Swarm Intelligence (2025) 19:295–316
https://doi.org/10.1007/s11721-025-00253-2

Evidence accumulation with adaptive weighting of social and personal information for collective perception

Gal Sajko[1,2] **· Jan Babič**[1]

Received: 15 November 2024 / Accepted: 27 August 2025 / Published online: 29 September 2025
© The Author(s) 2025

Abstract
This study advances collective perception in swarm robotics by introducing the Automated Swarm Opinion Diffusion Model (aSODM), which addresses the limitations of the original Swarm Opinion Diffusion Model (SODM). Both aSODM and SODM are evidence accumulation methods based on human-like decision-making processes. While SODM integrates social and personal information similarly to aSODM, its reliance on a manually predetermined social factor parameter limits its adaptability to diverse task difficulties. The Automated Swarm Opinion Diffusion Model eliminates this dependency by introducing an adaptive personal factor parameter, which automatically adjusts the weighting of personal and social information based on the information gathered about the environment. This automated approach improves robustness and reduces the dissemination of erroneous information in the early phases of a task. Comparative simulations against baseline methods (Voter Model and Majority Rule) demonstrate that aSODM enhances efficiency, particularly in tasks with higher difficulty levels. While aSODM outperforms SODM, its automated critical parameter selection makes it particularly well-suited for real-world applications where prior knowledge of the environment and task difficulty is lacking.

Keywords Collective perception · Decentralized decision-making · Swarm robotics · Evidence accumulation models · Social organisms

✉ Gal Sajko
gal.sajko@ijs.si

Jan Babič
jan.babic@ijs.si

[1] Department of Automatics, Biocybernetics and Robotics, Laboratory for Neuromechanics and Biorobotics, Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia

[2] Jožef Stefan International Postgraduate School, Jamova cesta 39, 1000 Ljubljana, Slovenia

1 Introduction

Swarm robotics methods (Hamann 2018) are typically inspired by simple social organisms such as ants, bees, and fish (Dorigo et al. 2020; Garnier et al. 2007). These methods seek to mimic the collective behavior of such species, leveraging the advantages of decentralized operation and applying them to robotic swarms. The objective is to define local rules for each agent in a swarm so that they collectively result in the desired behavior of the entire swarm.

Similarly, we considered human beings as another example of social organisms (Young 2008). In our previous work (Sajko and Babič 2024), we introduced the idea of using a decision-making method inspired by human group decision-making to solve a collective perception task in a robotic swarm (Valentini et al. 2016; Shan and Mostaghim 2020, 2021; Schranz et al. 2020; Bartashevich and Mostaghim 2019b). We employed the principles of the social drift diffusion model (Tump et al. 2020; Ratcliff and McKoon 2008) to develop the Swarm Opinion Drift Model (SODM), which falls under evidence accumulation methods (Boag et al. 2023). In SODM, each agent's decision is based on personal observations and information obtained from neighboring agents, referred to as social information. Together, these elements form an accumulated body of evidence that translates into a decision variable (Kira et al. 2015; Gold and Shadlen 2007), which represents the agent's current opinion. During decision-making, the agent accumulates evidence, and the decision variable drifts toward one of two decision boundaries (Myers et al. 2022; Thieu and Melnik 2023). Once the boundary is exceeded, the agent makes a decision in favor of one option or the other.

SODM incorporates a social factor parameter that amplifies the impact of social information. This parameter directly affects the trade-off between the speed and accuracy of consensus: a higher value increases speed but negatively impacts accuracy by facilitating the faster dissemination of incorrect opinions. A major drawback of SODM is that the social factor value must be manually determined before starting a task, which requires *a priori* knowledge about the task's difficulty–information that is usually unavailable. To address this limitation, we developed the original method, leading to a modified approach: aSODM, where "a" stands for "automated".

To overcome the limitations of a fixed social factor in the original SODM, we introduced a dynamic personal factor parameter in aSODM. This parameter enhances the influence of an agent's individual observations (personal information) during decision-making, particularly in the early stages of the process when agents lack sufficient reliable social input. The value of the personal factor is inversely related to the number of neighboring agents that have already formed a decision. At the beginning of the task, when agents have not yet crossed the decision threshold, the model assigns greater weight to personal information derived from local environmental observations. As more agents accumulate sufficient evidence and commit to a decision, the influence of the personal factor gradually diminishes. This shift allows social information to take on a more prominent role, accelerating consensus formation while mitigating the premature spread of incorrect opinions.

Fusing multiple pieces of evidence into a single representation is a well-established approach, notably used in the Dempster-Shafer Theory of Evidence (DST), (Dempster 1967; Shafer 1976). In DST, agents assign belief not only to individual hypotheses but to any subset of the hypothesis space–i.e., elements of the power set of hypotheses. Each agent can combine direct evidence (from its own observations) with indirect evidence (received

from neighboring agents). Using belief combination operators, such as Dempster's rule, these sources of evidence are fused, and the resulting belief masses are used to update the agent's belief function over the power set.

Despite its expressive power, DST suffers from several well-known limitations that hinder its practical use in distributed systems like swarm robotics. In particular, the computational complexity of operating over the power set grows exponentially with the number of hypotheses, making real-time application infeasible for large state spaces (Bauer 1997). Furthermore, DST is highly sensitive to conflicting evidence, often yielding unreliable or hard-to-interpret results when sources strongly disagree. Finally, the theory requires the a priori selection of a belief combination rule, but there is no universally accepted method, and different rules can lead to significantly different outcomes. These drawbacks–computational overhead, sensitivity to conflict, and ambiguity in rule selection–have motivated us to develop a more lightweight method that preserves the robustness of distributed evidence fusion while offering improved interpretability, and resilience to conflicting information in real-world swarm systems.

We evaluated aSODM by testing it on one of the most commonly used collective perception problems, as described in (Valentini 2017; Kaiser et al. 2023). Furthermore, we followed the methodology proposed by (Bartashevich and Mostaghim 2019a) to benchmark the effectiveness of collective perception methods. We conducted experiments across six different environmental patterns: random, band, block, star, stripe, and off-diagonal (Behrisch et al. 2017). This allowed us to assess how variations in cluster density (entropy E_c) and connectivity between clusters (Moran Index MI) affect aSODM's performance compared to three other methods: SODM, the Voter Model (Valentini et al. 2014), and Majority Rule (Valentini et al. 2015). Our comparison was based on consensus time and accuracy across 100 experimental trials. Since faster decisions typically lead to lower accuracy, we also calculated the inverse efficiency score (IES) (Statsenko et al. 2020), defined as the ratio of average consensus time to accuracy, where both factors were equally weighted.

This paper is structured as follows. Section 2 reviews related research. Section 3 describes the experimental setup and methods used. Section 4 presents our results, along with analyses and comparisons with other methods. Finally, Sect. 5 concludes the paper with a discussion on future research directions and the remaining limitations of our method.

2 Related work

To the best of our knowledge, no prior work applies a human-like collective decision-making mechanism to a swarm of robots to solve a collective perception task. However, several studies have employed evidence accumulation and information fusing principles to achieve decentralized collective decision-making.

Tump et al. (2020) propose a social drift diffusion model (social DDM) to explain how individuals integrate private and social information over time when making binary decisions. Their model extends classical evidence accumulation approaches by incorporating dynamically weighted social influence from others' choices, allowing for a nuanced analysis of collective decision dynamics. This approach, grounded in evidence accumulation theory, parallels our work in using accumulated signals to guide decision-making. Importantly, their method inspired our development of the Automated Swarm Opinion Diffusion Model

(aSODM), where we also modulate the influence of social and private information. Unlike (Tump et al. 2020), who focus on binary perceptual tasks in human subjects, we adapt the principle to distributed robotic systems with continuous belief updates and real-time consensus formation.

Ebert etal (2020) introduced a decentralized Bayesian method that enables the classification of the majority color, even when the difference between black and white tiles is only 4%. The method allows for adaptation of consensus time according to task difficulty through tunable parameters, but it requires some *a priori* knowledge.

Bartashevich and Mostaghim (2021) and Crosscombe et al. (2019) studied collective decision-making in environments where agents may receive conflicting local observations, particularly due to spatial correlations. Both approaches rely on Dempster-Shafer Theory (DST) to represent uncertainty and fuse two types of evidence: direct environmental observations and information shared by other agents. In their frameworks, evidence fusion is done in separate stages: each agent first updates its belief using newly collected local evidence from the environment, and then, during interactions, it combines this updated belief with that of a neighbor using one of several DST-based belief combination operators. This contrasts with our approach, where evidence from both sources is fused together with the prior belief in a single update step. Consequently, their method may involve higher computational overhead due to repeated fusion operations. Bartashevich and Mostaghim (2021) also evaluated eight different belief combination rules. Their simulations show that PCR5/6 performs particularly well in small swarms (e.g., up to $N = 20$ agents), achieving high decision accuracy and effective conflict resolution, especially in environments with strong spatial correlations and up to five decision options.

Bartashevich and Mostaghim (2019a) suggested extending the evaluation of collective perception methods to different tile patterns in 2D environments. Previously, methods were primarily tested on random patterns. The authors proposed to include additional patterns such as block, off-diagonal, star, stripe, band, band-stripe, bandwidth, and band-random-width. These patterns vary in entropy (cluster density) and the Moran Index (cluster connectivity). We incorporated six of these patterns to evaluate our method.

3 Methods

3.1 Experimental setup

The experiment was conducted in a simulation implemented using Python and the PyGame library. The task of the swarm was to determine the predominant tile color on a 2D surface composed of 20 × 20 tiles (2 m × 2 m), which could be either black or white. In addition to the random tile pattern (Valentini et al. 2016; Kaiser et al. 2023), five other patterns were included based on (Bartashevich and Mostaghim 2019a): band, block, off-diagonal, star, and stripe, as illustrated in Fig. 1. Table 1 summarizes the entropy and Moran Index values of each pattern (Bartashevich and Mostaghim 2019a). The block and off-diagonal patterns exhibit similar entropy and Moran Index values, while the differences among other patterns are more pronounced.

The task difficulty was defined by the ratio of black to white tiles, given as $\rho = \frac{N_{black}}{N_{white}}$. The closer N_{black} and N_{white} are to each other, the more difficult the task. The experiment

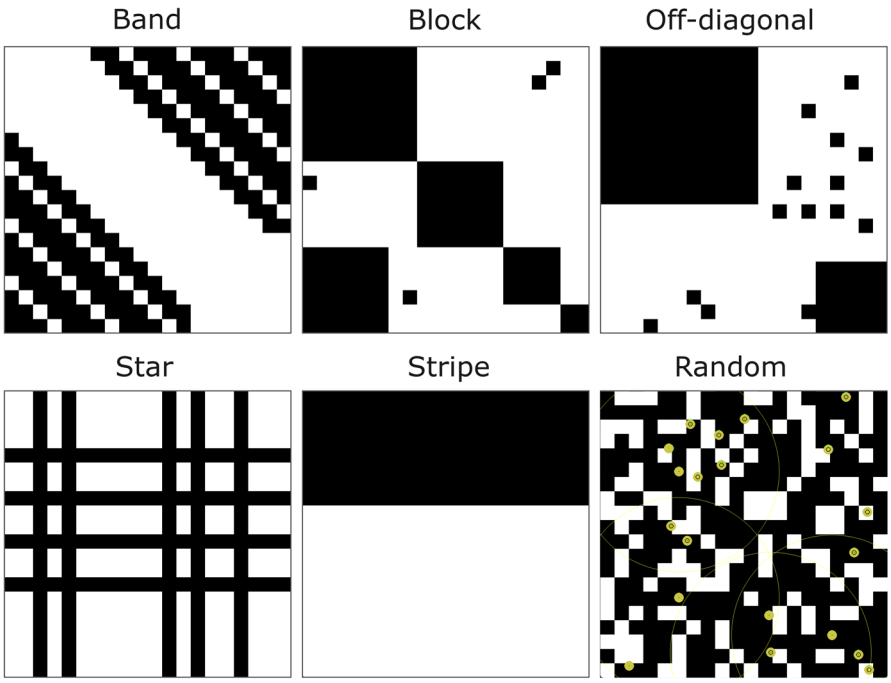

Fig. 1 Different patterns of the environment used for the evaluation of decision-making methods. Random pattern, shown in lower right image also includes agents in different states: small green circles representing the receiving state, large yellow circles the dissemination state with communication range and small black (possibly white) circles inside the agents representing their current opinions

Table 1 Entropy and Moran Index of each pattern

Pattern	E_c	MI
Random	0.51	0
Band	0.72	0.26
Star	0.8	0.4
Block	0.89	0.79
Off-diagonal	0.91	0.8
Stripe	0.98	0.82

was conducted at four difficulty levels, $\rho \in (0.25, 0.52, 0.67, 0.82)$, for each of the six patterns. The simulation frequency was set to 50 Hz.

The swarm consisted of 20 agents, each with a diameter of 6.8 cm, as shown in the lower right image of Fig. 1. Each agent transitioned between four possible states: (I) exploration, (II) receiving, (III) dissemination, and (IV) decision-making. The sequence of state transitions depended on the selected decision-making method.

The probabilistic finite-state machine design was based on (Valentini et al. 2016), in which agents move randomly within the space at all times. Movement was modeled as a series of linear paths in random directions for random durations (exponential distribution, mean time $\overline{t_{mov}} = 40$ s, velocity $v = 10$ cm/s), interspersed with stationary intervals simulating rotation (uniform distribution, $t_{rot} \in [0, 4.5]$ s, simulated angular velocity $\omega = 1.4$

rad/s). If an agent reached the environment boundary or encountered another agent, it stopped, turned, and continued moving in a new random direction.

In the exploration state, an agent gathered data about the environment by recording the time spent over each color (data was only recorded while moving, not while turning). The time in the exploration state, t_{exp}, followed an exponential distribution with a mean of $\overline{t_{exp}} = 10$ s. The receiving state lasted $t_{rec} = 8$ s, during which an agent received information from agents in the dissemination state within its communication radius. During dissemination, an agent broadcast its current opinion for a maximum duration of $T_{diss} = 15$ s, depending on its confidence level. This modeled positive feedback, where agents more confident in their opinions shared them for longer durations (Golman et al. 2015; Valentini 2017), increasing the likelihood of influencing other agents in the receiving state.

The final state of the cycle was the decision-making state, where agents updated their opinions based on the gathered information. Details on these methods are covered in the following subsections.

The experiments were evaluated by calculating the percentage of correct decisions across 100 trials, denoted as accuracy α, the average time required for the swarm to reach a consensus $\overline{T_c}$, and the inverse efficiency score (IES), defined as $IES = \frac{\overline{T_c}}{\alpha}$. Each trial was limited to 400 s.

3.2 Voter model and majority rule methods

Both the Voter Model (VM) and the Majority Rule (MR) model use a probabilistic finite-state machine to regulate positive feedback. In both models, agents cycle through the same sequence of states: (I) exploration, (II) dissemination, (III) receiving, and (IV) decision-making. The difference between the models lies in how agents process information received from others during the receiving state.

In the VM model, an agent randomly selects one of the collected opinions and aligns its own opinion accordingly. If the selected opinion matches the agent's current opinion, no change occurs. In contrast, in the MR model, the agent identifies the majority opinion of the collected opinions, including its own. If the majority opinion differs from the agent's current opinion, the agent updates its opinion accordingly. If there is a tie or if the majority opinion matches the agent's own, no change occurs. The VM model is considered more accurate and robust because selecting a single random opinion slows the spread of potentially incorrect opinions. Consequently, reaching consensus within the swarm typically takes longer in the VM model than in the MR model.

A common limitation of both methods is their inability to adapt to a dynamic environment once the swarm reaches consensus (Prasetyo et al. 2019). Since agents rely exclusively on the opinions of others when forming their own, once consensus is achieved, a single agent's opinion change cannot trigger further adaptation in others, making the swarm inflexible to environmental changes.

3.3 aSODM and comparison to SODM

aSODM stands for "Automated Swarm Diffusion Model," an extended version of the original SODM method described in our previous work (Sajko and Babič 2024). The primary

advantage of aSODM over SODM is its ability to automatically determine and adjust the balance between personal and social information while updating the decision variable in the decision-making state.

aSODM employs the same probabilistic finite-state machine as SODM, shown in Fig. 2. While similar to the finite-state machines used by VM and MR, aSODM differs in the sequence of state transitions and the modulation of positive feedback. The most significant difference, compared to VM and MR, lies in the type of information exchanged between agents and how each agent processes gathered information when updating its own opinion in the decision-making state.

During decision-making, agents collect evidence supporting one of the two options. This evidence consists of personal information, derived from an agent's own observations, and social information, obtained from neighboring agents during the receiving state. The agent's opinion is represented as the current value of the decision variable d_i. This variable can take values in the range $[-1\frac{2}{3} \cdot D, 1\frac{2}{3} \cdot D]$, where D represents the decision boundary value that d_i must exceed for an agent to make a decision. The boundary $+D$ represents a decision for black, while $-D$ represents a decision for white. To prevent infinite escalation and oscillation between making a decision and remaining undecided, the values of d_i are constrained to $1\frac{2}{3} \cdot D$.

All agents begin in the exploration state with an initial value of $d_i = 0$. This setup prevents premature assumptions about the environment that could bias toward an incorrect conclusion, increasing the likelihood of errors or delaying consensus. Even though early personal observations may be locally biased due to random movement and initial position, they still reflect actual environmental features–unlike random initialization of d_i, which

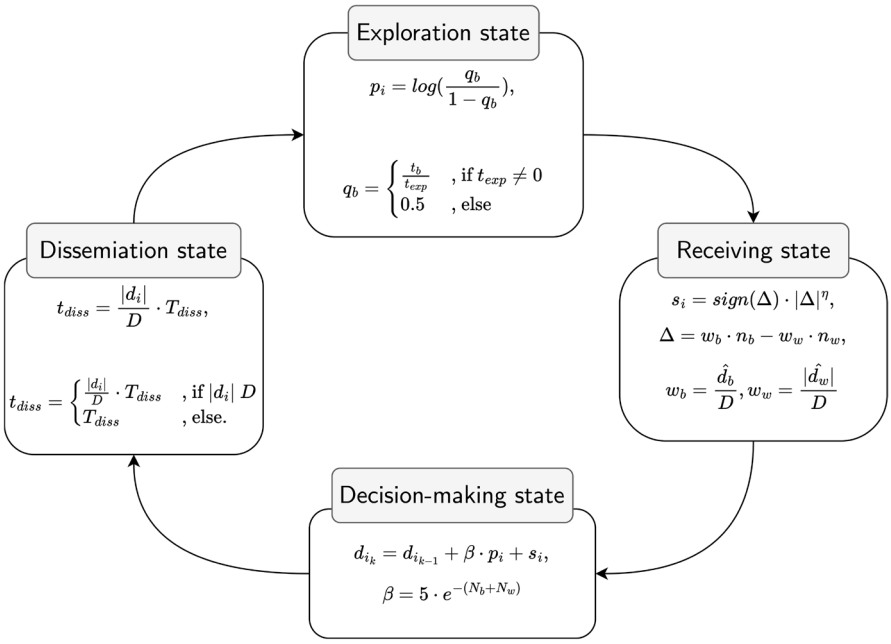

Fig. 2 Diagram of the finite-state machine used in the aSODM method. The loop begins with the exploration state

could introduce artificial bias unrelated to the true structure of the environment. At the end of the exploration state, each agent assesses personal information from gathered data, defined as:

$$p_i = log(\frac{q_b}{1-q_b}), \tag{1}$$

where the *i* is the agent's index, and q_b is defined as

$$q_b = \begin{cases} \frac{t_b}{t_{exp}} & \text{, if } t_{exp} \neq 0 \\ 0.5 & \text{, else} \end{cases} \tag{2}$$

with t_b representing the time spent by the agent over black tiles. If $t_b = \frac{1}{2} \cdot t_{exp}$, the agent has spent equal time over black and white tiles, yielding no insight into color prevalence, resulting in $p_i = 0$. If $t_b = 0$ or $t_b = t_{exp}$, p_i approaches $-\infty$ or $+\infty$, respectively. To prevent extreme values, q_b is constrained to $q_b \in [\epsilon, 1-\epsilon]$, with $\epsilon = 0.001$. Consequently, p_i is confined to $p_i \in [-6.907, 6.907]$.

Since the problem structure ensures the number of black tiles determines the number of white tiles, we can express:

$$P(x_i = \text{black}) = 1 - P(x_i = \text{white}), \tag{3}$$

where x_i represents an individual measurement by the agent, and $P(x_i)$ is the probability of obtaining a specific measurement based on the time spent over each color. This allows rewriting as:

$$\frac{P(x_i = \text{black})}{P(x_i = \text{white})} = \frac{\frac{t_b}{t_{exp}}}{t_{exp} - \frac{t_b}{t_{exp}}} = \frac{q_b}{1-q_b}. \tag{4}$$

In Eq. 1, this ratio is used in a logarithmic function (Marshall et al. 2017), offering two key advantages: (i) the function is symmetric around 0.5, meaning that if $q_b < 0.5$, p_i becomes negative, causing d_i to drift toward the lower decision boundary, and vice versa; (ii) at the end of the exploration phase, p_i will have a significantly larger absolute value for more definitive observations.

A significant improvement over SODM is introduced in the receiving and decision-making states. During the receiving state, each agent receives the current decision variable value exactly once from each agent in the dissemination state within its communication range. At the end of the receiving state, the agent calculates social information as:

$$s_i = sign(\Delta) \cdot |\Delta|^{\eta} \tag{5}$$

where

$$\Delta = w_b \cdot n_b - w_w \cdot n_w \tag{6}$$

with n_b (n_w) representing the number of received decision variables favoring black (white). Additionally, weights w_b and w_w are defined as:

$$\begin{aligned} w_b &= \frac{\hat{d_b}}{D} \\ w_w &= \frac{|\hat{d_w}|}{D} \end{aligned} \tag{7}$$

where $\hat{d_b}$ ($\hat{d_w}$) represents the average decision variable value favoring black (white). The parameter η amplifies the impact of majority opinions. Since it affects accuracy and consensus time, its value was empirically determined, as discussed in Sect. 4.4.

Instead of averaging collected opinions to compute social information, as done in SODM, we defined social information using the difference in counts of opinions favoring black (n_b) and white (n_w), aligning more closely with the social drift diffusion model (Tump et al. 2020). Additionally, n_b and n_w were weighted to reflect how close each group of neighbors is to making a decision. At the beginning of a task, weights w_b and w_w remain small, minimizing the influence of social information. As agents gather more evidence, opinions become more reliable, increasing the impact of social information. This approach mitigates the negative influence of erroneous information dissemination on swarm consensus accuracy.

When an agent enters the decision-making state, it updates its decision variable d_i based on personal and social information:

$$d_{i_k} = d_{i_{k-1}} + \beta \cdot p_i + s_i, \tag{8}$$

where the parameter β amplifies personal information and is defined as:

$$\beta = 5 \cdot e^{-(N_b + N_w)} \tag{9}$$

where N_b (N_w) stands for number of gathered decision variables that already form a decision (meaning the value of decision variable exceeds the decision boundary) in favor of black (white) color. This formulation ensures that at the beginning of the task, when few decisions have been made, personal information is weighted more heavily. As more agents reach a decision, β decreases, shifting emphasis to social information. This automatic adjustment eliminates the need for prior knowledge about the environment, which was required in SODM to set the optimal social factor parameter value.

At the same time, this approach automatically accounts for task difficulty. Our previous work demonstrated that stronger reliance on social information during the decision-making state negatively affects consensus accuracy, particularly in high-difficulty tasks. In such cases, the term $N_b + N_w$ in Eq. 8 remains lower for a longer duration, leading to prolonged reliance on personal information over social information. As previously mentioned, this mitigates the negative impact of disseminating incorrect opinions on consensus accuracy.

The probabilistic finite-state machine cycle concludes with the dissemination state. During this phase, an agent shares its current opinion, represented as the value of d_i, with other agents that are in the receiving state and within communication range. The duration of the dissemination state, t_{diss}, is defined as:

$$t_{diss} = \begin{cases} \frac{|d_i|}{D} \cdot T_{diss} & \text{, if } |d_i| < D \\ T_{diss} & \text{, else.} \end{cases} \tag{10}$$

The term $\frac{|d_i|}{D}$ in Eq. 10 modulates positive feedback. The closer the decision variable is to a decision boundary, the longer the agent will disseminate its opinion. This increases the likelihood that its opinion will be received and considered by other agents.

After the dissemination state, the cycle repeats as the agent returns to the exploration state. The process continues until all agents in the swarm reach consensus or the maximum allowed task duration is reached.

4 Results

In the following section, we present the results of the aSODM method and compare them with other methods in terms of accuracy, average consensus times, and inverse efficiency score (IES). Additionally, we analyze the impact of the parameters β and η on the decision-making process and examine how the communication radius size influences accuracy and consensus time. Furthermore, we compare the evolution of d_i throughout the decision-making process between the aSODM and SODM methods. For all comparisons, we set the SODM parameter γ to $\gamma = 0.2$, as this value yields the highest accuracy for the SODM method (see Sajko and Babič (2024) for more details).

4.1 Comparison of accuracy, average consensus times, and IES

The comparison of the results for all four methods is presented in Fig. 3. It includes accuracy α (top row), average consensus time $\overline{T_c}$ (middle row), and inverse efficiency score (IES) (bottom row) at all four levels of difficulty ρ for each of the six different environmental patterns. Higher values indicate better accuracy, whereas lower values are preferable for consensus time and IES.

In the accuracy graphs, we also include exit probability ϵ, which represents the percentage of trials in which methods reached a decision within the allowed time (regardless of correctness). The following subsections provide separate comparisons of accuracy and average consensus times. In Subsect. 4.1.3, we combine both metrics in the form of IES. Since there is a trade-off between accuracy and consensus time, we believe that IES provides the best overall evaluation of the methods' efficiency. Additionally, we include a table for each metric that illustrates how increasing difficulty affects the results of each method.

4.1.1 Accuracy

At difficulty $\rho = 0.25$, all methods achieved perfect accuracy, regardless of the pattern. As difficulty increases, the accuracy of all methods declines. For higher difficulty levels ($\rho \in (0.67, 0.82)$), the differences in accuracy among methods become more pronounced and are increasingly influenced by the environmental pattern.

In general, at higher difficulty levels, the VM performs worse than the other three methods, primarily due to its lower exit probability. Cases where a method fails to reach a decision within the allotted time are treated as incorrect final decisions. At $\rho = 0.82$, SODM

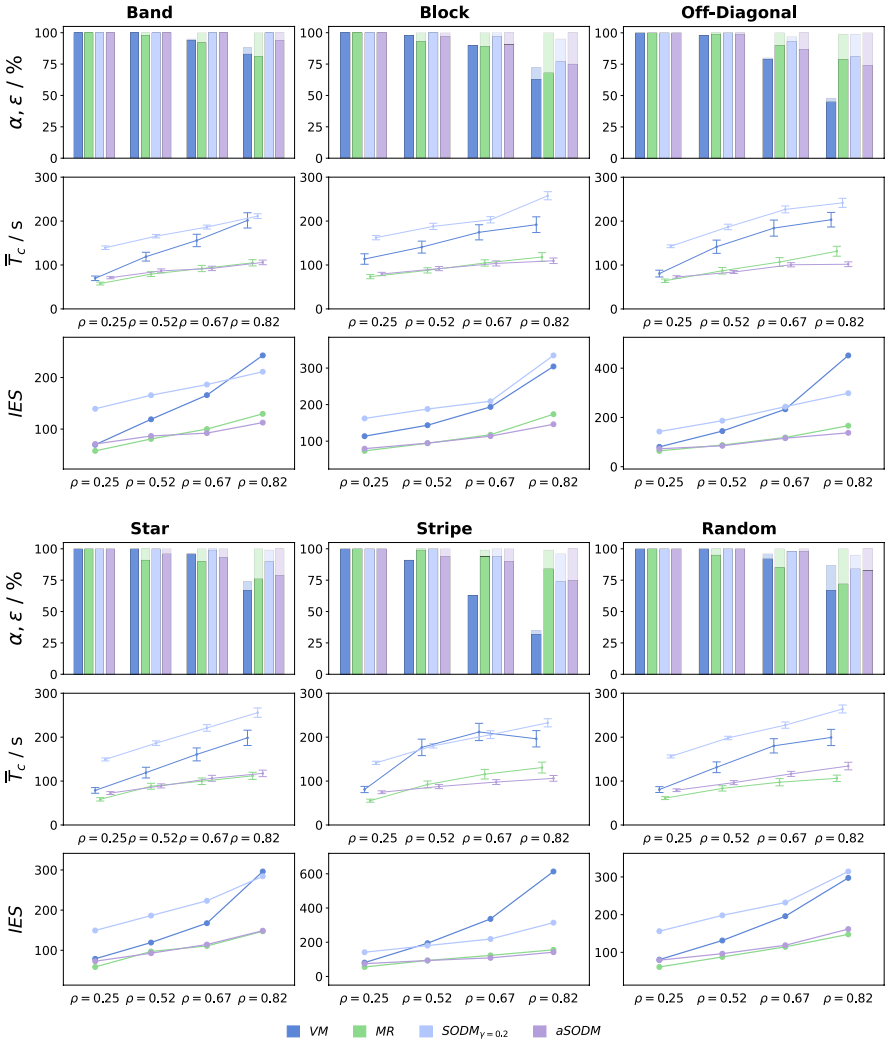

Fig. 3 Comparison of the performance of all four methods across six different environmental patterns. The title above each figure represents the pattern name. The first row of graphs shows accuracy and exit probability (bars with a more transparent color), the second row shows average consensus times with 95% confidence intervals presented with error bars, and the third row displays the inverse efficiency score (IES). The x-axis of each graph represents task difficulty

achieved the highest accuracy for most patterns, followed by aSODM and MR. However, for the stripe pattern, this order was reversed, with MR achieving the highest accuracy, followed by aSODM and SODM. Similarly, for the off-diagonal pattern, MR recorded the second-best accuracy, followed by aSODM.

The exit probabilities for aSODM and MR remain at their maximum levels across all difficulty levels. Meanwhile, the exit probability of SODM slightly decreases with increasing difficulty (except for the band pattern), although not as significantly as with VM.

Table 2 presents the difference between the best and worst recorded accuracy of each method. In all cases, the best accuracies were recorded at $\rho = 0.25$ and the worst at $\rho = 0.82$. This metric illustrates how increasing difficulty impacts the accuracy of each method.

We observe that the accuracy of the SODM method is the least affected by increasing difficulty levels for all patterns except the stripe pattern, where the MR method experiences the smallest accuracy drop. In general, the smallest accuracy drop across all methods occurs with the band pattern, except for the MR method, which experiences the smallest drop with the stripe pattern.

Conversely, the patterns associated with the largest accuracy drops vary by method. The largest accuracy drop for the VM and SODM methods is recorded with the stripe pattern, for the MR method with the block pattern, and for the aSODM method with the off-diagonal pattern (with the second-largest drop occurring in the stripe pattern, differing by only 1%).

The upper graph, shown in Fig. 4 illustrates how entropy and the Moran Index of patterns influence the accuracy of decision-making methods at difficulty $\rho = 0.82$. We observe that all methods except MR achieve the highest accuracy with the band pattern (which has the second-lowest values of E_c and *MI*). As E_c and *MI* increase, accuracy declines across all methods, particularly for the VM method. An exception is the MR method, whose accuracy increases for patterns with higher E_c and *MI* (off-diagonal and stripe).

4.1.2 Consensus times

A common characteristic among all four methods is that average consensus times are shorter at lower difficulty levels and increase as difficulty rises. Generally, the VM and SODM methods are slower in reaching consensus compared to MR and aSODM, with SODM being slower than VM in most cases. At $\rho = 0.82$, the aSODM method achieved the shortest average consensus times across all patterns, except for the star and random patterns, where MR had slightly shorter times.

Table 3 illustrates how increasing difficulty affects the average consensus times of each method. The metric is expressed as the percentage increase from the shortest average consensus time (recorded at $\rho = 0.25$) to the longest (recorded at $\rho = 0.82$).

Regarding average consensus time, the aSODM method is the least affected by increasing difficulty for all patterns except random, where SODM shows the smallest increase, with only a 0.2% difference between aSODM and SODM. In general, the smallest increase in average consensus times between $\rho = 0.25$ and $\rho = 0.82$ occurred in the block pattern, except for the SODM method, where the smallest increase was recorded with the band pattern.

Table 2 Differences between the best and worst recorded accuracies of all methods across the six patterns

	$\Delta_{\alpha_{VM}}$/%	$\Delta_{\alpha_{MR}}$/%	$\Delta_{\alpha_{SODM}}$/%	$\Delta_{\alpha_{aSODM}}$/%
Random	-33	-28	**-16**	-17
Band	-17	-19	**0**	-6
Star	-33	-24	**-10**	-21
Block	-37	-32	**-23**	-25
Off-diagonal	-55	-21	**-19**	-26
Stripe	-68	**-16**	-26	-25

Cells with a gray background indicate the method that was least affected by increasing difficulty for each pattern

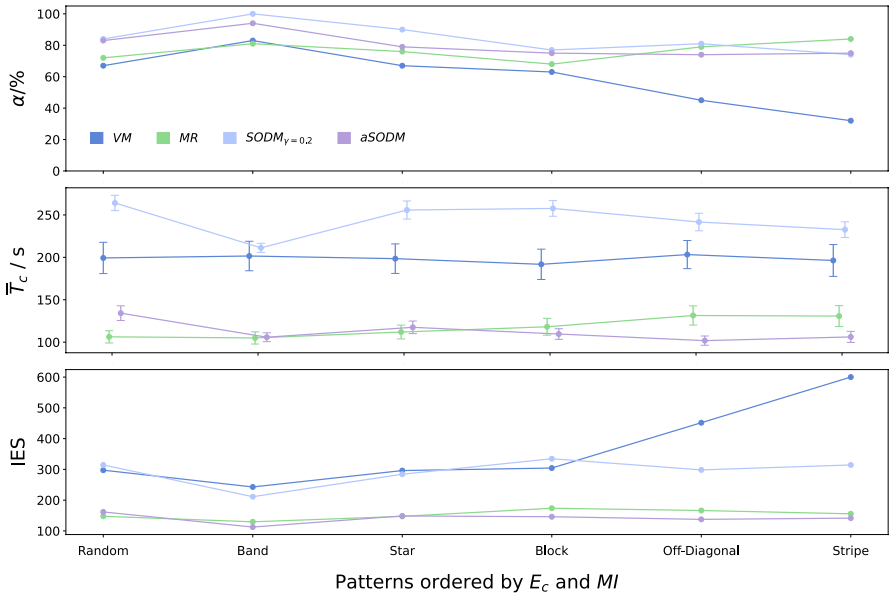

Fig. 4 Upper: accuracies of the methods across all patterns. Middle: average consensus times of the methods across all patterns. Error bars representing a 95% confidence interval. Lower: IES of the methods across all patterns. The x-axis orders patterns by increasing values of entropy and Moran Index. All results are recorded at $\rho = 0.82$

Table 3 Ratios between the longest and shortest average consensus times for each method across all six patterns

Cells with a gray background indicate the method least affected by increasing difficulty for each pattern

	$\Delta_{t_{VM}}$/%	$\Delta_{t_{MR}}$/%	$\Delta_{t_{SODM}}$/%	$\Delta_{t_{aSODM}}$/%
Random	247.0	173.6	**169.2**	169.4
Band	289.3	181.8	151.6	**149.0**
Star	252.1	191.6	171.4	**161.6**
Block	168.8	160.9	158.9	**137.9**
Off-diagonal	252.4	204.6	169.3	**139.9**
Stripe	242.9	236.2	164.5	**142.4**

The middle graph in Fig. 4 illustrates how increasing E_c and MI affect the average consensus times for each method. The results show that average consensus times remain relatively stable across patterns for all methods except SODM, which experiences a noticeable drop with the band pattern, followed by an increase with the star pattern. As E_c and MI increase further, SODM's consensus times slightly decrease, similar to aSODM. The MR and aSODM methods consistently achieve shorter consensus times than VM and SODM, with MR displaying a slight increase in consensus time as E_c and MI rise.

4.1.3 Inverse efficiency score

Inverse efficiency score (IES) is defined as $IES = \frac{\overline{T_c}}{\alpha}$. We use it to evaluate the overall efficiency of a decision-making method, where the goal is to minimize consensus time while maximizing accuracy. A lower IES value indicates greater efficiency.

In general, the IES of the aSODM and MR methods is better than that of the VM and SODM methods, with the difference becoming more pronounced as difficulty increases. At $\rho = 0.25$, VM achieves a similar IES to aSODM and MR for most patterns, except for the block pattern, where VM's IES score is already higher at $\rho = 0.25$. The IES of the SODM method is consistently worse than that of aSODM and MR across all cases.

Table 4 illustrates how increasing difficulty affects the IES of each method, expressed as the ratio between the worst and best recorded IES scores. The best IES scores were recorded at $\rho = 0.25$, while the worst occurred at $\rho = 0.82$. We observe that the SODM and aSODM methods are less affected by increasing difficulty compared to the VM and MR methods. Specifically, SODM is less affected in the band, star, and random patterns (patterns with lower E_c and *MI*), whereas aSODM is less affected in the block, off-diagonal, and stripe patterns (patterns with higher E_c and *MI*).

The lower graph in Fig. 4 illustrates how E_c and *MI* influence the IES scores of the methods. The IES scores of MR and aSODM remain relatively stable across all patterns, whereas the IES scores of VM and SODM increase with rising E_c and *MI*, particularly beyond the band pattern. The increase in IES for VM is especially pronounced in patterns with the highest E_c and *MI* (off-diagonal and stripe). Overall, all methods achieved their best IES score with the band pattern, indicating that this pattern represents the easiest task.

4.2 Comparison of d_i values between aSODM and SODM methods

The decision-making processes of both the aSODM and SODM methods are based on the accumulation of evidence, which is translated into the decision variable of each agent in the swarm. In this subsection, we compare the evolution of decision variables between the aSODM and SODM methods. The experiment was conducted in an environment with a random pattern at all four levels of difficulty. The results are shown in Fig. 5.

The key difference observed in the graphs is that the aSODM method requires fewer decision-making cycles to reach consensus compared to the SODM method, making it sig-

Table 4 Ratios between the best and worst IES scores for each method across all six patterns

	$\Delta_{IES_{VM}}$/%	$\Delta_{IES_{MR}}$/%	$\Delta_{IES_{SODM}}$/%	$\Delta_{IES_{aSODM}}$/%
Random	368.7	241.1	**201.5**	204.1
Band	348.6	224.4	**151.6**	158.5
Star	376.2	252.1	**190.4**	204.6
Block	268.0	236.6	206.3	**183.9**
Off-diagonal	259.0	204.6	209.0	**189.0**
Stripe	759.1	281.1	222.2	**189.9**

Cells with a gray background indicate the method least affected by increasing difficulty for each pattern

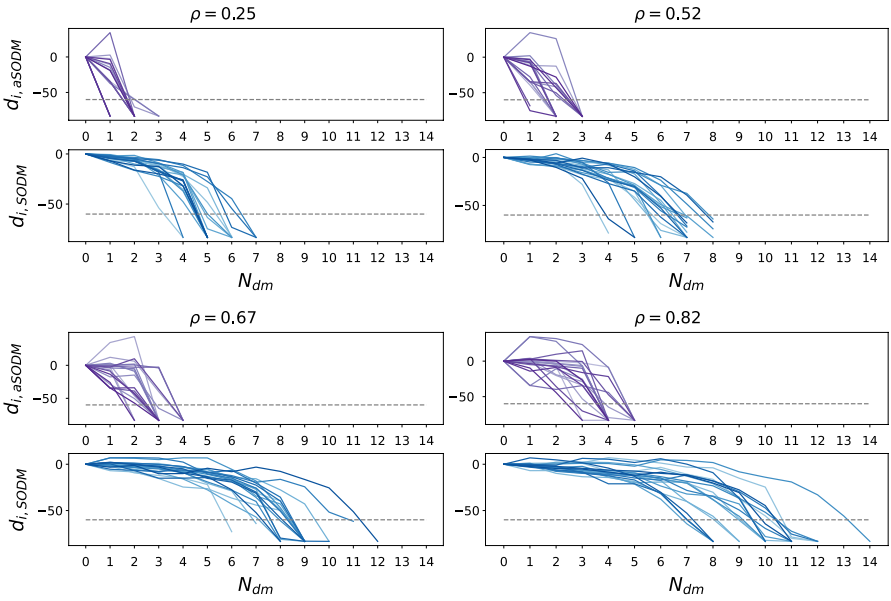

Fig. 5 Comparison of d_i values of the aSODM and SODM methods at all four difficulty levels, with the majority color being white. Purple graphs represent the aSODM method, and blue represents the SODM method. Grey dotted lines indicate the decision boundary in favor of the white color. The y-axis shows the decision variable values, while the x-axis represents the number of decision-making cycles

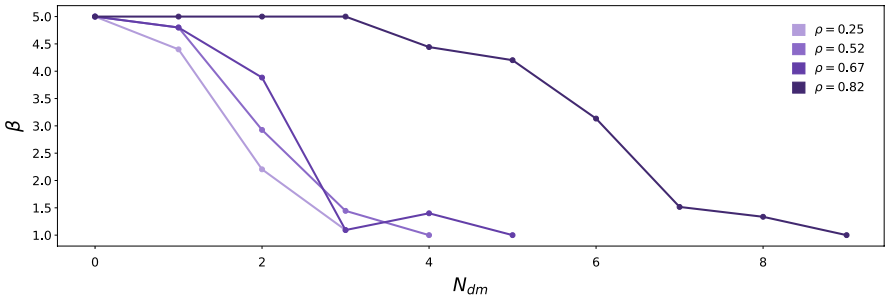

Fig. 6 Average β values of all 20 agents across all four levels of difficulty

nificantly faster. Additionally, in the early stages, the decision variable drifts in the aSODM method appear less monotonic and less directly oriented toward the decision boundary. However, they converge more rapidly in the later stages, once agents have gathered sufficient evidence from the environment to form reliable opinions. This behavior is driven by the β parameter, which ensures that agents initially prioritize personal over social information. This mitigates the risk of early cascades of unreliable opinions, a phenomenon commonly seen with the SODM method. The influence of the β parameter will be examined in more detail in the following subsection.

4.3 Effect of β parameter

The β parameter (Eq. 9) is used to amplify the personal component of information, as described in Chapter 3.3, thereby influencing the ratio between personal and social information. Figure 6 presents the evolution of the average β values of all agents throughout the decision-making process, recorded at all difficulty levels.

By definition, the value of the β parameter for an agent starts to decrease once a sufficient number of other agents (with which the given agent has communicated) have already formed a decision, meaning their d_i values have exceeded the decision boundary. Figure 6 shows that at higher difficulty levels, the β parameter remains at its maximum value (5) for a longer duration, as agents require more time to collect sufficient evidence to make a decision. This results in agents placing greater weight on personal information when calculating evidence for an extended period.

The higher the task difficulty, the more inconsistent the observations made by the agents, increasing the likelihood that these observations lead to incorrect opinions. By employing the β parameter, we mitigate the negative impact that these unreliable opinions can have on the decision-making process of an individual agent.

4.4 Effect of η parameter

The η parameter is used as an exponent in Eq. 5 during the receiving state while calculating social information. It amplifies each additional collected opinion that corresponds to the majority opinion. By doing so, agents assign more weight to social information derived from a larger number of unanimous collected opinions. The rationale behind this is that if

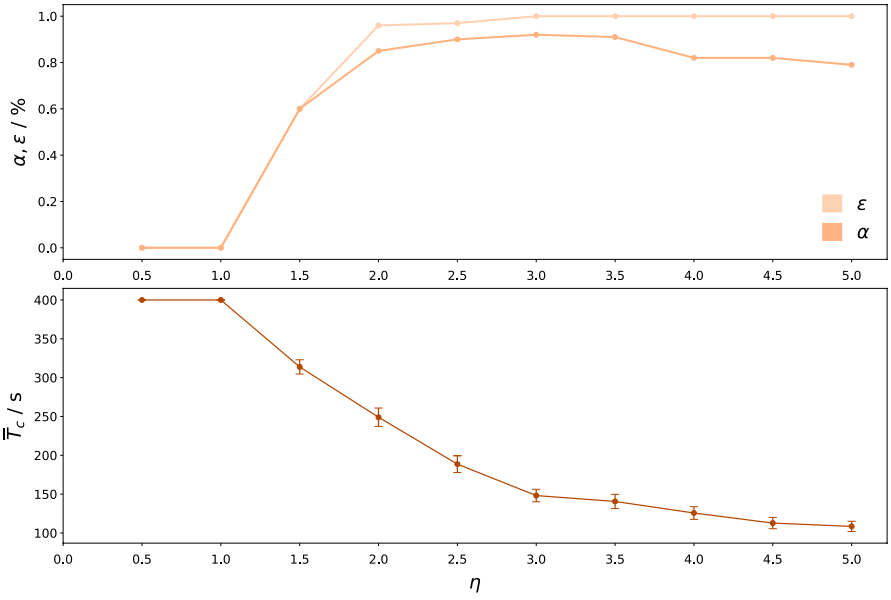

Fig. 7 Accuracy and exit probability (upper graph) and average consensus time with 95% confidence intervals (lower graph) as functions of the η parameter. The x-axes show different values of η

a greater number of agents disseminate the same opinion, it is more likely to be reliable, as the probability of all of them being mistaken decreases.

The value of the η parameter was determined experimentally in an environment with a random pattern at a task difficulty of $\rho = 0.82$. We conducted the experiment for η values ranging from 0.5 to 5.0 in increments of 0.5. The results, in terms of accuracy, exit probability, and average consensus time, are presented in Fig. 7.

For values of $\eta < 1.0$, the exit probability is 0, meaning the swarm cannot reach a consensus. As η increases, accuracy also increases, peaking at $\eta = 3.5$, beyond which further increases in η start to negatively affect accuracy. The best accuracies are observed for $\eta \in (2.5, 3.0, 3.5)$, and exit probabilities reach their maximum at $\eta \geq 3.0$. As η increases, the influence of social information relative to personal information grows during evidence accumulation. However, when $\eta \geq 4.0$, accuracy declines because early decisions by individual agents trigger cascades that lead the swarm to reach consensus rapidly but unreliably. At high task difficulty ($\rho = 0.82$), decisions formed in the early stages tend to be less reliable and more random due to environmental factors.

Figure 8 illustrates how η affects the evolution of averaged decision variables for the entire swarm. For lower values ($\eta \leq 1.5$), the drift is insufficient for decision variables to exceed the decision boundary, preventing consensus. As η increases, the drift becomes stronger and more directed. For $\eta \geq 2.0$, the drift is strong enough for the average decision variables to exceed the decision boundary. The drift strength continues to increase with higher values of η.

4.5 Effect of the size of the communication radius

In this subsection, we analyze the impact of the communication radius r_{comm} of agents on the efficiency of decision-making methods. Within the simulation, the communication radius can be easily modified to achieve optimal results. However, in real robotic applications, the communication range is hardware-dependent. A larger communication range allows the swarm to cover a broader area while maintaining effective communication or enables robots to collect opinions from agents further away. On the other hand, a larger

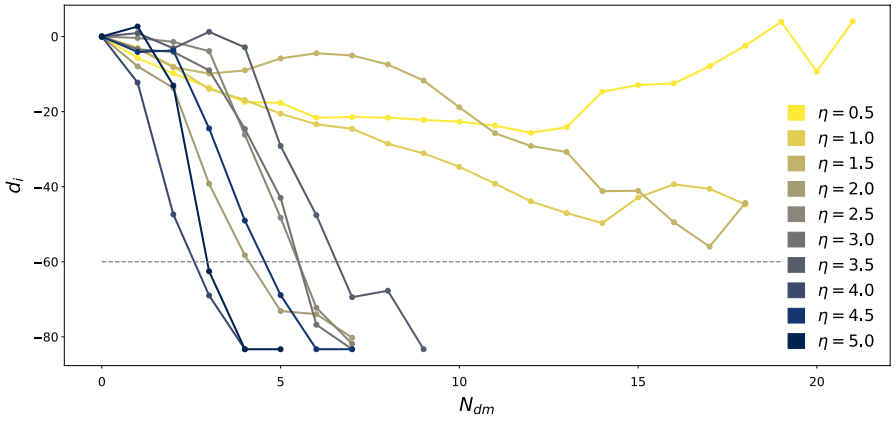

Fig. 8 Average d_i values at different values of the η parameter. The grey dotted line represents the decision boundary. The x-axis shows the number of decision-making cycles

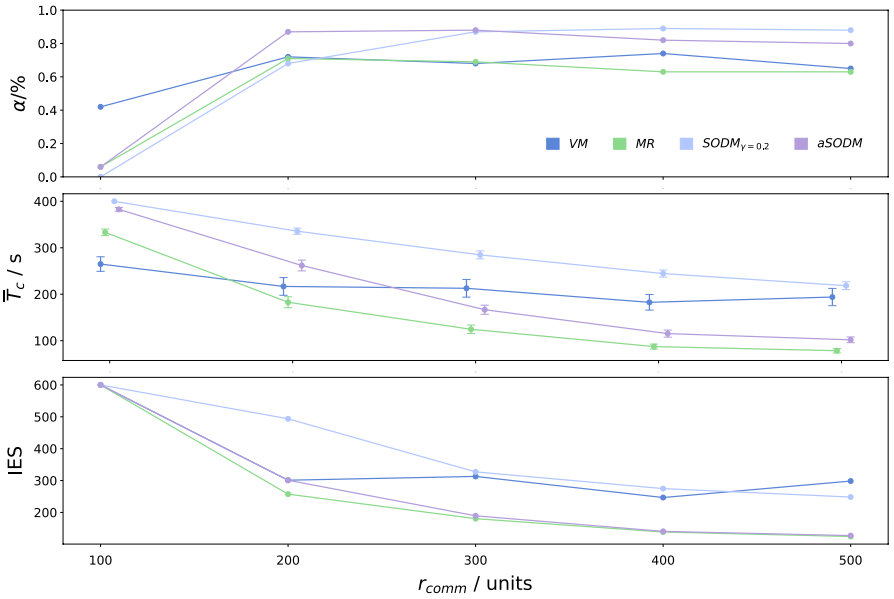

Fig. 9 Accuracies (upper graph), average consensus times with 95% confidence interval (middle graph), and IES values (lower graph) of all methods at different communication radius sizes (shown on the x-axis of each graph). Maximum IES values recorded at $r_{comm} = 100$ units were limited to 600 for better readability

communication range usually results in higher power consumption, leading to shorter battery life. This issue is particularly significant for drones, where power consumption is high relative to battery capacity.

Figure 9 presents the effect of the communication radius on accuracy, average consensus time, and IES across all methods. The experiment was conducted in an environment with a random pattern at a task difficulty of $\rho = 0.82$, with communication radius values ranging from 100 to 500 units in increments of 100 units. This corresponds to approximately 6 to 30 times the size of an agent's radius (set to 17 units). With these values, the communication area covers approximately 3% to 78% of the total environment area (1000 x 1000 units).

At $r_{comm} = 100$ units, all methods perform significantly worse compared to the results from Subsect. 4.1. This is because, during the receiving state, agents do not receive sufficient information from neighbors. Consequently, they must rely mostly on their own observations (using aSODM or SODM), which, especially in high-difficulty tasks, are not reliable enough. When r_{comm} increases to 200 units, all accuracies improve drastically. The aSODM and MR methods achieve their highest accuracies at $r_{comm} = 200$ units, which then slightly decrease with further increases in radius. The SODM method achieves its best accuracy at $r_{comm} = 300$ units and then remains relatively stable with increasing radius. For $r_{comm} \geq 200$, the accuracy of the VM method oscillates around $\alpha_{VM} = 0.7$.

Average consensus times decrease monotonically with increasing radius, except for the VM method, which experiences an increase at $r_{comm} = 500$ units. In general, a larger radius results in shorter consensus times. This is expected, as a greater communication range allows an agent to receive more information from its neighbors, enabling the swarm to reach consensus faster.

The IES score follows a similar pattern to average consensus time. For $r_{comm} \geq 200$ units, accuracy remains relatively constant, while average consensus time continues to decrease.

5 Discussion

This paper extends our previous work published at the ANTS2024 conference (Sajko and Babič 2024), where we introduced the SODM decision-making method based on the Social Drift Diffusion Model (Tump et al. 2020). The primary focus of this extension was to address the major drawbacks of SODM, particularly the need for *a priori* knowledge about the environment to select the γ parameter. This led to the development of aSODM, which eliminates the need for manual selection of γ by introducing the automatically calculated β parameter. This parameter amplifies the influence of personal information and is dynamically adjusted during each decision-making state based on opinions received from other agents. Specifically, it considers the number of opinions that have already formed a decision, meaning their decision variables exceed one of the two decision boundaries.

The β parameter is designed to assign greater weight to personal information at the start of the task when agents lack sufficient knowledge of the environment. As agents accumulate evidence and their opinions become more reliable, the β parameter gradually decreases, allowing social information to play a more significant role in updating their opinions. This approach mitigates the negative effects of incorrect information dissemination by preventing premature social cascades and enabling agents to rely more on personal observations in the early phase while facilitating faster consensus formation in the later phase.

In Sect. 4, we presented the results of aSODM in terms of accuracy, average consensus times, and IES, comparing them to SODM, the Voter Model, and Majority Rule, which served as baseline methods. The experiments were conducted on six different environmental tile patterns. Each trial was limited to 400 s, and a final decision was considered correct only if the swarm both correctly identified the majority color and reached consensus within the time limit.

The environmental patterns, ordered by increasing entropy and Moran Index values, were: random, band, star, block, off-diagonal, and stripe. The aSODM, SODM, and VM methods demonstrated higher accuracy in environments with lower entropy and Moran Index values (random, band, and star). Conversely, for these methods, the most challenging patterns to solve were those with the highest entropy and Moran Index values (off-diagonal and stripe). Interestingly, the MR method exhibited the opposite trend, achieving better accuracy with off-diagonal and stripe patterns compared to those with lower entropy and Moran Index. In terms of overall efficiency, as measured by IES, all methods performed best with the band pattern, which had the second-lowest entropy and Moran Index values. Generally, the aSODM method was the most efficient across all patterns, except in the random pattern, where MR performed best.[1]

[1] Although the aSODM method achieves a higher accuracy (83%), it requires more time to reach consensus (134.2 s). In contrast, the MR method, despite its lower accuracy (72%), completes consensus faster (106.29 s). As a result, MR achieves a higher IES score, showing that in this case, faster convergence outweighed the accuracy advantage of aSODM.

Additionally, we analyzed the effect of communication radius on decision-making efficiency. We initially hypothesized that a larger communication radius would improve overall efficiency by enabling agents to receive input from a greater number of neighbors, covering a larger area of the environment. However, as shown in Sect. 4.5, this was not necessarily the case. While IES and average consensus times decreased monotonically with increasing radius, this was primarily due to the reduced consensus time rather than improved accuracy. A larger communication radius allowed the swarm to reach consensus more quickly, but accuracy remained largely unchanged beyond a certain threshold.

For the smallest tested radius, accuracy was the lowest, primarily due to a reduced exit probability. However, once $r_{comm} = 200$ units was reached, most methods (except SODM) had already achieved their best or near-optimal accuracy. Increasing the radius beyond this point resulted in only minor fluctuations in accuracy, and in some cases (such as aSODM and MR), accuracy slightly decreased. The SODM method required a larger communication radius ($r_{comm} = 300$ units) to achieve its best accuracy.

While the SODM algorithm with $\gamma = 0.2$ still slightly outperforms aSODM in terms of accuracy, aSODM significantly improves consensus time and overall efficiency while eliminating the need for *a priori* knowledge about the environment, which was necessary for optimal parameter selection in SODM.

Currently, aSODM is designed to solve collective perception tasks requiring a decision between two options, formulated as a best-of-*n* problem with $n = 2$ in a static environment (where environmental properties, such as the predominant color, do not change over time). Future work will focus on expanding this capability to handle cases where $n > 2$ (Roxin 2019), making the method more applicable to real-world scenarios. Additionally, we aim to develop mechanisms that enable the swarm to revise its consensus dynamically, allowing adaptation to environmental changes. Another direction for future work is to compare our method with other, more closely related approaches to evidence accumulation and information fusion, such as those proposed by Bartashevich and Mostaghim (2021) and Crosscombe et al. (2019). Such comparisons would provide deeper insight into the efficiency of our method. Finally, we will look forward to eliminating the need for the finite-state machine that currently dictates each agent's operation. Removing this constraint could significantly reduce consensus time, enhancing swarm efficiency and making the method more suitable for real-world applications where factors such as battery life are crucial, particularly in robotics platforms like drones.

Author contributions G.S. wrote the main manuscript text, developed the described method, implement the software for the experiment and prepared all figures in text. All authors reviewed the manuscript.

Funding This work was partially supported by the Slovenian Research Agency (research core funding no. P2-0076).

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no Conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as

you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Bartashevich, P., & Mostaghim, S. (2019a). Benchmarking collective perception: New task difficulty metrics for collective decision-making. *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)* (Vol. 11804 LNAI, pp. 699–711). Springer Verlag.

Bartashevich, P., & Mostaghim, S. (2019b). Ising model as a switch voting mechanism in collective perception. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics),11805*, 617–629.

Bartashevich, P., & Mostaghim, S. (2021). Multi-featured collective perception with evidence theory: Tackling spatial correlations. *Swarm Intelligence, 15*, 83–110. https://doi.org/10.1007/s11721-021-00192-8

Bauer, M. (1997). Approximation algorithms and decision making in the Dempster-Shafer theory of evidence—An empirical study. *International Journal of Approximate Reasoning, 17*, 217–237.

Behrisch, M., Bach, B., Hund, M., Delz, M., Rüden, L. V., Fekete, J. D., & Schreck, T. (2017). Magnostics: Image-based search of interesting matrix views for guided network exploration. *IEEE Transactions on Visualization and Computer Graphics, 23*, 31–40. https://doi.org/10.1109/TVCG.2016.2598467

Boag, R. J., Strickland, L., Heathcote, A., Neal, A., Palada, H., & Loft, S. (2023). Evidence accumulation modelling in the wild: Understanding safety-critical decisions. *Trends in Cognitive Sciences, 27*, 175–188. https://doi.org/10.1016/j.tics.2022.11.009

Crosscombe, M., Lawry, J., & Bartashevich, P. (2019). Evidence propagation and consensus formation in noisy environments. In (p. 310-323).

Dempster, A.P. (1967). Upper and lower probabilities induced by a multivalued mapping. In *Classic works of the Dempster-Shafer theory of belief functions* (p. 57–72). Springer Berlin Heidelberg.

Dorigo, M., Theraulaz, G., & Trianni, V. (2020). Reflections on the future of swarm robotics. *Science Robotics, 5*(49), eabe4385. https://doi.org/10.1126/scirobotics.abe4385

Ebert, J.T., Gauci, M., Mallmann-Trenn, F., & Nagpal, R. (2020). Bayes bots: Collective bayesian decision-making in decentralized robot swarms. *2020 ieee international conference on robotics and automation (icra)* (p. 7186-7192). IEEE.

Garnier, S., Gautrais, J., & Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intelligence, 1*, 3–31.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30*, 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Golman, R., Hagmann, D., & Miller, J. H. (2015). Polya's bees: A model of decentralized decision-making. *Science Advances, 1*(8), Article e1500253. https://doi.org/10.1126/sciadv.1500253

Hamann, H. (2018). *Swarm robotics: A formal approach*. Springer.

Kaiser, T.K., Potten, T., & Hamann, H. (2023). Evolution of collective decision- making mechanisms for collective perception. *2023 ieee congress on evolutionary computation (cec)* (pp. 1–8).

Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of wald's sequential probability ratio test. *Neuron, 85*, 861–873. https://doi.org/10.1016/j.neuron.2015.01.007

Marshall, J. A., Brown, G., & Radford, A. N. (2017). Individual confidence-weighting and group decision-making. *Trends in Ecology and Evolution, 32*(9), 636–645. https://doi.org/10.1016/j.tree.2017.06.004

Myers, C., Interian, A., & Moustafa, A. (2022). A practical introduction to using the drift diffusion model of decision-making in cognitive psychology, neuroscience, and health sciences. *Frontiers in Psychology, 13*, 1039172. https://doi.org/10.3389/fpsyg.2022.1039172

Prasetyo, J., Masi, G. D., & Ferrante, E. (2019). Collective decision making in dynamic environments. *Swarm Intelligence, 13*, 217–243. https://doi.org/10.1007/s11721-019-00169-8

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation, 20*, 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Roxin, A. (2019). Drift-diffusion models for multiple-alternative forced-choice decision making. *Journal of Mathematical Neuroscience*. https://doi.org/10.1186/s13408-019-0073-4

Sajko, G., & Babič, J., et al. (2024). Achieving human-inspired drift diffusion consensus in swarm robotics. In H. Hamann (Ed.), *Swarm Intelligence* (pp. 29–41). Springer.
Schranz, M., Umlauft, M., Sende, M., & Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*. https://doi.org/10.3389/frobt.2020.00036
Shafer, G. (1976). *A mathematical theory of evidence*. Princeton University Press.
Shan, Q., & Mostaghim, S. (2020). Collective decision making in swarm robotics with distributed Bayesian hypothesis testing. *Lecture Notes in Computer Science, 12421*, 55–67.
Shan, Q., & Mostaghim, S. (2021). Discrete collective estimation in swarm robotics with distributed Bayesian belief sharing. *Swarm Intelligence, 15*, 377–402. https://doi.org/10.1007/s11721-021-00201-w
Statsenko, Y., Habuza, T., Gorkom, K. N. V., Zaki, N., & Almansoori, T. M. (2020). Applying the inverse efficiency score to visual-motor task for studying speed-accuracy performance while aging. *Frontiers in Aging Neuroscience*. https://doi.org/10.3389/fnagi.2020.574401
Thieu, T., & Melnik, R. (2023). *Social human collective decision-making and its applications with brain network models*. In (pp. 103–141).
Tump, A. N., Pleskac, T. J., & Kurvers, R. H. J. M. (2020). Wise or mad crowds? The cognitive mechanisms underlying information cascades. *Science Advances, 6*, eabb0266. https://doi.org/10.1126/sciadv.abb0266
Valentini, G., Brambilla, D., Hamann, H., & Dorigo, M. (2016). *Collective perception of environmental features in a robot swarm*. (pp. 65–76).
Valentini, G., Hamann, H., & Dorigo, M. (2014). Self-organized collective decision making: the weighted voter model. *Proceedings of the 2014 international conference on autonomous agents and multi-agent systems* (p. 45–52). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
Valentini, G., Hamann, H., & Dorigo, M. (2015). Efficient decision-making in a self- organizing robot swarm: On the speed versus accuracy trade-off. *Proceedings of the 2015 international conference on autonomous agents and multiagent systems* (p. 1305–1314). Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
Valentini, G. (2017). *Achieving consensus in robot swarms*. Springer.
Young, S. N. (2008). The neurobiology of human social behaviour: An important but neglected topic. *Journal of Psychiatry and Neuroscience, 33*, 391–392.

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