



Analysis and control of agreement and disagreement opinion cascades

Alessio Franci¹ · Anastasia Bizyaeva² · Shinkyu Park² · Naomi Ehrich Leonard²

Accepted: 19 April 2021 / Published online: 21 May 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

We introduce and analyze a continuous time and state-space model of opinion cascades on networks of large numbers of agents that form opinions about two or more options. By leveraging our recent results on the emergence of agreement and disagreement states, we introduce novel tools to analyze and control agreement and disagreement opinion cascades. New notions of agreement and disagreement centrality, which depend only on network structure, are shown to be key to characterizing the nonlinear behavior of agreement and disagreement opinion formation and cascades. Our results are relevant for the analysis and control of opinion cascades in real-world networks, including biological, social, and artificial networks, and for the design of opinion-forming behaviors in robotic swarms. We illustrate an application of our model to a multi-robot task-allocation problem and discuss extensions and future directions opened by our modeling framework.

Keywords Opinion dynamics · Centrality indices · Complex contagions · Networked control · Robot swarms · Task allocation

This research has been supported in part by NSF grant CMMI-1635056, ONR grant N00014-19-1-2556, ARO grant W911NF-18-1-0325, DGAPA-UNAM PAPIIT grant IN102420, Conacyt grant A1-S-10610, and by NSF Graduate Research Fellowship DGE-2039656. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

✉ Alessio Franci
afranci@ciencias.unam.mx

Anastasia Bizyaeva
bizyaeva@princeton.edu

Shinkyu Park
shinkyu@princeton.edu

Naomi Ehrich Leonard
naomi@princeton.edu

¹ Math Department, National Autonomous University of Mexico, 04510 Mexico City, Mexico

² Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

1 Introduction

Network cascades are classically modeled as the spread of a binary transition from an inactive (e.g., uninformed, not infected) to an active (e.g., informed, infected) state. Models have been used to examine a variety of phenomena, including the spread of fads (Gladwell 2006; Bikhchandani et al. 1992), competing technologies (Arthur 1989), diffusion of innovations (Valente 1995; Rosa and Giua 2013), and the spread of influence (Kempe et al. 2003). The linear threshold model (LTM), introduced in Granovetter (1978); Schelling (2006), defines a basic modeling hypothesis for studying network cascades, namely discrete time and binary states of agents that update as a threshold function of the states of neighbors (Watts 2002; Lim et al. 2015; Garulli et al. 2015; Rossi et al. 2017). Recently, the LTM was extended to multi-layer interaction networks (Yağan and Gligor 2012; Salehi et al. 2015; Zhong et al. 2020).

In an opinion formation context, the situation is more general. When agents form opinions about a set of options, the transition from an unopinionated to an opinionated state is not a binary process because the opinionated state can correspond to a variety of distinct opinion configurations. Even in the simplest case of two options, the opinionated state can evolve in two qualitatively distinct directions, that is, preferring option A or preferring option B. Thus, for opinion cascades, it not only does the active or inactive nature of an agent's neighbors matter but also the opinion state of neighbors, i.e., their preferences among options. That is, the opinion state to which an agent will converge during an opinion cascade depends not only on the agent's susceptibility to forming an opinion (i.e., its activation threshold) but also the agent's tendency to align with, or reject, its neighbors' opinions. Further, the transition from being unopinionated to holding a strong opinion in favor or disfavor of one or more options can be more or less continuous, both in time and in strength, and the graded transition from one opinion state to another cannot be captured in binary models. A continuous time and state-space approach to threshold cascade behavior was explored in Zhong and Leonard (2019). A multi-option, agent-based approach to opinion cascades was studied in Pilditch (2017). The approach followed in the present paper considers both the presence of multiple options and multiple possible opinion states and a process that is continuous both in time and in state.

In recent work (Bizyaeva et al. 2020a), we developed a general model of continuous time and state-space opinion formation for an arbitrary number of agents exchanging information over a communication network and forming opinions about an arbitrary number of options. Our model is both analytically tractable and general because it is motivated by a tractable and general model-independent theory of opinion formation (Franci et al. 2020). In particular, very general opinion-formation behaviors like consensus, dissensus, and various intermediate forms of agreement and disagreement can all be realized and controlled in our model with only a handful of key parameters with clear real-world interpretations (e.g., parameters that represent the agents' level of cooperativity or competitiveness).

In our model, opinion formation happens through a bifurcation. While much can be investigated with linear models (see, e.g., Mateo et al. (2019)), the nonlinearity of our model allows for richer possibilities. Notably, the resulting opinion-formation process can be continuous or switch-like, depending on the nature of the bifurcation and the way the model crosses it (Franci et al. 2020). Further, when options are close to equally valuable, there is multistability of either agreement or disagreement opinion configurations. In this case, the opinion-forming bifurcation breaks deadlock by pushing the state away from the neutral (deadlock) equilibrium and toward one of the stable opinion states. The

key (bifurcation) parameter of our model is the “attention” u_i that agent i pays to its neighbors in the communication network. This parameter can also be interpreted as the urgency to form an opinion that the agent perceives. When agents have low attention, they form opinions linearly according to the information-rich inputs they receive about the available options. Agents that receive no inputs (i.e., those that do not gather information about the options) remain neutral or unopinionated. Conversely, when agents have high attention, they form a robust opinion through nonlinear network interactions, independently of, but also influenced by, their inputs.

In this paper, we leverage the state-dependent attention dynamics introduced in Bizyaeva et al. (2020a) to show that a positive feedback coupling between attention and opinion strength at the single-agent level provides the basic mechanism underlying the emergence of tunable opinion cascades at the network level. The positive feedback makes single-agent opinion formation switch-like and provides it with a tunable threshold, like in binary models of opinion cascade, while also preserving the continuous-time, graded evolution of the opinion and attention states. We develop a detailed geometrical analysis of the single-agent attention-opinion dynamics for the case of three options, which extends the two-option geometrical analysis in Bizyaeva et al. (2020a). The two cases have similarities (i.e., hysteresis and switches between unopinionated and opinionated states) but also differences (i.e., the much richer menu of possible opinionated configurations for three options).

Because of the intrinsic multistability of opinion configurations, our model naturally captures “afterthoughts” during the opinion cascade process: an early opinionated agent can change its opinion under the influence of sufficiently strong signals from other agents. Binary threshold models do not exhibit the same degree of behavioral flexibility. Further, due to its continuous time and state-space nature, our model features both simple and complex contagion (in the sense of Centola and Macy (2007)). That is, along the same opinion cascade process, contagion can be simple for some agents and complex for others.

To analyze and control the emergent behavior of the cascade model on undirected network topologies, we introduce the notions of its agreement centrality, its disagreement centrality, and its signed disagreement centrality. Agreement centrality coincides with the classical notion of eigenvector centrality (Bonacich 1972, 2007). Disagreement and signed disagreement centralities are new. Our definitions are motivated by the tight connections existing between the patterns of opinion formation emerging at an opinion-forming bifurcation and the spectral properties of the underlying adjacency matrix (Bizyaeva et al. 2020b). We show that our definitions provide powerful tools to analyze and control opinion cascades and their sensitivity to cascade seed inputs.

Our opinion-formation and opinion-cascade modeling, analysis, and control tools are naturally suited for engineering, biological, and sociopolitical applications. For example, biological and artificial swarms in motion can make both agreement decisions, such as to move in the same direction like schooling fish (Couzin et al. 2005, 2011; Leonard et al. 2012), and disagreement decisions, such as to go in different directions during search and foraging (Traniello 1989; Zedadra et al. 2017). Sociopolitical networks, both at the citizen and political elite level, are made of agents that can more or less rapidly change between agreement and disagreement behaviors, depending on the topic, time, and context (McCarthy 2019; Macy et al. 2019).

Two main types of collective decision making, across applications, can broadly be classified as “best-of- n ” and “task allocation” (Valentini et al. 2017; Gerkey and Mataric 2004; Brambilla et al. 2013; Schranz et al. 2020). In the former, agents must reach a consensus about the best-of- n options, while in the latter, agents must allocate across the n options. Thus, solving a best-of- n problem requires a specific type of agreement opinion formation.

Indeed, an earlier special case of our model was originally inspired by best-of- n collective decision making in house hunting honeybees (Gray et al. 2018).

Alternatively, solving a task-allocation problem requires a specific type of disagreement opinion formation. In this paper, we illustrate this fact with an original example describing distributed allocation over three tasks by a swarm of robots. With this example, we highlight the importance of task-allocation cascades. The result is a dynamical task-allocation algorithm controlled by just a few parameters. Given its continuous time and state-space nature, the algorithm is naturally suited to implementation in low-level (e.g., analog) hardware, making it relevant for any multi-robot application with restricted computational power.

The paper is structured as follows. In Sect. 2, we introduce the general opinion dynamics model of Bizyaeva et al. (2020a) and we rigorously analyze its opinion-forming behavior. In Sect. 3, we frame a multi-robot task-allocation problem using our opinion formation dynamics. In Sect. 4, we show how an attention positive feedback mechanism at the agent and network scales engenders opinion cascade with tunable sensitivity in our opinion formation dynamics. To analyze and control opinion cascades, in Sect. 5 we introduce new notions of centrality and illustrate their power mechanistically in low-dimensional agreement and disagreement opinion cascade examples. In Sect. 6, we show how introducing a cascade mechanism makes the multi-robot task-allocation model introduced in Sect. 3 reactive to changes in task urgency and more robust to initial conditions and parameter uncertainties. Application to large scale complex networks is presented in Sect. 7. Section 8 provides a discussion of results, our notion of agreement and disagreement centralities, and the connection between our notion of opinion cascade and complex contagion. Section 8 also describes extensions and future research directions opened up by our modeling framework, including time-varying network topologies and communication losses, and applications beyond multi-robot task allocation.

2 A general model of agreement and disagreement opinion formation

We consider a network of N_a agents forming opinions about N_o options. Let $z_{ij} \in \mathbb{R}$ denote the opinion of agent i about option j . The greater z_{ij} is, the stronger is the preference of agent i for option j . We assume that

$$z_{i1} + \dots + z_{iN_o} = 0. \quad (1)$$

This assumption means that we are only interested in modeling an agent's *relative* preference for the various options. Indeed, to determine if agent i prefers option j over option l , $l \neq j$, only the value of z_{ij} relative to z_{il} matters. Condition (1) can also be interpreted as a linearized simplex condition, as detailed in Franci et al. (2020).

Let $\mathbf{Z}_i = (z_{i1}, \dots, z_{iN_o})$ be the opinion state of agent i and $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_{N_a})$ the opinion state of the agent network. We recall some useful definitions from Franci et al. (2020); Bizyaeva et al. (2020a). Agent i is neutral if $\mathbf{Z}_i = \mathbf{0}$. The origin $\mathbf{Z} = \mathbf{0}$ is called the network's neutral state. Agent i is unopinionated if its opinion state is small, i.e., $\|\mathbf{Z}_i\| \leq \vartheta$, for a fixed threshold $\vartheta \geq 0$. Agent i is opinionated if $\|\mathbf{Z}_i\| \geq \vartheta$. If $z_{ij} \geq z_{il} - \vartheta$ for all $l \neq j$, then agent i is said to favor option j . If $z_{ij} < z_{il} - \vartheta$ for all $l \neq j$, then agent i is said to disfavor option j . An agent is conflicted about a set of options if it has near equal and favorable opinions of them all. Agents can agree and disagree. Two agents agree when they are opinionated and have the same qualitative opinion state (e.g., they both favor

the same option). They disagree when they have qualitatively different opinion states. We refer the reader to Franci et al. (2020) for a geometric illustration of these definitions. The network is in an agreement state when all agents are opinionated and agree and in a disagreement state when at least two agents disagree.

Agents' opinions are assumed to evolve in continuous time. The networked opinion dynamics (Bizyaeva et al. 2020a) are the following:

$$\dot{z}_{ij} = F_{ij}(\mathbf{Z}) - \frac{1}{N_o} \sum_{l=1}^{N_o} F_{il}(\mathbf{Z}) \quad (2a)$$

$$F_{ij}(\mathbf{Z}) = -d_{ij}z_{ij} + u_i \left(S_1 \left(\sum_{k=1}^{N_a} A_{ik}^{jj} z_{kj} \right) + \sum_{\substack{l \neq j \\ l=1}}^{N_o} S_2 \left(\sum_{k=1}^{N_a} A_{ik}^{jl} z_{kl} \right) \right) + b_{ij}, \quad (2b)$$

where \dot{z}_{ij} is the rate of change, with respect to time t , of the opinion of agent i about option j . Observe that, by (2a), $\dot{z}_{i1} + \dots + \dot{z}_{iN_o} = 0$, which ensures that $z_{i1}(t) + \dots + z_{iN_o}(t) = 0$ for all $t > 0$, provided $z_{i1}(0) + \dots + z_{iN_o}(0) = 0$. The parameter $d_{ij} > 0$ models the resistance of agent i to form an opinion about option j . Exogenous inputs (e.g., prior biases or incoming information) influencing the opinion formation of agent i about option j are captured in the affine parameter b_{ij} . Network interactions are modeled in the *adjacency tensor* A . There are four types of network interactions represented by elements of the adjacency tensor as follows:

1. Intra-agent, same-option (with weights A_{ii}^{jj}), modeling the influence of agent i 's opinion about option j on itself;
2. Intra-agent, inter-option (with weights A_{ii}^{jl}), modeling the influence of agent i 's opinion about option l on the same agent's opinion about option j ;
3. Inter-agent, same-option (with weights A_{ik}^{jj}), modeling the influence of agent k 's opinion about option j on agent i 's opinion about the same option;
4. Inter-agent, inter-option (with weights A_{ik}^{jl}), modeling the influence of agent k 's opinion about option l on agent i 's opinion about option j .

Network interactions can be excitatory, if the associated weight is positive, or inhibitory, if the associated weight is negative. $S_q : \mathbb{R} \rightarrow [-k_{q1}, k_{q2}]$ with $k_{q1}, k_{q2} \in \mathbb{R}^{>0}$ for $q \in \{1, 2\}$ is a generic sigmoidal saturating function satisfying constraints $S_q(0) = 0$, $S'_q(0) = 1$, $S''_q(0) \neq 0$, $S'''_q(0) \neq 0$. The two sigmoids S_1 and S_2 saturate same-option interactions and inter-option interactions, respectively. They are the only nonlinearity of the model.

The *attention parameter* u_i tunes how strongly nonlinear network interactions affect agent i (Fig. 1):

- For small u_i , agent i 's opinion formation is dominated by a linear behavior determined by d_{ij} and b_{ij} ;
- For large u_i , agent i 's opinion formation is dominated by a nonlinear behavior determined by network interactions.




Fig. 1 Schematic model description.

In the linear regime, an agent develops an opinion proportional to the ratio of its inputs and its resistance, much like a classical leaky evidence integrator (Bogacz 2007). In the nonlinear regime, much richer network opinion formation behaviors can be expected. Model (2) was inspired by and specializes to the honeybee-inspired model in Gray et al. (2018). See also Bizyaeva et al. (2020a) for a discussion of the similarities and difference between model (2) and previous opinion formation models in the literature.

2.1 Agreement and disagreement opinion formations are determined by network spectral properties

Model (2) contains hundreds of intrinsic and interaction parameters when N_a and/or N_o are sufficiently large. As such, it is intractable in general. However, the model-independent theory in Franci et al. (2020) shows that even in the simplest case in which all four types of network interaction are all-to-all and homogeneous, inputs are zero, attention is uniform across the agents, and the resistance is the same for all the agents, a plethora of rich prototypical opinion formations are still possible in model (2).

To study the role of complex inter-agent network interactions, while preserving analytical tractability, we consider the *homogeneous regime* of model (2), defined by imposing the following restrictions on the model parameters:

$$\begin{aligned} d_{ij} &= d > 0, \quad u_i = u, \quad A_{ii}^{jj} = \alpha \in \mathbb{R} \\ A_{ii}^{jl} &= \beta \in \mathbb{R}, \quad A_{ik}^{jj} = \gamma \tilde{a}_{ik}, \quad A_{ik}^{jl} = \delta \tilde{a}_{ik} \end{aligned} \quad (3)$$

where $\alpha, \beta, \gamma, \delta \in \mathbb{R}$ and $\tilde{a}_{ik} \in \{0, 1\}$, for all $i, k = 1, \dots, N_a$, $i \neq k$, and for all $j, l = 1, \dots, N_o$, $j \neq l$. The resulting dynamics are

$$\dot{z}_{ij} = F_{ij}(\mathbf{Z}) - \frac{1}{N_o} \sum_{l=1}^{N_o} F_{il}(\mathbf{Z}) \quad (4a)$$

$$F_{ij}(\mathbf{Z}) = -dz_{ij} + u \left(S_1 \left(\alpha z_{ij} + \sum_{\substack{k \neq i \\ k=1}}^{N_a} \gamma \tilde{a}_{ik} z_{kj} \right) + \sum_{\substack{l \neq j \\ l=1}}^{N_o} S_2 \left(\beta z_{il} + \sum_{\substack{k \neq i \\ k=1}}^{N_a} \delta \tilde{a}_{ik} z_{kl} \right) \right) + b_{ij}. \quad (4b)$$

In this regime, the intra-agent network is all-to-all and homogeneous, with weight α for same-option interactions and weight β for inter-option interactions. The topology of the inter-agent network is encoded in the adjacency matrix $\tilde{A} = [\tilde{a}_{ik}] \in \mathbb{R}^{N_a} \times \mathbb{R}^{N_a}$. For all existing inter-agent links, i.e., those satisfying $\tilde{a}_{ik} = 1$, the inter-agent same-option weight is γ and the inter-agent inter-option weight is δ . It is natural to assume that

$$\alpha \geq 0, \quad \text{and} \quad \beta < 0, \quad (5)$$

which means that opinions are self-reinforcing ($\alpha \geq 0$) and mutually exclusive ($\beta < 0$). The two remaining parameters, γ and δ , determine whether agents cooperate or compete in forming opinions. More precisely, if $\gamma > \delta$, agents cooperate in forming opinions, since the opinion of any agent i about any option j is more strongly excited by the opinions of its neighbors about the same option j than by their opinions about all the other options $l \neq j$. Conversely, if $\gamma < \delta$, agents compete in forming opinions, since the opinion of any agent i about any option j is more strongly excited by the opinions of its neighbors about all the other options $l \neq j$ than by their opinions about the same option j .

Remark 1 Because intra-agent coupling is all-to-all and homogeneous in model (4), options are indistinguishable from a network topology perspective. A natural extension of model (4) would be to consider a topologically richer intra-agent network. This generalization is natural in the presence of distinguished options or “issues,” i.e., subgroups of more tightly related options.

For homogeneous networks, the cooperative or competitive nature of agent interactions is the key determinant of the opinion formation behavior. The proof of the following theorem follows by linearizing the opinion dynamics (4) at the neutral equilibrium and by studying the kernel of the resulting Jacobian matrix (see (Biziyaeva et al. 2020b, Theorem 1) for details). Let $\mathbb{R}\{v_1, \dots, v_k\}$ be the span of vectors $v_1, \dots, v_k \in \mathbb{R}^n$ and $\mathbf{1}_n^\perp$ be the orthogonal complement to the vector of ones in \mathbb{R}^n .

Theorem 1 Suppose that $b_{ij} = 0$ for all $i = 1, \dots, N_a$ and all $j = 1, \dots, N_o$.

- If $\gamma > \delta$, then in model (4) opinion formation emerges as a bifurcation from the neutral state for

$$u = u_a := \frac{d}{\alpha - \beta + \lambda_{\max}(\gamma - \delta)} \quad (6)$$

where λ_{\max} is the maximum eigenvalue of \tilde{A} . All bifurcation branches are tangent to the $(N_o - 1)$ -dimensional kernel $\mathbb{R}\{v_{\max}\} \otimes \mathbf{1}_{N_o}^\perp$, where v_{\max} is the positive (Perron–Frobenius) eigenvector associated with λ_{\max} . More precisely, each bifurcation branch is tangent to a one-dimensional subspace $\mathbb{R}\{v_{\max} \otimes v_{ax}\}$ of $\mathbb{R}\{v_{\max}\} \otimes \mathbf{1}_{N_o}^\perp$, for some

$v_{ax} \in \mathbf{1}_{N_o}^\perp$. In particular, because every agent's opinion state is directly proportional (according to the entries of v_{max}) to the same opinion configuration vector v_{ax} , all the opinion formation bifurcation branches are of agreement type.

- If $\gamma < \delta$ and the minimum eigenvalue of \tilde{A} is simple¹, then in model (4) opinion formation emerges as a bifurcation from the neutral state for

$$u = u_d := \frac{d}{\alpha - \beta + \lambda_{\min}(\gamma - \delta)} \quad (7)$$

where λ_{\min} is the minimum eigenvalue of \tilde{A} . All bifurcation branches are tangent to the $(N_o - 1)$ -dimensional kernel $\mathbb{R}\{v_{\min}\} \otimes \mathbf{1}_{N_o}^\perp$, where v_{\min} is the eigenvector associated with λ_{\min} . Each bifurcation branch is tangent to a one-dimensional subspace $\mathbb{R}\{v_{\min} \otimes v_{ax}\}$ of $\mathbb{R}\{v_{\min}\} \otimes \mathbf{1}_{N_o}^\perp$, for some $v_{ax} \in \mathbf{1}_{N_o}^\perp$. Because the vector v_{\min} is orthogonal to v_{max} , it has mixed-sign entries. In particular, some agents have opinion state directly proportional to the opinion configuration vector v_{ax} , whereas some agents have opinion state inversely proportional to it, and so all the opinion formation bifurcation branches are of disagreement type.

Proof Equations (6) and (7) were derived in Bizyaeva et al. (2020a) [Theorem 4.1] and Bizyaeva et al. (2020b) [Theorem 1]. The form of the one-dimensional subspaces along which agreement and disagreement bifurcations occur, in particular, the vectors v_{ax} , can be constructed using equivariant bifurcation theory methods (Golubitsky et al. 1988; Golubitsky and Stewart 2002). Because the intra-agent opinion network is all-to-all and homogeneous, model (4) is symmetric with respect to arbitrary option permutations, i.e., it is S_{N_o} -equivariant, where S_{N_o} is the symmetric group of order N_o . It follows from the Equivariant Branching Lemma (Golubitsky et al. 1988, Theorem XIII.3.3) that, generically, v_{ax} is the generator of the fixed point subspace of one of the axial subgroups of the action of S_{N_o} on $\mathbf{1}_{N_o}^\perp$. There is, of course, only one (modulo multiplication by a nonzero scalar) such vector for $N_o = 2$, i.e., $v_{ax} = (1, -1)$. Modulo multiplication by a nonzero scalar, there are three such vectors for $N_o = 3$, i.e., $v_{ax} = (1, 1, -2)$, $v_{ax} = (1, -2, 1)$, $v_{ax} = (-2, 1, 1)$. The form of v_{ax} for $N_o > 3$ can be similarly derived. \square

Although derived for the simplified homogeneous dynamics (4), Theorem 1 is predictive of the model behavior in the presence of small heterogeneities in the coupling parameter, i.e., when the homogeneous regime is weakly violated. Also, the same bifurcation techniques used to prove this theorem generalize to the case in which large heterogeneities are present. See, e.g., (Bizyaeva et al. 2020a, Sect. 3.C) for a general clustering result in the presence of large difference in the coupling parameter between two or more agent subpopulations.

Remark 2 In (4), we can bring the sums appearing inside S_1 and S_2 outside of the sigmoids without changing the qualitative behavior of the model. In particular, because the model Jacobian does not change if sums and sigmoids are swapped, Theorem 1 remains true.

¹ The assumption that the minimum eigenvalue of \tilde{A} has algebraic multiplicity one is usually true for all sufficiently large complex networks. The case in which this multiplicity is larger can also be analyzed, but for simplicity we do not address it here.




Fig. 2 Agreement (left) and disagreement (right) opinion formation in Watts–Strogatz networks with $N_a = 100$ agent nodes and **a** $N_o = 2$ and **b** $N_o = 3$ options. The Watts–Strogatz rewiring probability is 0.9. Other parameters: $\alpha = 0.2$, $\beta = -0.5$, $\gamma = -\delta = 0.1$ (agreement), $\gamma = -\delta = -0.1$ (disagreement). Initial conditions and inputs are drawn from a normal distribution with zero mean and variance 10^{-3} . In all the simulations, $S_1(x) = \tanh(x + 0.5 \tanh(x^2))$ and $S_2(x) = 0.5 \tanh(2x + \tanh(x^2))$.

In Figs. 2 and 3, we illustrate the prediction of Theorem 1 for two types of (undirected) complex networks: a Watts–Strogatz “small-world” network (Watts and Strogatz 1998) and a Barabasi–Albert “scale-free” network (Albert and Barabasi 2002), respectively. Panel A of both figures corresponds to $N_o = 2$ and panel B to $N_o = 3$. In the left plots of each panel, agents are cooperative, whereas in the right plots, agents are competitive. In all the simulations $N_a = 100$ and small Gaussian-distributed random perturbations were added to all inputs b_{ij} to verify the robustness of the theory. Theorem 1 is indeed valid for $b_{ij} = 0$ but, by normal hyperbolicity of the predicted bifurcation branches of agreement and disagreement equilibria, its predictions remain valid in the case where the b_{ij} are not exactly zero or other types of perturbations (e.g., coupling weights perturbations that weakly violate the homogeneous regime assumption) are present. Precise robustness bounds can be computed using the methods introduced in Wang and Michel (1994). This observation makes Theorem 1 relevant for real-world applications.

Each plot in Figs. 2 and 3 shows the temporal evolution of z_{i1} , i.e., agent i ’s opinion about the first option, as the attention parameter is slowly increased from below to above the agreement and disagreement opinion-formation critical values u_a and u_d defined in (6) and (7), respectively. The critical value is reached at half of the simulation time. The insets show the distribution of agent opinions about all the options at one third of the simulation time (i.e., before the opinion-forming bifurcation), two thirds




Fig. 3 Agreement (left) and disagreement (right) opinion formation in Barabási–Albert networks with $N_a = 100$ agent nodes and **a** $N_o = 2$ options and **b** $N_o = 3$ options. The Watts–Strogatz rewiring probability is 0.9. Other parameters: $\alpha = 0.2$, $\beta = -0.5$, $\gamma = -\delta = 0.1$ (agreement), $\gamma = -\delta = -0.1$ (disagreement). Initial conditions and inputs are drawn from a normal distribution with zero mean and variance 10^{-3} .

of the simulation time (i.e., right after the opinion-forming bifurcation), and at the end of the simulation time. The first (second, third) line shows the distribution of the agent opinions about the first (second, third) option. Colors indicate the strength of favor/dis-favor for a given option, with light blue (resp., red) indicating weak favor (resp. dis-favor) and dark blue (resp., red) indicating strong favor (resp., disfavor). Asymmetric distributions on the right (left) indicate that agents agree in favoring (disfavoring) the option. A roughly symmetric distribution indicates disagreement about an option, with some agents favoring it and other agents disfavoring it. A sharply bimodal distribution corresponds to *polarization* of opinions about the option.

The model behavior follows our theoretical predictions. When agents cooperate, the neutral equilibrium bifurcates at $u = u_a$ into an agreement branch where all the agents share the same qualitative opinion. The larger is u relative to u_a , the stronger are the emerging opinions. In Figs. 2a and 3a, left, all the agents favor Option 1 and disfavor Option 2. In Fig. 2b, left, all agents disfavor Option 2 and are conflicted between Options 1 and 3. In Fig. 3b, left, all agents favor Option 1 and are conflicted between Options 2 and 3. Conversely, when agents compete, the bifurcation branches emerging from the neutral equilibrium are of disagreement type and again opinion strength increases as u increases above u_d . In Figs. 2a and 3a, right, roughly half of the agents favor Option 1, while the other half favor Option 2. In Figs. 2b and 3b, right, for each of

the three options the group is roughly split into agents that favor the option and agents that disfavor it.

Remark 3 The precise effect of the inputs b_{ij} on the bifurcation and steady-state behavior can be computed using Lyapunov–Schmidt reduction techniques (Golubitsky and Schaeffer 1985). More precisely, the projection of the input vector onto the bifurcation kernel predicts, to first order, in which direction along the one-dimensional kernel’s subspaces, $\mathbb{R}\{v_{\max} \otimes v_{ax}\}$ (agreement case) or $\mathbb{R}\{v_{\min} \otimes v_{ax}\}$ (disagreement case) defined in Theorem 1, the agreement or disagreement bifurcation unfolds (see (Gray et al. 2018, Theorem 1) for the agreement, two-option case). At bifurcation, the unfolding direction determines which options are favored or disfavored by the inputs (see, e.g., (Gray et al. 2018, Fig. 2)). Here we generalize (Gray et al. 2018, Theorem 1) to the disagreement and three-option cases.

3 Multi-robot task allocation as disagreement opinion formation

We illustrate using a detailed example how our opinion dynamics (4) can be applied in robot swarm coordination and control, in particular, to a distributed task-allocation problem. Given its dynamical nature and the small number of parameters, the proposed distributed control algorithm can be implemented in analog circuits and is thus a good candidate for low-computational-power robot swarms that require dynamic task allocation. A single-agent motivation dynamics framework was also recently proposed and applied to a task coordination problem for a mobile robot (Reverdy and Koditschek 2018; Reverdy et al. 2021). Honeybee-inspired Gray et al. (2018); Reina et al. (2017) cross-inhibitory interactions were recently proposed Talamali et al. (2019) to improve decision accuracy in a multi-robot best-of- n decision problem. In the spirit of these works, the distributed opinion dynamics of Bizyaeva et al. (2020a) allows us to move beyond a single robotic unit and consider distributed task coordination in a robotic swarm.

Consider that a robot swarm is given N_o tasks to complete, in which each task has a distinct priority level and each robot may have its own preference to work on certain tasks. To formalize, we denote by a positive constant $\mu_j \in [0, 1]$, $j = 1, \dots, N_o$, satisfying $\sum_{j=1}^{N_o} \mu_j = 1$, the priority of each task j . The main objective in this example is to design a simple distributed decision-making algorithm that assigns robots to each task. The algorithm should be designed to respect task priority, i.e., the larger the task priority μ_j , the larger the number of robots assigned to that task by the algorithm.

Given an undirected communication network between the robots, with adjacency matrix $[\tilde{a}_{ik}]_{i,k=1}^{N_a}$, let \mathcal{N}_i be set of neighbors of robot i . We adapt opinion dynamics (4) to the task-allocation problem as follows. The preference that robot i gives to performing task j evolves accordingly to vector field

$$\dot{z}_{ij} = F_{ij}(\mathbf{Z}) - \frac{1}{N_o} \sum_{l=1}^{N_o} F_{il}(\mathbf{Z}) \quad (8a)$$

$$F_{ij}(\mathbf{Z}) = -z_{ij} + u \left(\mu_j (|\mathcal{N}_i| + v_{ij}) - \frac{1}{2} \sum_{k \in \mathcal{N}_i} S(2\tilde{\gamma} z_{kj}) \right). \quad (8b)$$

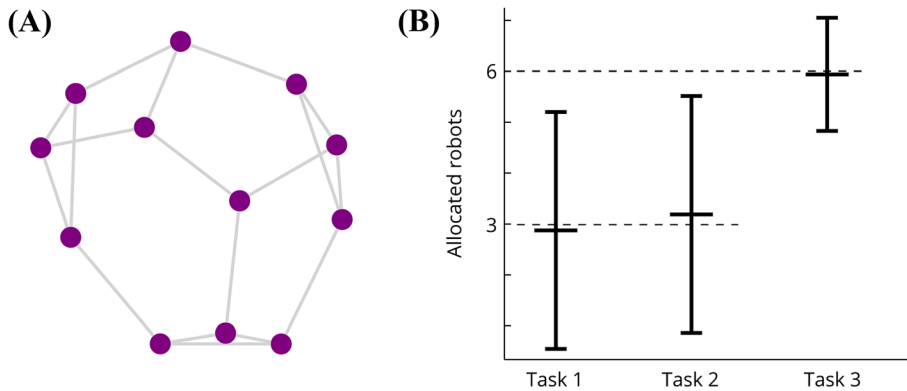


Fig. 4 **a** The Frucht graph. **b** Average (large bar) and standard deviation (short bar) of the number of agents allocated to Tasks 1,2,3 over one-hundred instances of our task-allocation algorithm. Details in the main text.

The positive parameter v_{ij} is the intrinsic zealotness of robot i to perform task j . As we will see, the zealotness parameter can be used as control input to the robot swarm to trigger the distributed task allocation. Opinion dynamics (4) specialize to model (8) by letting $b_{ij} = u\mu_j(|\mathcal{N}_i| + v_{ij})$, $\alpha = \beta = \delta = 0$, $\tilde{\gamma} = -\gamma > 0$, and by swapping the inter-agent interaction sum with the sigmoid S_1 (as discussed in Remark 2, this choice does not affect the qualitative behavior of the model).

The rationale behind the derivation of model (8) is the following. Let N_i^j be the number of robots in \mathcal{N}_i that express preference for task j , i.e., such that $z_{kj} > 0$. Because generically $z_{kj} \neq 0$ and if the upper and lower limits of S are ± 1 , respectively, for sufficiently large γ the term $\frac{1}{2} \sum_{k \in \mathcal{N}_i} S(2\gamma z_{kj}) \approx \frac{1}{2} (N_i^j - (|\mathcal{N}_i| - N_i^j)) = N_i^j - \frac{|\mathcal{N}_i|}{2}$. Imposing the simplex condition by subtracting the average opinion drift as in (8a), the constant $\frac{|\mathcal{N}_i|}{2}$ disappears from the opinion dynamics on the simplex and, for $v_{ij} = 0$ the effective interaction term reduces to

$$\mu_j |\mathcal{N}_i| - \frac{1}{2} \sum_{k \in \mathcal{N}_i} S(\epsilon z_{kj}) \approx \mu_j |\mathcal{N}_i| - N_i^j. \quad (9)$$

In other words, in our task-allocation model, each robot updates its preference for task j roughly proportionally to the priority of task j , scaled by the number of its neighbors, minus the number of its neighbors that already show a preference for the same task. A nonzero zealotness v_{ij} increases the tendency of robot i to choose task j .

Observe that the condition $-\tilde{\gamma} = \gamma < 0$ and $\delta = 0$ imposes that the model is in the disagreement regime for $u > u_d > 0$. Thus, in our model, the robots naturally develop different opinions about the tasks, i.e., they self-allocate across the task in a distributed fashion. Furthermore, robot task-allocation will reflect task priorities, which, from a bifurcation viewpoint, breaks the symmetry between the tasks.

We illustrate the performance of our task-allocation algorithm for a swarm of $N_a = 12$ robots allocating themselves across $N_o = 3$ tasks. We assume that the inter-agent communication graph is the Frucht graph (Frucht 1939) (see Fig. 4a), which is 3-regular ($|\mathcal{N}_i| = 3$ for all $i = 1, \dots, N_a$) but possesses no non-trivial symmetries, i.e., all robots can be distinguished by their position in the graph. Our algorithm is particularly suited in *a priori* dead-locked situations, i.e., those in which two or more tasks have the same objective priority. In

our example, we set $\mu_1 = \mu_2 = 0.3$ and $\mu_3 = 0.4$. We also set $\tilde{\gamma} = 1.0$ and $u = 2u_d$. For the sigmoid function we use $S(x) = \tanh(x + 0.5 \tanh(x^2))$, where the $\tanh(x^2)$ ensure genericity (i.e., nonvanishing even partial derivatives) of the S at the origin. We run one hundred repetitions of our example with small random zealotness parameters v_{ij} (modeling parametric uncertainties) and small random initial conditions, and computed the resulting average and standard deviation of the number of robots allocated to each option.

Robot i is considered allocated to task j if at the end of the simulation $\|Z_i\| \geq \theta$ and $z_{ij} > z_{il}$ for all $l \neq j$, where θ is a control threshold parameter that determines whether a robot became opinionated (in the simulations, $\theta = 0.1$). The final simulation time can be considered as a tunable decision deadline for the task-allocation decision. Figure 4b shows the results of our analysis. On average, a larger number of agents are allocated to the highest-priority task. A roughly equal number of agents is allocated on average to the less urgent tasks. The standard deviation of the allocation number is also smaller for the highest-priority task, which indicates that our algorithm consistently allocates the largest number of robots to the highest-priority task. Of course, there is a lot of room for improving the simple algorithm proposed in this section, in particular, in order to drastically reduce variability in the number of agents allocated to the various tasks, and thus increase robustness to initial conditions and parameter uncertainties. In some instances of our example, depending on initial conditions or the intrinsic zealotness parameters, no robots were allocated to either Task 1 or Task 2, which is clearly an undesirable situation in practice. We suggest a strategy to overcome these limitations in Sect. 6.

4 Opinion cascades with tunable sensitivity from feedback attention dynamics

In this section, we couple the opinion dynamics (4) with the attention dynamics introduced in Bizyaeva et al. (2020a). Using a detailed geometrical analysis for the case $N_o = 3$ (the case $N_o = 2$ is studied in Bizyaeva et al. (2020a)), we show that the resulting positive feedback between attention and opinion strength allows an agent to transition from being in a stable unopinionated state to a stable opinionated state in response to an input $\mathbf{b}_i = (b_{i1}, \dots, b_{iN_o})$ sufficiently large in magnitude. The associated activation threshold is implicitly defined by the geometric properties of the dynamics and is explicitly tunable with the model parameters. The threshold effect mediated by the positive feedback between attention and opinion strength scales up to networks of interconnected agents under the influence of inputs on a small number of “seed nodes.” At the network level, the coupling between opinion and attention dynamics creates a threshold for the input magnitude below which the network remains close to neutral and above which an opinion cascade is triggered.

4.1 Feedback attention dynamics and tunable sensitivity of opinion formation at the single agent level

To understand the effect of coupled opinion and attention dynamics, we first consider an agent without any neighbors. Even without an inter-agent communication network, the dynamics of each agent’s opinions are organized by a bifurcation at which the agent develops preferences for the available options. When $\tilde{a}_{ik} = 0$ for all $i, k = 1, \dots, N_a$, the $F_{ij}(\mathbf{Z})$ function in (4b) becomes

$$F_{ij}(\mathbf{Z}) = -dz_{ij} + u_i \left(S_1(\alpha z_{ij}) + \sum_{\substack{l=1 \\ l \neq j}}^{N_o} S_2(\beta z_{il}) \right) + b_{ij}. \quad (10)$$

When $b_{ij} = 0$, for all $j = 1, \dots, N_o$, and $u_i < \frac{d}{\alpha - \beta}$, the neutral opinion state $\mathbf{Z}_i = \mathbf{0}$ of agent i is stable. At $u_i^* = \frac{d}{\alpha - \beta}$ branches of opinionated solutions emerge in an S_{N_o} -equivariant bifurcation, which is studied in detail in Elmhirst (2004). We omit the formal details of this bifurcation analysis here and simply note that for $u_i < u_i^*$ the agent's opinion converges to the neutral equilibrium with $\|\mathbf{Z}_i^*\| = 0$, whereas for $u_i > u_i^*$ it converges to some nonzero magnitude equilibrium $\|\mathbf{Z}_i^*\| > 0$, whose norm grows monotonically with u_i . The stable opinionated equilibria that appear as primary bifurcation branches correspond either to agent i choosing one option and rejecting all others, or to it being conflicted between a subset of the available options (Fig. 5a, left).

The dynamics of an agent i 's opinion with nonzero input \mathbf{b}_i directly follow from the bifurcation behavior of its opinion dynamics with $\mathbf{b}_i = \mathbf{0}$. As discussed in Sect. 2, when u_i is small, the input \mathbf{b}_i and resistance d define the linear opinion formation behavior of agent i . For large u_i , the nonlinear behavior dominates and agent i 's opinion can become much larger in magnitude than its input \mathbf{b}_i . Generically, when $\mathbf{b}_i \neq \mathbf{0}$, for values of u_i near u_i^* , a single branch of stable opinionated solutions is selected accordingly to \mathbf{b}_i , while all other solutions are removed (Fig. 5a, right).

How the opinion bifurcation diagram changes in response to inputs motivates the design of an attention feedback dynamic that increases an agent's attention from below to above its bifurcation point when sufficiently large inputs are present. We define the closed-loop attention dynamics as

$$\tau_u \dot{u}_i = -u + u_{\min} + (u_{\max} - u_{\min}) S_u(\|\mathbf{Z}_i\|) \quad (11)$$

where τ_u is an integration timescale and $S_u : \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ is a smooth saturating function, satisfying $S_u(0) = 0$, $S_u(y) \rightarrow 1$ as $y \rightarrow \infty$, $S'_u(y) > 0$ for all $y \in \mathbb{R}_{\geq 0}$, $S''_u(y) > 0$ for all $0 \leq y < u_{\text{th}}$ and $S''_u(y) < 0$ for all $y > u_{\text{th}}$, with $u_{\text{th}} > 0$. For the sake of numerical illustrations we use $S_u(y) = \frac{y^n}{u_{\text{th}}^n + y^n}$. The larger the n , the steeper the sigmoid around its threshold u_{th} . We assume $u_{\min} < u_i^* < u_{\max}$ and n sufficiently large so that the saturation is sufficiently steep.

The closed-loop attention dynamics (11) create a localized positive feedback between an agent's attention and its opinion strength, as illustrated in Fig. 5b. We use the chain rule to compute partial derivatives at $\mathbf{Z}_i = \mathbf{0}$ of the time-derivative of z_{ij}^2 with respect to attention u_i of an agent whose dynamics are defined by (11) and, conversely, the time-derivative of u_i (11) with respect to z_{ij}^2 . The sign of the product of these two partial derivatives gives the sign of the feedback between attention and opinion strength at the neutral equilibrium. Computing, we get

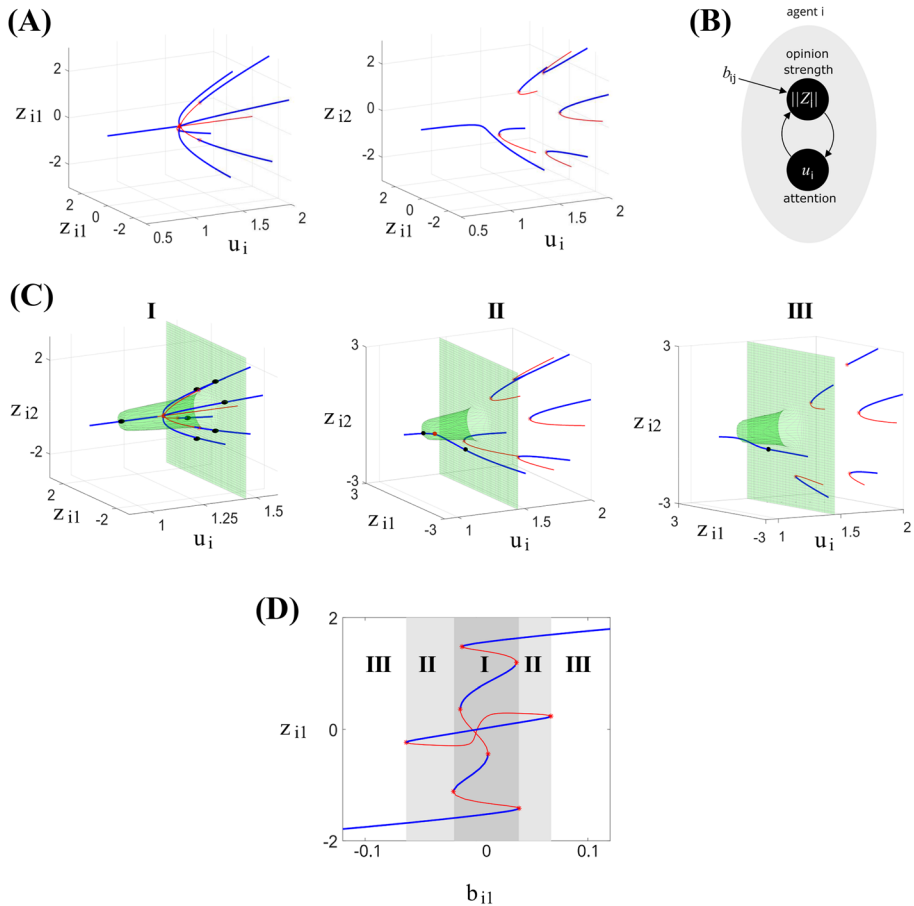


Fig. 5 **a** Bifurcation diagram of the single agent dynamics with respect to the attention parameter. Left: $b_i = 0$. Right: $b_i \neq 0$. **b** Schematic of the relationship between an agent's opinion strength, attention, and input. **c** Representative intersections of the attention nullclines (green surface) and opinion nullclines (red and blue lines) for the three parameter regimes with $N_o = 3$. **d** Bifurcation diagram of the coupled system (10),(11) with $N_o = 3$, projected to the b_{i1} , z_{i1} plane. In **d** blue (red) lines correspond to stable (unstable) equilibria of the coupled system. In **a**, **c** blue (red) lines correspond to stable (unstable) equilibria of the opinion dynamics (10) with a static attention parameter. Additional branches of unstable equilibria are not shown in **a**, **c–i** for diagram clarity. Black (red) spheres mark the stable (unstable) equilibria of the coupled system. Parameters: $d = 1$, $\alpha = 0.4$, $\beta = -0.4$, $u_{\min} = 1.05$, $u_{\max} = 1.45$, $n = 3$, $u_{th} = 0.15$. **B**: $b_{i2} = 0.005$, $b_{i3} = -0.002$, b_{i1} varied as illustrated. **A** (left) and **CI**: $\mathbf{b}_i = 0$; **A** (right) and **II**: $\mathbf{b}_i = (0.05, -0.01, 0.02)$; **III**: $\mathbf{b}_i = (0.1, -0.01, 0.02)$. Diagrams in **A**, **C**, **D** are generated with help of Mat-Cont numerical continuation package (Dhooge et al. 2008).

$$\frac{\partial(z_{ij}^2)}{\partial u_i} = 2z_{ij} \left(\frac{N_o - 1}{N_o} (S_1(\alpha z_{ij}) - S_2(\beta z_{ij})) - \frac{1}{N_o} \sum_{\substack{l \neq j \\ l=1}}^{N_o} (S_1(\alpha z_{il}) - S_2(\beta z_{il})) \right),$$

$$\frac{\partial \dot{u}_i}{\partial(z_{ij}^2)} = \frac{1}{\tau_u} (u_{\max} - u_{\min}) S'_u(\|Z\|).$$

Near the neutral equilibrium $Z_i = 0$ we can approximate $S_1(\alpha z_{il}) - S_2(\beta z_{il}) \approx (\alpha - \beta)z_{il}$, and using the simplex constraint on the opinions we find that for small Z_i and any $j = 1, \dots, N_o$, $\frac{\partial(z_{ij}^2)}{\partial u_i} \approx 2(\alpha - \beta)z_{ij}^2$. Therefore

$$\frac{\partial(z_{ij}^2)}{\partial u_i} \frac{\partial \dot{u}_i}{\partial(z_{ij}^2)} \approx \frac{2}{\tau_u} (\alpha - \beta)(u_{\max} - u_{\min}) z_{ij}^2 S'_u(\|Z\|). \quad (12)$$


The right-hand side of (12) is positive whenever $z_{ij} \neq 0$. Thus, a positive feedback loop between an agent's opinion strength and its attention is triggered as soon as its opinion deviates from zero.

The localized positive feedback between attention and opinion strength creates multi-stability and threshold behavior in the closed-loop opinion dynamics. Figures 5c and d illustrate the three-dimensional phase portrait and the steady-state behavior (bifurcation diagram) of the closed-loop opinion-attention dynamics for $N_o = 3$, respectively. The bifurcation diagram in Fig. 5d is projected onto the z_{i1}, b_{i1} plane. Note that the bifurcation diagram characterizes the input-output relationship, where the input is b_i and the output is the steady-state opinion.

We can distinguish three regimes (labels as in Fig. 5d):

1. *Multistability* $\|b_i^\perp\| < \varepsilon_1$. For small $\|b_i^\perp\|$, the opinion and attention nullclines intersect in many locations and a number of these intersections are stable equilibria. One of the intersections is near $u_i = u_{\min}$ and $Z_i = 0$ and is stable. The other stable equilibria correspond to intersections of opinionated solution branches with attention nullcline near $u_i = u_{\max}$. For small initial conditions, agent i 's opinion and attention remain small, attracted by the stable equilibrium near to $u_i = u_{\min}$. For large initial conditions, they approach one of the opinionated equilibria. In this parameter regime, the initial conditions determine whether agent i becomes opinionated, and also which opinion it forms.
2. *Bistability* $\varepsilon_1 < \|b_i^\perp\| < \varepsilon_2$. For intermediate $\|b_i^\perp\|$, the coupled dynamics admit only two stable equilibria: one is the unopinionated equilibrium near $u_i = u_{\min}$ and the other is an opinionated equilibrium on the solution branch favored by the input b_i . For small initial conditions, agent i 's opinion and attention remain small, whereas for large initial conditions they approach the stable equilibrium picked out by the input. In this regime, initial conditions determine whether or not an agent becomes opinionated, but the input governs the opinion formation outcome.
3. *Cascade* $\|b_i^\perp\| > \varepsilon_2$. For sufficiently large $\|b_i^\perp\|$, the coupled system has a unique stable equilibrium point which corresponds to an opinionated state favoring the opinion picked out by the input b_i . From any initial condition, an agent will reliably become opinionated in the direction informed by the input.

Fig. 6 Intra- and inter-agent positive feedback loops between the agents' opinion strengths and their attention.



The boundaries $\varepsilon_1, \varepsilon_2$ between these three regimes can be tuned by the parameters of the saturation function S_u in (11). In other words, the sensitivity of the closed-loop dynamics to forming opinions in response to inputs is tunable.

4.2 Network opinion cascades with tunable sensitivity

In the presence of symmetric network interactions, the intra-agent positive feedback loop between attention and opinion strength scales up to the network level, as illustrated in Fig. 6. Independent of their excitatory or inhibitory nature, network interactions create a positive feedback loop between the opinion strengths of any pair of interconnected agents, at least close to the neutral equilibrium where the interaction nonlinearities are steeper and thus the opinion interaction is stronger. To show this statement, we use the chain rule to compute partial derivatives at the neutral equilibrium of the time-derivative of z_{ij}^2 with respect to z_{kl}^2 for two agents i and k interacting according to (4) and for symmetric network coupling. Computing, we get

$$\left. \frac{\partial(z_{ij}^2)}{\partial z_{kl}^2} \right|_{Z=0} = \frac{1}{N_o} \frac{z_{ij}}{z_{kl}} u_i \tilde{a}_{ik} (\delta - \gamma), \text{ if } l \neq j, \quad \left. \frac{\partial(z_{ij}^2)}{\partial z_{kj}^2} \right|_{Z=0} = \frac{N_o - 1}{N_o} \frac{z_{ij}}{z_{kj}} u_i \tilde{a}_{ik} (\gamma - \delta)$$

Thus, if $\tilde{a}_{ik} = 1$, by symmetry,

$$\left. \frac{\partial(z_{ij}^2)}{\partial z_{kl}^2} \frac{\partial(z_{kl}^2)}{\partial z_{ij}^2} \right|_{Z=0} = \frac{1}{N_o^2} u_i u_k (\gamma - \delta)^2, \text{ if } l \neq j, \quad \left. \frac{\partial(z_{ij}^2)}{\partial z_{kj}^2} \frac{\partial(z_{kj}^2)}{\partial z_{ij}^2} \right|_{Z=0} = \left(\frac{N_o - 1}{N_o} \right)^2 u_i u_k (\gamma - \delta)^2, \quad (13)$$

which yields an element-wise positive feedback loop between the opinion strengths of agent i and agent k . As a consequence, when Eqs. (2) and (11) are coupled, the attention variables of agents i and k are also interconnected in an (indirect) positive feedback loop. This can easily be seen by computing the overall sign of the interconnection loop from u_i to u_k and back, passing through $||Z_i||$ and $||Z_k||$ (see Fig. 6). The existence of a networked positive feedback between the agents' attention variables and opinion strengths suggests the existence of a network threshold for the trigger of opinion cascades.

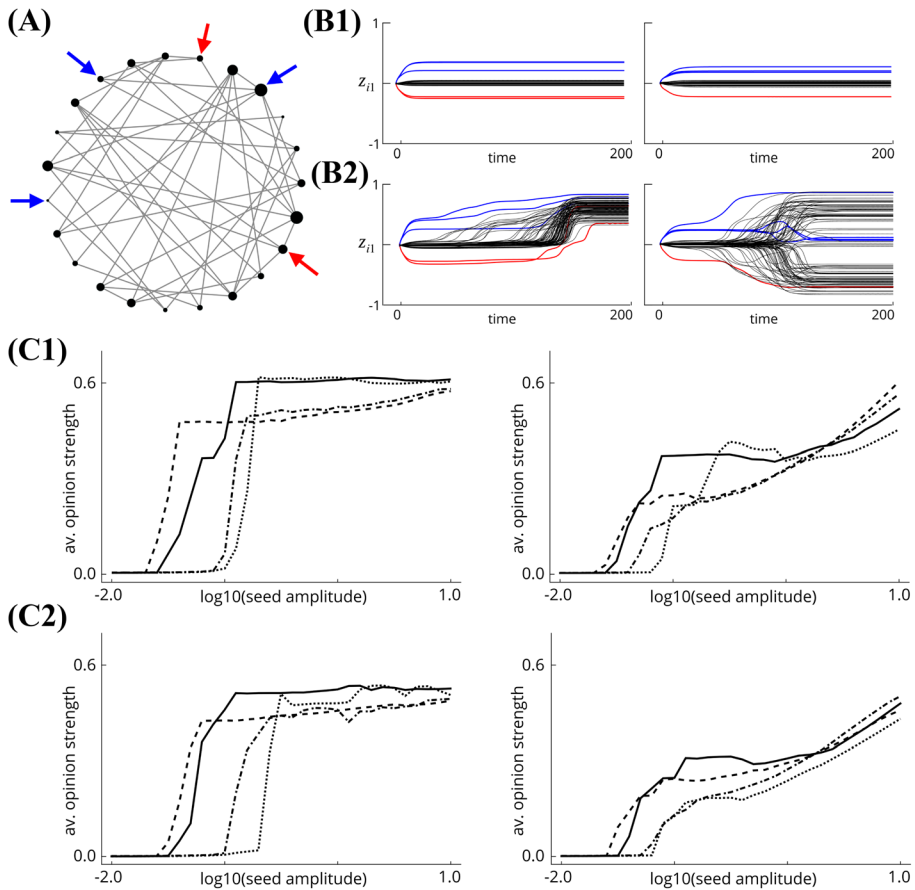


Fig. 7 Agreement and disagreement opinion cascades. **a** Schematic of the network structure. In a large network, only a small number of agents receive inputs, indicated by arrows. Inputs have the same strength, i.e., the same norm, but they can differ in the information they bring, i.e., which option they favor. In the numerical simulations, inputs are in favor (blue arrows) or in disfavor (red arrows) of Option 1. **b** Network response to inputs for cooperative (left) and competitive (right) agents. **b1** Small inputs (no cascade). **b2** Large inputs (cascade). **c** Cascade threshold analysis. The plots show how the average network opinion strength changes as a function of the seed input amplitude (in logarithmic scale) over ten network instances with random seed placement. $N_o = 2$ in **c1** and $N_o = 3$ in **c2** Left plots: Watts–Strogatz network. Right plots: Barabasi–Albert network. Full lines: cooperative agents, $u_{th} = 0.05$. Dashed line: competitive agents, $u_{th} = 0.05$. Dotted lines: cooperative agents, $u_{th} = 0.1$. Dashed-dotted lines: competitive agents, $u_{th} = 0.1$. Other parameters as in Figs. 2 (Watts–Strogatz case) and 3 (Barabasi–Albert case). In all simulations, $\delta_u = 0.2$.

An opinion cascade refers to the situation in which only a small number $N_{seeds} \ll N_a$ of agents, called seeds, receive nonzero inputs $b_1, \dots, b_{N_{seeds}}$ and yet a large part of the network develops a markedly nonzero opinion. In the network example of Fig. 7a, the cascade seeds are indicated with blue and red arrows depending on the options favored or disfavored by the input (see the figure caption for details). When the inputs to the seeds are small, only the seeds develop clear opinions aligned with their respective inputs, while the other agents remain weakly opinionated or unopinionated in such a way that the network

average opinion strength remains small (Fig. 7b1). When the inputs to the seeds are sufficiently strong, the seeds develop opinions that are sufficiently strong to engage the positive feedback of the network attention, which spreads to some or all of the rest of the agents with some variable delay. This yields an opinion cascade such that the network average opinion strength grows large (Fig. 7b2). Whether the cascade is toward agreement or disagreement depends on the cooperative (left plots) or competitive (right plots) nature of the agents. Observe that in both agreement and disagreement cases, a number of agents have afterthoughts: they develop an early opinion about Option 1 but either exhibit the opposite opinion or become neutral by the end of the cascade process.

In the following, we fix the number N_{seeds} and let all inputs $\mathbf{b}_1, \dots, \mathbf{b}_{N_{\text{seeds}}}$ have the same norm \bar{b} but still allow each to bring information either in favor or in disfavor of Option 1. In the attention dynamics (11), we use $u_{\min} = u_a - \delta_u$ and $u_{\max} = u_a + \delta_u$, where u_a is defined in (6), in the agreement (cooperative agents) regime, and $u_{\min} = u_d - \delta_u$ and $u_{\max} = u_d + \delta_u$, where u_d is defined in (7), in the disagreement (competitive agents) regime. This choice makes the opinion cascade behavior independent of the network-specific opinion-formation bifurcation point. The network topology, seed placement, seed input amplitude, and the attention threshold y_m are thus the only determinants of the observed opinion cascade behavior.

To verify the existence of a network threshold for the trigger of opinion cascades, in Fig. 7 we systematically vary network topology, attention threshold, and seed input amplitude, for randomly placed seeds. Figures 7c1 and 7c2 correspond to the case $N_o = 2$ and $N_o = 3$, respectively. All plots show the steady-state average network opinion strength $\frac{1}{N_a} \sum_{i=1}^{N_a} ||\mathbf{z}_i^*||$ as a function of the seed input amplitude in the coupled (4)–(11) system. Right plots are the average results for ten instances of a Watts–Strogatz network with randomly placed seeds and left plots are the average results for ten instances of a Barabási–Albert network with randomly placed seeds ($N_{\text{seeds}} = 5$ for both types of network). For each network and each value of N_o , we consider four different regimes: cooperative agents, $y_m = 0.05$ (full lines); competitive agents, $y_m = 0.05$ (dashed lines); cooperative agents, $y_m = 0.1$ (dotted lines); competitive agents, $y_m = 0.1$ (dashed-dotted lines). In all cases, one can detect a threshold value above which distinctively nonzero average opinions emerge, indicating a cascade as in Fig. 7b1 and b2.

The cascade threshold increases monotonically with u_{th} . The parameter u_{th} tunes cascade sensitivity at the single agent level, in the sense that smaller (larger) values of u_{th} decrease (increase) the minimum input strength necessary to trigger a cascade at the single agent level (Bizyaeva et al. 2020a, Theorem 4.6). In the presence of networked interaction, the tunable cascade sensitivity of the single agent level scales up to the network level: for large u_{th} , large seed input strengths are necessary to trigger a cascade; for small u_{th} , small seed input strengths are sufficient to trigger a cascade.

Observe that in our model, cascades can be partial or complete. For instance, the agreement cascade is complete in Fig. 7b2 left because all agents develop a markedly nonzero opinion. Conversely, the disagreement cascade is only partial in Fig. 7b2 right because some of the agents, including one of the seeds, possess a roughly neutral opinion at the post-cascade steady state. We further observe that, generically, when the bifurcation parameter is the seeds' input strength, as considered here, cascades happen through a simple saddle-node bifurcation. It follows that for supra-threshold inputs the time to reach the post-cascade steady state is inversely proportional to the input strength, a phenomenon known in the dynamical system literature as “traversing the ghost of the saddle-node bifurcation” (Izhikevich 2007, Sect. 3.3.5). A detailed analysis of both the transition from

partial to complete cascades and the time to the post-cascade steady state is left for future developments.

5 Control of agreement and disagreement opinion cascade by input centrality assignment

This section explores how the sensitivity of an opinion-forming network to trigger an agreement or disagreement cascade in response to a small number of inputs can be controlled by exploiting the emergent properties of the network. We introduce notions of agreement and disagreement centrality. These centrality indices can be used to guide the cascade seed placement toward maximizing or minimizing the network response to inputs, independent of the intrinsic node properties like, e.g., the attention threshold parameter u_{th} . In other words, our cascade model can be controlled independently and simultaneously in two ways: intrinsically at the agent level, by tuning the attention dynamic parameters, and extrinsically at the network level, by exploiting the network topology for efficient seed placement. We consider extrinsic control that enhances or diminishes sensitivity to inputs for agreement (disagreement) cascades by selecting as seeds the most or least central agents as defined by the agreement (disagreement) centrality index. The selected set of agents is optimal in the case centrality is modular; otherwise, it provides only an approximation to the optimal set of agents (see, for example, (Clark et al. 2014; Fitch and Leonard 2016)). We save the investigation of optimality for future work.

5.1 Agreement and disagreement centrality

A key prediction of Theorem 1 is that the absolute value of the entries of the adjacency matrix's eigenvectors v_{max} and v_{min} predict (up to linear order) the opinion strength (and sign) of each agent along agreement and disagreement bifurcation branches, respectively. This result suggests the following notion of centrality for agreement and disagreement opinion formation.

Let \tilde{A} be the unweighted adjacency matrix of an undirected, connected network. Let v_{max} be the normalized positive eigenvector associated with the unique maximal eigenvalue of \tilde{A} . Assume that the minimal eigenvalue is also unique and let v_{min} be the associated normalized eigenvector.

Definition 1 We define *agreement centrality of agent i* as $[v_{max}]_i$, *disagreement centrality of agent i* as $|[v_{min}]_i|$, and *signed disagreement centrality of agent i* as $[v_{min}]_i$

The proposed notion of agreement centrality is the same notion of eigenvector centrality introduced in Bonacich (1972). The proposed notion of disagreement centrality is new. Both notions of centrality will be key to analyzing and controlling opinion cascades over networks. It makes sense to talk about a signed disagreement centrality because the eigenvector v_{min} has mixed sign entries (whereas v_{max} can always be chosen to be positive). We will show that signed disagreement centrality is instrumental to controlling disagreement opinion cascades. For directed networks, one could use the eigenvectors of $\tilde{A} + \tilde{A}^T$ to define suitable centrality indices. We leave these ideas for future investigations.

5.2 Control of agreement opinion cascades by centrality assignment

Suppose that we can choose N_{seeds} agents to receive input, and suppose that each input has the same norm \bar{b} , but brings potentially different information (i.e., favoring or disfavoring) about, say, Option 1².

- *Toward maximization of agreement cascade sensitivity.* Let the control objective be to maximize the sensitivity of agreement opinion cascades to the N_{seeds} inputs. Then, an effective approach is to select the cascade seeds to be the N_{seeds} most central agents, according to the agreement centrality.
- *Toward minimization of agreement cascade sensitivity.* Let the control objective be to minimize the sensitivity of agreement opinion cascades to the N_{seeds} inputs. Then, an effective approach is to select the cascade seeds to be the N_{seeds} least central agents, according to the agreement centrality.

The rationale behind the proposed control strategy is the following. Maximally central agents are those that tend to develop stronger opinions at an agreement bifurcation. Triggering the network attention/opinion-strength positive feedback at maximally central agents therefore enhances the network response toward crossing the agreement cascade threshold. Likewise, minimally central agents tend to develop small opinions at an agreement bifurcation. Triggering the network attention/opinion-strength positive feedback at minimally central agents therefore diminishes the network response and keeps it away from crossing the agreement cascade threshold.

We illustrate the effectiveness of our agreement cascade control strategy in an eleven-agent network with $N_{\text{seeds}} = 2$ and $N_o = 3$. Figure 8a1 shows the network topology and the resulting node agreement centralities (visualized as the node sizes). The two black arrows point to the two most central agents. The two gray arrows point to the two less central agents. All agents have an intrinsic attention threshold $u_{\text{th}} = 0.1$. When the two inputs affect the two most central agents following our sensitivity-maximization algorithm, inputs with amplitudes as small as 0.1 are sufficient to trigger a full network cascade with large average opinion strength (Fig. 8b1—black curve) and in which all agents become opinionated (Fig. 8c1, top). Conversely, when the two inputs affect the two less central agents following our sensitivity-minimization algorithm, the network is still resisting to trigger a full cascade even for input amplitude as large as 1.0 (Fig. 8b1—gray curve). Only the two agents receiving an input are opinionated, while the rest of the agents remain close to neutral (Fig. 8c1, bottom).

5.3 Control of disagreement opinion cascades by centrality assignment

The control strategy described in the previous section to control the sensitivity of agreement opinion cascade can be generalized to the disagreement case. In this case, we exploit the notion of signed disagreement centrality introduced in Definition 1. Suppose again that we can choose N_{seeds} agents to receive input, and suppose that each input has the same

² The proposed agreement cascade control algorithm naturally generalizes to the case in which inputs bring mixed information about the various options.

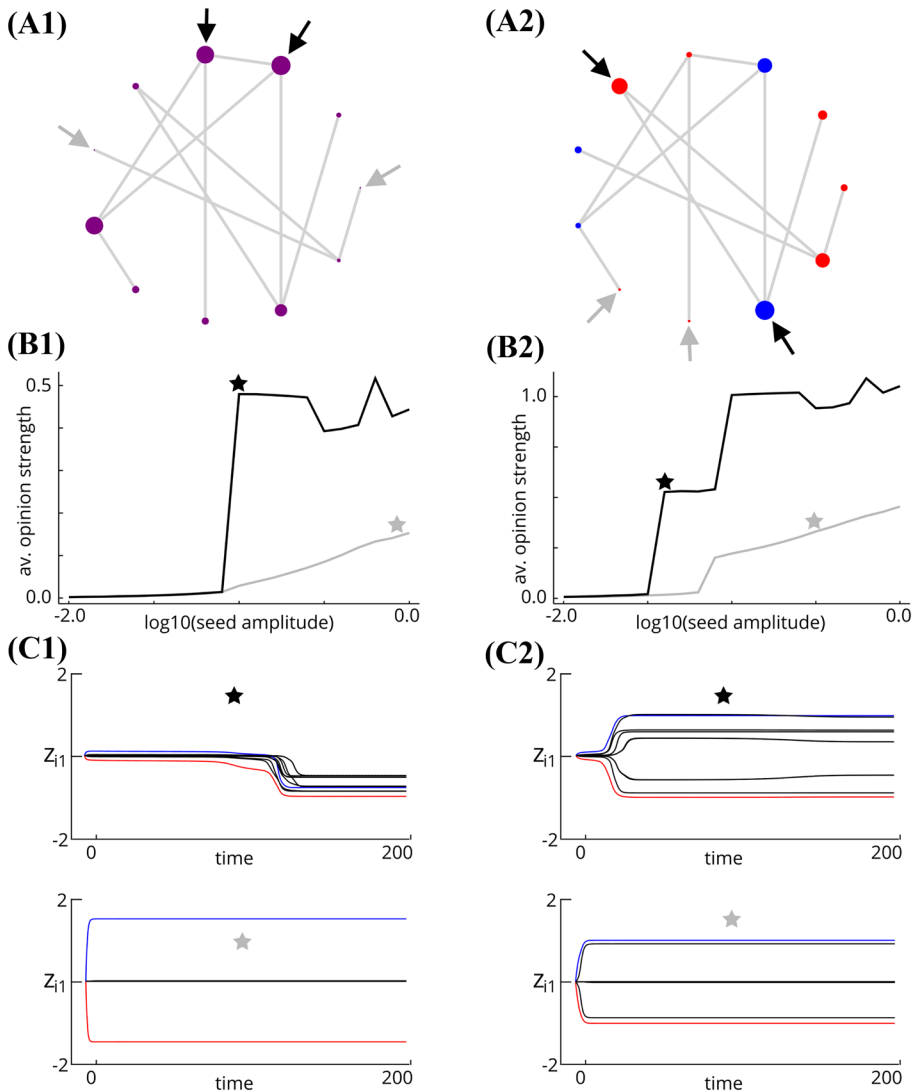


Fig. 8 Control of agreement and disagreement opinion cascade by centrality assignment. **a1** network topology and agreement centrality. The node size is proportional to its agreement centrality. **a2** network topology and signed disagreement centrality. The node size is proportional to its (absolute) disagreement centrality. Blue nodes have positive signed disagreement centrality. Red nodes have negative signed disagreement centrality. **b1** Average opinion strength as a function of the seed input amplitude under agreement cascade sensitivity maximization (black curve) and agreement cascade sensitivity minimization (gray curve). The network is in an agreement/cooperative regime. **b2** Average opinion strength as a function of the seed input amplitude under disagreement cascade sensitivity maximization (black curve) and disagreement cascade sensitivity minimization (gray curve). The network is in a disagreement/competitive regime. **c1** Evolution of agents' opinion about Option 1 in the agreement regime for the seed input amplitude indicated by the stars in **b1**. Top: under sensitivity maximization. Bottom: under sensitivity minimization. **c2** Evolution of agents' opinion about Option 1 in the disagreement regime for the seed input amplitude indicated by the stars in **b2**. Top: under sensitivity maximization. Bottom: under sensitivity minimization. Parameters: $\alpha = 0.2$, $\beta = -0.5$, $\gamma = -\delta = 0.1$ (agreement case), $\gamma = -\delta = -0.1$ (disagreement case), $u_{th} = 0.1$, $n = 5$, $\delta_u = 0.2$.

norm \bar{b} , but brings potentially different information (i.e., favoring or disfavoring) about Option 1³. The proposed control strategy is the following.

- *Toward maximization of disagreement cascade sensitivity.* Let the control objective be to maximize the sensitivity of disagreement opinion cascades to the N_{seeds} inputs. Then, an effective approach is to select the cascade seeds as follows. Assign inputs bringing information in favor of Option 1 to the agents with the largest signed disagreement centrality. Assign inputs bringing information in disfavor of Option 1 to the agents with the smallest signed disagreement centrality.
- *Toward minimization of disagreement cascade sensitivity.* Let the control objective be to minimize the sensitivity of disagreement opinion cascades to the N_{seeds} inputs. Then, an effective approach is to select cascade seeds to be the N_{seeds} least central agents, according to the (absolute) disagreement centrality index.

The rationale behind the proposed control strategy is the following. Two agents with large positive and large negative signed disagreement centrality, respectively, naturally tend to develop strongly disagreeing opinions at a disagreement bifurcation. Injecting two discordant inputs at two such agents, on the one hand, will contribute to the triggering of the network attention/opinion-strength positive feedback and, on the other hand, will amplify the disagreeing nature of the two agents. Taken together, these two effects greatly enhance the network response toward crossing the disagreement cascade threshold. Likewise, minimally central agents tend to develop small opinions at a disagreement bifurcation. Triggering the network attention/opinion-strength positive feedback at minimally central agents therefore greatly diminishes the network response and keeps it away from crossing the disagreement cascade threshold.

We illustrate the effectiveness of our disagreement cascade control strategy on the same eleven-agent network used for the agreement case. Figure 8a2 shows the resulting signed disagreement centralities, where red (blue) corresponds to negative (positive) signed disagreement centrality and the node size is proportional to the (absolute) disagreement centrality. As for the agreement case, our control strategy maximizes network sensitivity when inputs are placed on the most central agents (indicated by black arrows) according to our disagreement sensitivity-maximization algorithm, resulting in large disagreement cascades, with large average opinion strength (Fig. 8b2—black curve) and all agents exhibiting strong opinions (Fig. 8c2, top), for inputs with amplitude as small as $10^{-1.4} \simeq 0.04$. Conversely, when inputs are placed on the less central agents (indicated by gray arrows) according to our disagreement sensitivity-minimization algorithm, the network is still resisting to trigger a full disagreement cascade even for input amplitude as large as $10^{-0.5} \simeq 0.3$ (Fig. 8b2—gray curve). All but four agents remain close neutral (Fig. 8c2, bottom).

³ The proposed disagreement cascade control algorithm naturally generalizes to the case in which inputs bring mixed information about the various options as long as $\mathbf{b}_i = \pm \mathbf{b}_0$, $i = 1, \dots, N_a$.

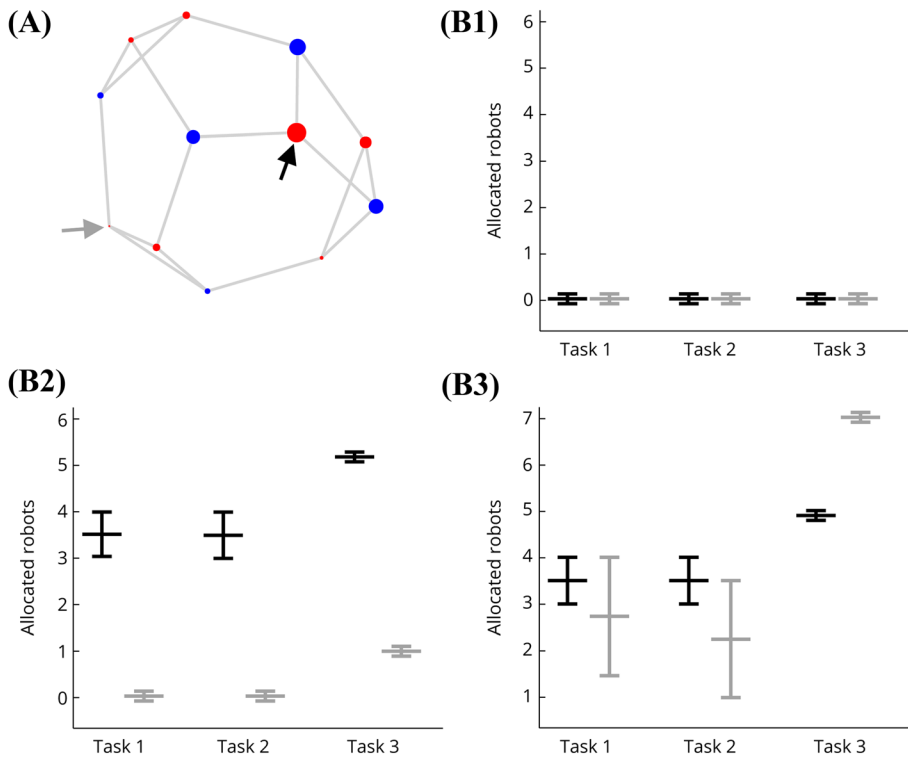


Fig. 9 Distributed allocation over three tasks by robot swarm with attention dynamics, one zealous robot, and communication defined by the Fruchterman-Reingold graph. The zealous robot perceives an increase in task urgency defined by ρ_3 . **a** The Fruchterman-Reingold graph and its node signed disagreement centralities graphically represented as the node size and color (the larger the node, the larger its disagreement centrality; blue means positive signed disagreement centrality and red means negative signed disagreement centrality). The black (gray) arrow indicates the most (least) central node. **b** Average (large bar) and standard deviation (short bar) of the number of agents allocated to Tasks 1, 2, 3 over one-hundred instances of our task-allocation cascade algorithm. Black: zealous robot is most central node. Gray: zealous robot is least central node. See main text for details. **b1** $\rho_3 = 10^{-0.9}$. **B2**: $\rho_3 = 10^{-0.6}$. **B3**: $\rho_3 = 10^{-0.3}$. Basal task urgency: $\mu_1 = \mu_2 = 0.3$ and $\mu_3 = 0.4$. Other parameters as Fig. 4.

6 Multi-robot task-allocation disagreement cascades

We now introduce attention dynamics (11) into the distributed multi-robot task-allocation dynamics (11) introduced in Sect. 3. The benefit of introducing attention dynamics in the task-allocation algorithm is to let the robot engage in task allocation only in the presence of a sufficiently large urgency to accomplish the tasks. We assume that only a limited number of robots (e.g., those equipped with suitable sensors, or those located in suitable positions) can sense real-time changes in task urgency above or below the basal urgency levels μ_i . We call urgency-sensing robots zealous. Zealous robots act as the seeds, or leaders, of task-allocation cascades.

An interesting property of the Fruchterman-Reingold graph is that, despite all agents having the same degree, node disagreement centralities are highly heterogeneous, as illustrated in Fig. 9a. In such a situation, our cascade-control strategy becomes crucial for efficient network placement of zealous robots.

To illustrate this fact, we consider a robot swarm as described by Fig. 9a and suppose that only one zealous robot is available, which detects an increase in the urgency of Task 3 of magnitude ρ_3 . More precisely, if i_z is the index of the zealous robot, then $v_{i_z,3} = 3\rho_3$, and, by (8), the effective urgency of Task 3 perceived by the zealous robot is $\mu_3 + \rho_3$. Furthermore, all robots are equipped with the attention dynamics (11), with $u_{\min} = u_d/2$, $u_{\max} = 2u_d$, $u_{th} = 0.1$, $n = 5$.

Figure 9b reproduces the model behavior in two cases: when the zealous robot is the least central node (according to disagreement centrality) and when the zealous robot is the most central node (according to signed disagreement centrality). As in Fig. 4b, we run one-hundred simulations with random initial conditions and small random zealousness parameter perturbations for three different values of ρ_3 , i.e., the urgency increase of Task 3 as perceived by the zealous robot. For small ρ_3 (Fig. 9b1), the zealous robot does not trigger a task-allocation cascade when it is in the least central or most central location in the network. All robots remain neutral (i.e., not allocated).

For intermediate values of ρ_3 (Fig. 9b2), only when the zealous robot is in the most central location does the task-allocation cascade get triggered. Observe that the standard deviation of the number of agents allocated to Task 3 drops to zero and that the standard deviations of the number of agents allocated to Tasks 1 and 2 also drop drastically as compared to the model behavior without attention (Fig. 4b). Furthermore, the ratio between the average number of robots allocated to Task 3 and the average number of robots allocated to Task 1 or Task 2 is closer to the ratio of the tasks' basal urgency $\mu_3/\mu_1 = \mu_3/\mu_2$ as compared to the model behavior without attention. We will analytically investigate these observations in future works.

Finally, for large values of ρ_3 (Fig. 9b3), a task-allocation cascade is triggered when the zealous robot is in the most central or least central location. However, only when the zealous robot is the most central node, do the robots allocate according to the task urgencies. When the zealous robot is the least central node, a disproportionate number of robots is allocated to Task 3. The allocation is not efficient. To summarize, disagreement centrality is key in placing urgency-sensing robots in a distributed task-allocation setting.

7 Application to large-scale complex networks

In this section, we provide numerical evidence that the analysis and control strategies introduced in this work scale up naturally to large networks with complex interactions. For comparison purposes, we focus on two types of complex networks: Watts–Strogatz with large rewiring probability and Barabasi–Albert. For the sake of illustration, we fix $N_a = 100$ and $N_o = 2$ or $N_o = 3$.

7.1 Centrality index predicts complex networks opinion formation behavior

The agreement and disagreement opinion formation behavior in Watts–Strogatz networks (Fig. 2) exhibits quantitative differences as compared to Barabasi–Albert networks (Fig. 3). Despite possessing the same average degree and for the same set of parameters, in the Barabasi–Albert network (Fig. 3) agents tend to develop weaker opinions as compared to the Watts–Strogatz network (Fig. 2). Also, in the disagreement regime, only in the Watts–Strogatz network do the group opinions organize into markedly bimodal (polarized)

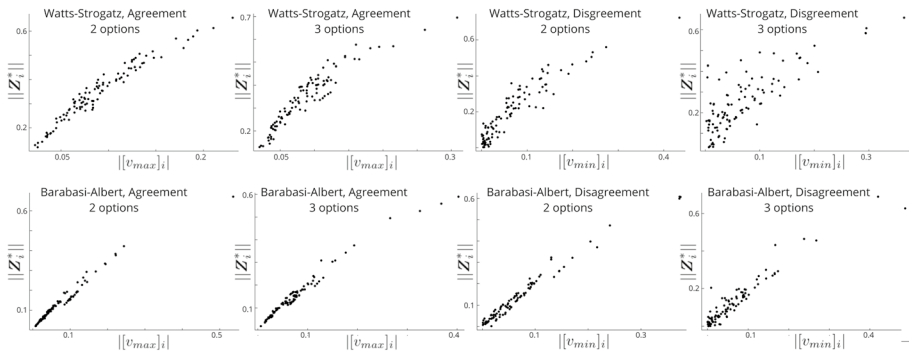


Fig. 10 First and second column: Scatter plots of agent opinion strengths $\|Z_i^*\|$ (vertical axes) at agreement steady-states against the absolute value of the entries of the adjacency matrix's eigenvector v_{\max} (horizontal axes). Third and fourth columns: Scatter plots of agent opinion strengths $\|Z_i^*\|$ (vertical axes) at disagreement steady-states against the absolute value of the entries of the adjacency matrix's eigenvector v_{\min} (horizontal axes) Plot titles indicate the network type, if the bifurcation was of agreement or disagreement type, and the number of options.




Fig. 11 Empirical distributions of the absolute values of the entries of v_{\max} (left) and v_{\min} (right) for Watts–Strogatz (black curves) and Barabasi–Albert (gray curves) networks.

distributions. These two observations can be explained and predicted in terms of the different centrality distributions in the two networks.

As predicted by Theorem 1, at an opinion-forming bifurcation, the agent states are distributed along a one-dimensional kernel subspace according to the entries of the adjacency matrix's eigenvectors v_{\max} and v_{\min} , corresponding to its maximum (agreement) and minimum (disagreement) eigenvalues, respectively. The opinion strength of agent i along an agreement (resp. disagreement) branch is thus expected to be roughly (i.e., modulo nonlinear terms in the center manifold expansion) proportional to the absolute value of the i -th entry of v_{\max} (resp. v_{\min}). Figure 10 illustrates this fact by plotting $\|v_{\max}\|_i$ (resp. $\|v_{\min}\|_i$) against $\|Z_i^*\|$, i.e., the opinion strength of agent i at the agreement (resp. disagreement) equilibrium reached for u slightly above the agreement (resp. disagreement) bifurcation point (6) (resp. (7)). In all cases, one can observe a monotone relationship between the two quantities, and the relationship is basically linear for the largest entries/opinion strengths. However, for such large networks ($N_a = 100$ in the numerical examples), the high-dimensionality of the state space makes nonlinear effects noticeable in some cases, particularly at small entries/opinion strengths.

The quantitative difference between agreement and disagreement opinion formation on Watts–Strogatz and Barabasi–Albert networks can thus be understood (and predicted) by looking at the distribution of the absolute values of the entries of v_{\max} and v_{\min} . Figure 11 shows the empirical distributions $f(|[v_{\max}]_i|)$ and $f(|[v_{\min}]_i|)$ of $|[v_{\max}]_i|$ and $|[v_{\min}]_i|$ obtained from one thousand realizations of Watts–Strogatz and Barabasi–Albert networks with the same size and parameters as those used in the simulation in Figs. 2 and 3. Both for v_{\max} and v_{\min} , the distributions are heavier at smaller values and lighter at larger values for the Barabasi–Albert case than for the Watts–Strogatz case. Both the median and expected value are larger in Watts–Strogatz networks than in Barabasi–Albert networks. In other words, both agreement and disagreement opinion formation in Barabasi–Albert networks tend to be characterized by a larger (resp. smaller) number of agents developing weak (resp. strong) opinions than in Watts–Strogatz networks. For the disagreement case, this translates into the fact that disagreeing agents tend to polarize more easily over a Watts–Strogatz topology than over a Barabasi–Albert topology.

7.2 Centrality index predicts complex networks opinion cascade behavior

Figure 7 highlights another quantitative difference between Watts–Strogatz and Barabasi–Albert networks. Only in Watts–Strogatz networks does the threshold behavior exhibit a marked all-or-none characteristic, in the sense that only for this kind of network does the average opinion switch from close to zero to a relatively constant nonzero value. This indicates that Watts–Strogatz networks are able to robustly trigger an opinion cascade as soon as the seed amplitude is above threshold, roughly independently of seed placement (which is random in the ten repetitions over which the results were averaged). Conversely, in Barabasi–Albert networks, random seed placement is incapable of robustly triggering a cascade. The averaged data does not exhibit a sharp threshold because different seed placements lead to sharply different cascade thresholds. As a consequence, the average opinion strength exhibits a strong linear dependence on the seed amplitude, reflecting the increasing fraction of cases in which even highly inefficient seed placements are able to trigger a cascade.

The quantitative differences between Watts–Strogatz and Barabasi–Albert networks can again be interpreted in simple algebraic terms from their agreement and disagreement centrality distributions (Fig. 11). Because Barabasi–Albert networks tend to have smaller centrality values than Watts–Strogatz networks, random agents used as seeds tend to develop weaker opinions in the former than in the latter. As a consequence, it is harder to spread the attention threshold-crossing behavior in Barabasi–Albert networks than in Watts–Strogatz networks because, on average, the attention–opinion positive feedback loop is less active in the former than in the latter.

Qualitatively, one could have predicted the same result by looking at more classical network properties, like the global clustering coefficient or the network diameter, which are both smaller in WS than BA networks. However, such statistical indices do not predict the patterns of opinion formation as our centrality indices do. In particular, they cannot distinguish between agreement and disagreement opinion formation and cascades. The crucial point of Theorem 1 is indeed that the bifurcation behavior at an agreement or a disagreement bifurcation is fully determined by the agreement and disagreement centralities, respectively. Local indices of clustering or shortest path length, like eccentricity, are also unsuitable to predict agreement and disagreement opinion formation and cascade behavior

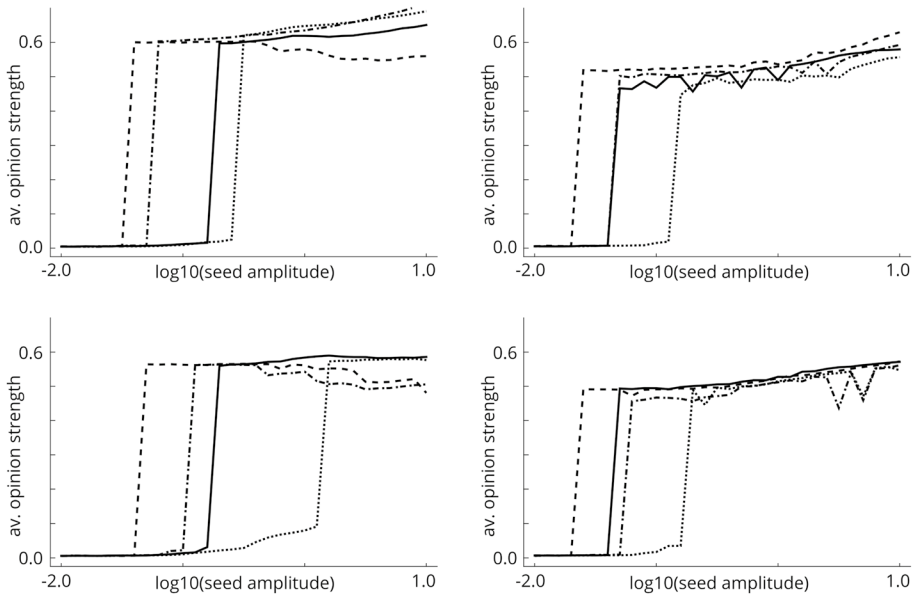


Fig. 12 Control of agreement (left plots) and disagreement (right plots) opinion cascades for $N_o = 2$ (top) and $N_o = 3$ (bottom) in a Watts–Strogatz network with $N_a = 100$ agents. Full and dashed lines correspond to $u_{th} = 0.05$ (i.e., relatively large global sensitivity) with minimizing and maximizing sensitivity control, respectively. Dotted and dashed-dotted lines correspond to $u_{th} = 0.1$ (i.e., relatively small global sensitivity) with minimizing and maximizing sensitivity control, respectively.

because (to the best of our knowledge) there are no general or simple connections between our notion of agreement and disagreement centrality and other local network properties.

7.3 Control of opinion cascade in complex networks

We finally illustrate the effectiveness of the control strategy introduced in Sects. 5.2 and 5.3 to maximize or minimize the sensitivity of agreement and disagreement opinion cascades to inputs on a large scale complex network. We use a Watts–Strogatz network in the case $N_o = 2$ and $N_o = 3$ and for two values of the intrinsic threshold u_{th} , which is assumed to be equal for all agents. Recall from Sect. 4.2 and Fig. 7 that the effect of increasing (resp., decreasing) u_{th} is to decrease (resp., increase) the global sensitivity of the network cascade to inputs.

Figures 12 top left (agreement, $N_o = 2$) and 12 bottom left (agreement $N_o = 3$) show that an agreement opinion network with a global small sensitivity ($u_{th} = 0.1$) can be made as sensitive as (or even more sensitive than) a network with relatively large global sensitivity ($u_{th} = 0.05$) by using our control strategy to maximize the agreement cascade sensitivity of the former (dashed-dotted lines) and minimize the agreement cascade sensitivity of the latter (full lines). For all cases, there is roughly an order of magnitude between the seed amplitudes required to trigger a cascade in the same network controlled toward maximizing versus minimizing its sensitivity.

Figures 12 top right (disagreement, $N_o = 2$) and 12 bottom right (disagreement $N_o = 3$) show that also in the disagreement case a network with a global small sensitivity ($u_{th} = 0.1$)

can be made as sensitive as (or even more sensitive than) a network with relatively large global sensitivity ($u_{th} = 0.05$) by using our sensitivity control strategy. As in the agreement case, for a fixed network topology, the seed amplitude required to trigger a cascade under sensitivity-maximizing control is much small (roughly half of an order of magnitude) than the seed amplitude required to trigger a cascade under sensitivity-minimizing control. As discussed at the end of Sect. 7.2, we are not aware of local network properties other than our agreement and disagreement centrality indices that could provide such an efficient agreement and disagreement opinion-cascade control.

8 Discussion and future directions

8.1 Summary of results

The possibility of forming an opinion in favor or disfavor of different options makes opinion cascades intrinsically different from binary network cascades (e.g., spread of fads or information) usually analyzed and modeled in the literature. Motivated by these differences, we introduced two types of centrality for opinion cascades: agreement centrality and disagreement centrality. The former predicts the pattern of opinion formation when agents cooperate in the opinion formation process (i.e., when they tend to align their opinion to that of other agents) and provide an efficient, although heuristic and possibly sub-optimal, way to control the sensitivity of agreement opinion cascades by seed placement. The latter predicts the pattern of opinion formation when agents compete in the opinion formation process (i.e., when they tend to reject other agents' opinions) and provide an efficient, although heuristic and possibly sub-optimal, way to control the sensitivity of disagreement opinion cascades by seed placement. We illustrated the practical validity of our theoretical predictions through an original example of distributed task allocation in multi-robot swarms and through numerical simulations in large random networks. Our centrality indices are grounded in the bifurcation behavior from a neutral to an opinionated state. To the best of our knowledge, this is the first time that centrality indices arise from intrinsically nonlinear bifurcation phenomena.

8.2 Agreement vs disagreement centrality

Our notion of agreement centrality is the same notion of eigenvector centrality introduced in (Bonacich 1972). Our notions of disagreement centrality are new. Crucially, the two indices of centrality can be sharply different for the same graph. The Frucht graph used for the task-allocation example in Sect. 3 reveals a key difference between agreement and disagreement centrality: whereas the former cannot distinguish nodes in a regular graph (eigenvector centrality is the same for all nodes in a regular graph), the latter can. We thus suggest that disagreement centrality might turn out to be a key to understanding complex networks where disagreement opinion formation is the norm, e.g., sociopolitical networks.

8.3 Simple and complex contagion

The seminal work (Centola and Macy 2007) revealed that the same network topology, in particular, in the presence of the bridges typical of small-world networks, can behave very differently during “simple” versus “complex” contagion. A simple contagion is one

in which a single active neighbor is sufficient to spread the contagion to a node. A complex contagion is one in which multiple active neighbors are necessary to spread the contagion to a node. Because time and state-space are continuous in our model, the difference between simple and complex contagions becomes nuanced, i.e., contagions can simultaneously be simple and complex. In the presence of neighbors with sufficiently strong opinions, the contagion is simple, because a single, strongly opinionated neighbor can be sufficient to let a given agent cross the cascade threshold. Conversely, if neighbors' opinions are weak the contagion is complex, because multiple opinionated neighbors are necessary to let a given agent cross the cascade threshold. The weaker the neighbors' opinions, the more complex the contagion. The results of this paper show that our notion of agreement and disagreement centrality are naturally suited to analysis and control of mixed, i.e., both simple and complex, contagions.

8.4 Extensions and future directions

8.4.1 Joint centrality and connection with other centrality indices

The proposed centralities exactly (up to nonlinear order in the center manifold expansion) characterize the pattern of agreement and disagreement opinion formation for *open-loop* changes of the attention parameter in the homogeneous system (4) without inputs. As such, they have been shown to provide a powerful heuristic for cascade seed placement. However, whether the resulting seed-placement algorithm is optimal remains an open question. A number of other centrality indices (e.g., information centrality (Stephenson and Zelen 1989), diffusion centrality (Banerjee et al. 2013)) are available for comparison. Most importantly, the results presented in Clark et al. (2014); Fitch and Leonard (2016) show that agent selection based on centrality becomes subtle if more than one agent must be selected. In the case of two agents, simply selecting the first and second most central agent generically does not provide the optimal solution, at least in highly regular networks. In particular, the notion of joint centrality introduced in Fitch and Leonard (2016) provides the tool to more rigorously tackle the multi-agent selection problem with respect to a given centrality index. We will explore how the notion of joint centrality applies to the proposed centralities in future works.

8.4.2 Heterogeneous agent thresholds

As in classical work on network cascades (Granovetter 1978; Schelling 2006; Watts 2002; Lim et al. 2015; Garulli et al. 2015; Rossi et al. 2017), the parameters of the closed-loop attention dynamics, and thus the resulting single-agent cascade threshold, can be assumed to be heterogeneous and drawn from some random distribution. By the intrinsic robustness of hyperbolic dynamical systems organized by bifurcations (Golubitsky and Schaeffer 1985), for small heterogeneity, we expect the result derived in this paper to remain robustly true. How large random heterogeneity interacts with the underlying network structure is an open question.

8.4.3 Non all-to-all intra-agent coupling

Another important extension will be to consider other than all-to-all intra-agent coupling topologies in the homogeneous regime (4). All-to-all intra-agent coupling models the case

in which options are a priori equivalent, with no hierarchies or clustering among them. Introducing more structured intra-agent coupling allows to group and order options into different issues, with or without hierarchies among them. This modeling option is particularly relevant when the opinion developed over an issue (e.g., at which restaurant to have dinner) influences and is influenced by opinion formation over another issue (e.g., which type of dish to eat). We will extend our modeling and analysis work to other than all-to-all intra-agent coupling in future work.

8.4.4 Time-varying network topologies and communication loss

We remark that, although not yet addressed by our theory, our modeling framework naturally allows to model and exploit time-varying network topologies. For instance, the network topology can change dynamically as a function of the agents' opinion state (as in classical bounded-confidence opinion-formation models Li et al. (2013)), which can naturally lead to clustered networks with clusters of agents having similar opinions, e.g., clusters of robots accomplishing similar tasks (see (Bizyaeva et al. 2020a, Theorem 3.5) for a static topology clustering result). Time-varying network topologies can also arise from communication loss between the agents, e.g., due to faults or when robots move too far away from each other. Crucially, as detailed in Sect. 4.1, our attention dynamics realizes a local multi-stable memory of the opinion developed during a cascade. This distributed memory might naturally be suited to make the opinion formation behavior robust to temporary loss of communication between two or more agent subpopulations.

8.4.5 Reversal to a neutral state through excitable attention-opinion dynamics

In some occasions, it might be useful for the swarm to revert from an opinionated state back to a neutral state, e.g., when a set of tasks or a foraging mission has been accomplished. This can be achieved by making the attention-opinion dynamics excitable, i.e., the transition to an excited (opinionated) state from the neutral state in response to inputs is only transient. Mimicking neural excitable dynamics (FitzHugh 1961; Hodgkin and Huxley 1952), the multi-stable attention-opinion dynamics present in each agent can be transformed into excitable dynamics by adding a slow adaptation variable, providing slow negative feedback on the attention-opinion dynamics. Fast positive feedback acting in conjunction with slow negative feedback constitutes the basic motif of excitability (Franci et al. 2018). Depending on the timescale of the slow adaptation variable, a transition to opinion through a cascade is returned to neutral after a tunable delay, allowing the swarm to become again reactive to new circumstances. The implementation of such an excitable attention-opinion dynamics will be the subject of future work.

8.5 Applications beyond multi-robot task-allocation

8.5.1 A flexible best-of- n and task-allocation model

Our model can rapidly and reliably switch between agreement and disagreement behavior. As such, it constitutes a dynamical opinion formation model that, when applied to robot swarms, can rapidly and reliably switch between best-of- n and task allocation. This flexibility can be key for adaptive behavior in robot swarms. For instance, during a foraging task (Zedadra et al. 2017; Pitonakova et al. 2018) the possibility to switch

between exploration, when the robots use the opinion dynamics in the disagreement regime to allocate them to explore different directions, and exploitation, when the robots use the opinion dynamics in the agreement regime to reach a consensus about the richest source location, is key to adaptive capabilities for the swarm. The parameters γ and δ , which determine the agreement ($\gamma > \delta$) or disagreement ($\gamma < \delta$) regime, can be equipped with input and state-dependent dynamics as in (Bizyaeva et al. 2020a, Sect. 6D), which allows for flexible transitions between exploration and exploitation as a function of the information collected about the environment and the state of each robot.

Opinion cascades, in both a best-of- n and a task-allocation setting, are useful when only a limited number of information-gathering agents are available. The information-gathering agents work as the seeds of an opinion cascade in the robot swarm. Through our threshold mechanism, a robot swarm can decide to remain in a neutral state (i.e., “do nothing” or “stay charging at the dock”) until a sufficiently large amount of information has been gathered by the seeds. Past this threshold, opinion formation spreads through the network. The resulting opinion (i.e., task allocation) state depends on the network topology, on the seed opinion states, and on possibly small (i.e., sub-threshold) biases of non-information gathering agents. Importantly, the seed-placement algorithm developed in Sect. 5 can be used to tune the swarm sensitivity to engage in the multi-task activity. When the sensitivity is small, the swarm behaves “lazily,” with relatively large inputs needed to trigger the transition to an active (i.e., opinionated) state. When the sensitivity is large, the swarm behaves “hyperactively,” with only relatively small inputs needed to trigger the beginning of the multi-task activity.

8.5.2 Biological and artificial swarms in motion

In both animal and robot swarms in motion, the decision to be made is about when and where to move. As discussed in Franci et al. (2020), our modeling framework naturally accommodates this situation, too, by letting the attention/bifurcation parameter be related to the geometry of the motion space, e.g., the group distance and relative position with respect to relevant objects like food sources, prey or predators, resources, or searched items. A number of works showed the relevance of similar approaches for biological systems (Couzin et al. 2005; Nabet et al. 2009; Couzin et al. 2011; Leonard et al. 2012; Pinkoviezky et al. 2018). Collision avoidance (Wang et al. 2017; Berg et al. 2011) can also be modeled and tackled in our opinion formation framework. A neutral opinion means in this setting moving straight, favoring Option 1 means turn right and favoring Option 2 means turn left. If in agreement regime, two colliding robots will either both turn right or turn left and collision will be avoided. In this simple setting, the attention parameter dynamics can be a function of the robots’ relative positions.

Similarly to the multi-task setting, the introduction of opinion cascades allows to model and design agent swarms in motion in which only a small number of agents (i.e., those that can sense objects and items in the environment) gather information about the geometry of the motion space. The same emergent behaviors described in the multi-task setting (e.g., tunable sensitivity of the swarm to informative inputs) are to be expected in the motion setting.

8.5.3 Sociopolitical networks and political polarization

The notion of “opinion” most naturally applies to networks of people discussing, expressing, and forming ideas (and opinions) about sociopolitical issues (e.g., governance, economics). Models of opinion dynamics have a long history (DeGroot 1974; Friedkin and Johnsen 1999; Cisneros-Velarde et al. 2019; Olfati-Saber and Murray 2004; Altafini 2013). Many modeling efforts focus on understanding the origin of polarized political views (McCarty 2019; Galam and Moscovici 1991; Macy et al. 2003; Robert 1997; Dandekar et al. 2013). Our modeling approach generalizes many existing opinion-formation models (Bizyaeva et al. 2020a) and provides the means toward a unified theoretical viewpoint on political polarization (Franci et al. 2020, Sec. 3.3).

Results from a recent experiment (Macy et al. 2019) support a ground-breaking hypothesis about the origin of polarized political views. Namely, in a network of initially neutral agents with heterogeneous political predispositions, political polarization arises from opinion cascades triggered by early movers (i.e., the cascade seeds) with the final opinion state unpredictable from the original political predispositions. Our model provides a rigorous explanation of these observations. The pattern of opinion formation reached at the end of a disagreement cascade is largely determined by the complex network topology and the seed placement, more than the intrinsic biases (i.e., political predispositions) of the agents. For fixed agent biases, different random distribution of seeds or changes in the random network topology can lead to completely different opinion formation patterns at the end of the cascade in our model. We are currently working on interdisciplinary sociopolitical research exploring these and other ideas.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

- Albert, R., & Barabasi, A.-L. (2002). Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1), 47.
- Altafini, C. (2013). Consensus problems on networks with antagonistic interactions. *IEEE Transactions on Automatic Control*, 58(4), 935–946. <https://doi.org/10.1109/TAC.2012.2224251>.
- Arthur, W. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). The diffusion of microfinance. *Science*, 341(6144).
- Berg, J. V. D., Guy, S., Lin, M., & Manocha, D. (2011). Reciprocal n-body collision avoidance. *Robotics research* (pp. 3–19). New York: Springer.
- Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992–1026.
- Bizyaeva, A., Franci, A., & Leonard, N. E. (2020a). A general model of opinion dynamics with tunable sensitivity. arXiv preprint, [arXiv:2009.04332](https://arxiv.org/abs/2009.04332).
- Bizyaeva, A., Matthews, A., Franci, A., & Leonard, N. E. (2020b). Patterns of nonlinear opinion formation on networks. arXiv preprint, [arXiv:2009.13600](https://arxiv.org/abs/2009.13600).
- Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences*, 11(3), 118–125.
- Bonacich, P. (1972). Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, 2(1), 113–120.
- Bonacich, P. (2007). Some unique properties of eigenvector centrality. *Social Networks*, 29(4), 555–564.

- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- Centola, D., & Macy, M. (2007). Complex contagions and the weakness of long ties. *American Journal of Sociology*, 113(3), 702–734.
- Cisneros-Velarde, P., Chan, K. S., & Bullo, F. (2019). Polarization and fluctuations in signed social networks. arXiv preprint, [arXiv:1902.00658](https://arxiv.org/abs/1902.00658).
- Clark, A., Bushnell, L., & 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. <https://doi.org/10.1109/TAC.2013.2281473>.
- Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025), 513–516.
- Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A. T., et al. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062), 1578–1580.
- Dandekar, P., Goel, A., & Lee, D. T. (2013). Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110(15), 5791–5796. <https://doi.org/10.1073/pnas.1217220110>.
- DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 121–132. <https://doi.org/10.1080/01621459.1974.10480137>.
- Dhooge, A., Govaerts, W., Kuznetsov, Y., Meijer, H., & Sautois, B. (2008). New features of the software matcont for bifurcation analysis of dynamical systems. *Mathematical and Computer Modelling of Dynamical Systems*, 14(2), 147–175.
- Elmhirst, T. (2004). Sn-equivariant symmetry-breaking bifurcations. *International Journal of Bifurcation and Chaos*, 14(03), 1017–1036.
- Fitch, K., & Leonard, N. E. (2016). Joint centrality distinguishes optimal leaders in noisy networks. *IEEE Transactions on Control of Network Systems*, 3(4), 366–378. <https://doi.org/10.1109/TCNS.2015.2481138>.
- FitzHugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophysical Journal*, 1(6), 445–466.
- Franci, A., Drion, G., & Sepulchre, R. (2018). Robust and tunable bursting requires slow positive feedback. *Journal of Neurophysiology*, 119(3), 1222–1234.
- Franci, A., Golubitsky, M., Bizyaeva, A., & Leonard, N. E. (2020). A model-independent theory of consensus and dissensus decision making. arXiv preprint, [arXiv:1909.05765](https://arxiv.org/abs/1909.05765).
- Friedkin, N. E., & Johnsen, E. C. (1999). Social influence networks and opinion change. In S. R. Thye, E. J. Lawler, M. W. Macy, & H. A. Walker (Eds.), *Advances in group processes, vol 16* (pp. 1–29). UK: Emerald Group Publishing Limited.
- Frucht, R. (1939). Herstellung von graphen mit vorgegebener abstrakter gruppe. *Compositio Mathematica*, 6, 239–250.
- Galam, S., & Moscovici, S. (1991). Towards a theory of collective phenomena: Consensus and attitude changes in groups. *European Journal of Social Psychology*, 21(1), 49–74. <https://doi.org/10.1002/ejsp.2420210105>.
- Garulli, A., Giannitrapani, A., & Valentini, M. (2015). Analysis of threshold models for collective actions in social networks. In *2015 European Control Conference (ECC)* (pp 211–216). IEEE.
- Gerkey, B., & Mataric, M. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research*, 23(9), 939–954.
- Gladwell, M. (2006). *The tipping point: How little things can make a big difference*. USA: Little, Brown.
- Golubitsky, M., & Schaeffer, D. G. (1985). Singularities and groups in bifurcation theory. *Applied mathematical sciences*. New York: Springer.
- Golubitsky, M., & Stewart, I. (2002). *The symmetry perspective: From equilibrium to chaos in phase space and physical space* (Vol. 200). New York: Springer.
- Golubitsky, M., Stewart, I., & Schaeffer, D. (1988). *Singularities and groups in bifurcation theory*. New York: Springer.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1420–1443.
- Gray, R., Franci, A., Srivastava, V., & Leonard, N. (2018). Multi-agent decision-making dynamics inspired by honeybees. *IEEE Transactions on Control of Networked Systems*, 5(2), 793–806.
- Hodgkin, A., & Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4), 500–544.
- Izhikevich, E. M. (2007). *Dynamical systems in neuroscience*. USA: MIT press.

- Kempe, D., Kleinberg, J., & Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD international conference on knowledge discovery and data mining* (pp 137–146).
- Leonard, N. E., Shen, T., Nabet, B., Scardovi, L., Couzin, I. D., & Levin, S. A. (2012). Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1), 227–232.
- Li, L., Scaglione, A., Swami, A., & Zhao, Q. (2013). Consensus, polarization and clustering of opinions in social networks. *IEEE Journal on Selected Areas in Communications*, 31(5), 1072–1083. <https://doi.org/10.1109/JSAC.2013.130609>.
- Lim, Y., Ozdaglar, A., and Teytelboym, A. (2015). A simple model of cascades in networks. Working Paper, Tech. Rep.
- Macy, M., Deri, S., Ruch, A., & Tong, N. (2019). Opinion cascades and the unpredictability of partisan polarization. *Science Advances*, 5(8), eaax0754.
- Macy, M., Kitts, J., Flache, A., & Benard, S. (2003). Polarization in dynamic networks: A hopfield model of emergent structure. In *Dynamic social network modeling and analysis* (pp. 162–173). National Academies Press, Washington DC.
- Mateo, D., Horsevad, N., Hassani, V., Chamanbaz, M., & Bouffanais, R. (2019). Optimal network topology for responsive collective behavior. *Science Advances*, 5, eaau0999.
- McCarty, N. (2019). *Polarization: What everyone needs to know*. Oxford: Oxford University Press.
- Nabet, B., Leonard, N. E., Couzin, I. D., & Levin, S. A. (2009). Dynamics of decision making in animal group motion. *Journal of Nonlinear Science*, 19, 399–435.
- Olfati-Saber, R., & Murray, R. M. (2004). Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9), 1520–1533.
- Pilditch, T. (2017). Opinion cascades and echo-chambers in online networks: A proof of concept agent-based model. *Cognitive science society*.
- Pinkoviezky, I., Couzin, I. D., & Gov, N. (2018). Collective conflict resolution in groups on the move. *Physical Review E*, 97(3), 032304.
- Pitonakova, L., Crowder, R., & Bullock, S. (2018). The information-cost-reward framework for understanding robot swarm foraging. *Swarm Intelligence*, 12(1), 71–96.
- Reina, A., Marshall, J. A. R., Trianni, V., & Bose, T. (2017). Model of the best-of- N nest-site selection process in honeybees. *Physical Review E*, 95(5), 052411.
- Reverdy, P., & Koditschek, D. (2018). A dynamical system for prioritizing and coordinating motivations. *SIAM Journal on Applied Dynamical Systems*, 17(2), 1683–1715.
- Reverdy, P., Vasilopoulos, V., & Koditschek, D. E. (2021). Motivation dynamics for autonomous composition of navigation tasks. *IEEE Transactions on Robotics*, 37(4). In press.
- Robert, A. (1997). The dissemination of culture: A model with local convergence and global polarization. *Journal of Conflict Resolution*, 41(2), 203–226. <https://doi.org/10.1177/0022002797041002001>.
- Rosa, D., & Giua, A. (2013). A non-progressive model of innovation diffusion in social networks. In *52nd IEEE conference on decision and control (CDC)* (pp. 6202–6207). IEEE.
- Rossi, W., Como, G., & Fagnani, F. (2017). Threshold models of cascades in large-scale networks. *IEEE Transactions on Network Science and Engineering*, 6(2), 158–172.
- Salehi, M., Sharma, R., Marzolla, M., Magnani, M., Siyari, P., & Montesi, D. (2015). Spreading processes in multilayer networks. *IEEE Transactions on Network Science and Engineering*, 2(2), 65–83.
- Schelling, T. (2006). *Micromotives and macrobehavior*. USA: WW Norton & Company.
- Schranz, M., Umlauf, M., Sende, M., & Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7, 36.
- Stephenson, K., & Zelen, M. (1989). Rethinking centrality: Methods and examples. *Social Networks*, 11(1), 1–37.
- Talamali, M. S., Marshall, J. A. R., Bose, T., & Reina, A. (2019). Improving collective decision accuracy via time-varying cross-inhibition. In *2019 International conference on robotics and automation (ICRA)* (pp. 9652–9659). IEEE.
- Traniello, J. (1989). Foraging strategies of ants. *Annual Review of Entomology*, 34(1), 191–210.
- Valente, T. (1995). Network models of the diffusion of innovations. *Models and Methods in Social Network Analysis*, 28, 98–116.
- Valentini, G., Ferrante, E., & Dorigo, M. (2017). The best-of- n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4, 9.
- Wang, K., & Michel, N. (1994). Robustness and perturbation analysis of a class of nonlinear systems with applications to neural networks. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 41(1), 24–32.

- Wang, L., Ames, A., & Egerstedt, M. (2017). Safety barrier certificates for collisions-free multirobot systems. *IEEE Transactions on Robotics*, 33(3), 661–674.
- Watts, D. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9), 5766–5771.
- Watts, D., & Strogatz, S. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684), 440–442.
- Yağan, O., & Gligor, V. (2012). Analysis of complex contagions in random multiplex networks. *Physical Review E*, 86(3), 036103.
- Zedadra, O., Jouandeau, N., Seridi, H., & Fortino, G. (2017). Multi-agent foraging: State-of-the-art and research challenges. *Complex Adaptive Systems Modeling*, 5(1), 1–24.
- Zhong, Y., & Leonard, N. E. (2019). A continuous threshold model of cascade dynamics. In *2019 IEEE 58th Conference on Decision and Control (CDC)* (pp. 1704–1709). IEEE.
- Zhong, Y., Srivastava, V., & Leonard, N. E. (2020). Influence spread in the heterogeneous multiplex linear threshold model. arXiv preprint, [arXiv:2008.04383](https://arxiv.org/abs/2008.04383).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.