# Brain Algorithms

Nancy Lynch

January 26, 2023

**Abstract**

This reading group will cover recent papers in the research area of Brain Algorithms. This area studies specific brain mechanisms for important brain tasks such as memory and recall, focus and attention, decision-making, intuitive and symbolic thinking, and prediction. This involves representing complex concepts in terms of patterns of neural firing, and using those representations to make decisions or produce other types of output. The area studies these mechanisms by modeling them formally as abstract distributed algorithms, and analyzing them using methods from analysis of algorithms.

Topics of interest this term may include recognition of hierarchically structured concepts, novelty detection, and neural assemblies. We will also study some mechanisms that are involved in interaction with the real world. Thus, we will consider representing notions such as position and motion, and using the representations to perform tasks such as orientation and navigation.

We will also consider general issues involved in modeling brain mechanisms, such as composition, abstraction, and general learning rules.

## 1 Introduction

This writeup attempts to map out the scope and contents of the research area of Brain Algorithms, putting the topics into a semi-coherent framework. It also provides ideas for the Brain Algorithms reading group in Spring, 2023.

The writeup contains references to many papers. Some are basic readings that we have already covered in our previous course and reading groups, but we may want to review them. The rest are new suggestions for the reading group.

This writeup is organized into three parts. Part I deals with specific brain mechanisms—how they work and to some extent, how they might be learned. Specific topics included so far are:

- Abstraction, memory, and recall.
- Decision-making and attention.
- Intuitive thinking.

- Symbolic thinking.
- World modeling and prediction.

Of course there are other mechanisms, such as mechanisms that process emotions. We might not get to those.

[N: Where does representation of structured concepts belong? Can be both intuitive and symbolic thinking. Perhaps also modularity.]

For each of the above, we can describe specific mechanisms, in the form of synchronous Spiking Neural Networks. In addition, we can consider how they might be learned. Here I am sprinkling the learning for the various specific mechanisms throughout, rather than trying to separate representation issues from learning issues. Also, some general learning issues appear at the end.

Part II deals with more general issues that span different kinds of mechanisms:

- Composition, modularity. This is a very important topic to study. For example, it appears that the important capability of generalization depends crucially on modularity.
- Different types of mechanisms, and different types of representations. Static vs. active mechanisms: active mechanisms have interesting timing properties, motion, etc. Discrete vs. continuous concepts/representations.
- High-level, abstract representations vs. low-level representations.
- Comparing Artificial Neural Networks (ANNs) to brain networks.
- General learning issues: Comparison of learning rules. Lifelong learning. Short-term and long-term learning.

Part III is a reading list, consisting of papers extracted fro Part I and Part II. It is not well organized in terms of topics; such an organization should emerge from our discussions. We can choose papers to present from this list. We may add others as the semester goes on.

# Part I: Specific Mechanisms

## 2 Memory and recall

Inputs to the brain are detailed representations, which can be of different modalities: vision, sound, touch, odor, taste... Inside the brain, these get converted to abstract internal representations. Such abstraction is an important part of memory/recall mechanisms.

These representations also yield outputs. An output can be some type of decision, or an output action involving, say, motion or speech.

Valiant's early work, which we have studied in our earlier reading groups, fits under this topic:

**Book:** Circuits of the mind. Valiant. 1994.

More recent work by Navlakha and his collaborators gives interesting abstract representations:

**Paper:** A neural algorithm for a fundamental computing problem. S Dasgupta, CF Stevens, S Navlakha Science 358 (6364), 793-796. 2017.
Covers dimensionality expansion, sparsification, clustering, dimensionality reduction.

**Paper:** A neural data structure for novelty detection. S Dasgupta, TC Sheehan, CF Stevens, S Navlakha. Proceedings of the National Academy of Sciences 115 (51), 13093-13098. 2018.
Maintains a data structure that can be used to classify inputs based on frequency of previous presentation.

**Paper:** Habituation as a neural algorithm for online odor discrimination. Y Shen, S Dasgupta, S Navlakha. Proceedings of the National Academy of Sciences 117 (22), 12402-12410. 2020.

**Paper:** Can a Fruit Fly Learn Word Embeddings? Y Liang, CK Ryali, B Hoover, L Grinberg, S Navlakha, MJ Zaki, D Krotov. arXiv preprint arXiv:2101.06887. 2021.

The following paper also deals with abstract representations:

**Paper:** Random sketching, clustering, and short-term memory in spiking neural networks. Hitron, Lynch, Musco, Parter. 2020.
Covers dimensionality expansion and reduction, and memorization based on setting up and maintaining firing patterns (no modification of edge weights).

[N: Find other papers that fit here. Specifically, look for papers that show how the abstract representations reflect natural structure. Clustering is discussed in the papers above. Also consider other types of structure such as hierarchy, composition, geometric relationships, timing, sequencing.]

**Paper:** Learning of hierarchical concepts. Lynch, Mallmann-Trenn. 2020.

**Paper:** Learning Hierarchically-Structured Concepts II: Overlapping Concepts, and Networks With Feedback. Nancy Lynch, Frederik Mallmann-Trenn. In progress.

## 3 Decision-making, focus, attention

The types of decisions we consider here include discrete decisions such as choosing among two, or several options, and continuous decisions such as choosing a direction of motion in 2-space or 3-space. We begin by revisiting Winner-Take-All algorithms.

**Paper:** Winner-Take-All Computation in Spiking Neural Networks. Nancy Lynch, Cameron Musco, Merav Parter. 2019.

**Paper:** Spike-based winner-take-all computation: Fundamental limits and order-optimal circuits. L Su, CJ Chang, N Lynch - Neural computation, 2019.

**Paper:** How fast is neural winner-take-all when deciding between many options? Birgit Kriener, Rishidev Chaudhuri, Ila R. Fiete
Proceedings of the National Academy of Sciences doi: 10.1073/pnas.1917551117

And we have other kinds of discrete decisions, such as Brabeeba's new work on decision-making in context:

**Paper:** Theory and model of thalamocortical processing in decision-making under uncertainty. Wang, Halassa, Lynch. 2022.

Sabrina's work on decision-making in normal and abnormal mice:

**Paper:** [N: Sabrina?]

Also include something on decisions from continuous spaces.

**Paper:** [N: Sabrina? Brabeeba?]

# 4 Intuitive thinking

In this section and the next, we deal with intuitive thinking and symbolic thinking, respectively.

Work by Papadimitriou, Vempala, and others on assemblies seems to mostly fit here. Here are several possible readings:

**Paper:** Assembly pointers for variable binding in networks of spiking neurons R Legenstein, CH Papadimitriou, S Vempala, W Maass. arXiv preprint arXiv:1611.03698 11. 2016.

**Paper:** Random projection in the brain and computation with assemblies of neurons. Papadimitriou, Vempala. 10th Innovations in Theoretical Computer Science Conference. 2019.

**Paper:** Brain computation by assemblies of neurons. CH Papadimitriou, SS Vempala, D Mitropolsky, M Collins, W Maass. Proceedings of the National Academy of Sciences 117 (25), 14464-14472. 2020.

**Paper:** Assemblies of neurons learn to classify well-separated distributions M Dabagia, SS Vempala, C Papadimitriou. Conference on Learning Theory, 3685-3717. 2022.

**Paper:** Emergent Computation and Learning from Assemblies of Neurons. S Vempala. 2022.

And here is a new paper by Vempala's student:

**Paper:** The $k$-cap process on geometric random graphs.
Reid, Vempala. https://arxiv.org/abs/2203.126802022.

# 5 Symbolic thinking

Now we move to symbolic thinking, which uses different mechanisms from intuitive thinking. For a preliminary position paper, see:

**Paper:** Symbolic Knowledge Structures and Intuitive Knowledge Structures.
Nancy Lynch.
arXiv:2206.02932, 2022.
This paper contains a preliminary description of a combination of structures that can be used for a combination of symbolic and intuitive thinking.

We could also consider learning of these structures.

We might want to include something here from the linguistics community, such as the following book. Berwick's linguistics course this semester might also provide some ideas.

**Book:** Why only us: Language and Evolution. Berwick, Chomsky. 2016.

We might also consider something from cognitive psychology. Or on systems for symbolic representation of concepts and relations, such as Shruti.

[N: Brabeeba said that the following probably relates to Keith's stuff. It looks like it's about symbolic reasoning, which is why I have put it here.]

**Paper:** Neural dynamics and geometry for transitive inference.
https://www.biorxiv.org/content/10.1101/2022.10.10.511448v1

## 6 World modeling, prediction

[N: Paper suggested by Sabrina. Find more papers that focus, not necessarily on autism, but on brain mechanisms used in prediction. I think it's interesting to understand how the brain does prediction. I am not convinced of its connections with autism or other disorders.]

**Paper:** Autism as a disorder of prediction. Sinha et al. PNAS 2014.

[N: The following is from Brabeeba He described it as being about neural learning rules for generating flexible predictions and computing the successor representation. It seems that it is about a particular prediction mechanisms (and how to learn it), rather than about some generic learning issue. So I put it here, with other work on prediction.]

**Paper:** Neural learning rules for generating flexible predictions and computing the successor representation. Dmitriy Aronov, L.F. Abbott, Emily Mackevicius doi: https://doi.org/10.1101/2022.05.18.492543

`https://www.biorxiv.org/content/10.1101/2022.05.18.492543v1.abstract`

# Part II: General Modeling Issues

## 7 Composition

Mechanisms in the brain are composable. Generalization naturally arises if the brain has already acquired some mechanisms that can serve as building blocks. What forms might this take?

We can consider basic models for composition in discrete, synchronous systems:

**Paper:** A compositional model for Spiking Neural Networks. Lynch, Musco. arXiv 2022.

**Paper:** A generalized compositional SNN model. Lynch. In progress.
Model extensions, to cover work with Frederik, Lili, Sabrina.

The following paper, recommended by Keith, shows composition of motifs from dynamical systems, which represent physical motion.

**Paper:** Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. Driscoll, Shenoy, Sussillo. 2022.

The above seem to be referring to special kinds of dynamical components. There should be other kinds. We should also consider composition of discrete components. Examples? One possibility is hierarchical concepts, which represent a kind of modularity:

**Paper:** Learning of hierarchical concepts. Lynch, Mallmann-Trenn. 2020.

**Paper:** Learning Hierarchically-Structured Concepts II: Overlapping Concepts, and Networks With Feedback. Nancy Lynch, Frederik Mallmann-Trenn. In progress.

But note that the above paper allows only limited overlap between high-level concepts. Overlap may be more important, and more widespread, in other settings that we want to consider, where the pieces are more like reusable building-blocks.

An important question here is how these composed representations are learned.

Possibly related is early work by Merav Parter on emerging modularity. Her work dealt with evolution rather than learning; it would be interesting to understand the essential differences:

**Paper:** Facilitated variation: how evolution learns from past environments to generalize to new environments. M Parter, N Kashtan, U Alon. PLoS computational biology 4 (11), e1000206. 2008.

# 8 Abstraction

Another general issue is abstract vs. lower-level representations of mechanisms. An abstract representations might, for example, have one reliable neuron for each concept being represented, whereas a lower-level representations might have multiple, less reliable, neurons. Lower-level, unreliable neurons might implement the more reliable higher-level neurons through redundancy. This type of lower level is sometimes called a "population model".

Another type of lower level representation might be more distributed, husing a ologram-like structure.

**References?**

# 9 Comparing ANNs to brain networks

The following was suggested by Sabrina:

**Paper:** Neural population geometry: An approach for understanding biological and artificial neural networks. SueYeon Chung and L.F. Abbott. ScienceDirect. 2021.

[N: This gives only one aspect of comparing ANNs to brains. Current neuroscience research uses different methods. We might explore some of these.]

# 10 Learning

So far, I have sprinkled learning issues throughout the study of the various mechanisms, but some general issues remain, not specifically tied to particular mechanisms. For instance, some research involves comparison of learning rules:

**Paper:** ODE-Inspired Analysis for the Biological Version of Oja's Rule in Solving Streaming PCA. Chi-Ning Chou, Mien Brabeeba Wang. COLT 2020: 1339-1343. 2020.

The following may also involve a study of learning rules:

**Paper:** A correspondence between normalization strategies in artificial and biological neural networks. Y Shen, J Wang, S Navlakha. Neural Computation 33 (12), 3179-3203. 2021.

The following may involve learning of representations:

**Paper:** A feedback control principle common to several biological and engineered systems. JY Suen, S Navlakha. Journal of the Royal Society Interface 19 (188), 20210711. 2022.

The following paper, suggested by Brabeeba, is on learning-to-learn / continual learning:

**Paper** Biological underpinnings for lifelong learning machines.
https://www.nature.com/articles/s42256-022-00452-0

[N: Where does modularity fit into the above? Surely part of being able to learn over time is learning pieces that can be reused.]

**Paper:** Meta-learning synaptic plasticity and memory addressing for continual familiarity detection.
https://www.cell.com/neuron/pdfExtended/S0896-6273(21)00947-8

# Part III: Tentative Reading List

I've collected papers from the discussion above, removing some and adding some. Right now, the organization isn't very coherent, first some groupings by authors, then by people who suggested the paper.

## 11 Our work

This is for background, showing our viewpoint(s).

Winner-Take-All Computation in Spiking Neural Networks. Nancy Lynch, Cameron Musco, Merav Parter. 2019.

Spike-based winner-take-all computation: Fundamental limits and order-optimal circuits. L Su, CJ Chang, N Lynch - Neural computation, 2019.

Random sketching, clustering, and short-term memory in spiking neural networks. Hitron, Lynch, Musco, Parter. 2020.
Covers dimensionality expansion and reduction, and memorization based on setting up and maintaining firing patterns (no modification of edge weights).

Learning of hierarchical concepts. Lynch, Mallmann-Trenn. 2020.

ODE-Inspired Analysis for the Biological Version of Oja's Rule in Solving Streaming PCA. Chi-Ning Chou, Mien Brabeeba Wang. COLT 2020: 1339-1343. 2020.

A compositional model for Spiking Neural Networks. Lynch, Musco. arXiv 2022.

Theory and model of thalamocortical processing in decision-making under uncertainty. Wang, Halassa, Lynch. 2022.

Symbolic Knowledge Structures and Intuitive Knowledge Structures. Lynch. ArXiv, 2022.
A preliminary description of a combination of structures that can be used for a combination of symbolic and intuitive thinking. We could also consider learning of these structures.

Learning Hierarchically-Structured Concepts II: Overlapping Concepts, and Networks With Feedback. Nancy Lynch, Frederik Mallmann-Trenn. 2023. In progress.

Lynch. In progress.
Model extensions, to cover work with Frederik, Lili, Sabrina.

## 12 Navlakha's work

A neural algorithm for a fundamental computing problem. S Dasgupta, CF Stevens, S Navlakha Science 358 (6364), 793-796. 2017.
Covers dimensionality expansion, sparsification, clustering, dimensionality reduction.

A neural data structure for novelty detection. S Dasgupta, TC Sheehan, CF Stevens, S Navlakha. Proceedings of the National Academy of Sciences 115 (51), 13093-13098. 2018.
Maintains a data structure that can be used to classify inputs based on frequency of previous presentation.

Habituation as a neural algorithm for online odor discrimination. Y Shen, S Dasgupta, S Navlakha. Proceedings of the National Academy of Sciences 117 (22), 12402-12410. 2020.

Yang Shen, Sanjoy Dasgupta, Saket Navlakha: Reply to Semelidou and Skoulakis: "Short-term" habituation has multiple distinct mechanisms. Proc. Natl. Acad. Sci. USA 117(34): 20373-20374 (2020)

Can a Fruit Fly Learn Word Embeddings? Y Liang, CK Ryali, B Hoover, L Grinberg, S Navlakha, MJ Zaki, D Krotov. arXiv preprint arXiv:2101.06887. 2021.

Yang Shen, Julia Wang, Saket Navlakha: A Correspondence Between Normalization Strategies in Artificial and Biological Neural Networks. Neural Comput. 33(12): 3179-3203 (2021)

Javier J. How, Saket Navlakha, Sreekanth H. Chalasani: Neural network features distinguish chemosensory stimuli in Caenorhabditis elegans. PLoS Comput. Biol. 17(11) (2021)

Yang Shen, Sanjoy Dasgupta, Saket Navlakha: Algorithmic insights on continual learning from fruit flies. CoRR abs/2107.07617 (2021)

A feedback control principle common to several biological and engineered systems. JY Suen, S Navlakha. Journal of the Royal Society Interface 19 (188), 20210711. 2022.

## 13 Papadimitriou, Vempala

Assembly pointers for variable binding in networks of spiking neurons.
R Legenstein, CH Papadimitriou, S Vempala, W Maass. arXiv preprint arXiv:1611.03698 11. 2016.

Random projection in the brain and computation with assemblies of neurons.
Papadimitriou, Vempala.
10th Innovations in Theoretical Computer Science Conference. 2019.

Brain computation by assemblies of neurons. CH Papadimitriou, SS Vempala, D Mitropolsky, M Collins, W Maass. Proceedings of the National Academy of Sciences 117 (25), 14464-14472. 2020.

Assemblies of neurons learn to classify well-separated distributions.
M Dabagia, SS Vempala, C Papadimitriou. Conference on Learning Theory, 3685-3717. 2022.

Emergent Computation and Learning from Assemblies of Neurons.
S Vempala. 2022.

The k-cap process on geometric random graphs.
Reid, Vempala.
https://arxiv.org/abs/2203.126802022.

# 14 Selections by students

And now we have selections by Brabeeba, Sabrina, and Keith. I've made an attempt to group by topic.

## 14.1 Brabeeba

Brabeeba has proposed some papers based on computing head directions and path integration, on mechanisms for using vector addition to calculate directions, on generating directions from goals.

He has also proposed some papers involving aspects of learning, including meta-learning, learning predictions, lifelong learning. I've added one on short-term synaptic plasticity.

### 14.1.1 Head direction, path integration

The Neuroanatomical Ultrastructure and Function of a Biological Ring Attractor.
https://www.sciencedirect.com/science/article/pii/S0896627320306139

Accurate angular integration with only a handful of neurons.
https://www.biorxiv.org/content/10.1101/2022.05.23.493052v1.abstract

From Brabeeba: The papers on the head direction are excellent because coming from the detailed connectome data, the theory of biological ring attractors can be mapped precisely onto the actual neuronal structures and their activities agree with the prediction. One can then begin to ask why the biological ring attractors are built this way. The second paper asked one such kind of question.

### 14.1.2 Vector addition

Building an allocentric travelling direction signal via vector computation.
https://www.nature.com/articles/s41586-021-04067-0

Transforming representations of movement from body-to world-centric space.
https://www.nature.com/articles/s41586-021-04191-x
This is another theory paper based on connectome and functional physiology. Based on the connectome knowledge and the electrophysiology recording, the authors convincingly demonstrate how the traveling directions are computed based on vector addition and how vector addition is implemented in actual neurons.

### 14.1.3 Control systems for navigation

Converting an allocentric goal into an egocentric steering signal.
https://www.biorxiv.org/content/10.1101/2022.11.10.516026v1
This is also a theory paper based on connectome and functional physiology. The authors convincingly demonstrate through experiments on how fruit flies implement a control system for navigation. There are lots of why types of questions remaining to be answered here.

### 14.1.4 Learning

Meta-learning synaptic plasticity and memory addressing for continual familiarity detection.
https://www.cell.com/neuron/pdfExtended/S0896-6273(21)00947-8

Neural learning rules for generating flexible predictions and computing the successor representation.
Dmitriy Aronov, L.F. Abbott, Emily Mackevicius
doi: https://doi.org/10.1101/2022.05.18.492543

`https://www.biorxiv.org/content/10.1101/2022.05.18.492543v1.abstract`

Biological underpinnings for lifelong learning machines:
https://www.nature.com/articles/s42256-022-00452-0
This paper is on Learning-to-learn / continual learning.

Continual learning in a multi-layer network of an electric fish.
https://www.sciencedirect.com/science/article/pii/S0092867419311705

A Biophysical Basis for Learning and Transmitting Sensory Predictions.

https://www.biorxiv.org/content/10.1101/2022.10.31.514538v1.full.pdf
This is another example of a great theory paper. Through electrophysiology recording on different cell types, the authors show how the interplay between complex spikes and simple spikes can serve as a mechanism to transmitting signals and teaching signals at the same time. A follow-up paper also discusses the theory of when this is possible.

Robust and brain-like working memory through short-term synaptic plasticity.
Leo Kozachkov, John Tauber, Mikael Lundqvist, Scott L. Brincat, Jean-Jacques Slotine, Earl K. Miller.
PLOS, December 27, 2022 https://doi.org/10.1371/journal.pcbi.1010776
https://journals.plos.org/ploscompbiol/article?id=10.137/journal.pcbi.1010776

#### 14.1.5 Related to Keith's ideas

Neural dynamics and geometry for transitive inference.
https://www.biorxiv.org/content/10.1101/2022.10.10.511448v1

## 14.2 Sabrina

Now we have papers suggested by Sabrina. They deal with the role of dopamine signalling, with motor control, with neural population geometry, with prediction, and with learning for pattern recognition. Related?

#### 14.2.1 Dopamine signals (not Reward-Prediction Error (RPE)

Mesolimbic dopamine release conveys causal associations.
https://www.science.org/doi/10.1126/science.abq6740
One of the long standing theories in neuroscience has been the idea of dopamine neurons signaling reward prediction error. In this recent paper, the authors provide a new theory to explain dopamine responses: causal learning. The authors show how their theory can reproduce the results of the current RPE experimental literature. Further, they design experiments to disentangle the two theories, showing that their new theory better predicts the experimental results. In summary, they have strong evidence for a different understanding of dopamine signals.

#### 14.2.2 Motor Control

Does the nervous system use equilibrium-point control to guide single and multiple joint movements?
https://www.cambridge.org/core/journals/behavioral-and-brain-sciences/article/abs/does-the-nervous-system-use-equilibriumpoint-control-to-guide-single-and-multiple-joint-movements/7516FD4A2D26B151095246920EA85527
[N: Sabrina? Better reference?]

Emilio Bizzi introduced an elegant theory for motor control that hypothesizes motor movement as transitions between equilibrium points.

#### 14.2.3 Neural population geometry and its usefulness for comparing models to brains

SueYeon Chung and L.F. Abbott. Neural population geometry: An approach for understanding biological and artificial neural networks. ScienceDirect. 2021. https://www.semanticscholar.org/paper/Neural-population-geometry

Williams, Kunz, Kornblith, Linderman. Generalized Shape Metrics on Neural Representations.
https://arxiv.org/pdf/2110.14739.pdf [N: Sabrina: Does this paper belong here?]

#### 14.2.4 Sensory perception and prediction problems in autism spectrum disorder

Jonathan Cannon, Amanda M. O'Brien, Lindsay Bungert, and Pawan Sinha. Prediction in Autism Spectrum Disorder: A Systematic Review of Empirical Evidence
https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8043993/

Sinha, Kjelgaard, Gandhi, Tsourides, Cardinaux, Partaxis, Diamond, Held. et al. Autism as a disorder of prediction. PNAS 2014.

`https://web.mit.edu/sinhalab/Papers/Magical_World_Theory2014.pdf`

And possibly some other papers from Pawan Sinha's group here at MIT and others.

#### 14.2.5 Learning for Pattern Recognition and Prediction

Adaptive Resonance Theory:
https://sites.bu.edu/steveg/files/2016/09/Encyclopedia-of-Machine-Learning-Data-Mining-ART-Carpenter-Grossberg-2016.pdf

Competitive Learning: From Interactive Activation to Adaptive Resonanc:.
https://www.sciencedirect.com/science/article/abs/pii/S0364021387800253
Stephen Grossberg has been a very influential theoretical neuroscientist. The above papers cover his theories of how neural networks can learn to tackle problems of pattern recognition and prediction. Further, Grossberg's ideas seem relevant to our group's interest in hierarchical learning.

### 14.3 Keith

Driscoll, Shenoy, Sussillo. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. 2022.

Special kinds of dynamical components. Need to study other kinds of components, for compositionality.

Generating Coherent Patterns of Activity from Chaotic Neural Networks.
https://www.sciencedirect.com/science/article/pii/S0896627309005479?via