Quorum-Based Social Interactions Improve Collective Decision Making

Fabio Oddi[1,2], Vito Trianni[2], and Andreagiovanni Reina[3,4]

[1] DIAG, Sapienza University of Rome, Italy
[2] ISTC, National Research Council, Rome, Italy
[3] Centre for the Advanced Study of Collective Behaviour, Universität Konstanz, Konstanz, Germany
[4] Department of Collective Behaviour, Max Planck Institute of Animal Behavior, Konstanz, Germany
oddi@diag.uniroma1.it, vito.trianni@istc.cnr.it, andreagiovanni.reina@gmail.com

1 Problem Description

In the Best-of-N decision-making problem, a swarm of M agents have to collectively select the best among N available options. To do so, agents need to combine environmental observations and social information. What we want to determine is if there is a benefit in changing the frequency with which social information is used to improve collective performance in terms of coherence and accuracy in a time-varying environment. Coherence indicates the average size of the subpopulation with the same opinion (regardless of being the correct or incorrect opinion). Accuracy indicates the probability of the largest subpopulation selecting the best option (regardless of the majority size). Ideally, we would like high coherence and high accuracy.

Our multi-agent model is non-spatial and well-mixed (i.e., any agent can potentially interact with any other agent at any timestep). At each simulation step, one agent is randomly selected to perform its actions. First, it reads the message of $m = 5$ random agents. Exchanged messages contain the agent's preferred opinion (i.e., the option they believe to be the best choice) and the unique agent's ID (i.e., the ID of the agent that sent the opinion). These opinions are stored in a buffer $\mathcal{B}$. The ID is stored to avoid having multiple messages from the same agent. Moreover, a timeout t_m is assigned to each message, indicating the number of agent's state updates, after which it expires. If one received message has the same sender ID as another message already stored in $\mathcal{B}$, then this buffer entry is updated only if their option IDs are different and t_m is reset. The combination of m (the number of messages read per timestep) and the timeout t_m (number of timesteps after a message is deleted) determines the maximum size of the buffer $\mathcal{B}$, i.e., buffer size $|\mathcal{B}| \leq \min(M-1, m\,t_m)$. In our experiments, $m = 5$ and $t_m = 6$, hence $|\mathcal{B}| \leq \min(M-1, 30)$.

Once the agent has updated its buffer $\mathcal{B}$, it updates its opinion either by using social information (with probability r) or by making an individual observation of the environment (with probability $1-r$). In practical terms, the agent makes

a random choice with probability r. If the agent chooses to update its opinion using social information, it applies the Majority Rule and selects the opinion appearing more often in the buffer $\mathcal{B}$. If two different opinions have the same maximum count, then they have the same probability of being randomly chosen. If the agent chooses to make a personal environmental observation, it selects the best option with probability $1 - \eta$ (correct observation) or selects any of the equally worst available options with probability $\eta/(N-1)$ (i.e., makes an incorrect observation with probability η, which can be seen as sensor noise).

Agent's opinion state can be represented as a state machine where the state directly switches to the selected opinion (Direct Switch model). Figure 1 shows how the agent switches or keeps its opinion, based on the probabilities r and η.

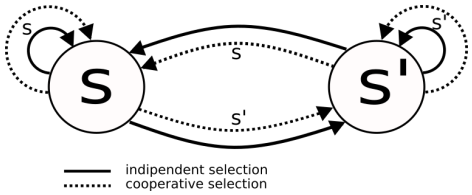


Fig. 1. State machine of the opinions updated through the Direct Switch model

2 Case Studies

We study two different initial conditions to investigate the impact of different levels for the control parameter r. In the first condition (A), the agents are initialised evenly committed to all the available options. This condition represents a fresh start in which agents have no a priori knowledge about the options and choose any one at random. Instead, in the second condition (B), we simulate a dynamic environment by initializing all the agents to the same option which is currently not the best one. This condition represents the case in which the agents previously reached a consensus on an option that is not the best any more (e.g., due to some environmental dynamics). Both the conditions were tested with $N = 2$ and $N = 5$ available options when the decision noise η is high enough to make the decision hard. Since the difficulty of the problem scales also with the number of available options we defined the problem difficulty as

$$\xi = \eta \frac{N}{N-1}, \qquad \xi \in [0, 1]. \tag{1}$$

Where $\xi = 0$ means agents always select the best option, and $\xi = 1$ means all options have the same probability of being selected. Obviously, $\eta \in [0, \frac{N-1}{N}]$, as otherwise agents observe more frequently one of the wrong options than the best option, and solving the best-of-n problem would be impossible. We illustrate the relationship among N, η, and ξ in Figure 2. In our tests, we use $\xi = 0.8$.

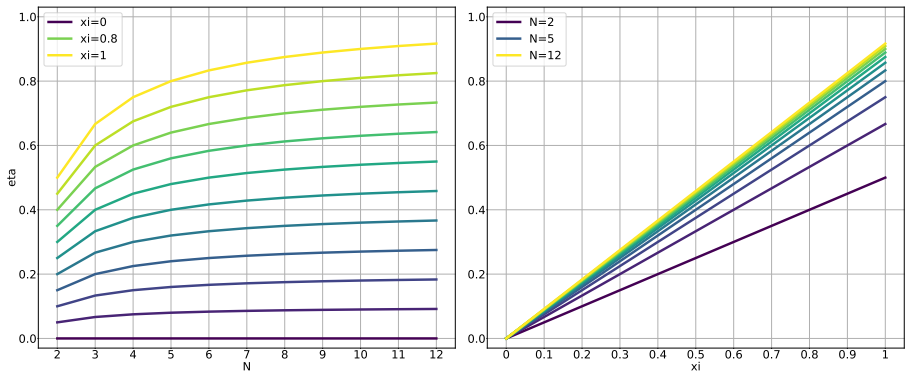


Fig. 2. Relation between η and N for fixed levels of difficulty ξ (left panel). Relation between η and ξ for fixed numbers of options N (right panel).

Static Control Parameter. In this case study we run a set of simulations with different values for the control parameter r which balances the use of social and personal information. The tested values are $r \in \{0.2, 0.4, 0.6, 0.8\}$.

Centralised Control Parameter. Here the quorum is computed in a centralised fashion, corresponding to the situation in which each agent knows exactly how many other agents selected the same option and can therefore compute the quorum Q as the ratio between the number of agents that selected its option plus itself and the size of the entire swarm. We tested two different functions for computing the parameter r based on the quorum level. The first function we propose is a simple linear relation $r = Q$. The second is a V-shaped function $r = V(Q)$. However, since the value $r = 1$ would prevent the system adaptation to a changing environment we decided to limit its values being $r \in [\epsilon, 1 - \epsilon]$. In our tests $\epsilon = 0.005$. Figure 3 illustrates the two functions, linear (L) and V-shaped (V).

Decentralised Control Parameter. This case is equivalent to the previous one but the agents compute the quorum Q by only looking at the messages stored in their buffer $\mathcal{B}$. The computation of the quorum level is allowed only if the buffer size is greater than a minimum fixed size $|\mathcal{B}| > b = 8$. If the minimum size b is not satisfied then there is not enough information to update the quorum level. Depending on the initial condition, the quorum level is initially set equal to zero or one.

3 Results

For each configuration, we run 25 simulations of 10'000 timesteps and use the last timestep to compute the group accuracy (proportion of runs where the

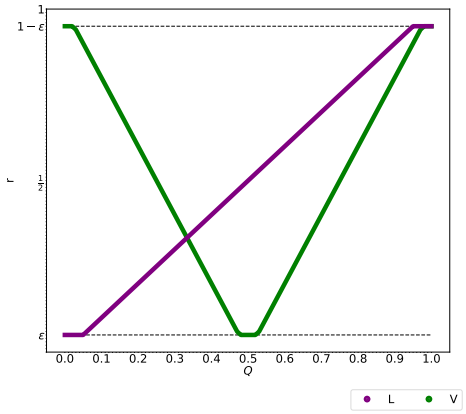


Fig. 3. Linear function (L) in purple and V-shaped function (V) in green of the social interaction frequency r, depending on estimated quorum Q.

largest subpopulation selected the best option) and coherence (average size of the largest subpopulation with the same opinion). The results in Figure 4 show the accuracy-cohesion trade-off for static-control-parameter experiments (i.e., constant r). We report results for both initial conditions: condition A (N equally-sized subpopulations committed to each option) and condition B (the whole population committed to one of the inferior options) for $N = 2$ and $N = 5$ options.

Low values of r (i.e., infrequent social interactions) give high levels of group accuracy, in both conditions and for both tested values of N. However, cohesion is relatively low, especially as the number of options increases (e.g., $N = 5$ and $r \leq 0.4$). Results for higher values of r (i.e., more frequent social interactions) show higher and more consistent group cohesion (see lower variability in the error bars). However, compared to low values of r, high r also leads to lower group accuracy: a moderate accuracy drop in condition A (circle markers) and zero accuracy in condition B (triangle markers). This means that relying too heavily on social interactions makes the swarm unable to change its collective commitment and adapt to environmental changes (as shown by experiments in condition B, where the swarm started committed from the incorrect option). The swarm adapts only through infrequent social interactions (low values of r), however with low cohesion.

Figure 5 shows the accuracy and cohesion levels for the adaptive control parameter with both centralised and decentralised quorum computation. The results show that when the frequency of social interactions is linearly proportional to the quorum (L), with $N = 2$ the swarm reaches very high accuracy and cohesion in condition A (initialising the agents evenly committed to all the available options). However, in condition B, the swarm is unable to adapt to a change and modify the group opinion (high cohesion but low accuracy). On the contrary, with $N = 5$ options, the linear relation $r = Q$ allows the swarm

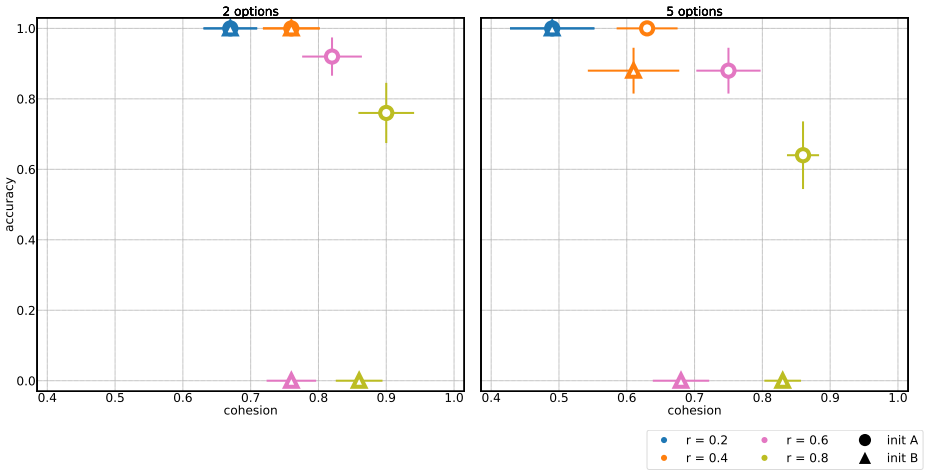


Fig. 4. Group cohesion and accuracy on the x and y axes for different probabilities of social interactions (control parameter r) and problem difficulty $\xi = 0.8$. The marker colour identifies the value of r while the marker shape identifies the initial condition A or B. The circles refer to condition A, where the agents are initialised evenly committed to all the available options, the triangles refer to condition B (adaptation to a dynamic environment), where the entire swarm is initialised committed to one of the inferior options. Each marker shows the accuracy and the median cohesion along with their standard deviation (vertical and horizontal bars). (left) Results for $N = 2$ options; (right) results for $N = 5$ options.

to adapt to changes consistently (100% accuracy) but with lower levels of cohesion (roughly 67%). On the other hand, by using the V-shaped function to link quorum and social interaction frequency (V), the swarm can more consistently adapt (even in the $N = 2$ options case) and obtain most of the time higher levels of accuracy and cohesion. Note, however, that the results from experiments using the V-shaped function have a higher cohesion variability (longer standard deviation bars) than experiments that used the linear function.

Figure 6 shows the comparison between the static control parameters' values, the linear function, and the V-shaped function. The adaptive computation of r, in particular using the V function, shows an improvement in both group accuracy and cohesion. Additionally, both L and V functions obtained more consistent results for different initial conditions and number of available options for both centralised and decentralised computation of the quorum.

To analyse further the swarm cohesion and ability to adapt to changes, depending on the values of r, we plot the average probability of one agent to change its opinion over time for the initial conditions A and B in Figures 7 and 8, respectively. The probability of changing opinion stays roughly constant over time for any constant value of r. On the contrary, with quorum-based r (both linear and V-shaped functions), the probability of an agent changing opinion varies over time. More precisely, Figure 7 shows that with $N = 2$ in condition A, all

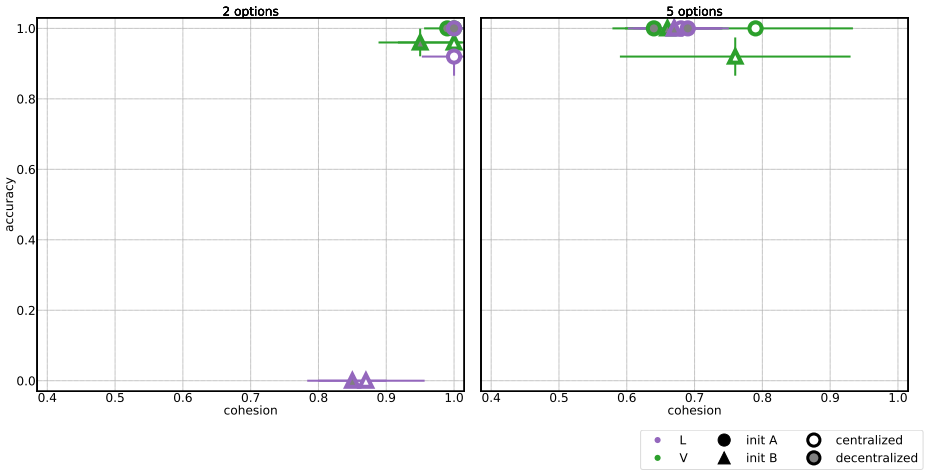


Fig. 5. Group accuracy and cohesion for quorum-dependent social interactions (control parameter r) and problem difficulty $\xi = 0.8$. The marker colour identifies different functions: purple markers are the results for the linear $r = Q$ function and green are for the V-shaped function $r = V(Q)$. The marker shape identifies the initial condition (see caption of Fig.4 for more details). Hollow markers refer to the centralised quorum computation, filled markers refer to the decentralised quorum computation. (left) Results for experiments with $N = 2$ options; (right) results for $N = 5$ options. Each marker shows the accuracy and median cohesion along with their standard deviation (vertical and horizontal bars).

tested quorum-based r functions (except the decentralised V-shaped configuration) decreased to zero, meaning that the swarm stabilised and agents stopped changing opinion after a fast transient. Differently, with $N = 5$ the probability stabilised over higher values, meaning that the agents kept changing their opinion, explaining the lower coherence values.

Figure 8 shows that with $N = 2$ the V-shaped function let to a slow increase of the probability of changing opinion before returning to a stable state in which the agents rarely change opinion. Differently, the linear function led to slower and more stable dynamics. Similarly to Figure 7, with $N = 5$, the probability of changing opinion stabilises to higher values than for $N = 2$ for both the linear and the V-shaped functions.

4 Other results

We tested a bunch of other things too:

Another way to select social information. We tested whether the observed results of improved accuracy-cohesion performance for adaptive social interactions hold in similar but different models of opinion dynamics. We replaced the Majority Rule with the Voter Model rule where each agent selects the opinion from a single random message from its buffer (rather than computing the

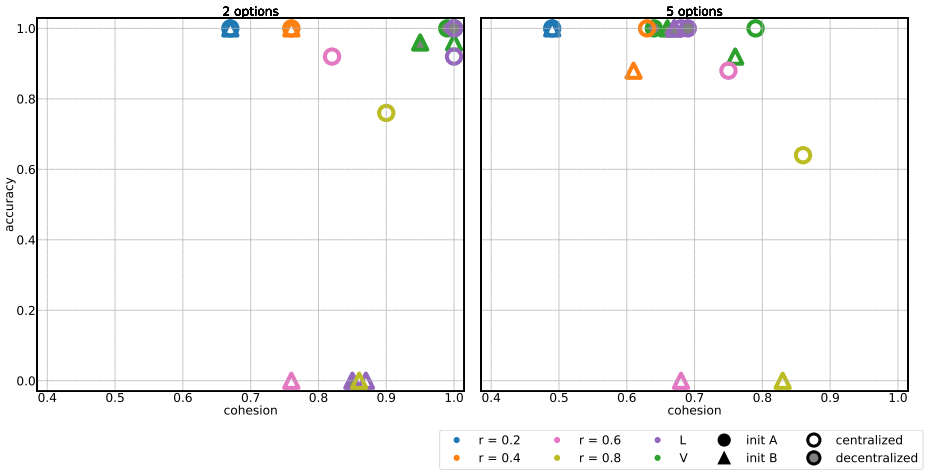


Fig. 6. The figure compares group accuracy and cohesion of swarm using static social interaction frequency (static r) and quorum-based r. For further details on the meaning of markers and colours see captions of Figures 4 and 5.

majority). However, preliminary results **did not** confirm that general trend and quorum-based social interactions did not improve the accuracy-cohesion trade-off. Figure 9 shows that using the Voter Model (dashed line markers) reduces the impact of computing r as a function of Q. The results show (nearly) maximum accuracy with both initial conditions (A and B) for both $N = 2$ and $N = 5$ cases, both with static r and quorum-based r. However, group cohesion does not improve and remains low in all tested conditions. Cohesion is particularly low when the number of available options N is larger than two. Comparing the $N = 2$ and $N = 5$ cases, cohesion decreases about 20% for any setting using the Voter Model.

Computing the quorum differently. We tested different functions to compute the quorum. In the results presented in Section 3, the quorum level computed by each agent is based on the size of the subpopulation selecting the same option of the agent. We also tested the quorum level as the ratio between the largest group of agents sharing any of the opinions (not necessarily the same as the agent) and the size of the swarm. The computation of the largest group ratio combined with the Majority Rule has different effects on the two tested functions, as shown in Figure 10. In the linear function case, there is a general improvement in group cohesion however, the swarm becomes unable to adapt with $N = 5$. Differently, the V-shaped function never reaches satisfactory levels of group cohesion with $N = 5$, which remains below 50% (roughly 45%) for both initial conditions. Using the "absolute" quorum computation did not vary the accuracy and cohesion of the Voter Model (not shown).

Different functions. Other than the linear and the V-shaped functions to compute r given Q, we also tested their inverse functions and their respective step

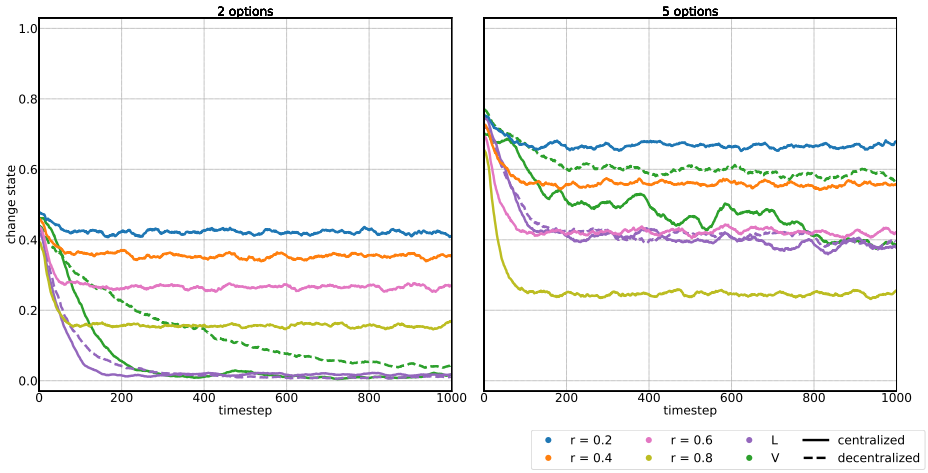


Fig. 7. Probability of agent opinion change over time with initial condition A. On the y-axis is shown the probability that one agent changes its opinion while on the x-axis is given simulation's duration. The blue, orange, pink and yellow lines refer to different values of the static control parameter. The green and the purple lines refer to the V-shaped and linear functions of r respectively. For the linear and V-shaped functions, solid lines refer to the centralised quorum computation and dashed lines to the decentralised quorum computation.

functions. The tested step functions are shown in Figure 11. In the step functions, we imposed that if $Q < 0.5$ then $r = \epsilon$ and if $Q \geq 0.5$ then $r = 1 - \epsilon$. However, compared to the results shown above, the outcome of these configurations was not satisfactory due to the low adaptability of the swarm (i.e., low performance in condition B).

Cross-Inhibition. Another considered model consists of replacing the Direct-Switch rule with the Cross-Inhibition rule, where receiving messages for different opinions leads to a transition to an uncommitted state. The state machine using the Cross-Inhibition rule is illustrated in Figure 12. The results obtained with the Cross-Inhibition model were a bit confusing and hard to analyse so we decided to focus on the Direct-Switch model instead.

5 Conclusions

In conclusion, we tested various functions to shape the control parameter r (social interaction frequency) through the quorum level for different state machine models and voting strategies. As shown in Section 3, quorum-based social interactions (i.e., changing the parameter r as a function of Q) can improve the collective performance and let the swarm adapt to changing environments. However, as indicated in Section 4, selecting the best function $r = f(Q)$ is not an easy task. Many of the tested configurations did not succeed and the system was

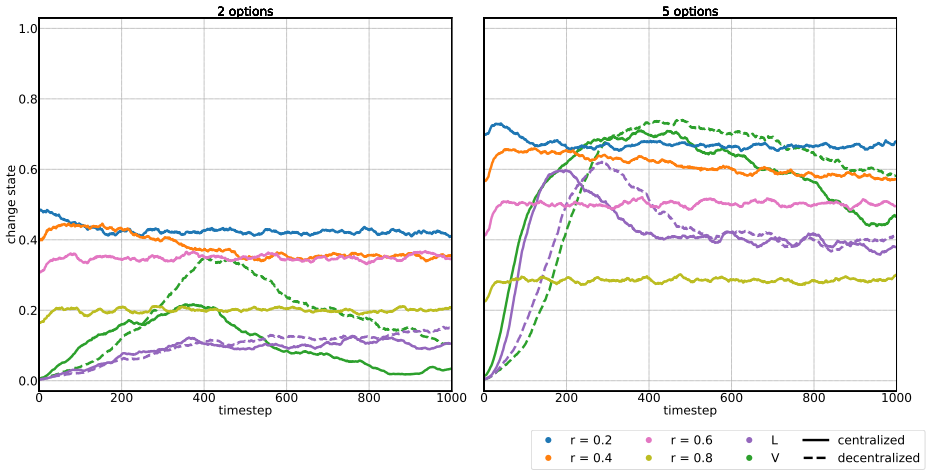


Fig. 8. Probability of agent opinion change over time with initial condition B. On the y-axis is shown the probability that one agent changes its opinion while on the x-axis is given simulation's duration. The blue, orange, pink and yellow lines refer to different values of the static control parameter. The green and the purple lines refer to the V-shaped and linear functions of r respectively. For the linear and V-shaped functions, solid lines are for the centralised quorum computation and dashed lines to the decentralised quorum computation.

not able to obtain acceptable accuracy and coherence levels with both initial conditions A and B. A rough interpretation of the results obtained with the implementation of the Direct-Switch model combined with a V-shaped function for r indicates that highly relying on social interaction is valuable in two situations: (i) when an agent has an opinion that is highly unpopular (i.e., it is not shared by the majority of agents, $Q \to 0$), supposedly enabling adaptation to environmental changes; and (ii) when the agent has a popular opinion (i.e., a large portion of the swarm has the same opinion, $Q \to 1$), supposedly increasing cohesion. Instead, it is preferable that the agent acts independently when an agent shares its opinion with an intermediate proportion of the population (i.e., $Q \approx 0.5$).

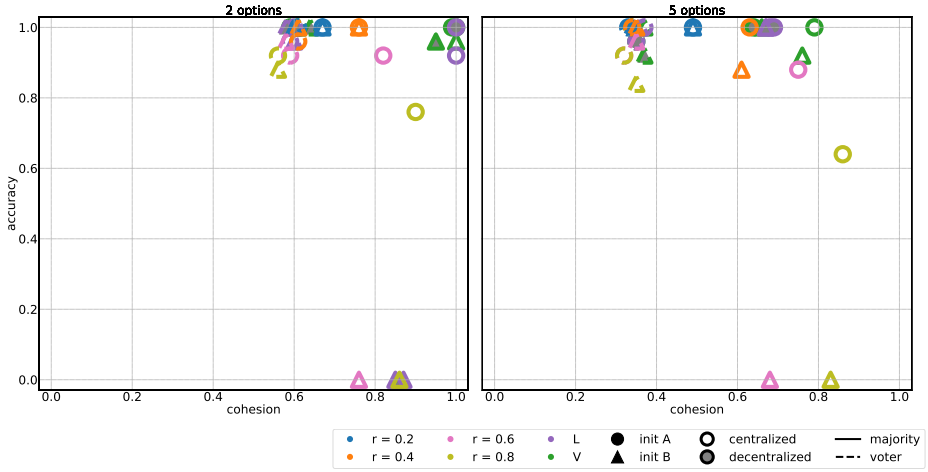


Fig. 9. The figure compares group accuracy and cohesion of swarm using static social interaction frequency (static r) and quorum-based r. Markers with solid borders refer to the algorithm implementing the Majority Rule, and markers with dashed borders refer to the algorithm implementing the Voter Model. For further details on the meaning of markers and colours see captions of Figures 4 and 5.

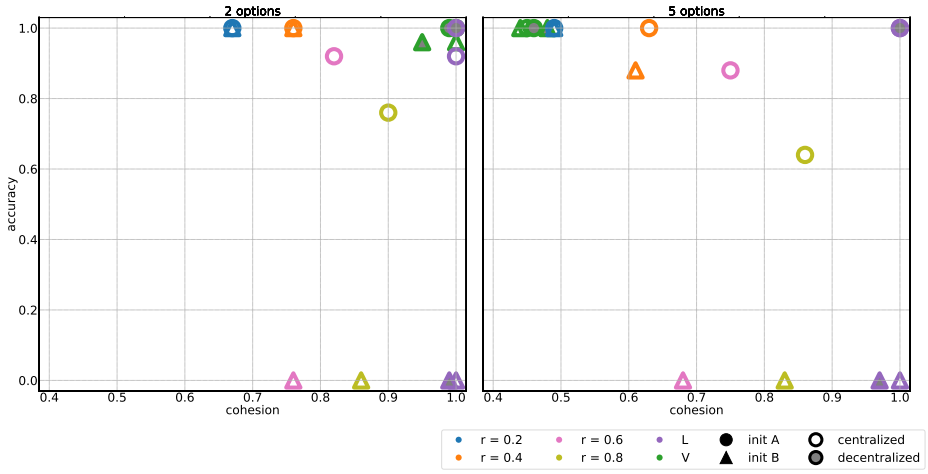


Fig. 10. The figure compares group accuracy and cohesion of swarm using static social interaction frequency (static r) and quorum-based r. Quorum computation is based on the largest group of agents sharing the same opinion. For further details on the meaning of markers and colours see captions of Figures 4 and 5.

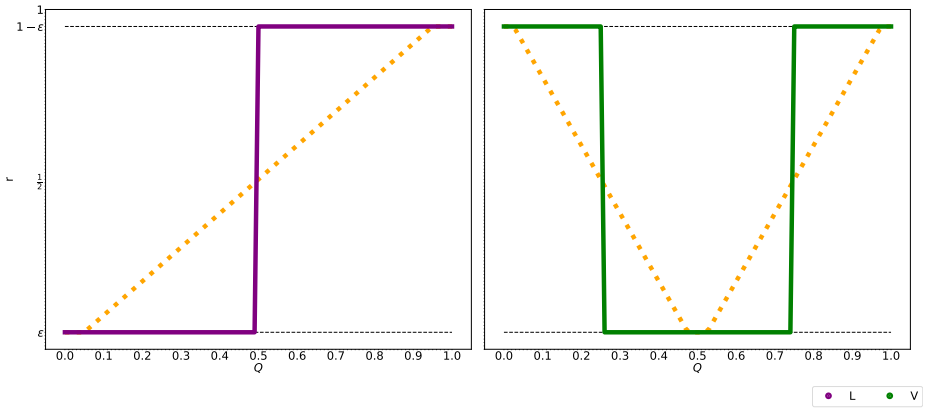


Fig. 11. Step function of L (left panel) and step function V (right panel) of the social interaction frequency r, depending on estimated quorum Q. The yellow dotted lines represent the non-step function, identical to the ones shown in Fig. 3.

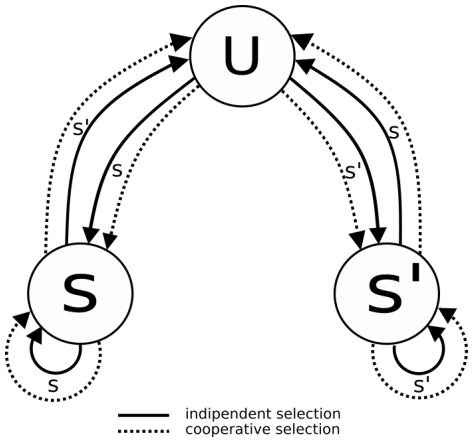


Fig. 12. Implemented Cross-Inhibition state machine model