

NIH BRAIN NeuroAI Workshop 2024

Program Book

This two-day hybrid workshop on November 12 and 13, 2024, will explore how BRAIN's data, tools, and technologies can accelerate scientific discovery and transformative advances at the intersection of neuroscience and AI. The program book includes the agenda, pre-workshop position paper, speaker bios, abstracts, and the BRAIN NeuroAI Early-Career Scholar poster abstracts. The position paper provides the scientific background for the workshop, outlines the opportunities for BRAIN and NeuroAI, and motivates the guiding questions for each of the scientific panel discussions.

NIH BRAIN NeuroAI Workshop 2024 — Program Book

Table of Contents

| | |
|---|-----------|
| NIH BRAIN NeuroAI Workshop 2024 — Program Book | 1 |
| BRAIN NeuroAI Workshop Organizers | 2 |
| NIH BRAIN NeuroAI Workshop Planning Committee | 2 |
| Scientific Planning Committee | 2 |
| BRAIN NeuroAI Workshop Sponsors | 3 |
| Federal, Institutional, and Private Sponsors | 3 |
| BRAIN NeuroAI Workshop Agenda | 4 |
| Day 1 Schedule | 4 |
| Day 2 Schedule | 5 |
| Pre-Workshop Overview and Participant Guidance | 6 |
| Shaping the future of BRAIN at the convergence of neuroscience and AI | 6 |
| Defining the scope of NeuroAI for consideration at the workshop | 6 |
| Overview of BRAIN NeuroAI Workshop sessions and goals | 6 |
| The next step in the mission to understand the brain | 7 |
| Guiding Questions for the Scientific Panel Discussions | 7 |
| Session 1 Guiding Questions | 7 |
| Session 2 Guiding Questions | 8 |
| Session 3 Guiding Questions | 8 |
| Session 4 Guiding Questions | 8 |
| References | 9 |
| Speaker Information & Biographies — Organized By Session | 12 |
| Overview and Keynote Speakers | 12 |
| Session 1 Speakers | 13 |
| Session 2 Speakers | 16 |
| Session 3 Speakers | 19 |
| Session 4 Speakers | 22 |
| Funders Panel — Moderator & Panelists | 25 |
| Discussants for Scientific Panels | 29 |
| Speaker Abstracts — Organized by Session Order | 31 |
| Session 1 Presentation Abstracts | 31 |
| Session 2 Presentation Abstracts | 34 |
| Session 3 Presentation Abstracts | 37 |
| Session 4 Presentation Abstracts | 40 |
| BRAIN NeuroAI Early-Career Scholar Poster Abstracts | 43 |

BRAIN NeuroAI Workshop Organizers

NIH BRAIN NeuroAI Workshop Planning Committee

- **Joseph Monaco**, NINDS (Co-Organizer)
- **Grace Hwang**, NINDS (Co-Organizer)
- **Jessica Mollick**, NIDA (Co-Chair, Training Subcommittee)
- **Courtney Pinard**, NIMH (Co-Chair, Training Subcommittee)
- **Nina Hsu**, NINDS (Co-Chair, Neuroethics Subcommittee)
- **Jay Churchill**, NIMH (Co-Chair, Neuroethics Subcommittee)
- **Elizabeth Powell**, NIAAA
- **Merav Sabri**, NIDCD
- **Mohd Anwar**, NIBIB
- **Susan Wright**, NIDA
- **Christina Hatch**, NIDA
- **Roger Miller**, NIDCD
- **Karen David**, NINDS
- **Bo-Shiun Chen**, NINDS
- **Pantea Moghimi**, NINDS
- **Sandra Molina**, NINDS
- **Leslie Osborne**, NINDS
- **Sudha Srinivasan**, NINDS
- **Clayton Bingham**, NLM
- **Michele Ferrante**, NIMH
- **Fernando Fernandez**, NIMH
- **Mauricio Rangel-Gomez**, NIMH
- **Sandeep Kishore**, NIMH
- **Dana Schloesser**, OBSSR
- **Chris Kinsinger**, OD

Scientific Planning Committee

- **Anthony Zador**, Cold Spring Harbor Laboratory
- **Doris Tsao**, University of California, Berkeley
- **Gina Adam**, George Washington University
- **Blake Richards**, Mila – Quebec Artificial Intelligence Institute
- **J. Brad Aimone**, Sandia National Laboratories

BRAIN NeuroAI Workshop Sponsors

Federal, Institutional, and Private Sponsors

- **Defense Advanced Research Projects Agency (DARPA)** — Christine Edwards
- **The Kavli Foundation** — Amy Bernard, Stephanie Albins
- **Johns Hopkins Kavli Neuroscience Discovery Institute** — Sarada Viswanathan
- **Yale Wu Tsai Institute** — Giovanna Guerrero-Medina

Thank you to the BRAIN NeuroAI Workshop sponsors for your support!

BRAIN NeuroAI Workshop

Day 1 Agenda

Tuesday, November 12, 2024

- 8:00 AM** **Registration & Badge Pick-Up** — Natcher Auditorium Lobby
Coffee and refreshments
- 9:00 AM** **Opening Remarks**
Andrea Beckel-Mitchener, NIH BRAIN Deputy Director
- 9:30 AM** **Introduction to the BRAIN NeuroAI Workshop**
Introduction & Workshop Overview — Joseph Monaco (NIH/NINDS)
Opening Keynote — Tony Zador (Cold Spring Harbor Laboratory)
- 10:00 AM** Break (15 mins)
- 10:15 AM** **Session 1: Defining NeuroAI for BRAIN: Gaps, Challenges, and Opportunities**
Invited Short Presentations & Panel Discussion
Chair: Tony Zador (CSHL)
NIH Co-Chairs: Joseph Monaco (NIH/NINDS), Susan Wright (NIH/NIDA)
- Noon** **Lunch**
Organizer Group Photo — Natcher Auditorium Stairs
- 1:15 PM** **Funder's Panel — DoE, DoD, NIH, NSF, Simons Foundation**
Moderated by Terrence Sejnowski (The Salk Institute)
- 2:15 PM** **Session 2: Exploring the Structural and Functional Convergence of Deep Neural Nets and Brains**
Invited Short Presentations & Panel Discussion
Chairs: Blake Richards (Mila), Doris Tsao (UC Berkeley)
NIH Co-Chairs: Jessica Mollick (NIH/NIDA), Clayton Bingham (NIH/NLM)
- 4:00 PM** Coffee Break (30 mins)
- 4:30 PM** **Moderated Discussion**
Day 1 Wrap-up — Moderated by Terrence Sejnowski
- 6:00 PM** Adjourn

BRAIN NeuroAI Workshop

Day 2 Agenda

Wednesday, November 13, 2024

- 8:00 AM** **Registration & Badge Pick-Up** — Natcher Auditorium Lobby
Poster Session Setup — Natcher Atrium
Coffee and refreshments
- 8:30 AM** **Day 1 Recap and Overview of Day 2 Sessions**
Joseph Monaco (NIH/NINDS)
- 8:45 AM** **BRAIN NeuroAI Early-Career Scholar Poster Blitz**
NIH Co-Chairs: Courtney Pinard (NIH/NIMH), Jessica Mollick (NIH/NIDA)
- 9:05 AM** **BRAIN NeuroAI Early-Career Scholar Poster Session** — Natcher Atrium
- 10:15 AM** **Session 3: Advancing Theory for BRAIN through Neuromorphic Computing, Embodiment, and Physical Intelligence**
Invited Short Presentations & Panel Discussion
Chair: J. Brad Aimone (Sandia National Labs)
NIH Co-Chairs: Joseph Monaco (NIH/NINDS), Leslie Osborne (NIH/NINDS)
- Noon** **Lunch**
Extended Poster Session — Natcher Atrium (Coffee and refreshments)
- 1:30 PM** **Session 4: Towards Reciprocal BRAIN NeuroAI Advances in Intelligent Computing, Robotics, and Neurotechnologies**
Invited Short Presentations & Panel Discussion
Chairs: Gina Adam (George Washington University), J. Brad Aimone (Sandia National Labs)
NIH Co-Chairs: Grace Hwang (NIH/NINDS), Roger Miller (NIH/NIDCD)
- 3:30 PM** Coffee Break (30 mins)
- 4:00 PM** **Moderated Discussions**
Day 2 Wrap-up — Moderated by Gina Adam
Workshop Synthesis & Next Steps — Moderated by Paul Middlebrooks
- 5:30 PM** **BRAIN NeuroAI Early-Career Scholar Poster Awards**
Closing Remarks — John Ngai, NIH BRAIN Director
- 6:00 PM** Adjourn

Pre-Workshop Overview and Participant Guidance

Shaping the future of BRAIN at the convergence of neuroscience and AI

As the potential benefits and limitations of artificial intelligence (AI) become clear, the mission to understand the brain and accelerate cures is converging with interdisciplinary efforts to discover fundamental principles of intelligence in brains and AI. The NIH *Brain Research Through Advancing Innovative Neurotechnologies*[®] (BRAIN) Initiative is poised to leverage its wealth of data and tools to advance new theories and catalyze emerging NeuroAI research directions at the intersection of neuroscience and AI. The BRAIN NeuroAI Workshop will bring together neuroscientists, physical scientists, and engineers; theorists, data scientists, and mathematicians; and clinicians, technologists, funders, and other stakeholders. Workshop participants at all career stages will identify prospects for novel NeuroAI research and reveal promising approaches and opportunities in this exciting field.

Defining the scope of NeuroAI for consideration at the workshop

NeuroAI is an emerging area of research at the intersection of neuroscience and AI. Current state-of-the-art approaches in NeuroAI research include, but are not limited to, the development of AI tools based on artificial neural networks (ANNs) for interpreting large-scale neural and behavioral datasets [1–3] and mapping models to brain function by using metrics to find representations predicted by ANN models [4–6].

For the purposes of the BRAIN NeuroAI Workshop, the BRAIN Initiative views NeuroAI expansively as an interdisciplinary field that leverages the convergence between neuroscience and AI to drive reciprocal advances in both domains. Scientific knowledge of the brain basis of intelligence has been increasing. Advances in neuroscience have been enabled over the previous decade by the BRAIN Initiative [7] and large BRAIN-supported datasets and knowledgebases could be leveraged to advance our understanding of multiscale neural representations and algorithms [8–11].

Given the increasing importance of naturalistic behavior to understanding the brain, participants are encouraged to discuss and evaluate approaches to understanding the role of embodiment and physical interaction [12–14] in supporting the cognitive capabilities of natural intelligence in humans and other animals [15–17]. This expanded view of the breadth of NeuroAI research may extend to questions of how living organisms—and artificial agents such as bio-inspired robots [18] and living neural networks [19]—learn continually, efficiently, and adaptively across the lifespan [20–22] or evolution [23]. Incorporating the role of embodiment and physical interaction may enable new transformative NeuroAI theories that connect brain data to real-world intelligent behavior across species.

Overview of BRAIN NeuroAI Workshop sessions and goals

The BRAIN Initiative seeks broad scientific and technological perspectives about the dual, reciprocal aims of deepening brain understanding and innovating intelligent computing models and technologies such as bio-inspired robotics and energy-efficient edge devices with potential applications in science and health.

On Day 1 (November 12), the opening keynote by Planning Committee member Anthony Zador and the Session 1 kick-off presentation by Ali Minai will set the stage for workshop discussions and provide working definitions for, respectively, “NeuroAI” [24] and “natural intelligence” [25]. The Session 1 and Session 2 scientific panels will explore how data, infrastructure, and computational tools enable advances in NeuroAI frameworks, theories, models, and metrics [26–28]. When comparing natural intelligence and brain data with AI models, panelists and participants are encouraged to discuss how to evaluate the success of NeuroAI approaches (e.g., in the language centers [3] or the visual system [29]) in disentangling the cognitive, perceptual, planning, and reasoning faculties generally considered to constitute natural intelligence [30–32]. Participants are encouraged to avoid entering into semantic, terminological, or philosophical debates (e.g., [33–36]) that, while potentially important for the field, may distract from core themes and goals of the workshop.

On Day 2 (November 13), the Session 3 and Session 4 scientific panels will consider how to potentially extend and translate NeuroAI research opportunities to incorporate approaches from complementary disciplines, such as embodied cognition, neuromorphic computing, and bio-inspired robotics. For example, large-scale neuromorphic computing systems [37–39] may hold promise for scalable neural simulation for testing new theories [40] and designing energy-efficient neuromorphic devices may enable future neural interfaces that learn and adapt [41–43]. Session 3 will focus on advancing theory-driven modeling and closed-loop neuroscience. Session 4 will look forward to reciprocal advances in NeuroAI-enabled technologies including robotics and health applications.

Expanding the scope of NeuroAI to include physical aspects of embodiment may simultaneously advance the science of natural intelligence, resilience, adaptability, and energy-efficiency in the brains of humans and other animals. Energy-efficiency is considered a key metric for evaluating both scientific and technological progress. Workshop participants should explore and discuss energy-efficiency and other relevant metrics or benchmarks.

Across both days of the workshop, the NIH BRAIN Initiative will hear from workshop panelists, discussants, and participants representing diverse NeuroAI-related research communities about promising NeuroAI opportunities. Throughout the workshop, participants are expected and encouraged to consider neuroethical implications, as appropriate, when evaluating or comparing potential approaches and priorities, particularly those involving human subjects research, data from human participants, or clinical populations [44–46]. The integration of neuroethics in BRAIN's mission to understand the brain is critical to translating innovative science into future health impact [7].

The next step in the mission to understand the brain

The scientific understanding of reasoning and planning in the brain has evolved through philosophical, theoretical, experimental, and computational phases [47]. With the advent of modern AI technology—enabled by massive pre-training datasets and large-scale parallel computing on high-throughput digital hardware—there is an opportunity to pursue complementary paths to transform the scientific understanding of intelligence through interdisciplinary approaches that bring together neuroscience and AI [48–50] with tools from fields such as cognitive science [51], network science [52], control theory [53–55], neuromorphic computing [56–59], and bio-inspired robotics [60–63]. Workshop participants are encouraged to identify and evaluate alternative, complementary, and interdisciplinary approaches.

Workshop discussions should identify current and future challenges, and explore how reciprocal NeuroAI approaches and priorities may be enabled or catalyzed by large-scale brain data such as cell atlases and connectomes, metrics and benchmarks for interpreting NeuroAI models, computational modeling tools and simulation infrastructure, conceptual and theoretical frameworks, developmental and evolutionary perspectives, and physical or in silico platforms and hardware devices. Participants are encouraged to consider cultural or institutional obstacles to interdisciplinary training and collaboration that may be faced by the next generation of NeuroAI scientists and engineers [64].

Disentangling fundamental principles of intelligence through BRAIN and NeuroAI could transform neuroscience and brain health. To [quote](#) Planning Committee member J. Brad Aimone, a neuromorphic computing researcher from Sandia National Labs: “We are entering a tremendously exciting era, and not just because of AI and neural networks. We have reached a point with the BRAIN Initiative and neurotechnologies that we can see the brain in deeper ways than we ever thought possible. We have data to constrain rich models of neural processes that we can map to diseases, and we have a growing set of interventions that could potentially revolutionize mental healthcare if we only had the data and strategies to personalize it.” By the end of the workshop, challenges and opportunities should be identified that BRAIN Initiative might consider to advance the emerging field of NeuroAI.

Guiding Questions for the Scientific Panel Discussions

Across a series of pre-workshop coordination meetings, the panelists (speakers and discussants) for each session met with workshop organizers to develop and refine overarching questions to guide each session's discussion. Those guiding questions are presented below.

Session 1 Guiding Questions

1. How can BRAIN's large-scale datasets be structured and leveraged to develop NeuroAI resources, such as neural foundation models or digital twins, that bridge multiple scales, while balancing the need for both hypothesis-driven science and high-entropy naturalistic data collection?
2. How should we expand our understanding of neural computation to incorporate broader biological systems (including glia, neuromodulation, and developmental/evolutionary perspectives) in ways that inform both theoretical advances and practical NeuroAI implementations?
3. How can experimental platforms and technologies support a "discovery loop" that integrates theory development, model validation, hypothesis generation, and data-driven approaches while meaningfully incorporating physical embodiment and real-world behavior?

4. What infrastructure, tools, and coordination mechanisms are needed to enable collection and analysis of naturalistic neural and behavioral data at scales beyond individual laboratories while maintaining scientific rigor and reproducibility?

Session 2 Guiding Questions

1. How can we develop and validate metrics for comparing biological and artificial systems that capture meaningful computational principles while avoiding overfitting to specific comparison methods or oversimplifying complex neural dynamics?
2. What frameworks and approaches can help identify meaningful comparisons between biological and artificial systems, considering different levels of abstraction from algorithmic principles to physical implementation? Which computational principles in brain systems are more or less difficult for NeuroAI models to capture?
3. How can BRAIN Initiative datasets be effectively leveraged to evaluate and validate theories about shared computational principles, while accounting for the different requirements of hypothesis-driven science and large-scale projects to develop NeuroAI models and resources?
4. What infrastructure, benchmarks, and standardized platforms will be needed to enable ethically and scientifically rigorous measurements comparing human or animal data to NeuroAI models across laboratories?

Session 3 Guiding Questions

1. How can neuromorphic approaches help us understand fundamental principles of brain computation while also advancing more efficient artificial systems? What determines whether neuromorphic computing serves primarily as a modeling and simulation platform versus providing emulation of biological processes to achieve deeper theoretical insights into neural computation?
2. Given the co-evolution of hardware and algorithms in both technology and biology, how do we ensure our choice of abstraction level and implementation approach reveals fundamental principles rather than artifacts of available technology? What biological mechanisms are essential to implement versus those that can be simplified?
3. How can different approaches to physical implementation—from neuromorphic hardware to bio-inspired robotics and physical or in silico model-in-the-loop systems—advance our theoretical understanding of neural computation? What role should embodied approaches or physical interaction play in the NeuroAI discovery loop?

Session 4 Guiding Questions

1. What foundational advances in NeuroAI are needed to enable energy-efficient neuromorphic computing, adaptive robotics, or intelligent neural interfaces with the potential to transform neuroscience and brain health? What kinds of metrics are needed to evaluate progress in both technical capabilities and translational impact?
2. What infrastructure and platforms are needed to enable innovative, scalable healthcare technologies that are both affordable and secure? How can BRAIN Initiative resources and cross-agency partnerships accelerate translation of NeuroAI advances to clinical applications?
3. How can NeuroAI approaches and technologies drive reciprocal advances between fundamental neuroscience and transformative health technologies while ensuring ethical use of neural data and meaningful incorporation of clinician/patient perspectives? How do we balance innovation with safety and accessibility?

References

1. Beam E, Potts C, Poldrack RA, Etkin A. A data-driven framework for mapping domains of human neurobiology. *Nat Neurosci.* 2021;24: 1733–1744.
2. Bzdok D, Thieme A, Levkovskyy O, Wren P, Ray T, Reddy S. Data science opportunities of large language models for neuroscience and biomedicine. *Neuron.* 2024;112: 698–717.
3. Tuckute G, Kanwisher N, Fedorenko E. Language in Brains, Minds, and Machines. *Annu Rev Neurosci.* 2024. doi:10.1146/annurev-neuro-120623-101142
4. Zhuang C, Yan S, Nayebi A, Schrimpf M, Frank MC, DiCarlo JJ, et al. Unsupervised neural network models of the ventral visual stream. *Proc Natl Acad Sci U S A.* 2021;118: e2014196118.
5. Margalit E, Lee H, Finzi D, DiCarlo JJ, Grill-Spector K, Yamins DLK. A unifying framework for functional organization in early and higher ventral visual cortex. *Neuron.* 2024. doi:10.1016/j.neuron.2024.04.018
6. Levy ERJ, Carrillo-Segura S, Park EH, Redman WT, Hurtado JR, Chung S, et al. A manifold neural population code for space in hippocampal coactivity dynamics independent of place fields. *Cell Rep.* 2023;42: 113142.
7. Ngai J. BRAIN @ 10: A decade of innovation. *Neuron.* 2024;112: 3003–3006.
8. Tosches MA, Lee HJ. Cellular atlases of the entire mouse brain. In: Nature Publishing Group UK [Internet]. 13 Dec 2023 [cited 13 Dec 2023]. doi:10.1038/d41586-023-03781-1
9. Langlieb J, Sachdev NS, Balderrama KS, Nadaf NM, Raj M, Murray E, et al. The molecular cytoarchitecture of the adult mouse brain. *Nature.* 2023;624: 333–342.
10. Shapson-Coe A, Januszewski M, Berger DR, Pope A, Wu Y, Blakely T, et al. A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science.* 2024;384: eadk4858.
11. Lappalainen JK, Tschopp FD, Prakhya S, McGill M, Nern A, Shinomiya K, et al. Connectome-constrained networks predict neural activity across the fly visual system. *Nature.* 2024; 1–9.
12. Cisek P, Kalaska JF. Neural Mechanisms for Interacting with a World Full of Action Choices. *Annu Rev Neurosci.* 2020;33: 269–298.
13. Jaeger H, Noheda B, van der Wiel WG. Toward a formal theory for computing machines made out of whatever physics offers. *Nat Commun.* 2023;14: 4911.
14. Jaeger J, Riedl A, Djedovic A, Vervaeke J, Walsh D. Naturalizing relevance realization: why agency and cognition are fundamentally not computational. *Front Psychol.* 2024;15. doi:10.3389/fpsyg.2024.1362658
15. Krakauer JW, Ghazanfar AA, Gomez-Marin A, Maclver MA, Poeppel D. Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron.* 2017;93: 480–490.
16. Cisek P. Resynthesizing behavior through phylogenetic refinement. *Atten Percept Psychophys.* 2019;81: 2265–2287.
17. Gomez-Marin A, Ghazanfar AA. The Life of Behavior. *Neuron.* 2021;104: 25–36.
18. Bartolozzi C, Indiveri G, Donati E. Embodied neuromorphic intelligence. *Nat Commun.* 2022;13: 1024.
19. Rouleau N, Murugan NJ, Kaplan DL. Functional bioengineered models of the central nervous system. *Nat Rev Bioeng.* 2023;1: 252–270.
20. Soures N, Dey J, Kudithipudi D. Learning continually in silicon. *Computer (Long Beach Calif).* 2024;57: 160–164.
21. Kudithipudi D, Aguilar-Simon M, Babb J, Bazhenov M, Blackiston D, Bongard J, et al. Biological Underpinnings for Lifelong Learning Machines. *Nature Machine.* 2022. Available: <https://www.nature.com/articles/s42256-022-00452-0>
22. Petrenko S, Brna A, Aguilar-Simon M, Wunsch D. Lifelong context recognition via online deep feature clustering. 2023. doi:10.36227/techrxiv.23653737.v1

23. Zador AM. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nat Commun.* 2019;10: 3770.
24. Zador A, Escola S, Richards B, Ölveczky B, Bengio Y, Boahen K, et al. Catalyzing next-generation Artificial Intelligence through NeuroAI. *Nat Commun.* 2023;14: 1597.
25. Minai AA. Deep Intelligence: What AI Should Learn from Nature's Imagination. *Cogn Comput.* 2023; 1–16.
26. Brunton BW, Beyeler M. Data-driven models in human neuroscience and neuroengineering. *Curr Opin Neurobiol.* 2019;58: 21–29.
27. Mathis MW, Perez Rotondo A, Chang EF, Tolias AS, Mathis A. Decoding the brain: From neural representations to mechanistic models. *Cell.* 2024;187: 5814–5832.
28. Li Q, Sorscher B, Sompolinsky H. Representations and generalization in artificial and brain neural networks. *Proc Natl Acad Sci U S A.* 2024;121: e2311805121.
29. Dyballa L, Rudzite AM, Hoseini MS, Thapa M, Stryker MP, Field GD, et al. Population encoding of stimulus features along the visual hierarchy. *Proc Natl Acad Sci U S A.* 2024;121: e2317773121.
30. Schrimpf M, Kubilius J, Lee MJ, Ratan Murty NA, Ajemian R, DiCarlo JJ. Integrative Benchmarking to Advance Neurally Mechanistic Models of Human Intelligence. *Neuron.* 2020;108: 413–423.
31. Richards B, Tsao D, Zador A. The application of artificial intelligence to biology and neuroscience. *Cell.* 2022;185: 2640–2643.
32. Doerig A, Sommers RP, Seeliger K, Richards B, Ismael J, Lindsay GW, et al. The neuroconnectionist research programme. *Nat Rev Neurosci.* 2023;24: 431–450.
33. Brette R. Is coding a relevant metaphor for the brain? *Behav Brain Sci.* 2018; 1–44.
34. Richards BA, Lillicrap TP. The Brain-Computer Metaphor Debate Is Useless: A Matter of Semantics. *Frontiers in Computer Science.* 2022;4. doi:10.3389/fcomp.2022.810358
35. Gershman SJ. What have we learned about artificial intelligence from studying the brain? *Biol Cybern.* 2024;118: 1–5.
36. Pillow JW. Cross Talk opposing view: Marr's three levels of analysis are not useful as a framework for neuroscience. *J Physiol.* 2024;602: 1915–1917.
37. Davies M, Srinivasa N, Lin T-H, China G, Cao Y, Choday SH, et al. Loihi: A Neuromorphic Manycore Processor with On-Chip Learning. *IEEE Micro.* 2018;38: 82–99.
38. Davies M, Wild A, Orchard G, Sandamirskaya Y, Guerra GAF, Joshi P, et al. Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proc IEEE Inst Electr Electron Eng.* 2021;109: 911–934.
39. Pehle C, Billaudelle S, Cramer B, Kaiser J, Schreiber K, Stradmann Y, et al. The BrainScaleS-2 accelerated neuromorphic system with hybrid plasticity. *Front Neurosci.* 2022;16. doi:10.3389/fnins.2022.795876
40. Aimone JB, Parekh O. The brain's unique take on algorithms. *Nat Commun.* 2023;14. doi:10.1038/s41467-023-40535-z
41. Qiao N, Mostafa H, Corradi F, Osswald M, Stefanini F, Sumislawska D, et al. A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128K synapses. *Front Neurosci.* 2015;9. doi:10.3389/fnins.2015.00141
42. Schuman CD, Kulkarni SR, Parsa M, Mitchell JP, Date P, Kay B. Opportunities for neuromorphic computing algorithms and applications. *Nat Comput Sci.* 2022;2: 10–19.
43. Aboumerhi K, Güemes A, Liu H, Tenore F, Etienne-Cummings R. Neuromorphic applications in medicine. *J Neural Eng.* 2023;20. doi:10.1088/1741-2552/aceca3
44. Yuste R, Goering S, Arcas BAY, Bi G, Carmena JM, Carter A, et al. Four ethical priorities for neurotechnologies and AI. *Nature.* 2017;551: 159–163.

45. Hendriks S, Ramos KM, Grady C. Survey of Investigators About Sharing Human Research Data in the Neurosciences. *Neurology*. 2022;99: e1314–e1325.
46. Monosov IE, Zimmermann J, Frank MJ, Mathis MW, Baker JT. Ethological computational psychiatry: Challenges and opportunities. *Curr Opin Neurobiol*. 2024;86: 102881.
47. Monaco JD, Hwang GM. Neurodynamical Computing at the Information Boundaries of Intelligent Systems. *Cognit Comput*. 2022. doi:10.1007/s12559-022-10081-9
48. Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ. Building machines that learn and think like people. *Behav Brain Sci*. 2017;40: e253.
49. Richards BA, Lillicrap TP, Beaudoin P, Bengio Y, Bogacz R, Christensen A, et al. A deep learning framework for neuroscience. *Nat Neurosci*. 2019;22: 1761–1770.
50. Sinz FH, Pitkow X, Reimer J, Bethge M, Tolias AS. Engineering a Less Artificial Intelligence. *Neuron*. 2019;103: 967–979.
51. Achler T. What AI, neuroscience, and cognitive science can learn from each other: An embedded perspective. *Cognit Comput*. 2023; 1–9.
52. Barabási DL, Bianconi G, Bullmore E, Burgess M, Chung S, Eliassi-Rad T, et al. Neuroscience needs Network Science. *arXiv [q-bio.NC]*. 2023. Available: <http://arxiv.org/abs/2305.06160>
53. Powers WT. Feedback: Beyond Behaviorism. *Science*. 1973;179: 351–356.
54. Sepulchre R, Drion G, Franci A. Control Across Scales by Positive and Negative Feedback. *Annual Review of Control, Robotics, and Autonomous Systems*. 2019;2: 89–113.
55. Li JS, Sarma AA, Sejnowski TJ, Doyle JC. Internal feedback in the cortical perception-action loop enables fast and accurate behavior. *Proc Natl Acad Sci U S A*. 2023;120: e2300445120.
56. Hasler J, Marr B. Finding a roadmap to achieve large neuromorphic hardware systems. *Front Neurosci*. 2013;7: 118.
57. Potok T, Schuman C, Patton R, Hylton T, Li H, Pino R. Neuromorphic computing, architectures, models, and applications. A beyond-CMOS approach to future computing, June 29-July 1, 2016, Oak Ridge, TN. USDOE Office of Science (SC) (United States). *Advanced Scientific Computing Research (ASCR)*; 2016 Dec. doi:10.2172/1341738
58. Aimone JB. A Roadmap for Reaching the Potential of Brain-Derived Computing. *Advanced Intelligent Systems*. 2021; 2000191.
59. Boahen K. Dendrocentric learning for synthetic intelligence. *Nature*. 2022;612: 43–50.
60. Sitti M. Physical intelligence as a new paradigm. *Extreme Mech Lett*. 2021;46: 101340.
61. Bogdan PA, Marcinnò B, Casellato C, Casali S, Rowley AGD, Hopkins M, et al. Towards a bio-inspired real-time neuromorphic cerebellum. *Front Cell Neurosci*. 2021;15: 622870.
62. Barter JW, Yin HH. Achieving natural behavior in a robot using neurally inspired hierarchical perceptual control. *iScience*. 2021;24: 102948.
63. Aubin CA, Gorissen B, Milana E, Buskohl PR, Lazarus N, Slipher GA, et al. Towards enduring autonomous robots via embodied energy. *Nature*. 2022;602: 393–402.
64. Luppi AI, Achterberg J, Schmidgall S, Bilgin IP, Herholz P, Sprang M, et al. Trainees' perspectives and recommendations for catalyzing the next generation of NeuroAI researchers. *Nat Commun*. 2024;15: 9152.

Speaker Information & Biographies — Organized By Session

Overview and Keynote Speakers

Opening Speakers



Anthony Zador, M.D., Ph.D.
The Alle Davis and Maxine Harrison
Professor of Neurosciences
Cold Spring Harbor Laboratory
zador.AT.cshl.edu



Joseph D. Monaco, Ph.D.
Scientific Program Manager
NIH BRAIN Initiative
joseph.monaco.AT.nih.gov

Biography

Dr. Zador received his MD and Ph.D. from Yale in 1994, where his focus was theoretical neuroscience and neural networks. He then did postdoctoral research in synaptic physiology at the Salk Institute in La Jolla, California. In 1999, he joined the faculty at Cold Spring Harbor Laboratory in New York, where he is now the Alle Davis Harris Professor of Biology and served as Chair of Neuroscience from 2008-2018. The goal of his research is to understand the computational principles whereby neural wiring enables complex behavior. His laboratory pioneered the use of rodents in complex sensory decision tasks, and also developed a revolutionary approach to determining brain wiring using high-throughput DNA sequencing. His current research interests include applying neuroscience to usher in the next generation of Artificial Intelligence.

Dr. Monaco is a scientific program manager in the Office of the BRAIN Director, where he coordinates the BRAIN Initiative's data sharing policy and provides guidance to BRAIN programs and transformative projects relating to theory, modeling tools, data integration, and artificial intelligence. He conducted theoretical and computational neuroscience research for over 20 years with a focus on the role of hippocampal circuits in spatial navigation and episodic memory. As a Postdoctoral Fellow and Research Associate at Johns Hopkins University, Dr. Monaco developed neurobehavioral analysis methods to link individual movements to memory formation in rats, built models of how brains compute with neural oscillations, and helped lead interdisciplinary collaborations to advance theoretical models for brain-inspired robotic control. He is a co-organizer of the NIH BRAIN NeuroAI Workshop.

Session 1 Speakers

Biography



Anton Arkhipov, Ph.D.
Investigator
Allen Institute
antona.AT.alleninstitute.org

Dr. Arkhipov joined the Allen Institute in 2013 as an Assistant Investigator in the Modeling, Analysis, and Theory group. He is leading efforts to carry out biophysically detailed simulations of individual neurons as well as large-scale neuronal circuits from the mouse visual system. The main focus of his research is on integration of experimental anatomical and physiological data to build sophisticated, highly realistic computational models of cortical circuitry, with the aim of elucidating mechanisms underlying processing of visual information in the cortex. Before joining the Allen Institute he was a Postdoctoral Fellow at D. E. Shaw Research in New York City, where he used a specialized supercomputing architecture to perform computational studies of structure-function relationships in proteins, with the emphasis on cancer-associated cell-surface receptors. Arkhipov received his B.S. and M.S. in Physics from Moscow Institute of Physics and Technology and a Ph.D. in Physics from the University of Illinois at Urbana-Champaign.



Bing W. Brunton, Ph.D.
Professor
University of Washington
bbrunton.AT.uw.edu

Dr. Brunton is currently a Professor of Biology and the Richard & Joan Komen University Chair at the University of Washington (UW) in Seattle, with affiliations at the eScience Institute for Data Science, the Paul G. Allen School of Computer Science & Engineering, and the Department of Applied Mathematics. She studied at Caltech (2006, B.S. in Biology, focus on biophysics) and then Princeton (2012, Ph.D. in Neuroscience). She is a computational neuroscientist with broad interests at the intersection of systems neuroscience, animal behavior, and artificial intelligence. Her research group focuses on developing data-intensive methods to understand and model neural function and behavior, using approaches from machine learning, deep reinforcement learning, computer vision, and physics-constrained simulations. She is drawn to understand how the nervous system solves challenges that are vital to the animal: sensing the environment, maneuvering in the physical world, planning and executing goals, and interacting with their societies.



Dominique Duncan, Ph.D.

Assistant Professor of Neurology and
Biomedical Engineering
University of South California
duncand.AT.usc.edu

Dr. Duncan is an assistant professor of Neurology, Neuroscience, and Biomedical Engineering at the USC Stevens Neuroimaging and Informatics Institute in the Laboratory of Neuro Imaging (LONI) at the University of Southern California. Dr. Duncan's background spans mathematics, engineering, and neuroscience. She received her Ph.D. at Yale University in Electrical Engineering where she analyzed intracranial EEG data using nonlinear factor analysis to identify pre-seizure states of epilepsy patients. Dr. Duncan is funded through both the National Institutes of Health (NIH) and the National Science Foundation (NSF). She has built international, multidisciplinary collaborations and developed novel analytic tools to analyze multimodal data, including imaging and electrophysiology, particularly in the areas of traumatic brain injury, epilepsy, and COVID-19. By creating large-scale data repositories and linking them with analytic, visualization, and quality control tools for multimodal data, her work aims to encourage collaboration across multiple fields.



Wolfgang Losert, Ph.D.

Professor
University of Maryland
wlosert.AT.umd.edu

Dr. Losert is MPower Professor of Physics and Interim Associate Dean for Research in the College of Computer Mathematical, and Natural Sciences at the University of Maryland. His research team investigates the dynamics of living systems at the convergence of biophysics and AI. He co-lead the Technology and Data Science Cores of an NIH BRAIN initiative U19 aimed at optogenetic measurements and control of the collective character of neurons in sensory processing of the brain. Prof. Losert's current research focuses on the multimodal electrical, chemical, and mechanical excitability of cells and tissues, which enable new paradigms for information flow and processing in living neural networks. Dr. Losert is a fellow of the AAAS and the American Physical Society.



Ali A. Minai, Ph.D.

Professor
University of Cincinnati
ali.minai.AT.uc.edu

Dr. Minai is Professor of Electrical & Computer Engineering at the University of Cincinnati, with a faculty appointment in the Neuroscience Graduate Program. He holds a Ph.D. in Electrical Engineering and completed postdoctoral training in neuroscience at the University of Virginia. Dr. Minai's research spans artificial intelligence, neural networks, computational neuroscience, and complex systems. His current focus is on place field-based models of robot navigation, analysis of stereo-EEG data using neural networks, representational interpretation in deep neural networks, and applications of large language models in cognitive tasks. Recently, he has engaged actively with philosophical issues in AI



Patrick Mineault, Ph.D.

NeuroAI Researcher
Amaranth Foundation
patrick.mineault.AT.gmail.com



Andreas Tolias, Ph.D.

Professor
Stanford University
tolias.AT.stanford.edu

through writings, discussions, and talks. Dr. Minai has been a member of the Board of Governors of the International Neural Network Society for several years, serving as President of the Society in 2015–16. He is currently an action editor for Neural Networks, and has served on the editorial boards of several other journals.

Dr. Mineault writes the NeuroAI archive. He is senior machine learning scientist working at the intersection of neuroscience and AI, with an adjunct appointment at the Math and Stats department at Université de Montréal as *chercheur invité*. He received his B.Sc. in Math and Physics and a Ph.D. in the computational neuroscience of vision at McGill, followed by a postdoc at UCLA. He was a software engineer at Google in Mountain View, CA and a research scientist in brain-computer interfaces at Meta. He was also the founding CTO of Neuromatch Academy and founded a NeuroAI startup called Blindsight Therapeutics. His research bridges neuroscience and AI, in particular modelling the dorsal stream of the visual cortex and building neural foundation models.

Dr. Tolias is a Professor at Stanford University, with affiliations in Bio-X, Wu Tsai Neurosciences Institute, Electrical Engineering, and the Institute for Human-Centered Artificial Intelligence. He holds degrees from the University of Cambridge (B.A., M.A.) and MIT (Ph.D.), with postdoctoral training at the Max-Planck Institute. Previously, he was Brown Endowed Professor of Neuroscience at Baylor College of Medicine and founding director of the Center for Neuroscience and Artificial Intelligence. Tolias has received numerous awards, including the NIH Director's Pioneer Award and McKnight Foundation Scholar Award. His research integrates large-scale neurophysiology and behavioral neuroscience with deep learning to understand visual intelligence mechanisms. He has led international DARPA and IARPA-funded teams, notably completing the IARPA MICrONS project, which generated a multi-petabyte dataset of co-registered neurophysiological and neuroanatomical brain data. Tolias developed the "inception loop" paradigm, combining neurophysiology with AI to decipher the neural code, leading to fundamental discoveries in visual cortex circuitry.

Session 2 Speakers

Biography



Carina Curto, Ph.D.
 Professor
 Brown University
carina_curto.AT.brown.edu

Dr. Curto received an A.B. in physics from Harvard in 2000 and a Ph.D. in mathematics from Duke in 2005. During my postdoctoral years at Rutgers and NYU, I transitioned to theoretical and computational neuroscience. She then held faculty positions in mathematics at UNL (2009–2014) and Penn State (2014–2024). Her current research focuses on the theory and analysis of neural networks and neural codes, motivated by questions of learning, memory, and sequence generation in cortical and hippocampal circuits. A big part of her research program involves developing novel applications of algebra, geometry, topology, dynamical systems, and combinatorics to neuroscience.



Evelina Fedorenko, Ph.D.
 Associate Professor
 Massachusetts Institute of Technology
evelina9.AT.mit.edu

Dr. Ev Fedorenko is a cognitive neuroscientist who studies the human language system. She received her bachelor's degree from Harvard in 2002, and her Ph.D. from MIT in 2007. She was then awarded a K99/R00 career development award from NIH. In 2014, she joined the faculty at MGH/HMS, and in 2019 she returned to MIT where she is currently an Associate Professor of Neuroscience in the BCS Department and the McGovern Institute for Brain Research. Dr. Fedorenko uses fMRI, intracranial recordings and stimulation, EEG, MEG, and computational modeling, to study adults and children, including those with developmental and acquired brain disorders, and otherwise atypical brains.



Panayiota Poirazi, Ph.D.
 Research Director
 IMBB-FORTH
poirazi.AT.imbb.forth.gr

Dr. Poirazi is a Director of Research at the Institute of Molecular Biology and Biotechnology, Foundation for Research and Technology-Hellas (FORTH) and head of the Dendrites lab (www.dendrites.gr). She received the B.S. in Mathematics from the University of Cyprus in 1996, M.S. and Ph.D. degrees in Biomedical Engineering in 1998 and 2000, respectively, from the University of Southern California. Her work focuses on understanding the role of dendrites in complex brain functions. She uses primarily computational modeling of neurons and networks, brain-inspired machine learning and recently in vivo experiments in mice. She has received several awards for academic excellence, including the EMBO Young Investigator award in 2005, two Marie Curie fellowships (2002 and 2008), an ERC Starting Grant in 2012, the Friedrich Wilhelm Bessel award of the Humboldt Foundation in 2018 and an EINSTEIN foundation visiting fellowship in 2019. She is a member of EMBO and currently serves as the Secretary General of FENS.



Blake Richards, Ph.D.

Professor

Mila – Quebec Artificial Intelligence Institute
blake.richards.AT.mila.quebec

Dr. Richards is an Associate Professor in the School of Computer Science and Montreal Neurological Institute at McGill University and a Core Faculty Member at MILA. Richards' research is at the intersection of neuroscience and AI. His laboratory investigates universal principles of intelligence that apply to both natural and artificial agents. He has received several awards for his work, including the NSERC Arthur B. McDonald Fellowship in 2022, the Canadian Association for Neuroscience Young Investigator Award in 2019, and a CIFAR Canada AI Chair in 2018. Richards was a Banting Postdoctoral Fellow at SickKids Hospital from 2011 to 2013. He obtained his Ph.D. in neuroscience from the University of Oxford in 2010 and his BSc in cognitive science and AI from the University of Toronto in 2004.



Karen S. Rommelfanger, Ph.D.

Director, Institute of Neuroethics Think and Do Tank
director.AT.instituteofneuroethics.org

Dr. Rommelfanger is a neurotech ethicist and strategist. She founded and directs the Institute of Neuroethics, the first think tank dedicated to neuroethics, working across sectors to promote trusted neuroscience for all. Pioneering neuroethics-by-design approaches, she launched the first neuroethics consultancy Ningen Neuroethics Co-Lab. Her early career as Ph.D.-trained neuroscientist, organically evolved into neuroethics research exploring how neuroscience challenges definitions of health across cultures and the ensuing societal implications of neurotechnology deployment. As a scholar, she maintains a professorship at Emory University in Neurology where she established a Neuroethics Program, has published extensively in neuroscience and neuroethics. She is a member of the NIH BRAIN Neuroethics Working Group and co-authored the BRAIN 2.0 Neuroethics Roadmap. A recognized global leader in neuroethics, she has collaborated with and advised policy, research, and diplomacy organizations such as the Council of Europe, DARPA, GESDA Science Diplomacy Anticipator, OECD, and World Economic Forum.

**Martin Schrimpf, Ph.D.**

Tenure-Track Assistant Professor
École Polytechnique Fédérale de Lausanne
martin.schrimpf.AT.epfl.ch

Dr. Schrimpf is a tenure-track assistant professor at EPFL where he builds artificial intelligence models of the brain. To achieve this goal, he bridges research in Machine Learning, Neuroscience, and Cognitive Science. He initiated the community-wide Brain-Score platform for evaluating models on their brain and behavioral alignment, and built state-of-the-art models such as CORnet and VOneNet. Martin completed his Ph.D. at MIT with Jim DiCarlo, following Bachelor's and Master's degrees in computer science at TUM, LMU, and UNA. Previously he worked at Harvard, MetaMind/Salesforce, Oracle, and co-founded two startups. His work has been published at top venues including PNAS, Neuron, Nature Human Behavior, NeurIPS, and ICLR. He has received numerous awards for his research, including the Neuro-Irv and Helga Cooper Open Science Prize, the McGovern and Takeda fellowships, and the Google.org Impact Challenge prize. Among others, Martin's work has been recognized in the news at Science magazine, MIT News, and Scientific American.

**Doris Tsao, Ph.D.**

Professor
University of California, Berkeley & HHMI
dortsao.AT.berkeley.edu

Dr. Tsao is a professor in the neurobiology division of the Department of Molecular & Cell Biology, and the Helen Wills Neuroscience Institute. She studies visual perception in primates in order to understand how the brain creates our sense of reality. She is widely recognized for her work on the neural system for face processing within the temporal lobe, clarifying its anatomical organization and coding principles. Most recently, her lab discovered that this system is part of a larger map of object space.

Session 3 Speakers

Biography



James "Brad" Aimone, Ph.D.
 Computational & Theoretical Neuroscientist
 Distinguished Member of Technical Staff
 Sandia National Laboratories
 jbaimon.AT.sandia.gov

Dr. Aimone is a Distinguished Member of Technical Staff in the Center for Computing Research at Sandia National Laboratories, where he is a lead researcher in leveraging computational neuroscience to advance artificial intelligence and in using neuromorphic computing platforms for future scientific computing applications. Brad currently leads a multi-institution DOE Office of Science Microelectronics Co-Design project titled *COINFLIPS* (which stands for CO-designed Influenced Neural Foundations Inspired by Physical Stochasticity) which is focused on developing a novel probabilistic neuromorphic computing platform. He also currently leads several other research efforts on designing neural algorithms for scientific computing applications and neuromorphic machine learning implementations.



Kwabena Boahen, Ph.D.
 Professor
 Stanford University
 boahen.AT.stanford.edu

Dr. Boahen is a Professor of Bioengineering and Electrical Engineering at Stanford University. His group models the nervous system computationally to elucidate principles of neural design at the cellular, circuit, and systems levels; and synthesizes neuromorphic electronic systems whose energy-use scales with their size as efficiently as the brain does. His research has resulted in over a hundred publications, including a cover story in *Scientific American* featuring his lab's work on a silicon retina and a silicon tectum that "wire together" automatically (May 2005). He has received several distinguished honors, including the National Institutes of Health Director's Pioneer Award (2006). He was elected a fellow of the American Institute for Medical and Biological Engineering (2016) and of the Institute of Electrical and Electronic Engineers (2016) in recognition of his lab's work on Neurogrid, an iPad-size platform that emulates a million neurons in the cerebral cortex in real time.



Frances Chance, Ph.D.
 Computational & Theoretical Neuroscientist
 Principal Member of the Technical Staff
 Sandia National Laboratories
 fschanc.AT.sandia.gov

Dr. Chance's research focuses on understanding how biological neural networks represent, transform, and transmit information in the brain. At Sandia Labs, she uses computational modeling and mathematical analysis of neurons and neural networks to understand the basic computations that underlie sensory processing and cognition.



SueYeon Chung, Ph.D.

Assistant Professor
New York University & Flatiron Institute
sueyeon.AT.nyu.edu

Prior to joining NYU, Dr. SueYeon Chung was a Postdoctoral Research Scientist in the Center for Theoretical Neuroscience at Columbia University, and a Fellow in Computation in the Department of Brain and Cognitive Sciences at MIT. Before that, She received a Ph.D. in applied physics at Harvard University. Before that, she studied physics and mathematics as an undergraduate at Cornell University.



Mitra Hartmann, Ph.D.

Professor
Northwestern University
hartmann.AT.northwestern.edu

Dr. Hartmann received a Bachelor of Science in Applied and Engineering Physics from Cornell University, a Ph.D. in Integrative Neuroscience from the California Institute of Technology and was a post-doctoral scholar at the Jet Propulsion Laboratory in the Bio-Inspired Technology and Systems group. She is currently a professor with a 50-50 joint appointment between the Departments of Biomedical Engineering and Mechanical Engineering at Northwestern University. She is the recipient of the Charles Deering McCormick Professor of Teaching Excellence award and an elected fellow of the American Institute for Medical and Biological Engineering (AIMBE).



Jennifer Hasler, Ph.D.

Regents Entrepreneur
Georgia Institute of Technology
jennifer.hasler.AT.ece.gatech.edu

Dr. Hasler is a Regents Professor in the School of Electrical and Computer Engineering at Georgia Institute of Technology. Dr. Hasler received her M.S. and B.S.E. in Electrical Engineering from Arizona State University in 1991, received her Ph.D. from California Institute of Technology in Computation and Neural Systems in 1997, and received her Master of Divinity from Emory University in 2020. Dr. Hasler received the NSF CAREER Award in 2001, and the ONR YIP award in 2002. Dr. Hasler has been involved in multiple startup companies, including GTronix, founded in 2002 and acquired by Texas Instruments in 2010. Dr. Hasler received the Paul Rapphorst Best Paper Award, IEEE Electron Devices Society, 1997, a Best paper award at SCI 2001, Best Paper at CICC 2005, Best Sensor Track paper at ISCAS 2005, Best paper award at Ultrasound Symposium, 2006, Best Demonstration paper award, ISCAS 2010, Best paper award at SCI 2001, 2nd Place, Student Paper Award, IEEE Sensors Conference. Dr. Hasler has been an author on over 400 journal and referenced conference papers.

**Dhireesha Kudithipudi, Ph.D.**

Professor

The University of Texas at San Antonio
dhireesha.kudithipudi.AT.utsa.edu

Dr. Kudithipudi, Ph.D. is Professor of Electrical and Computer Engineering and Computer Science; Robert F McDermott Chair in Engineering; and Director of the Neuromorphic AI (NuAI) lab at UTSA. She is also the Director of the MATRIX AI Consortium through which she serves the diverse population of San Antonio, Texas. Her research interests are in brain-inspired AI, neuromorphic computing, energy efficient ML, and AI accelerators. She received the Clare Booth Luce Scholarship in STEM for women in higher education (2018), Rochester's Technology Women of the Year (2018), ELATES Fellowship (2022), and San Antonio Lights Award (2022). Her teams' research work has been recognized with multiple best paper awards (CVPR-W, NICE, AI Summit) and featured in several outlets such as Nature Outlook. She actively leads Project Lovelace, supported by Xilinx Foundation among other initiatives. Kudithipudi is a first-generation and first Ph.D. graduate from UT San Antonio's Electrical Engineering program.

Session 4 Speakers

Biography



Gina Adam, Ph.D.
Associate Professor
George Washington University
ginaadam.AT.email.gwu.edu

Dr. Adam's lab develops novel hardware foundations at the intersection of materials, devices, and circuits to enable new ways of computing. Her research interests are focused on emerging nanoelectronic and nanoelectromechanical devices and their integration in beyond von Neumann systems such as computation-in-memory and neuromorphic platforms. Her group innovates at the design, simulation and nanofabrication level with a vision of system-level experimental demonstrations. Recent work has been investigating two-terminal non-volatile memory devices called memristors that have an electrical behavior similar to that of an artificial synapse and can be used for both data storage and processing.



Chiara Bartolozzi, Ph.D.
Senior Researcher
Istituto Italiano Di Tecnologia
chiara.bartolozzi.AT.iit.it

Dr. Bartolozzi is Researcher at the Italian Institute of Technology. She earned a degree in Engineering at University of Genova (Italy) and a Ph.D. in Neuroinformatics at ETH Zurich, developing analog subthreshold circuits for emulating biophysical neuronal properties onto silicon and modelling selective attention on hierarchical multi-chip systems. She is currently leading the Event-Driven Perception for Robotics group, with the aim of applying the "neuromorphic" engineering approach to the design of robotic platforms as enabling technology towards the design of autonomous machines.



Ralph Etienne-Cummings, Ph.D.
Professor
Johns Hopkins University
retienne.AT.jhu.edu

A pioneer for the past three decades in mobile robotics and legged locomotion, Dr. Etienne-Cummings' innovations have the potential to produce computers that can perform recognition tasks as effortlessly and efficiently as humans. He has developed prosthetics that can seamlessly interface with the human body to restore functionality after injury or to overcome disease. Etienne-Cummings is the Julian S. Smith Professor of electrical and computer engineering and is the vice provost for faculty affairs at Johns Hopkins University. He holds a secondary appointment in computer science. He previously served on JHU's Homewood Academic Council and is the former chair of the department of electrical and computer engineering.



Joseph Hays, Ph.D.
Robotics Research Engineer
Naval Research Laboratory
joe.hays.AT.nrl.navy.mil

Dr. Hays is a research scientist at the U.S. Naval Research Laboratory (2011–present) in Washington, DC. His research efforts focus on advancing Edge Intelligence capabilities for robotic systems through neuromorphic processing and low power AI accelerators, event-based sensing, artificial and spiking neural network algorithm development, and high-performance digital twins based modeling and simulation. Prior to NRL, Dr. Hays was a senior engineering manager at National Instruments in Austin, TX, (1998–2007) where he led software development efforts for technologies related to dynamical system hardware-in-the-loop simulation (HIL), control system design, system identification, dynamic system simulation, and real-time embedded computing. He received his Ph.D. degree from Virginia Tech (2007–2011), his MS degree from the University of Washington, Seattle (1996–1997) and a BS degree from Brigham Young University, Provo (1992–1996).



Giacomo Indiveri, Ph.D.
Professor
University of Zurich
giacomo.AT.ini.uzh.ch

Dr. Indiveri is a Professor at the Faculty of Science at the University of Zurich, Switzerland. He obtained an M.Sc. degree in electrical engineering and a Ph.D. degree in computer science from the University of Genoa, Italy. Indiveri was a post-doctoral research fellow in the Division of Biology at Caltech and at the Institute of Neuroinformatics of the University of Zurich and ETH Zurich. In 2006 he attained the “habilitation” in Neuromorphic Engineering at the ETH Zurich Department of Information Technology and Electrical Engineering. He won an ERC Starting Grant on “Neuromorphic processors” in 2011 and an ERC Consolidator Grant on neuromorphic cognitive agents in 2016. His research interests lie in the study of neural computation, with particular interest in spike-based learning and selective attention mechanisms, and in the hardware implementation of real-time sensory-motor systems using analog/digital neuromorphic circuits and emerging VLSI technologies.



Kai Miller, Ph.D., M.D., Ph.D.
Pediatric and Epilepsy Neurosurgeon
Mayo Clinic
miller.kai.AT.mayo.edu

Kai Miller is a pediatric and epilepsy neurosurgeon at Mayo Clinic in Minnesota. He attended the University of Washington for graduate school, obtaining a Ph.D. in Physics, an MD, and a second Ph.D. in Neuroscience. After completing his neurosurgery residency at Stanford University in California, Kai was named as the 2018 Van Wagenen fellow. He completed clinical fellowships at Stanford and Utrecht (Netherlands) in epilepsy, deep-brain stimulation, and tumor resection in children and adults. Dr. Miller joined the neurosurgery staff at Mayo Clinic in Rochester in 2019. In addition to his clinical



William Nourse, Ph.D.
Postdoctoral Scholar
Case Western Reserve University
wrn13.AT.case.edu

practice, he studies basic human neurophysiology and clinical translation for cybernetics, epilepsy and functional neurosurgery. His group, the Cybernetics and Motor Physiology Laboratory, is focused on the creation of new tools to 1) control cybernetic prostheses, 2) induce brain plasticity after injury, and 3) intervene with distributed circuits in neuropsychiatric disease and movement dysfunction.

Dr. Nourse received his Ph.D. in Electrical Engineering from Case Western Reserve University in 2024 and is currently a postdoctoral scholar at that same institution. Dr. Nourse's research aims to understand the fundamentals of neural control and decision-making in animals of different dynamic scales and how to translate these principles to neuromorphic control of legged robotic locomotion. He also acts as the Project Manager for the C3NS: Communication, Coordination, and Control in Neuromechanical Systems (<https://c3ns.org/>) network, funded under the NSF NeuroNex program and the BRAIN Initiative.

Funders Panelists



Terrence Sejnowski, Ph.D.
Moderator

Professor and Laboratory Head of the Computational
Neurobiology Laboratory
University of California at San Diego &
The Salk Institute

Hal Greenwald, Ph.D.

Program Officer
Air Force Office of Scientific Research
hal.greenwald.AT.us.af.mil



Chou Hung, Ph.D.
Program Manager
Army Research Office
chou.p.hung.civ.AT.army.mil

Biography

Dr. Sejnowski is a Professor and Laboratory Head of the Computational Neurobiology Laboratory. He is a pioneer in computational neuroscience and his goal is to understand the principles that link brain to behavior. He received a Ph.D. in Physics from Princeton. He is a Distinguished Professor at the University of California at San Diego and holds the Francis Crick Chair at The Salk Institute. He is a member of the National Academy of Sciences, the National Academy of Medicine and the National Academy of Engineering. In 2024, Dr. Sejnowski was awarded the Lundbeck Foundation's Brain Prize, alongside Larry Abbott and Haim Sompolinsky, for pioneering contributions to computational and theoretical neuroscience.

Dr. Greenwald has been a program officer at the Air Force Office of Scientific Research (AFOSR) since 2018. He manages the Cognitive & Computational Neuroscience program, funding basic research on perception, cognition, and behavior at the intersection of neuroscience and AI. Since 2021, he also oversees the Computational Cognition & Machine Intelligence program, supporting AI research on machine intelligence, autonomy, and human-machine teaming. Previously, Dr. Greenwald spent 10 years at MITRE, leading neuroscience and AI research, advising federal programs, and helping government agencies leverage neuroscience. He also worked for three years as a computer scientist/software engineer at Johns Hopkins University Applied Physics Laboratory. Dr. Greenwald holds a Ph.D. in Brain & Cognitive Sciences from the University of Rochester and dual bachelor's degrees in computer science and psychology from the University of Pennsylvania.

Dr. Hung is the Program Manager for Neurophysiology of Cognition at the DEVCOM ARL Army Research Office. Since 2015, he has been a researcher at the DEVCOM Army Research Laboratory, focusing on human cognition, human-machine interfaces, and bio-inspired AI development. Previously, he was a professor of neuroscience at Georgetown University and National Yang-Ming University in Taiwan, where he investigated neural circuits underlying visual perception. Dr. Hung's research interests span living neurons, circuits, mechanisms, and behaviors related to real-world and augmented perception and performance. His research has explored biological and AI-aided learning and decision-making as well as brain-inspired computational



Robinson Pino, Ph.D.
 Program Manager
 Department of Energy
 robinson.pino.AT.science.doe.edu

principles for novel AIs capable of complex reasoning. Dr. Hung obtained his Ph.D. in neuroscience from Yale University (2002) and completed a DARPA/ONR-funded BMI postdoctoral fellowship at MIT (2002–2005).

Dr. Pino is a Program Manager for the Advanced Scientific Computing Research (ASCR) program office in the U.S. Department of Energy's Office of Science. He previously served as Senior Advisor to the CHIPS Program Office at the National Institute of Standards and Technology, U.S. Department of Commerce. His portfolio focuses on revolutionary basic research and development in high performance computing, edge computing, neuromorphic computing, machine learning, artificial intelligence, photonics, microelectronics, and advanced wireless technologies. These efforts aim to maintain U.S. leadership in exascale computing and beyond, as well as in energy-efficient technologies. Prior to his current role, Dr. Pino was Director of Cyber Research at ICF International. He has a BE in Electrical Engineering, summa cum laude, from the City University of New York, City College, and obtained a MSc with honors and Ph.D. in Electrical Engineering from Rensselaer Polytechnic Institute.



Christine Edwards, Ph.D.
 NSA Representative at
 Defense Advanced Research Projects Agency
 Information Innovation Office
 christine.edwards.AT.darpa.mil

Dr. Edwards has served in diverse roles as a developer, researcher, leader, and technical advisor. As Deputy Chief of the National Security Agency's (NSA) Adaptive Cyber-Defense Systems Research Office, she led an interdisciplinary team investigating trustworthy AI-powered solutions. Previously, as Chief of Multimedia Processing Research, her team was recognized as the NSA Research Team of the Year. Dr. Edwards holds a BS in Electrical Engineering from the University of Maryland and MS degrees in Electrical and Computer Engineering and Applied Biomedical Engineering from Johns Hopkins University. She conducted graduate research at the Mayo Clinic Neural Engineering Laboratory and Department of Neurologic Surgery and obtained a Ph.D. from the Deakin University School of Engineering in Australia. Dr. Edwards' research interests focus on artificial intelligence and neuroscience, exploring their intersection for innovative solutions across multiple application domains.



Grace Hwang, Ph.D.
 Program Director
 National Institutes of Health
 National Institute of Neurological Disorders
 and Stroke
grace.hwang.AT.nih.gov

Dr. Hwang is a Program Director at the National Institute of Neurological Disorders and Stroke, where she manages BRAIN Initiative projects in the Technologies for Neural Recording and Modulation portfolio. Before joining NIH, she was a Program Director at the National Science Foundation while based at Johns Hopkins University with appointments at the Applied Physics Laboratory and Kavli Neuroscience Discovery Institute. At NSF, Dr. Hwang managed the Disability and Rehabilitation Engineering program while spearheading cross-agency initiatives including the Emerging Frontiers in Research and Innovation's Brain-Inspired Dynamics for Engineering Energy-Efficient Circuits and Artificial Intelligence (BRAID) program topic. Her research at Johns Hopkins spanned neuroscience, artificial intelligence, neuromodulation, and brain-machine interfaces. She served as a Principal Investigator on an NIH BRAIN award to investigate neural stimulation using sonogenetics and on an NSF award to develop a brain-inspired algorithm for multi-agent robotic control. She is a co-organizer of the NIH BRAIN NeuroAI Workshop.



Stephanie Gage, Ph.D.
 Program Director
 National Science Foundation
sgage.AT.nsf.gov

Dr. Gage is a Program Director in the Division of Computing and Communication Foundations (CCF) in the Computer and Information Sciences and Engineering directorate at the National Science Foundation (NSF). In 2021, she joined NSF as an AAAS Science and Technology Policy Fellow in the Division of Information and Intelligent Systems, focusing on neuroscience and artificial intelligence initiatives across the agency. In 2023, she became a cluster leader in the Division of CCF, supporting the Foundations of Emerging Technologies program and managing the biological systems portfolio. Before joining the NSF, Dr. Gage's research centered on neuromodulation and behavior in insects. She also completed fellowships with the Agricultural Research Service of the U.S. Department of Agriculture and the Georgia Institute of Technology. She holds a BS in Chemistry from Beloit College and obtained a Ph.D. in Neuroscience from the University of Arizona.

Funders Panelists



Steven Zehnder, Ph.D.
Program Director
National Science Foundation
szechnder.AT.nsf.gov



Jean-Luc Cambier, Ph.D.
Director of Technical Programs
Office of the Under Secretary of Defense,
Research & Engineering (OUSD(R&E))
Jeanluc.cambier.civ.AT.mail.mil



Alyssa Picchini Schaffer, Ph.D.
Vice President and Senior Scientist
Neuroscience Collaborations
Simons Foundation
apschaffer.AT.simonsfoundation.org

Biography

Dr. Zehnder is the Integrative Activities Program Director in the Division of Chemical, Bioengineering, Environmental, and Transport Systems (CBET) at the U.S. National Science Foundation. Previously he was the Associate Program Director for Engineering Biology and Health Programs in CBET and served as the Program Lead for Cellular and Biochemical Engineering, Biophotonics, and Disability and Rehabilitation Engineering. Steven received his Ph.D. in Mechanical Engineering from the University of Florida, studying cellular biomechanics and mechanobiology.

Dr. Cambier is a Program Director at the Office of the Under Secretary of Defense, Research & Engineering (OUSD(R&E)). He oversees the Vannevar Bush Faculty Fellowship (VBFF), the tri-service Laboratory-University Collaboration Initiative (LUCI), and serves as a technical advisor for a range of scientific fields of interest to the Basic Research Office, from mathematics, AI, quantum materials to biology and neuroscience. Prior to that, he was Program Officer, then Technical Advisor at the Air Force Office of Scientific Research (AFOSR). He started his professional career at NASA, the Aeronautical Research Institute in Sweden, and the US Air Force (USAF), and joined the civil service in 2001 at the Air Force Research Laboratory. Initially trained as a theoretical physicist, Dr. Cambier has conducted and directed research in many fields, including applied mathematics, quantum physics, plasma physics, hypersonics and combustion, and computer science.

Alyssa Picchini Schaffer, Ph.D., is a vice president and senior scientist at the Simons Foundation, directing the Neuroscience Collaborations and Pivot Fellowship. She has diverse expertise in neural stem cell biology, pharmacology, policy, and media across business, government, and academic sectors. Picchini Schaffer is passionate about fostering collaboration among multidisciplinary teams to address significant neuroscience questions and promote effective science communication. Previously, she was the scientific director of TEDMED, a TED division focused on science, health, and medicine. She earned her Ph.D. from Columbia University and is an alumna of the AAAS Science and Technology Policy Fellowship. Picchini Schaffer serves on the board of The IDEAL School of Manhattan, an inclusive independent school in NYC, and as board treasurer for the Heartbeat Music Project, which offers music education for Navajo (Diné) K-12 students on the Navajo Reservation in New Mexico.

| Session | Discussant Name, Title, & Affiliation |
|--|---|
| Session 1: Defining NeuroAI for BRAIN: Gaps, Challenges, and Opportunities | Dimitri Yatsenko — CEO, DataJoint |
| | Claire Pelofi — Research Scientist, Music and Audio Research Laboratory, NYU |
| | Frances Chance — Principal Member of Technical Staff, Sandia National Labs |
| | SueYeon Chung — Assistant Professor of Neural Science, NYU & Flatiron Institute |
| Session 2: Exploring the Structural and Functional Convergence of Deep Neural Nets and Brains | Paul Middlebrooks — Research Associate, Carnegie Mellon University & Host of Brain-Inspired Podcast |
| | Mark Histed — Investigator, NIMH Intramural Program |
| | Steven Zucker — David & Lucile Packard Professor of Computer Science & Biomedical Engineering, Yale University |
| | Joshua Vogelstein — Associate Professor of Biomedical Engineering (BME/Stats/Neuro/CS), John Hopkins University |
| | Jia Liu — Assistant Professor, School of Engineering and Applied Sciences, Harvard University |
| Session 3: Advancing Theory for BRAIN through Neuromorphic Computing, Embodiment, and Physical Intelligence | Terrence Sejnowski — Professor, Computational Neurobiology Laboratory, Francis Crick Chair, The Salk Institute for Biological Studies |
| | Chiara Bartolozzi — Senior Researcher, Istituto Italiano Di Tecnologia |
| | Carina Curto — Professor, Brown University |
| | Panayiota Poirazi — Research Director, IMBB-FORTH |
| | Dong Song — Associate Professor of Neurological Surgery and of Biomedical Engineering; Director of the Neural Modeling and Interface Laboratory, University of Southern California |

| Session | Discussant Name, Title, & Affiliation |
|--|--|
| <p>Session 4: Towards Reciprocal BRAIN NeuroAI Advances in Intelligent Computing, Robotics, and Neurotechnologies</p> | <p>Jennifer Hasler — Regents Professor for Entrepreneurship, School of Electrical and Computer Engineering; Georgia Institute of Technology</p> <p>Robinson Pino — Program Manager, Advanced Scientific Computing Research (ASCR), Department of Energy Office of Science</p> <p>Maryam Parsa — Assistant Professor, Electrical and Computer Engineering; George Mason University</p> <p>Bing Brunton (virtual) — Professor & Richard and Joan Komen University Chair; University of Washington</p> <p>Christopher Rozell — Julian T. Hightower Chair in Robotics, Automation and Control, Professor; Georgia Institute of Technology</p> <p>Karen Rommelfanger — Director, Institute of Neuroethics Think and Do Tank</p> |

Speaker Abstracts — Organized by Session Order

Session 1 Presentation Abstracts

Deep Intelligence: Why AI Must Learn from “Nature’s Imagination”

Ali A. Minai, University of Cincinnati

“The imagination of nature is far, far greater than the imagination of man.”
— Richard Feynman

Natural intelligence — the only existing example of general intelligence — emerges in multiscale, highly heterogeneous complex systems as a result of adaptation over many temporal scales — evolution, development, learning, and real-time self-organizing dynamics — instantiating a deep intelligence that is inherently integrated, grounded in experience, and built on billions of years of evolutionary engineering. In contrast, most AI models are initially naive systems with generic, relatively regular architectures that rely on training with huge amounts of data to force poorly constrained induction of complex functional mechanisms, making them inefficient, arbitrary, and vulnerable to having hidden failure modes. To make progress towards deep and scalable general intelligence, AI must actively and thoughtfully incorporate insights from biology, including, neurobiology, developmental learning, and the many functional principles and mechanisms that “Nature’s imagination” has discovered through evolution.

Biocomputing with Astrocytes

Wolfgang Losert, University of Maryland

The living neural networks of brains can rapidly adapt to new contexts and learn from limited data, a coveted performance characteristic that neuroscience aspires to explain and control and that the AI community has struggled to mimic. Progress in understanding how living brains achieve their unique performance has potential impact on both neuroscience and AI. Among several unique characteristics we have identified, here I will highlight the multimodal character of information, focusing on the potential role of astrocytes as carriers of analog information and as enablers of slow integrative processing of information in neural networks.

Closing the loop between neuroscience and with virtual neuroscience

Patrick Mineault, Amaranth Foundation

Recent advances in neurotechnology have provided unprecedented access into brain function, enabling the recording of more neurons with greater coverage and biophysical detail in naturalistic conditions than ever before. Simultaneously, AI has experienced exponential growth, evolving from specialized applications into broadly useful tools. This convergence presents a unique opportunity to create a virtuous circle between neuroscience and AI: building virtual neuroscience. I propose developing digital twins and foundation models that enable in silico experimentation and hypothesis generation to better understand perception, cognition, and behavior. These virtual models allow researchers to simulate neural activity and explore brain function beyond the limitations of traditional experimental methods. This shift toward virtual neuroscience is crucial for accelerating neuroscientific progress to match AI's rapid advancement, potentially giving insights into the development of flexible, safe, and human-compatible AI systems. Together, these complementary approaches have the potential to drive progress in both our understanding of the brain and the capabilities of artificial intelligence.

A Less Artificial Intelligence

Andreas Tolias, Stanford University

Neural activity fundamentally shapes our perceptions, behaviors, and cognition, propelling one of neuroscience's greatest quests: decrypting the neural code. This challenge is hindered by our limited ability to precisely record and manipulate extensive neuronal networks under complex conditions and to accurately model the relationships between stimuli, behaviors, and brain states within the natural world's complexity. Recent advancements have started addressing these barriers. Concurrently, advancements in AI now enable analysis of this complex data, facilitating the construction of brain foundation models. These models, akin to AI systems like Video-LLaMA, which decipher video and language relationships, can systematically compile large-scale neural and behavioral data from diverse natural settings. These digital twins of the brain allow for unlimited in silico experiments and the application of AI interpretability tools, enhancing our understanding of neural computations. By applying these insights to AI, we aim to develop more robust, energy-efficient, and comprehensible systems, advancing beyond Big Tech's practice of scaling models with just more behavioral data. Additionally, brain foundation models could revolutionize the diagnosis and treatments for neuropsychiatric disorders. To effectively build these models, we must now decisively move away from traditional hypothesis-driven neuroscience and commit to generating extensive, combined neural and behavioral data across a range of diverse natural tasks.

Bio-realistic modeling of brain circuits

Anton Arkhipov, Allen Institute

A central question in neuroscience is how the structure of brain circuits determines their activity and function. Answering this requires an ability to simulate the brain at the cellular level. Current developments in experimental techniques, AI, and bio-realistic modeling bring the field to the point where cellular-level simulations at the scale of the whole mammalian brain become feasible — and are already possible for other model species like worms and flies. Therefore, a realistic and highly impactful goal for the next decade is to combine dense reconstructions of the circuitry and neural activity across whole brains (in the mouse) or large portions of the brain (in non-human primate and human brain tissue) with bio-realistic modeling of these reconstructed circuits. Experimentally, this will leverage electron and light microscopy, expansion microscopy, spatial transcriptomics, large-scale optical- and electrophysiology, and associated AI tools. And modeling with AI-assisted training under biological constraints is becoming capable of reproducing not only function but also the structure and mechanisms of the brain circuits. Combining these approaches, the field will be able to create accurate simulations of brains — both the general 'foundation models' and 'digital twins' of individual animals — that will serve as computational platforms for discovery, investigation of diseases and treatments, and testbeds for hypotheses and theories.

Embodied intelligence through integrated neuromechanical models of natural behavior

Bing Brunton, University of Washington

A central question in neuroscience is how the structure of brain circuits determines their activity and function. Answering this requires an ability to simulate the brain at the cellular level. Current developments in experimental techniques, AI, and bio-realistic modeling bring the field to the point where cellular-level simulations at the scale of the whole mammalian brain become feasible — and are already possible for other model species like worms and flies. Therefore, a realistic and highly impactful goal for the next decade is to combine dense reconstructions of the circuitry and neural activity across whole brains (in the mouse) or large portions of the brain (in non-human primate and human brain tissue) with bio-realistic modeling of these reconstructed circuits. Experimentally, this will leverage electron and light microscopy, expansion microscopy, spatial transcriptomics, large-scale optical- and electrophysiology, and associated AI tools. And modeling with AI-assisted training under biological

constraints is becoming capable of reproducing not only function but also the structure and mechanisms of the brain circuits. Combining these approaches, the field will be able to create accurate simulations of brains — both the general ‘foundation models’ and ‘digital twins’ of individual animals — that will serve as computational platforms for discovery, investigation of diseases and treatments, and testbeds for hypotheses and theories.

Leveraging BRAIN Initiative Data Resources to Advance Novel NeuroAI Frameworks and Capabilities

Dominique Duncan, University of South California

The BRAIN Initiative's expanding multimodal data repositories, including the Data Archive for the BRAIN Initiative (DABI), represent a unique opportunity to advance fundamental NeuroAI theories and frameworks. This presentation will demonstrate how these extensive datasets can be leveraged to develop and validate novel NeuroAI approaches that bridge neural mechanisms across scales. Opportunities for developing new NeuroAI algorithms that can extract principles of neural computation from complex, multiscale datasets will be explored. The BRAIN Initiative increasingly brings together multidisciplinary expertise - from neuroscientists and engineers to theorists and data scientists - who would benefit from unified data infrastructure to advance NeuroAI research. Current gaps in existing data infrastructure that could enable transformative NeuroAI capabilities will be identified, including needs for standardized cross-modal data integration, common theoretical frameworks, and validation metrics for comparing biological and artificial neural networks. Future directions will focus on how enhanced data infrastructure across BRAIN Initiative resources could enable breakthrough NeuroAI theories about fundamental principles of intelligence and computation in biological systems.

NeuroAI Ethics: A proactive approach for the next 5–10 years (and beyond)

Karen S. Rommelfanger, Institute of Neuroethics Think and Do Tank

We are experiencing a growing convergence of neuroscience and AI methods. Neuroscience is informing AI techniques and AI is enhancing neuroscience discovery and enabling more sophisticated neurotechnologies. NeuroAI offers a clear promise for understanding brain function and dysfunction; advances in clinical diagnostics, treatment and restoration as well as brings potential for individual and human-AI enabled augmentation. The brain is a privileged site from which human identity, agency, autonomy, emotion, thought, and our overall lived experience arise. Therefore, each context in which we intervene with and explore the brain — from research and clinical applications to beyond the bench and clinic — will raise tensions and have potential ethical implications. AI applications alone have demonstrated why a reactive mitigation strategy will not suffice for NeuroAI. This short talk will offer a glimpse into the types of issues that arise in NeuroAI research and offer proactive ethical considerations for NeuroAI researchers.

Placing the field of NeuroAI in context — What is it, where does it come from, and where will it go?

Blake Richards, Mila – Quebec Artificial Intelligence Institute

Cognitive science promised that we could study intelligence without having to consider the link between implementation and algorithms, nor worry about low-level control tasks. But this promise did not pan out. AI ended up being most successful when researchers found algorithms that worked well with specific parallel hardware, and when we embraced feedback-based learning and control. Similarly, attempts to understand the human mind without considering the links between algorithms and neural circuits, and how they engage in feedback-based control, did not work out. I will propose that "NeuroAI" is essentially the original promise of cybernetics finally coming to fruition — a general science of intelligent systems that thinks hard about how to implement algorithms in parallel distributed hardware, whether it be in a brain or on a chip, and how to use feedback-based control to accomplish goals.

Fulfilling the potential of NeuroAI

Doris Tsao, University of California, Berkeley & HHMI

Current NeuroAI efforts primarily involve neuroscientists adapting off-the-shelf AI models to compare brain activity with network activation patterns. Although these efforts have advanced our understanding of brain structures like inferotemporal cortex, these approaches capture only a fraction of NeuroAI's potential. The most transformative impact of machine learning on neuroscience and vice versa is likely to occur at the conceptual level. Models like ChatGPT, for example, demonstrate how intelligence can emerge from simple objectives, inspiring new ways to think about cognition. Neuroscience, in turn, holds the potential to make profound contributions to AI by inspiring architectures that can achieve human-level perception and cognition, all while operating within the brain's efficient energy limits.

Recent research by Sejnowski and colleagues (Muller, et al., TINS, 2024; Gu & Dao, arXiv, 2023) exemplifies this potential, demonstrating how cortical waves—grounded in a wealth of neuroscience data—can implement functions similar to transformers. Likewise, foundational psychological principles, such as the brain's use of visual surfaces and object pointers, suggest structural priors that could inform the next generation of neural networks. These biologically-inspired frameworks challenge the prevailing trend toward large, generic architectures as advocated by Sutton's "The Bitter Lesson."

Achieving meaningful integration between AI and neuroscience demands collaboration across neuroscience, psychology, and machine learning. The need for theoretical integration is especially acute given the distinct goals of AI (focused on performance) and neuroscience (aimed at understanding the brain's principles). I propose cross-disciplinary foundation models co-designed by neuroscientists,

psychologists, and AI researchers that would leverage insights across disciplines to define, validate, and refine computational frameworks that explicitly reflect biological intelligence. This effort promises to create AI systems that not only scale but also embody the remarkable efficiency and adaptability of biological brains.

Towards advanced NeuroAI systems with dendrites

Panayiota Poirazi, Foundation for Research and Technology – Hellas

Dendrites, the receiving ends of neurons, are crucial for biological intelligence. Incorporating dendritic features into AI systems can enhance energy efficiency and address challenges like noise robustness and catastrophic forgetting. However, the specific dendritic aspects that empower brain functions and how to integrate them into AI remain unclear, hindering the development of advanced neuroAI systems. This challenge requires interdisciplinary collaboration and new experiments to explore the anatomical, biophysical, and plasticity properties of dendrites across various neuron types and species. Bio-realistic computational models can aid experimentation by integrating these properties and evaluating their joint effects at neuronal and circuit levels. Additionally, developing new mathematical formulations and metrics to capture key dendritic functionalities is essential for developing powerful dendritic AI systems. In return, by leveraging dendritic advantages, such AI systems can enhance our understanding of biological design principles and their evolutionary significance.

The future of NeuroAI: inspiration from insects and mathematics

Carina Curto, Brown University

The current state of artificial intelligence relies on very large networks, very large data sets (for training), and increasing amounts of energy. At the same time, these networks lack several key ingredients that seem important in natural intelligence. This includes neuromodulation, the special role of inhibition, rhythms and oscillations, non-synaptic signaling, and dendritic computation. How might these mechanisms make AI more powerful and efficient? With the recent completion of the fly connectome, the opportunity to learn from small, embodied brains is bigger than ever. A detailed understanding of how small brains can perform a rich variety of complex tasks will inspire new principles for designing artificial networks. Mathematical advancements will also be critical. Neural networks are complex, high-dimensional dynamical systems. This would be okay if they were also linear. However, the nonlinearities are essential in both natural and artificial settings, and they pose deep mathematical challenges. Advances on the math side promise to enable more efficient, more robust, and more interpretable AI.

Neural network language models as models of language processing in the human brain

Evelina Fedorenko, Massachusetts Institute of Technology

A network of left frontal and temporal areas in the human brain supports language processing. This “language network” a) is robustly dissociated from lower-level speech perception and articulation mechanisms, and from systems of reasoning (Fedorenko et al., 2024, NRN); and b) supports computations related to retrieving words from memory and building syntactic structures in the service of semantic composition (Shain & Kean, et al., 2024, JOCN). However, a mechanistic-level understanding of how we extract meanings from word sequences, or express meanings through language has remained elusive, in large part due to the limitations of human neuroscience approaches. Recently, a new candidate model organism emerged, albeit not a biological one, for the study of language — neural network language models (LMs). These models exhibit human-level performance on diverse language tasks, and their internal representations are similar to the representations in the human brain when processing the same linguistic inputs (Schrimpf, et al., 2021, PNAS). I will talk about how we can use LMs to evaluate hypotheses about language processing, development, and impairments at an unprecedented granularity and scale. I will also touch on how neural networks can be combined with symbolic architectures to investigate how the language system may interact with systems of thought.

Translating NeuroAI to Integratively Model the Brain Systems Underlying Cognitive Behavior

Martin Schrimpf, École Polytechnique Fédérale de Lausanne

The last decade has seen the rise of AI models to explain the brain and mind. Mapped onto brain regions, these models predict neural activity within sensory cortices and also higher-level systems like language; with model alignment typically driven by task performance such as object classification or next-word prediction. In the coming decade, I believe our field can make substantial progress by expanding the breadth and depth of NeuroAI models which will facilitate their translation into clinical applications. Realizing this vision requires large-scale, high-quality data made available across multiple brain systems, to enable integrated models that connect cognitive behaviors to multi-system neural mechanisms and perhaps even biophysical details. Advancing NeuroAI models will demand us to embrace learnings from machine learning, and further capture the intricacies of the brain such as topography and embodiment. Finally, I argue that we should leverage the best models towards clinical applications such as dyslexia and vision impairments, e.g. with targeted stimulus presentations and model-predicted stimulation patterns. By fostering integrated models of brain and behavior, NeuroAI will not only deepen our understanding of the neural mechanisms underlying cognition but also potentially transform clinical approaches to brain-related disorders.

How Neuromorphic Computing Can Help Us Understand the Brain

Brad Aimone, Sandia National Laboratories

Neuromorphic computing is increasingly being explored for use in artificial intelligence and energy-efficient computing, but its potential impact on neuroscience research has not been fully realized. That said, today's neuromorphic systems are reaching brain-like scales and can emulate more of the brain's complex dynamics than ever before. In my talk, I will briefly describe two potential paths by which neuromorphic computing can impact neuroscience research and neural health. First, can neuromorphic computing really help enable brain-scale simulations? Here, I will present early evidence that the Loihi 2 chip can simulate a full-scale computational neuroscience model. Second, can understanding the computational tasks that neuromorphic computing excels at provide clues to new perspectives of the brain's computations? To this end, growing evidence that neuromorphic hardware can implement complex numerical computing methods may indicate that the brain's computations are more sophisticated than previously considered.

Scaling Knowledge Processing from 2D Chips to 3D Brains

Kwabena Boahen, Stanford University

Artificial intelligence (AI) realizes a synaptocentric conception of the learning brain with dot-products and advances by performing twice as many multiplications every two months. But the semiconductor industry tiles twice as many multipliers on a chip only every two years. Moreover, the returns from tiling these multipliers ever more densely now diminish, because signals must travel relatively farther and farther, expending energy and exhausting heat that scales quadratically. As a result, communication is now much more expensive than computation. Much more so than in biological brains, where energy-use scales linearly rather than quadratically with neuron count. That allows an 86-billion-neuron human brain to use as little power as a single light-bulb (25 W) rather than as much as the entire US (3 TW). Hence, rescaling a chip's energy-use from quadratic to linear is critical to scale AI sustainably from 10^{12} parameters (mouse scale) today to 10^{15} parameters (human scale) in the near future. But this would require communication cost to be reduced radically. Towards that end, I will present a recent reconception of the brain's fundamental unit of computation that sparsifies signals by moving away from synaptocentric learning with dot-products to dendrocentric learning with sequence detectors.

To Silicon Columns and Beyond: Looking for a Computational Framework for Neuromorphic Systems

Jennifer Hasler, Georgia Institute of Technology

The neural roadmap paper (Hasler & Marr, 2013, Front Neurosci) showed an all silicon roadmap towards building a synthetic electronic structure to parallel the electronic structure of human cortex. The opportunity is to start building key components of that roadmap. Physical silicon implementation create the opportunity for energy efficient programmable and adaptable neural models that can build thousands and millions of realistic neurons on an IC or on stacks of ICs on a PC board.

The question is building this neuromorphic hardware & software infrastructure as well as building the computational models of significant groups of pyramidal cell and other neurons often found in a cortical column. The architecture needs to follow biological concepts, particularly that neural connectivity is local and sparse, primarily due to energy constraints.

Neural Primitives as the Missing (Synergistic) Link between Neuromorphic Computing and AI

Frances Chance, Sandia National Laboratories

Neuromorphic computing is an approach to engineering computer paradigms that mimic the structure and function of biological brains. My research has focused on identifying “computational primitives” that are widely used by biological nervous systems, and using underlying biological mechanisms to guide development of neuromorphic implementations of these primitives. For example, coordinate transformations are a fundamental computation essential for any behavior reliant upon sensorimotor processing (e.g., reaching and prey hunting). I will discuss how my collaborators and I have been developing neuromorphic emulations of different neuroscience models to identify neuromorphic approaches that are best suited for the constraints faced by artificial systems.

I hypothesize that neural computational primitives form the building blocks of “higher-level” behaviors including cognition. Exploring neuromorphic implementations of these neural primitives could therefore be critical for designing energy-efficient neural-inspired AI. Moreover, understanding how constraints facing engineered systems limit hardware implementations of AI-models may grant insight into the function of biological circuits and how they are optimized for biologically-specific constraints.

Learning from Neural Manifolds: From Biological Efficiency to Engineered Intelligence

SueYeon Chung, New York University & Flatiron Institute

Recent breakthroughs in experimental neuroscience and machine learning have revealed striking parallels in how biological and artificial systems process information across multiple scales. The next decade presents exciting opportunities to bridge neuroscience and AI. Our research proposes that geometric principles of neural representation and computation could revolutionize how we design AI systems while deepening our understanding of biological intelligence. This vision requires four key advances: new experimental technologies capturing the changes of neural manifolds across dynamical and learning timescales during behavior; theoretical frameworks that unite single-neuron properties with population-level computation while revealing principles of efficient information processing; a theory of cross-modal representations that explains how neural manifolds and their transformations preserve efficiency, and robustness principles across sensory, motor, and cognitive regions while supporting domain-specific adaptations; and scalable computational tools for analyzing massive-scale neural recordings across both biological and artificial systems to extract core efficiency principles. Our current work, combining statistical physics, machine learning, and geometry, lays the groundwork for this future. By understanding how neural representations evolve across scales — from individual neurons to population activities to cognitive functions — we can develop AI architectures that better reflect the efficiency and robustness of biological systems. This approach promises not just better models of the brain, but fundamentally new principles for artificial intelligence that capture the robust, embodied, adaptive nature of biological computation.

A Co-Design Approach to Continual Learning: Exploring Synergies Between Neuroscience and Neuromorphic Hardware

Dhireesha Kudithipudi, The University of Texas at San Antonio.

Presenter’s Note: In collaboration with Nicholas Soures, Fatima Tuz Zhora, Vedant Karia, Neuromorphic AI Lab UT San Antonio.

Emulation of neural processes can improve the ability to generate highly functioning continual or lifelong learning machines. Recent advances in understanding key mechanisms, such as neuromodulation, metaplasticity, reactivation, neurogenesis, and memory consolidation, are poised to inspire new learning algorithms. Despite progress, the interplay of these diverse mechanisms remains largely underexplored, representing a crucial avenue for advancing continual learning models. By identifying the core features

essential for these advanced models, we can create new architectures tailored to neuromorphic hardware. For instance, integrating probabilistic switching and the inherent variability of non-volatile memory to represent the plasticity mechanisms can lead to more adaptable and resilient architectures. Additionally, structural plasticity can be achieved through fine-grained runtime reconfigurability units within the memory. This specialized hardware facilitates rapid prototyping and generates hypotheses informed by experimental data, thereby advancing low-power adaptive applications in medical devices, sensors, and personalized AI assistants in healthcare. A long-term strategy should emphasize a plug-and-play modular approach to integrate various plasticity mechanisms into neuromorphic hardware architectures. This would enable us to design lifelong learning machines without the need to specify explicit end goals.

Embodied intelligence through the integration of biomechanics and neuroscience

Mitra Hartmann, Northwestern University

The nervous system of an animal species co-evolves with its sensory and motor systems. Understanding neural activity thus inherently involves understanding how neural responses are tuned to the sensory inputs they receive and the motor outputs they control. Moreover, because many perceptual processes rely on closed-loop sensorimotor control — in which perception is directly shaped by the animal's active control of sensory data acquisition — it is evident that understanding neural function will ultimately require integrating accurate biomechanical models of sensors and muscles with neurophysiological data. In this presentation, I discuss the rodent vibrissal (whisker) system as an effective model for studying “embodied intelligence,” integrating the fields of biomechanics and neuroscience. The long-term goal is to close the loop between whisker-based tactile sensing, the nervous system, and the muscles driving whisker movement. Both simulations and hardware models are essential components of the work.

An Interdisciplinary Vision in Neuromorphic Technologies for Computing

Gina Adam, George Washington University

Artificial intelligence systems are expected to consume increasing amounts of computing resources in the coming decades at significant financial and environmental costs. New devices and hardware alternatives are necessary to keep up with the increasing demand in complexity and energy efficiency required and make the transition to physical AI frameworks of relevance to robotics, neuro-control and prosthetics. In this brief talk, I will highlight some of the impressive innovations made in the development of neuromorphic hardware in the past four decades. I will also discuss our vertically-integrated approach to contribute to the incorporation of emerging technologies, such as memristors in neuromorphic computing. I will end with a summary of the interdisciplinary efforts at the 2024 Neuromorphic Computing for Science Workshop sponsored by the U.S. Department of Energy, Office of Advanced Scientific Computing Research in September 12–13, 2024. This workshop brought together a diverse range of experts in microelectronics, neuroscience and large-scale modeling and simulation. The participants discussed key research needs, challenges, and next steps necessary to develop scalable biologically-realistic neuromorphic circuits primitives that capture the functionality of neural systems found in nature and proposed four priority research directions in neuromorphic computing for science.

Neuromorphic Embodied Intelligence

Chiara Bartolozzi, Istituto Italiano Di Tecnologia

A pragmatic approach to the development of technology is to look at existing systems and capturing their working principles in artificial implementations. Neuromorphic engineering looks at the computational principles of the nervous system and is therefore suited to implement artificial systems that solve those tasks in which nervous systems excel. In the most general way, such tasks entail extracting information from the external world to produce appropriate behavior. Specifically, nervous systems are integrated in bodies, with sensors, to acquire information and limbs, to move and act.

Since the first prototypes of neuromorphic vision sensors and computing devices, part of the community focused its efforts in deploying neuromorphic systems that exploit neural computational principles in practical applications, e.g. robotics. There are examples of building blocks for sensing, perception, control and decision making, but only very few fully integrated systems, end-to-end neuromorphic that scale beyond proof of concepts. There is also the need to consider embodiment in the development of intelligent artificial agents, whereby the movement of sensors is not a nuisance to cancel, but a resource to generate useful information, and the morphology of body and sensors can simplify information processing.

Perspective On NeuroAI's Relationship With Edge Intelligent Embodied Continual Learning Agents

Joseph Hays, Naval Research Laboratory

The US Naval Research Laboratory (NRL) is actively researching intelligent autonomous systems, or embodied intelligence, which must function within physical constraints. These systems require significant edge computing power, but traditional methods increase size, weight, and power consumption (SWaP). To address this, NRL is investigating neuromorphic computing, which combines computation and sensing to reduce SWaP.

NRL is developing intelligent service robots for space and naval applications. These robots must be capable of constructing structures in space, performing ship maintenance, and continuously learning new skills in the field. By formulating these challenges as spiking neural networks, NRL aims to deploy them on low-SWaP neuromorphic hardware. To realize these ambitious goals, NRL is collaborating with a multidisciplinary research community. This collaboration will not only advance the development of edge

intelligent embodied continual learning agents but will hopefully assist the neuroscience research community in deepening their understanding of the principles and mechanisms involved in biological computation.

Towards Insect-Scale Intelligence for Robotics

William Nourse, Case Western Reserve University

With modern advances in control theory and artificial intelligence (AI), modern robots are capable of performing nearly any individual task, from climbing ladders to folding laundry. Looking ahead, how can we design systems which can not only perform a variety of tasks, but also decide which tasks to do using context-dependent decision-making? How can we do this using only onboard computation, without relying on the cloud? We currently have a connectome of all the neurons in the brain and ventral nerve cord of the fruit fly, and these animals are able to achieve a wide range of behaviors while autonomously switching between them based on both internal and external states. For these reasons, the insect nervous system is a prime template for creating autonomous machines. However, further work is needed before this will be possible. More connectomic information is needed, both in terms of detection resolution and across multiple individuals. Additionally, connectomes of insects with more sophisticated behavior such as mantises would help understand how these regions scale with intelligence. Moving beyond point-to-point graphs, dendritic and axonic structure may be necessary for some computations. Finally, more neuromorphic hardware is needed which can simulate millions of neurons while being easy to physically obtain.

Mixed-signal neuromorphic Systems for next-generation Brain-Computer Interfaces

Giacomo Indiveri, University of Zurich

Traditional Artificial Intelligent (AI) algorithms and technologies, while effective at analyzing large digital datasets, encounter limitations when applied to real-time processing of sensory data in closed-loop systems, particularly in the domain of reciprocal BRAIN–NeuroAI interfaces and neurotechnologies requiring real-time interaction with the nervous system. Next to the challenges related to the need for low-latency and secure local processing to mitigate privacy concerns, these limitations include critical power consumption constraints: as both wearable and implantable neural interface types of devices need to operate continuously for tasks such as real-time anomaly detection, they require extremely low power consumption, often within sub-milliwatt ranges. The requirements to minimize power consumption combined with the need to establish a continuous dialog with real neurons and the signals they produce naturally point to the adoption of a bottom-up physics approach, such as the one based on the use of analog neuromorphic electronic circuits and mixed-signal neuromorphic processing systems. In this presentation we will show how neuromorphic systems comprising passive subthreshold analog circuits and data-driven encoding and signal transmission methods can solve complex classification problems in the domain of epilepsy and biomedical signal processing, using microwatt power budgets.

Synthesis of Neuromorphic Principles in Biomedicine and Healthcare Workshop

Ralph Etienne-Cummings, Johns Hopkins University

The workshop brought communities together to create a new generation of biomedical and neuroengineering technologies that operate with extreme energy and data efficiency, adaptability, and performance advantages compared to current approaches, while staying informed by needs of constituent researchers, clinicians, prosthetists, medical devices developers and entrepreneurs. The two-day workshop included two keynote addresses, from a leading neuromodulation expert — Dr. Tim Denison — and a material scientist — Dr. Zhenan Bao — for wearable electronics. Twelve invited

presentations followed by moderated discussion sessions with 20 experts, and discussions with government stakeholders were also scheduled.

The discussion focused on the value of neuromorphics in the clinic, accentuating that their algorithmic richness, power efficiency and small size make it ideal for applications involving closed-loop, low latency, adaptive and mobile applications. Nonetheless, the group recognized that biological parsimony is not always necessary to solve many clinical problems. Further, given the emergence of exotic materials and sensors with biomedical applications, there needs to be a national strategy to encourage innovation by sharing expertise, design, manufacturing, testing and validation resources. Such a strategy, supported by a robust theoretical harness, is central to ensuring that the promise of neuromorphics is realized for biomedicine and healthcare.

Artificial Intelligence and the Functional Neurosurgeon

Kai Miller, Mayo Clinic

Functional neurosurgeons invasively interact with neural circuitry to measure & modify its computational purpose with electrodes, tissue destruction, and emerging genomic tools. There are a number of functionalities of artificial intelligence that may assist these doctors: identifying structure in biological measurements, documentation and chart synthesis, and clinical prediction. While these functionalities may help at many different points in the cycle of patient care, subsequent discussion will center on closed-loop devices, proposing first to match the measurement scale of implanted devices to the physical scale of the neurophysiological feature (Embodied measurement hardware), and then to implement neurologically-inspired algorithms to match the natural statistics and dynamic variation of brain circuitry (Neuromorphic computing). Each concept is illustrated in an intuitive way, and the presentation will present a concrete framework to facilitate discussion for how reciprocal advances in neuroscientific and algorithmic tools can benefit human patients.

BRAIN NeuroAI Early-Career Scholar Poster Abstracts

Learn more about the NeuroAI research described below by attending the BRAIN NeuroAI Early-Career Scholar Poster Blitz and Poster Session on the morning of Wednesday, November 13. The poster abstracts for presenters selected for the BRAIN NeuroAI Early-Career Scholar Poster Blitz are indicated by a star emoji at the end of the abstract title.

POSTER #1

Recurrent cortical networks encode natural sensory statistics via active filtering of sequences ☆

Ciana Deveau, National Institutes of Health/NIMH

In daily life, organisms interact with a sensory world that dynamically changes from moment to moment. Recurrent neural networks can generate dynamics, but in sensory cortex any dynamic role for the dense recurrent excitatory-excitatory network has been unclear. In this work we show a new role for recurrent connections in mouse visual cortex: they support powerful dynamical computations, but via filtering sequences of input instead of generating sequences.

Using two-photon optogenetics, we measure responses to natural images and play them back, showing amplification when played back during the correct movie dynamic context and suppression in the incorrect context. The sequence selectivity depends on a network mechanism: inputs to groups of cells produce responses in non-targeted local neurons, which interact with and change responses to later inputs. We confirm this mechanism by designing sequences of inputs that are selectively amplified or suppressed by the network. These observations support the idea that the visual cortex recurrent network is filtering sequences of input, specifically amplifying input sequences corresponding to natural vision.

Recurrent neural networks in artificial systems are also often used to create temporally-structured computations. Therefore, we examined an RNN trained to preferentially amplify some input sequences and found it showed the context-dependent effects seen in our experiment: a single pattern extracted from the sequence and played back in the correct context produced an amplified response, compared to when it was presented in the incorrect context. Our model data aligns with the understanding that recurrent artificial networks can learn temporal statistical structure, as seen also in transformers that generate highly complex natural language sequences. Thus, densely-connected recurrent networks seem to be useful for sequence processing both in artificial systems and in biological brains.

Together, these results suggest a novel function, sequence filtering, for recurrent connections. The implication is that the many recurrent excitatory-excitatory connections learn via development and experience the statistics of the natural world, encoding this information in recurrent synaptic weights.

A Rapid Adapting and Continual Learning Spiking Neural Network Path Planning Algorithm for Mobile Robots ☆

Harrison Espino, University of California, Irvine

Mapping traversal costs in an environment and planning paths based on this map are important for autonomous navigation. We present a neurobotic navigation system that utilizes a Spiking Neural Network (SNN) Wavefront Planner and E-prop learning to concurrently map and plan paths in a large and complex environment. We incorporate a novel method for mapping which, when combined with a Spiking Wavefront Planner (SWP), allows for adaptive planning by selectively considering any combination of costs. The SWP is compatible with neuromorphic hardware and could be used for applications requiring low size, weight, and power. The system is tested on a mobile robot platform in an outdoor environment with obstacles and varying terrain. Results indicate that the system is capable of discerning features in the environment using three measures of cost, (1) energy expenditure by the wheels, (2) time spent in the presence of obstacles, and (3) terrain slope. In just twelve hours of online training, E-prop learns and incorporates traversal costs into the path planning maps by updating the delays in the SWP. On simulated paths, the SWP plans significantly shorter and lower cost paths than A* and RRT*. Our algorithm is lightweight and has the potential for neuromorphic applications at the edge, which will be explored by our group in the near future.

Bio-Inspired Front-End for Deep Audio Processing ☆

R. Leslie Famularo, University of Maryland

While models in audio and speech processing are becoming deeper and more end-to-end, they as a consequence need expensive training on large data, and often lack robustness [1]. We build on a classical model of human hearing [2] and make it differentiable, so that we can combine traditional explainable biomimetic signal processing approaches with deep-learning frameworks. This allows us to arrive at an expressive and explainable model that is easily trained on as few as a few hours of data. Particularly, our model is differentiable all the way from the cochlear to the cortex (see figure), allowing parameters to be jointly fitted along with deep learning model parameters.

We apply this model to audio processing tasks, including classification and enhancement. Results show that our differentiable model surpasses black-box approaches in terms of computational efficiency and robustness, even with little training data. The advantage of our model was large especially when the dataset was small, making our model a better candidate in low-resource settings. Additionally, the trained model parameters are explainable, matching characteristics of the training audio data when the training data is controlled. Our approach also has clinical potential for hearing aid fitting.

Neuroethical concerns arise when AI models influence decision-making in clinical applications, such as hearing aid fitting, where transparency and interpretability are essential. Additionally, from an AI ethics perspective, our approach emphasizes the need for explainability, reducing the risks of bias and ensuring that the technology remains accessible and interpretable, particularly in sensitive healthcare contexts.

References:

[1] Wu, Haibin, et al. "Characterizing the adversarial vulnerability of speech self-supervised learning." ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2022.

[2] Chi, Taishih, Powen Ru, and Shihab A. Shamma. "Multiresolution spectrotemporal analysis of complex sounds." *The Journal of the Acoustical Society of America* 118.2 (2005): 887-906.

[3] Crosse, Michael J., et al. "The multivariate temporal response function (mTRF) toolbox: a MATLAB toolbox for relating neural signals to continuous stimuli." *Frontiers in human neuroscience* 10 (2016): 604.

Universality of Representation in Biological and Artificial Neural Networks

Eghbal Hosseini, Massachusetts Institute of Technology

Artificial neural networks (ANNs) have emerged as computational systems that align with behavior and underlying representations in biological neural networks (brains). Across domains, these feats are achieved by many different kinds of ANNs trained with ecologically valid objectives (Conwell, Prince, Kay, Alvarez, & Konkle, 2023; Schrimpf et al., 2021). Here we show that — akin to biological evolution where distinct organisms often converge on a similar solution to some target problem — models' ability to predict brain responses is a consequence of convergence onto universal representational axes that are shared both across high-performing models and between models and brains. First, we introduce model agreement as a measure of representation universality across ANNs (Golan, Raju, & Kriegeskorte, 2020; Platt, 1964). Second, we use model agreement to modulate the degree of match between individual ANNs and the brain for the language and visual systems (Allen et al., 2022), and show that convergence across ANNs leads to convergence across brain representation even without an alignment function (regression in this case). Third, we show that in the visual system agreement across brains modules both degree of agreement across ANN as well as their alignment with the brain. Finally, we begin to identify behavioral dimensions that distinguish between universal and model-specific representations, and show that in the language domain perceived frequency and meaning generality of stimuli correlates with their universality. These results in tandem establish the universality of representation as a core component in the alignment between ANNs and biological systems, thus providing a novel approach for using ANNs to uncover representations and computations in the brain.

References:

Allen, E. J., St-Yves, G., Wu, Y., Breedlove, J. L., Prince, J. S., Dowdle, L. T., ... Kay, K. (2022). A massive 7T fMRI dataset to bridge cognitive neuroscience and artificial intelligence. *Nature Neuroscience*, 25(1), 116–126.

Conwell, C., Prince, J. S., Kay, K. N., Alvarez, G. A., & Konkle, T. (2023). What can 1.8 billion regressions tell us about the pressures shaping high-level visual representation in brains and machines? (p. 2022.03.28.485868). <https://doi.org/10.1101/2022.03.28.485868>.

Golan, T., Raju, P. C., & Kriegeskorte, N. (2020). Controversial stimuli: Pitting neural networks against each other as models of human cognition. *Proceedings of the National Academy of Sciences of the United States of America*, 117(47), 29330–29337.

Platt, J. R. (1964). Strong Inference: Certain systematic methods of scientific thinking may produce much more rapid progress than others. *Science*, 146(3642), 347–353.

Schrimpf, M., Blank, I. A., Tuckute, G., Kauf, C., Hosseini, E. A., Kanwisher, N., ... Fedorenko, E. (2021). The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences of the United States of America*, 118(45). <https://doi.org/10.1073/pnas.2105646118>.

Exploring NeuroAI Models Of How Learned Behavior Can Evolve Into Instinct

Christos Karageorgiou Kaneen, Cold Spring Harbor Laboratory

From mate-calling to web-building, animals exhibit remarkable phenotypic diversity. However, the evolutionary origins of such behavioral innovations remain uncertain. The Baldwin effect offers a possible explanation, suggesting that behaviors initially acquired through learning can, over generations, become genetically assimilated. This theory posits that new environmental pressures drive the gradual inheritance of learned behaviors, which eventually become more innate. While computational studies have demonstrated the usefulness of the Baldwin effect in accelerating evolution (Hinton & Nowlan, 1986) and shaping neural network parameters (Fernando et al., 2018), models detailing the evolution of learned behaviors into instincts are still lacking. Here, we construct artificial agents whose genotypes encode neural network connectivity parameters (Lachi et al., 2024) and learning rates. By simulating Baldwinian evolution on these agent populations, we show that the time required to learn non-trivial tasks (e.g., MNIST) significantly decreases compared to classic Darwinian evolution, where selection is based solely on at-birth performance. We also evaluate our populations, evolved for learning propensity, on unseen datasets, not included in the optimization process. Again, we observe a notable increase in learning speed, highlighting the advantage of Baldwinian adaptation in transfer learning — a hallmark of flexible cognition not yet fully achieved by machine learning systems. Our findings offer insights into how genetically-hardwired traits can emerge without direct inheritance of plasticity-induced changes — a potential mechanism for the rapid emergence of complex cognitive phenomena such as language and abstract reasoning in humans.

G. E. Hinton & S. J. Nowlan. How learning can guide evolution. *Complex Systems* 1(3), 1987 pp. 495–502.

Fernando C, Sygnowski J, Osindero S, Wang J, Schaul T, Teplyashin D, Sprechmann P, Pritzel A, Rusu A. Meta-learning by the baldwin effect. *Proceedings of the Genetic and Evolutionary Computation Conference Companion*; 2018 Jul 15–19; Kyoto Japan.

Divyansha Lachi, Ann Huang, Augustine N. Mavor-Parker, Arna Ghosh, Blake Richards, Anthony Zador. Stochastic Wiring of Cell Types Enhances Fitness by Generating Phenotypic Variability. *bioRxiv* 2024.08.07.606541.

Neurogenesis-inspired Neuronal Models for Network Training and Neuromorphic Translation

Joseph Kilgore, George Washington University

The hippocampus, a core brain region for learning and memory, hosts a unique phenomenon in the adult mammalian brain: new neurons via neurogenesis in its front-end region, called the dentate gyrus (DG). While this phenomenon and its role in pattern separation is still under study in the neuroscience community, it could be critical for bio-inspired artificial neural networks with superior performance in continual learning. In this study, we take the first steps to investigate the impact of neuronal age and its firing behavior on network performance to inform future neurogenesis-inspired network learning. Currently, large-scale DG models are available but are computationally expensive and do not include neurogenesis (Fig. 1a) [1]. In this work we prepare to scale the model down for more efficient simulation with minimal compromise to biological-variety, and the incorporation of neurogenesis within the granule cell population (Fig. 1a). Using experimental recordings from the literature, we develop models to match the spiking behavior seen in young, and mature-aged DG granule cells, and set up a variety of energy efficient digital and analog hardware implementations.

A parameter tuning environment was set up to match firing behavior as categorized by Hippocampome.org and minimize the Van Rossum distance across multiple different models [1], [2]. Initial results show fitting of behavior between various model types including 9-parameter Izhikevich models, 4-parameter Izhikevich models, and a potential analog hardware implementation in CMOS 130nm circuitry (Fig. 1b) [3], [4]. These models can also be tuned to fit the experimental data of aging DG granule cells. Additionally, once tuned these parameters can be used to inform hardware development and implementation for age-informed hardware neuron models. These models are then used to build small-scale feed-forward spiking neural network variants. Network structures with a 100-neuron hidden layer containing a 90-10 and 80-20 ratio of mature-to-young aged neurons are built respectively inspired by the typical percentage of neurogenesis in the mammalian brain. Additional network variants with a hidden layer composed entirely of either young or mature-aged neurons serve as a benchmark. These networks are trained using surrogate gradient backpropagation through time, which we have adapted to Izhikevich model neurons. The results show that the higher threshold voltage of the mature-aged neurons causes the network to train slower, mirroring the lower synaptic plasticity in aging biological neural networks (Fig. 1c). Alternatively, the lower threshold young-aged neurons fire more often, expending more energy than their mature-neuron counterparts for comparable accuracy. These results will form the basis for incorporating neurogenesis-based neuron models in larger studies of bio-realistic neuromorphic systems. Further work will explore more sophisticated Izhikevich neurogenesis models as a function of age and benchmarking biologically realistic local learning rules against surrogate gradient training results in larger scale networks.

Linearly programmable halide perovskite memristors for brain-inspired computing ☆

Seung Ju Kim, University of Southern California

Neuromorphic hardware, which provides high-performance AI processing capability with low power consumption, is an attractive and challenging field designed to overcome the existing von Neumann computing systems. To implement high-performance training in neuromorphic hardware, it is essential to develop artificial synapses that exhibit linear and symmetric programmability with a bipolar operation, analog multi-states with a high dynamic range, a high yield, a long retention, a low variation, and a small footprint. To achieve these requirements, memristors, non-volatile memory devices that store data by their conductance, have been widely studied as artificial synapses. However, traditional memristors lack a reliable microscopic structure to confine ion migration during switching, resulting in commonly observed large variability (from device to device and switching cycle to cycle) and abrupt switching (instead of linear and symmetric programming). To address these issues, numerous approaches have been explored, such as modulating conductance by adding gate-terminal or optimizing programming schemes. Only limited success has been achieved so far, which, on the other hand, typically incurs substantial area, circuitry, time, and/or energy overheads. Recently, two-dimensional (2D) halide perovskites have arisen as a top candidate for artificial synapse due to their phase versatility, superior memristive properties, microstructural anisotropy in electrical and optoelectronic properties, and even excellent moisture resistance. Unfortunately, a common challenge in all memristors has also been identified in such halide perovskites, namely, asymmetric and nonlinear conductance change, which is a well-known roadblock for efficient training and accurate inference when such materials are used in neural networks.

Here, we achieve highly linear and symmetrical conductance changes (α_p : 0.002, α_d : -0.0015) in Dion-Jacobson 2D perovskites, which were unachievable previously in 2D perovskites⁵. We further build a crossbar array based on analog perovskite synapses, achieving a high (~100%) device yield, low variation (~1.85%) with synaptic weight storing capability, multilevel analog states with long retention (~104 s), and moisture stability over 7 months. We explore the potential of such devices in large-scale image inference via simulations and show an accuracy within 0.08% of the theoretical limit. The remarkable device performances are attributed to the homogenous migration of halide vacancies by eliminating gaps between inorganic layers, confirmed by first-principles calculations and experiments. Due to the Dion-Jacobson phase formed by changing large organic cations from monovalent to divalent ammonium cations ($A^{n-1}PbnX_{3n+1}$, Aⁿ is divalent ammonium cation), two hydrogen bonds are formed between organic and inorganic layers, eliminating van der Waals gaps, resulting in homogeneous interfacial ion migration through the entire region of vertically aligned layers. Our neuromorphic design rule is generally applicable to other memristive systems for achieving high-performance neuromorphic computing.

1. Kim, S. J., et al. Memristive Devices for New Computing Paradigms. *Advanced Intelligent Systems* 2000105, 2000105 (2020).
2. Kim, S. J., et al. Competing memristors for brain-inspired computing. *iScience* 24, 101889 (2021).
3. Kim, S. J., et al. Vertically aligned two-dimensional halide perovskites for reliably operable artificial synapses. *Materials Today* 52, 19–30 (2022).
4. Kim, S. J. et al. Halide Perovskites for Memristive Data Storage and Artificial Synapses. *Journal of Physical Chemistry Letters* 12, 8999–9010 (2021).
5. Kim, S. J., et al. Linearly programmable two-dimensional halide perovskite memristor arrays for neuromorphic computing. *Nature Nanotechnology* Accepted (2024).

Binding in hippocampal-entorhinal circuits enables compositionality in cognitive maps

Christopher Kymn, University of California, Berkeley

The hippocampal formation (HF), which includes hippocampus (HC) and entorhinal cortex (EC), is critical for forming memories and representing variables such as spatial position. Although it is believed that the same circuit mechanisms underwrite these capacities, it is far from clear what computational principles, wiring motifs, and cellular mechanisms are invoked. One high-level idea is that the HF constructs compositional representations of the world, in which complex memories or environments can be decomposed into their parts. Recent experimental findings have provided evidence of compositional structure in HF representations, such as novel recombinations of past experience occurring in replay [1]. In addition, compositional representations have practical advantages: they have high expressivity with lower dimensional storage requirements and can generalize to novel scenes with familiar parts. For these reasons, compositional representations have been of increasing interest in artificial intelligence [2].

We propose a normative model of the hippocampal formation that is explicitly compositional and consistent with observations from neuroanatomy and neural recordings [3]. Mechanistically, our suggestion is that binding operations, which can be mathematically formalized as compressed tensor products, are a fundamental primitive for rich compositional structure. In the model, spatial position is encoded in a residue number system, with individual residues represented by high-dimensional, complex-valued vectors. These are composed into a single vector representing position by a similarity-preserving, conjunctive vector-binding operation. Self-consistency between the representations of the overall position and of the individual residues is enforced by a modular attractor network whose modules correspond to the grid cell modules in entorhinal cortex. The vector binding operation can also associate different contexts to spatial representations, yielding a model for entorhinal cortex and hippocampus.

We show with mathematical analysis and empirical simulations that the model has strong representational efficiency. The modular attractor network achieves superlinear scaling of patterns with neural dimension, robust error correction, and a hexagonal, carry-free encoding of spatial position. These results build on existing theoretical studies of grid cells, including residue number systems developed for a linear track [4] and continuous attractor networks for a single module of grid cells (e.g., [5]). We show that these theoretical guarantees can be put to practical use in both navigation and memory tasks. Finally, we discuss the predictions and interpretations of our model that are relevant to experimental studies of grid cells and place cells.

This work does not directly contribute to neuroethics or AI ethics, but under some meta-ethical assumptions (e.g., ethical naturalism) would have implications for how we do and ought to understand ourselves.

[1] Kurth-Nelson et al. (2023), *Neuron*

[2] Greff et al. (2020), arXiv 2012.05208

[3] Behrens et al. (2018), *Neuron*

[4] Fiete et al. (2008), *J. Neuro.*

[5] Burak & Fiete (2009), *PLoS Comp. Bio*

Sequential predictive learning is a unifying theory for hippocampal representation and replay ☆

Daniel Levenstein, McGill University & Mila

The mammalian hippocampus represents an animal's position in the environment during active behavior, and generates “replay” simulations of plausible trajectories in the environment during periods of behavioral quiescence and sleep. While prominent models of the hippocampus as a continuous attractor network can produce both spatial representation and offline replay, they require specific wiring between units with pre-assigned spatial locations or learning from signals with pre-existing spatial tuning. Thus it's unclear how such a network can be learned from sensory information alone. Recently, it's been found that artificial neural networks trained to predict sensory inputs develop spatially tuned cells, aligning with predictive theories of hippocampal function. However, whether predictive learning also accounts for the ability to produce offline replay is unknown.

Here, we show that learning to predict sensory inputs can account for hippocampal representation and replay. By training recurrent neural networks to predict egocentric sensory input in a gridworld environment (Figure 1A), we show first that spatially tuned cells robustly emerge from multiple forms of predictive learning in recurrent neural networks (Figure 1B). However, we find that the presence of spatially-tuned cells does not guarantee the presence of a cognitive map with the ability to generate replay. Offline simulations only emerged in networks that used recurrent connections to predict multi-step observation sequences and received an orienting head direction signal from an upstream structure (Figure 1C), which promoted the formation of a continuous attractor reflecting the geometry of the environment (Figure 1D). These offline trajectories were able to show wake-like statistics, autonomously replay recently experienced locations, and could be directed by a virtual head direction signal. Further, we found that networks trained to make cyclical predictions of future observation sequences were able to rapidly learn a cognitive map and produced sweeping representations of future positions reminiscent of hippocampal theta sweeps (Figure 1E).

These results demonstrate how hippocampal-like representation and replay can emerge in neural networks engaged in predictive learning, and suggest that hippocampal theta sequences reflect a circuit that implements a data-efficient algorithm for sequential predictive learning. As such, sequential predictive learning is a candidate theory to unify three views of the hippocampus: 1) the hippocampus is a predictive map, 2) the hippocampus is a CANN, and 3) the hippocampus is a sequence generator. Together, this framework provides a unifying theory for hippocampal physiology, hippocampal functions and hippocampal-inspired approaches to artificial intelligence.

Attractor-based models for sequences and pattern generation in neural circuits ☆

Juliana Londono Alvarez, Brown University

Attractor neural networks, originally designed to model associative memory by storing static patterns as stable states, are useful for understanding how the brain processes information [1]. In the classical Hopfield paradigm (Figure 1A), memories are stored in the network as coexistent stable fixed points, each one accessible via distinct inputs [2]. However, while these networks handle static patterns well, more complex, dynamic behaviors, such as those in Central Pattern Generator (CPG) circuits that control rhythmic movements like walking or breathing, require dynamic attractors. Moreover, CPGs can encode multiple different, overlapping patterns, but achieving the coexistence of even static patterns in a single network is challenging [3]. With the rise of neural network theory [4, 5], its ubiquity as neural network models, and its hardware implementation it is convenient to unify locomotion models (typically modeled by coupled oscillators) with the attractor neural network framework.

To accomplish this, we use Threshold-Linear Networks to provide attractor models for three neural functions: First, we model a discrete neural integrator that can count inputs and is robust to noise, as a sequence of fixed point attractors, as shown in Figure 1B. Although similar to the classical Hopfield model in Figure 1A, it differs in that the sequence is internally encoded, with input pulses being identical and containing no information about which fixed point comes next. Second, we devise a network with attractors corresponding to five distinct quadruped gaits. These attractors coexist in the same network as distinct limit cycles in state space, as shown in Figure 1C. Despite the overlapping nodes between gaits, each one can be accessed through different initial conditions without changing the network's parameters. Lastly, we combine the approaches from panels 1B and 1C to develop a network capable of sequentially stepping through a set of dynamic attractors, as illustrated in Figure 1D. The resulting model has potential applications in robotics, particularly for tasks like robotic assembly lines, as it enables efficient reordering of elements in a sequence. Unlike "black box" AI systems, this model is theoretically grounded in the collective knowledge of brain function and mathematics. As a result, it poses no risks to privacy, mental health, or cognitive autonomy, is ethically designed, and requires no data collection, minimizing concerns of bias or discrimination. Because the ultimate aim is to automate mechanical tasks, there is concern about potential warfare exploitation—a direction I firmly do not support. My vision is for robots to assist with tasks humans find undesirable, such as housework or repetitive mechanical tasks. I seek to contribute to discussions within the research community on this issue. While this technology could displace certain jobs, it also has the potential to create supervisory roles where humans that keep humans in control of critical decisions.

References

- [1] Mikail Khona and Ila R. Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23(12):744–766, Dec 2022.
- [2] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.*, 79(8):2554–2558, 1982.
- [3] Alexander N. Pisarchik and Ulrike Feudel. Control of multistability. *Physics Reports*, 540(4):167–218, 2014. Control of multistability.
- [4] R. H. Hahnloser, H.S. Seung, and J.J. Slotine. Permitted and forbidden sets in symmetric threshold-linear networks. *Neural Comput.*, 15(3):621–638, 2003.
- [5] Carina Curto and Katherine Morrison. Graph rules for recurrent neural network dynamics. *Notices of the American Mathematical Society*, 70(04):536–551, 2023.

Investigating the Role of the two Pathways in Object Recognition and Grasping through the lens of deep neural networks

Aidasadat Mirebrahimi Tafreshi, Carnegie Mellon University

Daily interactions with objects rely heavily on visual processing. Anatomical and ablation studies, alongside research on non-human primates with cortical damage, suggest two distinct pathways in the visual cortex: the ventral pathway for object perception and identity, and the dorsal pathway for guiding actions (Mishkin, Ungerleider, & Macko, 1983; Goodale and Milner, 1992). Recent evidence focusing on tasks like object recognition for perception and grasping for action reveals significant overlap between these pathways, challenging this classical distinction (Freud, Plaut, & Behrmann, 2019; Ayzenberg, Simmons, & Behrmann, 2023). However in real-world tasks, perception and action are often inseparable in real-world tasks. For instance, grasping a hammer involves identifying it and choosing the correct end for its usage. This intertwining complicates isolating each pathway's role and may contribute to conflicting evidence in empirical studies.

Using an innovative framework based on artificial neural networks, we model the two-pathway theory of visual processing to isolate the relative contributions of each pathway and their joint activity. Our dataset includes 80 objects from 16 categories, each captured from 750 viewpoints. Each object has two types of grasp annotations: functional (based on object use) and generic (based on the center of mass). We train models to perform object recognition and grasping simultaneously, comparing single and dual-pathway architectures with varying levels of inter-pathway connectivity. We specifically explore how inter-pathway connectivity supports functional grasping, where perception and action must be integrated. Additionally, we can impose virtual lesions at different processing stages to evaluate how architectural configurations affect performance. This framework offers precise control over stimulus properties, task demands, and network architecture, bridging empirical evidence and theoretical models, to provide insights into how ventral and dorsal pathways support object recognition and action.

References

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in Neurosciences*, 6, 414–417. [https://doi.org/10.1016/0166-2236\(83\)90190-X](https://doi.org/10.1016/0166-2236(83)90190-X).

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, 15, 20–25. [https://doi.org/10.1016/0166-2236\(92\)90344-8](https://doi.org/10.1016/0166-2236(92)90344-8), PubMed: 1374953.

Freud, E., Plaut, D. C., & Behrmann, M. (2019). Protracted developmental trajectory of shape processing along the two visual pathways. *Journal of Cognitive Neuroscience*, 31, 1589–1597. https://doi.org/10.1162/jocn_a_01434, PubMed: 31180266.

Ayzenberg, V., Simmons, C., & Behrmann, M. (2023). Temporal asymmetries and interactions between dorsal and ventral visual pathways during object recognition. *Cerebral Cortex Communications*, 4, tgad003. <https://doi.org/10.1093/texcom/tgad003>, PubMed: 36726794.

Machine learning guided discovery of an intrinsic line attractor encoding aggression

Aditya Nair, Caltech

Internal affective states such as aggression and sexual drives are essential survival behaviors which share common properties such as persistence and variable intensity. The hypothalamus is a crucial hub that regulates diverse affective states and has been thought to function as a 'labeled-line' system, with populations of behavior-tuned neurons characterized by distinct transcriptional and connectomic profiles. However, neural imaging studies in hypothalamic regions such as the ventromedial hypothalamus (VMHvl), whose activation can causally trigger attack behavior [1], consistently fail to identify neurons specifically tuned to aggressive behaviors like attack [2].

To resolve this paradox between perturbation and representation in the subcortex and reveal the encoding of aggression in the hypothalamus, we apply data-driven machine learning (ML) models to approximate neural activity as a dynamical system. Analysis of the fit model uncovers an emergent computation of a line attractor in the VMHvl, where movement along the attractor correlates with increasing aggression [1]. To determine whether this line attractor is causally instantiated in the VMHvl, we conducted first-in-class closed-loop perturbation experiments in head-fixed mice, enabling model-guided activation of specific neuronal groups [2]. This revealed the capacity of VMHvl neurons to integrate along the line attractor, as well as selective recurrent functional connectivity among the ensemble contributing to the line attractor [2]. Finally, through a novel CRISPR-mediated knockout of neuropeptide receptors in VMHvl, combined with ML-enabled modeling, we provide evidence that functional connectivity is mediated by oxytocin and vasopressin receptors [3].

Together, these experiments, coupled with data-driven modeling, reveal a new motif for the computation of the persistence and intensity of an aggressive state, instantiated through intrinsic hypothalamic line attractors. Furthermore, they challenge dominant assumptions about subcortical computation and suggest that non-canonical mechanisms involving neuromodulation support emergent dynamics.

References:

1. Nair, A., Karigo, T., Yang, B., Ganguli, S., Schnitzer, M. J., Linderman, S. W., ... & Kennedy, A. (2023). An approximate line attractor in the hypothalamus encodes an aggressive state. *Cell*, 186(1), 178-193.
2. Vinograd, A., Nair, A., Kim, J., Linderman, S. W., & Anderson, D. J. (2024). Causal evidence of a line attractor encoding an affective state. *Nature*, 1-3.
3. Mountoufaris, G., Nair, A., Yang, B., Kim, D. W., Vinograd, A., Kim, S., ... & Anderson, D. J. (2024). A line attractor encoding a persistent internal state requires neuropeptide signaling. *Cell*.
4. Ságodi, Á., Martín-Sánchez, G., Sokól, P., & Park, I. M. (2024). Back to the Continuous Attractor. *ArXiv*.
5. Costacurta, J. C., Bhandarkar, S., Zoltowski, D. M., & Linderman, S. W. (2024). Structured flexibility in recurrent neural networks via neuromodulation. *bioRxiv*, 2024-07

Vision and language representations in multimodal AI models and human social brain regions during natural movie viewing

Hannah Small, Johns Hopkins University

Recent work in neuroAI suggests that representations in modern AI vision and language models are highly aligned with each other and human visual cortex. Some have found that even pure language model embeddings of image captions can predict visually-evoked activity in high-level visual areas [1, 2]. In addition, training AI vision models on language-aligned tasks (e.g., CLIP-style models) improves their match to visual cortex, particularly in regions involved in social perception[3], suggesting these brain regions may be similarly "language aligned". This prior work has primarily investigated only static stimuli without language, but in realworld settings, simultaneous visual and verbal semantic signals do not always share a commonly referenced semantic space (e.g., the subject of speech is not always visible). One area where this disconnect becomes obvious is naturalistic social processing and communication, which involves integrating converging but disparate visual and linguistic input. To understand the integration of vision and language during natural viewing, we fit an encoding model to predict voxel-wise responses to an audiovisual movie using visual representations from both purely visual and language-aligned vision transformer models and paired language transformer models (Figure 1). We first find that in naturalistic settings, there is remarkably low correlation between representations in vision and language models. Both of these model representations predict social perceptual and language region activity well. Next, we find that vision-language alignment does not improve a model's match to neural responses in visual, social perceptual, or language regions, despite social perceptual and language regions being well predicted by both vision and language embeddings. In fact, the language embeddings from the vision-language transformer perform worse than simple word-level embeddings. Our work demonstrates the importance of testing multimodal AI models in naturalistic settings and reveals differences between language alignment in modern AI models and the human brain. The limitation of current multimodal AI models in predicting brain responses to naturalistic stimuli calls for new approaches in modeling simultaneous vision and language, perhaps using recent open source high-quality video datasets [4, 5].

[1] Adrien Doerig, Tim C. Kietzmann, Emily Allen, Yihan Wu, Thomas Naselaris, Kendrick Kay, and Ian Charest. Semantic scene descriptions as an objective of human vision, September 2022. URL <http://arxiv.org/abs/2209.11737>. arXiv:2209.11737 [cs, q-bio].

[2] Colin Conwell, Jacob S. Prince, George A. Alvarez, and Talia Konkle. Language Models of Visual Cortex: Where do they work? And why do they work so well where they do? *Journal of Vision*, 23(9):5653, August 2023. ISSN 1534-7362. doi: 10.1167/jov.23.9.5653. URL <https://jov.arvojournals.org/article.aspx?articleid=2792615>.

[3] Aria Y. Wang, Kendrick Kay, Thomas Naselaris, Michael J. Tarr, and Leila Wehbe. Better models of human high-level visual cortex emerge from natural language supervision with a large and diverse dataset. *Nature Machine Intelligence*, 5(12):1415–1426, December 2023. ISSN 2522-5839. doi: 10.1038/s42256-023-00753-y. URL <https://www.nature.com/articles/s42256-023-00753-y>. Number: 12 Publisher: Nature Publishing Group.

[4] Miquel Farré, Andi Marafioti, Lewis Tunstall, Leandro Von Werra, and Thomas Wolf. Finevideo. <https://huggingface.co/datasets/HuggingFaceFV/finevideo>, 2024.

[5] Bria Long, Violet Xiang, Stefan Stojanov, Robert Z. Sparks, Zi Yin, Grace E. Keene, Alvin W. M. Tan, Steven Y. Feng, Chengxu Zhuang, Virginia A. Marchman, Daniel L. K. Yamins, and Michael C. Frank. The BabyView dataset: High-resolution egocentric videos of infants' and young children's everyday experiences, June 2024. URL <http://arxiv.org/abs/2406.10447>. arXiv:2406.10447 [cs].

Decomposing spiking neural networks with graphical neural activity threads ☆

Bradley Theilman, Sandia National Laboratories

To understand the computational capacities of the brain and how we might develop brain-inspired AI algorithms, we need powerful abstractions for neural computation. Ideally, these abstractions should be naturally adapted to the spiking and synaptic dynamics of real brains. In this poster, I will present an alternative approach to analyzing spiking neural networks that avoids many of the implicit assumptions in current approaches for spiking network analysis and offers a route to new computational abstractions. Current approaches for building computational abstractions for spiking dynamics begin by sorting spikes into time bins and constructing population activity vectors that trace the dynamics of neural activity in a high dimensional space over time. While fruitful, these approaches necessarily smear out intrinsic relations between spikes and may obscure computationally-relevant features of neural dynamics.

Our approach begins by constructing a directed acyclic graph directly from the synaptic relations between individual spikes. By definition, these synaptic relations must support the computations in the spiking network. The analysis combines spiking activity and the structure of the network into a unified mathematical object, without time bins. I will show how this directed graph naturally decomposes into weakly connected subgraphs we call Graphical Neural Activity Threads (GNATs). These GNATs are well-defined and provide a picture of information flow through a spiking network. Furthermore, GNATs are defined by the relative timings between spikes and are thus robust to spike timing variations. I will then describe an algorithm that can efficiently find isomorphic GNATs in large spiking neural datasets. By identifying isomorphic GNATs, we identify putatively isomorphic computations. I will show how GNATs arising in the dynamics of spiking network models are constructed out of other GNATs, analogous to sampling in music production. Thus, GNATs exhibit compositionality. Because of their naturalness, robustness, and compositionality, GNATs provide a powerful basis for computational abstraction in spiking neural networks.

Significant resources have been spent through the BRAIN Initiative in collecting large neural activity and connectivity datasets. The GNAT analysis is ideal for leveraging both resources because it combines activity and connectivity into a single mathematical structure. Through this abstraction, we may begin to leverage large-scale neuromorphic hardware to test hypotheses about neural computation with reduced reliance on animal models (an ethical challenge for NeuroAI development), or test hypotheses not possible in animal models. By understanding how activity and connectivity relate through GNATs, we can extract the fundamental computational principles from spiking neural networks and apply these principles to future NeuroAI architectures.

References:

Theilman, B.H., Wang, F., Rothganger, F., Aimone, J.B., Decomposing spiking neural networks with Graphical Neural Activity Threads. arxiv:2306.16684.

Quantum Materials-Enabled Associative Learning and Neuromorphic Computing ☆

Eric Wang, College Station High School

Associative learning is essential to human cognition, allowing us to connect distinct stimuli and adapt to our environment. The neurobiological foundations of associative learning involve the broad neuromodulating brainstem neurons (e.g., norepinephrine and serotonin) and brain astrocytes via volume transmission and release of neurotransmitters from the presynaptic neuron that bind to receptors on the postsynaptic neuron via wiring transmission [1,2]. This modulation includes synaptic connections, changes in neuronal properties, and altered transmission across neural networks. Volume transmission is less spatially specific, operating on slower timescales and allowing for broader neural activity modulation and coordination across brain regions, while wiring transmission enables precise communication within local areas.

This study incorporated sub-5nm quantum dots into ultrathin, a polyvinylpyrrolidone film, which was further sandwiched between silver and indium tin oxide to serve as the neuromorphic unit. This innovative design allows for forming an electric stimulus-induced Ag filament that diminishes for implementing resistive switching, effectively mimicking synapse wiring transmission. The key breakthrough, however, lies in the photo-stimulation that activates the quantum dots, initiating an additional redox reaction of migrating Ag ions to facilitate synaptic weight modulation, thereby emulating the volume transmission of the human brain. This bio-informed design led to unprecedented associative learning performance. The photo-responsive quantum materials significantly boost the associative memory due to multiple exciton generation and, thus, the learning effect due to the high quantum efficiency in converting photons to excitons to create a biomimetic association. Both individual artificial synapses and crossbar networks demonstrated rapid learning, reliable memory operations, and extended memory retention (> one day), providing a reliable foundation for further research and development. This work creatively incorporates an associative learning crossbar into an artificial neural network algorithm for hardware-accelerated machine learning and neuromorphic computing. Specifically, a dataset was created with 100 MNIST handwritten digit images for machine recognition. The dataset consists of 100-pixel digits containing original pixel-digit images and noised pixel-digit images (random 1 of 25 pixels was changed into a different value \approx 4% noise level) for training and testing. The associative learning crossbar-implemented neural network demonstrated an accuracy of >92%, much higher than the conventional neural network process (76%) under the same conditions, demonstrating power-efficient artificial intelligence with high accuracy at small training datasets, which is an inherent advantage of biological associative learning. This bio-mimic associative learning crossbar architecture demonstrates an exciting way to achieve neuromorphic intelligence.

In the study, I ensured that the emulated AI systems were designed and trained to be fair and unbiased. This included addressing issues of algorithmic bias and ensuring that the AI did not perpetuate existing inequalities.

Reference:

Carew, T.J., Walters, E.T., and Kandel, E.R. (1981). Associative Learning in *Aplysia* : Cellular Correlates Supporting a Conditioned Fear Hypothesis. *Science* (1979) 211, 501–504. [10.1126/science.7455692](https://doi.org/10.1126/science.7455692).

Yang, Q., Kuzyk, P., Antonov, I., Bostwick, C.J., Kohn, A.B., Moroz, L.L., and Hawkins, R.D. (2015). Hyperpolarization-activated, cyclic nucleotide-gated cation channels in *Aplysia* : Contribution to classical conditioning. *Proceedings of the National Academy of Sciences* 112, 16030–16035. [10.1073/pnas.1501731113](https://doi.org/10.1073/pnas.1501731113).

Dissecting the functional complexity of excitatory-inhibitory connectivity structures ☆

Qingyang Wang, Johns Hopkins University

A key aspect of intelligence is the ability to perform complex tasks. It remains elusive, on a principled level, what network structures are crucial in supporting complex functionalities. To provide insights on this front, we leverage theories from artificial intelligence (AI) and statistical learning and data of the whole-brain electron microscopy (EM) connectome. Along the process, we build AI that learns more efficiently; we also propose a functional complexity metric that is task-agnostic, learning-independent, and experiment-testable. We focus on the excitatory-inhibitory (E-I) structure (polarity). They are of importance because in biological systems, E-I connectivity is highly specified, E-I identities rarely switch post-development, and some features can be highly conserved across species, e.g. from invertebrates to vertebrates, excitatory neurons take up ~70-80%.

First, we show by experiments that highly specified and fixed E-I structure is not a biological constraint but an advantage. By adequately fixing the polarity structure a priori, deep neural networks (DNNs) learn with fewer samples, in less iterations [1]. Crucially, such efficiency requires the fixed E-I structure to be permissible to the task of interest. Then the key question is — what E-I structures are supportive of a wide range of complex tasks?

We start by mathematically proving networks without inhibitory connections (negative weights) are not universal approximators, thus they have extremely limited capacity to solve tasks in general [2]. Consequently in networks, extremely concentrated excitatory neurons lead to narrow set of functionalities. Next, we leverage the larva *Drosophila* whole brain EM connectome [3] to discuss more complicated E-I structures. Briefly, we built firing-rate models with the magnitudes given by the connectome synaptic counts and E-I identities randomly sampled. We use the proposed functional complexity metric to assess which sampled E-I configuration have higher functional complexity. In total, we sampled 8180 different E-I configurations. The optimal percentage of excitatory neurons that maximizes functional complexity is 75-81%. The optimality matches the true distribution observed via scRNA-seq [4]; it also matches the highly conserved E-I ratio across species. [5] Intriguingly, over-abundance of excitatory neurons show advantage in functional complexity only when the inhibitory neurons are biased to be highly connected; in contrast, when the E-I identities are sampled uniformly (not degree dependent), the optimal E-I ratio falls at balanced population size. All current AI models fall into the uniform scenario since they lack segregation of excitatory and inhibitory neurons — these point to an unexplored direction in building better AI.

By discussing what E-I configuration are supportive of diverse set of complex functions, we point out new directions of building more energy-efficient AIs; we also provide a normative explanation to a highly conserved biological phenomenon — why brains tend to have so many excitatory neurons — through our proposed functional complexity metric.

References and Notes

[1] Qingyang Wang, Michael Alan Powell, Eric W Bridgeford, Ali Geisa, and Joshua T Vogelstein. Polarity Is All You Need to Learn and Transfer Faster. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett, editors, Proceedings of the 40th International Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pages 36264–36284. PMLR, 11 2023. URL <https://proceedings.mlr.press/v202/wang23ae.html>.

[2] Qingyang Wang, Mike A Powell, Ali Geisa, Eric Bridgeford, Carey E Priebe, and Joshua T Vogelstein. Why do networks have inhibitory/negative connections? In Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pages 22551–22559, 10 2023.

[3] Michael Winding, Benjamin D. Pedigo, Christopher L. Barnes, Heather G. Patsolic, Youngser Park, Tom Kazimiers, Akira Fushiki, Ingrid V. Andrade, Avinash Khandelwal, Javier Valdes-Aleman, Feng Li, Nadine Randel, Elizabeth Barsotti, Ana Correia, Richard D. Fetter, Volker Hartenstein, Carey E. Priebe, Joshua T. Vogelstein, Albert Cardona, and Marta Zlatic. The connectome of an insect brain. *Science*, 379(6636), 2023. ISSN 10959203. doi: 10.1126/science.add9330.

[4] Marc Corrales, Benjamin T. Cocanougher, Andrea B. Kohn, Jason D. Wittenbach, Xi S. Long, Andrew Lemire, Albert Cardona, Robert H. Singer, Leonid L. Moroz, and Marta Zlatic. A single-cell transcriptomic atlas of complete insect nervous systems across multiple life stages. *Neural Development*, 17(1), 2022. ISSN 17498104. doi: 10.1186/s13064-022-00164-6.

[5] Qingyang Wang, Michael A. Powell, Eric Bridgeford, and Joshua T. Vogelstein. Why do we have so many excitatory neurons? In *Society for Neuroscience 2023*, Nov 2023.

CRIREL: A Hyperflexible and Reconfigurable Neural Circuit ☆

Alexander White, National Tsing-Hua University

Biological neural circuits exhibit remarkable flexibility, enabling rapid responses to dynamically changing environments [1]. Such fast response times imply that this adaptability cannot rely solely on synaptic plasticity, which operates on a much slower timescale. Instead, the adaptability suggests that neural circuits are inherently reconfigurable, allowing them to switch functionalities without synaptic modifications [2]. That is, multiple functions coexist within a neural circuit, and environmental and contextual stimuli trigger the appropriate response by dynamically shifting the circuit into the correct operational mode [1, 2].

It is suspected that flexibility in the nervous system stems from its highly recurrent nature [3, 4]. Recurrent networks operate near bifurcations, where changes in parameters trigger qualitative shifts in behavior, enabling the emergence of new dynamic states and thereby allowing rapid switches to occur depending on contextual inputs [3, 4].

To demonstrate the effectiveness of recurrent neural networks near bifurcations, we constructed a 4-neuron circuit dubbed CRIREL (Coupled Recurrent Inhibitory and Recurrent Excitatory Loop) and is capable of compressing 24 unique functions into a single circuit with a single set of fixed synaptic weights. We demonstrated that these circuits are flexible because we can control an underlying double cusp bifurcation in the network, allowing new stable or unstable states to emerge. We show that varying bias currents (baseline activity level, implemented as constant background input) and input, while keeping synaptic weights fixed, can unfold the double-cusp bifurcation and result in flexibility.

To systematically explore the circuit's functionality, we classify its output in terms of all 8 nontrivial logic gate operations (AND, OR, XOR, NAND, NOR, NXOR, IMP, NIMP) based on three input characteristics relevant to neuroscience: differences in magnitude, timing, and phase. Timing and phase introduce event-driven computing, requiring recurrent connections and extending beyond standard digital logic. As a concrete example, an XOR operation based on input timing produces an "on" signal only when two inputs arrive at slightly different times (approximately 1 ms time difference, or about a spike). If they arrive at the same time, the circuit is "off". Remarkably, we show that all 24 unique functions coexist with fixed synaptic weights, with changes in functionality driven solely by variations in bias current. Finally, we show (using the same circuit architecture) that logic can be performed downstream from the initial event-based layer and can be reconfigured using bias currents.

[1] D. N. Lyttle, J. P. Gill, K. M. Shaw, P. J. Thomas, and H. J. Chiel, "Neuromechanical bistability contributes to robust and flexible behavior in a model of motor pattern generation," *BMC Neuroscience*, vol. 16, p. P33, Dec. 2015.

[2] G. Hennequin, T. Vogels, and W. Gerstner, "Optimal Control of Transient Dynamics in Balanced Networks Supports Generation of Complex Movements," *Neuron*, vol. 82, pp. 1394–1406, June 2014.

[3] F. C. Hoppensteadt and E. M. Izhikevich, *Weakly Connected Neural Networks*. Applied Mathematical Sciences, New York: Springer-Verlag, 1997.

[4] B. Liu, A. J. White, and C.-C. Lo, "Augmenting flexibility: Mutual inhibition between inhibitory neurons expands functional diversity," *bioRxiv*, 2024.

[5] Hyper-Flexible Neural Networks: Rapidly Switching between Logic Operations in a Compact 4-Neuron Circuit Alexander James White, Belle Liu, Ming-Ju Hsieh, Kuo-An Wu, Chung-Chuan Lo *bioRxiv* 2024

Decoding brain intrinsic dynamics for NeuroAI

Xinhe Zhang, Harvard University

Recent advancements in neuroscience have enabled the stable, long-term tracking of single-cell neural activity. This marks a significant leap forward in our ability to capture the full spectrum of neural dynamics, from short-term event-driven responses to the brain's long-term intrinsic processes. With these new tools, we can collect data revealing the brain's intrinsic dynamics, such as representational drift — a phenomenon where neural patterns associated with consistent sensory inputs or motor actions change over time. Moreover, as AI increasingly draws inspiration from neural processes, capturing the full range of neural dynamics — including long-term adaptations and learning mechanisms — can provide valuable insights for creating more adaptive and efficient algorithms. Despite significant advancements in recording technology, existing analytical methods primarily focus on modeling short-term neural trajectories involved in task completion or stimulus-induced events. This focus limits our ability to grasp the long-term transformations that are fundamental to learning processes.

To address this gap, we introduce the Concurrent Hierarchical Representation of Neural Long-Short Term Dynamics (Chronos), a novel generative modeling framework designed to simultaneously infer both short-term neural dynamics within individual trials and the long-term evolution of these neural trajectories. Chronos employs a hierarchical approach to capture neural activity across multiple timescales, offering a comprehensive representation of both the transient, rapid dynamics associated with specific behaviors and the slower, gradual changes indicative of ongoing learning and adaptation.

We validated Chronos using large-scale neural recordings from mice subjected to repeated visual stimuli over extended periods. The model demonstrated robust performance in inferring neural activity patterns across different time points, effectively capturing both fast, stimulus-evoked responses and the slower representational drift observed in the data. When we compared the model's predictions to real neural recordings, Chronos exhibited a high degree of fidelity in predicting neural activity.

Additionally, we applied Chronos to decode and predict the visual stimuli presented to the mice, leveraging its dual-level representation of neural dynamics to mirror visual cortex activity. This dual-level decoding approach illustrates the potential to infer the natural progression of neural dynamics over both short- and long-term timescales.

Chronos advances our understanding of neural systems by capturing both rapid and gradual neural dynamics, with significant implications for neuroethics and AI ethics. By providing a comprehensive framework for modeling the evolution of neural representations over time, Chronos informs ethical considerations regarding interventions in neural circuits, such as brain-computer interfaces. Understanding these long-term dynamics is crucial to ensure that such technologies respect individual autonomy, consent, and privacy while preventing unintended consequences arising from neural modification. Furthermore, as AI systems increasingly draw inspiration from neural processes, Chronos offers insights that can guide the ethical development of adaptive AI systems, fostering alignment with cognitive processes and ethical standards.

References:

Pandarath, C., O'Shea, D.J., Collins, J. et al. Inferring single-trial neural population dynamics using sequential auto-encoders. *Nat Methods* 15, 805–815 (2018). <https://doi.org/10.1038/s41592-018-0109-9>.

Zhao, S., Tang, X., Tian, W. et al. Tracking neural activity from the same cells during the entire adult life of mice. *Nat Neurosci* 26, 696–710 (2023). <https://doi.org/10.1038/s41593-023-01267-x>.

Liberti, W., Markowitz, J., Perkins, L. et al. Unstable neurons underlie a stable learned behavior. *Nat Neurosci* 19, 1665–1671 (2016). <https://doi.org/10.1038/nn.4405>.

Tang, X., Shen, H., Zhao, S. et al. Flexible brain–computer interfaces. *Nat Electron* 6, 109–118 (2023).
<https://doi.org/10.1038/s41928-022-00913-9>.

Zhao, S., Shen, H., Qin, S., Jiang, S., Tang, X., Lee, M., Zhang, X., Lee, J., Chen, J. and Liu, J., Realigning representational drift in mouse visual cortex by flexible brain-machine interfaces. *bioRxiv Preprint* at <https://doi.org/10.1101/2024.05.23.595627> (2024).