



Collective intelligence as a public good

Collective Intelligence
August-September 2022: 1–18
© The Author(s) 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/26339137221083293
journals.sagepub.com/home/col



Naomi Ehrich Leonard

Department of Mechanical & Aerospace Engineering, Princeton University, USA

Simon A Levin

Department of Ecology & Evolutionary Biology, Princeton University, USA

Abstract

We discuss measures of collective intelligence in evolved and designed self-organizing ensembles, defining collective intelligence in terms of the benefits to be gained through the exchange of information and other resources, as well as through coordination or cooperation, in the interests of a public good. These benefits can be numerous, from estimating a hard-to-observe cue to efficiently searching for resource. The measures should also account for costs to individuals, such as in attention or energy, and trade-offs for the ensemble, such as the flexibility to respond to an important change in the environment versus stability that is robust to unimportant variability. When there is a tension between the interests of the individual and those of the group, game-theoretic considerations may affect the level of collective intelligence that can be achieved. Models of individual rules that yield collective dynamics with multi-stable solutions provide a means to examine and shape collective intelligence in evolved and designed systems.

Keywords

Collective intelligence, cooperation, selection, emergence, self-organization, nonlinear dynamics, multi-stability, decision-making, multi-robot systems

Introduction

Ensembles of individual units, in general, have the capacity to perform better in a variety of ways than individuals on their own, in part because individuals can share information and other resources, they can coordinate or collaborate on activities, and/or they have the potential for differentiation of function, any or all of which can be leveraged in the production or maintenance of a public good. However, for these opportunities to be leveraged in the service of a public good, the associated trade-offs that come from competing demands, their costs, and the limitations of individuals, must be well managed. Importantly, if the interests of the individual and the ensemble are in tension, the ability of the

ensemble to maintain a public good will depend on the ability of the individuals to cooperate.

A suitable measure of collective intelligence is thus the difference in performance between what can be achieved by the ensemble, and what can be achieved by individuals on their own, when there is an accounting for the relevant trade-offs and tensions. This measure could be the difference between all or nothing if the emergent functionality of the ensemble is absent for individuals, as in the case of cognition or the collective transport of an object too large for an individual to carry alone. Or it could be the difference in level of achievement, such as how frequently a threat is correctly identified, or a resource is successfully discovered, when individuals have limited sensing and the environment is uncertain.

Corresponding author:

Simon A. Levin, Department of Ecology & Evolutionary Biology, Princeton University, 106A Guyot Hall, Princeton, NJ 08544-1003, US.

Email: slevin@princeton.edu



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>) which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

Measuring collective intelligence in terms of performance requires that performance be defined to reflect relevant trade-offs and tensions. For example, when there are multiple critical tasks to be simultaneously managed, like identifying threats and discovering resource, and these compete for the limited energy and attention of individuals, performance should be defined in terms of the relative benefits and relative costs of every task. Similarly, for an ensemble in an uncertain or variable environment, the measure of performance should reflect how well a balance can be maintained between the robustness of the group's stability, that is, its ability to reject unimportant fluctuation or events, and its flexibility (adaptability), that is, its ability to respond to important signals or events, even if weak or rare. In neural systems, this is called the stability–flexibility dilemma (Liljenström, 2003).

In designed ensembles, in which rules for control of individual elements can be imposed to improve performance, the level of centralization of control is key. A centralized approach to optimizing coordinating control is possible when a central unit has access to all measurements and can decide for, and direct, all the elements of the ensemble, as in an air-traffic-control network (Gopalakrishnan and Balakrishnan, 2021) and in the coordination of the many actuators in an artificial hand (Matrone et al., 2010; Piazza et al., 2019). In the centralized case, collective intelligence may be maximized, at least locally, subject to tractability of the optimization problem, physical limitations of the system, and uncertainty about the system and the environment. However, centralization can become costly for large-scale systems and problematic when communication between individuals and the central unit is unreliable or the central unit experiences a failure.

In the case when centralized control is not possible or not practical, control can instead be decentralized across the many elements of the ensemble, and collective intelligence can emerge from local rules much as in ensembles in biology. For decentralized ensembles, collective intelligence is a public good, in the sense that it improves system-level performance relative to what can be achieved by individuals on their own. This is the case, for instance, for a team of autonomous mobile robots (Bullo et al., 2009; Parker et al., 2016). The team might be homogeneous, as for the construction robots in (Werfel et al., 2014). Homogeneity can provide robustness to failure since agents are interchangeable, and homogeneity does not preclude individuals taking on different roles, although facilitating differentiation may require artificial incentives. Alternatively, the team might be heterogeneous, as for the underwater archaeology robots of (Allotta et al., 2015), where, for instance, coordination of agents with heterogeneous means of sensing, actuation, and movement can serve the public good. Information sharing, coordination, and/or cooperation (Bizyaeva et al., 2022; Cao et al., 2013; Knorn et al., 2016;

Nedic et al., 2018; Olfati-Saber et al., 2007) will typically be necessary for collective intelligence, for example, in adaptive sampling of spatiotemporal processes (Paley and Wolek, 2020), explore-exploit decision-making (Kalathil et al., 2014; Madhushani and Leonard 2019; Landgren et al., 2021), and collective transport (Kube and Bonabeau, 2000; Farivarnejad and Berman, 2022).

In biological ensembles, in which evolutionary forces act to improve performance, the level of selection is key. In evolution, differentiation of function raises game-theoretic considerations since in general not all potential roles will confer the same payoffs. Differentiation may thus only be expected to occur through top-down enforcement, through frequency-dependent stabilization of payoffs, or through mechanisms like revenue-sharing that reward individuals that engage in what would otherwise be lower payoffs. When the level of selection is at the level of the ensemble, as for example, in the evolution of brain function in which the agents are individual neurons, one can expect that collective intelligence will be maximized, at least in a local sense mathematically, subject to constraints.

When the primary level of selection is at the level of individual agents, however, the situation is different. Bird flocks and fish schools provide examples. Collective intelligence in such circumstances is a public good similar to that for decentralized ensembles in design. As is often true for a public good in a social context, it may be the case for evolved systems (as well as for designed systems) that the only achievable Nash equilibria engender lower payoffs than the social optima, although second-best solutions may still provide improvements over what individuals might realize on their own. Key in such cases is the emergence of cooperation. That cooperation can take multiple forms; in animal aggregation, for example, the issues include whether or not to join a collective (Akçay et al., 2012; Pulliam et al., 1982; Axelrod and Hamilton, 1981; Pulliam, 1973), as well as matters of sharing of information (Torney et al., 2011), coordination, collaboration, and positioning within the aggregate (Hamilton, 1971).

We explore these and related issues in design and in biology in the rest of this paper. We show how the same measures of collective intelligence and the same models and methods for understanding collective intelligence can often be useful in both domains.

Collective intelligence and design

Collective intelligence can significantly enhance the performance of technological systems comprised of many individual units, from power grids (Liu et al., 2022), wind-turbine farms (Shapiro et al., 2022), and industrial process control (Christofides et al., 2013), to air traffic control (Gopalakrishnan and Balakrishnan, 2021) and teams of autonomous robotic vehicles (Bullo et al., 2009; Parker et al.,

2016). Teams of robots have enormous potential to address challenging problems in complex environments. Already, they are being deployed, often in collaboration with human partners, in a broad range of settings, including to search and remediate a polluted environment (Zahugi et al., 2012), identify and extinguish forest fires (Marjovi et al., 2009), monitor animal populations and their habitats (Park et al., 2019a), and provide search and rescue (Murphy et al., 2008; Queralt et al., 2020).

Level of centralization of control

Ideally, in design, an omniscient, centralized decision-making unit could solve for an evolving optimal solution for the system as a whole and direct the actions of each individual agent over time. This is the case in air-traffic-control systems where centralization is critical (Gopalakrishnan and Balakrishnan, 2021). However, in a great many applications, a centralized controller is often neither practical nor desirable. For a team of robots distributed in an uncertain environment, on land, in the ocean, or in the air, where sensing, communication, and energy are limited, relying on a single orchestrating unit may introduce delays and inaccuracies and compromise robustness to uncertainty and malfunction.

Instead, in the decentralized setting where individuals have agency but may have limited (e.g., only local) access to information about the environment and others in the group, each individual agent should make its own decisions in response to what it observes about the environment and what it learns from others. Learning about others could be acquired directly, for example, by sensing the activities of nearby neighbors or sharing information over a communication network, or it could be acquired indirectly, for example, through observing the influence of others on the environment. There is a large and growing literature (Bizyaeva et al., 2022; Bullo et al., 2009; Cao et al., 2013; Christofides et al., 2013; Knorn et al., 2016; Nedic et al., 2018; Olfati-Saber et al., 2007; Zhang et al., 2021) on design methodologies for decentralized groups that provide rules that individual agents should follow to make their own choices on-the-fly that account for environmental cues and the reward, state, and/or action of other agents, in the service of a system-level goal. Such a goal might be fast, accurate, and reliable estimation of a process from noisy measurements, coverage or search over a spatially distributed resource, decision-making among alternative options or strategies, dynamic allocation across tasks, synchronization of oscillatory activity, or stabilization of motion patterns and formations.

These design methodologies come with guarantees on group-level (and individual-level) performance, for example, accuracy, efficiency, and reliability, as a function of parameters that characterize the environment, such as

likelihood of an event, magnitude of uncertainty, or variability in resource landscape, as well as parameters that characterize the system, such as number of agents, differences among agents, and inter-agent network structure, that is, who is sensing or communicating with whom. The guarantees are derived from mathematical analyses of the influence of these parameters on the evolving behavior, possibly across multiple scales.

The analytical results can be further used to systematically formulate rules for individuals to adjust parameters (like dials) inside their behavioral response rules so that the system as a whole adapts as circumstances change. Formulating rules and guarantees can be done using mathematical models, and especially those that also help explain how biological groups adapt to changing environments, such as the harvester ants that regulate their foraging in response to changes in temperature and humidity in the desert (Gordon, 2013; Pagliara et al., 2018). In the multi-agent dynamic opinion model of Bizyaeva et al. (2022), discussed further in *Recent models for explaining and shaping collective intelligence in biology and design*, a rule is designed for the individual to modulate its attention (susceptibility) to the opinions of others, in response to its observations of how others are changing their engagement in the opinion-forming process. The analysis shows how parameters in this rule can be used to tune the sensitivity of the collective response to an external stimulus, that is, to adjust implicit thresholds that determine how large in magnitude or long-lasting an external stimulus detected by one or more individuals will need to be to elicit a group response.

Example: Decentralized estimation and learning of an environmental signal

The decentralized estimation problem (Speyer, 1979; Olfati-Saber, 2005; Spanos et al., 2005; Carli et al., 2008; Park and Martins, 2017; Nedić et al., 2018) illustrates well how an ensemble of individuals that share information, even in a limited capacity, can outperform those same individuals acting independently, in their ability to estimate a multi-dimensional signal from noisy measurements. The estimate for each individual in the ensemble will be more accurate because of what it learns from others; and even if it only measures part of the signal, it can still achieve an accurate estimate of the full signal, something it could not do on its own. Decentralized estimation is foundational for a variety of problems, like opinion-forming for decision-making in groups (Bizyaeva et al., 2022), when individuals may otherwise, on their own, know little about some of the options. The decentralized estimation problem also illustrates how a decentralized approach can be as good as

a centralized approach, but without the costs of centralization.

Consider the case in which each individual agent takes a noisy measurement, regularly over time t , of some part of a multi-dimensional signal \mathbf{x} , which represents stimuli in the environment that evolve over time t according to a known linear dynamic process. Assume that together the individuals measure the full signal \mathbf{x} . When the agents act independently, each can filter its own noisy measurements using a Kalman filter (Anderson and Moore, 1979) to get the most accurate estimate of \mathbf{x} , at least in part, given what it has measured. If, instead, there were a central unit with access to the noisy measurements from all the agents, a Kalman filter of the pooled measurements would provide an estimate of \mathbf{x} in full and more accurate than what any individual could do on its own.

In the decentralized case, to recover the improved accuracy and full estimate achieved in the centralized case, the centralized pooling of information can be replaced by the limited sharing of information over a communication network. In the approach of Olfati-Saber (2005), agents may only be able to share with a small number of other agents, but the result is guaranteed as long as the communication network is connected, that is, there is a path through the network between every pair of agents, and communication updates are sufficiently frequent. In this approach, each agent performs its own local Kalman filter and leverages the fusing of measurements and uncertainties obtained from a consensus filter, which is a process of averaging information over the network through individual communication updates. Spanos et al. (2005) showed how the rate of convergence of the estimate depends on the network structure. The approach of Olfati-Saber (2005) has been generalized in various ways, including to reduce the computational complexity of the distributed estimation algorithm for large-scale systems with a very high-dimensional signal \mathbf{x} (Khan and Moura, 2008) and to reduce the costs of frequent information sharing in networks with a large number of nodes (Park and Martins, 2017).

In a related decentralized problem, individuals share information over a network so that they all learn about an environmental signal, such as a route to be tracked or the direction of an approaching threat, with the difference being that only a small number of individuals can take measurements of the signal. The individuals that do not sense the signal rely on noisy observations of, or communication with, some of the others (their neighbors in the network). How accurately the group learns the signal depends on where in the network the individuals are that sense the signal—the best choice of the individuals that do the sensing is the solution of an optimization problem (Patterson and Bamieh, 2010; Clark et al., 2014; Lin et al., 2014). In Fitch

and Leonard (2016), it was shown that the individual with the highest “information centrality,” a measure derived from the network structure alone, is the optimal choice for a single sensing agent, and a measure of “joint centrality” determines the optimal choice for more than one sensing agent. If the network structure is not known, and any agent could in principle participate in the sensing, then in the decentralized setting, rules would be defined based on local information for individuals to determine if they should turn on their sensor. Or, if only a fixed subset of agents has sensing capability, then agents should have rules, which depend on local information, to change network connections so that these agents become more central, much like local rules designed to ensure network connectivity (Ando et al., 1999; Bullo et al., 2009) or to maximize robustness (Young et al., 2011).

Measures of collective intelligence

A suitable measure of collective intelligence in multi-agent technological systems, as in the case of ensembles of organisms, is the difference in performance between what can be achieved by the group of individual agents, and what can be achieved by the individuals on their own, where performance accounts for trade-offs and tensions. A related useful measure is the difference in performance between what can be achieved by the system when there is centralization and what can be achieved with decentralization. It then becomes a matter of defining measures of performance, which will depend on objectives and context.

An important problem for collective intelligence in decentralized multi-agent systems is in the coordinated monitoring of spatiotemporal processes (Paley and Wolek, 2020), such as pollution, forest fires, and animal populations. Fundamental to this problem is the distributed positioning of the agents so that they “cover” the region of interest (Cortes et al., 2004). Let q be a location in the region, p_i the location of agent i and $d(q, p_i)$ the distance between q and p_i . A larger distance $d(q, p_i)$ implies a greater degradation in the ability of agent i to monitor location q . Suppose that the information or probability of an event is uniformly distributed over the region. Then, a measure of coverage is the sum of the $d(q, p_i)$ over all locations q and all agents i ; the measure can be modified for nonuniform distributions by weighting the distance to q by the value of the distribution at q (Cortes et al., 2004) or a nonuniform distance metric can be applied (Lekien and Leonard, 2009). The optimal coverage problem is then to solve for the positions of the agents p_i that minimize the coverage metric, that is, that minimize the distances between agents and locations in the region. A uniformly (nonuniformly) distributed resource will yield a uniform (nonuniform) distribution of agents. In the decentralized approach of Cortes

et al. (2004), the optimal solution is achieved on-the-fly if each agent moves to the centroid of the domain comprised of locations q closest to it. Each agent can determine its “domain of dominance” merely from observations of the relative locations of its nearest neighbors.

In the related problem of adaptive sampling, such as a team of mobile robot sensors mapping the temperature, salinity, or concentration of pollutant over a fixed volume of the ocean, the goal is for the agents to move so that the samples (measurements) they take along the way minimize the uncertainty in the system-wide map over time and space (Leonard et al., 2007). The measure of uncertainty, which is inversely related to a measure of information in the data collected, can be defined much like the coverage metric. Because uncertainty decreases at the location where a sample is taken, the agent should move away from that location once it has taken the sample. If the field to be mapped changes over time, the uncertainty will grow back over time at any location that has been sampled and an agent will need to come back to that location after some time. Local rules that keep agents distributed over time and space, consistent with the temporal and spatial scales of the distributed field, minimize the uncertainty metric (and thereby maximize the information metric). In a field demonstration of a network of autonomous underwater gliders mapping temperature and salinity (Leonard et al., 2010), the decentralized rules for stabilizing motion patterns of agents over time and space (Sepulchre et al., 2007, 2008), to be consistent with the spatial and temporal scales of temperature distribution, derive from an extension of coupled oscillator dynamics.

In these coverage and mapping problems, collective intelligence helps ensure performance even as conditions change. For instance, in the adaptive sampling problem, if a robot is removed, the remaining robots should move to fill in the gap without sacrificing too much data elsewhere. Collective actions such as these are not easily accomplished by agents on their own, notably because of the lack of coordination. The metric provides a means to evaluate how well the robots manage their collective actions.

When agents monitor an uncertain environment in search of resource or reward, for example, robots finding and remediating pollutant, their collective intelligence can be measured by how well they jointly manage the explore-versus-exploit trade-off (Robbins, 1952). Exploiting is sampling a patch known to yield high reward. Exploring is sampling to find and utilize an unknown patch that might yield even higher reward. The most rewarding exploitation requires sufficient exploration, but too much exploration can slow down accumulation of reward. There is a vast literature on the multi-armed bandit (MAB) problem, which provides a mathematical formulation of the explore-exploit problem for a single decision-maker (Lai and Robbins, 1985; Auer et al., 2002). Kaufmann et al. (2012)

introduced the Bayesian MAB, where the decision-maker might have priors and use these in its decision-making according to Bayes’ rule. Reverdy et al. (2014) extended this to a Bayesian decision-maker that uses a soft-max (stochastic) decision rule for exploration with a cooling schedule that tunes the amount of stochasticity (like simulated annealing) and incurs costs for switching from one patch to another. In Reverdy et al. (2017), the decision-maker addresses an objective function that sacrifices rather than maximizes reward. Hills et al. (2015) provide a review of the explore-exploit trade-off in a wide diversity of contexts, from memory search to cultural innovation.

In the MAB problem, performance is measured by regret, which is the loss in reward that comes with sampling a suboptimal rather than an optimal option. So, for a group, minimizing the sum of each agent’s cumulative regret is equivalent to maximizing the total cumulative reward. An individual can be deployed to minimize its own accumulated regret. However, collective intelligence is needed for the agents to minimize total regret, for example, by sharing information to learn the environment better or taking on different roles to leverage differences in their capabilities or location in the network. Total cumulative regret is a measure of collective intelligence that provides a means to evaluate how well the agents manage their collective decisions. The measure can be augmented to include costs of communication (Madhushani and Leonard, 2020a).

To investigate the opportunities for collective intelligence in a group of explore-exploit decision-makers, the MAB has been derived for multiple decentralized agents that are faced with the same explore-exploit problem and seek to maximize their accumulation of reward, but take advantage of what they can learn from others. In Anandkumar et al. (2011); Gai and Krishnamachari (2014); Kalathil et al. (2014), the agents do not communicate but rather learn from others indirectly, such as when reward at an option is reduced because others have selected the same option. In Landgren et al. (2016, 2021); Martí nez-Rubio et al. (2019), agents can share estimates through a communication network, whereas in Chakraborty et al., 2017; Kolla et al., 2018; Madhushani and Leonard, 2019, 2020a, 2021, agents can only observe the choices and value of rewards received by their neighbors. Analytical results include a derivation of explore-exploit centrality measures (Landgren et al., 2016, 2021), based solely on network structure, that predict which network structures yield lower total regret and which individuals obtain lower regret by virtue of their location in the network. For stochastic interactions, agents that frequently observe neighbors do well if those neighbors make infrequent observations of others since then the neighbors do a lot of exploring, which benefits those who observe them (Madhushani and Leonard, 2019).

Other measures of collective intelligence are likewise defined in case of other objectives, trade-offs, costs, and benefits. Costs associated with risks to safety, energy consumption, computation, and communication can be included in general as warranted. When a group is to quickly and efficiently coordinate decision-making or actions to carry out a task together, performance, and thus collective intelligence, can be measured by how well the group manages the trade-off between speed and accuracy in the accomplishment of the task. The speed-versus-accuracy trade-off for collective decision-making in decentralized systems, and the influence of network structure, has been studied in, for example, [Srivastava and Leonard \(2014\)](#); [Valentini et al. \(2016\)](#).

For systems operating in a complex, changing environment, it can be critical for agents to manage the trade-off between being robustly stable to unimportant fluctuations, disturbances, and false-positives and responsive and adaptable to important environmental cues, even if they are weak or rare. The trade-off between robustness and adaptability is fundamental to many fields, including control and decision-making ([Bizyaeva et al., 2022](#)) and on-line learning ([Fukushima et al., 2021](#)); in neural systems, it is called the stability–flexibility dilemma ([Liljenström, 2003](#)). Collective intelligence for managing this trade-off can be measured in various ways. [Zhong et al. \(2021\)](#) study how a group of agents responds to an external stimulus, when agents share information over multiple network layers, each representing a different sensing modality, and when they differ in how readily they react. The trade-off is measured as the difference between the benefit of responding to true-positives and the cost of responding to false-positives. The methodology determines where in the given multi-layer network to locate more readily responsive individuals and where to locate less readily responsive individuals so that at the level of the group the measure is optimized ([Zhong et al., 2021](#)).

Cooperation and coordination

That agents have rules that enable them to come to a mutually cooperative or coordinated solution is often a critical part of collective intelligence, and especially so in the decentralized setting where agents do not have access to full information about the system, and may have conflicting individual goals, such as when each has its own rewarding task that it can do on its own ([Fisac et al., 2019](#); [Marden and Shamma, 2018](#); [Menache and Ozdaglar, 2011](#); [Parise and Ozdaglar, 2021](#)). The rules that govern the actions of agents may derive from objective functions that make them look much like a biological system with selection at the level of the individual with strategies that evolve over time. An important difference is the opportunity in design to impose coordination or cooperation, for example, by facilitating

explicit communication among agents or context-dependent “altruism,” such as when an agent that communicates with a large number of other agents is incentivized to do more exploring for the benefit of the group at a cost to its own reward ([Madhushani and Leonard, 2021](#)). This has a strong analogy with similar trade-offs in animal behavior. The design problem becomes more challenging when designed agents partner with people ([Dresner and Stone, 2008](#); [Mirsky et al., 2021](#); [Nikolaidis et al., 2017](#)), for example, in human–robot interactive tasks such as on a construction site, and the robots need rules to elicit cooperation or coordinated responses from their human partners.

A useful measure of collective intelligence in these contexts is the size of the region of attraction of the dynamics to a mutually cooperative or coordinated solution that benefits individuals and/or yields the highest payoff to the group. The region of attraction of an equilibrium is the set of all initial conditions that lead the system to converge over time to the equilibrium. The larger is the region of attraction for the cooperative or coordinated solution the more likely it is that the agents will converge on that solution. How local rules of response and interaction tune this region of attraction is a focus of the derivations and investigations of the analytically tractable model of multi-agent opinion dynamics presented in ([Bizyaeva et al., 2022](#)), which is discussed further in *Recent models for explaining and shaping collective intelligence in biology and design*.

In the model of [Bizyaeva et al. \(2022\)](#), individuals use a nonlinear rule for updating their opinions in response to observations of their neighbors’ opinions. This nonlinearity allows for the emergence of multi-stability—that is, more than one equilibrium is stable for the same set of parameter values. There can be multi-stability of the equilibria corresponding to all agents agreeing on one of the available options as the preferred option. This is useful for consensus decision-making since if there is no evidence to distinguish options, the agents will randomly select one option to agree on, rather than getting stuck in a deadlock about which to choose. And if there is evidence or internal bias that distinguishes one option over the others, the agents will agree on that option. Likewise, there can be multi-stability of the equilibria corresponding to disagreement among agents as to which options they prefer. This can be useful for task allocation since agents will distribute themselves over the tasks whether or not there is evidence or internal biases to help distinguish which agents should attend to which tasks ([Franci et al., 2021](#)).

The multi-stability emerges when the level of attention individuals pay to the opinions of others grows above a critical threshold. The relative size of the region of attraction for each of the stable solutions can be shaped by design parameters (dials), including relative attention levels,

network structure, and individual differences in response, such as whether an individual is attracted to, or repelled from, the opinions of its neighbors.

These nonlinear opinion dynamics and the opportunity to shape and grow the region of attraction of a coordinated or cooperative solution has been recently studied in the context of multi-agent games, when the game is repeated over time and there is the opportunity for agents to observe the opinions (strategies) of at least some of their opponents (Park et al., 2022). The approach addresses social dilemmas like the prisoner's dilemma and the public goods game, in which a sufficient level of cooperation is necessary for a rewarding outcome for all and avoidance of the non-cooperative Nash equilibrium, which, in the public goods game corresponds to no individual making an investment. The approach also addresses coordination problems like the stag-hunt game and the battle-of-the-sexes, where the payoff-dominant coordinated solution may be risky and/or have a small region of attraction.

In the dynamic modeling approach of Park et al. (2022), strategies are options and an agent's relative opinions about those options map to mixed strategies. The opinion update rule, which is driven by payoffs, and by observations of the opinions of others, can be shown to provide reciprocity, functioning much like a Tit-for-Tat strategy, which is well-known to lead to mutual cooperation (Axelrod, 1984). In games like the prisoner's dilemma and the public goods game, where mutual defection is the only Nash equilibrium, it is shown in (Park et al., 2022) how mutual cooperation emerges in the model as a second stable equilibrium for large enough attention to others and how the region of attraction to mutual cooperation grows with growing attention. Similarly, in coordination games, the region of attraction of coordination increases with increasing attention and other parameters. The level of an individual's attention can be interpreted as its expectation of the durability of the interactions, well-known to be a key determinant of cooperative behavior (Axelrod, 1984). For multiple agents making observations over a network, for example, in the public goods game, it is shown analytically how the network structure, that is, who is paying attention to the opinions/strategies of whom, can be varied to modify the region of attraction of a mutually cooperative solution (Park et al., 2022).

Cooperation can be designed to improve system-level performance, but it can introduce costs, such as those associated with communication. There are also scenarios in which information-sharing can diminish system-level performance, for example, when exploration is hampered by too much information-sharing about optimal options (Madhushani and Leonard, 2020b). And there are scenarios in which too much imposed altruism can lead to a deadlock, for example, when a set of mobile robots, driven by their individual objective functions,

find themselves headed for a collision and so they come to a stop. Thus, the most useful measures of collective intelligence account not only for the benefits of cooperation, but also for its costs, including hesitancy and deadlock. Deadlock (indecision) is overcome in (Biziyaeva et al., 2022), even if there is no evidence or bias to distinguish options, when the attention individuals pay to others in the group exceeds a critical threshold. The mechanism is inspired by the collective behavior of house hunting honeybees (Seeley et al., 2012; Pais et al., 2013) and animal groups on the move (Couzin et al., 2005; Nabet et al., 2009; Leonard et al., 2012).

Collective intelligence and biology

Collective intelligence is observed in biological ensembles at all scales wherein the group performs better than what individuals can do on their own. When evolutionary forces act to improve performance, the level of selection is key.

Emergence and levels of selection

Even in the ontogenetic development of an organism, natural selection must operate on the rules that guide development, rather than printing out end products according to blueprints; hence, errors and multiple end results are possible. Theories of development thus focus on how local rules of interaction can give rise to global patterns, and the influence of the local context on development (Turing, 1952; Waddington, 1957; Wolpert, 1969). Waddington's developmental landscape makes clear that, at the level of individual cells and tissues, there are decision points that determine the future fate of the tissue among alternative pathways; but there are higher-order constraints that regulate the relative proportions of different kinds of end-products, such as organs and appendages. There is differentiation of function among the genetically identical cells.

For ensembles of genetically distinct individuals, the individual agents are likely to respond more selfishly to payoffs, and hence, the collective performance will suffer. With the possible exception of selection acting on groups of genetically related individuals, such as the social insects (Hamilton, 1964, 1964b), natural selection for group performance will in general be weaker than selection for individual performance; nonetheless, for example, through reciprocal altruism, individual selection should favor behaviors, such as the tendency to cooperate with others, that will ultimately improve individual fitness. Collective intelligence and other collective properties can then emerge from those individual interactions, just as in the development of an individual organism from the behaviors of its cells; but top-down

constraints are likely to be looser in such cases than for individual development, and the potential for alternative outcomes greater.

Collective intelligence and selection at the system level

In the development of an organism, local rules give rise to global patterns. Those global patterns are the objective functions of natural selection, which must shape the global patterns by shaping the local rules. Much like in adaptive control (Åström and Wittenmark, 1995), opinion dynamics with time-varying attention (Bizyaeva et al., 2022), and other forms of iterated learning in artificial systems, where feedback laws are updated over time, local rules that give rise to more favorable outcomes are reinforced over evolutionary time, gradually improving the eventual outcome and making the ontogenetic process more reliable. In this way, reliable systems are built, involving cooperation among individual units, as for example, in brain function. System-level selection hence can impose local rules, though these are not in general uniquely determined.

Kauffman and Levin (Kauffman & Levin, 1987) illustrate this point by imitating an evolutionary dynamic to solve the traveling salesman problem, rewarding mutations and rearrangements that lead to shorter paths, and show that the dynamic can find any of multiple local optima; mechanisms analogous to simulated annealing or stochastic resonance (Levin and Miller, 1996; Gammaitoni et al., 1998) may allow more favorable optima to be found, but no guarantee of global optimality exists. In biological evolution, the number of such local optima may be very large, encompassing the rich diversity of ways organisms solve similar problems. Given sufficient geographical isolation, different outcomes may be realized in distinct areas, though species invasions, for example, due to the activities of humans can lead to displacement of the resident species, often because of the absence of the natural enemies of the invaders. Economists and political scientists explore similar phenomena, called Tiebout migration (Tiebout, 1956), in which individuals move among communities to select those most aligned with their preferences. In both cases, the global optimum may never be reached, especially if the differences in benefits realized among local optima are small.

This process is fundamentally different from say the emergence of trophic networks, nutrient cycling, or other system-level properties of ecosystems. In these cases, the primary selective pressures are at levels below the whole system, and the study of the emergence of cooperation, for example, in the maintenance of public goods such as critical nutrients, requires a game-theoretic approach rather than optimization. One can of course ask what the collective optimum would be, but one should expect the realized level of cooperation to result in something short of that optimum.

As an example, plants in a forest compete as well as cooperate through their above- and below-ground branching structures, sharing nutrients through biogeochemical cycling, while competing through shading, root grafting, and allelochemicals. These considerations apply as well to the situations discussed in the previous paragraph, once it is recognized that the movements of individuals among localities alters the distribution of fitnesses or utilities, due to density or frequency effects, as in the El Farol problem (Arthur, 1994). This means that in these cases as well, optimization considerations must give way to ones couched in game theory, including modern advances in the theory of mean-field games (Lasry and Lions, 2006; Guéant et al., 2011; Nourian et al., 2011; Huang et al., 2010).

Thus, in any ecological community, intelligence can be measured at various levels. Most species could not survive without “collective intelligence” as manifest through co-evolutionary processes. Indeed, collective intelligence is not limited to multispecies assemblages, but perhaps is even more obvious within species. Bacteria rely on “quorum sensing” to form biofilms, producing extracellular polymers that provide signals of abundance, and matrices for growth (Miller and Bassler, 2001; Nadell et al., 2008b; Drescher et al., 2014). Producing those polymers is costly, but essential for biofilm formation, introducing game-theoretic issues (Nadell et al., 2008a). Similar trade-offs can be found at all levels of biological organization, including especially aggregations of animals from insects to bird flocks and fish schools (Levin, 2014). We turn to consideration of such assemblages in the next section.

Collective intelligence and selection at the individual level

A wide range of animals live and move in groups, at least some of the time, even when selection is at the level of the individual (Parrish and Hammer, 1997; Krause and Ruxton, 2002; Sumpter, 2010). This implies that even a self-interested individual can benefit more from being part of the group than it could on its own (Pulliam, 1973). Groups of socially interacting individuals can do better than solitary individuals in various ways, including, for example, in maintaining vigilance for predators, foraging for food, migrating, conserving energy, heat, and water, searching for a mate, and reducing risk (Krause and Ruxton, 2002). That animals living in groups often manage all of these tasks and others, despite the trade-offs that arise from limited resources, for example, to search for predators versus to search for food (Rubenstein, 1978), provides evidence for a high level of collective intelligence.

However, the collective intelligence that makes group advantage possible can also impose costs on individuals. For example, the cohesive movement of a group is enabled only if individuals in the group invest effort in observing

the relative motion of their neighbors. Typically, there is uncertainty in the observations and in the decisions the animals make in response to what they observe. Cohesiveness of the group is more robust to uncertainty when individuals observe a greater number of neighbors, but this comes at a greater cost to the individuals. Costs can be expected to increase in more complex systems, where managing uncertainty in maintaining cohesiveness is just one of many critical tasks to be accomplished. In aggregations with regular arrangement, these costs may be uniformly distributed, as in the synchronized movement of a starling flock, where it has been argued that each starling pays attention to the same number (six or seven) of its nearest neighbors, independent of flock density (Ballerini et al., 2008; Bialek et al., 2012), at least under specific conditions. Yet even if the cost is the same for each individual, when selection is at the individual level, an exhibited collective intelligence suggests that the benefit to the individual from being part of the group outweighs the cost.

The data collected on starling flocks (Ballerini et al., 2008), moving in the absence of a threat, were analyzed in (Young et al., 2013) using an analytically tractable mathematical model of consensus-forming under uncertainty. The goal was to investigate the role of number of nearest neighbors attended to by each starling in the trade-off between the benefit to the flock in robustness to uncertainty and the cost to the individual in effort. A useful measure of collective intelligence, which accounts for the benefit and cost in this setting, is the per-neighbor contribution of an average bird in the flock to the robustness to uncertainty in the cohesiveness of the flock's motion. This can be interpreted as a public-goods game where robust cohesiveness of the group is the public good. Young et al. (2013) showed for the measured positions of between 440 and 2600 starlings in each of nearly 400 datasets, assuming each bird pays attention to the same number of nearest neighbors, that this measure of collective intelligence is maximized when that number is six or seven, the same number found by Ballerini et al. (2008).

While robustness to uncertainty benefits the group, too much of it can negatively impact the benefits of other kinds of collective intelligence, notably the ability of the group to respond rapidly to new environmental cues such as the detection of a predator or the discovery of a high-quality food source. This is much like the stability–flexibility dilemma of neural systems as discussed in *Collective intelligence and design*, although it can also be formulated as an explore-exploit trade-off (Cohen et al., 2007). For example, the number of neighbors attended to by each individual in the group for maximal robustness of cohesiveness (stability) to uncertainty may limit the flexibility needed for individuals and the group as a whole to be responsive to an environmental cue. Using an evolutionary

dynamic model with fitness defined at the level of the individual, Brush et al. (2016) showed how context matters in the evolution of individual behaviors, such as the number of neighbors attended to by individuals, and the balance of different kinds of collective intelligence. They show that the emergent optimal number of neighbors tracked varies greatly with the task at hand, for example, from foraging to predator avoidance or defense.

There are also costs to those individuals in the group that invest in the environmental cue, for example, in sensing the predator, identifying the migration route, or finding the high-quality patch of food. All members of a socially interacting group benefit when their collective intelligence yields effective group-level anti-predator vigilance, migration, and foraging. Yet, to achieve a level of collective intelligence that yields net reward, at least some of the (self-interested) individuals should invest in the cue rather than free-ride. The trade-off can be represented by the public goods game, as discussed in *Cooperation and coordination*. The problem is also closely related to the collective learning of an environmental signal discussed in *Example: Decentralized estimation and learning of an environmental signal*.

Observations of collective intelligence of animal groups in the wild tell us that this kind of cooperation exists. Just how much cooperation there is depends on environmental conditions. For example, how many individuals interrupt their feeding to scan for threats can depend on group size as well as on the abundance of food and the extent of obstacles to vigilance found in the group's habitat; see, for example, for birds (Beauchamp, 2008), Mongolian gazelle (Zhang et al., 2018), and plains zebra (Chen et al., 2021).

Likewise, the influence of environmental conditions on the evolution of collective migration of socially interacting groups in motion and the strategies individuals adopt for acquiring information were studied using numerical simulations in (Guttal and Couzin, 2010) and analytically tractable mathematical models in (Torney et al., 2010) and (Pais and Leonard, 2014). In these works, specialization of the migrating population into those that invest in the environmental cue and those that rely on social cues was shown to be a stable evolutionary outcome, meaning that it resists invasion by more exploitative strategies. But this kind of collective intelligence is fragile, which can be related to the sensitivity to environmental conditions of the size of the region of attraction for the cooperative solution. When costs are too high, that is, in a fragmented environment, there may not be enough individual investment in the cue and the group will lose the ability to migrate. Even worse, the models predict a hysteretic effect in the loss and recovery of migration ability (Pais and Leonard, 2014), that is, recovery requires a reduction of costs below the point at which migration ability was lost.

In an examination of cooperation in foraging populations (Torney et al., 2011), it was shown how a form of reciprocity

among communicating individuals, in which some individuals cooperatively signal to others when they find a resource, is an evolutionary stable solution for certain distributions of resource in the environment. The multi-agent multi-armed bandit framework, described in *Measures of collective intelligence*, has been defined and used to study social foraging, and the social explore-exploit trade-off more generally, when individuals are self-interested but cooperate through the sharing of information (Chakraborty et al., 2017; Kolla et al., 2018; Landgren et al., 2016, 2021, Madhushani and Leonard, 2019, 2020a, 2021; Martínez-Rubio et al., 2019). For example, Madhushani and Leonard (2020a) showed how dynamic signaling strategies can provide low-cost improvement to collective foraging.

The measure of collective intelligence is closely related to the question of whether individuals should join a group, and the related question of how any excess benefits to individuals in a group should be apportioned among the members. The simplest manifestation of this puzzle is that of two individuals, who must decide whether to cooperate in carrying out a task, or go it alone (Akçay et al., 2012). More broadly, however, this raises the issue of what constitutes an optimal group size (Hamilton, 1971; Brown, 1982; Pulliam and Caraco, 1984), a key trade-off being the possible conflict between the fitness advantages to an individual if it joins the group and the possible negative implications for current members of the group.

Transformational evolution, cultural evolution, and emergence of collective intelligence

The most familiar and robust mechanism shaping behaviors is evolution through natural selection operating on individuals, but cultural group selection is sometimes claimed as a complementary mechanism that selects for operative behaviors (Wilson and Wilson, 2007). Extreme care must be exercised in invoking such arguments to the extent that they rely upon genotypic mechanisms, because selection is much stronger at levels below the groups to which individuals belong, and because the lack of group integrity can easily undermine selection for group properties (Maynard Smith, 1974; Maynard Smith and Warren, 1982). However, transformational evolution (Lewontin, 1977, 1978, 1985), similar to what Lenton calls sequential selection (Lenton et al., 2018), can also affect what we observe in Nature. This is not evolution in the traditional sense of fitter agents leaving more offspring that carry their genes, but rather a filtering process over ecological time, in which units (like cultural groups) with favorable properties will last longer, and are thus more likely to be observed (Levin, 1999). In that collective intelligence can confer collective benefits to group members, we can expect that it may also emerge through a transformational process. The importance of group selection in cultural evolution

remains one of the most debated in evolutionary biology because of disagreement about the strength of selection processes at the group level (Fracchia and Lewontin, 1999) and is beyond the scope of this paper.

Recent models for explaining and shaping collective intelligence in biology and design

Measuring and assessing collective intelligence in both natural and technological groups requires a deep understanding of local rules, how they may evolve, how they can enable different kinds of global patterns, and how they may address the many trade-offs that arise both for the individual and for the group. Mathematical models of local rules that define the interactions of individuals with one another, and with the external environment, have proved enormously useful in uncovering principled explanations for collective intelligence and testable predictions for the sensitivity of collective intelligence to differentiation among individuals and possibly changing environmental conditions, and means for actively shaping collective intelligence; see, for example, Couzin et al. (2005); Sumpter (2010); Bialek et al., 2012; Seeley et al., 2012; Leonard (2014); Pagliara et al. (2018); Marden et al. (2009); Acemoglu and Ozdaglar (2011); Jia et al. (2015); and Ye et al. (2021).

Models well suited to this agenda allow for a wide range of alternative outcomes and the possibility for rapid and reliable collective response to even very weak stimuli or asymmetry. One such model, introduced in *Collective intelligence and design*, is given by the nonlinear multi-agent, multi-option opinion dynamics defined and studied in (Bizyaeva et al., 2022). This model, inspired in part by biophysical models of neuronal networks (Wilson and Cowan, 1972; Hopfield, 1982), supports multiple spatial and temporal scales, complex network interconnections, as well as cooperative and competitive interactions. The model exhibits nonlinear responsiveness and tunable sensitivity to input and uncertainty, allowing for a rigorous examination of the stability–flexibility trade-off. The model describes a fundamental mechanism for how agents that exchange information over a network can form an opinion, that is, avoiding indecision and deadlock even when all options are perceived as equal, and robustly come to one of multiple stable equilibria corresponding to agreement or disagreement solutions. Moreover, the model is analytically tractable, which is critical to a comprehensive identification, analysis, and shaping of mechanisms that influence collective intelligence.

The nonlinear dynamics of (Bizyaeva et al., 2022) describe how the multi-dimensional opinion state of each individual in a group evolves over time in response to the individual’s opinions, its observed opinions of others, and its predispositions or information about options that the individual may have acquired through payoffs or

environmental cues. Unlike many existing models based on the classical model of DeGroot (1974) in which each individual updates its opinions according to a weighted average of its neighbors' opinions (Altafini, 2013; Cisneros-Velarde et al., 2021; Friedkin and Johnsen, 1999; Olfati-Saber and Murray, 2004), this model imposes a saturation on the opinion exchanges. The saturation, motivated by models in biology, physics, and social science, provides a naturally smooth limit on social influence and avoids the paradox of the linear update wherein the stronger the difference in opinions between a pair of interacting agents, the greater the strength of their mutual response. The saturation also makes the update fundamentally nonlinear. As a result of the nonlinearity, opinions form through a point of ultra-sensitivity (a bifurcation point, which is a singular point in the dynamics), where a neutral or deadlocked opinion becomes unstable and multiple, simultaneously stable solutions of agreement or disagreement among individuals emerge. The ultra-sensitivity near the bifurcation point means that which of the alternative solutions is realized can depend on small and unpredictable signals or asymmetries. The relative size of the region of attraction for each of the stable solutions, under symmetric or asymmetric conditions, can be predicted (and thus shaped) by parameters that describe the system, interactions, and environment.

In Bizyaeva et al. (2022), each individual has associated to it a variable that represents its attention, or susceptibility, to social information. The authors study what happens when an individual's attention changes over time according to dynamics that depend on a leaky accumulation of the saturated measurement of the strength of opinions of its neighbors. They show how every member of the group can become engaged in the opinion-forming process through a cascade, even if only a few individuals are initially engaged or are in receipt of an external cue. Analysis of the coupled opinion-attention dynamics predicts how parameters in these dynamics tune implicit thresholds that govern the sensitivity and robustness of the formation of the collective state and its transitions among alternative solutions. The role of this sensitivity in enabling the kind of collective intelligence that allows for dynamic task allocation in a changing environment was studied in (Franci et al., 2021). As discussed in *Cooperation and coordination*, in the case of agents choosing among strategies in a coordination game or a social dilemma, such as the prisoner's dilemma or the public goods game, an increase in attention to social cues models an increase in reciprocal behavior and the emergence of a stable mutually cooperative solution (Park et al., 2022). The dynamics have been used to study the stability–flexibility dilemma in cognitive control (Musslick et al., 2019) and to investigate how things can go awry, as in the case of political polarization (Leonard et al., 2021).

There are other kinds of mathematical models well-suited to the examination of collective intelligence that are also analytically tractable. For example, the multi-agent N -armed-bandit model, discussed above and in *Collective intelligence and design*, represents the explore-exploit trade-off present for each individual making sequential choices over N options with uncertain reward, as in the case of a social group foraging over N patches in an environment with uncertain spatial and temporal distribution of resource. In one version of this model (Madhushani and Leonard, 2019), each agent k observes the choice and reward of each of its neighbors with probability p_k (defining its “sociability”) and combines this information with its own measurements to update its next choice of option using an algorithm that seeks to maximize its own cumulative reward (Auer et al., 2002). Collective intelligence can be analyzed in terms of bounds on the group's cumulative regret rate, which can be used to evaluate, for example, the influence of individual differences in sociability.

The replicator–mutator dynamics (Bürger, 1998; Page and Nowak, 2002) provide another example of an analytically tractable model that exhibits a rich set of alternative stable solutions. This model describes the evolutionary dynamics of the game-theoretic interactions of subgroups within a population as they modify their adoption of competing strategies where mutation, a key aspect of the theory of selection, is represented by the possibility of individuals randomly switching between strategies. The model has been applied in a number of contexts, notably to the evolution of grammar acquisition (Nowak et al., 2001; Komarova and Levin, 2010). Hofbauer and Sandholm (2009) developed a method for finding conditions on the payoff function in a class of evolutionary dynamics, called stable or contractive games, that guarantees convergence of strategies to the Nash equilibrium. Fox and Shamma (2013) showed that stable games are “passive” in the sense of feedback systems theory and generalized the method of Hofbauer and Sandholm using the tools of passivity theory (Willems, 1972a; 1972b). The results were further developed in (Park et al., 2019b) and extended to the case in which a networked group of agents experience communication and computational delays in learning the payoffs. Park et al. (2021) use a replicator model and the stabilization method to derive decision-making rules with performance guarantees for allocation of agents over tasks in real-time in a changing environment. The approach is demonstrated in a multi-robot trash collection problem, where robots make decisions for moving among patches to pick up trash that is accumulating nonuniformly over time so that trash volume is minimized.

Because mathematical models allow for the abstraction of principles that underlie collective intelligence, they can be used broadly, including for explorations in the performing arts. As part of a novel project to facilitate a kind of

creative collective intelligence (Özçimder et al., 2019), the replicator–mutator model with a nonlinear fitness function was used to find new ways to experiment with artistic composition in the collaborative making of a dance piece called “There Might Be Others,” which premiered in New York City in March 2016. The piece is a structured improvisation in which the dancers (and musicians) make compositional choices on-the-fly that are inspired and constrained by choreographed performance rules. Communicating only through motion, the dancers jointly perform their way through a sequence of dance modules, choosing anew in each performance how to order, juxtapose, and vary the modules to meet the aesthetic goals of the piece. The model was used to investigate an artistic explore-exploit trade-off identified in the piece and ways in which the performance rules could be designed to enhance the dancers’ ever-evolving invention of beautiful patterns and moments of human connection. A key to this was a nonlinear analysis of the replicator-mutator model with the nonlinear fitness function, which revealed a hysteresis among the multiple equilibria and a dial, representing how likely the dancers were to spontaneously switch among modules, that controlled the pacing and periodicity of the dancers’ explorations.

Conclusions

Mechanisms leading to collective intelligence of various forms have arisen multiple times during evolution, from bacteria and slime molds to human groups. In Nature, collective intelligence which is manifest both at the level of the individual and at the level of the group implies a degree of cooperation, including information sharing and activity coordination, among individuals, typically involving costs to individuals in exchange for the benefits of sharing information with others. When it does arise, in general, we expect that the expected payoffs to an individual engaging in cooperative behavior exceeds what that individual would receive by going it alone. Thus, one useful measure of collective intelligence is the size of the region of attraction to the mutually cooperative solution. A similar definition can be applied as well in engineered ensembles, such as a decentralized group of autonomous robots, where feedback laws that govern individual behaviors are designed to yield an intelligence at the level of the group that is unattainable by agents on their own.

For the designed system, unlike the evolved system, it might not be necessary that every individual is better off in the collective, especially if it is the performance of the ensemble that is the key design objective. However, in both natural and designed settings, useful measures of collective intelligence should account for trade-offs that arise among the many possible benefits, and costs, of collective intelligence, for example, between flexibility in the face of

change and robustness to disturbance or between accumulation of resource and reduction of risk. Mathematical models, like in Bizyaeva et al. (2022), that exhibit a wide range of outcomes, including multi-stability of solutions, and also allow for analytical tractability, can provide a systematic means for examining collective intelligence and the measures that define it.

One challenge for future work will be suitable definitions of collective intelligence for hierarchical systems, like corporations, in which there are trade-offs, for example, between the interests of the corporation, its executives, and its workers. More crucially, our cities and societies are facing a raft of challenges to their sustainability that are not solvable without collective intelligence, of a variety of forms: information retrieval, information sharing, decision-making, and—most importantly—finding the best of solutions in the coordination game that will determine our survival. Garrett Hardin (Hardin, 1968) highlighted the problem of the tragedy of the commons, and argued that the solution was in “mutual coercion, mutually agreed upon,” and Elinor Ostrom and her collaborators showed how those mutualisms could emerge from the individual up (Ostrom, 1990; Desouza, 2008). We need to understand what collective intelligence means for our heterogeneous world, and how in particular to incorporate the intelligence and desires of unborn future generations. This is perhaps the greatest open challenge for collective intelligence.

Acknowledgments

The first author thanks Shinkyu Park and Alessio Franci for helpful suggestions. Both authors thank Scott Page for his input. The authors gratefully acknowledge the support of the Army Research Office grant no. W911NF-18-1-0325 and Office of Naval Research grant no. N00014-19-1-2556.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by Army Research Office grant no. W911NF-18-1-0325 and Office of Naval Research grant no. N00014-19-1-2556.

References

- Acemoglu D and Ozdaglar A (2011) Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1(1): 3–49. DOI: [10.1007/s13235-010-0004-1](https://doi.org/10.1007/s13235-010-0004-1).
- Akçay E, Meirowitz A, Ramsay KW, et al. (2012) Evolution of cooperation and skew under imperfect information. *Proceedings of the National Academy of Sciences of the United States of*

- America* 109(37): 14936–14941. DOI: [10.1073/pnas.1212925109](https://doi.org/10.1073/pnas.1212925109).
- Allotta B, Costanzi R, Ridolfi A, et al. (2015) The ARROWS project: adapting and developing robotics technologies for underwater archaeology. *IFAC-PapersOnLine* 48(2): 194–199. DOI: [10.1016/j.ifacol.2015.06.032](https://doi.org/10.1016/j.ifacol.2015.06.032).
- Altafini C (2013) Consensus problems on networks with antagonistic interactions. *IEEE Transactions on Automatic Control* 58(4): 935–946. DOI: [10.1109/TAC.2012.2224251](https://doi.org/10.1109/TAC.2012.2224251).
- Anandkumar A, Michael N, Tang AK, et al. (2011) Distributed algorithms for learning and cognitive medium access with logarithmic regret. *IEEE Journal on Selected Areas in Communications* 29(4): 731–745. DOI: [10.1109/JSAC.2011.110406](https://doi.org/10.1109/JSAC.2011.110406).
- Anderson BDO and Moore JB (1979) *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall.
- Ando H, Oasa Y, Suzuki I, et al. (1999) Distributed memoryless point convergence algorithm for mobile robots with limited visibility. *IEEE Transactions on Robotics and Automation* 15(5): 818–828. DOI: [10.1109/70.795787](https://doi.org/10.1109/70.795787).
- Arthur WB (1994) Inductive reasoning and bounded rationality. *American Economic Review* 84(2): 406–411. <http://www.jstor.org/stable/2117868>
- Åström KJ and Wittenmark B (1995) *Adaptive Control*. Boston, Mass: Addison-Wesley.
- Auer P, Cesa-Bianchi N and Fischer P (2002) Finite-time analysis of the multiarmed bandit problem. *Machine Learning* 47(2–3): 235–256. DOI: [10.1023/A:1013689704352](https://doi.org/10.1023/A:1013689704352).
- Axelrod R (1984) *The Evolution of Cooperation*. New York, NY: Basic Books.
- Axelrod R and Hamilton WD (1981) The evolution of cooperation. *Science* 211(4489): 1390–1396. DOI: [10.1126/science.7466396](https://doi.org/10.1126/science.7466396).
- Ballerini M, Cabibbo N, Candelier R, et al. (2008) Interaction ruling animal collective behavior depends on topological rather than metric distance: evidence from a field study. *Proceedings of the National Academy of Sciences of the United States of America* 105(4): 1232–1237. doi: [10.1073/pnas.0711437105](https://doi.org/10.1073/pnas.0711437105). Available at: <https://doi.org/10.1073/pnas.0711437105>.
- Beauchamp G (2008) What is the magnitude of the group-size effect on vigilance? *BMC Ecology* 19(6): 1361–1368. DOI: [10.1093/beheco/arn096](https://doi.org/10.1093/beheco/arn096).
- Bialek W, Cavagna A, Giardina I, et al. (2012) Statistical mechanics for natural flocks of birds. *Proceedings of the National Academy of Sciences of the United States of America* 109(13): 4786–4791. DOI: [10.1073/pnas.1118633109](https://doi.org/10.1073/pnas.1118633109).
- Bizyaeva A, Franci A and Leonard NE (2022) Nonlinear opinion dynamics with tunable sensitivity. *IEEE Transactions on Automatic Control*. doi: [10.1109/TAC.2022.3159527](https://doi.org/10.1109/TAC.2022.3159527). Available at: <https://doi.org/10.1109/TAC.2022.3159527>.
- Brown JL (1982) Optimal group size in territorial animals. *Journal of Theoretical Biology* 95(4): 793–810. DOI: [10.1016/0022-5193\(82\)90354-X](https://doi.org/10.1016/0022-5193(82)90354-X).
- Brush ER, Leonard NE and Levin SA (2016) The content and availability of information affects the evolution of social-information gathering strategies. *Theoretical Ecology* 9(4): 455–476. DOI: [10.1007/s12080-016-0301-4](https://doi.org/10.1007/s12080-016-0301-4).
- Bürger R (1998) Mathematical properties of mutation-selection models. *Genetica* 102: 279–298. DOI: [10.1023/A:1017043111100](https://doi.org/10.1023/A:1017043111100).
- Bullo F, Cortés J and Martinez S (2009) *Distributed Control of Robotic Networks*. Princeton, NJ: Princeton University Press.
- Cao Y, Yu W, Ren W, et al. (2013) An overview of recent progress in the study of distributed multi-agent coordination. *IEEE Transactions on Industrial Informatics* 9(1): 427–438. DOI: [10.1109/TII.2012.2219061](https://doi.org/10.1109/TII.2012.2219061).
- Carli R, Chiuso A, Schenato L, et al. (2008) Distributed Kalman filtering based on consensus strategies. *IEEE Journal on Selected Areas in Communications* 26(4): 622–633. doi: [10.1109/JSAC.2008.080505](https://doi.org/10.1109/JSAC.2008.080505). Available at: <https://doi.org/10.1109/JSAC.2008.080505>.
- Chakraborty M, Chua KYP, Das S, et al (2017) Coordinated versus decentralized exploration in multi-agent multi-armed bandits. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial intelligence (IJCAI 2017), IJCAI, Melbourne, Australia, 164–170. DOI: [10.24963/ijcai.2017/24](https://doi.org/10.24963/ijcai.2017/24).
- Chen A, Reperant L, Fischhoff IR, et al. (2021) Increased vigilance of plains zebras (*Equus quagga*) in response to more bush coverage in a Kenyan savanna. *Climate Change Ecology* 1: 100001. doi: [10.1016/j.ecochg.2021.100001](https://doi.org/10.1016/j.ecochg.2021.100001). Available at: <https://doi.org/10.1016/j.ecochg.2021.100001>.
- Christofides PD, Scattolini R, Muñoz de la Peña D, et al. (2013) Distributed model predictive control: a tutorial review and future research directions. *Computers & Chemical Engineering* 51: 21–41. DOI: [10.1016/j.compchemeng.2012.05.011](https://doi.org/10.1016/j.compchemeng.2012.05.011).
- Cisneros-Velarde P, Chan KS and Bullo F (2021) Polarization and fluctuations in signed social networks. *IEEE Transactions on Automatic Control* 66(8): 3789–3793. doi: [10.1109/TAC.2020.3024967](https://doi.org/10.1109/TAC.2020.3024967).
- Clark A, Bushnell L and Poovendran R (2014) A supermodular optimization framework for leader selection under link noise in linear multi-agent systems. *IEEE Transactions on Automatic Control* 59(2): 283–296. DOI: [10.1109/TAC.2013.2281473](https://doi.org/10.1109/TAC.2013.2281473).
- Cohen JD, McClure SM and Yu AJ (2007) Should I stay or should I go? how the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362(1481): 933–942. DOI: [10.1098/rstb.2007.2098](https://doi.org/10.1098/rstb.2007.2098).
- Cortés J, Martínez S, Karatas T, et al. (2004) Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation* 20(2): 243–255. DOI: [10.1109/TRA.2004.824698](https://doi.org/10.1109/TRA.2004.824698).
- Couzin ID, Krause J, Franks NR, et al. (2005) Effective leadership and decision-making in animal groups on the move. *Nature* 433: 513–516. DOI: [10.1038/nature03236](https://doi.org/10.1038/nature03236).
- DeGroot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345): 121–132. DOI: [10.2307/2285509](https://doi.org/10.2307/2285509).

- Desouza KC (2008) Understanding knowledge as a commons: from theory to practice [Review of the book Understanding knowledge as a commons: From theory to practice, edited by C. Hess and E. Ostrom]. *Journal of the American Society for Information Science and Technology* 59(2): 333–334. DOI: [10.1002/Asi.20747](https://doi.org/10.1002/Asi.20747).
- Drescher K, Nadell CD, Stone HA, et al. (2014) Solutions to the public goods dilemma in bacterial biofilms. *Current Biology* 24(1): 50–55. DOI: [10.1016/j.cub.2013.10.030](https://doi.org/10.1016/j.cub.2013.10.030).
- Dresner K and Stone P (2008) A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research* 31: 591–656. doi: [10.1613/jair.2502](https://doi.org/10.1613/jair.2502). Available at: <https://doi.org/10.1613/jair.2502>.
- Farivarnejad H and Berman S (2022) Multirobot control strategies for collective transport. *Annual Review of Control, Robotics, and Autonomous Systems*. Advance Online Publication, 5(1). doi: [10.1146/annurev-control-042920-095844](https://doi.org/10.1146/annurev-control-042920-095844).
- Fisac JF, Bronstein E, Stefansson E, et al. (2019) Hierarchical game-theoretic planning for autonomous vehicles. *International Conference on Robotics and Automation (ICRA)*. Montreal, Canada: IEEE, 9590–9596. Available at: <https://doi.org/10.1109/ICRA.2019.8794007>.
- Fitch K and Leonard NE (2016) Joint centrality distinguishes optimal leaders in noisy networks. *IEEE Transactions on Control of Network Systems* 3(4): 366–378. DOI: [10.1109/TCNS.2015.2481138](https://doi.org/10.1109/TCNS.2015.2481138).
- Fox M and Shamma J (2013) Population games, stable games, and passivity. *Games* 4(4): 561–583. DOI: [10.3390/g4040561](https://doi.org/10.3390/g4040561).
- Fracchia J and Lewontin RC (1999) Does culture evolve? *History and Theory* 38(4): 52–78. DOI: [10.1111/0018-2656.00104](https://doi.org/10.1111/0018-2656.00104).
- Franci A, Bizyayeva A, Park S, et al. (2021) Analysis and control of agreement and disagreement opinion cascades. *Swarm Intelligence* 15(1): 47–82. DOI: [10.1007/s11721-021-00190-w](https://doi.org/10.1007/s11721-021-00190-w).
- Friedkin NE and Johnsen EC (1999) Social influence networks and opinion change. In: Thye R, Lawler EJ, Macy MW, et al. (eds) *Advances in Group Processes*. Bingley, UK: Emerald Group Publishing Limited, 16, 1–29.
- Fukushima S, Nitanda A and Yamanishi K (2021) *Online Robust and Adaptive Learning from Data Streams*. StatisticsArXiv. Available at: <https://arxiv.org/abs/2007.12160v2>.
- Gai Y and Krishnamachari B (2014) Distributed stochastic online learning policies for opportunistic spectrum access. *IEEE Transactions on Signal Processing* 62(23): 6184–6193. DOI: [10.1109/TSP.2014.2360821](https://doi.org/10.1109/TSP.2014.2360821).
- Gammaitoni L, Hänggi P, Jung P, et al. (1998) Stochastic resonance. *Revista médica de Panamá* 70(1): 223–287. DOI: [10.1103/RevModPhys.70.223](https://doi.org/10.1103/RevModPhys.70.223).
- Gopalakrishnan K and Balakrishnan H (2021) Control and optimization of air traffic networks. *Annual Review of Control, Robotics, and Autonomous Systems* 4(1): 397–424. DOI: [10.1146/annurev-control-070720-080844](https://doi.org/10.1146/annurev-control-070720-080844).
- Gordon DM (2013) The rewards of restraint in the collective regulation of foraging by harvester ant colonies. *Nature* 498(7452): 91–93. DOI: [10.1038/nature12137](https://doi.org/10.1038/nature12137).
- Guéant O, Lasry JM and Lions PL (2011) Mean field games and applications. In: Carmona RA, Çinlar E, Ekeland I, et al. (eds) *Paris-Princeton Lectures on Mathematical Finance 2010 Lecture Notes in Mathematics* Springer, 2003, 205–266. DOI: [10.1007/978-3-642-14660-2_3](https://doi.org/10.1007/978-3-642-14660-2_3).
- Guttal V and Couzin ID (2010) Social interactions, information use, and the evolution of collective migration. *Proceedings of the National Academy of Sciences of the United States of America* 107(37): 16172–16177. doi: [10.1073/pnas.1006874107](https://doi.org/10.1073/pnas.1006874107). Available at: <https://doi.org/10.1073/pnas.1006874107>.
- Hamilton WD (1964a) The genetical evolution of social behaviour: I. *Journal of Theoretical Biology* 7(1): 1–16. DOI: [10.1016/0022-5193\(64\)90038-4](https://doi.org/10.1016/0022-5193(64)90038-4).
- Hamilton WD (1964b) The genetical evolution of social behaviour: II. *Journal of Theoretical Biology* 7(1): 17–52. doi: [10.1016/0022-5193\(64\)90039-6](https://doi.org/10.1016/0022-5193(64)90039-6).
- Hamilton WD (1971) Geometry for the selfish herd. *Journal of Theoretical Biology* 31(2): 295–311. DOI: [10.1016/0022-5193\(71\)90189-5](https://doi.org/10.1016/0022-5193(71)90189-5).
- Hardin G. (1968) The tragedy of the commons. *Science* 162(3859): 1243–1248. doi: [10.1126/science.162.3859.1243](https://doi.org/10.1126/science.162.3859.1243). Available at: <https://doi.org/10.1126/science.162.3859.1243>.
- Hills TT, Todd PM, Lazer D, et al. (2015) Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences* 19(1): 46–54. DOI: [10.1016/j.tics.2014.10.004](https://doi.org/10.1016/j.tics.2014.10.004).
- Hofbauer J and Sandholm WH (2009) Stable games and their dynamics. *Journal of Economic Theory* 144(4): 1665–1693. e4 DOI: [10.1016/j.jet.2009.01.007](https://doi.org/10.1016/j.jet.2009.01.007).
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* 79(8): 2554–2558. doi: [10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554). Available at: <https://doi.org/10.1073/pnas.79.8.2554>.
- Huang M, Caines PE and Malhamé RP (2010) Social dynamics in mean field LQG control: egoistic and altruistic Agents. In: Proceedings of the 49th IEEE conference on decision and control, Atlanta, GA, US, 15–17 Dec. 2010, 3140–3145. DOI: [10.1109/Cdc.2010.5717299](https://doi.org/10.1109/Cdc.2010.5717299).
- Jia P, MirTabatabaei A, Friedkin NE, et al. (2015) Opinion dynamics and the evolution of social power in influence networks. *SIAM Review* 57(3): 367–397. DOI: [10.1137/130913250](https://doi.org/10.1137/130913250).
- Kalathil D, Nayyar N and Jain R (2014) Decentralized learning for multiplayer multiarmed bandits. *IEEE Transactions on Information Theory* 60(4): 2331–2345. DOI: [10.1109/TIT.2014.2302471](https://doi.org/10.1109/TIT.2014.2302471).
- Kaufmann E, Cappe O and Garivier A (2012) On Bayesian upper confidence bounds for bandit problems. *Proceedings of Machine Learning Research (PMLR)* 22: 592–600. Available at: <https://proceedings.mlr.press/v22/kaufmann12.html>.

- Kauffman S and Levin S (1987) Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology* 128(1): 11–45. DOI: [10.1016/S0022-5193\(87\)80029-2](https://doi.org/10.1016/S0022-5193(87)80029-2).
- Khan UA and Moura JMF (2008) Distributing the Kalman filter for large-scale systems. *IEEE Transactions on Signal Processing* 56(10): 4919–4935. DOI: [10.1109/TSP.2008.927480](https://doi.org/10.1109/TSP.2008.927480).
- Knorn S, Chen Z and Middleton RH (2016) Overview: collective control of multiagent systems. *IEEE Transactions on Control of Network Systems* 3(4): 334–347. DOI: [10.1109/TCNS.2015.2468991](https://doi.org/10.1109/TCNS.2015.2468991).
- Kolla RK, Jagannathan K and Gopalan A (2018) Collaborative learning of stochastic bandits over a social network. *IEEE/ACM Transactions on Networking* 26(4): 1782–1795. DOI: [10.1109/TNET.2018.2852361](https://doi.org/10.1109/TNET.2018.2852361).
- Komarova NL and Levin SA (2010) Eavesdropping and language dynamics. *Journal of Theoretical Biology* 264(1): 104–118. DOI: [10.1016/j.jtbi.2010.01.010](https://doi.org/10.1016/j.jtbi.2010.01.010).
- Krause J and Ruxton GD (2002) *Living in Groups*. Oxford, UK: Oxford University Press.
- Kube CR and Bonabeau E (2000) Cooperative transport by ants and robots. *Robotics and Autonomous Systems* 30: 85–101.
- Lai TL and Robbins H (1985) Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6(1): 4–22. DOI: [10.1016/0196-8858\(85\)90002-8](https://doi.org/10.1016/0196-8858(85)90002-8).
- Landgren P, Srivastava V and Leonard NE (2016). Distributed cooperative decision-making in multiarmed bandits: frequentist and bayesian algorithms. In: Proceedings of the 55th IEEE conference on decision and control, Las Vegas, NV, USA, 12–14 Dec. 2016. ComputerScienceArXiv, 167–172. DOI: [10.1109/CDC.2016.7798264](https://doi.org/10.1109/CDC.2016.7798264). Correction at: (2019). <https://arxiv.org/abs/1606.00911>
- Landgren P, Srivastava V and Leonard NE (2021) Distributed cooperative decision making in multi-agent multi-armed bandits. *Automatica* 125(4): 109445. DOI: [10.1016/j.automatica.2020.109445](https://doi.org/10.1016/j.automatica.2020.109445).
- Lasry J-M and Lions P-L (2006) Jeux à champ moyen. I-Le cas stationnaire. *Comptes Rendus Mathématique* 343(9): 619–625. DOI: [10.1016/j.crma.2006.09.019](https://doi.org/10.1016/j.crma.2006.09.019).
- Lekien F and Leonard NE (2009) Nonuniform coverage and cartograms. *SIAM Journal on Control and Optimization* 48(1): 351–372. DOI: [10.1137/070681120](https://doi.org/10.1137/070681120).
- Lenton TM, Daines SJ, Dyke JG, et al. (2018) Selection for gaia across multiple scales. *Trends in Ecology & Evolution* 33(8): 633–645. DOI: [10.1016/j.tree.2018.05.006](https://doi.org/10.1016/j.tree.2018.05.006).
- Leonard NE (2014) Multi-agent system dynamics: bifurcation and behavior of animal groups. *Annual Reviews in Control* 38(2): 171–183. DOI: [10.1016/j.arcontrol.2014.09.002](https://doi.org/10.1016/j.arcontrol.2014.09.002).
- Leonard NE, Lipsitz K, Bizyaeva A, et al. (2021) The nonlinear feedback dynamics of asymmetric political polarization. *Proceedings of the National Academies of the United States of America* 118(50): e2102149118. DOI: [10.1073/pnas.2102149118](https://doi.org/10.1073/pnas.2102149118).
- Leonard NE, Paley DA, Davis RE, et al. (2010) Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay. *Journal of Field Robotics* 27(6): 718–740. doi: [10.1002/rob.20366](https://doi.org/10.1002/rob.20366). Available at: <https://doi.org/10.1002/rob.20366>.
- Leonard NE, Paley DA, Lekien F, et al. (2007) Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE* 95(1): 48–74. DOI: [10.1109/JPROC.2006.887295](https://doi.org/10.1109/JPROC.2006.887295).
- Leonard NE, Shen T, Nabet B, et al. (2012) Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences* 109(1): 227–232. DOI: [10.1073/pnas.1118318108](https://doi.org/10.1073/pnas.1118318108).
- Levin JE and Miller JP (1996) Broadband neural encoding in the cricket cerebellar sensory system enhanced by stochastic resonance. *Nature* 380(6570): 165–168. DOI: [10.1038/380165a0](https://doi.org/10.1038/380165a0).
- Levin SA (1999) *Fragile Dominion: Complexity and the Commons*. New York, NY: Perseus Books.
- Levin SA (2014) Public goods in relation to competition, cooperation, and spite. *Proceedings of the National Academy of Sciences of the United States of America*, 111(Supplement 3): 10838–10845. DOI: [10.1073/pnas.140083011](https://doi.org/10.1073/pnas.140083011).
- Lewontin RC (1977). Adattamento. *Enciclopedia Einaudi* 1: 198–214.
- Lewontin RC (1978) Adaptation. *Scientific American* 239(3): 212–230. DOI: [10.2307/24955807](https://doi.org/10.2307/24955807).
- Lewontin RC (1985) Adaptation. In: Levins R and Lewontin RC (eds). *The Dialectical Biologist*. Cambridge, MA: Harvard University Press, 65–84.
- Liljenström H (2003) Neural stability and flexibility: a computational approach. *Neuropsychopharmacology* 28: S64–S73. DOI: [10.1038/sj.npp.1300137](https://doi.org/10.1038/sj.npp.1300137).
- Lin F, Fardad M and Jovanovic MR (2014) Algorithms for leader selection in stochastically forced consensus networks. *IEEE Transactions on Automatic Control* 59(7): 1789–1802. DOI: [10.1109/TAC.2014.2314223](https://doi.org/10.1109/TAC.2014.2314223).
- Liu T, Song Y, Zhu L, et al. (2022) Stability and control of power grids. *Annual Review of Control, Robotics, and Autonomous Systems*. Advance Online Publication 5(1). DOI: [10.1146/annurev-control-042820-011148](https://doi.org/10.1146/annurev-control-042820-011148).
- Madhushani U and Leonard NE (2019) Heterogeneous stochastic interactions for multiple agents in a multi-armed bandit problem. In: Proceedings of the 18th European control conference, 3502–3507. DOI: [10.23919/ECC.2019.8796036](https://doi.org/10.23919/ECC.2019.8796036).
- Madhushani U and Leonard NE (2020a) Distributed learning: sequential decision making in resource-constrained environments. *Practical ML for Developing Countries Workshop Learning Under Limited/Low Resource Scenarios*. Ethiopia, Virtual: ICLR: Addis Ababa. Available at: <https://pml4dc.github.io/iclr2020/papers.html>.
- Madhushani U and Leonard NE (2020b) It doesn't get better and here's why: a fundamental drawback in natural extensions of UCB to multi-agent bandits. In: I Can't Believe It's Not Butter Workshop: Bridging the Gap between Theory and Empiricism in Probabilistic Machine Learning. *Neural Information and Processing Systems (NeurIPS), Virtual*. Available at: <https://openreview.net/pdf?id=eK034ngO05Y>.

- Madhushani U and Leonard NE (2021) Heterogeneous explore-exploit strategies on multi-star networks. *IEEE Control Systems Letters* 5(5): 1603–1608. DOI: [10.1109/LCSYS.2020.3042459](https://doi.org/10.1109/LCSYS.2020.3042459).
- Marden JR, Arslan G and Shamma JS (2009) Cooperative control and potential games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 39(6): 1393–1407. DOI: [10.1109/TSMCB.2009.2017273](https://doi.org/10.1109/TSMCB.2009.2017273).
- Marden JR and Shamma JS (2018) Game theory and control. *Annual Review of Control, Robotics, and Autonomous Systems* 1: 105–134. doi: [10.1146/annurev-control-060117-105102](https://doi.org/10.1146/annurev-control-060117-105102). Available at: <https://doi.org/10.1146/annurev-control-060117-105102>.
- Marjovi A, Nunes JG, Marques L, et al (2009) Multi-robot exploration and fire searching. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), St. Louis, MO, USA, 1929–1934. DOI: [10.1109/IROS.2009.5354598](https://doi.org/10.1109/IROS.2009.5354598).
- Martinez-Rubio D, Kanade V and Rebeschini P (2019) Decentralized cooperative stochastic bandits. In: Wallach H, Larochelle H, Beygelzimer A, et al. (eds) *Advances in Neural Information Processing Systems*. Vancouver, BC: NeurIPS, 32, 4531–4542. Available at: <https://ora.ox.ac.uk/objects/uuid:89855f93-5a45-4554-9cda-7e21c22ab41d>.
- Matrone GC, Cipriani C, Secco EL, et al. (2010) Principal components analysis based control of a multi-dof underactuated prosthetic hand. *Journal of Neuroengineering and Rehabilitation* 7: 16. doi: [10.1186/1743-0003-7-16](https://doi.org/10.1186/1743-0003-7-16). Available at: <https://doi.org/10.1186/1743-0003-7-16>.
- Maynard Smith J (1974) The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology* 47(1): 209–221. DOI: [10.1016/0022-5193\(74\)90110-6](https://doi.org/10.1016/0022-5193(74)90110-6).
- Maynard Smith J, Warren N, Cavalli-Sforza LL, et al. (1982) Models of cultural and genetic change. [Review of the Book *Cultural Transmission and Evolution; Genes, Mind and Culture: The Coevolutionary Process* by L. L. Cavalli-Sforza, M. W. Feldman, C. J. Lumsden, & E. O. & Wilson]. *Evolution* 36(3): 620–627. <https://www.jstor.org/stable/2408109>
- Menache I and Ozdaglar A (2011) Network games: theory, models, and dynamics. *Synthesis Lectures on Communication Networks*. doi: [10.2200/S00330ED1V01Y201101CNT009](https://doi.org/10.2200/S00330ED1V01Y201101CNT009). Available at: <https://doi.org/10.2200/S00330ED1V01Y201101CNT009>.
- Miller MB and Bassler BL (2001) Quorum sensing in bacteria. *Annual Review of Microbiology* 55: 165–199. DOI: [10.1146/annurev.micro.55.1.165](https://doi.org/10.1146/annurev.micro.55.1.165).
- Mirsky R, Xiao S, Hart J, et al. (2021) Prevention and resolution of conflicts in social navigation – a survey. *Computer Science arXiv*. doi: [10.48550/arXiv.2106.12113](https://arxiv.org/abs/2106.12113). Available at: <https://arxiv.org/abs/2106.12113>.
- Murphy RR, Tadokoro S, Nardi D, et al. (2008) Search and rescue robotics. In: Siciliano B and Khatib O (eds). *Springer Handbooks of Robotics*. Cham, Switzerland: Springer, 1151–1173.
- Musslick S, Bizyaeva A and Agaron S (2019) Stability-flexibility dilemma in cognitive control: A dynamical system perspective. *Proceedings of the 41st Annual Meeting*. Montreal, Canada: Cognitive Science Society. Available at: <https://par.nsf.gov/servlets/purl/10125021>.
- Nabet B, Leonard NE, Couzin ID, et al. (2009) Dynamics of decision making in animal group motion. *Journal of Nonlinear Science* 19(4): 399–435. DOI: [10.1007/s00332-008-9038-6](https://doi.org/10.1007/s00332-008-9038-6).
- Nadell CD, Bassler BL and Levin SA (2008a) Observing bacteria through the lens of social evolution. *Journal of Biology* 7: 27. doi: [10.1186/jbiol87](https://doi.org/10.1186/jbiol87).
- Nadell CD, Xavier JB, Levin SA, et al. (2008b) The evolution of quorum sensing in bacterial biofilms. *Plos Biology* 6(1): 14–179. doi: [10.1371/journal.pbio.0060014](https://doi.org/10.1371/journal.pbio.0060014).
- Nedic A, Olshevsky A and Rabbat MG (2018) Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE* 106(5): 953–976. DOI: [10.1109/JPROC.2018.2817461](https://doi.org/10.1109/JPROC.2018.2817461).
- Nikolaidis S, Nath S and Procaccia AD (2017) Game-theoretic modeling of human adaptation in human-robot collaboration. *HRI '17: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. Vienna, Austria: ACM, 323–331. Available at: <https://doi.org/10.1145/2909824.3020253>.
- Nourian M, Malhamé R, Huang M, et al. (2011) Mean field (NCE) formulation of estimation based leader-follower collective dynamics. *International Journal of Robotics & Automation* 26(1): 120–212. DOI: [10.2316/Journal.206.2011.1.206-3426](https://doi.org/10.2316/Journal.206.2011.1.206-3426).
- Nowak MA, Komarova NL and Niyogi P (2001) Evolution of universal grammar. *Science* 291(5501): 114–118. DOI: [10.1126/science.291.5501.114](https://doi.org/10.1126/science.291.5501.114).
- Olfati-Saber R (2005) *Distributed Kalman filter with embedded consensus filters*. In: Proceedings of the 44th IEEE Conference on Decision and Control, Seville, Spain, 15-15 Dec. 2005. IEEE, 8179–8184. DOI: [10.1109/CDC.2005.1583486](https://doi.org/10.1109/CDC.2005.1583486).
- Olfati-Saber R, Fax JA and Murray RM (2007) Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE* 95(1): 215–233. DOI: [10.1109/JPROC.2006.887293](https://doi.org/10.1109/JPROC.2006.887293).
- Olfati-Saber R and Murray RM (2004) Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control* 49(9): 1520–1533. DOI: [10.1109/TAC.2004.834113](https://doi.org/10.1109/TAC.2004.834113).
- Ostrom E (1990) *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge University Press.
- Özçimder K, Dey B, Franci A, et al. (2019) Social decision-making driven by artistic explore-exploit tension. *Interdisciplinary Science Reviews* 44(1): 55–81. DOI: [10.1080/03080188.2018.1544806](https://doi.org/10.1080/03080188.2018.1544806).
- Page KM and Nowak MA (2002) Unifying evolutionary dynamics. *Journal of Theoretical Biology* 219(1): 93–98. DOI: [10.1006/jtbi.2002.3112](https://doi.org/10.1006/jtbi.2002.3112).
- Pagliara R, Gordon DM and Leonard NE (2018) Regulation of harvester ant foraging as a closed-loop excitable system. *PLoS Computational Biology* 14(12): e1006200. doi: [10.1371/journal.pcbi.1006200](https://doi.org/10.1371/journal.pcbi.1006200).

- journal.pcbi.1006200. Available at: <https://doi.org/10.1371/journal.pcbi.1006200>.
- Pais D, Hogan PM, Schlegel T, et al. (2013) A mechanism for value-sensitive decision-making. *PLoS ONE* 8(9): e73216–e732169. DOI: [10.1371/journal.pone.0073216](https://doi.org/10.1371/journal.pone.0073216).
- Pais D and Leonard NE (2014) Adaptive network dynamics and evolution of leadership in collective migration. *Physica D: Nonlinear Phenomena* 267: 81–93. DOI: [10.1016/j.physd.2013.04.014](https://doi.org/10.1016/j.physd.2013.04.014).
- Paley DA and Wolek A (2020) Mobile sensor networks and control: adaptive sampling of spatiotemporal processes. *Annual Review of Control, Robotics, and Autonomous Systems* 3: 91–114. DOI: [10.1146/annurev-control-073119-090634](https://doi.org/10.1146/annurev-control-073119-090634).
- Parise F and Ozdaglar A (2021) Analysis and interventions in large network games. *Annual Review of Control, Robotics, and Autonomous Systems* 4: 455–486. doi: [10.1146/annurev-control-072020-084434](https://doi.org/10.1146/annurev-control-072020-084434). Available at: <https://doi.org/10.1146/annurev-control-072020-084434>.
- Park S, Martins NC and Shamma JS (2019b) From population games to payoff dynamics models: a passivity-based approach. In: Proceedings of the 58th IEEE conference on decision and control (CDC), Nice, France, 11–13 Dec, 2019, 6584–6601. DOI: [10.1109/CDC40024.2019.9029756](https://doi.org/10.1109/CDC40024.2019.9029756).
- Park S, Zhong YD and Leonard NE (2021) *Multi-robot task allocation games in dynamically changing environments*. In: International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May–5 June, 2021: 8678–8684. DOI: [10.1109/ICRA48506.2021.9561809](https://doi.org/10.1109/ICRA48506.2021.9561809).
- Park H, Aschenbach KH, Ahmed M, et al. (2019a) Animal-borne wireless network: remote imaging of community ecology. *Journal of Field Robotics* 36(6): 1141–1165. doi: [10.1002/rob.21891](https://doi.org/10.1002/rob.21891).
- Park S, Bizyaeva A and Kawakatsu M (2022) Tuning cooperative behavior in games with nonlinear opinion dynamics. *IEEE Control System Letters* 6: 2030–2035. doi: [10.1109/LCSYS.2021.3138725](https://doi.org/10.1109/LCSYS.2021.3138725). Available at: <https://doi.org/10.1109/LCSYS.2021.3138725>.
- Park S and Martins NC (2017) Design of distributed LTI observers for state omniscience. *IEEE Transactions on Automatic Control* 62(20): 561–576. DOI: [10.1109/TAC.2016.2560766](https://doi.org/10.1109/TAC.2016.2560766).
- Parker LE, Rus D and Sukhatme GS (2016) Multiple mobile robot systems. In: Siciliano B and Khatib O (eds). *Springer Handbook of Robotics*. Berlin: Springer, 1335–1384. Available at: <https://www.springerprofessional.de/en/multiple-mobile-robot-systems/10539196>.
- Parrish J and Hamner W (eds) (1997) *Animal Groups in Three Dimensions: How Species Aggregate*. Cambridge, UK: Cambridge University Press.
- Patterson S and Bamieh B (2010) Leader selection for optimal network coherence. In: Proceedings of the 49th IEEE conference on decision and control, Atlanta, GA, USA, 15–17 Dec. 2010, 2692–2697. DOI: [10.1109/CDC.2010.5718151](https://doi.org/10.1109/CDC.2010.5718151).
- Piazza C, Grioli G, Catalano MG, et al. (2019) A century of robotic hands. *Annual Review of Control, Robotics, and Autonomous Systems* 2: 1–32. DOI: [10.1146/annurev-control-060117-105003](https://doi.org/10.1146/annurev-control-060117-105003).
- Pulliam HR (1973) On the advantages of flocking. *Journal of Theoretical Biology* 38(2): 419–422. DOI: [10.1016/0022-5193\(73\)90184-7](https://doi.org/10.1016/0022-5193(73)90184-7).
- Pulliam HR and Caraco T (1984) Living in groups: is there an optimal group size? In: Davies JR and Krebs NB (eds). *Behavioral Ecology: An Evolutionary Approach*. Sunderland, Mass: Sinauer Associates, 127–147.
- Pulliam HR, Pyke GH and Caraco T (1982) The scanning behavior of juncos: a game-theoretical approach. *Journal of Theoretical Biology* 95(1): 89–103. DOI: [10.1016/0022-5193\(82\)90289-2](https://doi.org/10.1016/0022-5193(82)90289-2).
- Queralta JP, Taipalmaa J, Can Pullinen B, et al. (2020) Collaborative multi-robot search and rescue: planning, coordination, perception, and active vision. *IEEE Access* 8: 191617–191643. DOI: [10.1109/ACCESS.2020.3030190](https://doi.org/10.1109/ACCESS.2020.3030190).
- Reverdy PB, Srivastava V and Leonard NE (2014). Modeling human decision making in generalized gaussian multiarmed bandits. *Proceedings of the IEEE* 102(4): 544–571. DOI: [10.1109/JPROC.2014.2307024](https://doi.org/10.1109/JPROC.2014.2307024). Correction at: (2019). ComputerScienceArXiv. <https://arxiv.org/abs/1307.6134v5>
- Reverdy P, Srivastava V and Leonard NE (2017) Satisficing in multi-armed bandit problems. *IEEE Transactions on Automatic Control* 62(8): 3788–3803. DOI: [10.1109/TAC.2016.2644380](https://doi.org/10.1109/TAC.2016.2644380). Correction at: (2021). IEEE Transactions on Automatic Control 66(1): 476–478. DOI: [10.1109/TAC.2020.2981433](https://doi.org/10.1109/TAC.2020.2981433).
- Robbins H (1952) Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society* 58(5): 527–535. DOI: [10.1090/S0002-9904-1952-09620-8](https://doi.org/10.1090/S0002-9904-1952-09620-8).
- Rubenstein DI (1978) On predation, competition, and the advantages of group living. In: Bateston PPG and Klopfer PH (eds). *Social Behavior*. New York, NY: Plenum Press, 3, 205–231. Available at: http://dx.doi.org/10.1007/978-1-4684-2901-5_9.
- Seeley TD, Visscher PK, Schlegel T, et al. (2012) Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science (New York, N.Y.)* 335(6064): 108–111. DOI: [10.1126/science.1210361](https://doi.org/10.1126/science.1210361).
- Sepulchre R, Paley DA and Leonard NE (2007) Stabilization of planar collective motion: all-to-all communication. *IEEE Transactions on Automatic Control* 52(5): 811–824. DOI: [10.1109/TAC.2007.898077](https://doi.org/10.1109/TAC.2007.898077).
- Sepulchre R, Paley DA and Leonard NE (2008) Stabilization of planar collective motion with limited communication. *IEEE Transactions on Automatic Control* 53(3): 706–719. DOI: [10.1109/TAC.2008.919857](https://doi.org/10.1109/TAC.2008.919857).
- Shapiro CR, Starke GM and Gayme DF (2022) Turbulence and control of wind farms. *Annual Review of Control, Robotics, and Autonomous Systems*. Advance Online Publication 5(1). DOI: [10.1146/annurev-control-070221-114032](https://doi.org/10.1146/annurev-control-070221-114032).
- Spanos DP, Olfati-Saber R and Murray RM (2005) Distributed kalman filtering in sensor networks with quantifiable performance. Proceedings of the fourth international symposium on

- information processing for sensor networks (IPSN'05), Boise, ID, USA, 15-15 April 2005, 133–139. DOI: [10.1109/IPSN.2005.1440912](https://doi.org/10.1109/IPSN.2005.1440912).
- Speyer J (1979) Computation and transmission requirements for a decentralized linear-quadratic-gaussian control problem. *IEEE Transactions on Automatic Control* 24(2): 266–269. DOI: [10.1109/TAC.1979.1101973](https://doi.org/10.1109/TAC.1979.1101973).
- Srivastava V and Leonard NE (2014) Collective decision-making in ideal networks: the speed-accuracy tradeoff. *IEEE Transactions on Control of Network Systems* 1(1): 121–132. DOI: [10.1109/TCNS.2014.2310271](https://doi.org/10.1109/TCNS.2014.2310271).
- Sumpter DJT (2010) *Collective Animal Behavior*. Princeton, NJ: Princeton University Press. Available at: <https://doi.org/10.1515/9781400837106>.
- Tiebout CM (1956) A pure theory of local expenditures. *Journal of Political Economy* 64(5): 416–424. <https://www.journals.uchicago.edu/doi/10.1086/257839>
- Torney CJ, Berdahl A and Couzin ID (2011) Signalling and the evolution of cooperative foraging in dynamic environments. *Plos Computational Biology* 7(9): e1002194. DOI: [10.1371/journal.pcbi.1002194](https://doi.org/10.1371/journal.pcbi.1002194).
- Torney CJ, Levin SA and Couzin ID (2010) Specialization and evolutionary branching within migratory populations. *Proceedings of the National Academy of Sciences of the United States of America* 107(47): 20394–20399. doi: [10.1073/pnas.1014316107](https://doi.org/10.1073/pnas.1014316107). Available at: <https://doi.org/10.1073/pnas.1014316107>.
- Turing AM (1952) The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences* 237(641): 37–72. DOI: [10.1098/rstb.1952.0012](https://doi.org/10.1098/rstb.1952.0012).
- Valentini G, Ferrante E, Hamann H, et al. (2016) Collective decision with 100 kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems* 30: 553–580. DOI: [10.1007/s10458-015-9323-3](https://doi.org/10.1007/s10458-015-9323-3).
- Waddington CH (1957) *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. London, UK: George Allen & Unwin. Available at: <https://wellcomecollection.org/works/nzwm3z65>.
- Werfel J, Petersen K and Nagpal R (2014) Designing collective behavior in a termite-inspired robot construction team. *Science (New York, N.Y.)* 343(6172): 754–758. DOI: [10.1126/science.1245842](https://doi.org/10.1126/science.1245842).
- Willems JC (1972a) Dissipative dynamical systems part I: general theory. *Archive for Rational Mechanics and Analysis* 45: 321–351. DOI: [10.1007/BF00276493](https://doi.org/10.1007/BF00276493).
- Willems JC (1972b) Dissipative dynamical systems part II: linear systems with quadratic supply rates. *Archive for Rational Mechanics and Analysis* 45: 352–393. DOI: [10.1007/BF00276494](https://doi.org/10.1007/BF00276494).
- Wilson HR and Cowan JD (1972) Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal* 12(1): 1–24. [10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5).
- Wilson DS and Wilson EO (2007) Rethinking the theoretical foundation of sociobiology. *The Quarterly Review of Biology* 82(4): 327–348. DOI: [10.1086/522809](https://doi.org/10.1086/522809).
- Wolpert L (1969) Positional information and the spatial pattern of cellular differentiation. *Journal of Theoretical Biology* 25(1): 1–47. DOI: [10.1016/S0022-5193\(69\)80016-0](https://doi.org/10.1016/S0022-5193(69)80016-0).
- Ye M, Zino L, Mlakar Ž, et al. (2021) Collective patterns of social diffusion are shaped by individual inertia and trend-seeking. *Nature Communications* 12: 5698. DOI: [10.1038/s41467-021-25953-1](https://doi.org/10.1038/s41467-021-25953-1).
- Young GF, Scardovi L, Cavagna A, et al. (2013) Starling flock networks manage uncertainty in consensus at low cost. *PLoS Computational Biology* 9(1): e1002894. doi: [10.1371/journal.pcbi.1002894](https://doi.org/10.1371/journal.pcbi.1002894). Available at: <https://doi.org/10.1371/journal.pcbi.1002894>.
- Young GF, Scardovi L and Leonard NE (2011) Rearranging trees for robust consensus. In: Proceedings of the 50th IEEE conference on decision and control and European control conference, Orlando, FL, USA, 12-15 Dec. 2011, 1000–1005. DOI: [10.1109/CDC.2011.6161270](https://doi.org/10.1109/CDC.2011.6161270).
- Zahugi EMH, Shanta M and Prasad T (2012) Design of multi-robot system for cleaning up marine oil spill. *International Journal of Advanced Information Technology* 2(4): 33–43. DOI: [10.5121/ijait.2012.2404](https://doi.org/10.5121/ijait.2012.2404).
- Zhang Y, Luo Z and Liu B (2018) The effect of group size on the vigilance of Mongolian gazelle (*Procapra gutturosa*). *Folia Zoologica* 63(2): 81–86. DOI: [10.25225/fozo.v63.i2.a5.2014](https://doi.org/10.25225/fozo.v63.i2.a5.2014).
- Zhang K, Yang Z and Başar T (2021) Multi-agent reinforcement learning: a selective overview of theories and algorithms. In: Vamvoudakis K. G., Wan Y., Lewis F. L., et al. (eds) *Handbook of Reinforcement Learning and Control (Studies in Systems, Decision and Control)*. Cham, Switzerland: Springer, 325, 321–384. Available at: https://doi.org/10.1007/978-3-030-60990-0_12.
- Zhong YD, Srivastava V and Leonard NE (2021) Influence spread in the heterogeneous multiplex linear threshold model. *IEEE Transactions on Control of Network Systems (Early access)* 1. <https://ieeexplore.ieee.org/document/9454288>