

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 24-12-2014		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 15-Sep-2011 - 14-Sep-2014	
4. TITLE AND SUBTITLE Final Report: Training a Multiagent Hive Brain for Coordinated UGV Operations (4.4.2)			5a. CONTRACT NUMBER W911NF-11-1-0489		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Kenneth O. Stanley			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Central Florida 12201 Research Parkway, Suite 501 Orlando, FL 32826 -3246			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 59670-NS.22		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This project builds on a learning method developed previously under DARPA support called multiagent HyperNEAT that evolves a set of neural controllers for a team of collaborating wheeled robots. The project was also supplemented by DARPA CSSG Phase 3 matching grant N11AP20003 during its first year. The research focused on three key directions: The first (1) is to extend multiagent HyperNEAT to allow evolving a team of robots that can send signals to one another over wireless connections directly from neurons in one agent to neurons in another, thereby facilitating tight coordination among robots can evolve without any explicit communication.					
15. SUBJECT TERMS multiagent learning, robot learning, robustness, real-world transfer, multiagent HyperNEAT					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU	UU		Kenneth Stanley
					19b. TELEPHONE NUMBER 407-823-4289

Report Title

Final Report: Training a Multiagent Hive Brain for Coordinated UGV Operations (4.4.2)

ABSTRACT

This project builds on a learning method developed previously under DARPA support called multiagent HyperNEAT that evolves a set of neural controllers for a team of collaborating wheeled robots. The project was also supplemented by DARPA CSSG Phase 3 matching grant N11AP20003 during its first year. The research focused on three key directions: The first (1) is to extend multiagent HyperNEAT to allow evolving a team of robots that can send signals to one another over wireless connections directly from neurons in one agent to neurons in another, thereby facilitating tight coordination among robots can evolve without any explicit communication language. The second direction (2) is a novel approach, called reactivity, which facilitates robust transfer from behaviors trained in simulation to robots in the real world. The third direction (3) is to add directionality to the communication system so that agents can efficiently decide and perceive from where in space signals originate. These three complementary ideas, plus enhancements to the underlying algorithms, have appeared so far in 8 conference and 4 journal articles. One paper, at IJCNN-2012, won a Best Student Paper Award out of a pool of 299 entries. Another was a Paper Award Finalist at ICIRA 2012.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received

Paper

06/03/2013 15.00	Kenneth O. Stanley, David B. D'Ambrosio. Scalable multiagent learning through indirect encoding of policy geometry, Evolutionary Intelligence, (06 2013): 1. doi:
06/13/2013 18.00	Joel Lehman, Kenneth O. Stanley. Evolvability is Inevitable: Increasing Evolvability Without the Pressure to Adapt, PLoS ONE, (04 2013): 0. doi:
08/28/2012 12.00	Sebastian Risi, Kenneth O. Stanley. An Enhanced Hypercube-Based Encoding for Evolving the Placement, Density and Connectivity of Neurons , Artificial Life, (11 2012): 0. doi:
12/24/2014 21.00	Joel Lehman , Sebastian Risi, David D'Ambrosio, Kenneth Stanley. Encouraging Reactivity to Create Robust Machines, Adaptive Behavior, (12 2013): 484. doi:

TOTAL: 4

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
04/17/2013 13.00	Justin K. Pugh, Kenneth O. Stanley. Evolving Multimodal Controllers with HyperNEAT, Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013).. 06-JUL-13, . : ,
04/18/2013 14.00	Sebastian Risi, Kenneth O. Stanley. Confronting the Challenge of Learning a Flexible Neural Controller for a Diversity of Morphologies, Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013). 06-JUL-13, . : ,
05/03/2012 6.00	Joel Lehman, Sebastian Risi, David B. D'Ambrosio, Kenneth O. Stanley. Rewarding Reactivity to Evolve Robust Controllerswithout Multiple Trials or Noise, The Thirteenth International Conference on the Synthesis and Simulation of Living Systems (Alife 13) . 19-JUL-12, . : ,
05/03/2012 5.00	Kenneth O. Stanley, Sebastian Risi. A Unified Approach to Evolving Plasticity and Neural Geometry, 2012 International Joint Conference on Neural Networks (IJCNN 2012) . 10-JUN-12, . : ,
06/13/2012 9.00	Joel Lehman, Kenneth O. Stanley. Beyond Open-endedness: Quantifying Impressiveness, The Thirteenth International Conference on the Synthesis and Simulation of Living Systems (Artificial Life 13). 19-JUL-12, . : ,
06/17/2014 20.00	Justin K. Pugh, Skyler Goodell, Kenneth O. Stanley. Directional Communication in Evolved Multiagent Teams, Genetic and Evolutionary Computation Conference (GECCO-2014). 12-JUL-14, . : ,
06/17/2014 19.00	Justin K. Pugh, Andrea Soltoggio, Kenneth O. Stanley. Real-time Hebbian Learning from Autoencoder Features for Control Tasks, The Fourteenth International Conference on the Synthesis and Simulation of Living Systems (ALIFE XIV). 30-JUL-14, . : ,
08/19/2012 11.00	Joel Lehman, Sebastian Risi, Kenneth Stanley, Skyler Goodell, David D'Ambrosio. Multirobot Behavior Synchronization through Direct Neural Network Communication, Proceedings of the 5th International Conference on Intelligent Roboticsand Applications (ICIRA-2012). 03-OCT-12, . : ,
TOTAL:	8

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
02/24/2012	2.00 Joel Lehman, Kenneth Stanley. Beyond Open-endedness: Quantifying Impressiveness, The Thirteenth International Conference on the Synthesis and Simulation of Living Systems (Alife 13) (02 2012)
02/24/2012	1.00 Sebastian Risi, Joel Lehman, David D'Ambrosio, Kenneth O. Stanley. Rewarding Reactivity to Evolve Robust Controllers without Multiple Trials or Noise, The Thirteenth International Conference on the Synthesis and Simulation of Living Systems (Alife 13) (02 2012)
02/24/2012	3.00 Sebastian Risi, Kenneth Stanley. Confronting the Challenge of Learning a Flexible Neural Controller for a Diversity of Morphologies, the Twenty-Sixth Conference on Artificial Intelligence (AAAI-12) (01 2012)
02/24/2012	4.00 Sebastian Risi, Kenneth Stanley. A Unified Approach to Evolving Plasticity and Neural Geometry, 2012 IJCNN International Joint Conference on Neural Networks (IJCNN 2012) (01 2012)
05/03/2012	7.00 David B. D'Ambrosio, Skyler Goodell, Joel Lehman, Sebastian Risi, Kenneth O. Stanley. Multirobot Behavior Synchronization through Direct Neural Network Communication, Proceedings of the 5th International Conference on Intelligent Robotics and Applications (04 2012)
06/13/2012	8.00 David B. D'Ambrosio, Kenneth O. Stanley. Scalable Multiagent Learning through Indirect Encoding of Policy Geometry, Evolutionary Intelligence (03 2012)
06/15/2012	10.00 Kenneth O. Stanley, Sebastian Risi. An Enhanced Hypercube-Based Encoding for Evolving the Placement, Density and Connectivity of Neurons, Artificial Life (04 2012)
TOTAL:	7

Number of Manuscripts:

Books

<u>Received</u>	<u>Book</u>
-----------------	-------------

TOTAL:

Received

Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

1) One of 5 (out of 198) Best Paper Award Finalists, 2012 International Conference on Intelligent Robotics and Applications (ICIRA 2012, Montreal, Canada), for D'Ambrosio, D., Goodell S., Lehman, J., Risi, S., and Stanley, K., Multirobot Behavior Synchronization through Direct Neural Network Communication. (October 2012)

2) Best Student Paper Award (out of 299 papers submitted with student first authors), 2012 International Joint Conference on Neural Networks (IJCNN 2012, Brisbane, Australia), for Risi, R. and Stanley, K., A Unified Approach to Evolving Plasticity and Neural Geometry.

3) 2013 UCF CECS Dean's Research Professorship Award, in recognition of research and mentorship.

4) Undergraduate independent study student Skyler Goodell (not supported monetarily by grant) won 2nd place in UCF's 2013 Showcase of Undergraduate Research Excellence (SURE) for ARO-related entry on "Multirobot Behavior Synchronization through Hive Brain Neuroevolution"

5) 2014 UCF Reach for the Stars Award, in recognition of highly successful research and creative activity accomplished by early-career university professionals.

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Justin Pugh	1.00	
Lisa Soros	0.22	
FTE Equivalent:	1.22	
Total Number:	2	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Kenneth Stanley	0.17	
FTE Equivalent:	0.17	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Skyler Goodell	0.00	Computer Science
FTE Equivalent:	0.00	
Total Number:	1	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 1.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 1.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 1.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PHDs

<u>NAME</u>
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Foreword

This report details results from the three-year ARO-sponsored project. The work from the initial one-year period was also complemented with support from a matching grant from the DARPA CSSG program (phase 3, grant N11AP20003, which expired in August 2012). Methods developed in earlier phases of the DARPA CSSG program are the foundation and inspiration for the work in the present project. In particular, this project extends the earlier multiagent HyperNEAT approach to training multiple robot or UGV agents to work together in a team.

List of Illustrations

Four figures are included in the attached file:

Figure 1. Hive Brain Substrate Architecture. A visualization of a hive brain architecture.

Figure 2. Hive Robots in the Patrol Domain. A picture of real Khepera III robots in the hive patrol domain.

Figure 3. Layout of New Multiagent HyperNEAT Directional Communication Neural Network Substrate. Schematic of new neural network layout for receiving directional communication.

Figure 4. Mazes. Visualizations of the four mazes in which robots were trained for reactivity.

Statement of the problem studied

This project focuses on two complementary problems. The first (1) is to develop a training algorithm to produce teams of autonomous robots or UGVs that can coordinate with each other seamlessly. Because such training algorithms usually run in a simulator, the second problem (2) is to ensure that the behaviors learned in simulation are as robust as possible when transferred to the real world. It is important to note that while these investigations are complementary, an ability to ensure robust simulation-to-real-world transfer can benefit a broad range of applications, both single-agent and multiagent.

(1) The problem of training a team of coordinated agents is a longstanding challenge in the fields of machine learning and reinforcement learning. My research group developed the multiagent HyperNEAT method during previous DARPA-funded research to address this challenge. The multiagent HyperNEAT algorithm could train larger heterogeneous teams than prior multiagent training methods and also could scale the number of agents on the team. It was demonstrated in e.g. room clearing and predator-prey domains in the past.

The idea for the present ARO-funded project is to augment multiagent HyperNEAT with a novel yet potentially powerful capability that could tighten the coordination of deployed teams of robots significantly: Instead of training each agent to assume a predetermined role that it executes autonomously throughout its deployment, the agents in the enhanced approach have the capability to communicate with each other through direct connections between the behavior policies of neural networks within each individual agent. In other words, put more simply, their brains are connected over a wireless network. That way, in principle, the agents can share their internal computations with each other and thereby act similarly to the fingers of an invisible hand, sharing the same thought process although separated by distance. For example, if one agent senses an important feature in the environment, it can notify all the other agents before any of them sense it as well; the team can then react to the new information in a coordinated fashion even though most of the team never actually sensed the feature of interest directly.

Another potential advantage of this approach, informally called the “hive brain,” is that because the communication flows over direct neural connections between the neural network in one agent and the neural network in another, the training algorithm (multiagent HyperNEAT) does not need to be provided any a priori communication language or formalism. The main challenge in the first year of this project was to develop the hive brain architecture and demonstrate its capabilities in the real world on Khepera III robots in a test domain. The second year then began research on directional communication and multimodality. In short, an important intermediate challenge recognized after the first year is that effective coordination in a geometric context requires an awareness of the direction from which signals originate. However, such directionality increased the amount of communication inputs, which triggered complementary research on how to represent large multimodal groups of inputs in HyperNEAT neural networks. In the third year the critical contribution of directional communication was fully investigated and published, and then applied to significantly more complex tasks.

(2) Most learning algorithms that train behaviors for autonomous vehicle and robots are run exclusively in simulators because training in the real world would be both too time-consuming and too expensive. That is, because training usually requires hundreds or thousands of trials that are often dangerous to the hardware involved, completing it safely and in a reasonable

timeframe is often unrealistic. Therefore, researchers build simulators that aim to replicate the conditions of the real world as accurately as possible. These simulators must include models of the robotic or UGV hardware and its physical responses to control signals from its learned controller. Because of the inherent complexity and presence of noise in the real world, simulators are rarely perfect replications of reality. Because of this discrepancy, any behavior learned in simulation likely will not behave equivalently in the real world. For example, a robot trained to travel down a hallway in simulation may succeed easily in the simulator, but imperfect wheels in the real world might cause a real robot with the same controller to veer slightly to the side when its controller commands it to move straight.

To address such discrepancies, the problem of simulation-to-real-world transfer has attracted significant attention in recent years. The typical solution to this problem is to test every candidate controller in the simulator over multiple trials with varying levels of noise. The idea is that the controller's average performance over many such trials gives a better idea of how consistent its behavior is in the presence of unpredictable noisy circumstances than a single trial can offer. However, in practice this approach is inconvenient because the need for testing every controller (recall there may be hundreds or thousands of such candidates in a single run) in multiple trials means training at best will be multiple times slower. Furthermore, often the noise itself actually slows down training further or even completely stalls progress because it makes it more likely that luck will play a role in the performance of candidate controllers. In other words, the gradient of improving performance becomes more difficult to follow for the learning algorithm due to the added noise in simulation.

Because the aim in this project is accordingly to train robots in simulators that will be transferred to the real world, our second goal is to develop an alternative to the expensive and unreliable approach of training with multiple noisy trials. In particular, we focused on designing an entirely new technique, inspired by the robustness of organisms in nature, for encouraging the robustness of trained controllers that would not require averaging over multiple noisy trials. Initial results in this initiative were published at the end of the first year, and a more extensive journal submission was completed in the second and finally published after revisions and improvements in the third year.

Summary of the most important results

First, initial investigation of the hive brain focused on establishing a preliminary communication architecture that facilitate coordination. Figure 1 in the attachment shows an example architecture. In this example, the "transmit" layer sends signals to the "receive" layer of adjacent agents in the hive setup. The receive layer then processes incoming transmissions and send the result to the hidden layer, where the agent decides what action to take. One challenge with determining hive connectivity is that agents in the world do not necessarily stay in the same geometric order as their hive architecture assumes. Thus establishing ordering guarantees during deployment (e.g. such that agents 1 is always near agent 2, agent 2 is always near agent 3, and so forth) significantly improves performance.

This architecture was optimized by multiagent HyperNEAT successfully to produce a working hive in a patrol synchronization task that requires communication because the robots do not see each other. In this task, a team of robots that are patrolling in an oscillatory left-right pattern within an enclosed area must gradually align such that they are moving in tandem. Evolved controllers that were trained in simulation were transferred to real Khepera III robot teams, which maintained the successful performance on the task. Seeing the real robots in action in this domain helps to illustrate the potential of the hive; for this purpose videos of the robot team with explanations are available at:

<http://eplex.cs.ucf.edu/demos/hive-brain-patrol>

A picture of the robot team performing the patrol task is also included in figure 2 of the attached file.

This result was published in the paper, "Multirobot Behavior Synchronization through Direct Neural Network Communication," at the 5th International Conference on Intelligent Robotics and Applications (ICIRA-2012), where it was one of five best paper finalists out of 198 submissions. This work provided the first working real-world prototype of the hive.

An important insight in the second year was that geometric coordination among communicating agents benefits from an ability to determine the directionality of a signal, i.e. the direction from which it originated. If such a sense is available, then agents do not need to encode independent clues to their location as part of their communication signals, thereby opening up communication to serve more sophisticated purposes and allowing coordinated behaviors to work with fewer a priori assumptions about the initial configuration of the team. For example, it becomes significantly easier to request that other agents come to a particular location because they can sense from where the signal originates.

To facilitate such directional communication, we reconfigured the standard hive architecture (figure 1 in the attachment) to include multiple directional communication inputs for each individual. However, interestingly, such additional information expands the size of the neural network substrate, which also includes inputs for sensing locations of teammates, walls, and other objects. A schematic of this new large substrate is shown in its entirety in figure 3 in the attachment. As the figure shows,

the substrate is becoming increasingly complex with multimodal inputs. Thus an important research contribution of the second year that laid the foundation for further research in directional communication was to develop a new method of encoding multimodal substrates to facilitate evolving them. This new method was published at GECCO-2013 with the title, "Evolving Multimodal Controllers with HyperNEAT."

At the same time, the ability now to know the origin of a signal opened up a new research direction in the advantages and disadvantages of different kinds of directional communication. For example, agents might not only know the direction of a signal, but also the identity of the agent transmitting it. Or they could be restricted only to knowing the sending agent's identity but not its direction. A comprehensive study of several such variants was published in the third year at GECCO-2014, titled, "Directional Communication in Evolved Multiagent Teams." One important result of this study is confirmation that some forms of communication (e.g. non-directional) can be no better than having no communication at all in some domains. Thus the critical role of directionality is confirmed.

Videos of real robots using direction communication are at:
<http://tinyurl.com/DirComVideo>

As the third and final year drew to a close we began applying this more powerful form of communication to more complicated geometric agent coordination problems. For example, in one challenging domain with potential DoD applicability, a team of robots must cover a number of critical points on a field but also occasionally come to the rescue of fellow robots who run out of power. The ability to maintain coverage while still keeping all robots operating requires sophisticated communication that preliminary results show the hive can learn to implement. Continuing to apply the new technologies developed in this project to increasingly complex multiagent coordination problems will remain a continuing direction of research in our group.

Second, we also succeeded in developing an entirely new method for ensuring robustness in simulation-to-real-world transfer. This approach, called reactivity, proved capable of eliminating the need for multiple trials of noise in four different robot maze-navigation scenarios (shown in figure 4 of the attached file), thereby reducing the number of trials needed for training by a factor of eight without losing any reliability. This approach is significant because it provides an entirely new avenue for thinking about how robots can be trained effectively in simulation. Because it is inspired by observations of the robust behavior of organisms in nature, we published a paper on the initial result ("Rewarding Reactivity to Evolve Robust Controllers without Multiple Trials or Noise") at the 13th International Conference on the Simulation & Synthesis of Living Systems (Alife 13, though it took place in 2012), which had only a 25% acceptance rate for oral presentations. In the second year we completed significantly more extensive investigations of reactivity that included extensive real-world tests. This expanded work was published in a comprehensive study in *Adaptive Behavior* journal in December 2013 after revisions and improvements in the third year.

To briefly summarize the idea behind reactivity, whereas the idea behind training with multiple noisy trials is that the problem of transfer requires a highly refined model of the actual environmental conditions, our hypothesis is that instead it may work better to assume that the model of the environment is always poor and therefore the robot should continually seek out additional information about the environment. That way, the robot can never assume that its sensors are really telling the truth. This information-seeking behavior is called "reactivity" because it encourages the robot to demonstrate a tendency to react to and seek out changes in its sensors. By rewarding evolving agents in part for demonstrating such behavior (i.e. as one objective in a multiobjective algorithm), they can be encouraged to be robust even though they are never subjected to multiple noisy trials.

Third, as a supporting effort behind the primary directions of the project, we also enhanced the core algorithm suite behind multiagent HyperNEAT, which improves its capabilities and robustness in general. For this purpose, we published a paper on an enhancement of the underlying HyperNEAT algorithm, "A Unified Approach to Evolving Plasticity and Neural Geometry," which won the Best Student Paper Award at the International Joint Conference on Neural Networks (IJCNN 2012) out of 299 total papers submitted with student first authors. A related 2012 work called "An Enhanced Hypercube-Based Encoding for Evolving the Placement, Density and Connectivity of Neurons," in *Artificial Life* journal, expands on this idea. We also explored the theoretical underpinnings of algorithms such as HyperNEAT that aim for high-level complexity in another paper at Alife 13 called, "Beyond Open-endedness: Quantifying Impressiveness."

In 2013, a comprehensive journal paper on the multiagent HyperNEAT algorithm was published, called "Scalable Multiagent Learning through Indirect Encoding of Policy Geometry." An exploration of the multiagent HyperNEAT approach in a quadruped walking task (where each leg is treated as an "agent") appeared in GECCO 2013, and a general study on evolvability in evolutionary algorithms appeared in the high-impact PLoS journal.

A persistent question throughout the work of this project was whether agents might be able to continue to adapt and change their policies after they are already deployed in the real world (i.e. after the formal training period is completed). An important factor in such a capability, and one studied in recent years in the field of deep learning in offline tasks, would be the ability to learn new features in real time, as the agents navigate the environment. Our paper, "Real-time Hebbian Learning from Autoencoder Features for Control Tasks," which appeared at the Fourteenth International Conference on the Synthesis and

Simulation of Living Systems (ALIFE XIV) in 2014, is the first to demonstrate that such features can indeed be learned online, while the agents acts in the environment.

In summary, 12 publications have resulted from this project. One such publication won an award out of 299 entries and another was one of five finalists for best paper out of 198 entries. Major achievements include successful demonstrations of the hive brain in the real world, the introduction of a new method for learning from multimodal sensors, a study of the contribution of a sense of directionality to communication, a new method for training for robustness without the need for multiple trials, and multiple supporting enhancements of the underlying HyperNEAT algorithmic infrastructure.

Bibliography

David B. D'Ambrosio, Skyler Goodell, Joel Lehman, Sebastian Risi, and Kenneth O. Stanley (2012). Multirobot Behavior Synchronization through Direct Neural Network Communication. To appear in: Proceedings of the 5th International Conference on Intelligent Robotics and Applications (ICIRA-2012). New York, NY: Springer-Verlag, 2012 (12 pages).

Joel Lehman, Sebastian Risi, David B. D'Ambrosio, and Kenneth O. Stanley (2012) Rewarding Reactivity to Evolve Robust Controllers without Multiple Trials or Noise. In: Proceedings of the Thirteenth International Conference on Artificial Life (ALIFE XIII). Cambridge, MA: MIT Press, 2012 (8 pages).

Sebastian Risi and Kenneth O. Stanley (2012). A Unified Approach to Evolving Plasticity and Neural Geometry. In: Proceedings of the International Joint Conference on Neural Networks (IJCNN 2012). Piscataway, NJ: IEEE, 2012 (8 pages). Winner of the Best Student Paper Award (out of 299 other entries with student first authors)

Joel Lehman and Kenneth O. Stanley (2012). Beyond Open-endedness: Quantifying Impressiveness. In: Proceedings of the Thirteenth International Conference on Artificial Life (ALIFE XIII). Cambridge, MA: MIT Press, 2012 (8 pages).

Sebastian Risi and Kenneth O. Stanley (2012). An Enhanced Hypercube-Based Encoding for Evolving the Placement, Density and Connectivity of Neurons. In: Artificial Life journal, volume 18, number 4. Cambridge, MA: MIT Press, 2012 (Manuscript 54 pages).

David B. D'Ambrosio and Kenneth O. Stanley (2013). Scalable Multiagent Learning through Indirect Encoding of Policy Geometry. In: Evolutionary Intelligence Journal 6(1). New York, NY: Springer-Verlag, 2013 (Manuscript 30 pages).

Justin K. Pugh and Kenneth O. Stanley (2013). Evolving Multimodal Controllers with HyperNEAT. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013). New York, NY: ACM, 2013 (8 pages).

Sebastian Risi and Kenneth O. Stanley (2013). Confronting the Challenge of Learning a Flexible Neural Controller for a Diversity of Morphologies. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2013). New York, NY: ACM (7 pages).

Joel Lehman and Kenneth O. Stanley (2013). Evolvability is Inevitable: Increasing Evolvability Without the Pressure to Adapt. In: PLoS ONE journal. (9 pages).

Joel Lehman, Sebastian Risi, David B. D'Ambrosio and Kenneth O. Stanley (2013). Encouraging Reactivity to Create Robust Machines. In: Adaptive Behavior journal 6 (21). London: SAGE, 2013 (17 pages).

Justin K. Pugh, Skyler Goodell, and Kenneth O. Stanley (2014). Directional Communication in Evolved Multiagent Teams. In: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2014). New York, NY: ACM, 2014 (8 pages).

Justin K. Pugh, Andrea Soltoggio, and Kenneth O. Stanley (2014). Real-time Hebbian Learning from Autoencoder Features for Control Tasks. In: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems (ALIFE XIV). Cambridge, MA: MIT Press, 2014 (8 pages).

Technology Transfer

On June 20th, 2014, we held an online videoconference with Stuart Young of ARL and his team, which included an overview of the project and its accomplishments, as well as technical details on the underlying algorithmic techniques that may be useful to ARL in the future. We hope this exchange of expertise and opportunity to answer questions can form the basis of further collaboration in the future.

It is also important to note that previous to this project, our group sent Stuart Young's group at ARL HyperNEAT-evolved neural networks to test in real ARL Packbots, which then successfully navigated corridors at ARL with their HyperNEAT controllers. The implication is that evolved multiagent strategies from this project can in principle transfer to real robots like the ones used at ARL. The simulator in which training was implemented in this project (which was built by our group in-house) includes a model of the Packbot in addition to the Khepera III robots that we used for real-world demonstrations over the course of the project.




Figure 1. Hive Brain Substrate Architecture. The individual neural network controllers of three separate agents are shown. The hive substrate includes input, output, and hidden layers. However two of the hidden layers are designated as transmitting and receiving layers that are used for communication. The flow of information between agents is shown by the dashed lines. The inputs are the left and right sensors and the output is interpreted as a motor command



Figure 2. Hive Robots in the Patrol Domain. The goal is to patrol back and forth in the enclosure until all the robots are horizontally synchronized. The robots can only see the wall and not each other, which is why the task cannot be solved without communication. Real Khepera III robots are shown performing the task with the hive brain architecture. They are communicating neural signals over wireless connections.




Figure 3. Layout of New Multiagent HyperNEAT Directional Communication Neural Network Substrate. The directional inputs are shown as two rows in the second plane from the right on the bottom (highlighted in red), one for signals from the front and the other for signals from the rear. With this representation, the agent can discern from where the signal originates based on which input in which row receives the communication signal. Other input planes (bottom planes) are for detecting a target, other friendly agents, and walls.




Figure 4. Mazes. The goal of the agent in the maze navigation domains is to navigate from the starting position (large circle) to the goal (small circle). Note that mazes are not drawn to scale.