Numenta

The Path to Machine Intelligence

White Paper

Executive Summary

The idea of machines that operate on the principles of the human brain has been around for more than fifty years. However, for most of the history of artificial intelligence, progress has been measured by how well machines solve particular problems, such as playing chess, driving cars, or passing the Turing Test. Relatively few artificial intelligence and machine learning techniques are based on an understanding of how the brain works and how it solves problems.

At Numenta we believe that the hallmark of intelligence is extreme flexibility, not the ability to solve any particular problem. The number of things humans can learn to do and the types of problems we can solve is vast. This versatility is a result of how our brains work. To build intelligent machines with equal flexibility, we need to understand the principles brains use to achieve this remarkable ability. Numenta's approach to machine intelligence is unique; we start with a deep understanding of how the neocortex learns and what makes it so flexible, and then we replicate those principles in software. Intelligent machines designed this way can be applied to a great variety of problems, large and small, problems both that humans can solve as well as those that are beyond human capabilities.

The impact of intelligent machines will rival and likely surpass the impact of computers operating under traditional principles, i.e. computers with pre-programmed rules, rather than learning systems. This endeavor will involve many people and many companies around the world. At Numenta, we are completely transparent about how our learning algorithms work, even placing our software in an open source project. There is a growing community around the world that contributes to this open source initiative and is committed to our approach to machine intelligence.

A focus on flexibility, learning from the brain, and adhering to open collaboration is the path to machine intelligence.

Machine Intelligence: A Focus on Flexibility

To be intelligent, a brain or machine must take in a stream of sensory data, automatically find patterns, adapt to changing conditions, make predictions about future events, and be able to act as required to get desired outcomes. Essentially, this automated pattern finding, learning, and behavior is what your brain does and what intelligent machines need to do.

Today's computers operate on entirely different principles. In a simple sense, we can think of them as "programmed machines" where brains are "learning machines". Since learning machines are often implemented on programmed computers, it is worth clarifying this distinction. To us, a "programmed computer" or just "computer" is one that executes a series of instructions where the programmer knows in advance what problem he or she is trying to solve and the algorithms for solving it. On the other hand, a learning machine does not know in advance exactly how to solve a problem. It has to learn from data. If a learning machine is implemented on a computer, the software is not solving the problem directly but instead implementing the learning rules and methods. A learning machine always has to be trained, where a programmed computer does not.

Programmed computers have many strengths. They can be programmed to execute any algorithm, they are fast, and they are reliable. The result is great performance for applications where the inputs and desired outcomes are known.

But programmed computers are unable to do many tasks that our brains perform easily, such as understanding language, analyzing a complex visual scene, planning, moving through a world filled with obstacles, or learning new solutions as the world changes.

Intelligent machines will accomplish tasks that humans cannot do. For example, intelligent machines can directly ingest data from non-human sensors such as GPS or radar. An intelligent machine using the same learning principles as the brain could automatically find patterns in a scanning radar data stream, make predictions, and identify anomalies. The explosion of sensors in every area of human endeavor will require automated learning systems in order to understand and make use of that data.

Throughout the evolution of programmed computers, no one could imagine which applications would be important even ten years in the future. Similarly, we expect there will be important applications for intelligent machines that we can't imagine today. This unclear future argues for flexibility as an essential component of machine intelligence. Intelligent machines designed around flexibility offer the promise of solving any problem where we have large amounts of data, the need for individualized models, and a need to understand data in a rapidly changing environment.

Finally, another important reason to have a flexible, general purpose architecture is the notion of "network effects". If each problem has a custom-built solution, the learning involved in solving that problem cannot be easily applied to the next problem. Moreover, the costs of crafting individual solutions to every problem are high, and are reliant on the availability of a small cadre of highly skilled data scientists. A universal, highly flexible approach will attract the greatest talent and resources. The accumulated value of shared applications, algorithms, utilities, tools and knowledge will enable the work to progress faster. Ultimately, this approach will yield lower cost solutions for a broader range of problems.

We believe that the right approach to creating intelligent machines is to focus on flexibility rather than applying specific solutions to specific problems. By combining a brain-like versatility with continuous learning we can design a single technology base that will be applicable to many different problems in many different domains.

The Brain as a Blueprint

If you talk to someone outside of the field of artificial intelligence or machine learning and suggest that the path to create intelligent machines is to first understand how the human brain works and then build machines that work on the same principles, they will invariably say "that makes sense."

However, this view is not held by everyone inside the fields of artificial intelligence and machine learning. A typical response you might hear is that "airplanes don't flap their wings", suggesting that it doesn't matter how brains work, or worse, that by studying the brain you will go down the wrong path -- like the people who tried to build planes with flapping wings. This analogy is mistaken. The Wright brothers understood well the difference between the aerodynamics of lift and the need for a method of propulsion. In fact, Orville Wright's motivating question was, "if birds can glide for long periods of time, then...why can't I?" Bird wings and airplane wings work on the same principles of lift. Those principles had to be understood before the Wright brothers could build an airplane. Wing flapping is a means of propulsion and there are several ways to create propulsion; the specific method used doesn't matter that much. By analogy, we need to understand the principles of intelligence before we can build intelligent machines. We might find that we deviate from the brain in some of our methods, but since the only example we have of a truly intelligent system is the brain, and since the principles of intelligence are not obvious, it is wise to first understand these principles before attempting to build intelligent machines.

The Brain's Flexibility

The neocortex is the center of intelligence in the brain. The neocortex (plus a few closely related brain regions) occupies about 75% of the volume of a human brain. The neocortex is the locus of everything we think of as intelligence. Language, music, math, vision, planned behavior, and a host of other capabilities are products of the neocortex. When you listen to someone speaking, or watch a movie, or brew a cup of coffee via thousands of coordinated movements, the neocortex is responsible. It is worth remembering that programmed computers are incredibly far from doing any of these things with the flexibility or ease of the neocortex. Despite progress in artificial intelligence and machine learning, we are not close to achieving the ability of a mouse, let alone a monkey or a human.

The most amazing and surprising capability of the neocortex is its flexibility. At birth, the neocortex has almost no knowledge. It has to learn to see, learn to hear, learn to speak, and learn how to act, and it continues to do so throughout life. The range of things the neocortex can learn is astounding.

We can learn any one of a thousand different languages or two or ten. We can learn to use any of thousands of tools. We can learn to program computers or build a bridge. We don't know the limits of what the neocortex can learn but the range of human ability is impressive. In essence, the neocortex is a universal learning machine.

This flexibility is possible because of an almost unbelievable fact: the neocortex uses one set of principles to achieve everything it does. This property, first described by Vernon Mountcastle in 1978 and well supported in neuroscience, is often referred to as the "common cortical algorithm". Human vision, hearing, and touch, in fact everything the neocortex does, are the result of essentially the same set of learning rules. The fact that the neocortex uses one uniform set of principles to solve all these problems makes it a compelling blueprint for creating a platform for machine intelligence.

Therefore the number one difference between neocortical learning algorithms and other learning algorithms is that the neocortical learning algorithms are more flexible and more universal. They can be applied to a broad range of problem types, and they continually learn and adapt as the nature of the problems change. At a minimum we know they are capable of everything humans can do, but we can be confident that they could be extended into new dimensions as well.

Principles of Neocortical Systems

Neuroscientists, using many new tools, have made tremendous advancements in understanding the physiology and connectivity of the brain. At Numenta, we use these insights to derive and articulate the underlying principles of cognition. We now understand many of these principles well enough to build practical systems, to test them, and to apply them to real-world problems. We still have a ways to go, but the scaffolding is in place and we have filled in many of the details. It is beyond the scope of this paper to describe all we have learned about the neocortex, but an introduction to a few major principles will help illustrate how Numenta's neocortical learning principles differ from other commonly used machine learning techniques.

The neocortex is primarily a memory of time-based patterns.
The neocortex learns how sensory patterns change over time. The brain relies on time-varying patterns to learn a model of the world, to infer what is happening, and to create behaviors. The memory of time-based patterns leads naturally to prediction, which is a ubiquitous property of the brain. To the brain, the world is more a video than a picture. Although not always obvious, this characteristic applies to all human senses. Most other machine learning techniques are not inherently time-based.

The neocortex learns continuously.
Most machine learning techniques require a training data set, often with labels. You train the system on the labeled data before applying it to novel data. In contrast, the brain learns continuously from unlabeled data. There is no separate training phase; the system continuously learns as each new input arrives. Continuous learning is important for flexibility and adaptability.

The language of the neocortex is sparse distributed representations.
The term "sparse distributed representation" refers to how the brain represents information about the world. At any point in time if you looked at 10,000 neurons, only 200 of them might be active (sparse), but since no individual neuron is essential, it is also a "distributed" representation. Sparse distributed representations have important and surprising properties that only recently have been understood well. With sparse distributed representations, the semantic meaning of something is part of the representation itself, not something stored in a separate data structure. Thus a simple comparison of two representations immediately reveals how they are semantically related. This characteristic is the basis of generalization in the brain. The mathematics of sparse distributed representations gives the brain a large capacity using surprisingly modest amounts of memory. We believe the use of sparse distributed representations is an essential component of intelligence.

In the neocortex, behavior and inference are integrated.
Every area of the neocortex integrates inference (or pattern recognition) with behavior. Even low level visual areas play a role in the movement of the eyes and head. This surprising observation informs us that learning to do something as simple as looking at a picture is dependent on behavior. To learn to see, we need to move.

These are just some of the core principles used in the brain, and they are essential for machine intelligence. Rather than using a behavioral definition of machine intelligence, such as the Turing test, we measure whether a machine is intelligent by noting how many essential neocortical principles it incorporates.

There are a variety of machine learning techniques in use today. Many of them do not incorporate biological constraints. For example, most deep learning implementations have no sense of time and do not learn continuously. However, an increasing number of researchers recognize that temporal structure and continuous learning should be added to these models. Thus, we anticipate a convergence towards biological principles. Our belief is that this convergence will happen much faster by starting with neocortical theory as a guide.

The Advantages of an Open Approach

Creating truly intelligent machines is a big idea, one that will drive the next era of computing. There are decades of work ahead to advance the learning algorithms, to build applications, and to create new hardware to support machine intelligence. The path to machine intelligence will parallel the path that was taken in creating programmable computers over fifty years ago. At the dawn of the computer era, nobody was quite sure what would be the important and valuable applications or how best to implement the ideas laid down by Turing and von Neumann. It took many individuals and teams working over many years to design and implement the key elements of modern computing. Nobody "owned" the computer back then and nobody owns it today. Although there are always proprietary components, it is clear that the computer's impact on the world was enhanced by open discourse and published standards, whether through operating systems, CPU architectures, or data formats.

We feel the development of machine intelligence will benefit from openness as well. Although Numenta is a for-profit business, we have explicitly designed our business architecture to enable broad outside participation across technical, scientific and business communities. The elements of our approach are detailed below.

Complete Documentation
Many of the neocortical principles described in this paper were initially introduced in the book *On Intelligence* and translated into 15 languages. Since the publication of the book, we have advanced the theory substantially, and have published white papers and videos documenting our progress. Our goal is to always document all aspects of the theory in multiple formats. We encourage others to independently implement the learning algorithms and we work to make it easy for them to do so.

Transparent Research
We share our work as it happens, not after a delay. We discuss our research online via email discussion groups and live video sessions. We post our experiments and code in open repositories so anyone can look at them. There is much to learn and uncover in cortical theory and learning algorithms, and we believe the path to discovery is through openness.

NuPIC Open Source Project
All the learning algorithms and cortical models we have developed are implemented and tested in software. We make this software available in the NuPIC (Numenta Platform for Intelligent Computing) open source project under the GPLv3 license. We monitor and contribute to the forums daily. The NuPIC community is growing and is actively moving the work forward.

Intellectual Property
We have been issued many patents associated with our work. In order to enable others to freely use this intellectual property, we chose to license our open source code under the GPLv3 license which includes a "patent peace" provision. In addition to our open source project, we have made a clear statement that we will not assert our patents against any non-commercial purpose, making access for scientists and researchers totally open.

The task of building intelligent machines will not be solved by one individual or even one company, but rather by the concerted efforts of contributors all over the world.

Conclusion

Machine intelligence is the ability for a machine to sense the world around it, to learn the patterns in the world, to learn continuously, to predict, and to act to achieve goals. Unlike programmed computers, intelligent machines will be flexible and able to solve fast-changing and complex problems without being programmed to do so explicitly. The brain is the best example we have of an intelligent system, and it provides a roadmap for building intelligent machines. The brain's center of intelligence, the neocortex, controls a wide range of functions using a common set of principles. We have made significant progress in discovering these principles and modeling them in software. We believe that the path to machine intelligence will draw on the talents of collaborators around the world. We have laid a solid foundation upon which intelligent machines will be built.

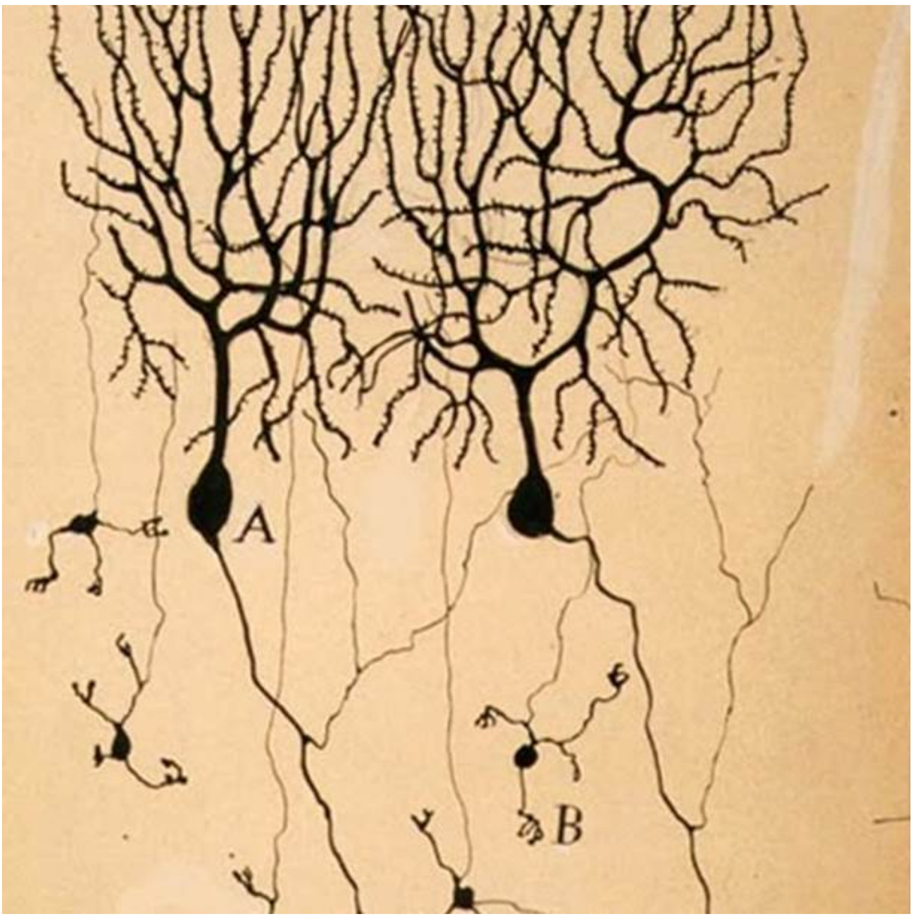

About Numenta

Numenta was founded in 2005 to be a leader in the emerging field of machine intelligence. Numenta builds solutions that automatically and intelligently act on data. Its biologically derived machine learning technology is based on a theory of the neocortex first described in co-founder Jeff Hawkins' book, *On Intelligence*. This technology is ideal for the analysis of continuously streaming data sets and excels at modeling and predicting patterns in all types of data. In addition, Numenta has created NuPIC (Numenta Platform for Intelligent Computing) as an open source project. Numenta is based in Redwood City, California.

Acknowledgements
Our work relies extensively on ideas and research from pioneering thinkers in both machine learning and neuroscience. While the cortical theory and cortical learning algorithms we have developed are unique and compelling, we have learned much from others, both from their successes and their setbacks. We have not attempted to credit every influencer in this paper, but we would like to acknowledge that many others have made our work possible.