A Macrophage-Inspired Approach To Fault Reaction In The Collective Perception Problem

Patrick Shortall[1], Chanelle Lee[2], and Jonathan Lawry[3]

University of Bristol, Bristol, UK
patrickshortall26@gmail.com[1], {c.l.lee[2], j.lawry[3]}@bristol.ac.uk

Abstract

The collective perception problem is commonly discussed in swarm robotics, with many proposed solutions. However, there has been less discussion on the impact of faulty agents on the efficacy of these decision making strategies, and few possible solutions to mitigate any negative effects. This paper introduces a decentralised, immuno-inspired 'check, track, mark' (CTM) routine, and tests its efficacy when used to mitigate the effect of faulty agents in the collective perception problem. The CTM routine is inspired by macrophages in the human immune system, and their use in preventing pathogens from infecting healthy cells. We test the routine using three previously established decision making strategies, and a model of a detection algorithm with saturating true-positive and false-positive rates. We find that the proposed approach improves the ability of agents to reach an accurate consensus in the presence of faulty agents across all three of the decision making strategies tested, with increases in accuracy between 15-213%. For one strategy, the CTM routine also allows for an accurate consensus to be reached in fewer timesteps, with a median decrease in time to consensus of 29%. Out of the parameters associated with the CTM routine, we find the interval between initial checks to be most significant in affecting the speed and accuracy of the group in reaching consensus.

Introduction

The need to reach consensus and make decisions is a fundamental requirement of multi-agent systems, with examples found in various domains including swarm robotics (Schranz et al., 2020), sociophysics (Galam, 2008), medicine (Yang et al., 2022), and behavioural biology (Franks et al., 2003). Group decision making and consensus formation can be split further into two categories; continuous decisions such as direction of travel, and discrete decisions (Valentini et al., 2017) such as an optimal nest site. The algorithms used to solve these problems find varying levels of success which increase or diminish as parameters and setups change. Key problems are often found in scalability (either as number of agents (Soma et al., 2023) or space of options (Lee et al., 2021) increases), dynamically changing environments (Prasetyo et al., 2018), and the presence of faulty or malicious agents (Dibaji and Ishii, 2015).

The collective perception problem is a discrete task whereby agents must explore an environment and reach a consensus on the environmental feature with the highest density (Morlino et al., 2010). In the collective perception problem, agents must explore the space and interact with neighbours in order to have a well-informed belief of the true state of the world. This can lead them to be vulnerable to faulty and malicious agents, as they may not have sufficient interactions with agents to be confident that they are faulty or malicious, and as a result, may be misinformed by their incorrect beliefs.

One approach is for agents to share their opinion on the reliability of other agents. However, this relies on agents either having some global knowledge of the number of total agents, or the ability to correctly identify and dynamically accommodate new agents in addition to sharing opinions on their reliability. As such, we employ the assumption of bounded rationality commonly used in collective social learning wherein agents exhibit rational behaviour according to the limits of their knowledge at the time (Zhan et al., 2022). This paper explores a decentralised, immuno-inspired 'check, track, mark' (CTM) approach to fault reaction, requiring no prior knowledge of group size, and assuming an imperfect detection algorithm. We refer to fault reaction in this context as the behaviours performed by agents *after* the detection of a faulty agent. In the CTM routine, non-faulty agents (henceforth named healthy agents) check neighbours for faults, track suspected faulty agents, and broadcast the unsuitability of interaction to neighbouring agents in order to mitigate the effect of faulty or malicious agents on group ability to reach consensus. The routine is inspired by previous work from Timmis et al. (2016), who in turn took inspiration from the formation of granulomas (groups of macrophages) in the human immune system, and their ability to isolate faulty cells from healthy cells. In our system, tracking agents can be likened to macrophages (key cells in granulomas) (Adams, 1976) creating a buffer between healthy agents and faulty agents, and neutralising the otherwise negative effect of faulty agents. The routine also takes inspiration from ant colonies, where ants often switch roles to benefit the greater need of the colony (Gordon, 1989).

The contributions of this paper are

- Evidence of the negative effect of faulty agents in the collective perception problem
- The introduction and exploration of a decentralised, immuno-inspired 'check, track, mark' (CTM) approach to fault reaction
- Evidence of the CTM routine's efficacy in mitigating the negative effect of faulty agents in the collective perception problem

Background

Faulty and malicious agents have been studied in many multi-agent models and systems; and can affect the system's ability to reach consensus with varying severity. So called 'dead' agents are identified by Szymanski et al. (2006) to be relatively harmless in their best-of-n shortest path finding experiments. They found that unresponsive agents were treated as any other obstacle in the maze, and so while possibly reducing the speed, did not diminish the accuracy of decision making. In contrast, the presence of locally bounded malicious agents in sampled-data double integrator multi-agent systems were found to cause divergence in position estimates by Dibaji and Ishii (2015), resulting in the system being unable to reach consensus in this continuous consensus problem. In social learning, Ntemos et al. (2021) explore truly malicious agents which deceive other agents by manipulating the likelihood functions used in the belief update process to launch 'inferential attacks'.

One approach to mitigating the effect of faulty agents is for healthy agents to be equipped with a detection method for use in testing their neighbours. If the neighbour is faulty, then the healthy agent can discard the neighbour's opinion, and possibly alert fellow healthy agents to the faulty agent. Millard et al. (2014) use a single embedded neural network classifier in autonomous healthy agents to detect the presence of faulty agents, and find that there is a trade-off between the number of false-positives and the latency of fault detection. The longer it takes to detect faulty agents, the more accurate the classifier, but the greater opportunity the faulty agent has to disrupt healthy agents. More generally, Miller and Gandhi (2021) found that fault detection methods invariably involve a trade-off in sensitivity between a high true-positive rate and a low false-positive rate. To reduce false-positives healthy agents often require more contextual data, and more observations of faulty agents.

Fault reaction is particularly important in decentralised systems, as it is likely only one or a few agents will detect a faulty agent at any given time, and it is unclear either how to effectively communicate this information to other healthy agents, or how to respond to the fault. In Timmis et al. (2016), a strategy of self-healing is employed, where a subset of healthy agents form a granuloma-like structure around detected faulty agents and repair the agent, allowing it to continue its part in the group task of reaching a fixed way-point. Timmis et al. (2016) also note the ability of macrophages to prevent 'cell to cell transmission'. In our paper we abstract this idea from macrophages blocking the transmission of pathogens from cell to cell, and follow a similar logic, using tracking agents to block the transmission of faulty opinions from agent to agent. Unlike in Timmis et al. (2016), where the act of repairing an agent requires multiple agents, and thus a granuloma-like structure is formed, for our purposes, a single agent likened more to a macrophage suffices.

Our paper focuses on fault reaction in the collective perception problem, which was first proposed formally by Morlino et al. (2010). It is a spatially explicit consensus formation task where a group of autonomous agents attempt to reach consensus on whether a tiled space is filled mostly with either white or black tiles. Each agent is able to move around and formulate a belief, based on their observations of the environment, and evidence received from neighbouring agents. The ability of the swarm to correctly perceive the dominant feature is measured by the speed and accuracy to which they reach consensus. Previous work has explored how this ability is affected by noisy sensors (Chin et al., 2023), non-binary feature options (Ebert et al., 2018), and the distribution of features (Bartashevich and Mostaghim, 2019).

Several strategies have been developed for the collective perception problem. For this paper we use the following three to test the effect of faulty agents and the efficacy of the proposed CTM solution.

1. Direct Modulation of Majority-based Decisions (DMMD) (Valentini et al., 2016)
2. Bayes Bots (Ebert et al., 2020)
3. Ev-ProdOp - Adapted from Lee et al. (2018)

DMMD is a simple, generalisable strategy proposed first in Valentini et al. (2016) and found to scale well. Bayes Bots and Ev-ProdOp are Bayesian strategies and both have been found to be successful in previous studies (Ebert et al., 2020) (Lee et al., 2018).

The remainder of this paper is outlined as follows. First, we introduce the experimental setup, including the basic setup of agents, the environment they move in, their goal, and the decision making strategies they utilise to achieve their goal. We then cover the imperfect detection model which healthy agents use to detect faulty agents, and introduce the novel CTM routine. After this, we show and discuss our simulation results, including the negative effect of faulty agents, and the extent to which the CTM routine mitigates faulty agents. Finally we present our conclusion, with a summary of the paper, limitations of our work, and possible further work.

Experimental setup

For this paper, we base our experimental setup on the collective perception problem as formalised by Valentini et al. (2016). A group of agents are initialised at random positions in a 200x200 unit environment where equal-sized square tiles are coloured either in black or white. The task of the group is to reach the correct consensus on the dominant tile colour, which they attempt to achieve through each agent exploring the space, collecting measurements, and sharing observations with neighbouring agents (defined in this paper as agents within a 50 unit radius). We impose the constraint that agents are only able to communicate locally, and the group is fully decentralised.

The problem difficulty is defined and controlled by the ratio between the number of black and white tiles, more commonly called the fill ratio, denoted by f. The task becomes more difficult the closer to 1:1 this ratio is (i.e. as f tends to 0.5). Figure 1 shows an environment where $f = 0.25$, which is the fill ratio used for our experiments. This fill ratio makes it relatively easy for a population of agents to reach the correct consensus where no agents are faulty, but is challenging with just a small number of faulty agents

Assuming $f \neq 0.5$, for each agent there are two possible hypotheses; H_1, that the area is mostly filled with white tiles, and H_2, that the area is mostly filled with black tiles. In our setup, agents can be of three types: healthy, faulty, or tracking. At each timestep every agent updates its position, and healthy agents perform actions associated with the decision making strategy they are following. The purpose of tracking agents is to follow suspected faulty agents, and shield healthy agents from their outputs.

A healthy or faulty agent's opinion ϕ, is a measure of their belief in the correct hypothesis, and can be discrete or continuous depending on the decision making strategy. We take $\phi = P(H_1)$, and conversely $1 - \phi = P(H_2)$. In our experimental setup healthy agents update their opinion throughout the task, faulty agents have a fixed, incorrect opinion, and tracking agents do not have an opinion. It should be noted that faulty agents in other experimental setups could have a fixed correct opinion, or an opinion that updates randomly. These types of faulty agents are not covered in this paper, but will be explored in further work. We define consensus as reached when 90% of healthy agents have $\phi > 0.9$ or 90% of healthy agents have $\phi < 0.1$. This corresponds with 90% of healthy agents holding a strong belief in either H_1, or H_2.

Tracking agents move at each timestep by updating their velocity to the difference between their current position (i.e. location) and the faulty agent's position which they are tracking. All other agents perform a random walk: travelling straight for a random number of timesteps sampled from an exponential distribution with a mean of 40 timesteps, then turning instantaneously a random angle sampled from a uniform distribution. Each agent moves at a constant speed of two units per timestep. At the boundaries of the space agents change the sign of their velocity. Each agent is modelled as a point particle and hence there is no collision avoidance motion.

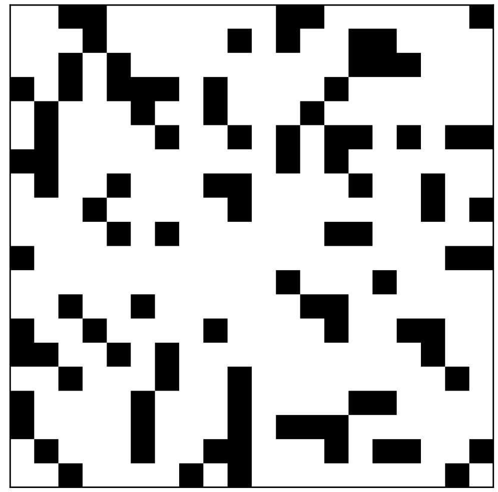

Figure 1: Collective perception environment comprised of a 200x200 unit grid of square tiles. The tiles are coloured either black or white, with a fill ratio $f = 0.25$, meaning 25% of the tiles are coloured black, the remainder are coloured white.

The algorithms for the three decision making strategies followed by healthy agents are laid out below, along with how opinion updating is affected by faulty agents.

Direct Modulation of Majority-based Decisions (DMMD)

Opinion for this decision making strategy is binary, either $\phi = 1$, or $\phi = 0$. For more information on the derivation of this strategy see Valentini et al. (2016).

Half of the healthy agents are initialised with a belief in H_1 ($\phi = 1$), half are initialised with a belief in H_2 ($\phi = 0$), and all begin in the 'exploration state'. In the exploration state, for some number of timesteps, sampled from an exponential distribution with mean τ_e (set to 30 in our experiments), the healthy agent perceives the environment. At each timestep, they record when they perceive evidence for their current belief. If their belief is in H_1, this is when on a white tile, if their belief is in H_2, this is when on a black tile. The resulting option quality υ, is then given by the number of timesteps where the agent perceives evidence for their current belief, divided by the total number of timesteps spent in the exploration state.

The healthy agent then enters the dissemination state. In this state, they broadcast their opinion to their neighbours for some number of timesteps, sampled from an exponential distribution with mean $\upsilon\tau_p$ (τ_p set to 90 in our experiments). The healthy agent also receives the broadcast of other agents in the dissemination state, including faulty agents who consistently broadcast the opinion corresponding to the incorrect belief. At the end of the dissemination state, the healthy agent assumes whichever opinion was broadcast to it more frequently. If the healthy agent receives an equal number of broadcasts of both opinions, the healthy agent retains its current opinion. The agent then re-enters the exploration state. Note that agents in the exploration state do not use any information received from agents in the dissemination state.

Evidential updating and ProdOp (Ev-ProdOp)

Opinion for this decision making strategy is continuous, $\phi \in (0,1)$. However, after each opinion update in the simulation, to avoid rounding errors, if an agent has an opinion $\phi > 0.95$, this is updated to $\phi = 0.95$. Similarly, if an agent has an opinion $\phi < 0.05$, this is updated to $\phi = 0.05$.

Each healthy agent is initialised with an opinion $\phi = 0.5$, and begins to collect evidence from their environment. Every τ_e (set to 30 in our experiments) timesteps, healthy agents update their opinion based on this evidence. If the number of timesteps an agent perceived a white tile directly below them is greater than the number of timesteps an agent perceived a black tile, then the agent takes this as evidence for H_1, denoted by E_1. If the number of timesteps perceiving a white tile is less than the number of timesteps perceiving a black tile, the agent takes this as evidence for H_2, denoted by E_2. The evidence is then used to update the agent's opinion using the following equation,

$$\begin{aligned}\phi \mid E_i &= P(H_i \mid E_i) \\ &= \frac{P(E_i \mid H_i)P(H_i)}{P(E_i \mid H_i)P(H_i) + P(E_i \mid \neg H_i)P(\neg H_i)}.\end{aligned}$$

Likelihood functions are defined as follows, with ω (set to 0.2 in our experiments) being a parameter which quantifies the trust the healthy agent has in the evidence:

$$P(E \mid H_i) = \begin{cases} 1-\omega & : E = E_i, \\ \omega & : E = \neg E_i. \end{cases}$$

In the case where the number of timesteps perceiving a white tile is equal to the number of timesteps perceiving a black tile, the agent does not update its opinion.

The healthy agent also pools opinions every τ_p (set to 20 in our experiments) timesteps with neighbouring agents. Pooling is performed using ProdOp as follows, where k is the number of neighbouring agents plus the healthy agent:

$$\phi_{t+1} = \frac{\prod_{i=1}^{k} \phi_{i,t}}{\prod_{i=1}^{k} \phi_{i,t} + \prod_{i=1}^{k} (1-\phi_{i,t})}.$$

Faulty agents have a fixed, faulty opinion of either $\phi = 0.05$ if H_1 is the correct hypothesis, or $\phi = 0.95$ if H_2 is the correct hypothesis. When a healthy agent includes a faulty agent's opinion in pooling, their opinion is skewed towards the incorrect hypothesis.

Bayes Bots

This decision making strategy allows three possible opinions, either $\phi = 1$, $\phi = 0.5$, or $\phi = 0$. For more information on the derivation of this strategy see Ebert et al. (2020).

Each healthy agent begins with $\phi = 0.5$, and models the unknown fill ratio as a beta distribution with parameters α and β. They are initialised with an equal number of observations $\alpha_0 > 0$ (set to 2 in our experiments) for α and β which controls the degree of uncertainty each agent begins with. The healthy agent then makes an observation of the tile they are currently on every τ_e (set to 30 in our experiments) timesteps. If observing a white tile, they add one to their α count, for a black tile they add one to their β count.

Every τ_p (set to 20 in our experiments) timesteps, the healthy agent also adds one to either their α or β count based on the last observations of their neighbours. Faulty agents broadcast an observation of whichever tile is the least frequent. Finally, at each timestep they test whether their decision criteria is met by evaluating p, the cumulative distribution function of their beta distribution at 0.5. If p or $1-p$ is above a certain threshold p_c (set to 0.9 in our experiments) then the healthy agent updates its opinion on which hypothesis it believes to be correct. If $p > p_c$ they form an opinion $\phi = 0$ of belief in H_2, if $1-p > p_c$ they form an opinion $\phi = 1$ of belief in H_1. Once healthy agents form an opinion they broadcast this to other agents in place of their last observation. This is the version of the algorithm proposed by Ebert et al. (2020) which makes use of positive feedback.

Imperfect Detection

We assume for this paper that each healthy agent is equipped with some imperfect detection algorithm, which can be characterised by the chance of producing a true-positive ρ_a, and the chance of producing a false-positive ρ_b. In other words, from the point of view of a healthy agent assessing a neighbouring agent:

$$\rho_a = P(\text{detect as faulty} \mid \text{is faulty}),$$

$$\rho_b = P(\text{detect as faulty} \mid \text{is healthy}).$$

Assuming that the performance of the detection algorithm improves as a healthy agent gathers more evidence, we can model the value of ρ_a, ρ_b as the output of some increasing or decreasing function of t_d (timesteps since detection) respectively. These rates vary as a function of timesteps spent by the healthy agent in close proximity to the suspected faulty agent. We model the true-positive rate as below, where $\rho_{a,0}$ is the initial true-positive rate based on an initial reaction between a healthy agent and a faulty agent, and μ is a learning rate constant which quantifies the effect of the healthy agent gathering more evidence on the suspected faulty agent:

$$\rho_a(t_d) = \frac{\rho_{a,0} + \mu t_d}{1 + \mu t_d}.$$

The false-positive rate is similarly modelled below, where $\rho_{b,0}$ is the initial false-positive rate:

$$\rho_b(t_d) = \frac{\rho_{b,0}}{1 + \mu t_d}.$$

The rates increase quickly from their initial value at first, before slowing and tending towards 1 as $t \to \infty$. The saturating shape of these functions is partially based on work in Millard et al. (2014), where empirical observation of the rates from their detection algorithm follow a similar shape. More sophisticated models might take into account noise, as would affect detection in embodied robots. We also assume for this model that healthy agents increase their detection ability only through proximity to the faulty agent they are currently tracking. As a result, the initial true and false-positive rates remain constant throughout the simulation. These simplifications allow for high-level modelling of a detection algorithm, and simplify the parameter space. This is as the focus of this paper is on fault reaction, not fault detection. Future work will explore this limitation.

CTM Mitigation

CTM is a high-level immuno-inspired control routine performed by healthy and tracking agents in the simulation to mitigate the effect of faulty agents. It is decentralised as only local communication is necessary.

The initial check part of the routine occurs every τ_c timesteps, defining τ_c as the interval between healthy agents checking their neighbours for faults. If the agent being checked is faulty, there is a $\rho_{a,0}$ chance of the healthy agent detecting it as faulty. If the agent being checked is healthy, there is a $\rho_{b,0}$ chance of the healthy agent detecting it as faulty. In either case, if the healthy agent identifies its neighbour as faulty, they switch their role from being a healthy agent to a tracking agent. Note that healthy agents do not check neighbours if a tracking agent is present, and if two healthy agents instantaneously detect an agent as faulty, only one of these becomes a tracking agent.

As a tracking agent, the agent stops its random walk motion. Instead, it updates its velocity to remain close to the detected agent as detailed in Experimental Setup. Whilst doing this, the tracking agent informs neighbouring healthy agents that they are in close proximity to a suspected faulty agent. Healthy agents do not update their opinion based on any evidence from their neighbours if at least one neighbour is a tracking agent. In this way, the tracking agent acts as a macrophage, effectively engulfing the faulty agent and preventing its faulty opinion from spreading.

The tracking agent will also recheck the agent they are tracking every τ_r timesteps. If the suspected agent is in fact faulty, the tracking agent has a $\rho_a(t_d)$ chance of sticking with its original assessment of faulty. If the suspected agent is healthy, the tracking agent has a $1 - \rho_b(t_d)$ chance of switching back to a healthy agent, and following the routines of a healthy agent.

Simulation Results

In our results, there are nine qualitatively different experimental setups. These are the combination of each of the three decision making strategies, and three distinct faulty agent scenarios: no faulty, where each agent in the simulation is healthy; unmitigated faulty, where a subset of the population is faulty and unmitigated against; and mitigated faulty, where a subset of the population is faulty, and healthy agents attempt to mitigate faulty agents by performing the CTM routine. For each setup 1000 simulations are run. As in Ebert et al. (2020); Valentini et al. (2016), we focus on time to consensus (speed) and accuracy as the key metrics to measure performance. In simulations with faulty agents, they comprise 10% of the total population. Each simulation is run for 1000 timesteps, or until the group reaches a consensus, whichever is sooner. Note that with CTM mitigation, the number of healthy agents varies due to agent switching. Consequently, the population in which consensus is measured over changes throughout the simulation. Parameter values not yet given are $\tau_c = 20$, $\tau_r = 10$, $\rho_{a,0} = 0.88$, $\rho_{b,0} = 0.03$, and $\mu = 0.05$.

Accuracy scores	DMMD	Ev-ProdOp	Bayes Bots
No faulty	0.90	1	0.98
Mitigated faulty	0.58	1	0.99
Unmitigated faulty	0.32	0.31	0.86

Table 1: Accuracy scores over 1000 simulations for each combination of decision making strategy and faulty agent scenario.

Table 1 shows the accuracy scores for each setup, where as in Valentini et al. (2016), accuracy is defined as the fraction of simulations in which the group of agents reach the correct decision. For each decision making strategy, accuracy scores are above 0.9 in the no faulty agent scenario. The introduction of faulty agents in the unmitigated scenario reduces the accuracy. This drop is most pronounced for Ev-ProdOp - from 1 to 0.31 - and DMMD - from 0.90 to 0.32. Bayes Bots is more robust to faulty agents, as the accuracy drops from 0.98 to 0.86. In each case, the introduction of CTM mitigation then improves the accuracy scores. For DMMD this is an increase from 0.32 to 0.58, and 0.31 to

1 for Ev-ProdOp. For Bayes Bots, the accuracy score rises above that of the no faulty scenario, from 0.86 to 0.99. This is likely due to the CTM mitigation slowing down the sharing of information between agents, which is beneficial to the group in simulation setups where initial conditions unfairly bias agents to the incorrect consensus. This hypothesis is supported by the following time to consensus results. For the time to consensus results, p-values measured in paired t-tests between results samples are given in brackets adjacent to median values.

As can be seen from Figure 2, the introduction of faulty agents increases the time to consensus across agent groups for all three decision making strategies. The median time to consensus more than triples for Ev-ProdOp: from 40 to 140 ($p = 3.7 \times 10^{-157}$), and almost doubles for Bayes Bots: from 100 to 180 ($p = 1.7 \times 10^{-42}$). For DMMD, the median increases from 268 to 393 ($p = 1.8 \times 10^{-70}$). The introduction of CTM mitigation increases the median time to consensus further for Bayes Bots and DMMD: from 180 to 220 ($p = 9.3 \times 10^{-9}$) and from 393 to 448 ($p = 2.4 \times 10^{-13}$) respectively. Interestingly, the introduction of CTM mitigation reduces the median time to consensus for Ev-ProdOp: from 140 to 100 ($p = 8.5 \times 10^{-39}$). Mitigation also reduces the spread of results for Ev-ProdOp. In contrast, the spread of results increases for Bayes Bots and DMMD.

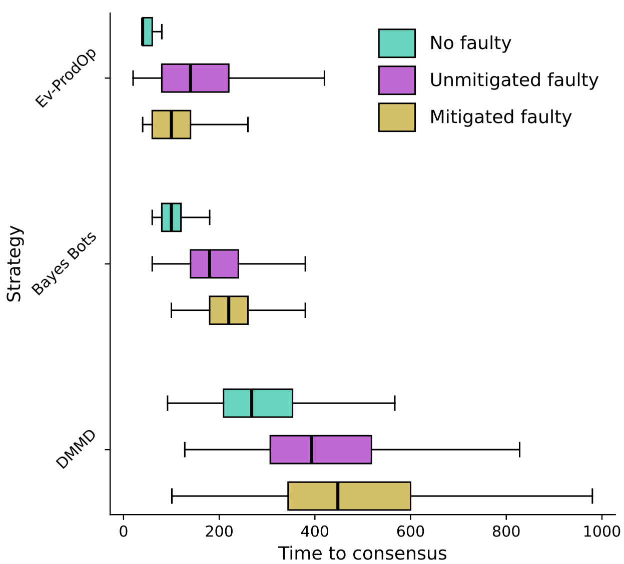

Figure 2: Boxplots showing time to consensus results across each of the nine experimental setups, grouped by decision making strategy. For each group of three, faulty agent scenarios are no faulty, unmitigated faulty, and mitigated faulty (top, middle, bottom)

To gain further insight into CTM mitgation we ran parameter sweeps across both the mitigation algorithm parameters, and the parameters of the modelled detection algorithm. For CTM mitigation, the parameters are the faulty search interval τ_c and the recheck interval τ_r. For the modelled detection algorithm, we vary the learning rate μ, and sensitivity, which is defined through combinations of initial true-positive and initial false-positive values. A sensitive algorithm is more likely to detect agents as faulty, with a higher true-positive, but also false-positive rate.

Figure 3 shows the results of these parameter sweeps, with accuracy plotted across the top, and boxplots below showing the spread in time to consensus. Values from a paired t-test between parameter variations are given in Table 2. Decreasing the faulty search interval increases the accuracy for each decision making strategy, but affects the time to consensus differently. For Ev-ProdOp, decreasing the interval reduces time to consensus. For Bayes Bots and DMMD, decreasing the interval appears to have little effect on the time to consensus. Decreasing the recheck interval has less of a discernible effect across each strategy. There is a decrease in accuracy for DMMD, however no significant change in time to consensus. For Bayes Bots, there is a slight decrease in time to consensus. Increased sensitivity results in a slight increase in time to consensus for Bayes Bots, but has no other significant effect on the other strategies. Finally, changing the learning rate has no statistically significant effect on any of the decision making strategies, other than a slight increase in accuracy for DMMD as the learning rate increases.

Conclusions

In this paper we have introduced CTM: a novel immuno-inspired agent-based routine for mitigating the effect of faulty agents in the collective perception problem. The routine is simple, tunable, and decentralised. In addition, we have also provided evidence of the negative effect of faulty agents in the collective perception problem, how they reduce the accuracy and speed of consensus.

Our results indicate that the efficacy of the routine varies when used in conjunction with different decision making strategies. When used with Bayes Bots and DMMD, there is a tradeoff between accuracy and speed. The routine increases the accuracy, but reduces the speed at which agents are able to reach consensus, and also the sensitivity to initial conditions with respect to the time to consensus. This tradeoff between speed and accuracy is well documented in many papers on algorithms used for solving the collective perception problem including Ebert et al. (2020). For these two decision making strategies, the CTM routine follows this trend. Interestingly, there is no such tradeoff with Ev-ProdOp. CTM mitigation improves accuracy, reduces time to consensus, and reduces the spread in time to consensus between simulation runs.

Varying parameters of both the CTM routine and the modelled detection algorithm, again produce different results between the three decision making strategies. For Bayes

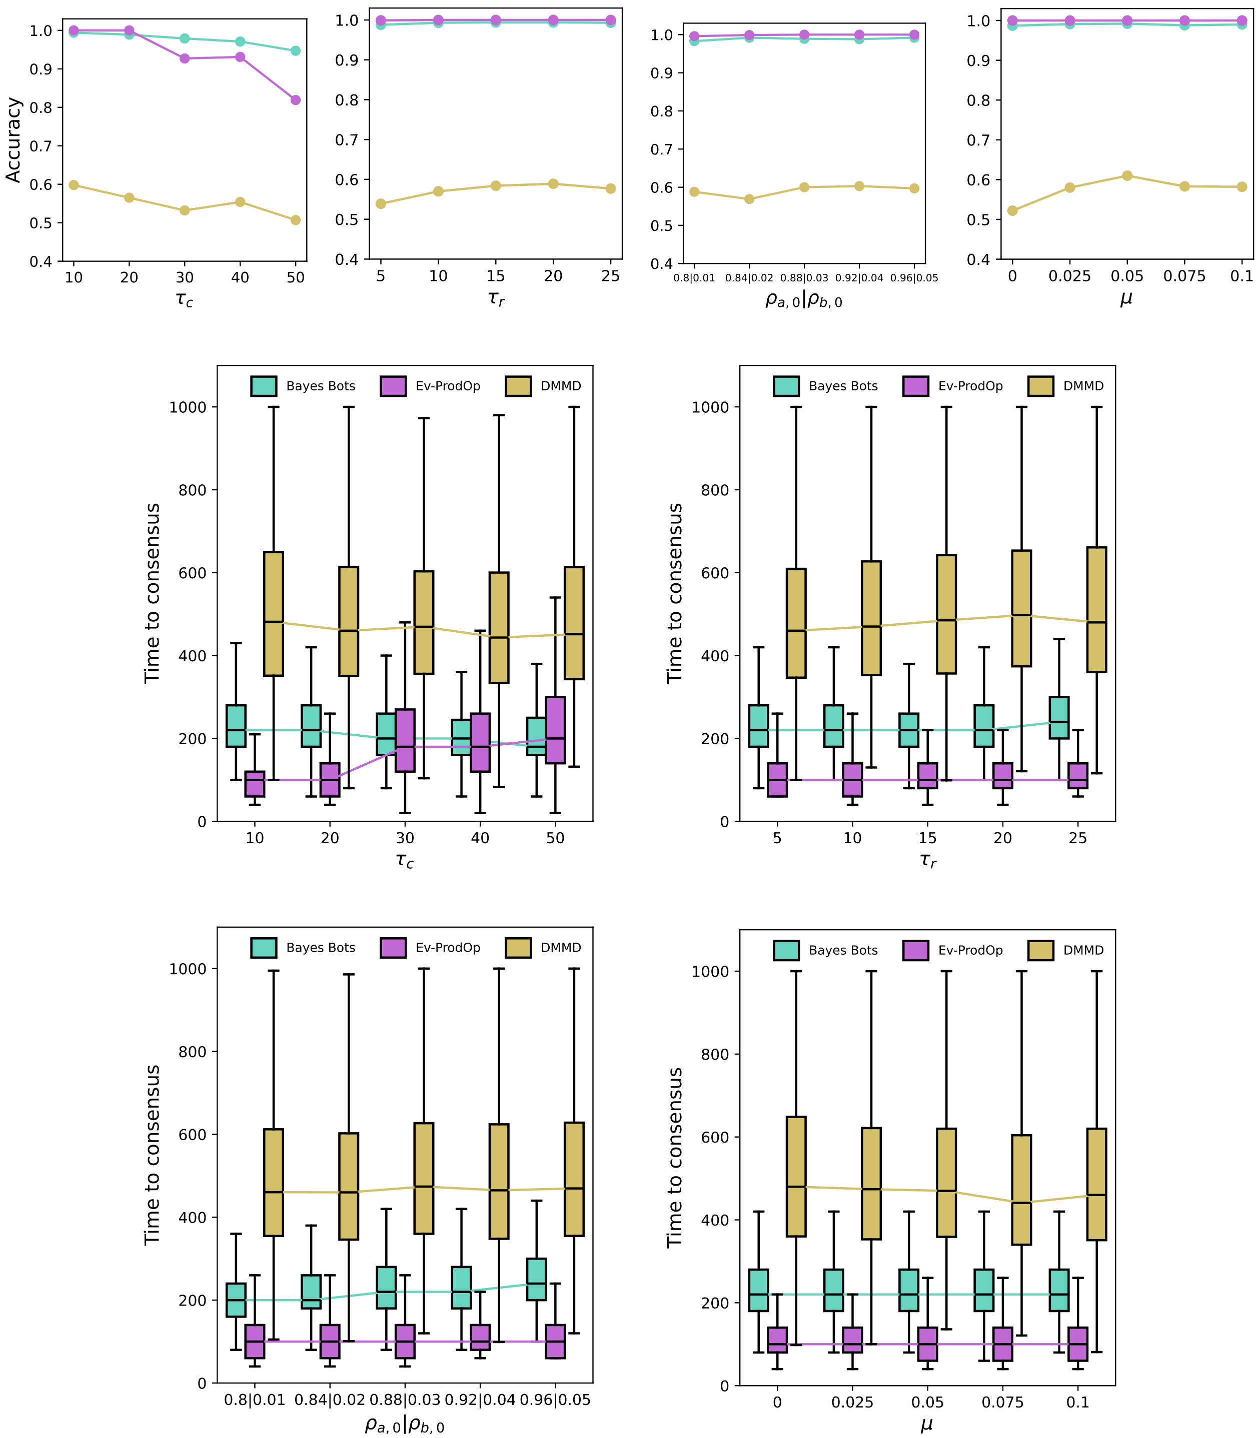

Figure 3: Parameter sweeps across faulty search intervals (τ_c), recheck intervals (τ_r), sensitivities ($\rho_{a,0}|\rho_{b,0}$), and learning rates (μ). Boxplots show the time to consensus results, and line graphs show accuracy results.

	τ_c				τ_r				$\rho_{a,0} \vert \rho_{b,0}$				μ			
	10, 20	20, 30	30, 40	40, 50	5, 10	10, 15	15, 20	20, 25	1	2	3	4	0, 0.025	0.025, 0.05	0.05, 0.075	0.075, 0.1
Bayes Bots	0.42	0.72	0.39	0.24	**2.5e-2**	0.59	0.83	**7.3e-9**	0.98	**9.5e-4**	0.11	5.3e-2	**4.1e-2**	0.59	0.81	0.95
Ev-ProdOp	0.11	**1.2e-10**	0.17	**5.5e-4**	0.76	0.61	**3.5e-3**	0.57	0.1	0.79	0.58	0.16	0.89	**3.2e-2**	0.93	0.36
DMMD	**7.4e-3**	0.50	0.11	0.17	0.06	0.30	0.14	0.47	0.75	0.12	0.58	0.87	0.13	0.88	0.26	0.44

Table 2: For each of the four parameter sweeps, five different values are tested across each of the three decision making strategies. This table shows the resultant p-values from paired t-tests between each consecutive set of results. For example, the p-value from comparing $\tau_c = 10$ and $\tau_c = 20$ Bayes Bots time to consensus results is 0.42. For space saving reasons, parameter pairs for sensitivity are simply referred to as 1, 2, 3 and 4. Significant results ($p < 0.05$) are highlighted in bold.

Bots and Ev-ProOp, only varying the faulty search interval produces an effect on accuracy. For DMMD, accuracy is more sensitive to parameter variations. Time to consensus is largely unaffected by any changes in parameters for DMMD. For Ev-ProdOp, varying the faulty search interval produces the most noticeable effect. Decreasing the interval reduces time to consensus. For Bayes Bots, time to consensus is affected most by the recheck interval, increasing as the interval increases.

A limitation of this work is the simplified modelling of a detection algorithm. Further work could include moving from a modelled detection algorithm to an actual detection algorithm, providing a fully realisable multi-agent system, and offering more opportunity to confirm how the sensitivity and learning rate of a detection algorithm affect the efficacy of the CTM routine. Another limitation is the simplified physics of the model. Experiments run with embodied robots or physics simulations would significantly aid in verifying the efficacy of the CTM routine. Finally, further work could also include testing and analysis on how the CTM routine reacts to larger swarm sizes, more challenging fill ratios, increased numbers of faulty agents, and alternative distribution patterns of white and black tiles.

Data sharing

The code used to generate all results in this paper can be found at `https://github.com/patrickshortall26/bon-sim`.

References

Adams, D. (1976). The granulomatous inflammatory response. a review. *The American journal of pathology*, 84(1):164.

Bartashevich, P. and Mostaghim, S. (2019). Benchmarking collective perception: New task difficulty metrics for collective decision-making. In *Progress in Artificial Intelligence: 19th EPIA Conference on Artificial Intelligence, EPIA 2019, Vila Real, Portugal, September 3–6, 2019, Proceedings, Part I 19*, pages 699–711. Springer.

Chin, K. Y., Khaluf, Y., and Pinciroli, C. (2023). Minimalistic collective perception with imperfect sensors. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8862–8868. IEEE.

Dibaji, S. M. and Ishii, H. (2015). Consensus of second-order multi-agent systems in the presence of locally bounded faults. *Systems & Control Letters*, 79:23–29.

Ebert, J. T., Gauci, M., Mallmann-Trenn, F., and Nagpal, R. (2020). Bayes bots: collective bayesian decision-making in decentralized robot swarms. In *2020 IEEE international conference on robotics and automation (ICRA)*, pages 7186–7192. IEEE.

Ebert, J. T., Gauci, M., and Nagpal, R. (2018). Multi-feature collective decision making in robot swarms. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1711–1719.

Franks, N. R., Mallon, E. B., Bray, H. E., Hamilton, M. J., and Mischler, T. C. (2003). Strategies for choosing between alternatives with different attributes: exemplified by house-hunting ants. *Animal behaviour*, 65(1):215–223.

Galam, S. (2008). Sociophysics: A review of galam models. *International Journal of Modern Physics C*, 19(03):409–440.

Gordon, D. M. (1989). Dynamics of task switching in harvester ants. *Animal Behaviour*, 38(2):194–204.

Lee, C., Lawry, J., and Winfield, A. (2018). Combining opinion pooling and evidential updating for multi-agent consensus. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 347–353. International Joint Conferences on Artificial Intelligence (IJCAI).

Lee, C., Lawry, J., and Winfield, A. F. (2021). Negative updating applied to the best-of-n problem with noisy qualities. *Swarm Intelligence*, 15(1):111–143.

Millard, A. G., Timmis, J., and Winfield, A. F. (2014). Run-time detection of faults in autonomous mobile robots based on the comparison of simulated and real robot behaviour. In *2014 IEEE/RSJ international conference on intelligent robots and systems*, pages 3720–3725. IEEE.

Miller, O. G. and Gandhi, V. (2021). A survey of modern exogenous fault detection and diagnosis methods for swarm robotics. *Journal of King Saud University-Engineering Sciences*, 33(1):43–53.

Morlino, G., Trianni, V., and Tuci, E. (2010). Collective perception in a swarm of autonomous robots. In *International Conference on Evolutionary Computation*, volume 2, pages 51–59. SciTePress.

Ntemos, K., Bordignon, V., Vlaski, S., and Sayed, A. H. (2021). Social learning under inferential attacks. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5479–5483. IEEE.

Prasetyo, J., De Masi, G., Ranjan, P., and Ferrante, E. (2018). The best-of-n problem with dynamic site qualities: Achieving adaptability with stubborn individuals. In *Swarm Intelligence: 11th International Conference, ANTS 2018, Rome, Italy, October 29–31, 2018, Proceedings 11*, pages 239–251. Springer.

Schranz, M., Umlauft, M., Sende, M., and Elmenreich, W. (2020). Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7:36.

Soma, K., Vardharajan, V. S., Hamann, H., and Beltrame, G. (2023). Congestion and scalability in robot swarms: A study on collective decision making. In *2023 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*, pages 199–206. IEEE.

Szymanski, M., Breitling, T., Seyfried, J., and Wörn, H. (2006). Distributed shortest-path finding by a micro-robot swarm. In *International Workshop on Ant Colony Optimization and Swarm Intelligence*, pages 404–411. Springer.

Timmis, J., Ismail, A. R., Bjerknes, J. D., and Winfield, A. F. (2016). An immune-inspired swarm aggregation algorithm for self-healing swarm robotic systems. *Biosystems*, 146:60–76.

Valentini, G., Brambilla, D., Hamann, H., and Dorigo, M. (2016). Collective perception of environmental features in a robot swarm. In *Swarm Intelligence: 10th International Conference, ANTS 2016, Brussels, Belgium, September 7-9, 2016, Proceedings 10*, pages 65–76. Springer.

Valentini, G., Ferrante, E., and Dorigo, M. (2017). The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9.

Yang, L., Jiang, J., Gao, X., Wang, Q., Dou, Q., and Zhang, L. (2022). Autonomous environment-adaptive microrobot swarm navigation enabled by deep learning-based real-time distribution planning. *Nature Machine Intelligence*, 4(5):480–493.

Zhan, M., Kou, G., Dong, Y., Chiclana, F., and Herrera-Viedma, E. (2022). Bounded confidence evolution of opinions and actions in social networks. *IEEE Transactions on Cybernetics*, 52(7):7017–7028.