

Hierarchical Communication Structures for Collective Decision Making in Robot Swarms

Kevin Wu

John Paul College, Whitworth Rd, Rotorua, Bay of Plenty, 3010, New Zealand; kevwu66@gmail.com

Mentor: Dr. Nisha Rawindaran

ABSTRACT: Swarm robots have potential in fields like search and rescue, space exploration, agriculture, and many other fields that are dangerous to humans or require large areas to be covered. The ability for a swarm to make decisions autonomously without a central controller plays a strong part in making them viable to implement in real-world situations. This research explores the potential for a hierarchical control method to be applied to the decision-making process by comparing it to a benchmark of a fully decentralized voter model-based method. This has been shown to improve performance in other areas of managing swarm robot behavior, which motivates this study into its application in collective decision-making. Using the Webots simulator, both of these methods were implemented using E-puck robots and compared to each other under the same conditions. The results showed that the hierarchical control method did not improve decision-making speed in the swarm measurably, but may have a beneficial effect on accuracy, which disproves the hypothesis of this research. The small scale of the simulation due to limitations encountered does, however, increase the effect of random noise on the results, which leaves the problem open to further investigation. Such further research for this could include the use of larger swarm sizes, more optimized behavioral algorithms, and/or real-world simulations.

KEYWORDS: Robotics and Intelligent Machines, Control Theory, Swarm Robotics, Collective Decision Making, Collective Perception, Hierarchies.

■ Introduction

Swarm robotics is the area of research into how large numbers of robots can cooperate to autonomously solve problems and complete tasks without a central controller. The difficulty is that without overall knowledge of the swarm, each robot acts based only on interactions with its neighbors and the environment. The idea is that from a swarm of many simple robots, complex behavior and abilities can emerge. This draws inspiration from existing structures in nature, such as social insects like bees,¹ ants, schools of fish, and bird flocking. This leads to systems that are more adaptive and robust than single robots,² as they can split tasks, replace damaged members, and add or lose members.

Challenges faced regarding the control of robot swarms include formation control, navigation, and the problem this research focuses on: decision-making, in this case, choosing one option out of multiple. In swarm robotics, this is known as the best-of-n problem, deciding the best option out of n possible ones. Typically, this means the swarm needs to reach either full consensus or have a majority agree on one option, at a threshold decided by the designer. A well-known example of this is the issue of nest-selection, seen in honeybees, which has given inspiration to models of the problem in robot swarms.³ This problem involves deciding the best location for a colony to move to out of multiple options, which all have traits that decide their quality, such as location, size, vulnerability, etc. Similarly, robots may need to decide the best location to gather resources in space or prioritize areas to search for injured people in emergencies.

Another concept being researched in swarm robotics is that of hierarchical control, a hybrid between centralized and decentralized control. The weakness of having a central controller is that there is a single point of failure because the single entity controlling the entire swarm could be damaged or compromised somehow. Reliance on decisions from the central controller also leads to reduced autonomy and adaptability to the situation, which is valuable in applications such as space exploration. The decentralized approach in swarm robotics removes these issues at the cost of being much more complex to engineer due to the fact that overall swarm behavior is emergent. This means that said behavior cannot be directly dictated, but instead arises from the interactions between individual robots. Solutions are being created to this problem, such as the use of neural networks,^{4,5} or combining pre-designed behaviors that are then fine-tuned.⁶ Hierarchical control provides a compromise where some robots in the swarm hold higher-level knowledge or abilities and have control over the lower-level members of the swarm, which means overall swarm behavior can be more easily controlled, which can also increase performance. These robots are part of the swarm and can be replaced by any other robot in it, removing the vulnerability to losing them. However, there may be potential trade-offs in that the selection of these higher-level robots adds complexity to an individual robot's behavior, which might take memory or processing power, which is valuable in swarms of smaller robots. In systems where the hierarchy cannot be reconfigured, higher-level robots would still be points of failure in the system, which negates the benefit of decentralized behavior.

In this research, the primary contribution made to the field is the proposal that a hierarchical control system, which is shown to have improvements in other areas regarding the control of robot swarms, can be applied to collective decision-making as well. This could potentially increase speed, accuracy, or both. Literature has been gathered to support this proposal, and currently, the literature lacks research exploring the connection between hierarchical control and decision-making, which motivates the work done in the study here. It will be explored by simulating a collective perception problem, where a swarm of robots must decide on a dominant feature of the environment. If applying hierarchical control to collective decision-making proves effective, the added efficiency in decision-making would allow complex tasks to be solved faster by robot swarms, as well as making the design of decision-making behavior simpler by allowing us to avoid the difficulty of engineering emergent behavior.

Aims and Objectives:

The objective of this research is to prove or disprove the hypothesis that a hierarchical decision-making strategy will perform better than the fully decentralized voter model in a collective perception task, in terms of accuracy and decision-making time.

— H_0 : The hierarchical decision-making method has no impact on decision-making speed/accuracy.

— H_1 : The hierarchical decision-making method has an impact on decision-making speed/accuracy.

To achieve this, we look into the best-of-two problem where the swarm decides between two options through experimental simulations. The robots, simulated in the Webots simulator,⁷ will navigate a floor of black and white tiles and must reach consensus on which color makes up the majority of the tiles. It applies the concept of hierarchical control structures to this with the aim of improving performance in swarm decision-making. Based on the literature, it is hypothesized that the application of a hierarchical control structure will improve the performance in decision-making of a swarm of robots as compared to a fully decentralized model.

The performance of the strategies will be measured using the specific metrics of:

- time to reach consensus in the swarm
- proportion of robots reaching the correct conclusion.

The constraints of this research include:

- Resources - the work required of this research, such as simulation design, implementation, and data analysis, falls within the ability of the researcher and the available resources.
- Relevance - This work fills in the observed gap in current swarm robotics research between hierarchical control methods and decision-making.

The specific contribution that this research aims to make to the field is to find whether it is possible to reduce the trade-off of speed and accuracy in collective decision-making in robot swarms by using the hierarchical decision-making method to increase both. This would allow both aspects to be improved and have neither one be sacrificed much.

■ Literature Review

This section will present work done in both collective decision-making and hierarchical control structures, showing some different proposed methods and related findings that inspire this paper. One prominent method of collective decision-making is known as the voter model, where a robot switches to the opinion of a randomly selected neighbor. Valentini *et al.* introduced the weighted voter model,⁸ inspired by the decision-making behavior of honeybees, where the group's opinion converges on the option of the highest quality by modulating the time of a robot's opinion dissemination based on its assessment of quality using environment-sourced information. This method of decision-making in a robot swarm is referred to as direct modulation. They found that this method increases accuracy with system size and is robust to noise from the input of quality assessments.

Valentini *et al.* introduce another method utilizing direct modulation of the majority vote method of decision-making,⁹ where a robot switches its opinion by taking that held by the majority of multiple neighboring robots. They find that the number of neighboring robots affects decision-making, showing a trade-off where lower numbers have more accuracy but slower consensus times, and higher numbers are faster yet less accurate. The present research aims to reduce the trade-off between speed and accuracy through the introduction of the hierarchical communication structure, which may potentially increase one or both aspects without sacrificing the other.

Another method of decision-making, inspired by behavior in honeybees,¹ is cross-inhibition. In the cross-inhibition method, a robot samples the opinion of its neighbor, and if the opinion is different, it enters an uncommitted state where it holds no opinion. It then takes the opinion of a neighbor in the next sampling phase. Zakir *et al.* find that compared to the voter model, the cross-inhibition method is better at breaking deadlocks between options of similar quality due to its resistance to noise from self-sourced information.¹⁰ Resistance to noise is an important aspect of decision-making in order for the swarm to navigate real-life scenarios, due to the environment's imperfections and due to potential issues such as malfunctioning robots gathering incorrect sensor data or malicious robots inserted into the swarm. Antonic *et al.*¹¹ tested a swarm of heterogeneous robots - robots that used different methods of decision-making internally - and found that this swarm was more robust against zealots/stubborn robots that do not change their opinion. The robustness of a swarm's decision-making against noise, stubborn robots, or deadlocks using a hierarchical control structure is not examined in the present paper, but may be a direction to investigate in future research.

These methods of decision-making operate on the idea of minimizing computational demands, which is useful for robots with fewer resources, such as swarm robots that are produced to be simple but high in number. This is the area in which the research presented here focuses. However, some methods require higher resources and have achieved good results, some of which are known as Bayesian decision making.¹²⁻¹⁴ Bayesian decision making is the use of self-sourced and socially-sourced information to calculate a probability of one option being cor-

rect, upon which a decision is then made. Shan *et al.* compare their method of Bayesian decision making with direct comparison and the weighted voter model, finding that it outperforms both methods in speed, accuracy, and robustness at the cost of being more complex. As a clarification, direct comparison is a method of decision-making that involves the robots directly sharing opinion quality estimates rather than just the opinions, and updating their opinions based on what they receive.

One more area of research into controlling swarms is with the use of artificial intelligence, such as neural networks and reinforcement learning. Research into the application of AI is somewhat less relevant to our specific case, as it looks into human-designed behavioral systems due to the fact that the work primarily focuses on movement-based problems, such as formation control, whereas we focus on decision-making. Also, we are specifically dictating rules of behavior, while systems like neural networks control behavior through internal workings, gained through training, that are difficult to understand. It carries the benefits of being simpler to create, simply by training the artificial intelligence on a task, but has drawbacks such as struggling when presented with the differences between simulation and reality, as well as being a black box that we cannot reverse-engineer at higher levels of complexity. As some examples, Li *et al.* used a small bio-inspired neural network to create virtual forces for their fish-inspired collective escape behaviors.⁴ Kim *et al.* used neural networks to generate interaction rules that cause robots to create desired collective formation structures, and both of these studies found these methods to be effective.⁵

The present research focuses on the application of a hierarchical control structure to the process of collective decision-making. This concept is supported in the literature, as applying a hierarchical control structure has been shown to aid performance in other areas of swarm robotics. Varadharajan *et al.*¹⁵ found that using 'guide' robots that had somewhat higher abilities, such as spatial awareness, to oversee simpler 'worker' robots, a model radiation cleanup task could be completed with far fewer robots than using only the worker robots, while still retaining the scalability and lack of a single point of failure, which a centralized system does not have. Shankar *et al.* apply this hierarchical control to formation control and movement of a swarm,¹⁶ presenting a method of controlling a swarm's shape and movement using guide robots to control worker robots, allowing easier design of swarm behavior while again avoiding the pitfalls of centralized control. Robots do not have to be physically superior; however, they may simply have a level of control over surrounding robots, using a higher level of information and/or intelligence to guide the others. This means that any robot in the swarm could become and act as a leader over other robots. One applicable case would be that if an existing leader gets damaged, the swarm can reorganize and select a new one. Some papers present methods that allow swarms to reconfigure hierarchies freely. For example, Zhu *et al.* describe a system where robots running the same algorithm establish a hierarchy by attempting to recruit each other and reaching an agreement where one surrenders control to another in a level above it, leading to a "brain" robot at

the top.¹⁷ Zhang *et al.* provide rigorous, mathematically analyzable algorithms for formation control using hierarchical frameworks, focusing on merging and splitting of frameworks and robot departure.¹⁸ These approaches form hierarchies that can still be changed according to the task, eliminating the issue that there still exist certain points of failure when having fixed leader robots, and allowing greater adaptability to the situation. These works show promise for applying a hierarchical control structure to collective decision making and the best-of-n problem. The present literature review found a lack of research into this specific concept, providing the inspiration for the research presented here.

Trabattoni *et al.*¹⁹ make some progress in filling the gap in the literature through their work, where they present a method for the best-of-2 problem. They describe a group of robots that acts as a central control entity that holds some overall swarm information. Any robot in this central control formation is replaceable with another member of the swarm, again giving advantages that a central controller provides without disadvantages such as having a single point of failure. It was found to also improve consensus time and accuracy. The work of Trabattoni *et al.* provides strong support that our research will find an improvement in decision-making using a hierarchical control method. The difference between the work of Trabattoni *et al.* and the present research is the problem being observed - Trabattoni *et al.* have investigated a site-selection problem where the robots must travel to a site and then convene in a central location to disseminate opinions, whereas we will investigate a collective perception problem where the robots will explore the environment and decide on the dominant feature, remaining dispersed through said environment. This different problem requires a different approach to the hierarchical control strategy.

It is reasonable to assume, and literature has shown, that there are differences in the results achieved by swarm simulations and actual robots.²⁰⁻²¹ Often, approximations of robot swarms, such as mean-field models, are used, but there are inaccuracies, such as the assumption that the swarm has opinions uniformly distributed throughout. To ensure a more accurate conclusion, it has been decided to simulate E-puck robots in the WeBots simulator,⁷ so that physical interactions are taken into consideration. The next section outlines the methodology used in the present research.

■ Methods

1. Basis of Simulation:

The intention was to create a simulation that would give results allowing a fair comparison between decentralized and hierarchical decision-making. The limits of what can be achieved consist of the available hardware - a Windows laptop with 16 GB RAM - to run the simulation, and time constraints in development.

The simulated task used for this research is inspired by the literature and primarily based on that of Zakir *et al.*¹⁰ They introduce the situation where there is a grid of floor tiles that can randomly be one of two colors, and the robot swarm must decide which color is most dominant. The method of

decentralized decision-making used in the simulation is their direct-switching model with modulated dissemination time based on opinion quality, also referred to in other literature as Direct Modulation of Voter-Based Decisions (DMVD). In the study done by Zakir *et al.*, they used DMVD as a baseline to compare their proposed implementation of cross-inhibition. In our simulation, an implementation of DMVD is also used as a baseline, but cross-inhibition is replaced with the newly proposed hierarchical control system.

To allow a fair comparison between the two methods of decision-making, some variables must be controlled. Each randomly generated grid will be explored by both types of decision-making, using the same number of robots and the same starting positions between the two tests for each board. Testing on multiple randomly generated grids will allow random error to be accounted for in the case that certain specific scenarios have one method perform better than the other.

2. Problem Definition:

The general task that this research focuses on about collective decision-making is that of collective perception. Collective perception refers to the idea that a swarm of robots with relatively weak capabilities can aggregate their observations to gain a more accurate picture of the environment as a whole. The decision-making aspect of this is that when robots gather contradictory information, they must ultimately decide that one or the other is the true observation. Gathering a better overall understanding of the environment aids in further decision-making, such as those regarding what actions should be performed afterward. So, finding methods to increase the efficiency of this process aids may increase performance in various tasks where robot swarms are deployed, hence the proposal of the hierarchical decision-making method for improving the process. The test of the effectiveness of a hierarchical decision-making method will be performed through a specific task designed with the parameters as described next.

Ten E-puck robots (Figure 1), each with a ground sensor module, are simulated in the Webots robot simulator.

They are randomly spread over an 8 by 8 randomly generated grid of black and white tiles, each tile being a square with a 10 cm side length. There are always more white tiles than black tiles, at varying ratios. The robots can send and receive signals from each other within a range of 30 cm. Using only communication with neighboring robots and the ground sensor to detect the color of tiles, they must reach consensus on the color that covers the largest area. Performance is measured by the time taken for all robots to reach consensus, and the proportion of robots reaching the correct conclusion.

The parameters of the simulation given are based on some trial and error to fit a few conditions: Firstly, the range of communication cannot extend to robots across the entire area, as this would defeat the purpose of creating methods to allow information to propagate across the swarm. Secondly, the total floor area and tile size are sized so that the robots will be able to cover the entire area and gather enough different samples while keeping multiple other robots within their communication ranges. This stops them from becoming isolated from

the swarm or ending up in feedback loops with only one other neighbor. This is constrained by the number of robots that can be simulated, which is limited to 10, as the hardware available struggled with physics calculations at swarm sizes beyond that. Changing the parameters given could be done intentionally to see the effectiveness of the hierarchical decision-making model when faced with more extreme situations, such as robots losing contact with the swarm, but this is beyond the scope of our study.




Figure 1: The setup of e-pucks in the Webots simulator.

3.1. Implemented Robot Behavior - Voter Model:

This section describes the implementation of the behavior of the simulated E-puck robots when testing the decision-making performance of the voter model.

Each E-puck robot has a random start position, orientation, and opinion on the grid. Their behavior is dictated as follows (broken down simply in Figure 2).




Figure 2: The basic robot behavior.

The robots move continuously using an implementation of Levy flight with basic obstacle avoidance, allowing a robot better coverage of the overall environment. The specifics as programmed in our simulation are as follows:

At any point:

If an obstacle is detected on the right, turn left

If an obstacle is detected on the left, turn right

Otherwise

Drive forward at a random speed for 3 seconds

Turn at a random speed for 1.5 seconds

Random distances and angles are achieved by changing the speed rather than the time of movement in order to reduce the random error in this simulation. In the exploration phase, a sample of the floor color is taken every time the robot begins driving forward, so randomly changing the time taken to move introduces error into the number of measurements a robot can take before the 18-second exploration phase ends. The amount of time used was chosen to allow the robots to move from one end of the arena to the other if the random speed happens to be maximum. Similarly, the turning time is selected to allow the robot to make a full rotation if chance dictates that it rotates at maximum speed. With every measurement of the floor color, the stored opinion quality - the proportion of samples taken which affirm the opinion - is recalculated based on the opinion that the robot currently holds. For example, if the robot's opinion is that the floor is predominantly white, the opinion quality is calculated as (# of white samples) / (# of samples taken)

After the exploration phase, the robot goes into the dissemination phase. Samples are no longer taken. The robot sends out its opinion to all neighboring robots within the specified range, in our case, 30 cm, every 0.1 seconds. Due to limitations in the abilities of the available hardware and the Webots simulator itself, the accuracy of timing does vary. The length of time this dissemination period lasts varies proportionally to the calculated opinion quality, with a minimum time of 2 seconds and a maximum of 10 seconds.

Once dissemination is finished, the e-puck then goes into the collection phase, where it receives opinions broadcast by neighboring robots. This period lasts for three seconds, after which a new opinion is randomly selected from the collected opinions and is then evaluated in the exploration phase again. If no opinions are collected after seven seconds, the robot proceeds with its previously held opinion.

The selection of the time for each phase is based on trial and error to create more realistic conditions, but it is not perfect and could be optimized in future work. The exploration phase timing of 18 seconds allows the robot to gather enough samples while not giving it full coverage of the area. The timing of the dissemination and collection phases is limited so that a robot will not receive or send information to every other robot on the floor within one cycle, while still collecting and disseminating enough information to impact the opinions of the swarm.

Once 9 of these loops pass with the opinion remaining the same, it is assumed that consensus has been reached, and the robot stops all processes. The number of loops was selected in order to ensure that every robot has reached consensus, giving any isolated robot time to find the rest of the swarm, while trying to reduce wasted time added to the process.

3.2. Implemented Robot Behavior - Hierarchical Model:

This section describes the implementation of the behavior of the simulated E-puck robots when testing the decision-making performance of the hierarchical model.

The majority of robots carry the same behavior as described previously for the voter model, but 2 out of the 10 have a "high-

er status" in that they don't directly take samples of the floor color, instead redistributing opinions received from robots around them. Because every robot carries the same hardware, any can take on this role, so the higher-level robots have been randomly selected in the simulation. They are referred to in the simulation as "managers". In the case of this simulation, it is unlikely that the specific selection of managers has much influence on results. The number of managers chosen is such that there are enough to usually make contact with every member of the swarm, while not being such a large proportion of the swarm that not enough samples are being taken to accurately represent the area.

The manager robots perform the same Levy flight to randomly move around the area. Instead of going through phases like those seen in the voter model algorithm, they store all received opinion broadcasts from other robots in a buffer, and broadcast and discard an opinion from the buffer every 0.1 seconds. This means they act to increase communication between robots, which in theory will improve the accuracy of the swarm's decision-making.

When the buffer is empty for long enough, the robot will assume that the other robots have ended their decision-making processes due to reaching consensus, thus ending its own process.

Refer to Appendix A for access to project files, including code.

■ Results and Discussion

1. Results:

The simulation was run for both models for tasks of different complexity. Task complexity is based on how many of the tiles are black. This is because the more white tiles there are, the easier it is for the swarm to then converge on the correct decision that white is the predominant floor color. Tests were run using randomly generated 8x8 grids with 32, 31, and 20 black tiles, as well as control runs of pure black and pure white to confirm the decision-making behaviors actually worked. For each number of black tiles, 3 tests were run on 3 random grids with that number of black tiles. The events of the simulation were then output to console logs, and data to be analyzed was extracted with the aid of a Python script written by Copilot. Tests using 10 black tiles per grid were planned; however, limited time resources restricted this.




Figure 3: Proportion of robots with the correct opinion over time for the voter model, tested with 31 black tiles.




Figure 4: Proportion of robots with correct opinion over time for the hierarchical model, tested with 31 black tiles.

Figures 3 and 4 show the proportion of robots holding the correct opinion over time for both the decentralized voter-model-based strategy and the proposed hierarchical strategy. As demonstrated in the graphs, which were consistent through all trials, the hierarchical model did not reach consensus significantly faster than the decentralized voter model. It should be noted that Run 2 of the hierarchical tests on 31 black tiles does not have a visible line due to random chance, causing all the robots to begin with the same opinion, so it is an outlier. The time taken to finish the decision-making process was actually longer for the hierarchical system that was implemented in this study, due to the “manager” robots using a time-based method to determine that decision-making had ended.

Despite this slower convergence, the hierarchical strategy sometimes reached a slightly higher final proportion of correct decisions than the decentralized method. However, this improvement was small, and the difference varied between trials. Overall, both systems successfully guided most robots toward the correct color classification. The hierarchical strategy did not demonstrate a clear advantage in speed and produced only marginal changes in accuracy.

Figure 3 also shows that both models sometimes converged on the wrong option, which is black having the larger area. One possible reason for this is the stochastic nature of the decision-making process, but it cannot be confirmed in this study as the swarm size and number of trials are limited.

Raw extracted data can be accessed by referring to Appendix A.

2. Discussion:

The simulations conducted in this study show that the proposed hierarchical decision-making method does not produce a clear improvement over the fully decentralized voter-model-based approach. These findings differ from the original hypothesis and more closely match the null hypothesis, H_0 : The hierarchical decision-making method has no impact on decision-making speed/accuracy. They do not support the alternative hypothesis H_1 : The hierarchical decision-making method has an impact on decision-making speed/accuracy.

Although the hierarchical system may have provided a slight increase in accuracy, this effect was small and inconsistent across trials. More notably, the total time required for the swarm to

reach consensus increased. This appears to result from the behavior of the manager robots, which wait for confirmation of convergence before terminating their decision-making process. When managers are ignored when considering timing measurements, the difference is not significant.

This is not expected based on trends reported in the literature, where hierarchical influence has enhanced swarm performance in tasks such as formation control, navigation, or coordinated resource collection. Multiple factors can explain this difference. First, the collective perception problem studied here is more spatially distributed than tasks previously tested with hierarchical structures. While hierarchies can facilitate top-down control, they may offer less benefit when robots must independently gather environmental data and when no consistent global information can be shared reliably. Second, the small swarm size used in the simulations may have prevented hierarchical advantages from emerging. Prior work suggests that hierarchical or hybrid methods often show larger gains at bigger scales. Finally, the randomness inherent in swarm behavior, especially at this small scale, limits the statistical strength of the observed trends.

2.1. Limitations and Challenges:

Other limitations faced were the limited capabilities of the hardware used, which held back the size of the simulation that could be run. As mentioned previously, the size of our simulation may not have been large enough to remove variation due to noise in our results. This does mean that the behavior of the robot swarm could be different at larger scales and that comparisons between the two decision-making methods here may not hold. It's also acknowledged that other sources of error may have arisen from the methods used in this research, which could not be resolved due to time constraints; however, this research still proposes a novel concept to collective decision-making and makes early steps in implementing it, which may be improved in future research.

Given these considerations, the results should not be interpreted as evidence that hierarchical structures are ineffective for collective decision-making in general. Instead, they highlight that the specific hierarchical strategy implemented here may not have been optimal, and that the constraints of the simulation (limited scale, computational constraints, relatively simple manager behavior, and time available for development) may have masked the potential benefits of such an approach.

2.2. Future Work:

Future work to determine the effectiveness of a hierarchical decision-making structure should include larger swarm sizes (e.g., more than 100 robots) and optimized parameters for their behaviors to ensure accurate comparisons with other models. Comparisons should also be made to models already proven to be more effective than the voter model, such as cross-inhibition as described by Zakir *et al.*¹⁰ or the majority vote model discussed by Valentini *et al.*⁹ The decision-making structure should also be applied to other decision-making scenarios, such as site-selection or task assignment. Alternative/more optimal implementations of the same concept should

also be tested, as the constraints mentioned in this study mean that it may not have used a well-designed behavior algorithm. Improvements might be made on the way hierarchies are implemented, but also in the movement or communication of information across the swarm. Finally, tests with physical robots are necessary to determine actual effectiveness in real-world applications.

■ Conclusion

This study investigated whether a hierarchical communication structure could improve the accuracy or speed of collective decision-making in a swarm of simulated E-puck robots performing a best-of-two perception task. The results did not show a clear performance advantage over the decentralized voter-model baseline. This means that the null hypothesis that the implemented method has no effect is not rejected, and there is insufficient evidence to prove the alternative hypothesis - that the implemented method does have an effect. While accuracy may have improved slightly, consensus times increased due to manager robots delaying termination. These findings suggest that the particular hierarchical mechanism implemented here does not enhance performance under the tested conditions, especially at small swarm sizes.

The study does still show that hierarchical structures can be integrated into collective decision-making processes, which gives room for further development in the area. The support shown in the literature for implementations of hierarchical structures still shows promise for this concept. Hierarchical control systems still may be able to measurably improve decision-making speed, accuracy, and algorithm design if implemented in an optimal way. By writing this paper, this concept might be pushed out to more capable authors. Improvements in decision-making and collective perception are crucial to the viability of implementing swarm robotics systems into real-world applications such as space exploration, agriculture, and search-and-rescue by allowing robot swarms to more effectively choose the right strategies to perform their tasks without human intervention. Future work should include changes such as improved communication methods, larger swarm sizes, more optimized behavior algorithms, and comparison to other decision-making models.

■ Acknowledgments

Acknowledgments to Dr. Nisha Rawindaran, who helped guide the research process with her knowledge.

■ References

- Seeley, T. D.; Visscher, P. K.; Schlegel, T.; Hogan, P. M.; Franks, N. R.; Marshall, J. A. R. Stop signals provide cross-inhibition in collective Decision-Making by honeybee swarms. *Science* 2011, 335 (6064), 108–111. <https://doi.org/10.1126/science.1210361>.
- Dorigo, M.; Theraulaz, G.; Trianni, V. Swarm Robotics: Past, Present, and Future [Point of View]. *Proceedings of the IEEE* 2021, 109 (7), 1152–1165. <https://doi.org/10.1109/jproc.2021.3072740>.
- March-Pons, D.; Múgica, J.; Ferrero, E. E.; Miguel, M. C. Honeybee-like Collective Decision Making in a Kilobot Swarm. *Physical Review Research* 2024, 6 (3). <https://doi.org/10.1103/physrevresearch.6.033149>
- Li, J.; Yang, S. X. Intelligent collective escape of swarm robots based on a novel Fish-Inspired Self-Adaptive approach with neurodynamic models. *IEEE Transactions on Industrial Electronics* 2024, 71 (11), 14460–14469. <https://doi.org/10.1109/tie.2024.3363723>.
- Kim, D.; Lee, J.; Kim, H.-Y. Navigating the Swarm: Deep Neural Networks Command Emergent Behaviors. *arXiv (Cornell University)* 2024. <https://doi.org/10.48550/arxiv.2407.11330>.
- Francesca, G.; Brambilla, M.; Brutschy, A.; Trianni, V.; Birattari, M. AutoMoDe: A Novel Approach to the Automatic Design of Control Software for Robot Swarms. *Swarm Intelligence* 2014, 8 (2), 89–112. <https://doi.org/10.1007/s11721-014-0092-4>.
- Cyberbotics. *Webots: robot simulator*. Cyberbotics.com. <https://cyberbotics.com/>
- Valentini, G.; Hamann, H.; Dorigo, M. Self-organized collective decision making: the weighted voter model. *Adaptive Agents and Multi-Agent Systems* 2014, 1, 45–52. <https://doi.org/10.5555/2615731.2615742>.
- Valentini, G.; Ferrante, E.; Hamann, H.; Dorigo, M. Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems* 2015, 30 (3), 553–580. <https://doi.org/10.1007/s10458-015-9323-3>.
- Zakir, R.; Dorigo, M.; Reina, A. Robot swarms break decision deadlocks in collective perception through Cross-Inhibition. In *Lecture Notes in Computer Science*, 2022; pp 209–221. https://doi.org/10.1007/978-3-031-20176-9_17.
- Antonic, N.; Zakir, R.; Dorigo, M.; Reina, A. Collective Robustness of Heterogeneous Decision-Makers against Stubborn Individuals; 2024. <https://www.ifaamas.org/Proceedings/aamas2024/pdfs/p68.pdf> (accessed 2025-08-31).
- Siemensma, T.; Chiu, D.; Ramshanker, S.; Nagpal, R.; Haghighat, B. Collective Bayesian Decision-Making in a swarm of miniaturized robots for surface inspection. In *Lecture Notes in Computer Science*, 2024; pp 57–70. https://doi.org/10.1007/978-3-031-70932-6_5.
- Chiu, D.; Nagpal, R.; Haghighat, B. Optimization and Evaluation of a Multi-Robot Surface Inspection Task through Particle Swarm Optimization. 2024, 8996–9002. <https://doi.org/10.1109/icra57147.2024.10611661>.
- Shan, Q.; Mostaghim, S. Collective Decision Making in Swarm Robotics with Distributed Bayesian Hypothesis Testing. In *Lecture Notes in Computer Science*, 2020; pp 55–67. https://doi.org/10.1007/978-3-030-60376-2_5.
- Varadharajan, V. S.; Soma, K.; Dyanatkar, S.; Lajoie, P.-Y.; Beltrame, G. Hierarchies define the scalability of robot swarms. *arXiv (Cornell University)* 2024. <https://doi.org/10.48550/arxiv.2405.02417>.
- Shankar, V. V.; Dyanatkar, S.; Beltrame, G. Hierarchical Control of Smart Particle Swarms. *arXiv (Cornell University)* 2022. <https://doi.org/10.48550/arxiv.2204.07195>.
- Zhu, W.; Oguz, S.; Heinrich, M. K.; Allwright, M.; Wahby, M.; Christensen, A. L.; Garone, E.; Dorigo, M. Self-organizing nervous systems for robot swarms. *arXiv (Cornell University)* 2024. <https://doi.org/10.48550/arxiv.2401.13103>.
- Zhang, Y.; Oğuz, S.; Wang, S.; Garone, E.; Wang, X.; Dorigo, M.; Heinrich, M. K. Self-Reconfigurable hierarchical frameworks for formation control of robot swarms. *IEEE Transactions on Cybernetics* 2023, 54 (1), 87–100. <https://doi.org/10.1109/tcyb.2023.3237731>.

10. Trabattani, M.; Valentini, G.; Dorigo, M. Hybrid control of swarms for resource selection. In Lecture Notes in Computer Science, 2018; pp 57–70. https://doi.org/10.1007/978-3-030-00533-7_5.
11. Reina, A.; Bose, T.; Trianni, V.; Marshall, J. A. R. Effects of spatiality on Value-Sensitive decisions made by robot swarms. In Springer Proceedings in Advanced Robotics; 2018; pp 461–473. https://doi.org/10.1007/978-3-319-73008-0_32.
12. Trianni, V.; De Simone, D.; Reina, A.; Baronchelli, A. Emergence of Consensus in a Multi-Robot Network: From abstract models to Empirical validation. IEEE Robotics and Automation Letters 2016, 1 (1), 348–353. <https://doi.org/10.1109/lra.2016.2519537>.

■ Author

Kevin Wu is a student of John Paul College in Rotorua, New Zealand, who has a long-standing love for the field of robotics. He hopes to major in mechanical engineering or mechatronics to work on robots that interface with humans.

■ Appendices

Appendix A – Simulation Files:

All Webots simulation files, including code and extracted data, can be found in a GitHub repo at the URL:

https://github.com/kevwano/swarm-hierarchy/tree/main/emitter_receiver/controllers/robot1

While Webots itself may be installed at:

<https://cyberbotics.com/>

The project can be accessed by opening the “worlds” folder and opening the file “emitterandreceivertesting_ver2.wbt” in Webots. It is recommended that, for optimal performance, the graphics settings be turned to the lowest.