Motifs to Models: Leveraging Biological Circuits toward Novel Computational Substrates

Elizabeth P. Reilly, Erik C. Johnson, Jordan K. Matelsky, Lucas J. Ziemba, Paul D. Hermann, Morgan V. Schuyler, Joan A. Hoffmann, and William R. Gray-Roncal

ABSTRACT

The Motifs to Models team at the Johns Hopkins University Applied Physics Laboratory (APL) leverages the existence proof provided by biological circuitry—of robustness, adaptability, and low-sample learning at very low size, weight, and power—to explore novel computational substrates toward critical sponsor needs in computation and artificial intelligence.

Achieving transformative capabilities through the development of novel computational models, such as those inspired by examination of neural circuits, is a core critical opportunity across diverse Department of Defense (DOD) domains. Current artificial intelligence and machine learning approaches often rely on a small number of heuristically constructed networks rather than exploiting a principled design strategy. Similarly, prototype neuromorphic processor architectures (e.g., Intel Loihi, IBM TrueNorth) have the capacity to statistically simulate individual neurons but draw little architectural inspiration from actual brains. In the work described in this article, we examine biological intelligence for inspiration. We do so by representing the brain as a graph where nodes are neurons and edges are synapses between them, identifying significant computational units, or motifs, within the data, and, finally, by incorporating those motifs into new computing architectures (see Figure 1). Brains provide an existence proof for efficient, robust computation and address many fundamental challenges in understanding, such as robustness, adaptability, and low-sample learning.

Connectomics is an emerging scientific pursuit focused on mapping neuronal circuitry, at the single-

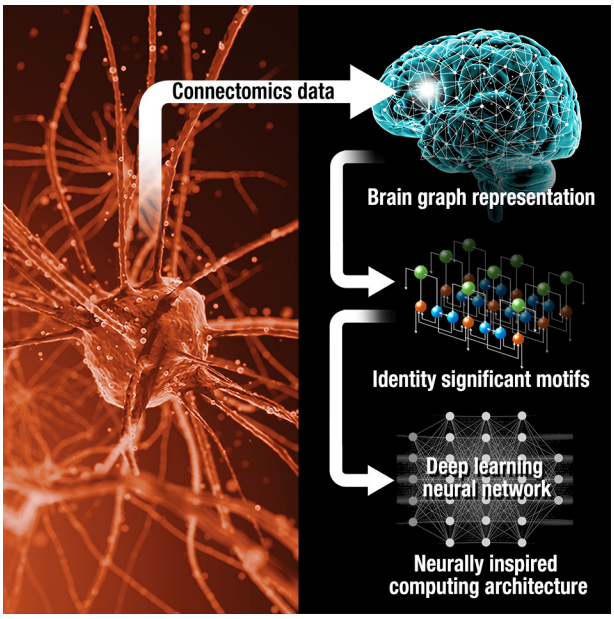

Figure 1. Motifs to Models conceptual diagram. We represent neuron-level brain data as a graph, identify significant motifs, and use those motifs to define neurally inspired neural network architectures.

neuron level, using both structural and functional connectivity information. Through internal investment, collaboration, and support for emerging data sets of unprecedented scale,[1,2] such as the structural and functional networks in the mouse and fly connectomes, APL is now positioned to discover and exploit fundamental micro- and mesoscale structure–function relationships.

Several core principles are believed to fuel the computational power and efficiency of the biological brain, including hierarchical networks, quantized computations, and learning algorithms. One underexplored concept is *motifs*—repeated computational units—that can be observed at different scales and modalities. Motifs likely account for key remaining gaps between today's artificial neural networks and future computational architectures that truly approximate the brain.

Neuroscientists have developed theories and looked for motifs within the brain at a smaller scale, often identifying candidate patterns, but without the techniques or data required to confirm a repeated layout.[3–5] Other researchers[6] detected small (two- to five-node) structural and functional motifs, testing the hypothesis that the brain employs a small number of structural motifs for efficient encoding and assembly while maximizing functional diversity. Researchers have also focused on data-driven motif discovery using community detection rather than searching for specific circuits.[7]

APL's Motifs to Models research develops the theoretical groundwork and computational tools needed to identify significant computational motifs in neural circuits, particularly in the presence of errors and at scale. Discovered neural motifs are then translated to novel architectures by adapting state-of-the-art evolutionary architecture search algorithms. In short, the goals of this project are threefold: (1) to identify candidate motifs within connectomics data sets, (2) to verify the incidence of these motifs in large connectomics data sets, and (3) to demonstrate their algorithmic significance (e.g., on visual perception tasks).

To identify candidate motifs within inherently noisy data, the team is developing a probabilistic approach for identifying significant repeated computational units within the brain, even in the presence of errors common to automatically reconstructed networks. The approach defines a random graph model based on uncertainties within the reconstruction and counts the expected number of a particular subgraph rather than an exact count. These expected values can then be compared with those of a purely (non-data-driven) random model to see whether they appear to occur purposefully or at random. Counting triangles in a simple tripartite graph illustrates how a probabilistic approach provides better information for comparison, showing the distribution of possible count values as opposed to a single static count value without knowledge of how reconstruction errors may have biased that count (see Figure 2).

Once candidate motifs are identified, we use our deterministic motif-finding tool, DotMotif,[8] which provides a neuroscience-oriented paradigm for efficient subgraph search. See Figure 3, a–c, for a description of the subgraph search problem. Incidence of motifs can be quickly identified across large volumes of both structural and functional data for further verification of significance.

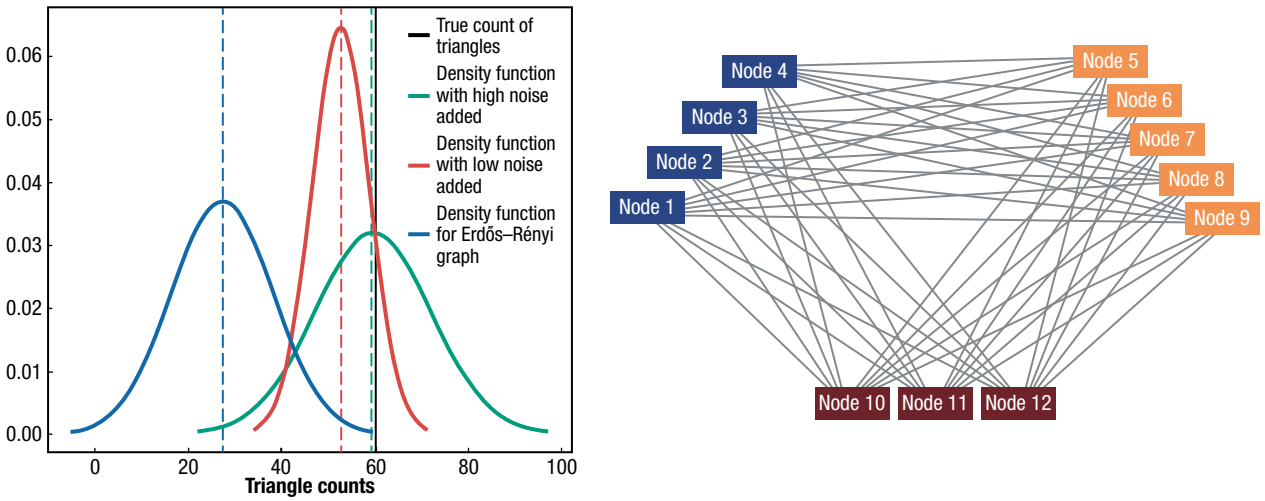

Figure 2. Demonstration of preliminary results from probabilistic motif-finding algorithm. This approach can be used to find probable computational patterns for seeding network search, even in the presence of noise. On the left, the green and red curves estimate the number of triangles in a noisy version of the tripartite graph on the right for two different levels of noise. The mean values of these curves are good approximations of the true count (the black vertical line in the plot on the left) despite significant levels of noise in the data. The blue curve represents the number of triangles that you might expect to see at random (in an Erdős–Rényi random graph model of the same density).

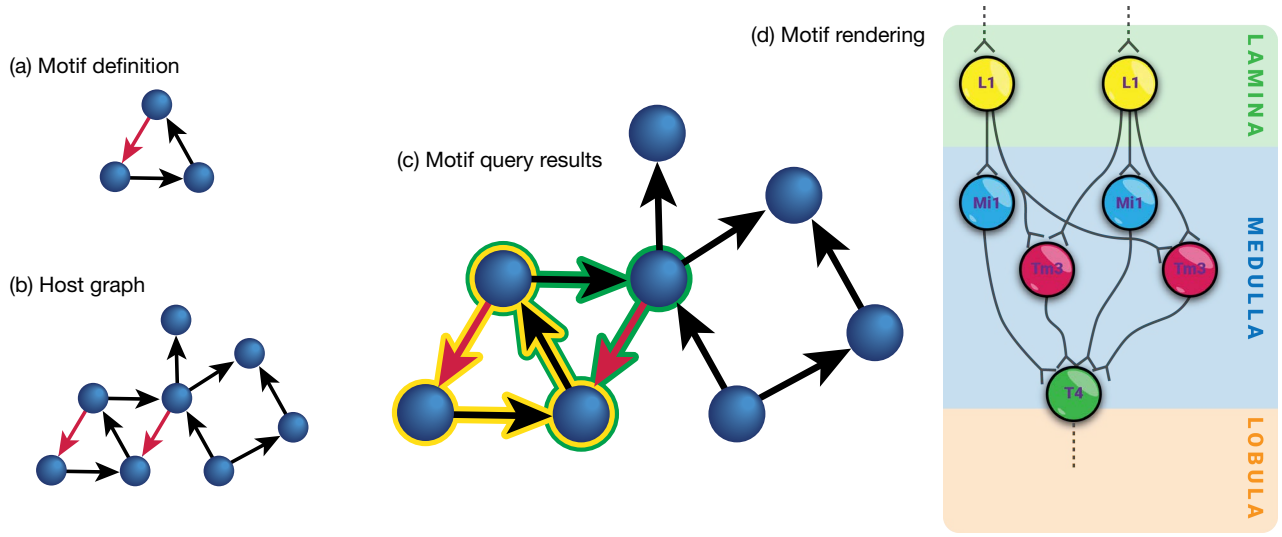

Figure 3. Subgraph search and DotMotif examples. A subgraph of interest is identified (a), as well as the host graph (b) in which the search will occur. All instances of the subgraph are identified in the host graph (c). Also shown is a rendering of a motif hypothesized by Takemura et al.[5] (d) to be important for computation in the *Drosophila* data set. Using DotMotif, the APL team was able to quickly identify all instances of this subgraph, as well as many permutations of the subgraph, to verify that it does, in fact, occur most frequently within the data set. This demonstrates a capability to automatically identify relevant computational units in large neuroanatomical data sets, which will greatly accelerate motif search.

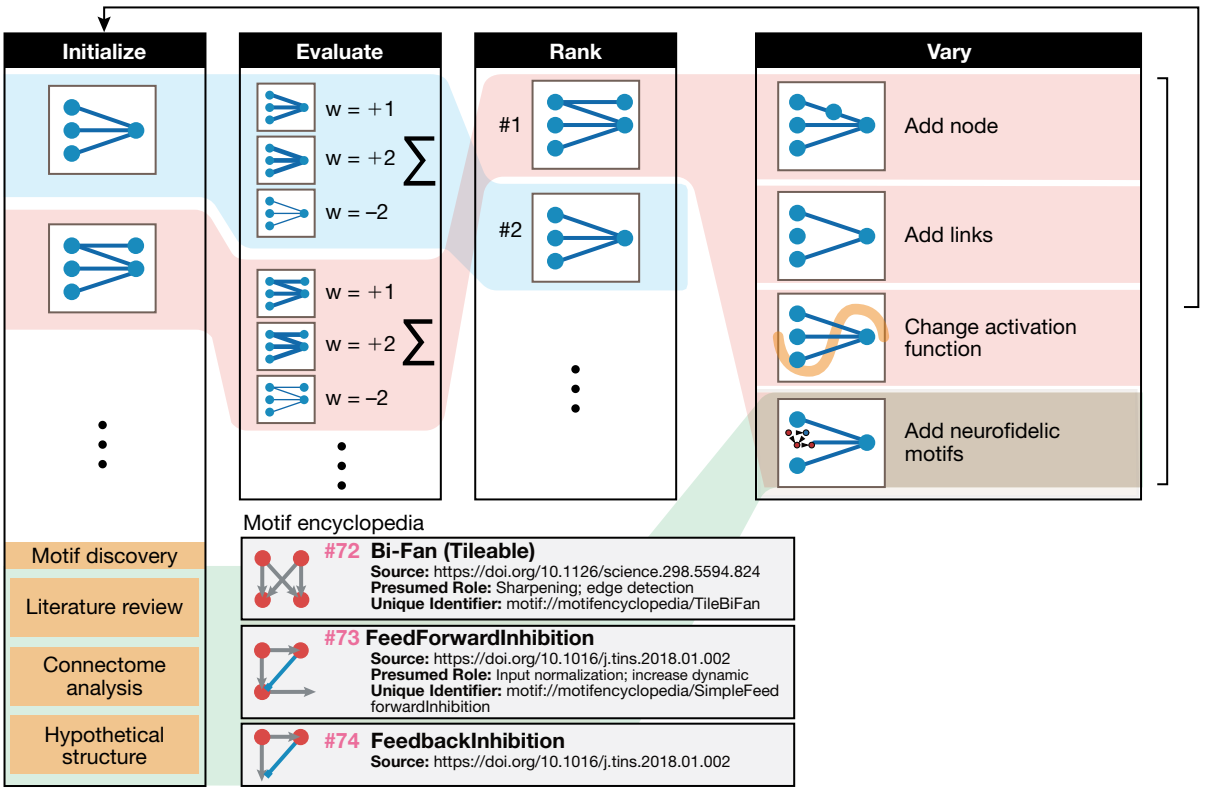

Figure 4. WANN and motifs concept. While the WANN algorithm evolves neural network architectures through simple actions like node insertion, edge addition, or change of activation function, our approach allows inclusion of more sophisticated computational units discovered in neural data using our probabilistic and deterministic approaches to motif finding.

As an example of this work, the team verified the repeated existence of the motion detection circuit in a large volume of *Drosophila* data,[5] a circuit that was previously confirmed only through intense manual labor (see Figure 3d for a rendering of the hypothesized motif of interest). Though work thus far has focused solely on motif structure, leveraging existing functional attributes of nodes and edges may further constrain or guide computational motif discovery. The emergence of these enriched connectomes will enable search through both structural and functional modalities.

Initial computational tests of candidate motifs will build on weight-agnostic neural networks (WANNs).[9] A WANN is a neural evolutionary algorithm designed to produce novel neural network architectures to solve reinforcement learning and classification tasks. It focuses on producing high-performing neural networks by iteratively evolving architecture as opposed to training weights. The initial implementation of WANNs is seeded with a small set of motifs. The APL team will seed and regularize WANNs with neurally inspired motifs to identify architectures with desirable properties of brain networks, such as robustness, efficient computing, and more. Initial results are promising, with seeded motifs appearing frequently within the resulting architectures (see Figure 4).

When applying these approaches to biological data, we hope to eventually embody human-like cognition in machines to address a variety of challenges related to knowledge acquisition, understanding, and action in complex environments. Toward this goal, the techniques described in this article, developed using internal investment, will enable exploration of topics such as (1) recurrent architectures and improved learning rules; (2) reconfigurable and adaptable networks for continual learning; (3) neuroscience rules for neuromorphic architectures; (4) improvements in size, weight, power, and robustness; and (5) an understanding of neural circuit transfer functions and their application to regularizing simulated and robotic agents. We believe this work provides a key differentiator for the future of artificial intelligence research.

REFERENCES

[1]E. P. Reilly, J. S. Garretson, W. R. Gray-Roncal, D. M. Kleissas, B. A. Wester, M. A. Chevillet, and M. J. Roos, "Neural reconstruction integrity: A metric for assessing the connectivity accuracy of reconstructed neural networks," *Front. Neuroinf.*, vol. 12, article 74, pp. 1–15, 2018, https://doi.org/10.3389/fninf.2018.00074.

[2]E. P. Reilly, E. C. Johnson, M. J. Hughes, D. Ramsden, L. Park, B. Wester, and W. Gray-Roncal "Connecting neural reconstruction integrity (Nri) to graph metrics and biological priors," in *Complex Networks XI. Springer Proceedings in Complexity*, H. Barbosa, J. Gomez-Gardenes, B. Gonçalves, G. Mangioni, R. Menezes, M. Oliveira, Eds. Springer, Cham, 2020, pp. 182-193, https://doi.org/10.1007/978-3-030-40943-2_16.

[3]A. Borst and M. Helmstaedter, "Common circuit design in fly and mammalian motion vision," *Nat. Neurosci.*, vol. 18, no. 8, pp. 1067–1076, 2015, https://doi.org/10.1038/nn.4050.

[4]O. Braganza and H. Beck, "The circuit motif as a conceptual tool for multilevel neuroscience," *Trends Neurosci.*, vol. 41, no. 3, pp. 128–136, 2018, https://doi.org/10.1016/j.tins.2018.01.002.

[5]S.-y. Takemura, A. Nern, D. B. Chklovskii, L. K. Scheffer, G. M. Rubin, and I. A Meinertzhagen, "The comprehensive connectome of a neural substrate for 'ON' motion detection in *Drosophila*," *Elife*, 6, e24394, 2017, https://doi.org/10.7554/eLife.24394.

[6]O. Sporns and R. Kötter, "Motifs in brain networks," *PLoS Biol.*, vol. 2, no. 11, e369, pp. 1910–1918, 2004, https://doi.org/10.1371/journal.pbio.0020369.

[7]V. Lyzinski, M. Tang, A. Athreya, Y. Park, an C. E. Priebe "Community detection and classification in hierarchical stochastic blockmodels," *IEEE Trans. Netw. Sci. Eng.*, vol. 4, no. 1, pp. 13–26, 2016, https://doi.org/10.1109/TNSE.2016.2634322.

[8]J. K. Matelsky, E. P. Reilly, E. C. Johnson, J. Stiso, D. S. Bassett, B. A. Wester, and W. Gray-Roncal, "DotMotif: An open-source tool for connectome subgraph isomorphism search and graph queries," bioRxiv 2020.06.08.140533, 2021, https://doi.org/10.1101/2020.06.08.140533.

[9]A. Gaier and D. Ha, "Weight agnostic neural networks," in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. San Diego, CA: Neural Information Processing Systems Foundation, 2019.

Elizabeth P. Reilly, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Elizabeth P. Reilly is a project manager and assistant group supervisor of the Artificial Intelligence Group in APL's Research and Exploratory Development Department. She has a BS in mathematics from Wake Forest University, an MA in mathematics from the University of South Carolina, and a PhD in applied mathematics and statistics from Johns Hopkins University. Elizabeth has a special interest in graph theory, connectomics, graph signal processing, and artificial intelligence for climate change. She is the co-lead of the Motifs to Models efforts and has focused on probabilistic approaches to identifying motifs in noisy networks. Her email address is elizabeth.reilly@jhuapl.edu.

Erik C. Johnson, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Erik C. Johnson is the assistant supervisor of the Neuro-AI and Connectomics Section in APL's Research and Exploratory Development Department. He has a BS, an MS, and a PhD in electrical engineering, all from the University of Illinois at Urbana-Champaign. His thesis work was on mathematical modeling of energy-constrained neural coding and information processing in sensory systems. Erik has an interest and background in computational neuroscience, signal processing, information theory, brain–machine interfaces, and robotics. Before joining APL, he worked at a consumer robotics start-up studying motion modeling, navigation, and control of wheeled mobile robots. His MS work was on off-grid sparse signal recovery with applications to array signal processing. His email address is erik.c.johnson@jhuapl.edu.

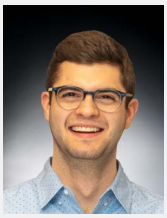

Jordan K. Matelsky, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Jordan K. Matelsky is a project manager and big-data scientist in APL's Research and Exploratory Development Department. He holds a BS in neuroscience from Johns Hopkins University. Jordan's current research focus areas include big-data neuroscience connectomics and machine learning for health care environments. His projects include large-scale network analyses on brain maps and central-nervous-system microbiome detection, and he currently serves as a precision-medicine data scientist in collaboration with the Johns Hopkins Hospital system. Before joining APL, Jordan was the chief technology officer and cofounder of FitMango, a personalized fitness software company. His email address is jordan.matelsky@jhuapl.edu.

Lucas J. Ziemba, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Lucas J. Ziemba is a software engineer in APL's Research and Exploratory Development Department. He has a BS in computer science and an MS in applied computer science, both from Dakota State University. He is currently assisting researchers working in machine perception and computer vision and developing skills in machine learning, data analysis, and software engineering. His email address is lucas.ziemba@jhuapl.edu.

Morgan V. Schuyler, Air and Missile Defense Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Morgan V. Schuyler is a mathematician in APL's Air and Missile Defense Sector. She has a BS in mathematics from the University of Maryland, Baltimore County. Morgan is experienced in working across organizational lines on a wide array of diverse programs; her specialties include development of MATLAB and Python implementations of simulation and analysis tools for infrared and electro-optical imagery, graph-based connectomic (neural) data sets, and simulated (Monte Carlo) flight trajectories. Her email address is morgan.schuyler@jhuapl.edu.

Joan A. Hoffmann, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Joan A. Hoffmann is a program manager for alternative computing paradigms in APL's Research and Exploratory Development Department. She has a BA in physics and math from Swarthmore College and an MS and a PhD in physics from the University of California, Berkeley. Joan is responsible for strategic direction and execution of a diverse portfolio encompassing research in quantum computing, trustworthy computing, neuromimetic computing, and computational physics. Over the last 4 years, she has worked with an exceptional group of investigators to build APL's core quantum information science team, now composed of nearly 20 scientists engaged in robust experimental, theoretical, and computational efforts. Her broad background in experimental physics, with emphasis on nanoscale and device physics, has led to a long track record of successful technical leadership—building and guiding strong, cross-disciplinary teams to attack critical problems. Her email address is joan.hoffmann@jhuapl.edu.

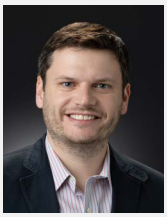

William R. Gray-Roncal, Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

William R. Gray-Roncal is an electrical engineer, a project manager, and the supervisor of the Neuro-AI and Connectomics Section in APL's Research and Exploratory Development Department. He has a BS in electrical engineering from Vanderbilt University, an MS in electrical engineering from the University of Southern California, and a PhD in computer science from Johns Hopkins University. William has 17 years of professional experience in diverse domains, including neuroscience, space, and submarines. His current projects focus on brain mapping and applied neuroscience toward next-generation artificial intelligence and health care applications. He has experience in computer vision, signal and image processing, machine learning, and systems engineering. His email address is william.gray.roncal@jhuapl.edu.