

---

# COLLECTIVE DECISION MAKING BY EMBODIED NEURAL AGENTS

---

✉ **Nicolas Coucke\***

Consciousness, cognition and computation group  
IRIDIA  
Université Libre de Bruxelles  
Brussels, Belgium  
PPSP team, CRCHUSJ  
Université de Montréal  
Montréal, Canada  
Moral and social brain lab  
Universiteit Gent  
Ghent, Belgium  
nicolas.coucke@ugent.be

✉ **Mary Katherine Heinrich**

IRIDIA  
Université Libre de Bruxelles  
Brussels, Belgium  
mary.katherine.heinrich@ulb.be

✉ **Axel Cleeremans**

Consciousness, cognition and computation group  
Université Libre de Bruxelles  
Brussels, Belgium  
axel.cleeremans@ulb.be

✉ **Marco Dorigo**

IRIDIA  
Université Libre de Bruxelles  
Brussels, Belgium  
mdorigo@ulb.ac.be

✉ **Guillaume Dumas**

PPSP team, CRCHUSJ  
Mila - Quebec AI Institute  
Université de Montréal  
Montréal, Canada  
guillaume.dumas@umontreal.ca

## ABSTRACT

Collective decision making using simple social interactions has been studied in many types of multi-agent systems, including robot swarms and human social networks. However, existing multi-agent studies have rarely modeled the neural dynamics that underlie sensorimotor coordination in embodied biological agents. In this study, we investigated collective decisions that resulted from sensorimotor coordination among agents with simple neural dynamics. We equipped our agents with a model of minimal neural dynamics based on the *coordination dynamics* framework, and embedded them in an environment with a stimulus gradient. In our single-agent setup, the decision between two stimulus sources depends solely on the coordination of the agent's neural dynamics with its environment. In our multi-agent setup, that same decision also depends on the sensorimotor coordination between agents, via their simple social interactions. Our results show that the success of collective decisions depended on a balance of intra-agent, inter-agent, and agent-environment coupling, and we use these results to identify the influences of environmental factors on decision difficulty. More generally, our results demonstrate the impact of intra- and inter-brain coordination dynamics on collective behavior, can contribute to existing knowledge on the functional role of inter-agent synchrony, and are relevant to ongoing developments in neuro-AI and self-organized multi-agent systems.

### \*Corresponding author:

Nicolas Coucke, Department of Experimental Psychology, Henri Dunantlaan 2, 9000 Gent, Belgium,  
nicolas.coucke@ugent.be

**Keywords** NeuroAI · Embodied cognition · Neurodynamics · Cooperation · Multi-agent · Symmetry-breaking · Collective decision making · Swarm intelligence

**Classification**

Physical Sciences: Biophysics and Computational Biology

Biological Sciences: Psychology and Cognitive Sciences

## Significance statement

Collective behaviors require the spatial and temporal coordination of actions by many individuals. The neural mechanisms that enable such coordination among embodied biological agents are currently not well understood. By using simulations of simple embodied agents equipped with biologically plausible neural dynamics, we demonstrated how collective decision making can result from adaptive coupling between an agent’s neural dynamics, its environment, and other agents. Our findings make the case for the inclusion of intrinsic neural dynamics in the development of artificial intelligence and multi-agent systems, as a means to expand their ability for social interactions and collective tasks.

## 1 Introduction

Collective decision making is important to the normal functioning of human and animal groups [Bang and Frith, 2017, Conradt and List, 2008] and is also used in groups of artificial agents such as robots [Hamann et al., 2010, Montes de Oca et al., 2011, Valentini et al., 2017]. Decisions can refer to physical actions, such as the direction of movement of animal groups [Couzin et al., 2011], or symbolic questions that are disconnected from physical situations, such as those studied in collective estimation tasks in humans [Becker et al., 2017, Centola, 2022]. Collective decisions that are made by the group itself without external intervention typically require that a consensus emerge in the group. Consensus entails that all or at least a large majority of individuals agree, either on an approximate continuous value (e.g., a position in continuous space [Couzin et al., 2005, Yoo et al., 2021]) or a discrete option (e.g., voting for an arbitrary item from a list [Suzuki et al., 2015]).

Consensus is achieved through a distributed process that is not under the control of any single agent [Valentini et al., 2017]. Multi-agent models have been instrumental in investigating how the distributed interactions of individuals can result in a consensus. The individual agents used in most models behave according to rather simple rules or heuristics. For example, in the well-known ‘opinion dynamics’ models, agents typically update their opinion according to the majority or to the voter rule [Hegselmann and Krause, 2002, Galam, 2008, Flache et al., 2017]. Some recent models aim to replicate human cognitive processes more closely by using neuro-inspired approaches such as the drift-diffusion model [Srivastava and Leonard, 2014, Tump et al., 2020, Lokesh et al., 2022, Reina et al., 2023]. These opinion dynamics models have greatly advanced our understanding of how peer-to-peer interactions can give rise to collective phenomena such as polarization or consensus.

Another class of multi-agent models addresses collective dynamics in physical space rather than opinion space [Vicsek et al., 1995]. Such models pertain to, for example, animals choosing a new nest site or humans finding an exit during an emergency evacuation [Helbing and Molnar, 1995, Ma et al., 2016, Reina et al., 2017]. Also, in these more embodied models, behaviors are typically governed by simple rules or heuristics [Couzin et al., 2005, Moussaïd et al., 2011]. These models have illuminated, for example, how animals can resolve differences in initial movement directions and move to a single location using only local, implicit communication.

The simple behavioral rules that agents use in these models are approximations of more elaborate brain processes that underlie behavior in biological organisms. Thanks to recent advances in multi-brain neuroscience, it is now possible to simultaneously measure the brain activity of multiple interacting agents during collective behaviors [Kingsbury and Hong, 2020]. Despite these innovations, it remains unclear how the combined neural activity of multiple agents is involved in producing collective behavior. Computational models can be of help, by simulating how certain patterns of neural activity across agents produce collective behaviors [Moreau and Dumas, 2021]. However, current multi-brain models typically do not link the neural activity of agents to their behavior in an environment; so far, only a handful of multi-agent models use brain-inspired mechanisms rather than simple heuristics to generate collective movements [Reséndiz-Benhumera et al., 2021, Sridhar et al., 2021, Heins et al., 2024]. In this study, we propose a multi-agent model of collective decision making by embodied agents that are controlled by an oscillatory model of brain dynamics. Our goal with this approach is to pave the way for computational approaches that bridge neuroscience and the burgeoning field of collective behavior.

Considerations of the brain–body–environment interplay have gradually permeated cognitive science, culminating in the *4E cognition framework*, which sees cognitive processes as being embodied, enactive, embedded, and extended [Clark, 1998, Newen and Gallagher, 2018]. Similar considerations have become increasingly common in computational neuroscience and artificial intelligence research [Steels and Brooks, 2018, Colas et al., 2022]. Largely inspired by the enactive approach to embodied cognition, we constructed minimalist agents that simulate important attributes of biological agents: (1) intrinsic neural dynamics produced autonomously by the agent and (2) a constant sensorimotor loop with the environment [Varela et al., 1992, Thompson and Varela, 2001, Froese and Paolo, 2011]. We studied how such agents can reach a consensus by continuously adjusting to one another’s movements in a simple environment. We are not proposing that other approaches to modeling collective decision making are invalid. Rather, we wish to complement prior approaches, by providing a way of incorporating important and understudied aspects of embodied cognitive processes.

In order to successfully operate in an environment, the neural activity of a biological agent must be attuned to the characteristics of that environment. In neuroscience, brain activity is typically studied in terms of oscillations. Neural oscillations (also referred to as brain rhythms) have been linked to perception, movement, and even abstract cognition [Ward, 2003, Buzsáki, 2019]. Across taxa, evolution has selected a subset of rhythms that allow organisms to adequately interact with their environment [Buzsáki et al., 2013]. Moreover, brain rhythms can rapidly shift to accommodate changing environments and task demands [Senoussi et al., 2022, Charalambous and Djebbara, 2023].

Large-scale brain rhythms are produced by the coordinated oscillatory activity of many interacting brain regions. The *coordination dynamics* framework has been widely used to study how the activity of these dynamically interacting components is coordinated [Kelso, 1997, Kelso et al., 2014, Tognoli et al., 2020]. One important advantage of the coordination dynamics approach is that it can be used to study the metastable regime in which the brain usually operates [Tognoli and Kelso, 2014]. If coordination among brain regions were always stable, its dynamics could not be adequately modulated by the environment so as to allow the agent to engage in any adaptive behavior. On the other hand, overly unstable dynamics would result in the brain being too easily overwhelmed by environmental input, again preventing adaptive behavior. A metastable regime resolves this problem by allowing the brain to dynamically switch between several stable oscillatory states, thereby being neither completely stable nor completely unstable.

The Haken-Kelso-Bunz (HKB) equations provide a straightforward way to model metastable dynamics among interacting components, such as two populations of oscillating neurons in the brain [Haken et al., 1985, Kelso et al., 2014]. Two oscillating components, when modeled with the HKB equations, show in-phase attraction (similar to the Kuramoto model [Kuramoto, 1984]) and also anti-phase attraction. The simultaneous existence of in-phase (symmetrical) and anti-phase (asymmetrical) attraction produces a simple form of metastability [Kelso, 2013]. The HKB equations were first implemented as a neural controller of an embodied agent by Aguilera et al. [2013], to model the sensorimotor interactions of a situated agent with its environment. The authors illustrated the importance of taking into account embodied interactions when studying brain dynamics, by showing that the simulated neural dynamics of the agent were qualitatively different when sensorimotor interactions with the environment were disrupted. In this study, we adopted the HKB agent of Aguilera et al. [2013] and modified it so that it could support embodied collective decision making. By making our agent sensitive to both the environment and other agents, we could study the oscillatory neural dynamics of agents when coordinating with both each other and the environment.

In studies of continuous collective decision making, an often-studied question is: under which conditions can agents with different preferred movement directions reach a consensus? For example, two influential modeling approaches have shown that reaching a consensus is facilitated by the presence of a subgroup of unopinionated individuals [Couzin et al., 2005, Leonard et al., 2011]. In this study, we investigated this ability to reach consensus is modulated by the agents’ neural dynamics. In line with the enactive approach to embodied cognition, we expected that an agent can operate in its environment when it balances (a) its intrinsic neural activity with (b) its sensorimotor coordination with the environment. To investigate this, we first assessed how a single agent’s movement towards one of two local optima in its environment depended on the relative strength of the internal coupling of its brain oscillators versus its coupling to the environment. Agents involved in collective decisions must additionally balance interaction with other agents. Therefore, when we embedded a group of agents with varying initial states in an environment, we assessed how the success of collective decisions depended on a balance of not only intra-agent and agent–environment coupling but also inter-agent coupling (i.e., social influence). Lastly, we connected our model to previous models of collective decision making, by assessing how consensus is influenced by differences in (a) the quality of the two stimulus sources in the environment and (b) differences in individuals’ initial states.

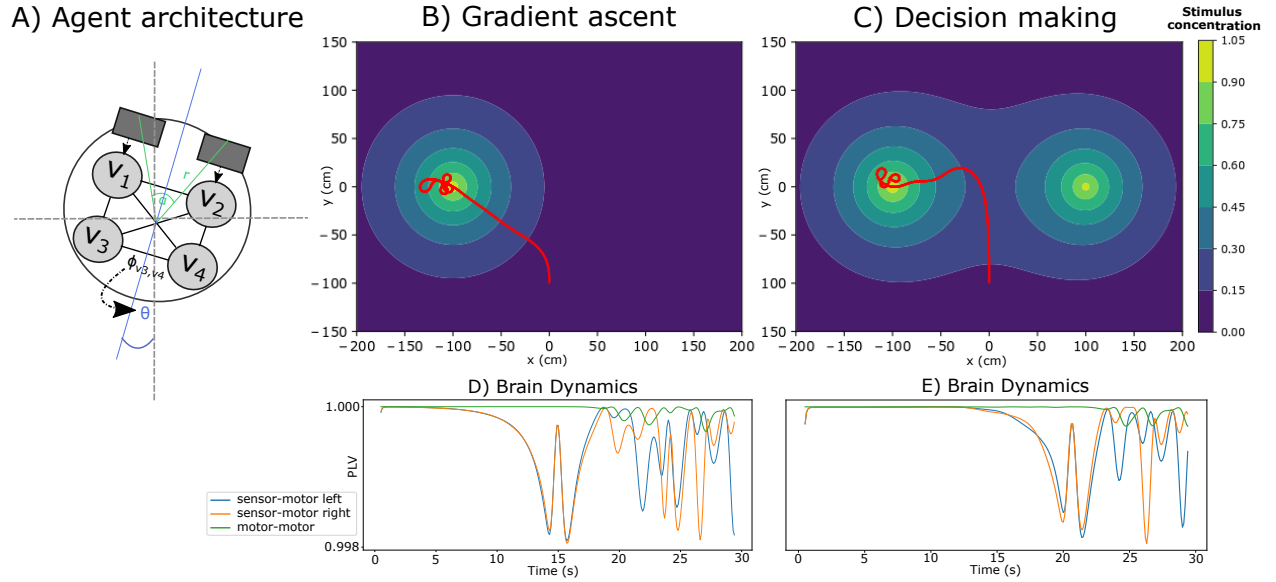


Figure 1: Single-agent behavior and neural dynamics. (A) The agent architecture: two sensors, each connected to a sensory oscillator, nodes 1 and 2 ( $v_1$  and  $v_2$ ), that are each in turn connected to a motor oscillator, nodes 3 and 4 ( $v_3$  and  $v_4$ ). The traveling orientation  $\theta$  of the agent is determined by the angle difference  $\phi_{v_3, v_4}$  between motor oscillators. (B) Gradient ascent: the agent’s trajectory (red) in the environment (brighter colors indicate higher stimulus concentration). (C) Decision making: two stimulus sources are present in the environment and the agent’s performance is measured by its ability to approach one of the two. (D) Internal phase locking of oscillators (contralateral sensor–motor and motor–motor) in B. (E) Internal phase locking of oscillators in C.

## 2 Model

### 2.1 Task environment

We created a simple environment in which agents could move and sense a stimulus. The environment contained one or more stimulus sources (i.e., sites), at which stimulus concentration is maximal (Fig. 1B-C). The stimulus concentration in the environment was inversely proportional to the distance from the stimulus source. Thus, the stimulus concentration followed a gradient from low to high concentration when approaching the stimulus source. A simple task that agents could perform in this environment was that of gradient ascent, i.e., following the gradient towards maximal stimulus concentration (Fig. 1B) Aguilera et al. [2013]. If one imagines the stimulus as ‘food,’ and the stimulus source as a food source, then this behavior reflects the food-seeking behavior of many simple organisms. When two sources of stimulus are present, the scenario could be considered a binary decision-making task (Fig. 1C). In this scenario, an agent could successfully reach a stimulus source if it could ‘decide’ between the two sources. In the multi-agent scenario, the task became a collective binary decision-making task. In this scenario, 10 agents started at the same position, but had different initial orientations (Fig. 2B). Due to these different initial orientations, agents could end up at different sites (Fig. 2C). However, agents had some social information about each other’s position (Fig. 2A; see below), and their task was to use this information to aggregate at the same site. We quantified performance of collective decision making according to how closely the agents collectively approached a single candidate site (see Methods).

### 2.2 Agent

We modeled an agent with minimal neural dynamics that could use sensorimotor coordination with the stimulus sources and the movements of other agents in order to move towards a candidate site. Our agent architecture was based on a minimal Braitenberg vehicle [Braitenberg, 1986], which is a self-driven agent with a very simple architecture: two sensors directly control two motors. To give our agent intrinsic neural dynamics, we connected two oscillator nodes to the sensors (loosely representing sensory brain regions; nodes 1 and 2 in Fig. 1A) and two oscillator nodes connected to the direction of the movement of the agent (loosely representing motor regions; nodes 3 and 4 in Fig. 1A). This design resembles the *situated HKB agent* of Aguilera et al. [2013], which had two oscillator nodes (one sensory and

one motor). Our agents have four nodes, so that they can use stereovision and differential drive to move directly to a stimulus source, rather than approaching it in a spiraling motion [Aguilera et al., 2013, cf.].

To model the interaction between the oscillators, we used an update rule for the phase of each oscillator, based on the following version of the HKB equation [Zhang et al., 2019]:

$$\dot{\phi}_{v_i} = \delta\omega_{v_i} + cI_{v_i} - \sum_{j=0}^N a_{v_i, v_j} \sin \phi_{v_i, v_j} - 2b_{v_i, v_j} \sin 2\phi_{v_i, v_j} , \quad (1)$$

where  $\dot{\phi}_i$  is the phase change of node  $v_i$ , and  $\omega_{v_i}$  is the intrinsic frequency of oscillator  $v_i$ . Parameters  $a_{v_i, v_j}$  and  $b_{v_i, v_j}$  represent the contribution of, respectively, in-phase attraction, and anti-phase attraction between oscillators  $v_i$  and  $v_j$ . Lastly,  $c$  parameterizes how strongly the oscillator phase is modulated by sensory input  $I_{v_i}$ . This parameter is set to zero for each motor oscillator, because it is not connected to a sensor. (See Methods for the version of the update equation used for each oscillator.)

Our agent moves at a constant speed and the activity of the motor oscillations is linked to the agent’s movement direction. The heading is updated according to the phase angle between the two motor oscillators, such that

$$\dot{\theta} = \eta\phi_{v_3, v_4} , \quad (2)$$

where  $\theta$  is the orientation of the agent in the environment and  $\eta$  is a scaling factor. Together, these equations create a closed sensorimotor loop between the agent’s internal oscillator dynamics and the external environment.

In the multi-agent scenario, we gave agents the added behavior of emitting the same stimulus that they observed to be present in the environment. The stimulus concentration emitted by social agent  $j$  was perceived by agent  $i$  as:

$$I_{ij} = S * e^{-\lambda_a D_{ij}} , \quad (3)$$

where  $D_{ij}$  is the Euclidean distance between agent  $i$  and agent  $j$ ,  $S$  is the strength of social influence between agents (identical for all agents), and  $\lambda_a$  is the decay rate of the emitted stimulus. Note that an agent did not perceive its own emitted stimulus.

### 3 Results

We performed both single-agent simulations and multi-agent simulations. For the single-agent simulations, we quantified neural coordination dynamics in terms of integration and metastability. The integration of brain regions by means of phase-locked activity is a central mechanism of brain function [Varela et al., 2001, Avena-Koenigsberger et al., 2018], and can be quantified using the phase-locking value (PLV). Brain function supportive of adaptive behavior relies on switching between different brain states. To quantify this aspect of neural dynamics, we used the standard deviation of the Kuramoto order parameter SD(KOP) [Strogatz, 2000, Cabral et al., 2022].

For the multi-agent simulations, we additionally quantified the coordination dynamics occurring across the different agents. We analyzed agents’ movement trajectories using the KOP as a measure of alignment, and SD(KOP) as a measure of alignment variability between agents’ movements [Strogatz, 2000, Cabral et al., 2022]. Lastly, we also quantified the degree of coordinated activity between the neural dynamics of the different agents by using the weighted phase-lag index (wPLI), a measure of phase locking that discards zero-phase coupling and can be interpreted as the co-variance between two signals [Vinck et al., 2011].

#### 3.1 Single-agent simulations

We first performed a series of single-agent simulations to assess how an individual agent’s neural dynamics are related to its ability to move towards a stimulus source in its environment. In the single-agent setup (see Fig. 1), an agent tried to climb a gradient towards a global maximum. To simulate different types of neural dynamics, we varied the internal coupling strength between the agent’s oscillator nodes ( $a_{v_i, v_j}$  in Eq. 1), and varied whether or not it was sensitive to external stimuli. For each agent configuration, we performed 50 runs with different random initial oscillator phases. In Fig. 3, we characterize the neural dynamics associated with each agent configuration. The top panel shows the average level of integration (i.e., phase locking) between the agent’s oscillators, measured by the mean PLV, and the bottom panel shows the degree of metastability among the agent’s oscillators (measured by SD(KOP); see Methods).

It is notable that, in the absence of sensory input, the system quickly found a stable state with minimal variation in oscillator dynamics: agents without stimulus input (squares in Fig. 3) consistently had values close to PLV= 1 and SD(KOP)= 0. Conversely, agents with sensory input (circles in Fig. 3) had a broad range of parameter values that

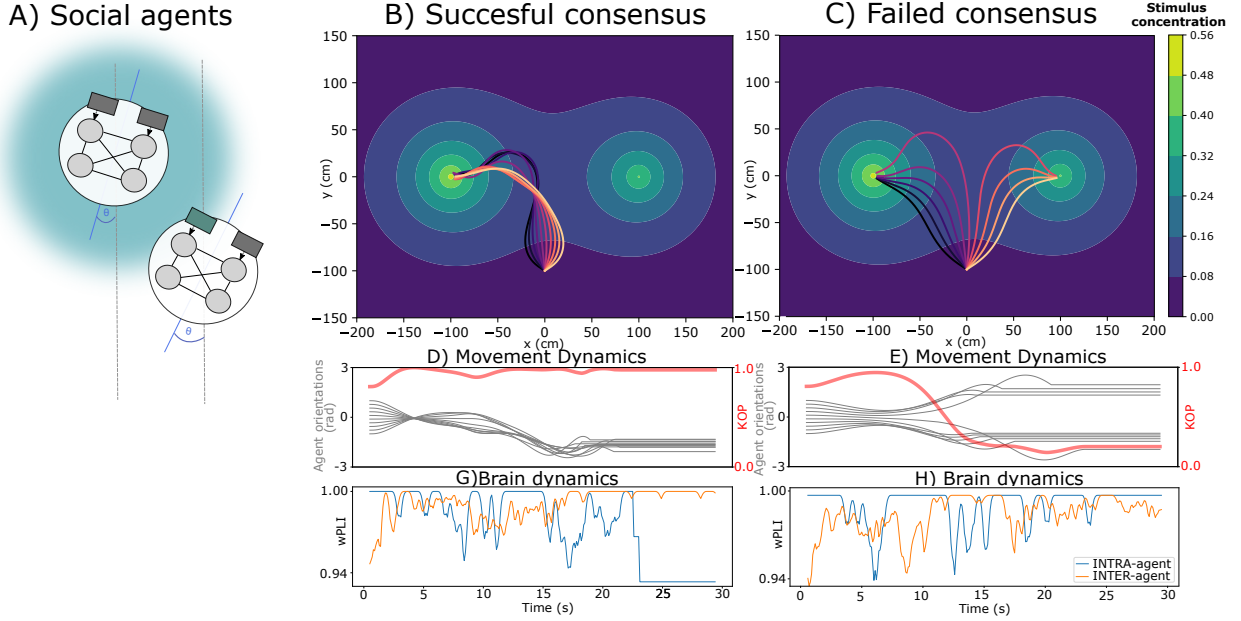


Figure 2: Agent behavior and intra-agent neural dynamics during collective decision making. (A) Agents emit stimulus that can be perceived by other agents. (B) Higher social stimulation allows agents to converge onto the same stimulus source in their environment. (C) With lower social stimulation, agents do not converge on the same source. (D-E) Movement angles of agents (gray lines) and KOP of the group, indicating the degree of alignment (red line). Alignment increases when all agents are moving towards the same stimulus source and decreases when they are not. (G-H) Intra- and inter-agent neural dynamics: average intra-agent wPLI (blue line) and average inter-agent wPLI (orange line).

resulted in lower PLV and higher SD(KOP). This shows that, as expected, sensory input can alter the coordination regime of the neural dynamics. In the presence of sensory input, PLV decreases as internal coupling increases from 0 to 1, indicating that the neural dynamics at low internal coupling are mostly driven by stimulus input, without being significantly modulated by the interactions among the agent's own oscillators. Simultaneously, SD(KOP) remained relatively high, indicating that, at low internal coupling, sensory input caused the system of oscillators to quickly cycle between oscillatory states.

As the agent's internal coupling increases further, the interactions between the agent's oscillators become strong enough to meaningfully modulate sensory input, which results in a lower level of apparent oscillator integration, with PLV decreasing to 0.75 while the degree of metastability plateaus. Beyond an internal coupling level of  $a_{v_i, v_j} = 1.4$ , the internal coupling of the agent's oscillators started to dominate, which resulted in highly integrated oscillators, indicated by higher PLV. At an internal coupling level of  $a_{v_i, v_j} = 1.7$ , the effect of internal coupling became strong enough that it nullified the effect of any sensory input, resulting in PLV = 1 (indicating no variation in inter-oscillator dynamics). This increase in integration was accompanied by a sharp drop in metastability, indicating that the system tends to get stuck in a single stable state. Such a stable state of high integration between oscillators precludes changes in movement direction in response to sensory input, inhibiting the agent from approaching the stimulus source.

The colors of the data points in Fig. 3 indicate agent performance: brighter colors indicate that the agent was better able to approach the stimulus. The distribution of colors shows that agents performed best in the range  $a_{v_i, v_j} \in \{0.8, \dots, 1.5\}$ . At these intermediate coupling values, there was a decrease in oscillator integration (as shown by PLV) and moderately high metastability (as shown by SD(KOP)). In short, these results show that at low internal coupling, neural dynamics are mostly driven by sensory input, and oscillatory coordination states change quickly. At moderate internal coupling, neural dynamics modulate sensory input without fully dominating it. In this intermediate range with relatively low integration and high metastability, the agent displays adaptive behavior. At high internal coupling, neural dynamics nullify the effect of sensory input and agent behavior cannot change in response to the environment. (See the supplementary text S2 for a more elaborate analysis of gradient climbing and decision making by single agents, as well as its relation to neural dynamics.)

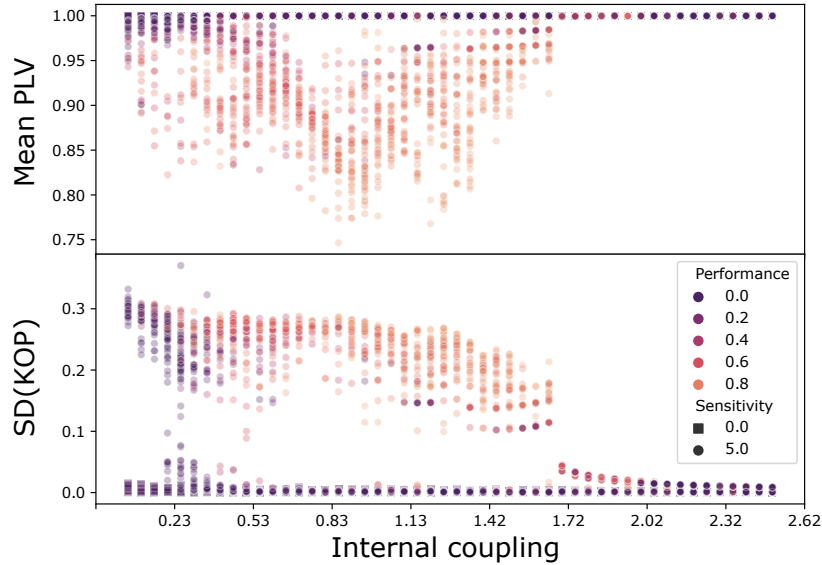


Figure 3: Intra-agent neural dynamics: the mean PLV and SD(KOP) of each agent during its run ( $y$  axis), according to the internal coupling degree ( $a_{v_i, v_j}$  and  $b_{v_i, v_j}$ ) in its neural controller ( $x$  axis). Squares represent agents without sensory input and dots represent agents with sensory input. Lighter colors represent higher performance. For each configuration (i.e., internal coupling degree and stimulus sensitivity), 50 runs were performed with random initial phases of the oscillators. Each data point represents the average of one run.

### 3.2 Multi-agent simulations

We ran two sets of multi-agent simulations. In the first set of simulations, we varied parameters related to the agent configuration (internal coupling, environmental sensitivity, social influence) and observed the effects on decision-making performance and collective dynamics. In the second series of experiments, we kept the agent configuration constant and modified the quality difference between the two stimulus sources in their environment, as well as the starting angles between the agents.

#### 3.2.1 Consensus as a function of internal, environmental, and social influences

We conducted a series of simulations with groups of 10 agents in an environment with a fixed quality (brightness) ratio of  $r = 0.8$  between the two stimulus sources (Fig. 2). For each simulation run, we quantified the performance as the degree to which agents could approach the same stimulus source (as in Fig. 2B) rather than going to different sources (as in Fig. 2C; see Methods). Throughout each simulation, we track the coordination dynamics among agents' movements (Fig. 2D-E), as well as measures of coordination between the neural dynamics of the different agents (Fig. 2G-H).

We conducted simulations for different parameter values of internal coupling strength ( $a_{v_i, v_j}$  in Eq. 1), sensitivity to the environment ( $c$  in Eq. 1), and the degree of social influence ( $S$  in Eq. 3). For each combination of parameter values, we display the final performance in a ternary plot (Fig. 4A). Each of the three corners of the ternary plot corresponds to one of the parameters being maximal and the others zero. We also display the corresponding measures of movement coordination and neural coordination dynamics for each parameter combination in adjacent plots (Fig. 4B-E).

We assessed movement coordination in terms of the movement alignment (KOP) and the variability in movement alignment (i.e., 'alignment variability'; SD(KOP)). To assess the coordination dynamics within and between agents' neural dynamics, we used a measure of phase covariance (wPLI), rather than the phase-locking value used for the single-agent case. In contrast to the PLV, the wPLI does not take into account zero-lag coupling between oscillators. Measures with this property are preferred in multi-brain neuroscience, since they discount spurious coordination due to, e.g., common input from the environment [Czeszumski et al., 2020, Schwartz et al., 2022]. In our simulations, spurious coordination between oscillators could similarly have been caused by identical initial phases of agents' oscillators (see Methods).

The middle region of the plots in Fig. 4 corresponds to a parameter range in which internal, environmental, and social influences are appropriately balanced for reaching consensus, as indicated by the bright yellow area in Fig. 4A. This region was accompanied by high movement alignment and low alignment variability, indicating that agents could use environmental and social information to coherently move towards the same stimulus source. Part of this region corresponds to a narrow area of increased inter-brain covariance, indicating that this aligned movement was accompanied by coordinated neural dynamics across agents. The lower right corner of the plots corresponds to a region of increased internal influences and decreased social influences. As long as social influences are non-zero, performance remains relatively high in this parameter range. Movement alignment is decreased and alignment variability is increased relative to the middle part of the plot. This indicates that, when internal coupling increases, agents' movements become less aligned, but can still result in consensus. Interestingly, the lower right corner of the plot corresponds to a decrease in both intra-brain covariance (between different oscillator nodes within the same agent) and inter-brain covariance (between the same oscillator nodes across different agents). More alignment variability among the agents' movements is thus accompanied by more independent by brain dynamics that are more independent.

A last and interesting observation can be made at the left edge of the plots, where internal coupling is low. In this region, agents are driven entirely by a combination of social and environmental influences. This region in parameter space was accompanied by high movement alignment and low alignment variability, but was associated with a decrease in performance. This suggests that agents moved in a highly aligned manner, but failed to collectively approach either of the two stimulus sources. This outcome highlights the importance of balancing external influences with sufficient internal coupling. When agents are overly coupled to external stimuli without enough counteracting internal coupling, they struggle to move towards an increasing stimulus concentration. Supplementary figure S2 illustrates how increased social influence, without a corresponding increase in internal coupling, leads to decreased performance.

Taken together, these results show that agents reach a consensus when their configuration facilitates a balanced integration of environmental, social, and internal influences, and this is reflected in neural and behavioral dynamics.

### 3.2.2 Consensus as a function of environment configuration

If the ability to reach a consensus depended on the features of the environment, we would expect that binary decision making should be easier (and thus performance higher) when the difference between the two stimulus sources is larger, and when the initial angle between the agents is smaller. We performed simulations with 10 agents with a fixed architecture, and varied the initial starting angle between agents and the brightness ratio between the two stimulus sources (see Methods).

The results in Fig. 5 show that, overall, performance depended on a combination of starting angle and stimulus ratio. Performance was maximal when agents had identical starting orientations and only one stimulus source was present (top left of Fig. 5). In accordance with our expectations, performance decreases as the second stimulus source became brighter and the starting angle between agents increased.

It should be noted that performance did not increase linearly as the decision-making task became easier. Rather, there was a repeating pattern of sharp decreases in performance followed by short plateaus. This was most likely due to a combination of our performance measure and the relatively small number of agents (see Methods). Each drop in performance was caused by one of the ten agents moving away from the global maximum and towards the competing local maximum (i.e., to the stimulus source with lower brightness). Supplementary figure S3 provides a more detailed account of this pattern. Overall, these simulations show that the ease with which the agents reach a consensus depends not only on their architecture but also on the environment in which they operate.

## 4 Discussion

We have modeled collective decision making with embodied neural agents that are controlled by simple oscillatory brain dynamics. We first showed that the ability of an agent to move towards a stimulus source was reflected in the intra-brain dynamics of that agent. In an intermediate parameter range where agents were neither overwhelmed by environmental stimulus nor insensitive to it, neural dynamics could become sufficiently uncoupled and metastable to allow the agent to approach its target. Furthermore, multiple agents with different initial heading directions could overcome their differences and converge on one of the two available options. Agents were able to do this within a parameter range that adequately weighted environmental, social, and internal influences. When one influence was too high with respect to the others, the agents failed to converge on a decision.

In this regard, our model differs from more disembodied multi-agent models [Flache et al., 2017]. In disembodied cognition models of collective decision making, increasing social influence to an arbitrarily high degree would lead to a fast and efficient convergence on any of the possible decision options. The option selected may be of greater or



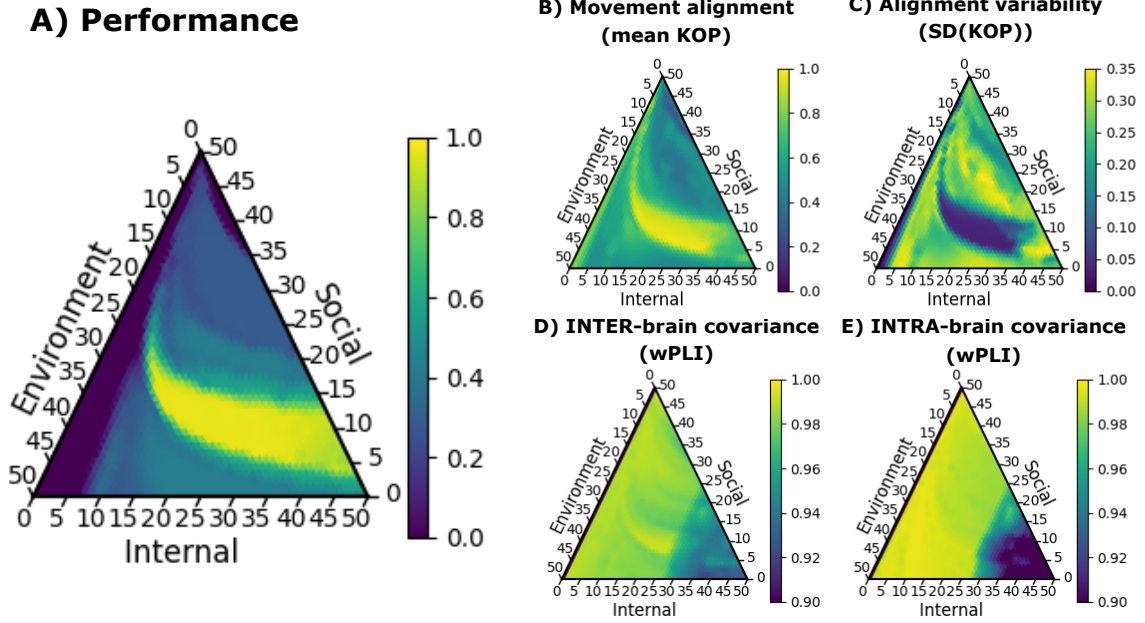


Figure 4: Ternary plots illustrating how collective behavior and neural dynamics depend on the agent configuration. Each point in the triangle corresponds to a certain weighting of environmental stimulus, social information, and internal motor coupling. In each simulation, the parameters fulfill the condition  $stimulus\ sensitivity + social\ sensitivity + internal\ coupling = 100$ . The scale  $[0, 50]$  for each dimension corresponds to respective parameter values of  $c \in \{0, \dots, 10\}$  for stimulus sensitivity,  $S \in \{0, \dots, 5\}$  for social sensitivity, and  $a_{v_3, v_4} \in \{0, \dots, 1\}$  for internal coupling. The top corner corresponds to maximal social sensitivity, the left corner to maximal environmental sensitivity, and the right corner to maximal internal coupling. The brightness (yellowness) in panel A indicates the performance of collective decision making. A performance of 0 indicates that agents failed to reach either of the two stimulus sources. A performance of 0.5 indicates that half of the agents reached the same stimulus source. A performance of 1 indicates that all agents reached the same stimulus source and thus that a consensus was reached. The brightness in panels B-E indicate the strength of, respectively, the movement alignment, alignment variability, inter-brain covariance, and intra-brain covariance.

lesser quality, but the ability of agents to reach any option at all has rarely been studied. Our results showed that, when increasing social sensitivity too much, the agent’s neural dynamics may become saturated with social information—to such a degree that it cannot adequately interact with the environment and move towards one of the options. This is reminiscent of real-world situations in which agents that are too consumed by interacting socially with one another lose adaptive interactions with the environment (see Strasbourg dancing plague [Waller, 2008] or circular milling in army ants [Couzin and Franks, 2003]).

Our model also differs from the more embodied multi-agent models of movement-based collective decisions by animals. In most such models, agents have a parameter that explains their preferred target location or movement direction [Couzin et al., 2011, Leonard et al., 2011, Sridhar et al., 2021]. In our simulations, agents had different initial movement directions but did not have a parameter representing a preferred movement direction. Rather, their movement direction emerged from their interactions with each other and with their environment. Our results showed that the more an agent’s initial movement direction was between the two sources rather than pointing to one of them, the more likely agents were to reach a consensus. These results are somewhat in line with previous findings that a larger proportion of unopinionated individuals promotes consensus in models of movement-based decisions made by groups of animals? [Couzin et al., 2011, Leonard et al., 2011].

Our simulations also showed that a larger difference between the quality of stimulus sources in the environment generally resulted in a higher proportion of agents reaching a consensus. This is in accordance with studies of discrete decision making between a few options in an environment, where models have shown how the speed and accuracy of collective decisions depend on differences between environmental stimuli [Valentini et al., 2015, 2017]. Models that do not take into account the quality of options in the environment, or that consider options of equal quality, often observe agents traveling in a compromise direction [Leonard et al., 2011]. Since agents’ neural dynamics were strongly

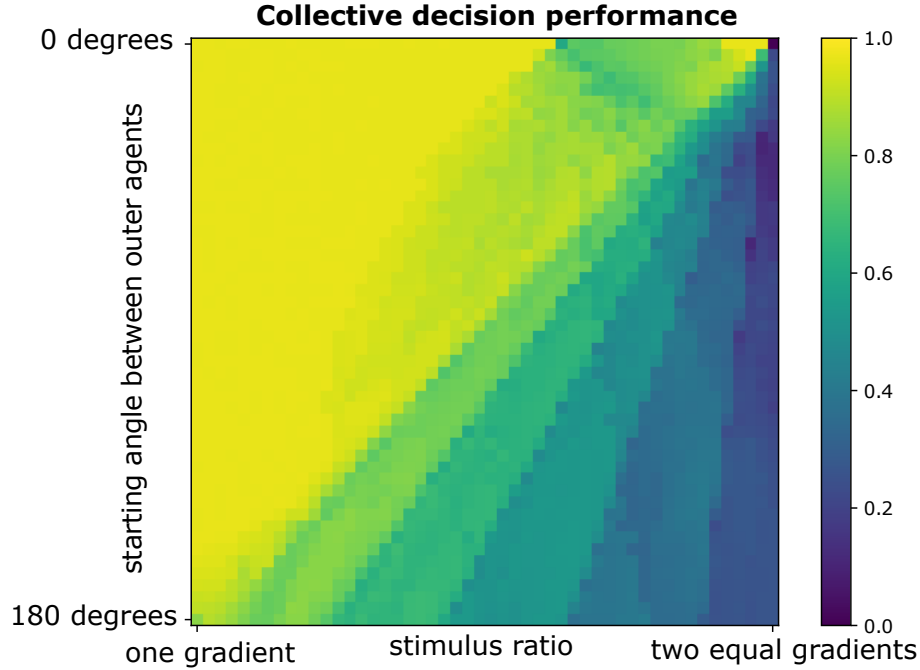


Figure 5: Dependence of the collective decision-making performance on the environment and initial orientations of agents. The leftmost extreme of the  $x$ -axis represents the cases with only one stimulus source present in the environment. Moving towards the right, the brightness of a second stimulus source increases until the two have equal brightness. The agents always start with equal angles between them. At the bottom of the  $y$ -axis, the agents are spread so that the outermost two of the ten agents are at a  $180^\circ$  angle. At the top of the  $y$ -axis, all agents start with angles of  $0^\circ$  between them. All agents have identical parameters; stimulus sensitivity is  $c = 3$ , social sensitivity is  $S = 1$ , and internal coupling is  $a_{v_i, v_j} = 0.5$ .

influenced by the environment in our simulations, such compromises only occur in a small region of the parameter space. The closer agents came to a stimulus source, the higher the stimulus concentration they observed, causing agents to almost always go to one of the two stimulus sources instead of taking a compromise direction. Our simulations also showed some surprising emergent collective behaviors, such as ‘overshooting’ in response to social influence and, as a result, moving towards an option that corresponds neither with the initial movement direction nor with the option chosen by other agents (see Supplementary figure S2A).

In the current paper, we performed deterministic simulations of agents that started from the same position and moved towards one of two stimulus sources. In future work, our model could be studied in environments with a higher number of stimulus sources, without changes to the agent architecture [Sridhar et al., 2021]. Furthermore, agents might be allowed to visit multiple options of stimulus sources before converging on one, as is common in collective decisions of ants and honeybees [Hölldobler and Wilson, 1990, Reina et al., 2017]. Another extension of our simulations could be to let agents start from different spatial locations. Using the HKB equations to maintain asymmetric patterns of coordination between different agents’ oscillators, future work could study a rudimentary allocentric way of using social information [Pickavance et al., 2018]. Future work could also study the influence of noise on collective performance of embodied neural agents, especially in environments with many local optima, as random fluctuations are often an important aspect of self-organized systems and collective intelligence [Eric Bonabeau, 1999, Kahneman et al., 2022].

Inspired by the brain dynamics of biological agents, we used oscillator models to study neural dynamics in a multi-agent system. Recent *swarmalator* models have also combined collective movement with oscillator dynamics [O’Keefe et al., 2017, Ceron et al., 2023]. In these models, an agent’s behavior is based on the directly observed oscillator phases of the surrounding agents. In our model, the oscillators represent brain dynamics that are not directly available to an outside observer. The brain dynamics of the agents can only become coordinated by intermediary of their behavior. Moreover, since the stimulus emitted by agents was indistinguishable from stimulus originating in the environment, agents could not selectively react to instantaneous social stimulation. Yet, agents could interact by mutually reacting to local changes in stimulus concentration caused by their respective movements in the environment. This situation reflects a form of social interaction in which agents cannot use any social cognition capabilities other than those involved in the

interaction itself. Such situations are also studied with human participants in perceptual crossing experiments, and have led some to argue that social interaction can be constitutive of social cognition [Auvray et al., 2009, De Jaegher et al., 2010].

Using oscillators to model brain dynamics allows us to use phase-locking and phase-covariance measures to quantify the degree of coordination between the brain dynamics of different agents [Czeszumski et al., 2020]. Experimental studies with multiple animals or humans have suggested that coordination between brain dynamics of different agents, typically quantified in terms of inter-brain synchronization (IBS), might have an important role in supporting collective behaviors [Dumas et al., 2010, Yang et al., 2021]. A few computational models have provided initial mechanistic explanations for the emergence of such inter-brain synchrony by using the Kuramoto model, showing that the strength and frequencies at which IBS takes place depend on a combination of agents’ individual brain dynamics and their inter-agent coupling [Dumas et al., 2012, Moreau et al., 2022]. Although our results cannot conclusively show whether collective decision-making performance depended on inter-agent synchrony, our models could provide a way to study the complex brain–brain behavior dynamics that can give rise to IBS. Furthermore, our results replicate an interesting finding of Kuramoto models of interpersonal synchrony, namely that some degree of intra-agent coupling is required to achieve rich patterns of interpersonal coordination [Heggli et al., 2019]. While previous studies (e.g., [Heggli et al., 2019]) have shown this requirement when a pair of agents were coupled to each other directly, we have confirmed it for multiple agents embedded in a spatial environment.

A major challenge in the development of artificial agents is coordinating social interactions with both humans and other artificial agents. Recent developments in Social NeuroAI attempt to bring social interaction into the realm of AI by advancing artificial agents’ social embodiment, temporal dynamics, and biological plausibility [Bolotta and Dumas, 2022]. In this work, we accommodate 1) social embodiment, as collective decisions are movement-based and agents can be reciprocally influenced by each other’s movements; 2) temporal dynamics, through continuous intra-agent, inter-agent, and agent–environment interactions; and 3) biological plausibility, by using oscillations to control agent behavior. Our approach could be a starting point for developing social-neural agents that collaborate on a wider range of collective tasks through the implicit coordination of their neural dynamics.

## 5 Methods

### 5.1 Experiment setup

Each experiment took place in a 2D environment in which every position had an associated stimulus concentration. Depending on the experiment type, each environment contained one or more stimulus sources of different quality. The stimulus concentration at a certain position was exponentially proportional to its closeness to the stimulus source:

$$I(x, y) = e^{-\lambda D(x, y)}, \quad (4)$$

where  $D(x, y)$  is the Euclidean distance to the stimulus source and  $\lambda = 0.02$  is the exponential decay rate of the environmental stimulus. In setups with two stimulus sources, each had stimulus concentrations defined by eq. 4, and the overall stimulus concentration at a certain position was the combination of the two:

$$I(x, y) = e^{-\lambda D_1(x, y)} + r e^{-\lambda D_2(x, y)}, \quad (5)$$

where  $r$  indicates the quality ratio of the two stimulus sources. When there were multiple agents in an environment, the stimulus level that an agent perceived was a combination of the stimulus concentration in the environment and the stimuli concentrations emitted by other agents, such that

$$I_i(x, y) = e^{-\lambda D_1(x, y)} + r e^{-\lambda D_2(x, y)} + S * \sum_j e^{-\lambda_s D_{ij}}. \quad (6)$$

In all experiments, the environment was 300 by 400 cm, the radius of each agent’s body was 2.5 cm, and each agent has a fixed velocity of 10 cm/s. Simulations were performed with a timestep of 0.01 s. All experiments ended after 30 s, which provided sufficient time for agents to reach a stimulus source in the environment. All simulations were performed in Python version 3.9.2 [Van Rossum and Drake, 2009] and the agents were implemented in Pytorch version 1.12.0 [Paszke et al., 2019]. The code is available in an open-source code repository: <https://github.com/ppsp-team/PyHKBs>.

### 5.2 Agent design

In our agents, sensory input did not directly control motor activity. Sensory information (in the form of stimulus concentration) was first integrated into the oscillator phase of two sensory nodes. These sensory nodes were dynamically connected to two motor nodes.

The situated agent designed by Aguilera et al. [2013] consisted of one motor oscillator and one sensory oscillator, and thus could only perform gradient ascent with spiraling movement. We resolved this by giving our agent two sensory oscillators for stereovision ( $v_1$  and  $v_2$ , see nodes 1 and 2 in Fig. 1A) and two motor oscillators for differential drive steering ( $v_3$  and  $v_4$ , see nodes 3 and 4 in Fig. 1A). The sensors are directionless and are placed at the front of the agent,  $90^\circ$  apart as measured from the agent’s center (see Fig. 1a). The orientation  $\theta$  of the agent in the environment is determined by the angle between the two motor oscillators ( $v_3$  and  $v_4$ , see nodes 3 and 4 in Fig. 1A).

Altogether, the dynamics of the agent are governed by the following set of equations:

$$\dot{\phi} \begin{cases} \dot{\phi}_{v_1} = \omega_{v_1} + cI_l - \sum_j a_{v_1,v_j} \sin \phi_{v_1,v_j} - \sum_j b_{v_1,v_j} \sin 2\phi_{v_1,v_j}, \\ \dot{\phi}_{v_2} = \omega_{v_2} + cI_r - \sum_j a_{v_2,v_j} \sin \phi_{v_2,v_j} - \sum_j b_{v_2,v_j} \sin 2\phi_{v_2,v_j}, \\ \dot{\phi}_{v_3} = \omega_{v_3} - \sum_j a_{v_3,v_j} \sin \phi_{v_3,v_j} - \sum_j b_{v_3,v_j} \sin 2\phi_{v_3,v_j}, \\ \dot{\phi}_{v_4} = \omega_{v_4} - \sum_j a_{v_4,v_j} \sin \phi_{v_4,v_j} - \sum_j b_{v_4,v_j} \sin 2\phi_{v_4,v_j} \end{cases}, \quad (7)$$

$$\dot{\theta} = \eta \phi_{v_3,v_4} = \eta(\varphi_{v_3} - \varphi_{v_4}), \quad (8)$$

where we fixed the ratio  $k = \frac{b}{a} = 2$  so that the HKB equations are bistable (see supplementary text S1).

In our neural controller, the oscillators influenced each other over the following connections: the contralateral ones ( $a_{v_1,v_4} = a_{v_4,v_1}$  and  $a_{v_2,v_3} = a_{v_3,v_2}$ ) the one between the motor regions ( $a_{v_3,v_4} = a_{v_4,v_3}$ ), as well as their antiphase counterparts. Thus, we kept the two sensory oscillators independent and incorporated the contralateral sensorimotor connections present in the Braitenberg vehicles [Braitenberg, 1986] and in many of the biological neural organizations [Sterling and Laughlin, 2017]. The intrinsic frequencies of all oscillators were set to 5 Hz, to resemble the frequency of the theta oscillations in biological brains. In our model, the next phase of each oscillator is calculated at each time step, by integrating the differential equations using the fourth-order Runge-Kutta method.

### 5.2.1 Single-agent experiments

In the single-agent gradient ascent setup, the agent initiated movement at  $xy$  position (0, -100) and the stimulus source was located at  $xy$  position (-100, 0). To study the link between a single agent’s behavior and its intra-agent neural dynamics, we varied the sensory sensitivity ( $c = 0$  or 5, in Eq. 7) and the coupling strength of all connections ( $a_{v_i,v_j}$  values from 0.05 to 2.5, in steps of 0.05, in Eq. 7). For each variation combination, we performed 50 runs with random initial phases of the oscillators.

In the single-agent binary decision making setup, the first stimulus source was located at  $xy$  position (-100, 0), the second at  $xy$  position (100, 0), and the brightness (i.e., quality) ratio of the two stimulus sources is  $r = 0.95$  (see Eq. 6). The agent initiated movement equidistant to the two stimulus sources, at  $xy$  position (0, -100) with all internal oscillators starting as in-phase. To evaluate the dependence of performance on agent behavior, we varied the stimulus sensitivity of the agent ( $c$  values from 0 to 10, in steps of 1, in Eq. 7) and the internal coupling ( $a_{v_i,v_j}$  values from 0.05 to 2.5, in steps of 0.05, in Eq. 7). To evaluate the importance of internal coupling, we also varied whether the motor regions were connected or not ( $a_{v_3,v_4} = 0$ , in Eq 7). We ran one simulation for each variation combination, since we did not introduce an element of randomness in the simulation.

### 5.2.2 Multi-agent experiments

Each multi-agent experiment had a group of 10 agents and an environment with two stimulus sources, located at  $xy$  positions (-100, 0) and (100, 0). To study consensus achievement under divergent starting opinions, we evenly distributed the initial orientations of the agents (between angle  $-\theta_{\max}$  and  $\theta_{\max}$ ,  $\theta_{\max} \in [0^\circ, 180^\circ]$ ). Thus, each agent faced a different initial direction, with half facing more towards the lefthand stimulus source and half facing more towards the righthand one. Following Nabet et al. [2009], Leonard et al. [2011], the agent behavior in these experiments was deliberately deterministic. Although noise can be highly beneficial to the self-organization of complex systems [Kahneman et al., 2022], our focus in this study was specifically on the relationship between inter-agent dynamics and consensus, rather than exploring how noise might modulate these dynamics.

We ran two groups of multi-agent experiments. In the first group, to study the influence of the intra- and inter-agent coordination regimes, we varied the degree of internal coupling between the motor oscillators ( $a_{v_i,v_j}$  values from 0 to 1, in steps of 0.02, in Eq. 7), the social sensitivity ( $S$  values from 0 to 5, in steps of 0.1, in Eq. 3), and the stimulus sensitivity ( $c$  values from 0 and 10, in steps of 0.5, in Eq. 7). In these experiments, the starting angle between agents was  $10^\circ$  and the brightness (i.e., quality) ratio of the two stimulus sources was  $r = 0.8$ . This ratio is lower than that in the single-agent case, to facilitate a wider range of collective behaviors. With a higher ratio, agents initially oriented towards the least bright stimulus source did not deviate enough from their initial movement path for collective dynamics to occur.

In the second group, to study the influence of the environmental and initial conditions, we varied the brightness (i.e., quality) ratio of the two stimulus sources ( $r$  values from 0 to 1, in steps of 0.02) and the starting angles of the agents (from  $0^\circ$  to  $18^\circ$ , in steps of  $0.36^\circ$ ). In these experiments, the stimulus sensitivity was  $c = 3$ , social sensitivity is  $S = 1$ , and internal connection was  $a_{v_i, v_j} = 0.5$ .

### 5.3 Evaluation

In single-agent setups, performance is based on the agent’s end position relative to a stimulus source. Note that, in these experiments, the agent could continue moving after reaching a source, so the performance metric includes how well the agent remained close to a source after initially approaching it. For gradient ascent, we evaluated performance based on how closely the agent approaches the stimulus source:

$$\text{performance} = 1 - \frac{D(t_{\text{end}})}{D(t_0)}, \quad (9)$$

where  $D(t_0)$  and  $D(t_{\text{end}})$  represent the agent’s distance to the stimulus source at the beginning and end of the simulation. For binary decision making, we evaluate how closely the agent approaches its closest stimulus source, regardless of the source’s brightness level:

$$\text{performance} = 1 - \frac{\min \{D_{\text{source}_1}(t_{\text{end}}), D_{\text{source}_2}(t_{\text{end}})\}}{D(t_0)}, \quad (10)$$

where  $D_{\text{source}_1}(t_{\text{end}})$  and  $D_{\text{source}_2}(t_{\text{end}})$  represent the agent’s distance to `source1` and `source2` at the end of the simulation. In multi-agent setups, performance is based on whether agents reach a consensus. Note that, in these experiments, an agent could no longer move once it came within 5 cm of a stimulus source, so any changes in agent angle and position due to agents circling around the stimulus source after arrival did not influence the decisions of the other agents. We evaluated performance based on the smallest average distance to one of the two stimulus sources:

$$\text{performance} = \min \left\{ \frac{1}{N} \sum_{n=1}^N \left[ 1 - \frac{D_{\text{source}_1}^n(t_{\text{end}})}{D(t_0)} \right], \frac{1}{N} \sum_{n=1}^N \left[ 1 - \frac{D_{\text{source}_2}^n(t_{\text{end}})}{D(t_0)} \right] \right\}, \quad (11)$$

with  $N$  being the number of agents and  $D_{\text{source}_1}^n(t)$  being the Euclidean distance from `source1` to agent  $n$  at time  $t$ .

#### 5.3.1 Measures of coordination dynamics

To evaluate the intra-agent neural dynamics, inter-agent neural dynamics, and inter-agent behavioral (i.e., movement) dynamics, we used the following measures: Kuramoto order parameters (KOP) [Strogatz, 2000], phase locking value (PLV) [Lachaux et al., 1999], and weighted phase-lag index (wPLI) [Vinck et al., 2011].

First, we calculated KOP (i.e., the parameter  $R(t)$  [Strogatz, 2000]) as

$$Z(t) = R(t)e^{i\Theta t} = \frac{1}{N} \sum_{i=1}^N e^{i\varphi_i(t)}, \quad (12)$$

where  $N$  is the number of oscillators and  $\varphi_i$  is the phase angle of each oscillator (which, in this study, can be either the oscillator nodes of the neural controller or the movement directions of agents). KOP quantifies the extent to which several oscillating components are in phase. If KOP is 1, all components are completely in phase, whereas low KOP values indicate an absence of synchronization between components. KOP values remaining constant over time indicate that the system has resorted to a stable dynamic (whether synchronized or not), whereas variation in the parameter indicate that the system is passing through various coordination states. Therefore, the standard deviation (SD) of KOP can be used as a measure of the metastability of coordination between oscillating components [Shanahan, 2010, Cabral et al., 2022]. We used metastability to assess the agents’ neural dynamics but also to evaluate the collective movements in the multi-agent simulations. In the latter case, we used the measure of metastability to quantify the degree of ‘alignment variability’, with which we mean the degree to which the collective switches between aligned and unaligned movement directions.

Based on Lachaux et al. [1999], we calculated sliding PLV <sub>$ij$</sub>  for the connection between  $v_i$  and  $v_j$  as

$$\text{PLV}_{ij} = \frac{1}{T} \left| \sum_{t=1}^T e^{i\phi_{ij}(t)} \right|, \quad (13)$$

where  $T$  is the number of samples in a window. PLV is different from KOP in that it is maximal if the phases of the two oscillators are ‘locked’, i.e., the relative phase of oscillators remains constant over time. We use PLV as a measure to indicate the degree of *integration* of the oscillating components, i.e., the degree to which their phases are co-determined. PLV is often used as a measure of connectivity between brain components [Varela et al., 2001] as well as functional connectivity between the brains of different individuals during social interaction [Dumas et al., 2010].

Finally,  $w\text{PLI}_{ij}$  for the connection between  $v_i$  and  $v_j$  is calculated as follows [Vinck et al., 2011]:

$$w\text{PLI}_{ij} = \frac{\frac{1}{T} \left| \sum_{t=1}^T |I_{ij}(t)| \operatorname{sgn}(I_{ij}(t)) \right|}{\frac{1}{T} \sum_{t=1}^T |I_{ij}(t)|}, \quad (14)$$

with

$$I_{ij}(t) = \operatorname{Imag}(e^{i\phi_{ij}(t)}). \quad (15)$$

Like the PLV, the  $w\text{PLI}$  characterizes to what degree different oscillators are integrated [Vinck et al., 2011] and has been used in several hyperscanning studies to quantify synchronization between brain regions [e.g., Schwartz et al., 2022].  $w\text{PLI}$  differs from PLV in that its weighting of phases puts more emphasis on the covariance of phases than simple ‘locking’. When zero-phase locking (i.e., completely synchronized activity) driven by common input needs to be distinguished from locking between other phases,  $w\text{PLI}$  provides a more robust characterization of oscillator integration than PLV.

## Acknowledgments

We thank Andreagiovanni Reina for useful suggestions to the manuscript. This work was partially supported by the program of Concerted Research Actions (ARC) of the Université libre de Bruxelles and a Mitacs Globalink Research Award to N. Coucke. M.K. Heinrich, A. Cleeremans, and M. Dorigo acknowledge support from the F.R.S.-FNRS, of which they are, respectively, postdoctoral researcher and research directors.

## References

- Dan Bang and Chris D. Frith. Making better decisions in groups. *Royal Society Open Science*, 4(8):170193, aug 2017. doi:10.1098/rsos.170193.
- Larissa Conrard and Christian List. Group decisions in humans and animals: a survey. *Philosophical Transactions of the Royal Society B*, 364(1518):719–742, dec 2008. doi:10.1098/rstb.2008.0276.
- Heiko Hamann, Thomas Schmickl, Heinz Wörn, and Karl Crailsheim. Analysis of emergent symmetry breaking in collective decision making. *Neural Computing and Applications*, 21(2):207–218, apr 2010. doi:10.1007/s00521-010-0368-6.
- Marco A. Montes de Oca, Eliseo Ferrante, Alexander Scheidler, Carlo Pinciroli, Mauro Birattari, and Marco Dorigo. Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3–4):305–327, 2011.
- Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4, mar 2017. doi:10.3389/frobt.2017.00009.
- Iain D. Couzin, Christos C. Ioannou, Güven Demirel, Thilo Gross, Colin J. Torney, Andrew Hartnett, Larissa Conrard, Simon A. Levin, and Naomi E. Leonard. Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580, dec 2011. doi:10.1126/science.1210280.
- Joshua Becker, Devon Brackbill, and Damon Centola. Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, page 201615978, jun 2017. doi:10.1073/pnas.1615978114.
- Damon Centola. The network science of collective intelligence. *Trends in Cognitive Sciences*, 26(11):923–941, 2022.
- Iain D. Couzin, Jens Krause, Nigel R. Franks, and Simon A. Levin. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516, feb 2005. doi:10.1038/nature03236.
- Seng Bum Michael Yoo, Benjamin Yost Hayden, and John M. Pearson. Continuous decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1819):20190664, jan 2021. doi:10.1098/rstb.2019.0664.
- Shinsuke Suzuki, Ryo Adachi, Simon Dunne, Peter Bossaerts, and John P. O’Doherty. Neural mechanisms underlying human consensus decision-making. *Neuron*, 86(2):591–602, apr 2015. doi:10.1016/j.neuron.2015.03.019.
- Rainer Hegselmann and Ulrich Krause. Opinion dynamics and bounded confidence models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5, 07 2002.

- Serge Galam. Sociophysics: A review of galam models. *International Journal of Modern Physics C*, March 2008. doi:10.1142/S0129183108012297.
- Andreas Flache, Michael Mäs, Thomas Feliciani, Edmund Chattoe-Brown, Guillaume Deffuant, Sylvie Huet, and Jan Lorenz. Models of social influence: Towards the next frontiers. *Journal of Artificial Societies and Social Simulation*, 20(4), 2017. doi:10.18564/jasss.3521.
- Vaibhav Srivastava and Naomi Ehrich Leonard. Collective decision-making in ideal networks: The speed-accuracy trade-off. *IEEE Transactions on Control of Network Systems*, 1(1):121–132, mar 2014. doi:10.1109/tcns.2014.2310271.
- Alan N. Tump, Timothy J. Pleskac, and Ralf H. J. M. Kurvers. Wise or mad crowds? the cognitive mechanisms underlying information cascades. *Science Advances*, 6(29), jul 2020. doi:10.1126/sciadv.abb0266.
- Rakshith Lokesh, Seth Sullivan, Jan A. Calalo, Adam Roth, Brenden Swanik, Michael J. Carter, and Joshua G. A. Cashaback. Humans utilize sensory evidence of others’ intended action to make online decisions. *Scientific Reports*, 12(1), may 2022. doi:10.1038/s41598-022-12662-y.
- Andreagiovanni Reina, Thomas Bose, Vaibhav Srivastava, and James A. R. Marshall. Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders. *Royal Society Open Science*, 10(3), mar 2023. doi:10.1098/rsos.230175.
- Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, aug 1995. doi:10.1103/physrevlett.75.1226.
- Dirk Helbing and Peter Molnar. Social force model for pedestrian dynamics. *Physical review E*, 51(5):4282, 1995.
- Yi Ma, Richard Kwok Kit Yuen, and Eric Wai Ming Lee. Effective leadership for crowd evacuation. *Physica A: Statistical Mechanics and its Applications*, 450:333–341, may 2016. doi:10.1016/j.physa.2015.12.103.
- Andreagiovanni Reina, James AR Marshall, Vito Trianni, and Thomas Bose. Model of the best-of-n nest-site selection process in honeybees. *Physical Review E*, 95(5):052411, 2017.
- Mehdi Moussaïd, Dirk Helbing, and Guy Theraulaz. How simple rules determine pedestrian behavior and crowd disasters. *Proceedings of the National Academy of Sciences*, 108(17):6884–6888, 2011.
- Lyle Kingsbury and Weizhe Hong. A multi-brain framework for social interaction. *Trends in neurosciences*, 43(9): 651–666, 2020.
- Quentin Moreau and Guillaume Dumas. Beyond “correlation vs. causation”: multi-brain neuroscience needs explanation. *Trends Cogn. Sci.*, 25:542–543, 2021.
- Georgina Montserrat Reséndiz-Benhumea, Ekaterina Sangati, Federico Sangati, Soheil Keshmiri, and Tom Froese. Shrunken social brains? a minimal model of the role of social interaction in neural complexity. *Frontiers in Neurobotics*, 15:634085, 2021.
- Vivek H. Sridhar, Liang Li, Dan Gorbonos, Máté Nagy, Bianca R. Schell, Timothy Sorochkin, Nir S. Gov, and Iain D. Couzin. The geometry of decision-making in individuals and collectives. *Proceedings of the National Academy of Sciences*, 118(50), dec 2021. doi:10.1073/pnas.2102157118.
- Conor Heins, Beren Millidge, Lancelot Da Costa, Richard P Mann, Karl J Friston, and Iain D Couzin. Collective behavior from surprise minimization. *Proceedings of the National Academy of Sciences*, 121(17):e2320239121, 2024.
- Andy Clark. *Being There*. The MIT Press, 1998. ISBN 9780262531566.
- Albert Newen and Shaun Gallagher. *The Oxford Handbook of 4E Cognition*. Oxford University Press, 2018. ISBN 9780198735410.
- Luc Steels and Rodney Brooks. *The artificial life route to artificial intelligence: Building embodied, situated agents*. Routledge, 2018.
- Cédric Colas, Tristan Karch, Clément Moulin-Frier, and Pierre-Yves Oudeyer. Language and culture internalization for human-like autotelic AI. *Nature Machine Intelligence*, 4(12):1068–1076, dec 2022. doi:10.1038/s42256-022-00591-4.
- Francisco J. Varela, Evan T. Thompson, and Eleanor Rosch. *The Embodied Mind*. The MIT Press, 1992. ISBN 9780262720212.
- Evan Thompson and Francisco J. Varela. Radical embodiment: neural dynamics and consciousness. *Trends in Cognitive Sciences*, 5(10):418–425, oct 2001. doi:10.1016/s1364-6613(00)01750-2.
- Tom Froese and Ezequiel A. Di Paolo. The enactive approach. *Pragmatics and Cognition*, 19(1):1–36, jul 2011. doi:10.1075/pc.19.1.01fro.

- Lawrence M Ward. Synchronous neural oscillations and cognitive processes. *Trends in cognitive sciences*, 7(12): 553–559, 2003.
- György Buzsáki. *The Brain from Inside Out*. Oxford University Press, jun 2019. doi:10.1093/oso/9780190905385.001.0001.
- György Buzsáki, Nikos Logothetis, and Wolf Singer. Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. *Neuron*, 80(3):751–764, oct 2013. doi:10.1016/j.neuron.2013.10.002.
- Mehdi Senoussi, Pieter Verbeke, Kobe Desender, Esther De Loof, Durk Talsma, and Tom Verguts. Theta oscillations shift towards optimal frequency for cognitive control. *Nature Human Behaviour*, 6(7):1000–1013, 2022.
- Efrosini Charalambous and Zakaria Djebbara. On natural attunement: shared rhythms between the brain and the environment. *Neuroscience & Biobehavioral Reviews*, 155:105438, 2023.
- J. A. Scott Kelso. *Dynamic Patterns The Self-organization Of Brain And Behavior*. Bradford Book, 1997. ISBN 9780262611312.
- J. A. Scott Kelso, Emmanuelle Tognoli, and Guillaume Dumas. Coordination dynamics: Bidirectional coupling between humans, machines and brains. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, oct 2014. doi:10.1109/smc.2014.6974258.
- Emmanuelle Tognoli, Mengsen Zhang, Armin Fuchs, Christopher Beetle, and J. A. Scott Kelso. Coordination dynamics: A foundation for understanding social behavior. *Frontiers in Human Neuroscience*, 14, aug 2020. doi:10.3389/fnhum.2020.00317.
- Emmanuelle Tognoli and J. A. Scott Kelso. The metastable brain. *Neuron*, 81(1):35–48, jan 2014. doi:10.1016/j.neuron.2013.12.022.
- Hermann Haken, J. A. Scott Kelso, and H. Bunz. A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51(5):347–356, feb 1985. doi:10.1007/bf00336922.
- Y. Kuramoto. *Chemical Oscillations, Waves, and Turbulence*. Springer Berlin Heidelberg, 1984. ISBN 9783642696916.
- J. A. Scott Kelso. Coordination dynamics. In *Encyclopedia of Complexity and Systems Science*, pages 1–41. Springer New York, 2013. doi:10.1007/978-3-642-27737-5\_101-3.
- Miguel Aguilera, Manuel G. Bedia, Bruno A. Santos, and Xabier E. Barandiaran. The situated HKB model: how sensorimotor spatial coupling can alter oscillatory brain dynamics. *Frontiers in Computational Neuroscience*, 7, 2013. doi:10.3389/fncom.2013.00117.
- Naomi E. Leonard, Tian Shen, Benjamin Nabet, Luca Scardovi, Iain D. Couzin, and Simon A. Levin. Decision versus compromise for animal groups in motion. *Proceedings of the National Academy of Sciences*, 109(1):227–232, dec 2011. doi:10.1073/pnas.1118318108.
- Valentino Braitenberg. *Vehicles, Experiments in Synthetic Psychology*. M.I.T. P., 1986. ISBN 0262521121.
- Mengsen Zhang, Christopher Beetle, J. A. Scott Kelso, and Emmanuelle Tognoli. Connecting empirical phenomena and theoretical models of biological coordination across scales. *Journal of The Royal Society Interface*, 16(157): 20190360, aug 2019. doi:10.1098/rsif.2019.0360.
- Francisco Varela, Jean-Philippe Lachaux, Eugenio Rodriguez, and Jacques Martinerie. The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, 2(4):229–239, apr 2001. doi:10.1038/35067550.
- Andrea Avena-Koenigsberger, Bratislav Misic, and Olaf Sporns. Communication dynamics in complex brain networks. *Nature reviews neuroscience*, 19(1):17–33, 2018.
- Steven H. Strogatz. From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1-4):1–20, sep 2000. doi:10.1016/s0167-2789(00)00094-4.
- Joana Cabral, Francesca Castaldo, Jakub Vohryzek, Vladimir Litvak, Christian Bick, Renaud Lambiotte, Karl Friston, Morten L. Kringelbach, and Gustavo Deco. Metastable oscillatory modes emerge from synchronization in the brain spacetime connectome. *Communications Physics*, 5(1), jul 2022. doi:10.1038/s42005-022-00950-y.
- Martin Vinck, Robert Oostenveld, Marijn van Wingerden, Francesco Battaglia, and Cyriel M.A. Pennartz. An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias. *NeuroImage*, 55(4):1548–1565, apr 2011. doi:10.1016/j.neuroimage.2011.01.055.
- Artur Czeszumski, Sara Eustergerling, Anne Lang, David Menrath, Michael Gerstenberger, Susanne Schuberth, Felix Schreiber, Zadkiel Zuluaga Rendon, and Peter König. Hyperscanning: A valid method to study neural inter-brain underpinnings of social interaction. *Frontiers in Human Neuroscience*, 14, feb 2020. doi:10.3389/fnhum.2020.00039.



- Linoy Schwartz, Jonathan Levy, Yaara Endevelt-Shapira, Amir Djalovski, Olga Hayut, Guillaume Dumas, and Ruth Feldman. Technologically-assisted communication attenuates inter-brain synchrony. *NeuroImage*, 264:119677, dec 2022. doi:10.1016/j.neuroimage.2022.119677.
- John Waller. *A time to dance, a time to die*. Icon Books, 2008. ISBN 9781848310216.
- I. D. Couzin and N. R. Franks. Self-organized lane formation and optimized traffic flow in army ants. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1511):139–146, jan 2003. doi:10.1098/rspb.2002.2210.
- Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, and Marco Dorigo. Collective decision with 100 kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3):553–580, dec 2015. doi:10.1007/s10458-015-9323-3.
- Bert Hölldobler and Edward Osborne Wilson. *The Ants*. Belknap Press, 1990. ISBN 9780674040755.
- John Pickavance, Arianne Azmoodeh, and Andrew D. Wilson. The effects of feedback format, and egocentric & allocentric relative phase on coordination stability. *Human Movement Science*, 59:143–152, jun 2018. doi:10.1016/j.humov.2018.04.005.
- Guy Theraulaz Eric Bonabeau, Marco Dorigo. *Swarm Intelligence*. Oxford University Press Inc, 1999. ISBN 0195131592. URL [https://www.ebook.de/de/product/3253990/eric\\_postdoctoral\\_fellow\\_postdoctoral\\_fellow\\_santa\\_fe\\_institute\\_bonabeau\\_marco\\_researcher\\_researcher\\_free\\_university\\_of\\_brussels\\_dorigo\\_guy\\_researcher\\_researcher\\_cnrs\\_university\\_paul\\_sabatier\\_theraulaz\\_swarm\\_intelligence.html](https://www.ebook.de/de/product/3253990/eric_postdoctoral_fellow_postdoctoral_fellow_santa_fe_institute_bonabeau_marco_researcher_researcher_free_university_of_brussels_dorigo_guy_researcher_researcher_cnrs_university_paul_sabatier_theraulaz_swarm_intelligence.html).
- Daniel Kahneman, David C Krakauer, Olivier Sibony, Cass Sunstein, and David Wolpert. An exchange of letters on the role of noise in collective intelligence. *Collective Intelligence*, 1(1):263391372210785, aug 2022. doi:10.1177/26339137221078593.
- Kevin P. O’Keefe, Hyunsuk Hong, and Steven H. Strogatz. Oscillators that sync and swarm. *Nature Communications*, 8(1), nov 2017. doi:10.1038/s41467-017-01190-3.
- Steven Ceron, Kevin O’Keefe, and Kirstin Petersen. Diverse behaviors in non-uniform chiral and non-chiral swarmalators. *Nature Communications*, 14(1), feb 2023. doi:10.1038/s41467-023-36563-4.
- Malika Auvray, Charles Lenay, and John Stewart. Perceptual interactions in a minimalist virtual environment. *New ideas in psychology*, 27(1):32–47, 2009.
- Hanne De Jaegher, Ezequiel Di Paolo, and Shaun Gallagher. Can social interaction constitute social cognition? *Trends in Cognitive Sciences*, 14(10):441–447, oct 2010. doi:10.1016/j.tics.2010.06.009.
- Guillaume Dumas, Jacqueline Nadel, Robert Soussignan, Jacques Martinerie, and Line Garnero. Inter-brain synchronization during social interaction. *PLoS ONE*, 5(8):e12166, aug 2010. doi:10.1371/journal.pone.0012166.
- Yiyuan Yang, Mingzheng Wu, Abraham Vázquez-Guardado, Amy J. Wegener, Jose G. Grajales-Reyes, Yujun Deng, Taoyi Wang, Raudel Avila, Justin A. Moreno, Samuel Minkowicz, Vasin Dumrongprechachan, Jungyup Lee, Shuangyang Zhang, Alex A. Legaria, Yuhang Ma, Sunita Mehta, Daniel Franklin, Layne Hartman, Wubin Bai, Mengdi Han, Hangbo Zhao, Wei Lu, Yongjoon Yu, Xing Sheng, Anthony Banks, Xinge Yu, Zoe R. Donaldson, Robert W. Gereau, Cameron H. Good, Zhaoqian Xie, Yonggang Huang, Yevgenia Kozorovitskiy, and John A. Rogers. Wireless multilateral devices for optogenetic studies of individual and social behaviors. *Nature Neuroscience*, 24(7):1035–1045, may 2021. doi:10.1038/s41593-021-00849-x.
- Guillaume Dumas, Mario Chavez, Jacqueline Nadel, and Jacques Martinerie. Anatomical connectivity influences both intra- and inter-brain synchronizations. *PLoS ONE*, 7(5):e36414, may 2012. doi:10.1371/journal.pone.0036414.
- Quentin Moreau, Lena Adel, Cairiona Douglas, Ghazaleh Ranjbaran, and Guillaume Dumas. A neurodynamic model of inter-brain coupling in the gamma band. *Journal of Neurophysiology*, sep 2022. doi:10.1152/jn.00224.2022.
- Ole Adrian Heggli, Joana Cabral, Ivana Konvalinka, Peter Vuust, and Morten L. Kringelbach. A kuramoto model of self-other integration across interpersonal synchronization strategies. *PLOS Computational Biology*, 15(10):e1007422, oct 2019. doi:10.1371/journal.pcbi.1007422.
- Samuele Bolotta and Guillaume Dumas. Social Neuro AI: Social Interaction as the “Dark Matter” of AI. *Frontiers in Computer Science*, 4, 2022. ISSN 2624-9898. URL <https://www.frontiersin.org/articles/10.3389/fcomp.2022.846440>.
- Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009. ISBN 1441412697.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala.

Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.

Peter Sterling and Simon Laughlin. *Principles of Neural Design*. The MIT Press, 2017. ISBN 9780262534680.

Benjamin Nabet, Naomi E. Leonard, Iain D. Couzin, and Simon A. Levin. Dynamics of decision making in animal group motion. *Journal of Nonlinear Science*, 19(4):399–435, jan 2009. ISSN 1432-1467. doi:10.1007/s00332-008-9038-6.

Jean-Philippe Lachaux, Eugenio Rodriguez, Jacques Martinerie, and Francisco J. Varela. Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4):194–208, 1999. doi:10.1002/(sici)1097-0193(1999)8:4<194::aid-hbm4>3.0.co;2-c.

Murray Shanahan. Metastable chimera states in community-structured oscillator networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(1):013108, mar 2010. doi:10.1063/1.3305451.

# Supplementary information for “Collective decision making by embodied neural agents”

Nicolas Coucke, Mary Katherine Heinrich, Axel Cleeremans, Marco Dorigo,  
and Guillaume Dumas

## S1 The agent design and HKB equations

Assume that an oscillator has intrinsic frequency  $\omega$  and phase  $\varphi$ . In a system with two uncoupled oscillators, the phase angle between the oscillators is governed by

$$\dot{\phi} = \delta\omega , \quad (\text{S1})$$

where  $\delta\omega$  is the difference between the intrinsic frequency of the oscillators and  $\phi = \varphi_1 - \varphi_2$  is the instantaneous phase difference between the two oscillators. Realistically, biological brain regions are coupled and influence each other’s activity. The simplest possible interaction between two components is in-phase coupling, i.e., the oscillators will be attracted to synchronous oscillations. This relationship has been modeled by the Kuramoto equation [3]:

$$\dot{\phi} = \delta\omega - a \sin \phi . \quad (\text{S2})$$

Whether a stable coordination state exists for the system depends on the relative sizes of  $\delta\omega$  and  $a$ ; no stable in-phase coordination will occur when  $\delta\omega$  is larger than  $a$ .

Biological components can also have a more complex form of coordination. To explain bistable coordination between hand movements and accommodates antiphase coupling between the oscillators, the Haken-Kelso-Bunz (HKB) equation was introduced [2]:

$$\dot{\phi} = \delta\omega - a \sin \phi - 2b \sin 2\phi , \quad (\text{S3})$$

where  $a$  and  $b$  determine the strength of phase and antiphase coupling, respectively. Whether stable antiphase coupling is present in this system depends on the ratio  $k = \frac{a}{b}$ . Stable antiphase coupling is only possible when  $k < 4$ . Note that even when no stable coordination pattern is present, both the phase and the antiphase coupling might still contribute to creating rich metastable dynamics in which the system continuously moves from phase to antiphase, with durations proportional to the coupling strengths [6].

### S1.1 Situated HKB agents

The broader context in which a system is situated can profoundly affect the coordination behaviors it can exhibit [4]. For example, adding a source of external influence can enhance and inhibit stable coordination regimes [5]. The HKB equation can be applied to situated contexts and used to model a cognitive system’s reaction to sensory inputs from the external environment. In the *situated HKB agent* introduced by [1], the HKB equation was used to model the interaction between sensory and motor brain regions

of the agent. The agent was coupled to the environment so that its intra-agent neural dynamics (phase angle between sensory and motor regions) were modulated by sensory information, such that:

$$\dot{\phi} = \delta\omega + cI - a \sin \phi - 2b \sin 2\phi , \quad (\text{S4})$$

where  $I$  is the intensity of the stimulus and  $c$  is the stimulus sensitivity of the agent. A closed sensorimotor loop was created by controlling the agent's movement in the environment according to the phase angle between the two oscillators, such that

$$\dot{\theta} = \eta\phi , \quad (\text{S5})$$

where  $\theta$  is the orientation of the agent in the environment and  $\eta$  is a scaling factor.

### S1.1.1 HKB equations for $N$ components

To enable modeling of  $N + 1$  interacting components, [7] introduced variations on the HKB equations in which the phase of each oscillator is explicitly modeled, rather than only the phase difference. After rescaling the parameters  $a_{ij}$  and  $b_{ij}$  (i.e., the parameters associated with the connections between  $v_i$  and  $v_j$ ), the HKB equation becomes:

$$\dot{\phi}_i = \delta\omega - \sum_{j=0}^N a_{ij} \sin \phi_{ij} - 2b_{ij} \sin 2\phi_{ij} , \quad (\text{S6})$$

where  $\omega_i$  is the intrinsic frequency of the  $i^{\text{th}}$  component.

## S1.2 Our controller design: social HKB agents

Combining the equations above, we achieve the set of update equations reported in the main text:

$$\dot{\phi} \begin{cases} \dot{\phi}_{v_1} = \omega_{v_1} + cI_l - \sum_j a_{v_1,v_j} \sin \phi_{v_1,v_j} - \sum_j b_{v_1,v_j} \sin 2\phi_{v_1,v_j}, \\ \dot{\phi}_{v_2} = \omega_{v_2} + cI_r - \sum_j a_{v_2,v_j} \sin \phi_{v_2,v_j} - \sum_j b_{v_2,v_j} \sin 2\phi_{v_2,v_j}, \\ \dot{\phi}_{v_3} = \omega_{v_3} - \sum_j a_{v_3,v_j} \sin \phi_{v_3,v_j} - \sum_j b_{v_3,v_j} \sin 2\phi_{v_3,v_j}, \\ \dot{\phi}_{v_4} = \omega_{v_4} - \sum_j a_{v_4,v_j} \sin \phi_{v_4,v_j} - \sum_j b_{v_4,v_j} \sin 2\phi_{v_4,v_j} \end{cases} , \quad (\text{S7})$$

## S1.3 Performance evaluation

The performance for multi-agent consensus achievement is defined as:

$$\text{performance} = \min \left\{ \frac{1}{N} \sum_{n=1}^N \left[ 1 - \frac{D_{\text{source}_1}^n(t_{\text{end}})}{D(t_0)} \right], \frac{1}{N} \sum_{n=1}^N \left[ 1 - \frac{D_{\text{source}_2}^n(t_{\text{end}})}{D(t_0)} \right] \right\} ,$$

with  $N$  being the number of agents and  $D_{\text{source}_1}^n(t)$  being the Euclidean distance from  $\text{source}_1$  to agent  $n$  at time  $t$ .

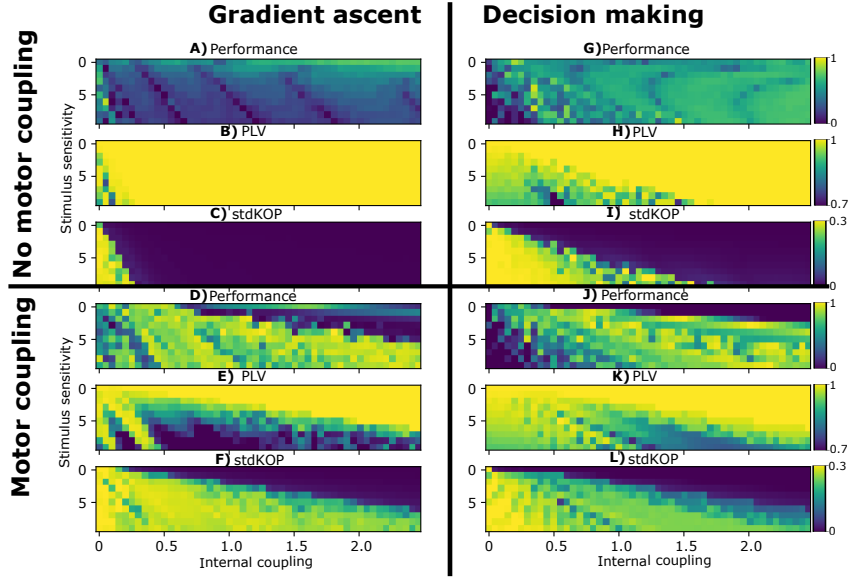


Figure S1: Agent performance and intra-agent neural dynamics in both types of single-agent experiments: (A-F) gradient ascent and (G-L) binary decision-making. Agent variations are either (D-F;J-L) with with motor coupling or (A-C;G-I) without motor coupling. In all agents, the performance and intra-agent dynamics (PLV and SD(KOP)) depend on the stimulus sensitivity  $S$  and the internal coupling  $a_{v_i, v_j}$  of the agent. The 50 different internal coupling values in the plots correspond to  $a_{v_i, v_j} \in \{0, \dots, 2.5\}$  and the 10 values of stimulus sensitivity correspond to  $c \in \{0, \dots, 10\}$  in Eq. S7

## S2 Individual gradient ascent and binary decision-making performance

Also in single-agent experiments, we analyze the effect of varying stimulus sensitivity and internal coupling on agent performance in gradient ascent and binary decision-making. We also analyze how agent performance is associated with intra-agent dynamics. A first observation is that agents without any coupling between motor oscillators (Fig. S1A-I) have, in general, lower performances and a lower degree of metastable dynamics. Adding motor coupling (Fig. S1D-L) increases performance and metastability in the agent’s intra-agent dynamics. It might seem counterintuitive that an extra coupling link in the agent’s architecture leads to lower functional integration of the oscillator dynamics (indicated by lower PLV in E than B, in Fig. S1). Note however, that adding a link between the motor oscillators allows the agent to have a motor response to environmental stimuli that is coordinated between the left and right sensorimotor connections.

In general, the analysis in Fig. S1 shows that the best performance occurs for agents that have a balance of stimulus sensitivity and internal coupling. A high degree of stimulus sensitivity without high internal coupling results in a hypersensitive agent that cannot maintain a relatively stable movement direction. Conversely, high internal coupling without high stimulus sensitivity results in an agent that cannot adequately incorporate sensory information in its internal dynamics, which impedes changing movement direction in response to sensory input.

In our agents, sensory input does not directly control motor activity. Sensory information (in the form of stimulus concentration) is first integrated into the oscillator phase of two sensory nodes. These sensory nodes are dynamically connected to two motor nodes.

## S3 The nonuniform effect of social influence on collective decision-making

The ability of agents to move towards one of the stimulus sources depends on their interactions with the environment: the sensory input related to the stimuli needs to have sufficient influence on the intra-agent dynamics of the agent. The agent also processes social stimuli, which should be integrated with environmental stimuli. If the social stimulation becomes too strong, it can nullify the processing of the environmental stimuli. Indeed, along the right edge of Fig. 5A in the main paper, we see that increased social influence does not uniformly increase performance; beyond some threshold, it inhibits agents’ ability to reach a consensus.

Fig. S2A shows a cross-section of the results shown in Fig. 5A. Increased social influence initially increases the decision-making performance (compare Fig. S2C and Fig. S2D), but with further increased social influence, the effect is strong enough that agents initially overshoot, in the direction that the most social influence was sensed (Fig. S2E). Due to this overshooting, agents split into two groups and end up beyond the range where they can effectively influence the other group’s direction. When social influence is increased even further (Fig. S2G), the effect is so strong that agents’ actions are dependent only on each other, not on the environment, and most agents get into a deadlock and cannot approach a stimulus source.

The neural measures for this cross-section (see Fig. S2B) show that intra-agent wPLI decreases with respect to inter-agent wPLI in the range with the highest performance.

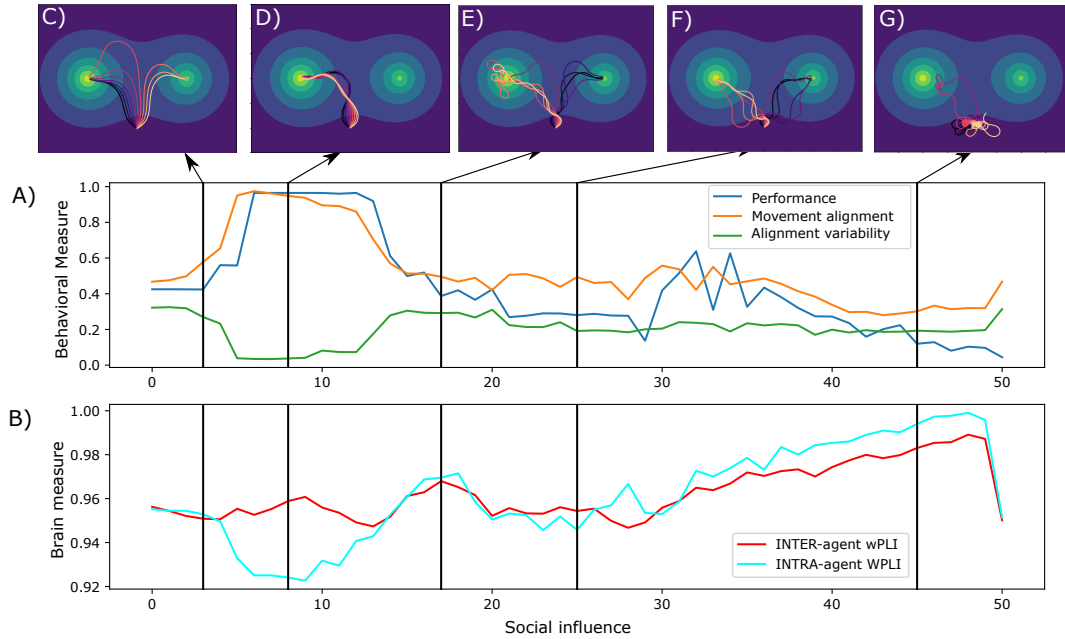


Figure S2: The effect of social influence on decision-making performance. (A) Performance, average alignment (mean KOP) and alignment variability (SD(KOP)), as a function of the degree of social influence (social sensitivity  $S$  of agents). (B) intra- and inter-agent covariance of neural dynamics (wPLI). (C-G) Trajectories of collective decisions taken with increasing social influence.

This indicates that agents achieved consensus when their intrinsic dynamics were perturbed (i.e., uncoupled) by stimulus but remained similar (i.e., coupled) across agents. With sensitivity to social influence increases beyond this range (while internal coupling and environmental sensitivity decrease), both intra- and inter-agent wPLI increase, as social influence starts to drive each agent’s internal dynamics.

## S4 Nonlinear decrease in performance due to environmental variation

To explain the apparent ‘plateaus’ in performance in Fig. 6 of the main paper, Fig. S3 shows the results from one horizontal row of data in Fig. 6. When only the left stimulus source is present (i.e., stimulus ratio 0), all agents converge to it (Fig. S3B). When the stimulus source on the right attains a certain level of brightness (Fig. S3C), one agent is sufficiently attracted to it to be diverted from the group and ‘catapulted’ far away, causing a sharp performance decrease. This continues as the right source gets brighter, causing a rough plateau. At some threshold along the plateau (Fig. S3D), the right source is bright enough to attract the diverted agent to its center, which results in a very slight performance bump (since the agent is also closer to the left stimulus, as a side effect).

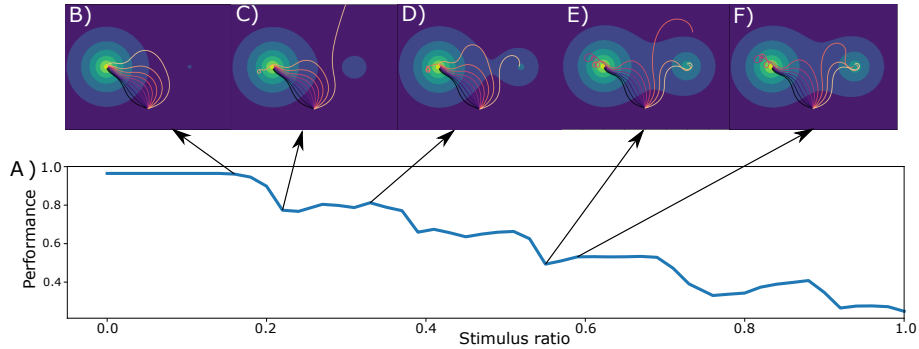


Figure S3: Decision-making performance decrease with increased stimulus ratio. (A) Performance data from Fig. 6, for the  $140^\circ$  starting angle. (B-F) Agent trajectories for different stimulus ratios.

As the right stimulus becomes brighter, more agents are attracted to it, each additional agent causing a sharp decrease in performance (e.g., Fig. S3E) followed by a short plateau (Fig. S3F).

## S5 Qualitative description of unexpected behaviors

When some agents are individually able to reach a stimulus source, our model also shows how the spatial and time-dependent nature of the decision-making process can sometimes impede groups with higher social sensitivity from reaching a consensus. For example, in Fig. S2E, agents start out with different movement directions and perceive high social information relative to stimulus information, which changes their travel directions to face towards the others. If all agents do this simultaneously and strongly, agents can *overshoot*, arriving where other agents previously were, but have since left. By the time agents react to the new conditions, they can have moved so close to the different stimulus sources that they are too far apart to socially influence each other and realign towards a consensus. If social influence were less strong, the initial overshoot would be smaller, allowing agents to remain closer together and better balance social and environmental information (as in Fig. S2d). In extreme cases, some agents might experience a combination of social influence that leaves them moving away from the group and undecided for too long, so that they have moved past the window in which stimulus information was salient and moved further away from both stimulus sources, illustrating how certain patterns of social influence can lead some individuals to isolation.

## References

- [1] Miguel Aguilera, Manuel G. Bedia, Bruno A. Santos, and Xabier E. Barandiaran. The situated HKB model: how sensorimotor spatial coupling can alter oscillatory brain dynamics. *Frontiers in Computational Neuroscience*, 7, 2013.
- [2] Hermann Haken, J. A. Scott Kelso, and H. Bunz. A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, 51(5):347–356, feb 1985.



- [3] Y. Kuramoto. *Chemical Oscillations, Waves, and Turbulence*. Springer Berlin Heidelberg, 1984.
- [4] Joseph McKinley, Mengsen Zhang, Alice Wead, Christine Williams, Emmanuelle Tognoli, and Christopher Beetle. Third party stabilization of unstable coordination in systems of coupled oscillators. *Journal of Physics: Conference Series*, 2090(1):012167, nov 2021.
- [5] G. Schöner and J.A.S. Kelso. A dynamic pattern theory of behavioral change. *Journal of Theoretical Biology*, 135(4):501–524, dec 1988.
- [6] Emmanuelle Tognoli and J. A. Scott Kelso. The metastable brain. *Neuron*, 81(1):35–48, jan 2014.
- [7] Mengsen Zhang, Christopher Beetle, J. A. Scott Kelso, and Emmanuelle Tognoli. Connecting empirical phenomena and theoretical models of biological coordination across scales. *Journal of The Royal Society Interface*, 16(157):20190360, aug 2019.