PHILOSOPHICAL TRANSACTIONS B

royalsocietypublishing.org/journal/rstb

Research

Cite this article: Vining WF, Esponda F, Moses ME, Forrest S. 2019 How does mobility help distributed systems compute? *Phil. Trans. R. Soc. B* **374**: 20180375. http://dx.doi.org/10.1098/rstb.2018.0375

Accepted: 1 March 2019

One contribution of 15 to a theme issue 'Liquid brains, solid brains: How distributed cognitive architectures process information'.

Subject Areas:
systems biology

Keywords:
collective computation, multi-agent systems, mobility, consensus

Author for correspondence:
William F. Vining
e-mail: wfvining@cs.unm.edu

How does mobility help distributed systems compute?

William F. Vining[1,†], Fernando Esponda[2,4,†], Melanie E. Moses[1,3] and Stephanie Forrest[4,3]

[1]University of New Mexico, Albuquerque, NM, USA
[2]Instituto Tecnológico Autónomo de México, Mexico DF, Mexico
[3]Santa Fe Institute, Santa Fe, NM, USA
[4]Arizona State University, Tempe, AZ, USA

WFV, 0000-0001-7477-2212

Brains are composed of connected neurons that compute by transmitting signals. The neurons are generally fixed in space, but the communication patterns that enable information processing change rapidly. By contrast, other biological systems, such as ant colonies, bacterial colonies, slime moulds and immune systems, process information using agents that communicate locally while moving through physical space. We refer to systems in which agents are strongly connected and immobile as *solid*, and to systems in which agents are not hardwired to each other and can move freely as *liquid*. We ask how collective computation depends on agent movement. A liquid cellular automaton (LCA) demonstrates the effect of movement and communication locality on consensus problems. A simple mathematical model predicts how these properties of the LCA affect how quickly information propagates through the system. While solid brains allow complex network structures to move information over long distances, mobility provides an alternative way for agents to transport information when long-range connectivity is expensive or infeasible. Our results show how simple mobile agents solve global information processing tasks more effectively than similar systems that are stationary.

This article is part of the theme issue 'Liquid brains, solid brains: How distributed cognitive architectures process information'.

1. Introduction

Natural systems process information in a wide variety of ways. For example, in the vertebrate brain, relatively immobile neurons transmit information over networks that enable both long- and short-distance communication. By contrast, ant colonies collectively process information by individual ants that communicate with other ants as they move through space. We define a *brain* to be a collection of agents that act in concert to process information. A solid brain is one whose agents have fixed position with respect to one another and mostly fixed network connections that are unlikely to change during a computation. In addition to the human brain, other examples include plants and microprocessors, which have components (i.e. cells or transistors) that are fixed in space but use communication networks to exchange information with neighbouring and possibly distant components. By contrast, agents in a liquid brain lack fixed connections and can move independently with respect to one another, as is the case with social insects, cells in the immune system, and many other living systems.

The existence of both liquid and solid brains in nature raises questions regarding the evolutionary pressures that led to these different designs, the algorithms appropriate to each, and the classes of problems that each is best at solving. Solid brains, owing to their static nature, can form hierarchies and other stable network structures that enable multiple levels of abstraction for a given input, and these physical structures make them suitable for processing images and symbols. The stability of solid brains may enable higher levels of

[†]These authors contributed equally to the study.

Electronic supplementary material is available online at https://dx.doi.org/10.6084/m9.figshare.c.4441700.

THE ROYAL SOCIETY PUBLISHING

cognition. Consider, for example, the huge advances made in image processing with deep artificial neural networks [1], where deep refers to a highly structured, fixed relationship between individual agents, loosely mimicking the neural structures found in the animal brains that inspired these models. On the other hand, the dynamic structure of a liquid brain allows it to cover large and variably shaped physical spaces. For instance, ants can defend a large territory from attack from multiple directions, a swarm of robots can search a large physical area, and the immune system can initiate a response in any location in the organism. In a solid brain, complex connectivity such as small-world structure [2] provides scalability through the ability to move information between any two agents using a small number of hops; in a liquid brain, agent mobility can provide similar scalability [3,4]. Slavkov *et al.* demonstrated how the feedback between computation and movement in a robot swarm produces in robust and adaptable morphogenesis [5]. Kao and Couzin [6] describe how animal groups that are spatially dispersed but maintain some internal structure can both sense large amounts of information about the environment and make more accurate collective decisions. In general, liquid brains are constrained by how agents can move through physical space, involve concurrency, and require dynamic reconfiguration to enforce a policy or compute a feature in a spatially distributed environment.

We focus on how agent movement and communication range affect the ability of liquid brains to solve global consensus problems. Because motion is the most striking difference between these two types of brains, we focus on how agent mobility affects information processing. We consider simple agents with limited communication range in simple unstructured environments, and demonstrate how the presence, speed and direction of motion affect simple computational tasks. We argue that movement is an important mechanism by which agents in a liquid brain share information with one another, perhaps also collecting information from, and acting on, the environment. In essence, mobility enables agents to sample many other spatially dispersed agents in order to compute global consensus or to redirect resources to regions that require attention.

In the following sections, we examine mobility as the mechanism by which agents in a liquid brain communicate with other agents and propagate information throughout the system. Using agent-based and mathematical models, we investigate speed of movement, showing that small increases in speed can dramatically improve information processing capabilities. In §2, we review earlier work on consensus problems; §3 describes the liquid cellular automata (LCA) and shows how movement affects its performance on the density classification problem. Section 4 presents a mathematical model that predicts a key feature of LCAs and shows how speed and communication range facilitate information propagation. Sections 5 and 6 vary initial conditions and communication fidelity to show how these factors affect optimal agent speed. The discussion in §7 highlights how these results unify observations from many disparate fields. In §8, we conclude by summarizing how tuning parameters governing physical movement and communication range can generate effective computation in a variety of circumstances. We suggest that the myriad forms of liquid brains in nature is evidence that evolution has discovered many ways to tailor movement and communication for collective computation.

2. Consensus problems

We consider the role of mobility in multi-agent systems trying to achieve global consensus, using majority computations as a canonical example. In this problem, the agents are initially in one of two possible states, `zero` or `one`, and the goal is that all agents adopt the state that initially was in the majority. Agents update their state using local rules taking as input the states of other nearby agents in their neighbourhood. This problem is inspired by the 'density classification' problem in cellular automata (CA) [7], where each cell must converge to one of two states depending on which state was initially in the majority (i.e. the state that had the highest initial density) in the CA.

Mitchell *et al.* analysed the density classification problem to demonstrate how global information processing could occur using only local communication. Because of the fixed structure of a CA (analogous to the fixed connections in a solid brain), a successful solution to the density classification task requires complex spatially distributed representations of intermediate state. These intermediate representations travel from cell to cell via communication in the form of emergent patterns called 'particles' [8] which carry information about partial solutions to distant cells. The fundamental challenge for a CA to solve the density classification problem is overcoming the need for distant cells to share information in order to reach consensus. CA rules that do not produce intermediate representations to communicate across the system cannot move information to reach global consensus. CAs that do produce intermediate representations are usually fragile, for example in the game of life, small changes to initial conditions break the patterns required for computation [9]. Subsequent work showed that more complex connectivity patterns between cells, such as small-world networks, can replace intermediate representations by providing links to propagate information quickly to distant cells [10–12]. Such topologies allow some direct communication between spatially distant cells, removing the need for intermediate representations. Given such topologies, a simple rule, the 'majority rule,' where a cell takes on the same state as the majority of its neighbours, is sufficient to solve the problem. By contrast, the majority rule performs poorly in a traditional CA that lacks long-range communication.

Majority problems are interesting because they demonstrate how a global property of a system (the majority state of all agents) can be communicated and altered by its individual constituents. The need to achieve global consensus through local decision-making occurs in many settings, for example, an ant colony reaching consensus on a new nest site [13] or a swarm of robots selecting an encryption key [14]. If we interpret the initial state of the agents as representing what each agent has sensed from its environment, then density classification serves as a model for how a global environmental property can be computed and shared, e.g. indications of danger or resource availability. In this vein, the CA density classification problem has been used to model how leaf stomata converge on the binary decision to open or close [15] based on the light sensed by each cell in the leaf. Many animals have also been found to use some form of majority voting to come to consensus [16].

We consider a simple liquid brain where the constituent agents resemble the cells of a CA in that they apply a simple rule based on their states and those of their neighbours, but the agents move with respect to one another. Such a model was

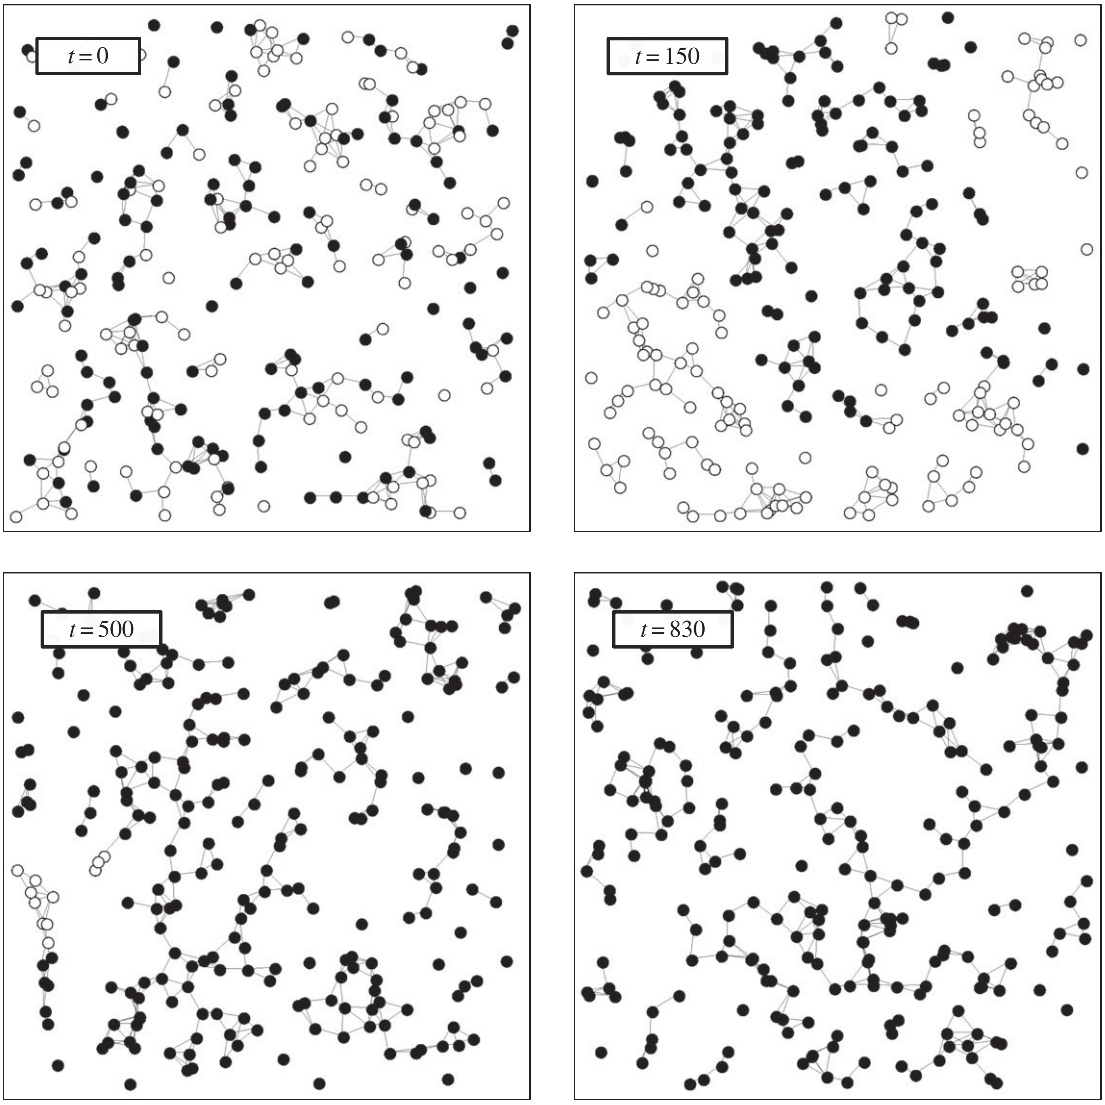

Figure 1. An example LCA run beginning with an initial density of ones (hollow circles) $\rho = 0.48$, agent speed $s = 1$, $N = 255$, $L = 80$ and $r = 5$. Each frame shows a single instant in time. Edges are drawn between agents that are close enough to communicate. Boundaries are closed so that agents neither move nor communicate across them. There is no wrap around. By $t = 150$, two white clusters of `ones` are separated by a single black cluster of `zeros`. The clusters of `ones` shrink slowly until consensus is reached at between $t = 820$ and $t = 830$. Movement eventually combines clusters of local consensus into a global consensus.

referred to by Langton as a 'colony automaton,' [17] and a more complex version resembling a neural network was described by Miramontes and Solé [18–20]. In this work, the authors found that movement facilitates synchronization of agents that act entirely asynchronously [18] and that such a system is capable of simulating logic gates which make it capable of universal computation [19]. In this type of system, information can spread through local communication, and it can be physically carried through space as agents move. We hypothesize that physical movement is a mechanism for spreading information which allows global consensus to be achieved without complex network connectivity. However, agent movement comes at a cost: it is difficult, or perhaps impossible, to preserve spatially distributed intermediate representations when agents constantly change their locations and their neighbours.

To demonstrate the effect of movement, we present a simple agent-based model (ABM) and ask it to solve the density classification problem using mobile agents. We examine how well the problem is solved if each agent uses only the simple local 'majority rule,' where each agent switches state to the majority state of its local neighbourhood. To reiterate, this rule fails to solve the density classification problem in traditional CAs. The density classification problem is most difficult when the initial state of the system is close to 50% `ones` and 50% `zeros`. Such cases are especially challenging if agents have only local communication and are stationary. We show that when agents can move, the majority rule is highly effective, even for the most difficult initial conditions, and its performance increases as agents increase speed.

3. Mobile agents solve global consensus using only local communication in liquid cellular automata

We implemented LCA[1] as an ABM that simulates simple interactions such as those found in a robotic swarm or mobile sensor network. LCA consists of N agents moving in a two-dimensional square arena of size L. Each agent is analogous to a cell in a CA, containing an internal state variable that is updated synchronously according to a rule table shared by all agents. Instead of a fixed neighbourhood, agents have a fixed communication radius r and broadcast their current state to all other agents within that radius. Screenshots of an LCA example run are shown in figure 1.

royalsocietypublishing.org/journal/rstb *Phil. Trans. R. Soc. B* **374**: 20180375

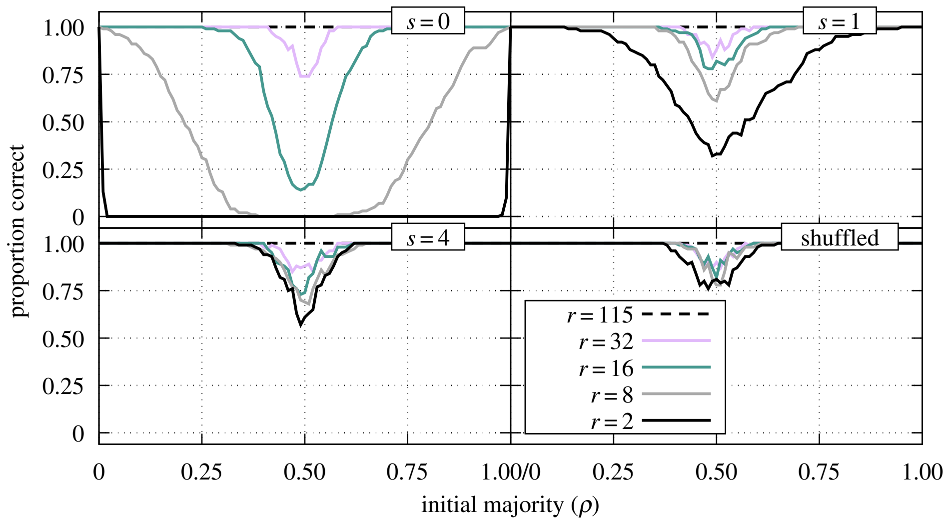

Figure 2. Performance of the majority rule for the density classification task with varying communication radius and speed ($N = 255$, $L = 80$). The proportion correct is the fraction of initial conditions that are correctly classified out of 100 random initial conditions with the given ρ. A correct classification requires that all agents have converged to the correct initial majority state within time T. In 'shuffled', agents are randomly rearranged in space at each time-step. (Online version in colour.)

The simulation proceeds in discrete time steps. At each time step, each agent moves a fixed distance, determined by its speed s, in a straight line along its current heading (if an agent attempts to move beyond the arena boundaries it is reflected back into the arena). The agents then read the states of all other agents within their communication range, and each updates its own state according to the majority rule. Each agent then selects a new heading uniformly at random from the interval $[0, 2\pi)$, and the process repeats so that agents move in a random walk. All agents move at the same constant speed for the duration of the simulation. LCA resembles earlier models that have been used to demonstrate the benefits of mobility for achieving cooperation in the prisoner's dilemma [21,22] and emergence of consensus in a minimal naming game [14]. LCA provides a simple framework in which the affect of movement speed and communication range can each be investigated by adjusting a single parameter.

Each agent uses the majority rule to update its state at every time step to the majority state of its neighbours and itself. When there is no majority (a tie) the agent simply switches to the opposite state deterministically. (By contrast, Watts [12] and others define models where agents select a state at random in the event of a tie.) We replicated multiple simulations using the parameters listed in table 1 in order to determine the effect of changing speed and communication on performance.

The results shown in figure 2 confirm that the majority rule performs poorly with no movement. The system rarely reaches the correct consensus even when agents have a relatively large communication radius (r). Only when r is large enough to cover almost the entire arena ($r = 32$) is a correct classification made by static agents in more than 50% of trials for the hardest case ($\rho = 0.5$).

Introducing a small amount of movement dramatically improves classification performance, especially for smaller communication radii. This is illustrated in figure 2 ($s = 1$) that shows a large improvement in performance even with a slow speed. While performance is still below 50% in the hardest case when $r = 2$, for three-quarters of the initial majorities, agents converge to the correct answer in more than 75% of trials. As s and r increase, performance improves. For speed $s = 4$, the performance for all communication radii converges to almost the best possible case. This best case is illustrated in the 'shuffled' panel where agents are randomly rearranged in space at every time step simulating a situation where the spatial distribution of agents is irrelevant. In this case, consensus is reached through random sampling rather than local consensus.

Table 1. LCA simulation parameters and typical values used in experiments. (T is set to a sufficiently high value to allow agents with a small communication radius (r) and speed (s) enough interactions to have a chance to reach consensus. $\rho = 0.2$ indicates the initial fraction of `ones` is a binomial random variable with an expected value of 0.2.)

model parameter	description	values (unitless)
L	arena side length	80
N	number of agents	255
T	max time	5000
r	communication radius	{2, 8, 16, 32}
s	agent speed	{0, 1, 4, 8}
ρ	expected initial fraction of `ones`	0–1

Mobility lets each agent sample a large fraction of the population over time despite having few, if any, communication partners at any given moment. Combined with the majority rule this allows the formation of local consensus clusters of agents in the same state. Through movement, these clusters grow or shrink until consensus is reached. This is illustrated in figure 3 which shows that despite having very few neighbours at any given time, mobility enables communication with many new agents that move into communication range over time. Similar to the 'effective range of perception' discussed by Couzin [23], movement expands agents' range of communication. As agents move,

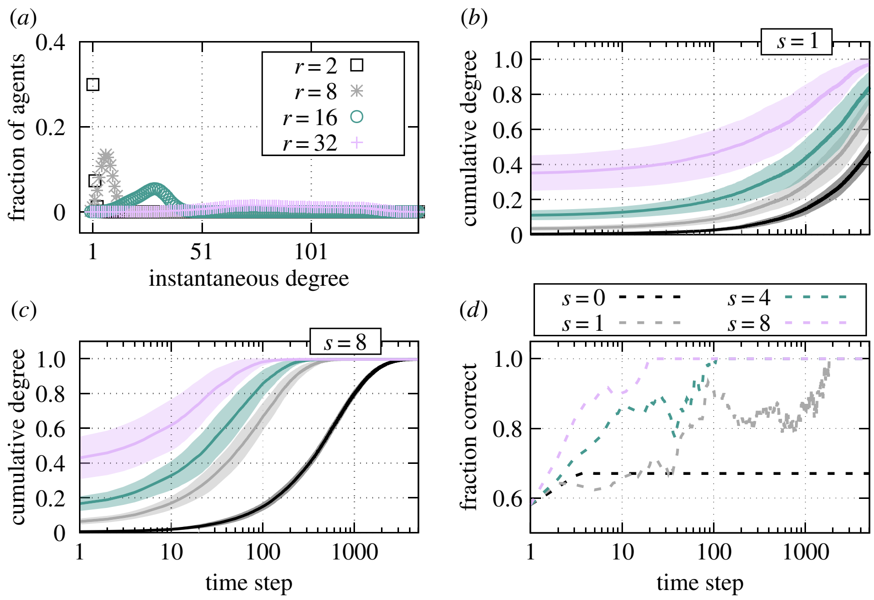

Figure 3. (*a*) Instantaneous degree distribution showing the fraction of agents with a particular degree (number of agents in communication) at a particular time step. As communication radius increases, more agents have higher degree. Degree 0 is not shown for figure clarity, but approximately 60% of all agents with communication radius $r = 2$ have degree 0. (*b*,*c*) Average cumulative degree over time (the cumulative number of connections each agent has had until the indicated time) for two speeds ($s = 1$ and $s = 8$). The shaded region represents plus or minus one standard deviation. (*d*) The fraction of agents in the correct state over time for an example experiment with $r = 8$ and an initial majority close to 60%. With no movement, global consensus is never achieved. With movement, the correct classification is achieved. At slow speeds, there is a long plateau due to clusters of agents in the minority state; these converge to the correct state more quickly with faster movement. (Online version in colour.)

they add edges to the cumulative interaction network, acquiring a substantial sample of the other agents very quickly even with a small communication radius.

These results show how movement enables a liquid system with mobile agents to solve the global density classification problem with the simple local 'majority rule.' The majority rule is ineffective when there is no movement because clusters of local consensus have no mechanism for combining to form a global consensus. When agents can move, they sample (and influence) a larger fraction of the other agents, which resolves conflicts at the boundaries between clusters. This allows clusters to combine into larger clusters. As speed increases, clusters in the minority state persist for less time (figure 3*d*).

Beyond a certain point, however, increased speed has little effect on performance, as shown in figure 2. We confirmed this for speeds as high as $s = 300$ (data not shown) with similar results to those described by Baronchelli *et al.* [14] for the minimal naming game. At low speeds, consensus is reached through competition between clusters of agents that have reached consensus locally, but for higher speeds consensus is reached by a process that is effectively global sampling.

Speed is not the only defining characteristic of movement; another important consideration is *direction*. The previous simulations used a simple random walk, which results in very slow diffusion of agents. To investigate the effect of other movement patterns, we implemented a correlated random walk. Following methods in [24], agents select a new heading h_{t+1} from a normal distribution $\mathcal{N}(h_t, \sigma)$ at each time step where h_t is the heading at time t and σ is the standard deviation of the turning angle. With $\sigma = 0$ agents move with ballistic motion turning only when they reflect off arena boundaries; $\sigma = \pi$ approximates the random walk described above.

Our experiments show that a correlated random walk with $\sigma \leq \pi/4$ improves performance. Higher correlations produce faster consensus times. Because a correlated walk results in greater displacement than the diffusive random walk, agents quickly sample a larger number of other agents. Electronic supplementary material, figure S3 shows how mobile agents with relatively low speed can achieve similar performance to faster agents simply by moving in a straighter walk.

4. Predicting how quickly agents accumulate novel interactions

To better understand how information flows among agents, we estimate how the cumulative degree of the average agent changes over time. This indicates how quickly each agent is expected to communicate with new agents that it has not communicated with before.

We begin by counting the number of novel agents (those not seen in the previous time step) that each agent interacts with at each time step. We estimate the number of novel agents as a function of an agent's communication radius (r), its speed (s) and its type of movement. We then use this value to calculate the expected cumulative degree of the interaction graph at a particular time step as a function of the number of agents and the size of the environment. To simplify this problem, we assume that the arena wraps around at the boundaries forming a torus. Because of this assumption, the predictions we develop in this section apply when r is very small relative to the arena size.

Let $M = N/L^2$ be the density of agents and let $A_r = \pi r^2$ be the area covered by a single agent's communication range. Then the expected number of agents within the communication range of an arbitrary agent is $A_r M = \pi r^2 N/L^2$.

We assume that the density of agents is constant and that agents are uniformly spread across the arena at all times. We predict the number of novel agents encountered at each time

royalsocietypublishing.org/journal/rstb *Phil. Trans. R. Soc. B* **374**: 20180375

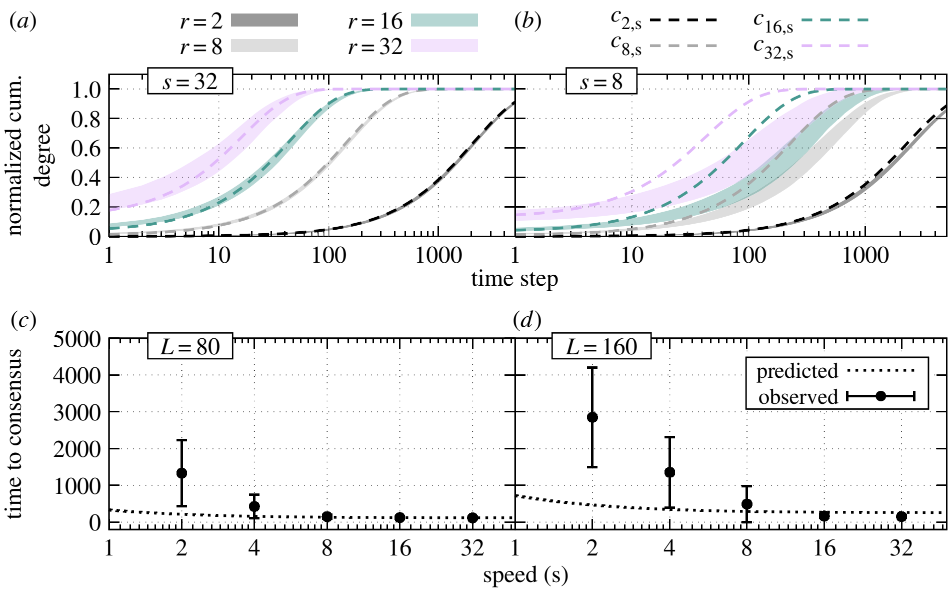

Figure 4. Predictions made by equations (4.2) and (4.3). (*a*,*b*) The predicted cumulative degree as a fraction of the total number of agents for a large arena ($L = 160$ and $N = 1023$) at speeds $s = 8$ (*b*) and $s = 32$ (*a*). The large arena mitigates the effect of the arena boundaries. The shaded region is the observed average cumulative degree plus or minus one standard deviation in the LCA. (*c*,*d*) The predicted time to reach consensus, based on consensus occurring on average at a cumulative degree of $0.21N$ in the small arena (*c*, $L = 80$, $N = 255$) and $0.12N$ in the large arena (*d*). The observed time to consensus in LCA is shown for comparison with error bars indicating one standard deviation. (Online version in colour.)

step. First, we estimate the number of agents that remain within A_r of a focal agent after one time step by computing the expected distance an agent currently in A_r can travel away from the focal agent at speed s. This distance is $r + d_s$, where d_s is the expected straight-line distance travelled by an agent in one time step at speed s. Note that d_s depends on the type of motion agents undertake between each time step. For the straight line motion used in our LCA experiments $d_s = s/2$.

Assuming an agent within the area of influence can be in any location within $\pi(r + d_s)^2$ after one time step, the probability that an agent remains within the area of influence of another is the ratio of the area of influence of an agent (A_r) and the area an agent could be in the next time step $p^{\text{inside}}_{r,s} = r^2/(r + d_s)^2$.

By the assumption of uniform density, the number of novel agents, $m_{r,s}$, in an area at a particular time step is equal to the number of agents that left the area from the previous time step. As a fraction of the total number of agents this is $f_{r,s} = (A_r/L^2)(1 - p^{\text{inside}}_{r,s})$. As a first-order approximation to sampling without replacement, we subtract from the initial population the number of agents that are initially within the area of influence, $N - A_rM$, and obtain the estimated number of novel agents:

$$m_{r,s} = f_{r,s}(N - A_rM). \tag{4.1}$$

Note that if an agent is not moving ($s = 0$) then it will always interact with the same agents and $m_{r,s} = 0$. Conversely, when an agent moves at extremely high speed ($s \gg L$) $m_{r,s} = A_rM$—the expected number of agents to be found within its area of influence sampled at random from all other agents.

Given the number of novel agents $m_{r,s}$, the probability of any two agents interacting is

$$p_{r,s} = \frac{\binom{N}{m_{r,s}} - \binom{N-1}{m_{r,s}}}{\binom{N}{m_{r,s}}} = \frac{m_{r,s}}{N}.$$

Then the probability of two agents interacting at least once in t time steps is approximately $1 - e^{-p_{r,s}t}$.

Finally, we include an extra term to account for the initial A_rM neighbours. This yields the following expression for $c_{r,s}$, the probability that two agents have interacted as a function of time:

$$c_{r,s} \approx 1 - \left(\left(1 - \frac{A_r}{L^2}\right)e^{-p_{r,s}t}\right). \tag{4.2}$$

Figure 4 shows that the predicted $c_{r,s}$ approximates the normalized cumulative degree in the LCA model, particularly when r is small and the arena size is relatively large. However, at slower speeds, as communication range increases, the fit deteriorates owing to edge effects in the simulated arena. When agents have large r and low s, some agents have a reduced area of communication that is cut off by the edge of the arena. They stay near the edge for long periods of time because speed is slow. This shows that we can predict cumulative degree under ideal cases, but we need an ABM to simulate more complicated circumstances where edge effects might dominate, which may often be the case in real biological systems.

From equation (4.2), we can derive the estimated number of time steps needed to get an expected cumulative degree of $Nc_{r,s}$:

$$t = -\frac{N}{m_{r,s}} \ln\left(\frac{1 - c_{r,s}}{1 - A_r/L^2}\right). \tag{4.3}$$

Consensus is reached as agents gain a larger sample of the states of other agents. We measured the median cumulative degree at the time consensus was reached for $s \in \{2, 4, 8, 16\}$. We used a maximum of 100 random initial conditions for each speed and only those runs which reached consensus are included. The median normalized cumulative degree over all speeds was $0.12N$ in an arena with $N = 1023$ and $L = 160$ and $0.21N$ for the small arena ($N = 255$, $L = 80$). These results are summarized in the electronic supplementary material, figure

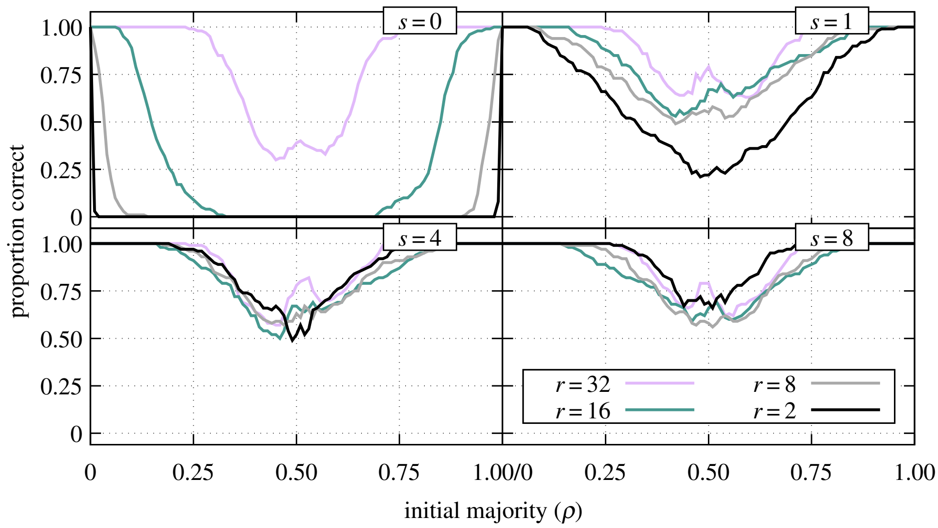

Figure 5. Proportion of LCA simulations that correctly classify an initial condition where agents in different states are spatially segregated with all agents in state `one` to the left and all agents in state `zero` to the right. $r = 2$ performs best at high speeds because the agents are able to mix before sharing much information. (Online version in colour.)

S4. Equation (4.3) should, therefore, correlate with the average time to reach consensus, as shown in figure 4.

5. Movement overcomes challenging initial conditions

Some initial conditions are particularly challenging for the density classification problem. In the previous section, we considered only initially well-mixed agents, a relatively easy set of conditions to classify. Here, we assume that the states are spatially segregated. All agents in state `one` are on the left and all agents in state `zero` are on the right, and the initial density ρ is controlled by moving the dividing line left or right. Figure 5 illustrates that the proportion of correctly classified problems is lower than when agents are well mixed, but movement allows agents to overcome difficult initial conditions and make a reasonably accurate classification much of the time.

Surprisingly, figure 5 illustrates cases where low communication range is beneficial. While the previous results showed that the LCA performed better as the communication radius increases, for segregated initial conditions, the lowest r can outperform the largest r when speed is high. Because of the low communication radius, agents rarely interact, and agents in different states become well mixed more rapidly than they share information. This demonstrates how movement provides a powerful and flexible tool that can be exploited to facilitate solutions to a variety of problems under a variety of conditions. In this case, fast movement allows for the use of a small communication range, making the system robust to pathologically constructed initial conditions. The following section illustrates another case where the parameters of movement can be adjusted to solve the density classification problem under different conditions.

6. Speed balances local versus global information exchange

We consider a second liquid brain model that simulates more complex interactions between agents.[2] We assume that more complex interactions may require more time to complete, and that constrains how fast the agents can move. Examples of varying interaction times include: ants can quickly sense each other to determine nest membership but take more time to perform allogrooming and to form a bridge-like structures [25]; nectar exchange in honeybees takes only a few seconds [26] while it takes several minutes for guards to examine bees joining a nest [27].

We preserve most of the features of LCA to solve the density classification problem using local majority rule; agents move at a constant speed on a two-dimensional grid, and agents exchange information with other agents within a fixed radius. However, we include a probability that the state of an agent is miscommunicated. When an agent polls its neighbours to compute the majority, there is a fixed probability of error. Each agent has the same independent probability p of reporting its real state on each query, thus if an agent is queried repeatedly on K occasions it is expected to report its true state pK times; in this way, for an agent to increase the chance of obtaining the correct state of its neighbours, it must query each of them repeatedly.

Each agent can poll its neighbours up to 10 times in each time step. The agents remember the reported states and identities of their neighbours. In our implementation, each individual maintains a running estimate of the mode of the answers for each agent within its radius. The memory is cleared after computing the local majority at each time step. We use the number of queries required for an accurate estimate of an agent's state as a proxy for the complexity required of agent–agent interactions: the closer p is to 0.5, the more polling will be required to reach a correct consensus. However, agents can move instead of polling their neighbours, so the more an agent needs to poll its neighbours before moving, the lower its speed. We implement this by giving each agent a fixed movement budget per time step that is composed of its speed (travel distance per time step) and the polling cycles allowed for its current neighbours: movement budget = speed + polling cycles. In this way, the number of polling cycles increases as speed decreases. This implementation forces a trade-off similar to the sampling versus speed trade-off seen in immune cells [28].

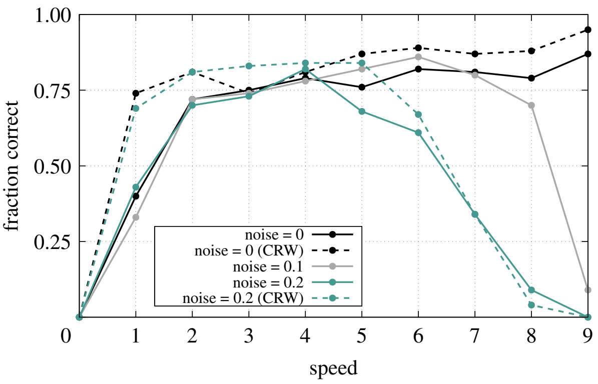

Figure 6. The effect of increasing agent speed in a noisy environment on the performance of the majority rule for the majority classification task with the initial density $\sigma = 0.55$ and a communication radius of three. A diffusive random walk is shown as a solid and the improvement achieved by a correlated random walk (CRW) is shown as a dashed line. (Online version in colour.)

Figure 6 shows that for any given level of noise there is an optimum intermediate speed that combines depth (more polling) with the breadth (more agent interactions) for maximum performance. Moving too fast prevents agents from accurately determining the correct states of their neighbours, while moving too slowly prevents agents from interacting sufficiently to reach global consensus. Even when we run longer simulations for slow-moving agents, there is still an optimal intermediate speed that results from the dynamic nature of the majority rule. Moving slowly creates large spatial clusters bearing, on occasion, the minority state, which are difficult to flip and can ultimately dominate the population. Moving faster prevents these large clusters from forming as in figure 3*d*. Moving using a correlated random walk that generates more directed movement with greater displacement improves performance in this model, similar to the improvement described in §3.

Although our second model abstracts away details of agent interactions, it exhibits similar results as those reported in [21], where moving agents interact by playing the iterated prisoner's dilemma. That work showed that there is an ideal intermediate movement speed that enhances the evolution of cooperation because it can create complex structure in the network of interactions between players. This is beneficial in the prisoner's dilemma because complex structure facilitates both the emergence of cooperation through repeated interaction and the spread of cooperation by gradually changing which agents interact. In other words, moderate speed balances the need for stable network structure and the spread of information.

The mathematical model presented in §4 can be extended to incorporate noisy communication. This is accomplished by making the number of novel agents inversely proportional to speed so $m_{r,s} = f_{r,s}(N - A_rM)/(\alpha s + 1)$ where $0 < \alpha \le 1$. This gives the same result that intermediate speeds maximize $m_{r,s}$, consistent with our observation that intermediate speed maximizes performance in the presence of noise.

7. Discussion

The abundance of both solid and liquid brains in nature suggests that evolution has found advantages with both approaches. Traditionally, brains are thought to be composed of components such as neurons that are fixed in space with communication among components through networks. For other distributed biological systems such as ant colonies, information exchange occurs locally between agents as they move through space. We suggest that there is probably a trade-off: mobile agents have more flexible communication patterns determined through movement but lack dedicated communication networks owing to the difficulty of maintaining fixed communication structures as agents move. In the light of this trade-off, two different solutions have emerged—agents can move information directly in liquid brains or through communication channels in solid brains. Liquid brains may have evolved movement for other reasons (i.e. the need to sample environments across space), but solid brains have the advantage of being able to store and use structures to process more complex information.

Our analysis focuses on how mobility affects the ability to compute global consensus. We show that for a set of simple agents able to communicate only within a small local radius, many instances of the global density classification problem can be solved by the majority rule if the agents are mobile, but not if they are immobile. Additionally, more challenging instances of the problem can be solved by mobile agents by adjusting the way agents move (figures 5 and 6). Problems that require longer or more complex communication among agents are best solved by agents that move at an intermediate speed, balancing time for communication with the current neighbourhood against communication with agents in other locations. Problems that require more mixing can be solved by moving faster and reducing the communication range. Movement makes it simple to adapt the LCA to work effectively in a variety of situations.

In systems without movement, information can be transmitted from one part of a spatial domain to another using long-distance communication channels. These long-distance connections combined with many local connections are a key feature of human and animal brains [29]. Rent's rule describes both these biological systems and technological systems such as computer microprocessors with a similar communication network between transistors—as sharing similar mathematical properties. In both cases the probability of connection between two components decays with distance [30–33]. Furthermore, many systems can be characterized by more general small-world topologies [2] in which agents are connected primarily to their local neighbours, but there is some small amount of long-distance communication. Such complex topologies allows scalable communication between the agents in a solid brain.

Even without complex connectivity, simple solid brains, such as traditional CAs, are able to move information through space in the form of sophisticated intermediate representations, such as patterns of states, that enable global communication. Take, for example, the gliders in Conway's game of life or the rules evolved for density classification by Crutchfield and Mitchell [8] which facilitate global communication by way of 'particles' that move through space as representations of partial solutions. These rules depend not just on the regular structure of the network, but also on the periodicity of the network. A particle must be able to travel through every cell and back to where it started by moving only in one direction. Without a periodic topology, particles reach a dead end preventing the CA from reaching consensus. Relying on a particular topology can be a limitation for solid brains. Liquid brains, however, are robust to changes in network topology by their very nature.

There are many ways solid brains can support global information processing of the type discussed in this paper. Complex connectivity in the form of small-world networks is a good example. We investigated the performance of our model with random rewiring to add long-distance connections. Performance using these random connections is roughly the same for $\rho = 0.5$, but worse for all other initial majorities (see the electronic supplementary material, figure S2). Moreira *et al.* investigated conditions under which the majority rule can solve the density classification problem [34]. They found that the presence of noise in reading a neighbour's state, combined with a small amount of randomness in the topology of the CA, were sufficient to achieve good classification performance. This is notable because the mechanism by which noise helps the CA is similar to the mechanism by which mobility helps—it breaks down the boundaries between clusters of local consensus enabling a global solution.

Other types of problems, such as the emergence of global cooperation in the prisoner's dilemma and its variants, have been studied in the context of liquid brains [21,22,35,36]. Because the prisoner's dilemma presents a fundamentally different problem from global consensus, the effect of movement is different. In the prisoner's dilemma, all information used to produce cooperative behaviour is found in the local region of an agent, so slow movement helps to maintain cooperation. By contrast, the density classification problem requires the computation of a global property distributed over all of the agents. Because information needed for this computation is distributed across the population of agents, increased sampling through faster movement is beneficial. Other types of problems such as estimating the number of agents in the system can be solved by mobile agents [37] and can serve as building blocks for more complex tasks such as quorum sensing.

An important feature of some liquid brains, not addressed by the LCA, is their ability to build and use physical structures to limit some interactions and facilitate others. For example, ants concentrate interactions at the nest entrance [38], while immune cells are concentrated in lymph nodes to facilitate more frequent interactions, and robot swarms use structure to determine a shortest path [39]. Even the spatial distribution of features in the environment can be exploited by liquid brains to enable optimal decision-making comparable to decision-making in solid brains [40].

The LCA provides a platform for investigating other important characteristics of liquid brains. For example, quorum sensing has been shown to lead to better decision accuracy than straight majority voting [41] and can prevent poor decisions based on copying the behaviour of neighbours [42]. The LCA model could easily be extended to investigate these types of questions by substituting a quorum-based rule for the majority rule. Additionally, as discussed in §2, further work is needed to refine what types of movement are beneficial for which types of computation. Many other types of movement have been observed in natural systems, including: the lognormal walk of T-cells [43]; Lévy walks [44]; and two-phase walks [45], which are optimal for some search tasks and could be studied from the perspective of communication and computation in a liquid brain. Additional open questions include the differences in learning algorithms employed by liquid or solid brains, differences in how they achieve robustness and resilience, and the role of diversity and specialization of the agents.

Computing in distributed systems that are similar to liquid brains—with frequently changing connections between components—is increasingly important in computer science as new technologies such as robotics and the Internet of Things become prevalent. Recent theoretical work in this area is summarized in [46,47]. Kuhn *et al.* [48] proved that many foundational distributed algorithms can be implemented on these dynamic networks under certain connectivity assumptions.

A significant amount of theoretical research in population protocols [49] studies a model of computation similar to the LCA, including the specification of protocols for solving the majority problem [50,51]. LCA differs from population protocols in several ways: LCA communication is always through bidirectional broadcast in which all agents within communication range send and receive messages, while population protocols model pairwise interactions between an identified initiator and receiver. Population protocols are inherently asynchronous in contrast to the synchronous LCA model presented here. An asynchronous version of LCA is worthy of future investigation. The most important distinction is that while the population protocol model facilitates detailed mathematical analysis, it is built on an interaction network that does not capture the constraints on interactions imposed by agents that can only communicate locally as they move through space. While spatial interactions among mobile agents are difficult to capture in a simple mathematical model, LCA shows that mobile agents can compute global consensus even under realistic spatial constraints. Furthermore, our analysis of cumulative degree shows that some aspects of how agents interact are predictable by a simple mathematical model.

Until now, most computational systems have been solid, for example, microprocessors have static components and fixed wires, even field-programmable gate arrays have components fixed in space with flexible communication but no mobility. Neural networks emulate the idea of immobile components and flexible communication networks inspired by solid human brains. However, many forms of unconventional computing are much closer to the liquid end of the spectrum, for example, DNA computing, ant-inspired algorithms, swarm robotics, and mobile *ad hoc* networks (MANET) are all systems in which movement through physical space is part of the problem and an inherent aspect of the computation. A robot swarm, for example, can exploit movement to achieve a global map of sensor readings [52] and a MANET can exploit movement to increase its capacity [53]. We have shown how those systems which lack the traditional computational infrastructure for storing and moving information can take advantage of mobility in order to compute. Understanding how consensus can be reached by a collection of loosely connected mobile agents highlights distributed computational approaches in nature and suggests how these approaches can be implemented in artificial systems.

8. Conclusion

Although many previous models examine communication in distributed systems, in most prior work the locations of agents in physical space is decoupled from their probability of communicating. In this paper, we combined simulations and mathematical modelling to quantify how movement through physical space affects the ability of a distributed system to compute global consensus.

A complete mathematical treatment of liquid brains is challenging for several reasons: (i) naturally occurring distributed systems are rarely well-mixed, (ii) physical environments have boundaries that create edge effects and complicate mathematical predictions, (iii) physical topographies are complex, and (iv) real systems are often asynchronous. In this paper, we address the first two of these complexities, leaving the second two for future work.

A general theory of distributed communication with spatial constraints may remain elusive, but evolution does not need a general theory. It has succeeded by discovering specific solutions tailored to specific environmental situations. Our models show how tuning a small number of parameters governing movement and communication can allow systems to find global consensus given different initial conditions, movement patterns, communication error rates, and agent densities. Our approach is a step towards a more general theory of distributed computing under spatial constraints, but more importantly it provides insight into how evolution leverages different movement and communication strategies to compute in liquid brains.

Data accessibility. This article has no additional data.

Authors' contributions. F.E. and W.V. contributed equally to this work. W.V. conceived of and implemented the LCA model, carried out the experiments in §§3 and 5. F.E. derived the mathematical model in §4 and implemented and carried out experiments in §6 and wrote the initial draft. M.M. and S.F. conceived of the paper, and contributed analysis and text throughout. The authors thank two anonymous reviewers for very helpful comments and suggestions.

Competing interests. We declare we have no competing interests.

Funding. F.E. acknowledges the support provided by ASU's Biodesign Institute and the Asociación Mexicana de Cultura, A. C. S.F. acknowledges the partial support of NSF(1518878), DARPA(FA8750-15-C-0118), AFRL(FA8750-15-2-0075) and the Santa Fe Institute. M.M. is funded by a James S. McDonnell Foundation Complex Systems Scholar Award and W.V. is funded by NASAMUREP no. NNX15AM14A.

Acknowledgements. We thank Ricard Solé and the participants of the Liquid Brains, Solid Brains workshop at the Santa Fe Institute. W.V. thanks Jed Crandall for guidance and helpful feedback.

Endnotes

[1]Source code available from: github.com/wfvining/liquid-ca.
[2]Source code available from: github.com/fesponda/noisy-liquid-brain.

Reference

1. LeCun Y, Bengio Y, Hinton G. 2015 Deep learning. *Nature* **521**, 436–444. (doi:10.1038/nature14539)
2. Watts DJ, Strogatz SH. 1998 Collective dynamics of 'small world' networks. *Nature* **393**, 440–442. (doi:10.1038/30918)
3. Miramontes O, Luque B. 2002 Dynamical small-world behavior in an epidemical model of mobile individuals. *Physica D* **168**, 379–385. (doi:10.1016/S0167-2789(02)00525-0)
4. Tang J, Scellato S, Musolesi M, Mascolo C, Latora V. 2010 Small-world behavior in time-varying graphs. *Phys. Rev. E* **81**, 055101. (doi:10.1103/PhysRevE.81.055101)
5. Slavkov I, Carrillo-Zapata D, Carranza N, Diego X, Jansson F, Kaandorp J, Hauert S, Sharpe J. 2018 Morphogenesis in robot swarms. *Sci. Rob.* **3**, eaau9178. (doi:10.1126/scirobotics.aau9178)
6. Kao AB, Couzin ID. 2019 Modular structure within groups causes information loss but can improve decision accuracy. *Phil. Trans. R. Soc. B* **374**, 20180378. (doi:10.1098/rstb.2018.0378)
7. Mitchell M, Crutchfield JP, Hraber PT. 1994 Evolving cellular automata to perform computations: mechanisms and impediments. *Physica D* **75**, 361–391. (doi:10.1016/0167-2789(94)90293-3)
8. Crutchfield JP, Mitchell M. 1995 The evolution of emergent computation. *Proc. Natl Acad. Sci. USA* **92**, 10 742–10 746. (doi:10.1073/pnas.92.23.10742)
9. Gardner M. 1970 Mathematical games: the fantastic combinations of John Conway's new solitaire game 'life'. *Sci. Am.* **223**, 120–123. (doi:10.1038/scientificamerican1070-120)
10. Sipper M, Ruppin E. 1997 Co-evolving architectures for cellular machines. *Physica D* **99**, 428–441. (doi:10.1016/S0167-2789(96)00172-8)
11. Tomassini M, Giacobini M, Darabos C. 2005 Evolution and dynamics of small-world cellular automata. *Complex Syst.* **15**, 261–284.
12. Watts DJ 1999 *Small worlds: the dynamics of networks between order and randomness. Princeton studies in complexity*. Princeton, NJ: Princeton University Press.
13. Franks NR, Pratt SC, Mallon EB, Britton NF, Sumpter DJ. 2002 Information flow, opinion polling and collective intelligence in house–hunting social insects. *Phil. Trans. R. Soc. Lond. B* **357**, 1567–1583. (doi:10.1098/rstb.2002.1066)
14. Baronchelli A, Díaz-Guilera A. 2012 Consensus in networks of mobile communicating agents. *Phys. Rev. E* **85**, 016113. (doi:10.1103/PhysRevE.85.016113)
15. Peak D, West JD, Messinger SM, Mott KA. 2004 Evidence for complex, collective dynamics and emergent, distributed computation in plants. *Proc. Natl Acad. Sci. USA* **101**, 918–922. (doi:10.1073/pnas.0307811100)
16. Conradt L, Roper TJ. 2003 Group decision-making in animals. *Nature* **421**, 155–158. (doi:10.1038/nature01294)
17. Langton CG. 1986 Studying artificial life with cellular automata. *Physica D* **22**, 120–149. (doi:10.1016/0167-2789(86)90237-X)
18. Miramontes O, Solé RV, Goodwin BC. 1993 Collective behavior of random-activated mobile cellular automata. *Physica D* **63**, 145–160. (doi:10.1016/0167-2789(93)90152-Q)
19. Solé R, Delgado J. 1996 Universal computation in fluid neural networks. *Complexity* **2**, 49–56. (doi:10.1002/(ISSN)1099-0526)
20. Solé RV, Miramontes O. 1995 Information at the edge of chaos in fluid neural networks. *Physica D* **80**, 171–180. (doi:10.1016/0167-2789(95)90075-6)
21. Meloni S, Buscarino A, Fortuna L, Frasca M, Gómez-Gardeñes J, Latora V, Moreno Y. 2009 Effects of mobility in a population of prisoner's dilemma players. *Phys. Rev. E* **79**, 3–6. (doi:10.1103/PhysRevE.79.067101)
22. Sicardi EA, Fort H, Vainstein MH, Arenzon JJ. 2009 Random mobility and spatial structure often enhance cooperation. *J. Theor. Biol.* **256**, 240–246. (doi:10.1016/j.jtbi.2008.09.022)
23. Couzin I. 2007 Collective minds. *Nature* **445**, 715. (doi:10.1038/445715a)
24. Hecker JP, Moses ME. 2015 Beyond pheromones: evolving error-tolerant, flexible, and scalable ant-inspired robot swarms. *Swarm Intell.* **9**, 43–70.
25. Hölldobler B, Wilson EO. 1990 *The ants*. Cambridge, MA: Harvard University Press.
26. Goyret J, Farina WM. 2005 Non-random nectar unloading interactions between foragers and their receivers in the honeybee hive. *Naturwissenschaften* **92**, 440–443. (doi:10.1007/s00114-005-0016-7)
27. Stabentheiner A, Kovac H, Schmaranzer S. 2002 Honeybee nestmate recognition: the thermal behaviour of guards and their examinees. *J. Exp. Biol.* **205**, 2637–2642.
28. Krummel MF, Bartumeus F, Gérard A. 2016 T cell migration, search strategies and mechanisms. *Nat. Rev. Immunol.* **16**, 193. (doi:10.1038/nri.2015.16)
29. Sporns O, Chialvo DR, Kaiser M, Hilgetag CC. 2004 Organization, development and function of complex brain networks. *Trends Cogn. Sci.* **8**, 418–425. (doi:10.1016/j.tics.2004.07.008)
30. Bezerra GB, Forrest S, Moses M, Davis A, Zarkesh-Ha P. 2010 Modeling NoC traffic locality and energy consumption with rent's communication probability distribution. In *Proceedings of the 12th ACM/IEEE International workshop on system level interconnect*

prediction, SLIP '10, pp. 3–8, New York, NY, USA: ACM Press.

31. Christie P, Stroobandt D. 2000 The interpretation and application of rent. *IEEE Trans. Very Large Scale Integr. VLSI Syst.* **8**, 639–648. (doi:10.1109/92.902258)
32. Moses M, Bezerra G, Edwards B, Brown J, Forrest S. 2016 Energy and time determine scaling in biological and computer designs. *Phil. Trans. R. Soc. B* **371**, 20150446. (doi:10.1098/rstb.2015.0446)
33. Solé RV, Valverde S, Casals MR, Kauffman SA, Farmer D, Eldredge N. 2013 The evolutionary ecology of technological innovations. *Complexity* **18**, 15–27. (doi:10.1002/cplx.v18.4)
34. Moreira AA, Mathur A, Diermeier D, Amaral LAN. 2004 Efficient system-wide coordination in noisy environments. *Proc. Natl Acad. Sci. USA* **101**, 12 085–12 090. (doi:10.1073/pnas.0400672101)
35. Antonioni A, Tomassini M, Buesser P. 2014 Random diffusion and cooperation in continuous two-dimensional space. *J. Theor. Biol.* **344**, 40–48. (doi:10.1016/j.jtbi.2013.11.018)
36. Enquist M, Leimar O. 1993 The evolution of cooperation in mobile organisms. *Anim. Behav.* **45**, 747–757. (doi:10.1006/anbe.1993.1089)
37. Musco C, Su H-H, Lynch NA. 2017 Ant-inspired density estimation via random walks. *Proc. Natl Acad. Sci. USA* **114**, 10 534–10 541. (doi:10.1073/pnas.1706439114)
38. Pinter-Wollman N, Wollman R, Guetz A, Holmes S, Gordon DM. 2011 The effect of individual variation on the structure and function of interaction networks in harvester ants. *J. R. Soc. Interface* **8**, 1562–1573. (doi:10.1098/rsif.2011.0059)
39. de Oca MAM, Ferrante E, Scheidler A, Pinciroli C, Birattari M, Dorigo M. 2011 Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intell.* **5**, 305–327. (doi:10.1007/s11721-011-0062-z)
40. Marshall JA, Bogacz R, Dornhaus A, Planqué R, Kovacs T, Franks NR. 2009 On optimal decision-making in brains and social insect colonies. *J. R. Soc. Interface* **6**, 20080511. (doi:10.1098/rsif.2008.0511)
41. Marshall JA, Kurvers RH, Krause J, Wolf M. 2018 Quorums enable optimal pooling of independent judgements. *bioRxiv* 394460. (doi:10.1101/394460)
42. Sumpter DJ, Krause J, James R, Couzin ID, Ward AJ. 2008 Consensus decision making by fish. *Curr. Biol.* **18**, 1773–1777. (doi:10.1016/j.cub.2008.09.064)
43. Fricke GM, Letendre KA, Moses ME, Cannon JL. 2016 Persistence and adaptation in immunity: T cells balance the extent and thoroughness of search. *PLoS Comput. Biol.* **12**, 1–23. (doi:10.1371/journal.pcbi.1004818)
44. Viswanathan GM, Buldyrev SV, Havlin S, Da-Luz MGE, Raposo EP, Stanley HE. 1999 Optimizing the success of random searches. *Nature* **401**, 911–914. (doi:10.1038/44831)
45. Bénichou O, Loverdo C, Moreau M, Voituriez R. 2006 Two-dimensional intermittent search processes: an alternative to Lévy flight strategies. *Phys. Rev. E* **74**, 020102. (doi:10.1103/PhysRevE.74.020102)
46. Casteigts A, Flocchini P, Quattrociocchi W, Santoro N. 2012 Time-varying graphs and dynamic networks. *Int. J. Parallel Emergent Distrib. Sys.* **27**, 387–408. (doi:10.1080/17445760.2012.668546)
47. Michail O, Spirakis PG. 2018 Elements of the theory of dynamic networks. *Commun. ACM* **61**, 72–81. (doi:10.1145/3181977)
48. Kuhn F, Lynch N, Oshman R. 2010 Distributed computation in dynamic networks. In *Proceedings of the 42nd ACM symposium on Theory of computing - STOC '10*, p. 513, New York, New York, USA, 2010. ACM Press.
49. Angluin D, Aspnes J, Diamadi Z, Fischer MJ, Peralta R. 2006 Computation in networks of passively mobile finite-state sensors. *Distrib. Comput.* **18**, 235–253. (doi:10.1007/s00446-005-0138-3)
50. Alistarh D, Gelashvili R, Vojnović M. 2015 Fast and exact majority in population protocols. In *Proceedings of the 2015 ACM Symposium on principles of distributed computing - PODC '15*, pp. 47–56, New York, New York, USA: ACM Press.
51. Angluin D, Aspnes J, Eisenstat D. 2008 A simple population protocol for fast robust approximate majority. *Distrib. Comput.* **21**, 87–102. (doi:10.1007/s00446-008-0059-z)
52. Winfield AF. 2000 Distributed sensing and data collection via broken ad hoc wireless connected networks of mobile robots. In *Distributed autonomous robotic systems 4* (eds LE Parker, G Bekey, J Barhen), pp. 273–282. Tokyo, Japan: Springer.
53. Grossglauser M, Tse DNC. 2001 Mobility increases the capacity of ad-hoc wireless networks. *IEEE/ACM Trans. Netw.* **10**, 477–486. (doi:10.1109/TNET.2002.801403).