larger groups then associate with ~ 3 similarly structured groups, etc. If they are right that various human groups conform to this scaling property, this would suggest that human social systems tend to form self-similar structures at multiple levels of organization.

This scaling property has been confirmed in the social organization of contemporary hunter-gatherers (Hamilton et al. 2007) as well as in online multiplayer games (Fuchs et al. 2014) indicating that it may be a universal feature of human sociality. Perhaps this kind of structure may have developed to deal with the demands of maintaining efficient resource and information exchange (Hamilton et al. 2007), and to preserve collective adaptability in the face of changing environments (Flack et al. 2013). If basic forms of human sociality are adapted to this structure, in a way that parallels the phenomena in macaques that we discussed above, then perhaps the ability to exploit the information that is encoded in the dynamics of social interaction may prove to be an essential feature of efficient information processing across human groups. If this is right, then hierarchical network structure may function analogously to a hierarchical neural network, with information acquired at the individual level being aggregated and refined as it spreads through the network. Put differently, something like the following might be the case: network structure may provide constraints on the flow of information through human groups of various sizes; patterns of informational exchange within these structures may then yield local sources of knowledge, which can be fed upward through the group (using an aggregation function that strikes a balance between independence and interdependence); finally the information encoded in these aggregated signals may feed downward into individual-level computations,

allowing individuals to locate themselves within larger structures, and to update their behavior in light of this information. This process would yield a bidirectional flow of information, which would allow groups of interacting individuals to search for the linked set of hypotheses that would make the most sense of the group's current situation.

If this analogy goes through, then collective computation in humans may reflect the implementation of the same type of hierarchical information processing scheme that is found in the brain. If so, it may be possible to develop a unified theory of human collective computation that is continuous with computational neuroscience. More research would be necessary to confirm this hypothesis; but no matter how the data turn out, we will gain a much clearer understanding of the computational structures that support patterns of stability, and patterns of variation in human sociality. Discovering a basic computational structure would have significant ramifications for the design of effective social interventions, and for the development of efficient social institutions. Likewise, if no basic computational structure can be found, understanding the diversity inherent in human collective computation will have equally important consequences for good social design.

Works Cited:

- Anderson, M. L., Richardson, M. J., & Chemero, A. (2012). Eroding the Boundaries of Cognition: Implications of Embodiment. *Topics in cognitive science*, 4(4), 717-730.
- Apicella, C. L., Marlowe, F. W., Fowler, J. H., & Christakis, N. A. (2012). Social networks and cooperation in hunter-gatherers. *Nature*, 481(7382), 497.
- Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. arXiv preprint arXiv:1606.00753
- Berdahl, A., C. Torney, C. Ioannou, J. Faria & I. Couzin (2013). Emergent sensing of complex environments by mobile animal groups. *Science*, 339(6119), 574-576.
- Buck, J., Buck, E., Case, J. F., & Hanson, F. E. (1981). Control of flashing in fireflies. *Journal of comparative physiology*, 144(3), 287-298.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Colombo, M. (2015). Review of B. Huebner, *Macro-cognition: A Theory of Distributed Minds and Collective Intentionality*. *Minds and Machines*, 25, 103-109.
- Couzin, I. (2007). Collective minds. *Nature*, 445 (7129), 715.
- Couzin, I. (2009). Collective cognition in animal groups. *Trends in cognitive science*, 13, 1, 36-43.

- Dale, R., Fusaroli, R., Duran, N., & Richardson, D. C. (2013). The self-organization of human interaction. *Psychology of learning and motivation*, 59, 43-95.
- DeDeo, S. (2013). Collective phenomena and non-finite state computation in a human
- Derex, M., & Boyd, R. (2015). The foundations of the human cultural niche. *Nature communications*, 6.
- Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences*, 113, 2982-2987.
- Dikker, S., Wan, L., Davidesco, I., Kaggen, L., Oostrik, M., McClintock, J., Rowland, J., Van Bavel, J.J., Ding, M., & Poeppel, D. (2017). Brain-to-brain synchrony tracks real-world dynamic group interactions in the classroom. *Current Biology*, 27, 1375-1380.
- Dyer, J. R., Ioannou, C. C., Morrell, L. J., Croft, D. P., Couzin, I. D., Waters, D. A., & Krause, J. (2008). Consensus decision making in human crowds. *Animal Behaviour*, 75(2), 461-470.
- Flack, J. C. (2012). Multiple time-scales and the developmental dynamics of social systems. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367(1597), 1802-1810.
- Flack, J. C., Erwin, D., Elliot, T., & Krakauer, D. C. (2013). Timescales, symmetry, and uncertainty reduction in the origins of hierarchy in biological systems. *Evolution cooperation and complexity*, 45-74.
- Flack, J. C., Girvan, M., De Waal, F. B., & Krakauer, D. C. (2006). Policing stabilizes construction of social niches in primates. *Nature*, 439(7075), 426-429.
- Flack, J. C., & Krakauer, D. C. (2006). Encoding power in communication networks. *The American Naturalist*, 168(3), E87-E102.
- Friston, K., & Frith, C. (2015). A duet for one. *Consciousness and cognition*, 36, 390-405.
- Fuchs, B., Sornette, D., & Thurner, S. (2014). Fractal multi-level organisation of human groups in a virtual world. *arXiv preprint arXiv:1403.3228*.
- Gallotti, M., Fairhurst, M. & Frith, C. (2017) "Alignment in Social Interactions", *Consciousness and Cognition*, 48 (2017): 253-261.
- Goldstone, R. & G. Theiner (2017). The Multiple, Interacting Levels of Cognitive Systems (MILCS) Perspective on Group Cognition. *Philosophical Psychology*, 30(3), 334-368.
- Grofman, B., & Feld, S. L. (1988). Rousseau's general will: a Condorcetian perspective. *American Political Science Review*, 82(02), 567-576.
- Hamilton, M. J., Milne, B. T., Walker, R. S., Burger, O., & Brown, J. H. (2007). The complex structure of hunter-gatherer social networks. *Proceedings of the Royal Society of London B: Biological Sciences*, 274(1622), 2195-2203.

- Harris, C. B., Keil, P. G., Sutton, J., Barnier, A. J., & McIlwain, D. J. (2011). We remember, we forget: Collaborative remembering in older couples. *Discourse Processes*, 48(4), 267-303.
- Harris, C. B., Barnier, A. J., & Sutton, J. (2013). Shared encoding and the costs and benefits of collaborative recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(1), 183.
- Harris, C. B., Barnier, A. J., Sutton, J., & Keil, P. G. (2014). Couples as socially distributed cognitive systems: Remembering in everyday social and material contexts. *Memory Studies*, 7(3), 285-297.
- Hein, A. M., Rosenthal, S. B., Hagstrom, G. I., Berdahl, A., Torney, C. J., & Couzin, I. D. (2015). The evolution of distributed sensing and collective computation in animal populations. *eLife*, 4, e10955.
- Hill, K. R., Walker, R. S., Božičević, M., Eder, J., Headland, T., Hewlett, B., Hurtado, A.M., Marlowe, F., Wiessner, P. & Wood, B. (2011). Co-residence patterns in hunter-gatherer societies show unique human social structure. *Science*, 331(6022), 1286-1289.
- Huebner, B. (2014). *Macrocognition: A theory of distributed minds and collective intentionality*. Oxford University Press.
- Hutchins, E. (1995a). *Cognition in the Wild*. Cambridge: MIT press.
- Hutchins, E. (1995b). How a cockpit remembers its speed. *Cognitive science*, 19(3), 265-288.
- Jebari, J. & Huebner, B. (in press). From objectivized morality to objective morality. Commentary on P. K. Stanford, "The difference between ice cream and nazis: Moral externalization and the evolution of human cooperation". *Behavioral and Brain Sciences*.
- Kao, A, N. Miller, C. Torney, A. Hartnett & I. Couzin (2014). Collective learning and optimal consensus decisions in social animal groups. *PLoS Computational Biology* 10, 8, e1003762.
- Kelso, J.A.S. (2008) Haken-Kelso-Bunz model. *Scholarpedia*, 3(10):1612.
- Leonard, N. E., Shen, T., Nabet, B., Scardovi, L., Couzin, I. D., & Levin, S. A. (2012). Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1), 227-232.
- List, C., Elsholtz, C., & Seeley, T. D. (2008). Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society B*, 364, 755–762.
- Marshall, J. A. R., Bogacz, R., Dornhaus, A., Plaque, R., Kovacs, T., & Franks, N. R. (2011). On optimal decision-making in brains and social insect colonies, *Journal of the Royal Society Interface*, 6, 1065-1074.

- Michaelian, K., & Sutton, J. (2013). Distributed cognition and memory research: History and current directions. *Review of philosophy and psychology*, 4(1), 1-24.
- Mirollo, R. E., & Strogatz, S. H. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6), 1645-1662.
- Moiseff, A., & Copeland, J. (1994). Mechanisms of synchrony in the North American firefly *Photinus carolinus* (Coleoptera: Lampyridae). *Journal of insect behavior*, 8(3), 395-407.
- Palermos, S. O. (2016). The Dynamics of Group Cognition. *Minds & Machines*
- Piccinini, G. (2008). Computation without representation. *Philosophical studies*, 137(2), 205-241.
- Piccinini, G., & Bahar, S. (2013). Neural computation and the computational theory of cognition. *Cognitive science*, 37(3), 453-488.
- Piccinini, G., & Scarantino, A. (2011). Information processing, computation, and cognition. *Journal of biological physics*, 37(1), 1-38.
- Pratt, S. C., Mallon, E. B., Sumpter, D. J., & Franks, N. R. (2002). Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *Leptothorax albipennis*. *Behavioral Ecology and Sociobiology*, 52(2), 117-127.
- Ramírez-Ávila, G., Guisset, J. L., & Deneubourg, J. L. (2003). Synchronization in light-controlled oscillators. *Physica D: Nonlinear Phenomena*, 182(3), 254-273.
- Rand, D. G., Arbesman, S., & Christakis, N. A. (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences*, 108(48), 19193-19198.
- Reid, C. R., Lutz, M. J., Powell, S., Kao, A. B., Couzin, I. D., & Garnier, S. (2015). Army ants dynamically adjust living bridges in response to a cost–benefit trade-off. *Proceedings of the National Academy of Sciences*, 112, 49, 15113-15118.
- Ren, Y., & Argote, L. (2011). Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences. *The Academy of management annals*, 5(1), 189-229.
- Riesenhuber, M. & Poggio, T. (1999) Hierarchical models of object recognition in cortex, *Nature Neuroscience*, 2, 11: 1019-1025.
- Roepstorff, A., & Frith, C. (2004). What's at the top in the top-down control of action? Script-sharing and 'top-top' control of action in cognitive experiments. *Psychological Research*, 68(2-3), 189-198.
- Rosenthal, S., C. Twomey, A. Hartnett, H. Wu & I. Couzin (2015). Revealing the hidden networks of interaction in mobile animal groups allows prediction

- of complex behavioral contagion. *Proceedings of the National Academy of Sciences*, 112, 15, 4690-4695.
- Salali, G. D., Chaudhary, N., Thompson, J., Grace, O. M., van der Burgt, X. M., Dyble, M., ... & Vinicius, L. (2016). Knowledge-Sharing Networks in Hunter-Gatherers and the Evolution of Cumulative Culture. *Current Biology*, 26(18), 2516-2521.
- Sasaki, T., Granovskiy, B., Mann, R. P., Sumpter, D. J., & Pratt, S. C. (2013). Ant colonies outperform individuals when a sensory discrimination task is difficult but not when it is easy. *Proceedings of the National Academy of Sciences*, 110, 34, 13769-13773.
- Schmidt, R. C., & Richardson, M. J. (2008). Dynamics of interpersonal coordination. In *Coordination: Neural, behavioral and social dynamics* (pp. 281-308). Springer Berlin Heidelberg.
- Seeley, T. (1995). *The wisdom of the hive*. Cambridge: Harvard University Press.
- Seeley, T. (2010). *Honeybee Democracy*. Princeton: Princeton University Press.
- Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D., & Crofoot, M. C. (2015). Shared decision-making drives collective movement in wild baboons. *Science*, 348(6241), 1358-1361.
- Seeley, T. & Buhrman, S. (1999) Nest-Site Selection in Honey Bees: How Well do Swarms Implement the 'Best-of-N' Decision Rule? *Behavioral Ecology and Sociobiology* 49: 416-27
- Seeley, T. & Visscher, P. (2003). Choosing a Home: How the Scouts in a Honey Bee Swarm Perceive the Completion of Their Group Decision Making,” *Behavioral Ecology and Sociobiology* 54: 511-20.
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064), 108-111.
- Strandburg-Peshkin, A., Farine, D. R., Couzin, I. D., & Crofoot, M. C. (2015). Shared decision-making drives collective movement in wild baboons. *Science*, 348(6241), 1358-1361.
- Surowiecki, J. (2005). *The wisdom of crowds*. Norwell, MA: Anchor Press.
- Sumpter, D. J., Krause, J., James, R., Couzin, I. D., & Ward, A. J. (2008). Consensus decision making by fish. *Current Biology*, 18(22), 1773-1777.
- Sumpter, D. J., & Pratt, S. C. (2009). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1518), 743-753.
- Theiner, G. (2013). Transactive memory systems: A mechanistic analysis of emergent group memory. *Review of Philosophy and Psychology*, 4(1), 65-89.
- Theiner, G. (2014). A Beginner's Guide to Group Minds. In *New Waves in Philosophy of Mind* London: Palgrave Macmillan UK:301-322.

- Theiner, G. (2017). Group-sized distributed cognitive systems. *The Routledge Handbook of Collective Intentionality*. New York: Routledge.
- Torney, C. J., Lorenzi, T., Couzin, I. D., & Levin, S. A. (2015). Social information use and the evolution of unresponsiveness in collective systems. *Journal of the Royal Society Interface*, 12(103), 20140893.
- Walton, A., Richardson, M. J., & Chemero, A. (2014). Self-Organization and Semiosis in Jazz Improvisation. *International Journal of Signs and Semiotic Systems (IJSSS)*, 3(2), 12-25.
- Zhou, W.-X., Sornette, D., Hill, R. A., & Dunbar, R. I. (2005). Discrete hierarchical organization of social group sizes. *Proceedings of the Royal Society of London B: Biological Sciences*, 272(1561), 439-444.