Swarm Robotics: Nature-Inspired Design and Error-Detection in Autonomous Behavior

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Taylor Kramer
Spring, 2020

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature ______________________ Date 05/07/2020
Taylor Kramer

Abstract

In this paper I present a comprehensive analysis of swarm robotics from its foundation in biological patterns to its future in society. Highly bio-inspired from existing insect communication patterns, robot swarms are characterized by large quantities of simple robots working together on a task in a decentralized control structure. This base in nature also serves as a starting point for problem solving through the analogy to nature, which prompts engineers and stakeholders alike to consider ways in which a problem they encounter could have already been solved in nature. This also encourages the formation of trading zones of knowledge regarding areas of subject matter expertise, and promotes the interdisciplinary nature of robot swarms.

I present this view as one of the many STS frameworks used to analyze the future of swarm robotics. As an emerging technology, swarms have the potential to become both a great solution and a great problem as well. The decentralized communication structure, along with the autonomy of swarms, presents many areas of vulnerability. The actions taken now, and in the near future, by engineers and stakeholders will determine the long-term future of swarm robotics. Whether or not swarms will pose a greater ethical risk is up to the anticipatory governance and ethical oversight conducted in parallel to research and development. Working together, we can all provide a safer future for ourselves and robotics.

Swarm Robotics: Nature-Inspired Design and Error-Detection in Autonomous Behavior

Robotics as a field of study has been interdisciplinary from the start. This feat of engineering requiring the interaction of multiple types of engineers and programmers has found success in forming innovative solutions to complex problem spaces. As research has evolved in this field, and the idea of automation has left the structured factory environment, various research areas within robotics have emerged with focus areas requiring the incorporation of even more types of scientists (Bauer et al., 2008). In order to better represent the cooperation necessary to many complex problems, the study of swarm intelligence emerged within robotics. This developing technology is characterized by decentralized and autonomous control, use of multiple communication channels, and homogeneity, all of which will be explored more in depth in the next section. This paper explores the unique background factors surrounding swarm behavior, and explores the ethical concerns and security steps necessary to implement a widespread adoption of this technology in the future.

The concept of swarm intelligence has garnered an immense amount of traction in recent years, chiefly influenced by the growing public interest in artificial intelligence, machine learning, and autonomous systems. Although an easily identifiable feature of swarms is the large number of robots involved, the particular definition of a robotic swarm is much more intricate in nature. Publications surrounding swarm robotics tend to vary in their specificity, with casual authors often portraying swarms as simply any system with a large number of robots. What this basic definition really describes instead is the concept of a multi-robot system, which operates in a centralized control structure. Multi-robot systems are most commonly used in transportation, sensing, or robotic team games such as sports (Tan & Zheng, 2013). A centralized control system for multi-robot systems means that all of the robots operate under a hierarchy of control, with smaller subgroups within the group usually taking on separate and specific roles. In such a way, the robots may collaborate together on a task more effectively than a singular robot could, but there remains a low level of flexibility or scalability within the system. The centralized nature also limits the group's ability to attack or malfunction, due to the fact that individual robots cannot easily replace each other's roles in the group. Thus, one defective robot may disable the functionality of the swarm as a whole, which is not ideal for risky environments and presents little opportunity for advanced autonomous behaviors.

Swarms differ from standard multi-robot systems in that they operate under decentralized control, which allows the swarm to function with a higher level of scalability and flexibility. In essence, a swarm can also be thought of as one small robot cloned a hundred or a thousand times, with all of the identical robots instructed to communicate with the clones nearest to themselves. This example represents the homogeneity of swarms.

Although the distinction between multi-robot systems and swarm robotics may seem small to a casual observer, it is important to maintain the strict boundaries around these definitions. As swarms become more prevalent in communication with the public, everyone should be able to know the difference between these technologies in order to limit the spread of misinformation moving forward.

As with any new technology, part of limiting the public's ethical concerns is giving them the proper information to proceed with. Potential trading zones of knowledge related to ethical concerns in technology increasingly need this proper definition structure in order to prevent miscommunication, such as was the case when Mark Zuckerburg recently addressed Congress. There was a scenario in which Zuckerburg was questioned on the revenue method for facebook, given that users are not charged a fee (*Zuckerberg Explains the Internet to Congress*, n.d.). This became a point of concern for Congress, given that the concept of ad revenue in conjunction with internet traffic had not properly been defined or translated prior to the interaction in a way which was easily absorbed by non-experts. This lack of communication between subject matter experts and the general public also exists with the emerging buzzwords of machine learning, artificial intelligence, and deep learning. All of these phrases convey something slightly different within the technology base, which should be properly defined in order to prevent misrepresentation of systems. In this way, scientists and the public alike can work together to address all proper ethical concerns to the specific technology, saving time that would’ve been wasted discussing characteristics that are not truly relevant to the technology’s case.

Going back to the discussion on robotics, trading zones have been instrumental in the creation of advanced robotics techniques. Usually when considering the concept of robots, nature is the last thing to come to mind. In a lot of applications, however, thinking in terms of nature is extremely beneficial. A commonly-recognized use of bio-inspired design in robotics is its use in locomotion. This area of robotics became a hot-topic before the implementation of robotic swarms, but contains a similar innovative mindset, and has been steadily developing in parallel with swarms since then. As robotics evolved past the predictable environments of an automation-friendly industrial environment, a variety of nature-inspired design innovations have emerged. Although wheeled robots are efficient in ideal factory environments, bio-inspired locomotion can have better results on rough, unpredictable terrain (Yanagida et al., 2017). This area of robotics is preferable in certain types of situations where a flexible range of motion is required for odd terrain which wheels or rotorcraft cannot easily traverse. This innovation was primarily influenced by the physical difficulties encountered in unknown environments.

Similarly, scientists and engineers worked together to discover an innovation in robotics inspired by natural group behavior dynamics of humans and animals, capitalizing on the benefits of communication in unpredictable problem spaces. In some scenarios – both robotic and naturally occurring – a task is more easily accomplished through a group effort. This multi-robot coordination effort with decentralization has been termed "swarm robotics" based upon the similar insect behavior in nature. In fact, swarm robotics is highly bio-inspired in this regard (Kolling et al., 2016). Since the cost of building simple robots has decreased in the past few years, creating simple robots at a large quantity level has gotten cheaper (*Swarm Robotics -- from Local Rules to Global Behaviors | Magnus Egerstedt | TEDxEmory*, n.d.). Simple robots on their own may not be able to accomplish much, but even in nature, simple creatures such as ants or bees are able to accomplish much as a group. This effect has been useful for the application of swarm robotics, since the strength of a swarm is in its numbers, rather than the individual's contribution. When it comes to swarm robotics, a new dimension of interdisciplinary collaboration becomes essential to successfully coordinate the behaviors and communication methods of a massive group of robots, where technology has aided the creation of swarm robotics in many ways.

In such a way, insect behavior experts have been able to go from studying the social interactions of ants to understanding how wasps build nests based on localized information in the environment rather than instructions directly communicated to them for a specific area and a global goal (Garnier et al., 2007). This idea itself naturally expanded to the work of computational biologists attempting to simulate virtual wasps building the same structures. This modelling and simulation meant to confirm an understanding learned from observation, but later served as an algorithmic inspiration for computer scientists and engineers. The resulting work

accomplishes a goal for the biologist, but also translates the work into an easily recognizable algorithmic format which is more easily understood and recognized by engineers.
Once there is an algorithmic or mathematical pattern observed in science, engineers and computer scientists will often implement fundamental principles learned from this pattern as an algorithmic advantage to solving a complex problem.

Advances in modelling and simulation techniques allow computational biologists to act as translators in the trading zone. In addition, the prevalence of research available online has facilitated the ability of an interdisciplinary trading zone to expand through time and space, serving as a path for engineers to access the biological translations. Scientists can conduct incremental research, slowly adding functionality and use cases to insights learned through prior work.

Although trading zones are important within the team for producing solutions, another set of influential knowledge-holders are the stakeholders. Users and stakeholders need to be included in the trading zone as well, providing their knowledge of needs which the engineers may not be able to meet or understand separately. This interaction will inform the unique needs of the swarm application, and influence the direction of development. In particular, an example of a unique stakeholder need includes the interaction with military personnel. Swarm robotics is becoming a large area of research in conjunction with military objectives. For example, the Army Research Laboratory has partnered with Northwestern and Georgia Institute of Technology to conduct research on swarm robotics in use for Army applications. An interesting lesson learned from this interaction is that although autonomous swarm technologies can technically be used for a variety of applications, the military can only use fully-autonomous systems in non-lethal applications (Osborn, 2019). Pentagon doctrine dictates that use of lethal force requires a “human-in-the-loop” segment (Osborn, 2019). As a result, research working with the Army focuses primarily on search and rescue operations, which means that advancements in this area will optimize for interpreting the environment and tracking objectives, rather than weaponized actions. In addition, collaboration with the Army encourages engineers to meet defined objectives and timelines for short-term, mid-term, and long-term robotic and autonomous systems (RAS) goals as defined by the Army’s RAS Strategy documents (Osborn, 2019).

Although the current development in military applications views the swarm as a non-weaponized instrument, that does not mean that future implementations of swarms would never pose harm to civilians if put into wide use. For example, a popular application area for swarms in the future may be in conjunction with self-driving cars. With the cars acting as members of the swarm, the vehicle itself may pose risk to humans if a control algorithm goes awry or a cyber attack is conducted. Although the specific agent in the swarm is not a weaponized device, the nature of the size and weight of the vehicle may pose a safety risk regardless.

Although swarms currently consist of mostly “simple” robots, the ever-expanding nature of technology makes it inevitable that this definition will soon be stretched to include mobile agents capable of higher cognition, as was the case in the expansion of mobile phones to include smartphones. In the scenario of self-driving cars, the agents within the swarm are no longer mindless worker ants. This means that each member of the swarm has the potential to provide more complex information, enabling them to become an “expert” regarding their immediate surroundings. In such a way, the swarm effectively forms a trading zone of knowledge regarding the surrounding environment. This proves to be useful in long-distance travel scenarios, where cars in the swarm are able to know important factors such as road conditions, road blockages, and dangerous weather scenarios before they are encountered by the individual vehicle. This communication method for real-time information regarding upcoming hazards is a component which can be developed separately from the autonomous driving itself and incorporated as a feature while human drivers are still at the wheel. In fact, Honda recently announced a safety feature of this exact type called SafeSwarm, which is currently undergoing field tests (Etherington, n.d.).

Future-forward considerations are extremely important for emerging technologies such as swarm robotics. Although current research and application methods may not pose much risk to users or the general public, as the popularity of the technology increases and the number of people impacted by the technology rises, the risk of vulnerabilities scales as well. This has been the case with many key technologies in the past such as internet and mobile devices. In those cases, and as noted in a 2009 paper from Higgins, Tomlinson and Martin, many widespread technologies which are popular today were actually retrofitted with security to address threats after the fact (Higgins et al., 2009).

In the past, this has presented challenges with data breaches and viruses, but these attacks have become more advanced in nature and the interconnectivity of devices in this current age brings a higher level of device vulnerability. As the prevalence of cyber crime and cyber attacks increases, emerging technologies pose bigger risks for large-scale devastation. In the case of swarm robotics, and as with other autonomous systems as a whole, special care should be taken during the development process to address security threats before they occur.

Beyond security threats, there still remains a concern regarding the ethicality of autonomous systems themselves. If applied to the large-scale future concept of swarm fleets of self-driving vehicles, this describes the difference between concern regarding the self-driving behavior itself versus concern regarding the swarm's behavior as a whole. Since swarm intelligence relies on autonomous behavior as a baseline to function, this means that the integrity of machine learning must be ethically challenged as well.

Although the establishment of machine learning has made autonomous machinery more capable, opportunities may emerge where the machinery's real actions deviate from behavior the designer intended. These are cases of robots not going "rogue" but rather finding more efficient ways of completing a task that a human did not think of and which may avoid the intended purpose of the action as well. For example, a 2017 experiment from Stanford and Google researchers revealed an AI "cheating" at its task and avoiding the intended purpose of the task itself (Chu et al., 2017). In this experiment, a system designed to transform between aerial photographs and street maps back and forth discovered a cheat which was imperceptible to its human supervisors. Since the system was tested for accuracy between the final double-transformed image and the original presented image, the system attempted to optimize information retention through any means it could. This led the system to hide information in tiny bits of the transformed image, allowing it to store away hints useful for transforming back to the original image.

Although this action was useful for the specific task assigned to the system, it defied the purpose of the experiment as the system could not successfully transform a street map into an aerial map without the hidden information (non-native to the image) which the system had stored for itself. Other experiments in AI and robotics have revealed similar unintended consequences, particularly when including reinforcement learning, which encourages a system to behave as close to perfect as possible (Vamplew, 2004).

This tendency of machine learning to follow the designer's specific instructions verbatim may make it difficult to set regulations of testing these autonomous systems, since humans relay information in a way which does not always match up with the logical and literal methodology of a computer. This is particularly concerning in scenarios where the algorithm uses reinforced learning, since the method could be encouraged to do a task efficiently rather than ethically. This could lead to holes in the designer’s knowledge of the algorithms functionality, and pose issues in the future for untested use cases. As a result, special care should be taken in the future to test reinforced learning systems more rigorously.

With regards to communication, most multi-robot systems employ a direct robot to robot communication system. The actual implementation of this may vary based on the system, but often consists of RF or IR communications (Higgins et al., 2009). Beyond this however, swarm technology is unique in that it capitalizes on a communication method learned from insects called stigmergy. Stigmergy is a way of leaving clues and learning from the environment around you (Tang et al., 2017). This type of communication is new to the world of robotics, so it poses both new advantages and novel challenges. When it comes to security, a large uncertainty lies with indirect communication. RF and IR communication methods have been thoroughly tested and secured around prior applications, so direct communication includes the same security risks as with any other robotic system. Since the indirect communication relies on context clues in the environment however, it is difficult to define the security of it in swarm robotics.

In terms of error-detection algorithms designed specifically for swarm robotics, there is a silver lining in that recent work has shown that data-driven error-detections are better in dynamic environments than model-driven detections (Higgins et al., 2009). If this trend continues, it bodes well for the testability of these detection algorithms. A small downside however, is that the data driven systems of course have their own verification difficulties as well, which might be dependent on the sensors. All this would mean in terms of testing however, is to expand the range of components tested. This would likely require incorporation of a technology auditing team, who could conduct interviews with the device's subject matter experts and use their own expertise in evaluating risk mitigating controls to provide a determination of the device's risk.

With the recent emphasis on AI and self-driving cars in the media, researchers have been under pressure to adhere to a certain "arms race" within autonomous vehicles (Desmond, 2018). Although this trend is useful for advancing technology quickly, it may lead to normalized deviance in allowing small imperfections in environment recognition, which can build up and result in a snowball effect. As such, engineers must take proper steps when optimizing autonomous systems for efficiency, as improper diligence may result in unintended shortcut behaviors. The self-driving "arms race" may also have an undue effect on the development of proper security protocols for autonomous vehicles. In preventing this normalized deviance in shortcuts, an internal auditing team may be appropriate here as well.

Speaking to cybersecurity, auditing, and ethics, all of this is far outside the normal scope of what is usually considered in swarm robotics. Although this discrepancy may seem worrisome, the reality is that swarm intelligence covers much more ground than what one discipline can adequately address. Emerging technologies have become so advanced and intertwined that trading zones beyond the usual applied scientist roles in engineering must be utilized.

Swarm robotics in analysis presents many similar obstacles to widespread adoption as in the case with nanotechnology. Using the terminology defined by Gorman in the Encyclopedia of Nanoscience and Society, further conclusions may be made. While the academic area of swarm robotics can be defined as an interdisciplinary collaboration, the problem space presented by societal adoption of swarm technology requires multi-disciplinary sharing (Gorman, 2010). This type of multi-disciplinary sharing is best facilitated by a local trading zone, where experts in autonomy, robotics, and cyber security can exchange knowledge in progress towards the common goal of ethical implementation. At current state, these subject areas coalesce into the intellectual trading zone of cyber-physical systems. In time, once cyber-physical systems as a study becomes more established, it may also evolve into its own interdisciplinary collaboration that will inform new expertise areas in other emerging technology areas (Gorman, 2010).

For example, the University of Virginia recently unveiled the Link Lab in 2018, part of a 2015 cyber-physical systems initiative to establish a multi-disciplinary center to bring together researchers to "develop and deploy systems that link the cyber and physical worlds," a key portion of which is autonomous robotics (*Link Lab*, 2018). This physical trading zone serves as a facilitator for the exchange of information required to not only develop and deploy robust intelligent systems, but to also inspire and educate the next generation of engineers to follow suit. Even within the past two years since the lab's opening, many autonomy-focused courses are taught through faculty associated with the Link Lab. This provides another step in the anticipatory governance process which may often fly under the radar, where anticipatory governance principles are embedded not only in implementation but in education as well. Students with this type of background will be better trained to approach multi-disciplinary issues in the future in a way which limits ethical risks. When it comes to cyber-physical systems and swarm robotics, it is important to understand both the security vulnerabilities and the ways to address them moving forward. With the security and vulnerabilities of swarm robotics identified, many error-detection algorithms are already being researched at the university scale to improve a swarm’s resistance to attacks.

These methods of error detection could identify anomalies in robot behavior, and use this to prevent error propagation to the rest of the swarm. This is an extremely important piece of providing an ethical assurance to the actions taken by robots in the swarm, because it allows observers to diagnose malignant behavior before it can be transmitted across the swarm. This malignant behavior could appear as the result of a bad actor performing an attack or infiltration upon the swarm, or simply a physical sensor or actuator malfunction within a device itself.

When approaching deviant behavior at the swarm level, two possibilities emerge: either a rogue robot may influence a global normalized deviance effect amongst its peers, or the rogue robot will only be able to have a local effect, and the majority of the system as a whole will remain unaffected (Kolling et al., 2016). With the latter possibility, impact due to deviance may be minimal due to the small scale and capability of an individualized robot. Many research publications have touted the resiliency of swarm robotics to faults in individual robots (Lau et al., 2011). The theory behind this conclusion is that a failure of one robot out of one hundred still leaves ninety-nine functioning robots. Those yet-functioning robots are able to communicate and work together to complete the task of the failed robot. In such a way, swarms have been used in situations which require dynamic and robust solutions. Military applications of swarms are the most promising, since the swarm as a whole is able to survive individualized attacks.

When individual robot error comes into play, a key component of reconciliation and error detection is to ensure state-communication across robots of a defined neighborhood. This is achieved through a method called collective self-detection (CoDe) where a robot tries to identify itself as faulty in comparison to the healthy neighborhood behavior around it, rather than trying to identify faults in other robots. Interestingly enough, this method is inspired by the observed self-isolation behavior of dying ants, where "some species of ants infected by parasites tend to isolate themselves to die" rather than be identified as sick and isolated by other ants (Lau et al., 2011). This method brings into play the idea of using nature as inspiration for both the technology and its solution. What this study fails to secure however, is the possibility of multi-robot failure and robot fault due to environmental factors. Both of these considerations may bring into effect a complete failure of the system as a whole if the system lacks a robust error-detection scheme.

When multi-robot failure occurs at one time, the ant method above lacks the ability to properly combat the deviant behavior. Since detection is founded in comparing the individual self to nearest neighbors, it is generally operating under a majority-rule scheme. If the majority of neighbors are acting similarly to the self, behavior continues. If the majority of neighbors instead present a differing behavior model, the individual's behavior ceases. In an ideal world, the correct robots would always outnumber the incorrect members of the swarm. However, what if a well-functioning robot becomes surrounded by neighbors which are all exhibiting unintended deviant behavior? This presents a vulnerability to the system. Under the comparison policy, the good robot will consider itself "wrong" compared to its peers. This would usually lead the robot to either remove itself from the swarm, or adapt to match the new normal of the neighbors' behavior. In such a way, if deviant behavior is able to affect multiple robots at once, it may become normalized within the swarm unintentionally.

In most stationary groups, the errors presented by multi-robot failure may lead to a global normalized deviance effect. In a swarm however, the magnitude of impact is unknown. This is due to the shifting nature of mobile robots in a swarm. As the robots move around, an individual’s “neighborhood” will change, reducing the likelihood of multiple rounds of false “healthy” conclusions. Although this provides a bit of added security, more robust error-detection algorithms should be built to provide multi-robot error detection. The beauty of the scientific research process is that a few groups have already built upon the CoDe method to directly address the multi-robot problem. This new method is unique in that it takes inspiration in part from the natural immune system. This shows that even nature-inspired solutions indirectly related the swarms themselves may provide insights to successful error-detection programs. The algorithm presented in this 2013 paper highlights the use of a Receptor Density Algorithm (RDA) for classification, inspired by the methods of T-cell receptor signalling in the immune system (Lau et al., 2013). This, combined with the CoDe method allowed for the researchers to successfully identify instances of multi-robot failure within a swarm. Although all of the studies mentioned above do not provide a robust solution platform for swarm security, they provide a sense of certainty that researchers are actively working on this problem, continuing to look towards nature for inspiration. These studies will continue to build upon each other, and provide a sound base for predicting the sustainability of swarms in a wider-spread use case.

The overall subject of swarm robotics is a very complicated and intricate problem space. It combines key knowledge from a variety of expertise areas, and the potential for wide-spread use with other emerging technologies requires an advanced look at how the technology interacts with and requires even more prevalence of trading zones. As with any exciting new technology, the opportunities for growth are endless but so are the possibilities for ethical concerns. The way to adequately address these concerns in the future and provide a base for anticipatory governance is to begin and sustain transparent collaboration between many key stakeholders and technologists. Simultaneously, researchers need to be increasingly aware of the security ramifications surrounding swarm robotics, and work diligently on error-detection algorithms. The work already being done and trade zones already being established shows a promising view for the future of swarm robotics, but a constant vigilance needs to be maintained to keep up with changing stakeholder needs and system vulnerabilities.

References

Bauer, A. M., Wollherr, D., & Buss, M. (2008). Human-Robot Collaboration: A Survey. *Int. J. Humanoid Robotics*. https://doi.org/10.1142/S0219843608001303

Beer, J. M., Fisk, A. D., & Rogers, W. A. (2014). Toward a framework for levels of robot autonomy in human-robot interaction. *Journal of Human-Robot Interaction*, *3*(2), 74–99. https://doi.org/10.5898/JHRI.3.2.Beer

Chu, C., Zhmoginov, A., & Sandler, M. (2017). CycleGAN, a Master of Steganography. *ArXiv:1712.02950 [Cs, Stat]*. http://arxiv.org/abs/1712.02950

Couceiro, M. S., Rocha, R. P., & Martins, F. M. L. (2015). Towards a predictive model of an evolutionary swarm robotics algorithm. *2015 IEEE Congress on Evolutionary Computation (CEC)*, 2090–2096. https://doi.org/10.1109/CEC.2015.7257142

Desmond, J. (2018, March 13). Normalization of Deviance Endangers AI Self-Driving Cars. *AI Trends*. https://www.aitrends.com/ai-insider/normalization-of-deviance-endangers-ai-self-driving-cars/

Etherington, D. (n.d.). *Honda's Safe Swarm concept has cars mimicking fish for safer driving | TechCrunch*. Retrieved April 21, 2020, from https://techcrunch.com/2017/01/05/hondas-safe-swarm-concept-has-cars-mimicking-fish-for-safer-driving/?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xlLmNvbS8&guce_referrer_sig=AQAAAGZGwZY2ICyGEInensfc2yQAfwi2t-r0eZPKRhe7IxlHmQI3YDb3Juzhh4bVrnh7kIogPissje2AGdoH3JGBRbLkVEK9t46nzFCB5KVR1xjagZZaRFMwyNalAhuIwflHgvfUpTyOjfNY_mhwzux5vymjPrdT3rO8718_vwLMHQKM

Garnier, S., Gautrais, J., & Theraulaz, G. (2007). The biological principles of swarm intelligence. *Swarm Intelligence*, *1*, 3–31. https://doi.org/10.1007/s11721-007-0004-y

Gorman, M. E. (2010a). Multidisciplinarity. In *Encyclopedia of Nanoscience and Society* (pp. 447–447). SAGE Publications, Inc. https://doi.org/10.4135/9781412972093

Gorman, M. E. (2010b). Trading Zones. In *Encyclopedia of Nanoscience and Society* (pp. 768–768). SAGE Publications, Inc. https://doi.org/10.4135/9781412972093

Higgins, F., Tomlinson, A., & Martin, K. (2009). *Survey on Security Challenges for Swarm Robotics*. 307–312. https://doi.org/10.1109/ICAS.2009.62

Javaid, A. (2013, April). *Understanding Dijkstra's Algorithm by Adeel Javaid: SSRN*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2340905

Kolling, A., Walker, P., Chakraborty, N., Sycara, K., & Lewis, M. (2016). Human Interaction With Robot Swarms: A Survey. *IEEE Transactions on Human-Machine Systems*, *46*(1), 9–26. https://doi.org/10.1109/THMS.2015.2480801

Lau, H., Bate, I., & Timmis, J. (2013). Immune-Inspired Error Detection for Multiple Faulty Robots in Swarm Robotics. *Artificial Life Conference Proceedings*, *25*, 846–853. https://doi.org/10.1162/978-0-262-31709-2-ch124

Lau, H., Timmis, J., & Bate, I. (2011). Collective Self-detection Scheme for Adaptive Error Detection in a Foraging Swarm of Robots. In P. Liò, G. Nicosia, & T. Stibor (Eds.), *Artificial Immune Systems* (pp. 254–267). Springer. https://doi.org/10.1007/978-3-642-22371-6_23

Li, W. (2019). An Improved Artificial Potential Field Method Based on Chaos Theory for UAV Route Planning. *2019 34rd Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, 47–51. https://doi.org/10.1109/YAC.2019.8787671

Link Lab. (2018, December 17). University of Virginia School of Engineering and Applied Science. https://engineering.virginia.edu/link-lab-0

Monaco, J. D., Guzman, R. M. D., Blair, H. T., & Zhang, K. (2019). Spatial synchronization codes from coupled rate-phase neurons. *PLOS Computational Biology*, *15*(1), e1006741. https://doi.org/10.1371/journal.pcbi.1006741

Ohkura, K., Yasuda, T., & Matsumura, Y. (2013). Coordinating the Collective Behavior of Swarm Robotics Systems Based on Incremental Evolution. *2013 IEEE International Conference on Systems, Man, and Cybernetics*, 4024–4029. https://doi.org/10.1109/SMC.2013.687

Osborn, K. (2019, November 1). *The Army Has a Crazy Idea To Win the Next Conflict: Ant Colony Warfare* [Text]. The National Interest; The Center for the National Interest. https://nationalinterest.org/blog/buzz/army-has-crazy-idea-win-next-conflict-ant-colony-warfare-92631

Swarm robotics—From local rules to global behaviors | Magnus Egerstedt | TEDxEmory. (n.d.). Retrieved October 3, 2019, from https://www.youtube.com/watch?v=ULKyXnQ9xWA

Tan, Y., & Zheng, Z. (2013). Research Advance in Swarm Robotics. *Defence Technology*, *9*(1), 18–39. https://doi.org/10.1016/j.dt.2013.03.001

Tang, Q., Yu, F., Zhang, Y., Ding, L., & Eberhard, P. (2017). A Stigmergy Based Search Method for Swarm Robots. In Y. Tan, H. Takagi, Y. Shi, & B. Niu (Eds.), *Advances in Swarm Intelligence* (pp. 199–209). Springer International Publishing. https://doi.org/10.1007/978-3-319-61833-3_21

Thomas, S., Christoph, M., Ronald, T., & Karl, C. (n.d.). *Bio-inspired Navigation of Autonomous Robots in Heterogenous Environments*.

Tsvetkova, M., Yasseri, T., Meyer, E. T., Pickering, J. B., Engen, V., Walland, P., Lüders, M., Følstad, A., & Bravos, G. (2017). Understanding Human-Machine Networks: A Cross-Disciplinary Survey. *ACM Computing Surveys*, *50*(1), 12:1–12:35. https://doi.org/10.1145/3039868

Vamplew, P. (2004). *Lego Mindstorms Robots as a Platform for Teaching Reinforcement Learning*. Undefined. https://www.semanticscholar.org/paper/Lego-Mindstorms-Robots-as-a-Platform-for-Teaching-Vamplew/555d2283f4bef4eda4c7237cfdc368a535890ab9

Verlekar, H., & Joshi, K. (2017). Ant bee inspired foraging swarm robots using computer vision. *2017 International Conference on Electrical, Electronics, Communication, Computer, and Optimization Techniques (ICEECCOT)*, 191–195. https://doi.org/10.1109/ICEECCOT.2017.8284663

Yanagida, T., Mohan, R. E., Pathmakumar, T., Elangovan, K., & Iwase, M. (2017). Design and Implementation of a Shape Shifting Rolling–Crawling–Wall-Climbing Robot. *Applied Sciences*, *7*, 342. https://doi.org/10.3390/app7040342

Zuckerberg explains the internet to Congress. (n.d.). Retrieved April 8, 2020, from https://www.youtube.com/watch?v=ncbb5B85sd0