

Memetic Engineering as a Basis for Learning in Robotic Communities

Walter F. Truskowski, Christopher Rouff, Mohammad H. Akhavannik

NASA Goddard Space Flight Center (Emeritus, Code 580), Lockheed Martin, The Inspire Project
walt.truskowski@nasa.gov, christopher.rouff@lmco.com, hasan_a@gwmail.gwu.edu

Abstract

This paper represents a new contribution to the growing literature on memes. While most memetic thought has been focused on its implications on humans, this paper speculates on the role that memetics can have on robotic communities. Though speculative, the concepts are based on proven advanced multi agent technology work done at NASA - Goddard Space Flight Center and Lockheed Martin.

The paper is composed of the following sections :

- An introductory section which gently leads the reader into the realm of memes.
- A section on memetic engineering which addresses some of the central issues with robotic learning via memes
- A section on related work which very concisely identifies three other areas of memetic applications, i.e., news, psychology, and the study of human behaviors.
- A section which discusses the proposed approach for realizing memetic behaviors in robots and robotic communities.
- A section which presents an exploration scenario for a community of robots working on Mars.
- A final section which discusses future research which will be required to realize a comprehensive science of robotic memetics.

Introduction

NASA is investigating autonomous and autonomic robotic systems able to act independently in dynamic and uncertain environments (Ambrose et al., 2010, Truskowski 2009, Truskowski 2011). Robots and autonomous systems are already at work in all of NASA's Mission Directorates. NASA will see even more pervasive use of these systems in its future . Key attributes of such autonomy for a robotic system include the ability for complex decision making, including the ability to self-adapt as the environment in which the system is operating changes, and the ability to understand system state and react accordingly (Ambrose et al., 2010). The robots can do exploration in areas that would otherwise be very costly, dangerous or impractical due to travel time. As seen in recent robotic missions, single robots

have limited abilities to explore large areas efficiently. NASA has shown increased interest in robotic teams for exploration and servicing. In these teams different robots would have various roles and responsibilities. To work together in a team, robots will need to exchange information on their current status, needs, capabilities and findings. Using memes as a unit of transmission could facilitate this needed information exchange.

Memetics has been a recent subject of interest as a new method for information exchange (Silby 2000, Hougen et al., 2003, Auger 2002, Blackmore 2000, Wilson et al. 2011) . In communities, memes have been studied to understand and enhance group learning (Hougen et al., 2003). Richard Dawkins first defined memes as a unit of cultural transmission (Dawkins, 1989). Essentially, memes are ideas that evolve according to the same principles as biological evolution (Silby, 2000). Memetic learning works by transmitting units of cultural ideas or symbols from one mind to another. All ideas that exist within an individual's mind are examples of memes. Memes that are good at replicating leave more copies of themselves in minds. Examples of memes are catch phrases, musical themes, scientific ideas and sayings. In robotics, examples of memes are algorithms, observations, and instructions.

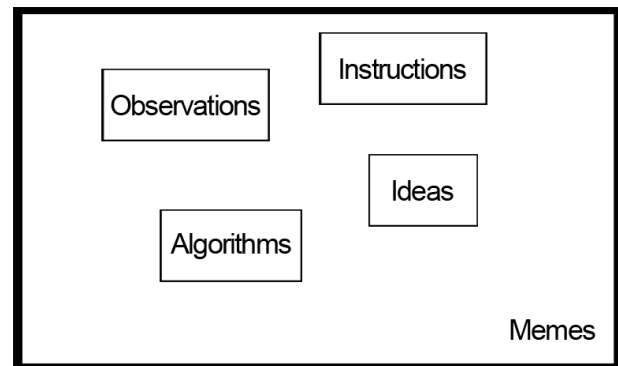


Figure 1 - Visual Representation of Memes

Memes benefit from having behavior similar to genes (Gunders, 2010). Like genes, memes are under constant selection pressures. Memes are in constant competition to be absorbed and evolved from the collection of memes in the community. Memes are competing to be learned and those that are better at reproducing are those that are successful in accomplishing their intended behavior. Memes give communities more power and knowledge to accomplish their tasks. There are still many unknowns regarding the representation and transmission of memes in a robotic community.

Memetic Engineering

For robots, memes have been defined as sets of instructions that can be followed to evolve behavior. Instructions can be encoded as written text and visible or vocal action (Silby, 2000). To allow for memetic learning, memes should also include observations of the environment. A robot that is able to observe and intelligently imitate the behavior of others is able to participate in memetic learning. In order to perform intelligent imitation, a robot needs to be able to process memetic information. This process involves evaluating models, examples, and patterns which the robot observes (Hougen et al., 2003). In addition, the robot is expected to analytically compare its current knowledge to the new information it is observing.

A robot has modified its individual knowledge base when it learns a new meme. Each robot is expected to evaluate active memes in the community knowledge base for strengths and weaknesses when deciding whether to learn them. It can be expected that each individual robot will benefit from the aggregation of other robots which are also participating in memetic learning. When the community knowledge base and size expands, there is a larger selection of memes which can be evaluated and learned. With a larger community knowledge base, a robot has a larger selection of memes to modify to develop novel memes. All individuals capable of participating in memetic learning are able to generate new memes. Also, individuals who are able to broadcast observations are capable of generating new memes.

The memes that are in the knowledge base of a robot are in constant competition with all other memes in the meme pool (Silby, 2000). The meme pool is the collection of all existing memes that are accessible to the other individuals in the community. An individual robot may develop new memes that become candidates for imitation in the community meme pool. The community meme pool increases with the addition of novel memes generated by individual robots in the community. The connection between an individual knowledge base and the community meme pool is similar to the structure of a distributed cloud network. A distributed cloud network is structured so that each individual is connected to the cloud where they can

access the knowledge bases of others in the community pool. Individuals will be able to quickly access, process, and analyze the collective knowledge within the cloud (Krutz et al., 2010).



Figure 2 - Visualization of Memetic Knowledge Base

Related Work

Memes can be utilized for many other applications other than robotics. One example has been tracking memes to develop a coherent representation of the news cycle - the daily rhythms in the news media that have previously been difficult to perform quantitative analysis (Leskovec, Backstrom, & Kleinberg, 2008). By using memes, Leskovec et al were able to develop a framework for tracking short and distinctive phrases that travel relatively intact in media. Further, the change in information as it propagates was able to be observed. Using memetics in this model has provided a method to see how a particular idea moves within and between groups.

Memetics also has the potential to enhance our study of psychology. Memes can be used to discover the origin of certain psychological conditions. By modeling the mind as a memetic construct, conditions such as depression or addiction might be explained by memetic viruses that influence the behavior of an individual (Silby, 2000). This method may provide a path to pinpoint methods to prevent the spread of malicious psychological conditions.

The most powerful adaptation of memetic systems is the ability to predict behavior and evolution of future memetic structures. In the future, psychologists may be able to model memetic learning within communities and predict what will happen when individuals are exposed to combinations of memes. With this knowledge, memes will become a driving force in the study of human behavior (Silby, 2000).

Proposed Approach

A functional and efficient meme is more likely to flourish within a community than a meme which is not as practical.

When a meme does not replicate from one source to another, it has lost in the struggle for survival of the fittest (Dawkins, 1989). This evolutionary component provides a distinct advantage to memetic learning algorithms. The better meme is the one which is evaluated and chosen to be replicated, executed, and evolved by the other members of the community.

If a new meme is determined to be beneficial to the individual, then the individual will update its knowledge base. After updating its knowledge base, the individual will alter its current memes and subsequently change its behavior. As a result, the new behavior exhibited by the individual will create new observations and analysis made by its peers.

The effect that a new meme has on a community of robots can be seen by observing the change in behavior within the community. When new memes are introduced to an individual's knowledge base, the individual performs analysis and determines whether or not it needs to adjust its current behavior. This determination is dependent on comparing the effectiveness of the new and current meme to accomplishing its goal. For example, the new meme could be information that will improve the individual's ability to accomplish its goal.

The new memes generated from these observations may spread and change the memes and behavior of other individuals in the community. This cycle repeats itself while there are new memes being generated. Therefore, the community meme pool evolves as there are better memes generated due to the observations and experiences of the community and increases the intelligence of the community.

In order for a meme to be selected for learning, its individual strengths and weaknesses need to be evaluated. Robots that are participating in memetic learning are capable of intelligent imitation and can determine the success of each meme it can observe in accomplishing its goal. The individual robot may look at the success other individuals have to determine which to imitate (Hougen et al., 2003). Each robot will have to report the success it has had with the meme. With multiple robots reviewing each meme, the community will have a way of knowing which memes are best and functional. Robots can report the background condition, function, procedure, and outcome of the memes so other members in the community can analyze the effectiveness of each meme. Additionally, if an individual fails to do something, then that meme should also be reported. Reporting failures will allow other individuals to know not to attempt something which they have not yet done.

Mars Exploration Scenario

The benefits of memetic learning become increasingly evident when considering a scenario with a team of robots on Mars. In the following scenario the flow of the memes

through the robotic team is shown in Figure 4. This team may consist of exploration robots, logistic robots and analytic robots. Each robot has a specific role in the team and is initially programmed with intelligence prior to the start of the mission. The intelligence provided to the robots prior to the start of the mission is not complete but is sufficient to complete the mission if the initial intelligence is all correct. For instance, the memes provided to the team suggests that in order to find a specific sample the exploration robots need to search or excavate in a geographical area where the sample is thought to be. The logistics robots have the responsibility of setting up a base camp for refueling, resupply and directing communications. The analytic robots will analyze the samples gathered by the exploration robots to determine whether or not they are gathering the correct sample.

The robots start searching for the sample in a predetermined area that was defined by the scientific team on earth (Figure 4, box 1). Upon searching in that area for a length of time, the team has no success in finding anything that matches the sample specification (Figure 4, Box 2). However, now realizing that the area does not contain the sample which they are looking for, the exploration robots go through their knowledge base and alter their current search meme to determine a new area to search (Box 3). The exploration robots use memes they have learned after the exploration has begun along with its current knowledge base to determine a new search area. Although the sample is not known to be in this new area the robot has determined the team may have more success by searching there because it also matches some of the geological characteristics that the robot team was given at the start of the mission. Because of this decision, the exploration robots move to the new area (Box 4). This action is a direct result of the memetic learning algorithms used by the team.

After searching in the new area the analytic robots examine the samples obtained by the exploration robots and gets positive results (Box 5). Upon examination, the analytic robots develop a new meme which confirms that searching in this new area provides success in obtaining the target sample. This meme is now in the community knowledge pool and directly impacts the behavior of the other robots (Box 6). Now other exploration robots may move to this new area to also search for additional samples. Also, the exploration robots may check similar nearby locations with the same geological features for future samples. Based on the movement of the other robots to this new area, the logistic robots may plan to move the base camp to a site closer to the new area (Box 7). The logistic robots may also have to do new calculations on the best areas to perform communications based on the new area's topology (Box 8). The behavior of the analytic robots may not change since this new location does not affect how they will do their analysis on incoming samples. However, the meme generated by the analytic robot is what




Figure 4 - The memetic system level flow diagram for the scenario. In this figure, memes are represented in process blocks while decision making processes are represented in decision blocks.

spurred the memetic learning process within the team and brought success to the mission.




Figure 3 - Example of Mars Colony

As another example, suppose during the movement of the robots to the new location a sand storm occurs. Because of the severity of the sandstorm, the robots' vision is reduced to the point where it is no longer safe for the robots to travel so they have to stop. Since the robots are not able to get to a shelter, they are exposed to the blowing sand. The result is that some of the robots are damaged. In this case, a new meme is also learned that the robots should seek shelter as soon as possible (Box 10,11), which is when there is a weather report of a possible sandstorm (Box 9). In this case the new meme enhances safety of the robotic community through the shared knowledge. It may also inform the scientific team back on earth that future robots sent to Mars need to be better protected from sand, and thus change the design of future robots that will join their community.

Future Work

An area for future investigation is to modify the meme's structure so that the memes themselves are intelligent. Memes could be encapsulated with intelligent software, similar to a mobile agent, that can make the meme an active instead of a passive entity. Active memes could monitor a robot's state, listen to communications between robots or look at other memes that are being passed between robots to determine if it may be needed. If so, the meme could then push itself to a robot that needs it or insert itself into the active reasoning being done by the host robot. Memes that are not needed after a period of time could decay and destroy themselves if they are outdated or no longer needed. This could help limit the proliferation of memes in a system. Active memes could also automatically update themselves based on changes in the environment and observed learning in the robot host (e.g., seeking protection when a sandstorm is forecasted). They could also seek out similar memes and combine with them to form better memes, or even use techniques like genetic programming to improve or transform themselves into something new. This could speed up the evolution of new memes.

An intelligent and mobile meme would have to be lightweight so it could easily move between robots without

having a large communications overhead. A large number of heavy weight memes being sent between robots over a limited bandwidth network, such as might be on Mars, could overload the network. Memes could also use swarming or other behaviors to increase their individual impact and to quickly react to new situations a robot may encounter where there is no one meme that has all of the information. Swarming could also help to keep individual memes small since the swarm would provide all of the needed knowledge. Since each situation could be different, new swarms would be spawned based on the situation at hand.

An additional area which requires further investigation is the implementation of such a communication and analytical model to communities (Wang, 2008). There are unknowns regarding future requirements and specifications for a group of robots to be fitted with the intelligence to perform memetic learning. The communication structure of modern systems may be challenged and need to be modified to allow future systems to be capable of implementing this technology.

There are many other questions which need to be investigated to obtain a better understanding of memetic engineering. For example, What is an intelligent meme? How can web intelligence be mapped into meme pools (Stuckenschmidt, 2005)? What are the other potential impacts that memetic engineering will have on other areas of research and development? Lastly, formal models of memes and memetic behaviors in robotic communities must be developed.

Conclusions

Robotic communities can benefit from introducing memetic engineering into their infrastructures. Memes provide a method for robots to share ideas and for communities to evolve. As memes evolve so do their communities and meme pools. Memes can empower robots and their algorithms. In the scenario, a community of robots used memes to learn from each other while in a dynamic environment. Space environments, such as Mars, can be very difficult to predict beforehand and programmed into the robots. As shown in the scenario, memetics could have an extremely practical role in the future of space exploration and the development of artificial intelligence for unpredictable environments and in sharing what is learned. A number of potential future research problems were also introduced that shows that memetic engineering is an area that will surely remain in scientific discourse and research.

Acknowledgements

Walt Truskowski wishes to thank Thomas Flatley and Edward Hicks in Code 587, Goddard Space Flight Center, for continuing support of his NASA emeritus status.

Mohammad Akhavannik wishes to thank his mentors and the INSPIRE project for sponsoring his NASA internship.

References

Ambrose, R., Wilcox, B., Reed, B., Matthies, L., & Lavery, D. National Aeronautics and Space Administration, (2010). Draft robotics, tele-robotics and autonomous systems roadmap (Technology Area 04). Retrieved from http://www.nasa.gov/pdf/501622main_TA04-Robotics-DRAFT-Nov2010-A.pdf

Aunger, R. (2002). The electric meme: a new theory of how we think. Simon and Schuster.

Blackmore, S. (2000). The meme machine. Oxford University Press, USA.

Dawkins, R. (1989). The selfish gene. Oxford University Press, USA.

Gunders, J. (2010). The complete idiot's guide to memes. Alpha.

Hougen, D.F., Carmer, J., & Woehrer, M. (2003). Memetic learning: a novel learning method for multi-robot systems. Robotic Intelligence and Machine Learning Laboratory, School of Computer Science University of Oklahoma, Norman, OK.

Krutz, R. L., & Vines, R. D. (2010). Cloud security, a comprehensive guide to secure cloud computing. Wiley.

Leskovec, J., Backstrom, L., & Kleinberg, J. (2008). Meme-tracking and the dynamics of the news cycle. Retrieved from <http://www.memetracker.org/>

Miller, J. G. (1995). Living systems. Univ Pr of Colorado.

Silby, Brent. (2000). What is a meme?. Retrieved from http://www.def-logic.com/articles/what_is_a_meme.html

Stuckenschmidt, H (2005). Information sharing on the semantic web. Springer.

Truskowski, W. (2011). Memes Research Proposal, Goddard Space Flight Center, Internal Document

Truskowski, W., Hallock, H., Sterrit, R., Rouff, C., Karlin, J., Hinchey, M., Rash, J. (2009). Autonomous and Autonomic Systems with Applications to NASA Intelligent Spacecraft Operations and Exploration Systems. NASA Monographs in Systems and Software Engineering Series. Springer.

Wang, Y (2008). Foundations- a software science perspective. Auerbach Publications.

Wilson, E., & Unruh, W. (2011). The art of memetics. lulu.com.