Crosscombe, M., Lawry, J., Hauert, S., & Homer, M. (2018). Robust distributed decision-making in robot swarms: Exploiting a third truth state. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2017): Proceedings of a meeting held 24-28 September 2017, Vancouver, British Columbia, Canada* (pp. 4326-4332). Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/IROS.2017.8206297>

Peer reviewed version

License (if available):
Unspecified

Link to published version (if available):
[10.1109/IROS.2017.8206297](https://doi.org/10.1109/IROS.2017.8206297)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <http://ieeexplore.ieee.org/document/8206297> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Robust Distributed Decision-Making in Robot Swarms: Exploiting a Third Truth State

Michael Crosscombe¹, Jonathan Lawry¹, Sabine Hauert¹ and Martin Homer¹

Abstract—In this paper, we investigate the best-of- n distributed decision problem in robot swarms. In this context, we compare the weighted voter model [25] with a three-valued model that incorporates an intermediate belief state meaning either ‘uncertain’ or ‘indifferent’. We focus particularly on the trade-off between speed of convergence to a shared belief, and robustness to the presence of unreliable individuals in the population. By means of both simulation and embodied experiments in real robot swarms of 400 Kilobots, we show that the three-valued model is much more robust than the weighted voter model, but with decreased speed of convergence.

I. INTRODUCTION

Reaching an optimal shared decision by applying only decentralised algorithms is a key aspect of many swarm robotic applications. For example, the best-of- n problem [16] requires a multi-agent swarm to select the best option from n mutually exclusive alternatives, based only on localised feedback and learning. In fact, this generic problem underpins a wide variety of distributed decision-making tasks (see [23] for an overview). In this paper we investigate the best-of- n problem in robot swarms, with an emphasis on fault tolerance and robustness. In particular, we propose a simple local belief updating rule which exploits a third intermediate truth state between ‘true’ and ‘false’, and show that when employed in swarm decision-making the result is system level behaviour which is robust to the presence of faulty or malfunctioning robots. Furthermore, through comparison with the weighted voter model [25] we suggest that there is a clear trade-off between speed of convergence and robustness.

In many applications it is important for robot swarms to be robust to a variety of different types of noise as well as hardware and software failure. In particular, the lack of calibration and the use of low-cost hardware can sometimes cause catastrophic failure [21]. In [27] five distinct ways are identified in which a swarm can be robust, including being tolerant to noise and uncertainties in the environment, or because it has no common-mode point of failure. The notion of robustness that concerns us here relates to “individual robots who fail in such a way as to thwart the overall desired swarm behaviour”. For the best-of- n problem, the desired swarm behaviour is that of convergence to the best decision, and we will investigate the effect of malfunctioning robots with the potential to disrupt this desired behaviour by making decisions on the basis of random beliefs.

One way of building fault tolerance into robot swarms is to enable individual robots to detect faults in their neighbours

so that they can compensate for them. This approach is referred to as *exogenous* fault detection [15]. For example, [4] propose an approach inspired by the synchronised flashing behaviour of fireflies in which each robot flashes by lighting up its on-board LEDs. Neighbouring robots then flash in synchrony unless they are malfunctioning. The non-periodic flashing of faulty individuals can then be detected by other members of the swarm. However, in the following we do not allow for exogenous fault detection and we assume that individual robots have no way of distinguishing between neighbours which are functioning correctly and those which are malfunctioning. Instead, robustness to the presence of faulty individuals should be inherent to the distributed decision-making algorithm employed.

An outline of the rest of the paper is as follows. In Section II we give an overview of the relevant existing literature on decentralised decision-making in swarm robotics, and also on the use of a third truth state in opinion dynamics. Then, in Section III we introduce our proposed three-valued voter model, and in Section IV we describe the robotic platform on which we implement it. Sections V and VI present the results of our investigation of both the three-valued and the weighted voter model for the $n = 2$ case, in simulation and experiments, respectively. Section VII then extends these models to the $n > 2$ case and presents simulation results for $n = 3$ and $n = 5$. Finally, in Section VIII we provide some concluding remarks.

II. RELATED WORKS

The weighted voter model [25] has been proposed as an extension of the classic voter model [6], [9], taking into account agent motion as well as feedback on the value or quality of the n different options. In this latter respect it is partly inspired by social-insects searching for nest sites e.g. honey bees and *Temnothorax* ants [2], [12]–[14], [22]. Network science methods are commonly applied to voter models, to understand the coupling between the states of the agents and the dynamics of the network which connects them [3], [10]. In contrast, here the interaction network is relatively straightforward, but we will explicitly include the effect of the agents’ motion. The algorithm has strong convergence properties as proven analytically by means of steady state analysis [24]. It is assumed that agents continually choose between the n alternatives based on their current beliefs, receiving feedback on the quality of their choices. Agents then signal their beliefs to their neighbours for a length of time which is proportional to the feedback they received for their latest choice. In order to update its beliefs an agent then

¹Department of Engineering Mathematics, University of Bristol, United Kingdom {m.crosscombe, j.lawry, sabine.hauert, martin.homer}@bristol.ac.uk

TABLE I

TRUTH TABLE FOR THE THREE-VALUED UPDATING OPERATOR.

updating agent	signalling agent		
	1	$\frac{1}{2}$	0
1	1	1	$\frac{1}{2}$
$\frac{1}{2}$	1	$\frac{1}{2}$	0
0	$\frac{1}{2}$	0	0

randomly selects a signalling agent within its communication range and simply adopts its beliefs.

Studies of decentralised decision-making applied to Kilobot swarms include recent work in [26] investigating the trade-off between speed of convergence and accuracy in the context of the majority rule. The effect of spatiality on the best-of- n decision problem is investigated in [18] with experiments being conducted on a swarm of 150 Kilobots. Other related work on robot swarms includes a honey bee nest-site inspired decision model implemented on Jasmine micro-robots [12]. An extensive overview of swarm decision-making research can be found in [23]. Also, [19] presents a general model of decentralised decision-making for the best-of- n problem. Interestingly this also employs a third truth state representing ‘uncommitted’. The updating model proposed, however, is inherently probabilistic with probabilities dependent on quality values. This is in contrast to our approach in which updating is a purely logical operation (as given in Table I). Another significant difference concerns the case in which $n > 2$, for which according to [19] an agent must be in one of $n + 1$ states; one for each option and an overall uncommitted state. However, in our approach there are $2^n - 1$ states which include the cases where the agent rules out certain options but is uncertain about the remaining ones.

In this paper we extend the weighted voter model to incorporate a third truth state representing either ‘unknown’ or ‘borderline true-false’ (see [5] for a discussion of the subtle difference between these two interpretations, which we do not dwell on here as it does not impact the system-level properties of interest). There have been a number of studies of three-valued models in the opinion dynamics literature. For example, [1] define the three truth states by applying a partitioning threshold to an underlying real value. Updating is pairwise between agents, and takes place incrementally on the real values, except that the magnitude and sign of the increments depends on the current truth states of the agents involved. An alternative three-valued operator is proposed in [17], and extended to incorporate feedback in [7]. It is applied directly to truth states and assigns the third truth state as a compromise between two opinions with strictly opposing truth values. We will adapt this operator to form the core of our proposed weighted model.

Fig. 1. State transition diagram for the three-valued voter model, with states S (signalling), U (updating), A and B (choose A and B respectively).

III. A THREE-VALUED VOTING MODEL

We now introduce a three-valued version of the weighted voter model, based on the belief combination operator proposed in [17] and [7]. Initially, we consider the $n = 2$ case of the best-of- n decision problem, and label the two options A and B (we discuss the generalisation to $n > 2$ in Section VII). Each agent is then in one of three possible belief states, 0, $\frac{1}{2}$ or 1, where 0 means “ B is the best option”, 1 means “ A is the best option” and where $\frac{1}{2}$ means either “I have no preference between A and B ” or “I am uncertain whether A or B is the best option”. Based on its current belief, an agent chooses between A and B as follows. If in belief state 1 then the agent chooses A , or if in belief state 0 it chooses B ; if in belief state $\frac{1}{2}$ it chooses between A and B at random. Feedback is given in the form of a positive integer; ρ_A for choosing A , and ρ_B for choosing B . Agents then enter a signalling state during which they broadcast their current belief to all other agents within communication range, for a time directly proportional to the value of their latest feedback. After signalling, an agent then updates its belief by randomly selecting a signalling agent within its radius of communication and applying the truth-table shown in Table I. The updating agent’s new belief state is thus a function both of its current belief and that of the selected signalling agent. This is in contrast to the weighted voter model in which an agent’s current belief does not directly influence their updated belief. The broad intuition underlying Table I is that a strong belief dominates except where there is a direct conflict between beliefs, in which case the intermediate (compromise) state of $\frac{1}{2}$ is adopted. For example, if the updating agent’s current belief state is 0 and the signalling agent’s belief state is 1, then the updating agent will adopt the new belief of $\frac{1}{2}$.

Figure 1 shows a state transition diagram for the three-valued model. In this simple model there are 4 states: S is the signalling state in which agents broadcast their current belief. The time that an agent spends in this state depends directly on the feedback they received when they made their latest choice; U is an updating state in which agents update their beliefs by applying the operator in Table I; and A and B are choice states corresponding to choosing A and B respectively. Agents in state A receive the quality ρ_A while those in state B receive ρ_B . Choices are made on the basis of an agent’s belief immediately after updating as described above. Furthermore, given the form of the

three-valued updating operator in Table I, this results in the following probabilities for choosing A and B , denoted P_A and P_B , based on the belief states of signalling agents within communication range: Letting Σ_i denote the proportion of other agents in an updating agent's radius of communication with belief state i for $i = 0, \frac{1}{2}$ or 1 , then the probabilities that the agent will choose A or B after updating are given by

$$P_A = \begin{cases} \Sigma_1 + \Sigma_{\frac{1}{2}} + \frac{1}{2}\Sigma_0 : \text{current belief} = 1, \\ \Sigma_1 + \frac{1}{2}\Sigma_{\frac{1}{2}} : \text{current belief} = \frac{1}{2}, \\ \frac{1}{2}\Sigma_1 : \text{current belief} = 0, \end{cases}$$

$$P_B = \begin{cases} \frac{1}{2}\Sigma_0 : \text{current belief} = 1, \\ \Sigma_0 + \frac{1}{2}\Sigma_{\frac{1}{2}} : \text{current belief} = \frac{1}{2}, \\ \Sigma_0 + \Sigma_{\frac{1}{2}} + \frac{1}{2}\Sigma_1 : \text{current belief} = 0. \end{cases}$$

Note that we have not included any exploration state. Instead we focus only on the part of distributed decision-making in which agents are mixing together, exchanging information, and updating beliefs. In this abstraction we assume that feedback is received immediately when a choice is made. This simplification means that the consensus times are considerably reduced, which enables us to run multiple versions of both simulation and embodied experiments.

In Sections V and VI we will compare the three-valued voting model with a variant of the weighted voting model proposed in [25]. The latter has two belief states 0 and 1 , meaning “ B is the best option” and “ A is the best option” respectively. The state transition model has the same structure as Figure 1, except that updating simply requires an updating agent to adopt the signalling agent's belief so that $P_A = \Sigma_1$ and $P_B = \Sigma_0$.

IV. KILOBOT SWARMS

Here we use Kilobots as a robotic platform for studying swarm decision-making. These are three-legged robots 33 mm in diameter and 34 mm tall, specifically designed to interact in large collectives, or ‘swarms’ [20]. Each Kilobot is an independent unit possessing, amongst other features, two motors providing left/right turning and forward motion, an RGB LED indicator for signalling to an observer (e.g. an overhead camera) and an infrared transceiver (see Figure 2). Kilobots have a communication range of approximately 10 cm, over which they can send and receive messages of up to 9 bytes in length. However, the simulator allows for communication radii exceeding this limit and we exploit this feature to explore a range of communication radii r between 0 and 20 cm. Given that the number of Kilobots and the size of the arena are fixed, r can serve as a proxy to allow us to vary the density of Kilobots involved in the updating process. We consider this effect in Section V (Figure 5). Alternatively, by varying r we can also study directly the robustness of the two algorithms to different constraints on communications as might be relevant to different robotic platforms.

For both the simulation (Section V) and embodied experiments (Section VI), a swarm of 400 Kilobots is deployed

Fig. 2. **Left:** A Kilobot robot with an LED light for signalling, an underside IR transceiver for communication and side-mounted motors for movement. **Right:** Arena used for experiments in which Kilobots interact while moving amongst one another.

in a square 1.2m^2 arena. Whilst in the signalling state S , Kilobots move randomly¹ by either turning left or right, moving forward, or remaining stationary, i.e. at each time one of these 4 options is chosen with equal probability. At initialisation Kilobots are distributed randomly across the arena but then, as a result of random motion, they may collide and cluster together. Simulations are implemented according to the state transition diagram shown in Figure 1.

As described in Section III, the experiments only model the mixing and information sharing part of the decision process, and the Kilobots do not visit specific physical locations or take other actions on the basis of their current beliefs. Instead, feedback is received immediately on the basis of their latest choice. While this is clearly a simplification, we believe that it still allows us to explore properties of the decision-making algorithms thanks to the reduction in run-time which allows us to repeat experiments multiple times.

V. SIMULATION EXPERIMENTS

In this section we describe experiments in which the Kilobots are simulated in a virtual environment and interact at random, iteratively updating their beliefs using the rules described in Section III, in order to form a consensus about which is the best option, A or B . Here we assume that option A is of higher quality than option B with respective quality values $\rho_A = 9$ and $\rho_B = 7$. This is an opinion-based approach with asymmetric qualities and symmetric costs [23].

We employ a Kilobot simulation environment [11] which captures many of the physical properties of a Kilobot swarm including motion, collisions, and communication between robots. The simulator also uses the same API² as used on the actual robots which makes it easier to transfer code from simulations to the real world. We present results for both the weighted voter model [25] and the three-valued model, comparing their convergence to consensus (i.e. population-wide convergence to a single option) for both different communication radii, and also in the presence of malfunctioning Kilobots. Results are averaged over 50 independent runs,

¹Due to the use of uncalibrated Kilobots, movement speed varies across the population.

²<https://www.kilobotics.com/docs/index.html>

Fig. 3. Percentages of signalling and updating Kilobots for both the weighted voter model and the three-valued model after 1000 iterations for different communication radii $r \in [0, 20]$.

Fig. 4. Average belief states after 1000 iterations for different communication radii $r \in [0, 20]$.

each of which terminate after 1000 iterations; we found this to be a sufficient number of iterations for the system to reliably reach a steady state in which a consensus is achieved.

Figure 3 shows the percentages of Kilobots in the signalling and updating states after 1000 iterations for a range of communication radii, for both the weighted voter model and the three-valued model. In this Figure, the lines labelled A and B respectively refer to the percentage of Kilobots who are currently in the signalling state having previously chosen A or B prior to entering that state. For $r \geq 5$ cm we see that a clear majority have chosen A for both the weighted voter and the three-valued models.

A more direct way of measuring convergence is to evaluate the average belief state of the Kilobots. In a population of k individuals we define the average belief state in a given iteration as follows: Let $\mathcal{B}_i$ denote the belief state of agent i for that iteration, then;

$$\text{average belief state} = \frac{1}{k} \sum_{i=1}^k \mathcal{B}_i$$

This corresponds to a weighted average of 0 and 1 in the weighted voter model and of 0, $\frac{1}{2}$ and 1 in the three-valued model. For the weighted model there is a direct relationship between the average belief state and the percentage of the population choosing either A or B . This is because an agent chooses A (B) if and only if their belief state is 1 (0). For the three-valued model, however, the relationship between these two measures is less direct since, while an agent will

Fig. 5. Average number of messages received by each Kilobot per iteration for different communication radii $r \in [0, 20]$.

Fig. 6. Average belief states against iterations for a communication radii of 10 cm.

definitely choose A when in belief state 1, they may also choose A (with probability 0.5) when in the intermediate belief state $\frac{1}{2}$. Hence, on average we would expect the percentage of agents choosing A to be proportional to the number of agents in belief state 1 plus 50% of the number in belief state $\frac{1}{2}$.

The average belief states are shown in Figure 4 where for $r \geq 5$ cm we can see that both models result in almost all Kilobots adopting the belief state 1 after 1000 iterations. More precisely, for $r = 10$ cm the average belief states for the weighted voter model and the three-valued model after 1000 iterations are 1.00 and 0.99 respectively. It is interesting to note that for the three-valued model the intermediate truth state is also totally abandoned, suggesting that there is convergence to total certainty that A is the best option. The average number of messages per unit time received by each Kilobot in the updating state as a function of communication radius is shown in Figure 5. Notice that for $r = 5$ cm Kilobots receive just under 2 messages per unit time suggesting that both algorithms are robust to a relatively low population density of Kilobots.

If we measure quality of convergence either by the percentage choosing A or by the average belief state, then Figures 3 and 4 all suggest that convergence to A is slightly better for the weighted voter model than for the three-valued model, although the difference is very small for $r \geq 5$ cm. Furthermore, speed of convergence also appears to be faster for the weighted voter model. For example, Figure 6 shows the trajectory of average beliefs against iterations for both models when the communication radius is 10 cm. In this

Fig. 7. State transition diagram for malfunctioning Kilobots. R denotes randomize beliefs.

Fig. 8. Average belief in the best site after 1000 iterations for malfunction rates $\lambda \in [0, 100]$ and a communication radius $r = 10$ cm.

case the weighted voter model converges after about 200 iterations, while the three-valued model needs around 600 iterations to converge.

We now use the simulation environment to investigate the robustness of both models to the presence of malfunctioning Kilobots amongst the population. We assume that a certain percentage λ of the Kilobots malfunction by selecting their beliefs at random as shown in the state transition diagram in Figure 7. Here R refers to a state in which the Kilobot simply selects its new belief state at random by picking uniformly from $\{0, 1\}$ in the case of the weighted voter model and from $\{0, \frac{1}{2}, 1\}$ for the three-valued model. Consequently, for both models there is then a probability of 0.5 that they will choose either A or B and receive the associated feedback value. As for functioning Kilobots, malfunctioning Kilobots then enter the signalling state and remain there for time ρ_A or ρ_B depending on their latest option. We have adopted this particular model of malfunction as one which is likely to disrupt convergence to the desired belief state, by broadcasting randomized belief states to functioning Kilobots when the latter are updating their beliefs.

Figure 8 shows the average belief states after 1000 iterations for the weighted voter and the three-valued model respectively, with different percentages of malfunctioning agents (i.e. $\lambda \in [0, 100]$) for a communication radius of 10 cm. Here the belief states are averaged across functioning Kilobots only³, as the population will never appear to fully converge while the randomly-signalling, malfunctioning agents are included. From these Figures it is apparent that the

³Except for the case where $\lambda = 100\%$, in which results are averaged across all Kilobots due to the lack of any functioning Kilobots present in the population.

Fig. 9. Percentage of Kilobots signalling for A and B against iterations.

Fig. 10. Percentages of Kilobots signalling for A and B against iterations and with $\lambda = 10\%$.

three-valued model is more fault tolerant than the weighted voter model in that it achieves average belief state values closer to 1 for each of the values of λ . For example, given a communication radius of 10 cm and assuming that 10% of the population is malfunctioning, then the three-valued model converges to an average belief state of 0.99 in the highest-value option while the weighted voter model converges to an average belief state of 0.87. Indeed, even if 50% of the population is malfunctioning then the three-valued model still converges to an average belief of 0.83 while the average belief of the weighted voter model drops to 0.67.

VI. KILOBOT SWARM EXPERIMENTS

We now describe a series of experiments conducted on actual swarms of 400 Kilobots which follow the same template as the simulation studies in Section 4. Figure 2 shows the $1.2m^2$ arena used. Note that it has a smooth and reflective surface so as to allow good communication between Kilobots and to enable motion. During the experiments each Kilobot in the signalling state displays a coloured light using its LED to indicate its most recent choice; blue for A and red for B . A video was made of every experiment and analysed using standard image processing algorithms (OpenCV) to identify the different coloured lights and to determine a time series of the proportion of robots favouring options A and B . Each experiment was run independently 10 times with mean and percentiles (10% and 90%) then being determined. These are shown in Figures 9 and 10 with error bars indicating the 10th and 90th percentiles.

An overhead controller (OHC) was used to upload programs and initialisation instructions to each Kilobot. This resulted in non-uniform starting times across the population, leading to high variance in the results for the first 60 iterations. Although unsynchronised, each Kilobot updates approximately once every $\frac{1}{4}$ second, and hence we take

Fig. 11. Average belief in the best site after 1000 iterations for $n = 3$, malfunction rates $\lambda \in [0, 100]$ and a communication radius $r = 10$ cm.

this time period as corresponding to an iteration so that an experiment conducted over 1000 iterations lasts just over 4 minutes.

Figure 9a shows the percentage of Kilobots signalling A or B as a function of time for the weighted voter model. In this case we can see that the swarm converges on option A after approximately 240 iterations. In contrast, the three-valued model only fully converges to A after 800 iterations as can be seen in Figure 9b. Hence, as is consistent with the simulation studies we see that the weighted voter model significantly outperforms the three-valued model in terms of speed of convergence. However, after 1000 iterations the level of convergence is the same for both models.

We also conducted experiments to test how fault-tolerant the two models were to the presence of malfunctioning Kilobots in the population. Here we introduced faulty Kilobots which malfunctioned according to the state transition diagram in Figure 7 and which made up $\lambda = 10\%$ of the population. As in the simulation experiments, the Kilobots signalling for each option are recorded as a percentage of the functioning individuals only. Figure 10a shows the percentage of functioning signalling Kilobots which have chosen A and B , as a function of time, for the weighted voter model. After 1000 iterations we see that 86.5% of the functioning Kilobots have chosen A . In contrast, Figure 10b shows that the three-valued model still maintains almost total convergence to A , notwithstanding the 10% of Kilobots that are malfunctioning, with 99.7% of functioning Kilobots choosing A after 1000 iterations. On the other hand, Figure 10a shows that the weighted voter model achieves steady-state after about 120 iterations, while from Figure 10b we can see that the three-valued model requires around 700 iterations to achieve steady-state. Hence, as is consistent with the simulation experiments in Section 4, these results suggest that while the weighted voter model converges more quickly than the three-valued model, the latter is more fault tolerant than the former.

VII. BEST-OF- n PROBLEM FOR $n > 2$

Much of the existing literature on the best-of- n problem for swarms concerns the $n = 2$ case, and so far we have dealt exclusively with this case. However, while we can directly extend the weighted voter model to the $n > 2$

Fig. 12. Average belief in the best site after 1000 iterations for $n = 5$, malfunction rates $\lambda \in [0, 100]$ and a communication radius $r = 10$ cm.

case, it is less immediately clear how best to extend the three-valued model. One natural approach is to define belief states as n -dimensional vectors in $\{0, \frac{1}{2}, 1\}^n$, so that, for example, the belief state $\langle 0, \frac{1}{2}, \frac{1}{2}, 0, \dots, 0 \rangle$ is interpreted as meaning that options 2 and 3 are believed to be better than all the other options, but that there is no preference between them. The updating operator in Table I could then be applied independently to each dimension of the relevant belief states. For example, updating $\langle 0, \frac{1}{2}, \frac{1}{2}, 0, \dots, 0 \rangle$ given the signalled state $\langle \frac{1}{2}, \frac{1}{2}, 0, \dots, 0 \rangle$ results in the updating agent adopting the new state $\langle 0, \frac{1}{2}, 0, \dots, 0 \rangle$. However, the latter is a belief state in which, although the agent has ruled out all except the second option, they still remain uncertain that this is the best option. In effect they are not taking account of the fact that in the best-of- n problem the n options are assumed to be exhaustive. Our approach is then to incorporate a form of normalisation into the model, so that, for example, $\langle 0, \frac{1}{2}, 0, \dots, 0 \rangle$ is normalised to $\langle 0, 1, 0, \dots, 0 \rangle$.

Using this approach we now present preliminary results for the $n = 3$ and $n = 5$ cases using the simulation environment. For $n = 3$ we assumed that the options were A , B and C with quality values $\rho_A = 11$, $\rho_B = 8$, $\rho_C = 5$ and for $n = 5$ the options were A , B , C , D and E with quality values $\rho_A = 25$, $\rho_B = 20$, $\rho_C = 15$, $\rho_D = 10$ and $\rho_E = 5$. Figures 11 and 12 show the average belief values for both algorithms plotted against the percentage of malfunctioning Kilobots $\lambda \in [0, 100]$. Overall, in both cases the three-valued model is more robust to malfunction than the weighted voter model. Although, as can be seen in Figure 12 the three-valued model performs worse for lower malfunction rates where $\lambda \leq 15\%$. This may be a result of reduced overall convergence of the three-valued model as the number of options, n , increases. Further research is required in order to explore this effect more fully and, in general, to provide a more extensive analysis of the $n > 2$ case.

VIII. CONCLUSIONS

In this paper we have proposed a three-valued model for belief updating in the best-of- n problem, and compared it to the weighted voter model. We have focussed on robustness to individual malfunction or error, and also on speed of convergence. Experiments were conducted using a realistic

simulation environment, as well as on actual Kilobot swarms of 400 robots. The results from both sets of experiments agree that the three-valued model is more robust to the presence of malfunctioning or noisy individuals in the population than the weighted voter model, but that the weighted voter model has the advantage of converging more quickly to the best option.

We note that in both models, belief updating is based on the belief state of only one signalling agent. This property may be advantageous in scenarios where there is either very limited communications or low swarm density. Nonetheless, future work on robustness should consider decision algorithms which take account of larger samples of belief states drawn from the signalling agents within an individual's radius of communication. This could include majority rule models as studied in [26] as well as probabilistic pooling operators of the kind reviewed in [8]. One might hypothesise that by taking account of a larger sample of signalling agents, models would tend to be more robust to noise, error and malfunction. However, this robustness still needs to be considered in a broader context which also takes account of speed and overall level of convergence.

In our experiments we have investigated robustness to the presence of a particular type of malfunctioning agent, in which error results from a proportion of the population continually selecting their beliefs at random rather than as part of the belief updating process. Clearly there are other models of error which should also be studied. For example, we might consider the errors resulting when some agents constantly broadcast the same fixed beliefs or when agents maliciously broadcast the 'wrong' belief. Furthermore, while we have focussed on the case where there is a fixed proportion of malfunctioning or erroneous agents, and all other agents are error-free, it is also important to consider noise resulting from generic errors or sensing and processing limitations to which all agents are equally susceptible.

ACKNOWLEDGMENT

We are grateful to the reviewers for their insightful comments and suggestions.

This research was partially supported by the EPSRC Engineering Complexity Resilience Network Plus (ENCORE) EP/N010019/1. Research data and supporting material can be found at <http://bit.ly/ENCORE-EP-N010019-1>.

REFERENCES

- [1] P. Balenzuela, J. P. Pinasco, and V. Semeshenko. The undecided have the key: Interaction-driven opinion dynamics in a three state model. *PLOS ONE*, 10(10):1–21, 10 2015.
- [2] N. F. Britton, N. R. Franks, S. C. Pratt, and T. D. Seeley. Deciding on a new home: how do honeybees agree? *Proceedings of the Royal Society of London B: Biological Sciences*, 269(1498):1383–1388, 2002.
- [3] L. Chen, C. Huepe, and T. Gross. Adaptive network models of collective decision making in swarming systems. *Phys. Rev. E*, 94:022415, Aug 2016.
- [4] A. L. Christensen, R. OGrady, and M. Dorigo. From fireflies to fault-tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766, Aug 2009.
- [5] D. Ciucci, D. Dubois, and J. Lawry. Borderline vs. unknown: comparing three-valued representations of imperfect information. *International Journal of Approximate Reasoning*, 55(9):1866 – 1889, 2014. Weighted Logics for Artificial Intelligence.
- [6] P. Clifford and A. Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.
- [7] M. Crosscombe and J. Lawry. Exploiting vagueness for multi-agent consensus. pages 67–78. Springer Singapore, Singapore, 2017.
- [8] F. Dietrich and C. List. Probabilistic opinion pooling generalized. part one: general agendas. *Social Choice and Welfare*, 48(4):747–786, Apr 2017.
- [9] R. A. Holley and T. M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *Ann. Probab.*, 3(4):643–663, 08 1975.
- [10] P. Holme and M. E. J. Newman. Nonequilibrium phase transition in the coevolution of networks and opinions. *Physical Review E*, 74:056108, 2006.
- [11] S. Jones, M. Studley, S. Hauert, and A. Winfield. *Evolving behaviour trees for swarm robotics*. Springer Tracts in Advanced Robotics. Springer, 9 2016.
- [12] S. Kernbach, R. Thenius, O. Kernbach, and T. Schmickl. Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. *Adaptive Behavior*, 17(3):237–259, 2009.
- [13] C. List, C. Elsholtz, and T. D. Seeley. Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1518):755–762, 2009.
- [14] J. A. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, and N. R. Franks. On optimal decision-making in brains and social insect colonies. *Journal of The Royal Society Interface*, 2009.
- [15] A. G. Millard, J. Timmis, and A. F. T. Winfield. *Towards Exogenous Fault Detection in Swarm Robotic Systems*, pages 429–430. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [16] C. A. C. Parker and H. Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251, April 2009.
- [17] E. Perron, D. Vasudevan, and M. Vojnovic. Using three states for binary consensus on complete graphs. In *IEEE INFOCOM 2009*, pages 2527–2535, April 2009.
- [18] A. Reina, T. Bose, V. Trianni, and J. A. Marshall. Effects of spatiality on value-sensitive decisions made by robot swarms. In *Proc. 13th Int. Symp. Distrib. Auton. Robot. Syst.*, 2016.
- [19] A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLOS ONE*, 10(10):1–18, 10 2015.
- [20] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal. Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems*, 62(7):966 – 975, 2014. Reconfigurable Modular Robotics.
- [21] M. Rubenstein, A. Cornejo, and R. Nagpal. Programmable self-assembly in a thousand-robot swarm. *Science*, 345(6198):795–799, 2014.
- [22] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–111, 2012.
- [23] G. Valentini, E. Ferrante, and M. Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9, 2017.
- [24] G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo. Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3):553–580, May 2016.
- [25] G. Valentini, H. Hamann, and M. Dorigo. Self-organized collective decision making: The weighted voter model. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14*, pages 45–52, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
- [26] G. Valentini, H. Hamann, and M. Dorigo. Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, pages 1305–1314, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.
- [27] A. F. T. Winfield and J. Nembrini. Safety in numbers: fault-tolerance in robot swarms. *IJMIC*, 1:30–37, 2006.