

MACHINE LEARNING AND NEUROSCIENCE. COMBINED

We leverage neurobiological and neurophysiological principles to develop mathematical models and algorithms that push the boundaries of machine learning and artificial intelligence.

We are building mathematical models, algorithms and software that enable new forms of machine computation, learning, and intelligence, by leveraging the current state of knowledge of neurobiological processes and computation. Informed by experiments and data science.

Technical goals. Our focus is on achieving what machine learning cannot with the current state of the art: real-time, robust, and adaptive contextual learning and analytics, with minimal to no prior training.

Industry Value. The goal is to build neurobiological-derived engineered systems that can be integrated into non-biological engineered technologies of value to industry.

Neuroscience. And by thinking about neural derived machine learning algorithms, we are gaining a new perspective on understanding the biological brain.

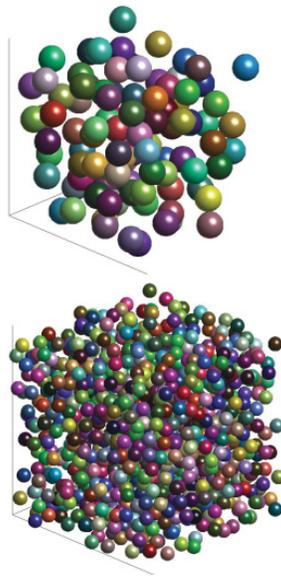


PUSHING THE BOUNDARIES INTO NEW MACHINE LEARNING

We are interested in pushing the limits and understanding of existing machine learning algorithms through our understanding of systems and computational neuroscience. Our lead research is focused on exploring fundamentally new machine learning methods. These technologies have the potential to impact many industry and application verticals.

Generating Machine 'Ideas'.

We are developing algorithms and methods intended to enable machines to 'think' creatively. This means that the system will be capable of producing original and creative machine-generated data about a specific problem *not originally present in existing associations or data patterns it may have learned*. Our goal is to enable this through the system's own internal recourses.



Our approach involves a multi-step process. In a first stage, we have developed algorithms capable of carrying out real time unsupervised learning (encoding) of input data *without prior training*. This is based on a theoretical framework that encodes learned data as dynamic states in geometric networks derived from models of neural signaling. In a second stage, manipulation of encoded internal representations in these networks associated with learned data produces machine-generated data that represent novel machine 'ideas'. This is achieved by a dynamic morphing process of the input data. These morphed structures encode new machine-generated data patterns to form a subset of the solution space that reflect new patterns (ideas) in the machine's brain not associated with the original input data. However, machine-generated data is related to (bounded by) the input data it is operating on in order to arrive at meaningful ideas. The third stage involves reducing the high dimensional dynamical network states that encode learned data into lower dimensional topological latent spaces.

These topological spaces have a number of properties we can take advantage of. We will be able to perform real time contextual adaptation by shifting the classification boundaries for input and machine-generated data in response to novel data, tasks, or demands external to the system.

The same algorithms also allow us to perform real time (or near real time) contextual mining of a knowledge base (e.g. association graph or matrix) in the absence of a specific question or query. This is a form of proactively monitoring context given a starting set of situations or state. It is inference of an actualized subset of the total set of associations (the total solution space) given a starting state. Independent of what a query might be. Our methods allow us to zero in and focus only on the actualized or realized subsets of the solution space (i.e. the possible associations) from which the right answer must be drawn. We are able to shift the context anchor before computing dynamic pathways of information flow through the network.

Our Approach

Although the biological brain exhibits an ability to learn that we want to leverage, this does not necessarily mean that we have to reverse engineer the brain to the point that we are modeling every aspect of how the biology itself implements the brain's internal algorithms. Grounded in empirical data and experiments, our goal is to understand the emergent computational properties of the brain by abstracting away as much of the biological details as needed in order to capture the core algorithms, i.e. the underlying rules.

Faculty

Gabriel A. Silva

Founding Director

Professor of Bioengineering and Neurosciences

Jacobs Faculty Endowed Scholar in Engineering

Email: gsilva@ucsd.edu

Tel. 858.822.4591

Henry D.I. Abarbanel

Associate Director

Distinguished Professor of Physics and

Scripps Institution of Oceanography

Email: habarbanel@ucsd.edu

Gert Cauwenberghs

Professor of Bioengineering and Neurobiology

Co-Director, Institute for Neural Computation

Email: gert@ucsd.edu

Jeffrey L. Elman,

Distinguished Professor of Cognitive Science

Chancellor's Associates Endowed Chair

Email: jelman@ucsd.edu

Timothy Q. Gentner

Professor of Psychology and Neurobiology

Director, Neurosciences Graduate Program

Email: tgentner@ucsd.edu

Fan Chung Graham

Distinguished Professor of Mathematics and

Computer Science and Engineering

Email: fan@ucsd.edu

Terrance J. Sejnowski

Distinguished Professor of Neurobiology and

Salk Institute for Biological Sciences

Francis Crick Chair

Co-Director, Institute for Neural Computation

Email: tsejnowski@ucsd.edu

Industry Relationships

Lon McPhail

Director for Corporate Research Partnerships

Email: lmcp@mail@eng.ucsd.edu

Adapting Network Structure to Data Resolution.

We are developing methods to modify the structure of artificial neural networks, i.e. number of nodes, number of layers and edges, and the structural connectivity, in order to accommodate input data of differing and changing resolution or quality, and to modulate the computational load on the network so that it adapts to the changing availability of computational resources. The quality and resolution of input data can vary dynamically (moment to moment) dependent on many practical factors. For example, due to the bandwidth of sensors recording real time measurements subject to environmental and contextual conditions, or varying degrees of resolution and noise as a function of pass through from other data sources. This implies that the degree or 'resolution' of the learning required by the network necessary to encode the data will vary also. We are working on a method and technologies that will allow the structure of the network to scale dynamically in real time to match the input data resolution.

Similarly, practical limitations associated with the computational resources available to the network at any given moment in time may vary. For example, machine learning on a mobile device or on-board an autonomous system versus a high bandwidth dedicated connection to the cloud. Geography, environmental conditions, and infrastructure are all considerations that can affect the availability of computational resources. We are working on methods that will allow networks to adapt to such changes in real time as the situation and context dictate.

We anticipate that these technologies will have an important impact on autonomous systems, including robotics, gaming, virtual and augmented reality, edge computing, mobile AI (smart phones and tablets), and national security (where connection reliability and risks of information interception in the field are important).

Accountable Algorithms and Explainable Artificial Intelligence

The methods and algorithms we are developing go beyond the black box of existing machine learning and artificial intelligence. Because our work is grounded in deep mathematical theory and empirical neuroscience, the algorithms we develop learn and make decisions with mechanistic and logical transparency. We understand why our algorithms behave the way they do and how they arrive at decisions. We can also back trace the causal sequence of events computed output to the input data.

Applications.

The machine learning and artificial intelligence algorithms we are developing are designed to learn, adapt, and make decisions under highly constrained conditions with limited data (which is often noisy and messy). We are interested in practical applications and use-cases which exhibit these hard constraints for which existing machine learning methods fall short. In particular, situations where a system must learn on the fly (e.g. streaming sensor data that dynamically changes), has access to limited or no training data, must adapt to new data or computational constraints, and requires limited or low power. Related applications of interest are situations where contextual demands and shifting perspectives driven by a query or problem external to the system must operate on the same learned data to arrive at different decisions autonomously. Examples include cybersecurity, genomic and other -omic data analysis, autonomous systems such as commercial drones and vehicles, robotics, contextual data analytics to find deep patterns and relationships, systems neuroscience and neurological disorders, and the design of communication and transportation networks that must accommodate specific operational constraints.

Adoption and Integration Workflow.

Our core strength is in out-of-the-box thinking and the fundamental development of unique machine learning algorithms and methods. Our typical workflow goes from mathematical descriptions to initial coding and validating in test environments such as Matlab and Python. We follow this with proof of concept engineering applications often designed to test a specific property or capability. Following these development and testing phases, adoption and eventual application to industry grade problems and customer specific uses are done in collaboration with our partners. This typically includes software and code development and optimization for integration into specific systems or products, followed by a jointly defined industry-grade test case or real world problem that if solved adds clear value to the customer's bottom line.

Machine Learning from Neuroscience.

Why the brain? And can our understanding of neuroscience contribute to next generation machine learning? We believe so. The biological brain does not use feed forward deep networks and back propagation. There are distinct advantages to designing algorithms based on our understanding of how the biological brain learns and manipulates data and information. There is a level of dynamical and structural organization to the brain that allows it to execute computational algorithms that we are only beginning to understand. The Center for Engineered Natural Intelligence exists to be at the forefront of applying this knowledge to the future of machine learning.

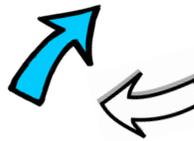
How complex *really* is the human brain computationally?

The total number of neurons in the human brain is estimated at about 85 billion. There are an additional roughly 86 billion non-neuronal cells (including astrocyte neural glial cells)

Individual neurons have been estimated to have order of magnitude 10^4 to 10^5 synapses per neuron

As a first order of magnitude approximation the total number of synapses in the human brain is roughly about 10^{16} , or ten quadrillion synapses

The 'total solution space' of the neuronal networks that make up the human brain



- Spatial summation on incoming synapses
- Temporal summation of incoming synapses
- Massive temporal offsets of PSP's
- Heterogeneous distribution of IPSP's vs EPSP's
- Neuromodulation
- Synaptic strength of PSP's
- Back propagating Ca^{2+} potentials

All to determine the next action potential!

By comparison, there are an estimated 10^{21} stars in the observable universe

How We Engage with Industry Partners

We are a Dean's Agile Center in the Jacobs School of Engineering. We work collaboratively with industry partners and investors to identify and solve problems of practical relevance to their businesses and individual interests. Industry partners join the Center and have access to faculty, students, and resources as mandated by the needs of the work. Companies also have an opportunity to embed their own engineers and scientists as part of the academic team working on the problem. And there are various opportunities and models for directed research and the use, licensing, and commercialization of the resultant work.

- We are responsive to our industry partner's needs problems, and vision
- We to build a tight feedback loop between industry and academic research towards solving far-reaching problems in engineered Intelligence

