# Opinion dissemination in a swarm of simulated robots with stubborn agents: a comparative study

1st Guillaume Maître
*Faculty of Computer Science*
*University of Namur*
Namur, Belgium
guillaume.maitre@student.unamur.be

2nd Elio Tuci
*Faculty of Computer Science*
*University of Namur*
Namur, Belgium
elio.tuci@unamur.be

3rd Eliseo Ferrante
*Department of Computer Science*
*Vrije Universiteit Amsterdam*
Amsterdam, NL
e.ferrante@vu.nl
*Technology Innovation Institute*
Abu Dhabi

***Abstract*—Classic opinion dissemination models such as Majority model, and Voter model are particularly important in swarms robotics because they regulate the interactions among the agents while the swarm is engaged in collective decision-making processes requiring the consensus of the large majority or the unanimity of the group's members. In this paper, we compare the effectiveness of three different opinion dissemination models in a specific scenario where consensus is searched on an opinion disseminated within the swarm by a small number of legitimate agents. The task of the legitimate agents is hindered by few adversarial agents which disseminate within the swarm an "invalid" message. This scenario is meant to model a data communication manipulation attack in a swarm of robots. By comparing the dissemination models in different experimental conditions, the results of our study inform us on which model is more effective in supporting the case of legitimate agents, while reducing the disruptive effects of the data communication manipulation attack.**

## I. Introduction

Opinion dissemination models concern the way in which single individuals change opinion based on the interactions with other individuals. Opinion dissemination dynamics have been extensively studied to understand the mechanisms that regulate the dissemination of opinions in human society [4], [16]. They have also been studied in ethology to unveil the mechanisms underpinning the processes by which social insects such as ants and bees make decisions collectively: i.e., in a way that, once made, the choice is no longer attributable to any of the individual members of the group [5]. A collective decision is made when the large majority or the unanimity of the group members favour the same option.

In this paper, we focuses on opinion dissemination in swarm robotics systems. A robot swarm is a group of robots that co-operate to accomplish a mission that a single robot is unable to accomplish alone. A robot swarm operates in a self-organised and distributed manner: there is no leader and coordination is obtained via interaction between the individual robots [2]. A robot swarm does not rely on any external infrastructure: each individual robot acts on the basis of local information obtained through its sensors or provided via local communication by neighbouring robots. Opinion dissemination is particularly important in swarm robotics, since in the absence of a group leader, any decision in presence of alternative options, has to be taken collectively, and without relying on concepts such as reputation and confidence in group mates [13]. This is because the robots of a swarm are "anonymous". Thus, they communicate through modifications of the local environment (e.g., by emitting sound or by generating other types of signals that are eventually detected by other agents located within signals' range). These characteristics make the swarm extremely flexible and robust but also particularly vulnerable to intentional and systematic disruptions from an adversarial source [12]. In this paper, we focus on an "attack scenario" in which malicious adversarial agents disrupt a data communication process in a swarm of mobile robots. Our study aims to evaluate the extent to which different opinion dissemination models shield the swarm from such malicious and potentially disruptive attack.

Within swarm robotics, opinion dissemination mechanisms are generally studied with respect to the best-of-$n$ problem, where a swarm is required to chosen the best option out of the $n$ available. As extensively discussed in [14], quality and cost of the options can be used to further describe the nature of the best-of-$n$ decision-making problem. For example, in a foraging scenario, where food availability is the option quality and time to reach the food patch is the option cost, the problem can be symmetric/asymmetric for quality (all food patches have/have not the same amount of food), and/or symmetric/asymmetric for cost (all food patches require/do not require the same amount of time to be reached). When both costs and quality are asymmetric, we can have scenarios in which the option costs and quality are synergic (e.g., the best option has maximum quality and minimum cost) and scenario in which they are antagonistic (e.g., the best option has maximum quality and highest costs).

The opinion dissemination scenario for mobile robots investigated in this paper is adopted from the one described in [11] to model a data communication manipulation attack in a swarm of robots. The scenario consists in a very simple and basic form of the best-of-$n$ problem, where quality of the options does not need to be evaluated, and the environment is symmetric with respect to cost of accessing the options. In particular, robots of a swarm randomly move in a close arena and communicate each other their current opinion state.

978-1-7281-6929-3/20/$31.00 ©2020 IEEE

The opinion space is binary: there is option $0$ and option $1$. Moreover, a certain number of robots in the swarm are stubborn: that is, they never change their opinion. Stubborn robots differentiate in legitimate and adversarial. Legitimate robots disseminate the “correct” opinion, and adversarial agents disseminate the “wrong” opinion. The non-stubborn robots in the swarm can change their opinion based on interactions with nearby stubborn and non-stubborn robots and according to the mechanisms of the opinion dissemination model in place. In similar symmetry breaking scenario, the predicted outcome is usually consensus to the piece of information held by the majority of stubborn agents [8], [15]. Thus, since in [11] the swarm is made in a way to have more legitimate than adversarial agents, the opinion dissemination process is expected to converge on a consensus on the opinion of legitimate agents. However, in [11], the authors show that the addition of adversarial agents in a smaller number than the legitimate agents hinders reaching consensus to the majority opinion. The contribution of the study described in [11] is in the development of a new opinion dissemination model (referred to as Probabilistic model in this paper) where each member of the swarm has a probability $p$ to change its opinion based on new incoming messages from the neighbours. The parameter $p$ is initially set to 1 for all non-stubborn agents, but then, during the course of the simulation, each agent adjusts this individual probability according to the following information updating rules: the individual probability to change opinion $p$ is decreased anytime a robot interacts with an agent committed to the same opinion; $p$ is increased anytime a robot interacts with an agent committed to a different opinion. The results illustrated in [11] show that the Probabilistic model managed to contain and, in several experimental conditions to suppress the dissemination of wrong information by malicious adversarial agents, even if its effectiveness progressively vanishes with the number of malicious agents approaching that of legitimate agents.

This study aims to further evaluate the effectiveness of the Probabilistic model in a more informative comparative robot-based setting. First, we compare the Probabilistic model with classic opinion dissemination models such as the Majority and Voter model in the same data communication manipulation scenario illustrated in [11]. In this scenario, legitimate, malicious, and non-stubborn robots interact and communicate their opinions for long enough until the opinions’ distribution in the swarm does not change any longer. The results of our comparative study offer a more informative perspective to evaluate the extent to which the Probabilistic model shields a swarm of robots from the malicious intentions of agents that act in order to manipulate communication and eventually to hinder the swarm from achieving consensus on the “correct” opinion. Second, in [11], the Probabilistic model is evaluated in a 2D toroidal grid world, in which agents are just points that communicate anytime they find themselves in the same grid cell. In [11], the use of this simple environment has been dictated by the necessity to reduce the computational costs related to the evaluation of the different parameters of the opinion dissemination model. However, the lack of properties such as the embodiment of the robots and their situatedness in world that complies with the laws of physics largely simplifies the agents’ interactions and the dynamics of the opinion dissemination process. In this study, we run our tests in a simulation environments in which robots have (simulated) bodies based on the characteristics of the physical robot Foot-bot (see section II for details). The motion of the Foot-bots in the simulated environment is modelled by taking into account the physical robots’ kinematics properties. The interactions between the robots and their physical and social environment is based on the the characteristics of physical Foot-bot sensors. Thus, our opinion dissemination dynamics happen in a simulated word in which communication between robots does not ignore the physical properties of the agents and of their environment. On the contrary, it is made possible by them. In this way, we offer a comparative evaluation of the effectiveness of different opinion dissemination models which is more pertinent for the swarm robotics community since it is not biased by untenable assumptions about how the robots move and interact.

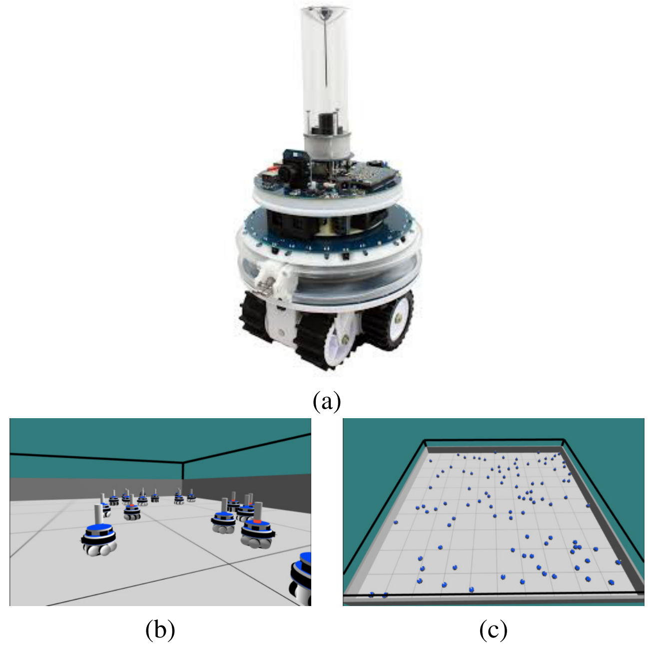

Fig. 1. In (a), image of the physical Foot-bot mobile robot. Image downloaded from www.swarmanoid.org. In (b) and (c) the simulated arena with the simulated Foot-bots.

In the next section, we describe the data communication manipulation scenario used to compare the opinion dissemination models and we illustrate these models. In section III, we describe the results of our study, and in section IV, we draw our conclusions.

## II. Methods

A swarm of $100$ simulated robots is located in a square arena (10 m x 10 m, see Figure 1b and 1c). The robots moves according to a isotropic random walk, with a fixed step length

(5 seconds, at 20 cm/s), and turning angles chosen from a wrapped Cauchy probability distribution characterised by the following PDF:

$$f(\theta,\mu,\rho) = \frac{1}{2\pi}\frac{1-\rho^2}{1+\rho^2-2\rho\cos(\theta-\mu)},\ 0<\rho<1, \quad (1)$$

where $\mu = 0$ is the average value of the distribution, and $\rho$ determines the distribution skewness [7]. For $\rho = 0$ the distribution becomes uniform and provides no correlation between consecutive movements, while for $\rho = 1$ a Dirac distribution is obtained, corresponding to straight-line motion. In this study $\rho = 0.5$. While moving around, the robots continuously perform an obstacle avoidance behaviour. To perform obstacle avoidance, first a robot stops, and then it keeps on changing its headings of a randomly chosen angle uniformly drawn in $[0, \pi]$ until no obstacles are perceived.

To simulate the above mentioned environment we use ARGoS multi engine simulator [9]. The simulation environment models the square arena as detailed above, and the kinematic and sensors readings of the Foot-bot mobile robot, a circular based robot of 13 cm diameter [1]. For our experiments, we made use of only a subset of all possible sensors mounted on the Foot-bots. In particular, the sensory apparatus of the simulated robots used in this study includes twelve proximity sensors positioned in the front half of the robot's circular body, and the omni-directional camera. The proximity sensors have a range of few centimetres and are used for sensing and avoiding the arena's walls and also for sensing and avoiding collisions with other robots. The camera, whose view is limited to a circular field centred on the robot body with radius 1.2 m, is used to detect robots within filed of view. Each robot has a LED located on the top face of its body which can emit red or blue light. The perception of nearby robots through the camera, is achieved by detecting coloured blobs generated by the LEDs' light.

In our study, robots can have different opinions. Each robot disseminates its current opinion by appropriately setting the colour of its LED. In particular, a blue LED signals opinion $1$ and a red LED signals opinion $0$. While the LED is used to disseminate the robots own opinions to nearby agents, the camera is used to sample the opinion of nearby robots. The robots differ not only in terms of the opinion they are committed to (i.e., $0$ or $1$) but also in terms of the capability to change their opinion. The robots of type "stubborn" never change opinion during the simulation. They can be fully committed to the opinion labelled $1$ (hereafter, we refer to these robots as $S1$) or to the opinion labelled $0$ (hereafter, we refer to these robots as $S0$). The non-stubborn robots are initially committed to an opinion that can be either $0$ (hereafter, we refer to these robots as $NS0$) or $1$ (hereafter, we refer to these robots as $NS1$). Differently from stubborn robots, non-stubborn robots can change their opinion (from $0$ to $1$ or vice-versa) based on the opinion dissemination model that regulates the robots' interaction. We consider robots $S1$ as legitimate, and robots $S0$ as adversarial. The objective of this study is to evaluate, in different experimental conditions, which dissemination model facilitates the dissemination of the legitimate robots' opinion to non-stubborn robots, thus limiting the influence of adversarial robots.

All robots disseminates their opinion all the time. Only non-stubborn robots observe and make use of the opinion of nearby robots to update their own opinion. The observation of the opinions of the nearby robots happens at irregular and individual time intervals which, for each robot, are computed by sampling an exponential function with $\lambda = 10$. With this exponential function, the average time interval between two consecutive opinion sampling events (or observations) corresponds to $1.4$ s. During an observation, each non-stubborn robot samples the opinions of nearby robots and based on the dissemination model in place it chooses to either change or not its current opinion. The observation and opinion changing process is regulated by three different opinion dissemination models: the voter model ($Vm$), the majority model ($Mm$), and the probabilistic model ($Pm$).

When the opinion dissemination is regulated by the voter model, at the time of observation each non-stubborn robot randomly selects a robot among those within camera view. If the selected neighbour is signalling a opinion different from the one of the observing robot (e.g., the selected robot is signalling red and the observing robot is signalling blue, or vice-versa), the observing robot changes its opinion (i.e., from $0$ to $1$ or vice-versa). If observing and selected robot share the same opinion, or if no neighbours are within camera view of the observing robot, the later agent keeps its current opinion.

When the opinion dissemination is regulated by the majority model, at the time of the observation each non-stubborn robot randomly selects two robots among those within camera view. Then, it applies the majority rule. According to this rule, the opinion that the observing robot uses for comparison with its own opinion is the most represented in a group of three agents made by the two selected robots plus itself. If the most represented opinion is the same as the one of the observing robot, this later agent does not change its opinion. If instead, the most represented opinion is different to the one of the observing robot, this later changes its opinion to the most represented one in the group of three agents. If at the time of the observation there is not enough neighbours within camera view (i.e., less than two robots), the observing robot does not change opinion.

The model we referred to as Probabilistic has been firstly introduced and tested in [11]. When the opinion dissemination is regulated by the Probabilistic model, at the time of the observation each non-stubborn robot randomly selects a robot among those within camera view. Then, if the selected and observing robots have different opinions (e.g., the observing robot is currently committed to $1$ and the selected robot is committed to $0$, or vice-versa), the observing robot changes to the opinion of the selected robot with probability $p$. At the same time it updates the individual parameter $p \in [0, 1]$ in the following: $p = \lfloor p/z \rfloor$ with $z = 0.4$. Since $p$ is bounded in $\in [0, 1]$, anytime $p$ becomes bigger than 1, $p$ is set to 1. If instead the selected and observing robots share the same opinion, the observing robot keeps its opinion, and at the same time it updates the individual parameter $p$ in the following: $p = \lfloor p \times k \rfloor$ with $k = 0.8$. In other word, the individual $p$ of a robot observer increases anytime observer and selected robot have different opinions, and it decreases anytime observer and selected robot share the same opinion. Thus, observing robots with the same opinion reduces the observer probability to change its opinion in the subsequent observation; observing robots with a different opinion increases the observer probability to change its opinion in the subsequent observation.

The parameter $p$ is initialised to $p = 1$ for all non-stubborn robots at the beginning of each trial. A trial starts when the robots are randomly placed in the arena, and lasts 50.000 time steps. At trial start, half of the non-stubborn robots are randomly chosen to be committed to 1 and half to 0. At each time steps the position of each robot is updated, and for those non-stubborn robots that are in an observation state, the opinion is checked and eventually updated. The updating of the parameter $p$ is executed regardless of the outcome of the stochastic opinion changing process regulated by $p$. If at the time of the observation there is not neighbours within camera view, the observing robot does not change opinion and the parameter $p$ is not updated. The effects of different values of the parameters $k$ and $z$ on the opinion dissemination process have been tested in [11]. In the set of simulations described in this study, we fixed the value of these parameters to $k = 0.8$ and $z = 0.4$ for all non-stubborn robots. This is because according to the results shown in [11], these are the parameters' values that return the highest number of simulations with robots converging to the legitimate robots' opinion, even in conditions in which the difference between the number of legitimate and adversarial robots in the swarm is just one.

## III. RESULTS

Our study compares the effects of the three different opinion dissemination models described in section II in swarms of robots that differ for the initial number of stubborn and non-stubborn robots. In particular, we adopted the experimental design originally described in [11], with two sets of 15 different experimental conditions. In the first set, the number of legitimate robots (i.e., stubborn and committed to opinion 1, robots $S1$) is fixed to five. The number of adversarial robots (i.e., stubborn and committed to opinion 0, robots $S0$) is varied from $S0 = 1$ to $S0 = S1 = 5$, by increasing them of one agent at the time. In the second set of simulations, the number of legitimate robots is fixed to ten. The number of adversarial robots is varied from $S0 = 6$ to $S0 = S1 = 10$ by increasing them of one agent at the time. Differently from [11], we have also run a third set of simulations, in which the number of legitimate robots is fixed to twenty. In this set of simulations, the number of adversarial robots is varied from $S0 = 10$ to $S0 = S1 = 20$ by increasing them of two agents at the time. Since the swarm size is fixed to 100 in all experimental conditions, the total number of non-stubborn robots ($NS0 + NS1$) in each experimental condition is given by $100 - S0 - S1$. For each pair of values $S1 - S0$ (i.e., first set 5-1, 5-2, 5-3, 5-4, 5-5, second set 10-5, 10-6, 10-7, 10-8, 10-9, 10-10, and third set 20-10, 20-12, 20-14, 20-16, 20-18, 20-20), we have studied the opinion dissemination dynamics with the three opinion dissemination models described in section II. For each condition, given by the number of robots type $S1$, the number of robots type $S0$, and the opinion dissemination model, we have run 50 differently seeded simulation trials. We remind the reader that the objective of this study is to find out which opinion dissemination model guarantees the maximum dissemination of the opinion held by legitimate robots ($S1$) by limiting the influence of adversarial robots ($S0$). To assess the statistical significance of these results, we utilised the Mann-Whitney-Wilcoxon test to compare the proportion of $NS1$ robots generated at the end of the runs by each dissemination model in each experimental condition, and a generalised linear model with a binomially distributed response, as each robot in each run can only converge to two possible outcomes. With this later model, we tested the effect of all the factors, which include the utilised model, the number of stubborn robots $S0$ and $S1$, and all possible interaction combinations among them.

Figure 2 shows the results of our simulations. Each graph shows the proportion (i.e., $\lfloor NS1/(NS1 + NS0) \rfloor$) of non-stubborn robots $NS1$ at the end of each trial. The white boxes refer to the simulations with the Majority model, the light grey boxes to the Probabilistic model, and the dark grey boxes to the Voter model. The x-axes refer to the number of adversarial robots $S0$. Figure 2a refers to the set of simulations with five legitimate robots. First, we notice that when $S0 = 1$, the medians of the proportion of non-stubborn robots committed at opinion 1 at the end of a trial is close to one for all dissemination models. This means that when the ratio $\lfloor S0/S1 \rfloor$ is $\lfloor 1/5 \rfloor$ the large majority (if not the totality) of the non-stubborn robots end the trial committed to the opinion of legitimate robots. With the increment of the number of adversarial robots $S0$ the median of the proportion of non-stubborn robots committed to opinion 1 at the end of the simulation decreases, and the variability increases, at least for the Majority and the Probabilistic models. The median decreases slower when the opinion dissemination is regulated by the Majority model (see Figure 2a, white boxes), than when the dissemination process is regulated by the Probabilistic and the Voter model (see Figure 2a, light and dark grey boxes). This is supported by the fact that the effect of the interactions between the models and the number of adversarial robots $S0$ was found to be significant, which means that the two regression curves decrease at a significantly different rate. The Probabilistic model tends to generate results closer to those generated by the Majority model (see Figure 2a, white and light grey boxes).

When the ratio $\lfloor S0/S1 \rfloor$ is one ($S0 = S1 = 5$), three different end results are the most frequently observed. For $Mm$ and $Pm$, the boxes spread along the entire range of the y-axis's values (see Figure 2a white and light grey boxes). For $Vm$ the box is more condensed (see Figure 2a dark grey box). A closer look at the data shows that, for $Mm$ and

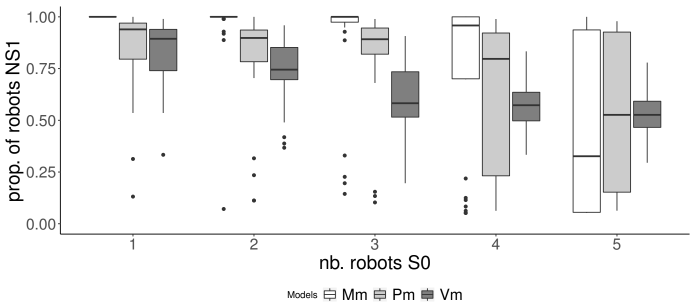


(a)

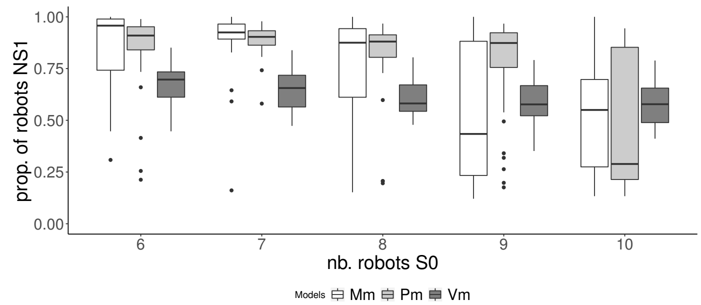


(b)

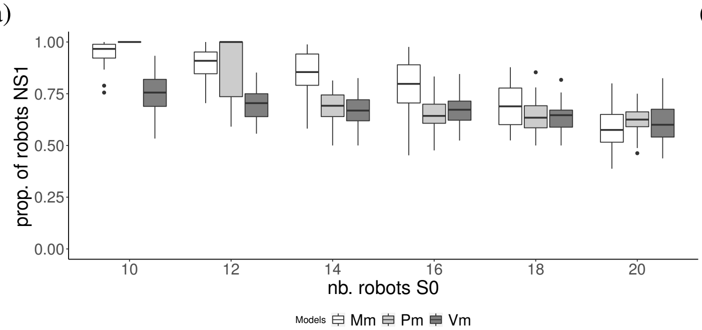


(c)

Fig. 2. Graphs showing the proportion of non-stubborn robots committed to opinion 1 (i.e., $\lfloor NS1/(NS1+NS0)\rfloor$) at the end of the simulation runs per experimental condition. The x-axes refers to the number of stubborn adversarial robots $S0$ committed to opinion 0. The dissemination process is regulated by the Majority Model in the white boxes, by the Probabilistic Model in the light grey boxes, and by the Voter Model in the dark grey boxes. Each box is made of 50 observations. Boxes represent the inter-quartile range of the data, while the horizontal line in the middle point of the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. The number of legitimate stubborn robots $S1$ committed to 1 is five in (a) and ten in (b), and twenty in (c).

$Pm$, the number of runs that terminates with a very large proportion of robots $NS1$ is roughly equal to the number of runs that terminate with a very large proportion of robots $NS0$. These two possible end results, that is a large dissemination of opinion 1 and a large dissemination of opinion 0 in the swarm tend to be equally likely and frequently observed. For $Vm$, the simulations very frequently terminate with a roughly equal proportion of robots $NS1$ and $NS0$. In summary, when the number of stubborn robots in the swarm is less than 10 for swarm of 100 robots, with legitimate robots $S1 = 5$ and adversarial robots $S0 \in [1, 2, 3, 4, 5]$, the Majority is the dissemination model that guarantees the largest dissemination of the opinion held by the majority of the stubborn robots (in our case, legitimate robots $S1$). For each of the five experimental conditions, the proportion of $NS1$ in swarms using the Majority model is significantly different from the proportion of $NS1$ in swarms using the Probabilistic and the Voter model (Mann-Whitney-Wilcoxon Test, $p < 0.001$).

Figure 2b refers to the set of simulations with legitimate robots $S1 = 10$. For all three dissemination models, we observe similar trends to the one illustrated in Fig 2a. Also in this set of simulations the median of the proportion of robots $NS1$ decays at a different rate for the three dissemination models, as supported by the statistical analysis. The most interesting difference from what observed in Figure 2a, is that the Probabilistic model performs closer to the Majority model (see Figure 2b white and light grey boxes), while the Voter model appears to be less reliable in supporting the cause of legitimate robots $S1$ even for the smallest $\lfloor S0/S1\rfloor$ ratio (i.e., $S0 = 6$ and $S1 = 10$). Interestingly, when $S0 = 9$, the Probabilistic model generates better results than the Majority model (see Figure 2b, for $S0 = 9$, white and light grey boxes). At 0.001 significant level (Mann-Whitney-Wilcoxon Test), we conclude that the proportion of $NS1$ for the Probabilistic and Majority model are non-identical populations. Moreover, differently from the previous set of experimental conditions with $S1 = 5$, when $S0 = S1 = 10$ the Majority and the Probabilistic models appear to generate results closer to the one generated by the Voter models. That is, the dissemination process ends more frequently with swarms in which the two opinions are roughly equally represented rather than swarms in which one opinion strongly dominates over the other (see Figure 2b, for $S0 = 10$).

In addition to these two set of experimental conditions, we have also run simulations in which $S1 = 20$, and $S0$ is set to $S0 \in [10, 12, 14, 16, 18, 20]$. That is, in this set of simulations the total number of stubborn robots varies from 30% to 40% of the entire swarm. The results of these tests are shown in Figure 2c. Looking at the graph, we see that both the Majority and the Probabilistic models are very effective in generating a consensus for opinion 1 when the $\lfloor S0/S1\rfloor$ ratio is $\lfloor 10/20\rfloor$ and $\lfloor 12/20\rfloor$. When $S0 = 14$, we observe a clear divergence between the results of the Majority and those of the Probabilistic model. The later model performs similarly to the Voter model, with a large number of trials ending with swarms that tend to have a roughly equal numbers of non-stubborn robots committed to $1$ and to $0$. It seems to us that, for $Pm$ and $Vm$ this particular final state, in which both opinions are roughly equally represented is the only observed end state of the dissemination process when the total number of stubborn robots in the swarm is bigger than a given threshold. This effect was already observed in [10] in a slightly different settings where dissemination times of the two options was unequal. However, subsequent work [3] has shown that this effect increases as the difference between the differential dissemination times between the two options disappears. Therefore, such an effect is particularly expect in the settings of this experimental work where there is no difference between the dissemination times of the two options. Contrary to the Probabilistic and the Voter models, the Majority model appears to be less affected by this phenomenon. Indeed, the graph in Figure 2c indicates that, for the Majority model (see Figure 2c, white boxes) the median of the proportion of non-stubborn robots committed to $1$ tends to decreases without observable drops, reaching the point $0.5$ when the number of legitimate robots is equal to the number of adversarial robots.

## IV. Conclusions

We have shown the results of a comparative study aimed at evaluating, in a simulated swarm robotics scenario, the effectiveness of an opinion dissemination model, referred to as Probabilistic, originally illustrated in [11]. This model has been compared with the Majority and Voter model in a data communication manipulation scenario, in which legitimate robots disseminate the "correct" opinion and a fewer number of adversarial robots disseminate an "invalid" opinion. Our simulations show that, within this simple data communication manipulation scenario, the Majority model is a better dissemination model than the Probabilistic model. The Majority model better supports the case of legitimate agents and more effectively limits the influence of adversarial agents in multiple experimental conditions characterised by a different number of stubborn robots in the swarm and by a different number of legitimate and adversarial robots. The results of our simulations allowed us to notice a phenomenon which was not observed in [11]. That is, the effectiveness of the Probabilistic vanishes suddenly as soon as the proportion of stubborn robots in the swarm reaches a certain threshold, regardless of the ratio of legitimate versus adversarial agents. This phenomenon, also discussed in [3], [10], does not seem to concern the Majority model in the same way. However, further research on this is needed.

Since, in various experimental conditions, the results with the Probabilistic model are close to those of the Majority model, and in one condition even better (see Figure 2b, for $S0 = 9$), we can not exclude that with a different setting of the model's parameters (i.e., the parameters $z$ and $k$ which regulate the increment/decrement of the individual probability $p$ to change opinion) the Probabilistic model would perform better than what observed in our tests. To verify this hypothesis on a swarm robotics setup, a large number of simulations are required. Unfortunately, the use of a simulation environment like the one modelled with the simulator ARGoS, used in this study, may require a too large computational resources and/or computational time. Thus, for the future, we intend to rely on mathematical models similar to the one used in [6], to be able to evaluate and to predict, for a large range of parameters' value (including swarm size, total proportion of stubborn robots, and the ratio between legitimate versus adversarial agents), how the Probabilistic model would perform compared to the other dissemination models.

## References


[1] M. Bonani, V. Longchamp, S. Magnenat, P. Rétornaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada, "The marxbot, a miniature mobile robot opening new perspectives for the collective-robotic research," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010, pp. 4187–4193.

[2] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.

[3] G. de Masi and E. Ferrante, "Quality-dependent adaptation in a swarm of drones for environmental monitoring," in *2020 Advances in Science and Engineering Technology International Conferences (ASET)*, 2019, p. to appear.

[4] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch, "Mixing beliefs among interacting agents," *Advances in Complex Systems*, vol. 3, no. 1–4, pp. 87–98, 2000.

[5] N. Franks, S. Pratt, E. Mallon, N. Britton, and D. Sumpter, "Information flow, opinion polling and collective intelligence in house-hunting social insects," *Philos. Trans. R. Soc. Lond. B*, vol. 357, pp. 1567–1583, 2002.

[6] Y. Gillet, E. Ferrante, Z. Firat, and E. Tuci, "Guiding aggregation dynamics in a swarm of agents via informed individuals: an analytical model," in *Proc. of the Artifcial Life Conference*, H. Fellermann, J. Bacardit, A. Goñi-Moreno, and R. Füchslin, Eds. MIT Press, 2019, pp. 590–597.

[7] S. Kato and M. Jones, "An extended family of circular distributions related to wrapped cauchy distributions via brownian motion," *Bernoulli*, vol. 19, no. 1, pp. 154–171, 2013.

[8] M. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo, "Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making," *Swarm Intelligence*, vol. 5, pp. 305–327, 2011.

[9] C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L. Gambardella, and M. Dorigo, "ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems," *Swarm Intelligence*, vol. 6, no. 4, pp. 271–295, 2012.

[10] J. Prasetyo, G. de Masi, and E. Ferrante, "Collective decision making in dynamic environments," *Swarm Intelligence*, vol. 13, no. 3-4, pp. 217–243, 2019.

[11] G. Primiero, E. Tuci, J. Tagliabue, and E. Ferrante, "Swarm attack: A self-organized model to recover from malicious communication manipulation in a swarm of simple simulated agents," in *Proc. of the* $11^{th}$ *Int. Conf. on Swarm Intelligence*, M. Dorigo, M. Birattari, C. Blum, A. Christensen, A. Reina, and V. Trianni, Eds. Springer, 2018, pp. 213–224.

[12] V. Strobel, F. Castello, and M. Dorigo, "Managing byzantine robots via blockchain technology in a swarm robotics collective decision making scenario," in *Proc. of the* $17^{th}$ *Int. Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2018, pp. 541–549.

[13] V. Trianni, E. Tuci, K. Passino, and J. Marshall, "Swarm cognition: an interdisciplinary approach to the study of self-organising biological collectives," *Swarm Intelligence*, vol. 5, no. 1, pp. 3–18, 2015.

[14] 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*, vol. 4, p. 9, 2017.

[15] 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*, vol. 30, no. 3, pp. 553–580, 2016.

[16] G. Weisbuch, G. Deffuant, F. Amblard, and J. Nadal, "Meet, discuss, and segregate!" *Complexity*, vol. 7, no. 3, pp. 55–63, 2002.