Neuro-Inspired Dynamic Replanning in Swarms—Theoretical Neuroscience Extends Swarming in Complex Environments

Grace M. Hwang, Kevin M. Schultz, Joseph D. Monaco, and Kechen Zhang

ABSTRACT

In the NeuroSwarms framework, a team including researchers from the Johns Hopkins University Applied Physics Laboratory (APL) and the Johns Hopkins University School of Medicine (JHM) applied key theoretical concepts from neuroscience to models of distributed multi-agent autonomous systems and found that complex swarming behaviors arise from simple learning rules used by the mammalian brain.

INTRODUCTION

The development of multi-agent platforms with small-scale robotic vehicles is an exciting target of state-of-the-art autonomous systems engineering: many new applications may emerge from controlling large, distributed groups of inexpensive but agile vehicles. However, current communication and control frameworks need to be improved to provide the adaptiveness, resilience, and computational efficiency required for operating in complex and dynamically changing real-world conditions.[1–4] This project explored concepts in theoretical neuroscience to bridge the collective interactions and timescales of brain activity and animal behavior to emergent spatial and temporal patterns of groups of mobile autonomous agents. Fast Company, who recognized APL as one of the top three best workplaces for innovators, highlighted this project in an article and accompanying video about its 2020 list.[5]

In our operating metaphor, autonomous agents are neurons, agent-based communication is the phase synchronization of neuronal spiking, and the swarm as a whole is a neuronal network in which emergent network behaviors map to emergent swarming behaviors. Although neural representations of the hippocampal formation have motivated prior approaches to spatial mapping, planning, and navigation of robotic platforms,[6–9] these neuromimetic approaches have relied on the representations from spatial neurons—including place cells, head direction cells, border cells, and/or grid cells[10]—to drive spatial computations in support of single-platform robotic control.[11–14] However, these approaches do not apply temporal coding mechanisms from the recently discovered phaser cells[15,16] or spatial self-organization concepts from attractor map theory.[17–19] It has remained unclear how the spatiotemporal dynamics of these neural representations might inform advances in autonomous control.

Prior work has also applied biomimetic approaches to swarming problems, which require collective behaviors to accomplish spatially distributed tasks. One such approach, inspired by animal groups with oscillatory communication patterns, was generalized as the “swarmalators” formalism,[20] in which an agent’s internal phase is governed by local Kuramoto synchronization and swarming attraction and repulsion are phase-coupled.

However, swarmalator systems naturally relax to static states or simple cycling behaviors. Thus, our team, made up of researchers from APL and the Johns Hopkins School of Medicine (JHM), investigated how the spatiotemporal dynamics of a subset of spatial neurons (i.e., place cells and phaser cells) might drive useful navigational behaviors in distributed groups of mobile oscillators via swarming.[21]

The team demonstrated, for the first time, that spatial attractor dynamics and temporal phase-based organization can be driven, in parallel, by a form of Hebbian learning modified to operate on, and indirectly control, inter-agent distances.[22] Further, the link from learning to swarming is a fast online process, unlike existing pretrained or slowly adapting neural network controllers. In this article, we discuss the NeuroSwarms controller framework with analogies to neuroscience and provide example demonstrations. We also discuss how neuroscience concepts of oscillatory phase coding and generalization of phase states to computationally relevant manifolds further inspired our development of a metastable swarming framework—i.e., Stiefelators.

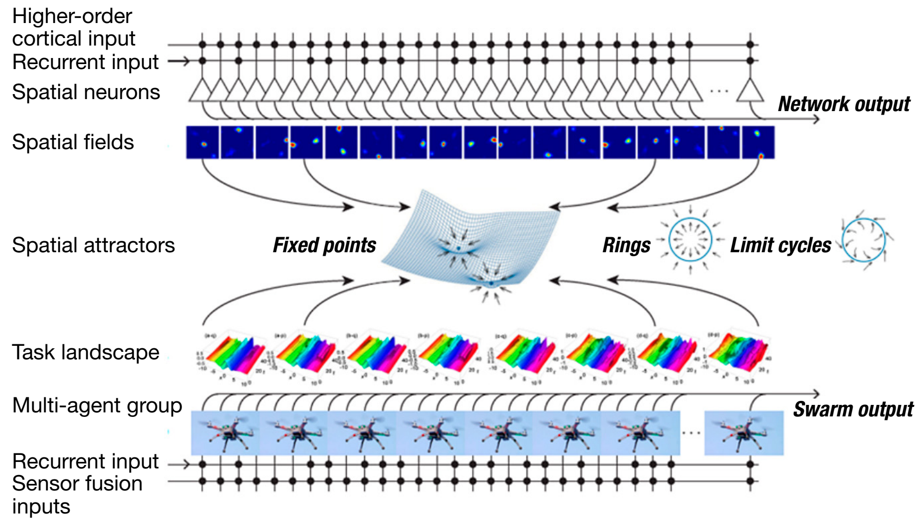

Figure 1. Brain-to-swarm analogy. A recurrent hippocampal neuronal network and its place fields are shown to collectively form spatial attractors. Other types of spatial neurons and network connectivity could be applied to produce rings, or limit cycles, of activity. In NeuroSwarms, distances between each artificial agent within the multi-agent group is updated at each time step of the simulation while agents locate rewards. Other task landscapes could be applied such that collectively these artificial fields could participate in attractor dynamics in the local communication subnetworks of swarms. This operating analogy suggests that high levels of distributed control and autonomy could be supported by ~1-bit "spike-phase" channels with minimal energy footprints to achieve a range of swarm functions (e.g., load balancing, consensus, dynamic replanning).

On the basis of decades of neuroscience research, hippocampal place cells are known to fire within a contiguous region of the animal's learned environment, or place field.[23] Our key insight was that an individual agent could be represented as a spatial neuron (e.g., a place cell) whose preferred location, or place field, indicates the agent's desired position in the environment. It thus follows that a multi-agent group would be analogous to a neuronal network (e.g., the recurrently connected place cells of hippocampal subregion CA3; Figure 1). Connections between neurons may be characterized by the "synaptic weight" that acts as a multiplicative gain on neuronal inputs. We thus further suppose that mutually visible agent pairs are reciprocally connected and that the distance between these agent pairs maps to the symmetric synaptic weight of those connections. Consequently, relative agent motion corresponds to changes in connectivity and weights. Thus, a spatial configuration of the group constitutes an attractor map network[17–19] and relative motion (i.e., swarming) constitutes learning based on synaptic modification.[24,25] Put simply, swarming motion can be treated as learning based on synaptic plasticity within a memory network.

METHODS AND RESULTS

NeuroSwarms

Applying methods described in our previous work that implemented a rate-based implementation of NeuroSwarms,[22] we observed several novel, yet unexpected, emergent dynamical behaviors in simulations of both multi-agent swarming (Figure 2a) and the single-entity reward approach (Figure 2b). The most notable and persistent behaviors included the emergence of phase-sorted spatial formations (as shown by the color-sorting of adjacent agents) such as line segments, rings, or concentric loops. When we added a reward memory to the single-entity NeuroSwarms implementation, the single-entity was able to secure each of the three rewards (shown as yellow stars in Figure 2) in 30 out of 40 trials.

Stiefelators

We developed a swarming framework for a distributed "cocktail party" problem, wherein swarming agents are searching for emitters and simultaneously attempting to both localize the emitters in physical space and isolate them in spectral space. We exploited the fact that the subspace spanned by a signal is an element on a Grassmannian manifold, and thus is the natural auxiliary space to generalize our swarmalator model to generate a metastable system (Figure 3a, multimode) that had substantially improved performance over the baseline generalized swarmalators (Figure 3a, single mode) in a dynamic environment. Metastability and load balancing are also demonstrated while the multimode Stiefelator agents emergently alternate on three targets located at X = 0.95, –0.1, and –0.5 (Figure 3b).

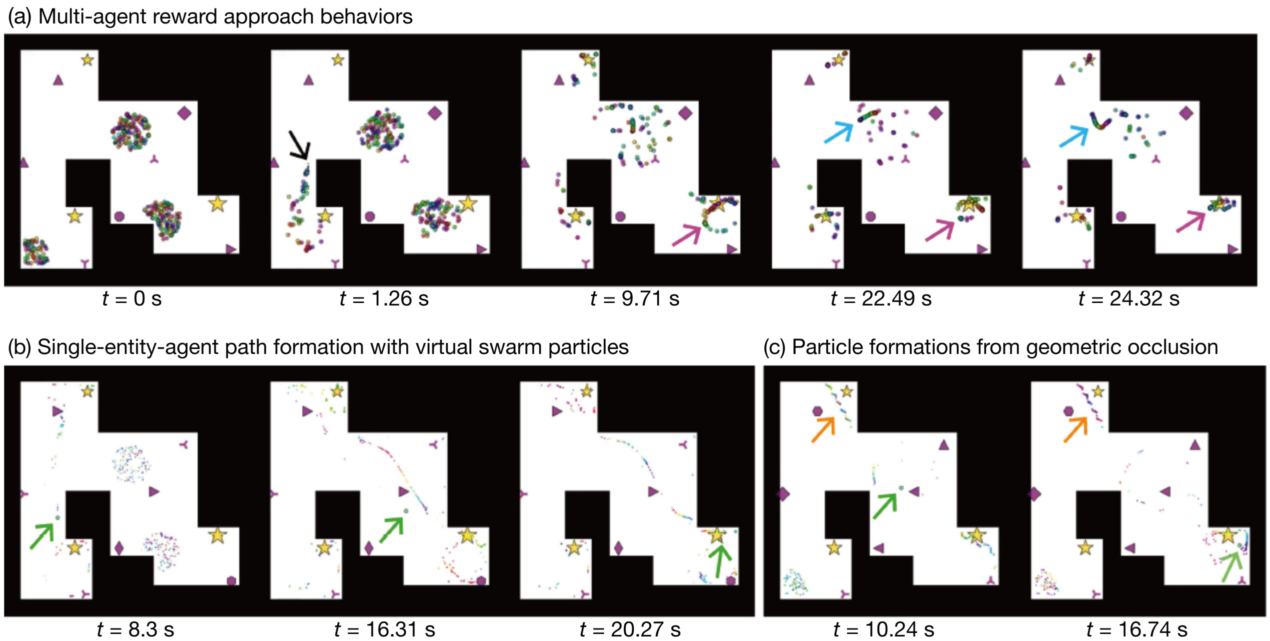

Figure 2. Temporal evolution of swarming and single-entity approaches to rewards. (a) Three agent-clusters were initially populated in the multi-reward arena. The internal place-field location of each agent is indicated by a small black dot (e.g., t = 1.26 s, black arrow). Phase sorting is indicated by sequentially ordered colors of the circle markers representing agent positions. A reward-centered phase ring was created (t = 9.71 s) with a decreasing diameter over time (t = 22.49 s and t = 24.32 s; magenta arrows). A phase-sorted line segment formed and moved around a corner (t = 22.49 s and t = 24.32 s; blue arrows). (b) A single-entity agent (larger green circle with green arrow) was guided by 300 virtual particles (phase-colored dots). Swarm particles formed phase sequences leading the agent from the southwest corner to the reward location in the southeast corner of the arena by t = 20.3 s. (c) Step-like patterns of particles (orange arrows) appeared near rewards that were occluded from the perspective of the single agent (green arrows) by corners in the environmental geometry. While the agent became "indecisive" around t = 10.24 s because it was pulled simultaneously in both directions, the agent ultimately found its way to the southeast reward by t = 16.74 s.

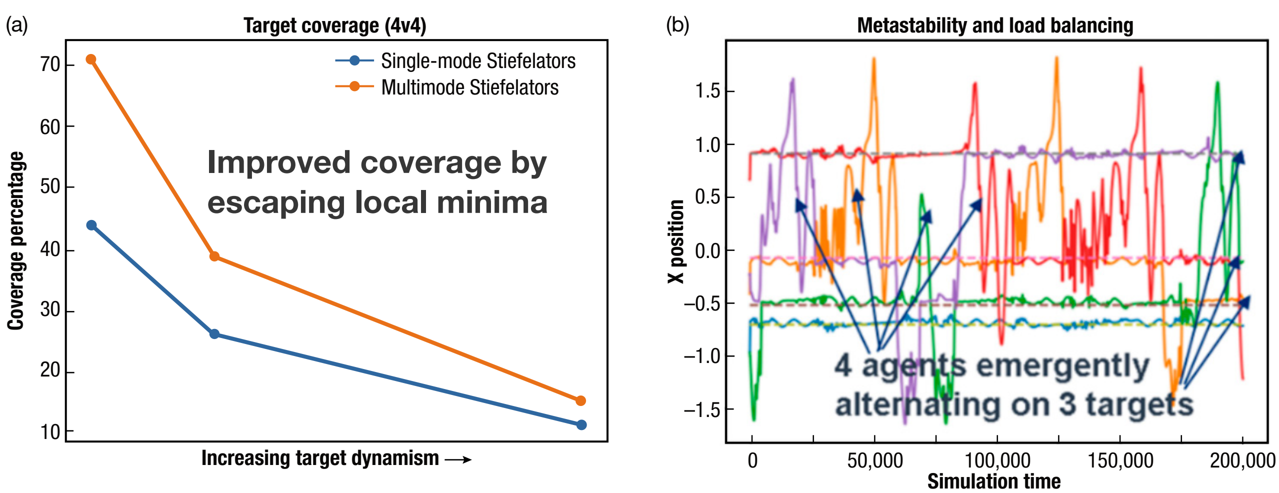

Figure 3. Cocktail party simulation results. (a) The addition of an auxiliary variable for search that is modulated by the commonality with neighbors' Grassmannian variables results in improved ability to escape local minima. (b) Demonstration of metastable behavior exhibited by four swarming agents.

OUTLOOK

Encouraged by our promising preliminary results on NeuroSwarms and Stiefelators, future work includes the addition of alternative learning rules, conversion from rate-based dynamics to spiking communications, the incorporation of sharp-wave/ripple[26–28] dynamics to accommodate multi-agent deliberation, and the use of graph signal processing[29] for stability analysis. Current funding (through an APL Propulsion Grant, a program that aims to advance bold, high-risk, and transformational ideas) should allow us to evaluate structure and dynamics on time-varying graphs. Collectively, if these next steps are successful, NeuroSwarms may become a generalized controller for dynamic replanning in unstructured environments in GPS-denied areas.

CONCLUSIONS

By analogizing agents and swarms to neurons and networks, we showed that a high-level neural approach to distributed autonomous control produces complex dynamics with navigational value. This analogy permitted the tools of theoretical neuroscience to be leveraged in developing a model controller of an artificial swarming system. Our key insight was that swarm motion can be interpreted as a mobile variation of Hebbian learning, given a natural translation between spatial relationships in a swarm and connectivity relationships in a neuronal network. This insight that "swarming is learning" further allowed us to demonstrate advances in generalized swarmalator systems to solve the distributed cocktail party problem by achieving metastability, which is a preliminary form of dynamic replanning that we will further pursue.

ACKNOWLEDGMENTS: This material is based on work supported by (while serving at) the National Science Foundation. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This work was supported by National Science Foundation Award NCS/FO 1835279, APL independent research and development awards, and the Johns Hopkins University Kavli Neuroscience Discovery Institute.

REFERENCES

[1]R. M. Murray, "Recent research in cooperative control of multivehicle systems," *J. Dyn. Syst., Meas., Control.*, vol. 129, no. 5, pp. 571–583, 2007, https://doi.org/10.1115/1.2766721.

[2]G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, et al., "The grand challenges of science robotics," *Sci. Robot.*, vol. 3, no. 14, eaar7650, pp. 1–14, 2018, https://doi.org/10.1126/scirobotics.aar7650.

[3]H. Hamann, K. Yara, B. Jean, D. S. Mohammad, F. Eliseo, et al., "Hybrid societies: Challenges and perspectives in the design of collective behavior in self-organizing systems," *Front. Robot. AI*, vol. 3, pp. 1–14, 2016, https://doi.org/10.3389/frobt.2016.00014.

[4]S.-J. Chung, A. A. Paranjape, P. Dames, S. Shen, and V. Kumar, "A survey on aerial swarm robotics," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 837–855, 2018, https://doi.org/10.1109/TRO.2018.2857475.

[5]J. Woodruff, "Johns Hopkins University APL is one of Fast Company's Best Workplaces for Innovators" (article and video), Fast Company, Jul. 29, 2020, https://www.fastcompany.com/90529833/best-workplaces-for-innovators-2020-johns-hopkins-university-apl.

[6]M. J. Milford, G. F. Wyeth, and D. Prasser, "RatSLAM: A hippocampal model for simultaneous localization and mapping," in *Proc. ICRA'04: 2004 IEEE Int. Conf. on Robot. Automat.*, 2004, vol. 1, pp. 403–408, https://doi.org/10.1109/ROBOT.2004.1307183.

[7]N. Cuperlier, M. Quoy, and P. Gaussier, "Neurobiologically inspired mobile robot navigation and planning," *Front. Neurorobot.*, vol. 1, no. 3, pp. 1–15, 2007, https://doi.org/10.3389/neuro.12.003.2007.

[8]M. J. Milford and G. F. Wyeth, "Mapping a suburb with a single camera using a biologically inspired SLAM system," *IEEE Trans. Robot.*, vol. 24, no. 5, pp. 1038–1053, 2008, https://doi.org/10.1109/TRO.2008.2004520.

[9]A. Barrera and A. Weitzenfeld, "Biologically-inspired robot spatial cognition based on rat neurophysiological studies," *Auton. Robot.*, vol. 25, no. 1–2, pp. 147–169, 2008, https://doi.org/10.1007/s10514-007-9074-3.

[10]S. Poulter, T. Hartley, and C. Lever, "The neurobiology of mammalian navigation," *Curr. Biol.*, vol. 28, no. 17, pp. R1023–R1042, 2018, https://doi.org/10.1016/j.cub.2018.05.050.

[11]M. J. Milford, J. Wiles, and G. F. Wyeth, "Solving navigational uncertainty using grid cells on robots," *PLoS Comput. Biol.*, vol. 6, no. 11, e1000995, pp. 1–14, 2010, https://doi.org/10.1371/journal.pcbi.1000995.

[12]G. Tejera, M. Llofriu, A. Barrera, and A. Weitzenfeld, "Bio-inspired robotics: A spatial cognition model integrating place cells, grid cells and head direction cells," *J. Intell. Robot. Syst.*, vol. 91, no. 1, pp. 85–99, 2018, https://doi.org/10.1007/s10846-018-0852-2.

[13]R. Kreiser, M. Cartiglia, J. N. Martel, J. Conradt, and Y. Sandamirskaya, "A neuromorphic approach to path integration: A head-direction spiking neural network with vision-driven reset," in *Proc. 2018 IEEE Int. Symp. on Circuits and Systems (ISCAS)*, 2018, pp. 1–5, https://doi.org/10.1109/ISCAS.2018.8351509.

[14]P. Gaussier, J. P. Banquet, N. Cuperlier, M. Quoy, L. Aubin, et al., "Merging information in the entorhinal cortex: what can we learn from robotics experiments and modeling?," *J. Exp. Biol.*, vol. 222, jeb186932, pp. 1–13, 2019, https://doi.org/10.1242/jeb.186932.

[15]J. D. Monaco, R. M. De Guzman, H. T. Blair, and K. Zhang, "Spatial synchronization codes from coupled rate-phase neurons," *PLoS Comp. Biol.*, vol. 15, no. 1, e1006741, pp. 1–42, 2019, https://doi.org/10.1371/journal.pcbi.1006741.

[16]J. D. Monaco, H. T. Blair, and K. Zhang, "Spatial theta cells in competitive burst synchronization networks: Reference frames from phase codes," *bioRxiv*, 211458, pp. 1–67, 2017, https://doi.org/10.1101/211458.

[17]K. Zhang, "Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory," *J. Neurosci.*, vol. 16, no. 6, pp. 2112–2126, 1996, https://doi.org/10.1523/JNEUROSCI.16-06-02112.

[18]M. Tsodyks, "Attractor neural network models of spatial maps in hippocampus," *Hippocampus*, vol. 9, no. 4, pp. 481–489, 1999, https://doi.org/10.1002/(SICI)1098-1063(1999)9:4<481::AID-HIPO14>3.0.CO;2-S.

[19]A. Samsonovich and B. L. McNaughton, "Path integration and cognitive mapping in a continuous attractor neural network model," *J. Neurosci.*, vol. 17, no. 15, pp. 5900–5920, 1997, https://doi.org/10.1523/JNEUROSCI.17-15-05900.1997.

[20]K. P. O'Keeffe, H. Hong, and S. H. Strogatz, "Oscillators that sync and swarm," *Nat. Comm.*, vol. 8, no. 1, 1504, pp. 1–13, 2017, https://doi.org/10.1038/s41467-017-01190-3.

[21]J. D. Monaco, G. M. Hwang, K. M. Schultz, and K. Zhang, "Cognitive swarming: An approach from the theoretical neuroscience of hippocampal function," in *Proc. SPIE 10982, Micro- and Nanotechnology Sensors, Systems, and Applications XI*, May 2019, vol. 10982, pp. 109822D-1–109822D-10, https://doi.org/10.1117/12.2518966.

[22]J. D. Monaco, G. M. Hwang, K. M. Schultz, and K. Zhang, "Cognitive swarming in complex environments with attractor dynamics and oscillatory computing," *Biol. Cybernet.*, vol. 114, no. 2, pp. 1–16, 2020, https://doi.org/10.1007/s00422-020-00823-z. Source code available at https://doi.org/10.5281/zenodo.4281048.

[23]J. O'Keefe and J. Dostrovsky, "The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat," *Brain Res.*, vol. 34, no. 1, pp. 171–175, 1971, https://doi.org/10.1016/0006-8993(71)90358-1.

[24]D. O. Hebb, *The Organization of Behavior*, New York: John Wiley and Sons, Inc., 1949.
[25]E. Oja, "Neural networks, principal components, and subspaces," *Int. J. Neural Syst.*, vol. 1, pp. 61–68, 1989, https://doi.org/10.1142/S0129065789000475.
[26]B. E. Pfeiffer and D. J. Foster, "Autoassociative dynamics in the generation of sequences of hippocampal place cells," *Science*, vol. 349, no. 6244, pp. 180–183, 2015, https://doi.org/10.1126/science.aaa9633.
[27]B. E. Pfeiffer and D. J. Foster, "Hippocampal place-cell sequences depict future paths to remembered goals," *Nature*, vol. 497, no. 7447, pp. 74–79, 2013, https://doi.org/10.1038/nature12112.
[28]B. E. Pfeiffer, "The content of hippocampal 'replay,'" *Hippocampus*, vol. 30, no. 1, pp. 6–18, 2017, https://doi.org/10.1002/hipo.22824.
[29]K. Schultz and M. Villafane-Delgado, "Graph signal processing of indefinite and complex graphs using directed variation," *arXiv:2003.00621*, 2020, https://arxiv.org/abs/2003.00621.

Grace M. Hwang, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Grace M. Hwang is a project manager in APL's Research and Exploratory Development Department. She is also a member of the Johns Hopkins Kavli Neuroscience Discovery Institute. She has a BS in civil engineering from Northeastern University, an MS in civil and environmental engineering from the Massachusetts Institute of Technology (MIT) and an MS and a PhD in biophysics from Brandeis University. Grace has an extensive background in experimental design, computational neuroscience, biophysics, biosensors, biomarker discovery, and instrumentation. She provides neuroscience expertise to brain–computer interface (BCI) projects including noninvasive ultrasound neuromodulation in rodents with the goal of accessing deep regions of the brain with spatial selectivity to recover neuronal plasticity from cell-specific brain circuits. She conceptualized the initial National Science Foundation (NSF) grant that drew inspirations from neural science for the control of self-organized multi-agent autonomous systems. She wrote this manuscript and generated results for Figure 2. She is a member of the Society for Neuroscience. Her email address is grace.hwang@jhuapl.edu.

Kevin M. Schultz, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Kevin M. Schultz is a project manager and senior staff scientist in APL's Research and Exploratory Development Department. He has a BS, an MS, and a PhD in electrical and computer engineering, all from the Ohio State University. His research interests include the control, characterization, and modeling of quantum systems, as well as distributed control and signal processing applications outside of quantum information. His most recent research focused on characterization and control problems in quantum information, including experimental validation of characterization protocols, neuromimetic approaches to distributed control problems, and applications of graph signal processing to emergence in distributed systems. As a principal investigator or project manager, Kevin has executed research funded by various sponsors, as well as numerous internal efforts His email address is kevin.schultz@jhuapl.edu.

Joseph D. Monaco, Johns Hopkins University School of Medicine, Baltimore, MD

Joseph D. Monaco is a computational neuroscientist who uses theory, modeling, and analysis to address questions of biological function and mechanism arising from a dynamical perspective of neural circuits and behavior. He has a BA in cognitive science and a BA in mathematics from the University of Virginia and a PhD in neurobiology and behavior from Columbia University. His research has focused on spatial cognition and memory in mammals, including the cells, circuits, neural codes, and dynamics of the hippocampus and related structures. His current research focuses on models of memory reactivation during waking and sleeping states to help understand the neural dynamics that are crucial to lifelong learning and behavioral adaptation. In collaboration with APL researchers, he has been developing potential technological applications for these neural dynamics in self-organized control of artificial autonomous systems. In particular, he created the NeuroSwarms theory and associated source code. His email address is jmonaco@jhu.edu.

Kechen Zhang, Johns Hopkins University School of Medicine, Baltimore, MD

Kechen Zhang is an associate professor of biomedical engineering at the Johns Hopkins University School of Medicine. He has a BS in biophysics and physiology, an MS in neurobiology from Peking University, and a PhD in cognitive science from the University of California San Diego. He completed a postdoctoral fellowship in computational neurobiology at Salk Institute. The research in his laboratory focuses on theoretical/computational neuroscience using mathematical analysis and computer simulations to study the nervous system at multiple levels, from realistic biophysical models to simplified neuronal networks. Several of the lab's current research projects involve close collaborations with experimental neuroscience laboratories. His email address is kzhang4@jhmi.edu.