

Jeff Hawkins - On Intelligence - 2004

Tom Rochette <tom.rochette@coreteks.org>

November 2, 2024 — [36c8eb68](#)

0.1 Context

0.2 Learned in this study

0.3 Things to explore

- Is there any research/proof that cortical regions learn and store sequences of pattern? How does that work at the neuron level?
- It is claimed that neuron? cells remain active for a relative long duration of a sequence, what is considered “active”? Is it the firing of the neuron on its axon?
- Is there any practical/theoretical limit to the size of an intelligent memory system?

1 Overview

1.1 One hundred-step rule exercise

Hawkins suggests that the brain is able to recognize a cat in a picture in less than half a second. He names this the one hundred-step rule, meaning processing an image to determine the presence or absence of a cat can go through neural network of 100 neurons deep (each neuron takes about 5ms to fire and reset, thus a chain of 100 neurons can be fired in 500ms).

Here we do a little bit of thinking about the network itself. If it was a single input neuron connected to all of the neurons in the brain, then in a single step, all neurons would have been reached. However this design is highly unlikely as it would require every single neuron to be connected to that input neuron.

Then we think of the brain as a tree structure. With 100 levels, it is very likely for the brain to be fully covered. However, if the signal goes from the root to the leaves, then the tree structure implies that there would be no consolidation of information. Neurons would be fired, but at no point would a neuron say “yeap, that’s a cat alright”. It would be more like “I saw a paw”, “I saw some ears”, “I saw a tail”, but nothing aggregating all of these indices. If, on the other hand we let the signal to flow from leaves to the root, and also consider that the signal may terminate at any internal node and act as “conscious” output, then this architecture becomes a lot more interesting.

1.2 AGI model

- System that integrates and differentiate inputs
- The differences are what trigger/attract attention

1.3 Some critique

- The books mentions “experiments show that”, yet no sources are given

2 Notes

2.1 Preface

- The biggest mistake is the belief that intelligence is defined by intelligent behavior
- It is the ability to make predictions about the future that is the crux of intelligence
- A good theory should be easy to comprehend, not obscured in jargon or convoluted argument

2.2 1 Artificial Intelligence

- Computers and brains are built on completely different principles
 - One is programmed, one is self-learning
 - One has to be perfect to work at all, one is naturally flexible and tolerant of failures
 - One has a central processor, one has no centralized control
- The AI approach was born with the digital computer
- A key figure in the early AI movement was the English mathematician Alan Turing
- A thesis of this book is that understanding cannot be measured by external behavior, it is instead an internal metric of how the brain remembers things and uses its memories to make predictions
- When you don't know how to proceed, often the best strategy is to make no changes until your options become clear

3 2 Neural Networks

- Linguists talked of intelligence in terms such as “syntax” and “semantics”. To them, the brain and intelligence was all about language
- Vision scientists referred to 2D, 2.5D and 3D sketches. To them, the brain and intelligence was all about visual pattern recognition
- Computer scientists talked of schemas and frames, new terms they made up to represent knowledge
- Anatomists and neurophysiologists wrote extensively about the structure of the brain and how neurons behave, but they mostly avoided any attempt at large-scale theory
- A neural network is unlike a computer in that it has no CPU and doesn't store information in a centralized memory. The network's knowledge and memories are distributed throughout its connectivity - just like real brains (the connectivities/weight values is the distributed memory, which, unless “stored”/“persisted”, means that the brain would be void of any memory)
- Hawkins 3 essential things to understanding the brain:
 - The inclusion of time in brain function: Real brains process rapidly changing streams of information. There is nothing static about the flow of information into and out of the brain.
 - The importance of feedback: Feedback dominates most connections throughout the neocortex.
 - Any theory or model of the brain should account for the physical architecture of the brain: Any neural network that didn't acknowledge . . . the repeating hierarchical structure of the neocortex . . . was certainly not going to work like a brain.
- The most fundamental problem with most neural networks is a trait they share with AI programs. Both are fatally burdened by their focus on behavior. Whether they are calling these behaviors “answers,” “patterns,” or “outputs,” both AI and neural networks assume intelligence lies in the behavior that a program or a neural network produces after processing a given input. The most important attribute of a computer program or a neural network is whether it gives the correct or desired output. As inspired by Alan Turing, intelligence equals behavior
- Intelligence is not just a matter of acting or behaving intelligently. Behavior is a manifestation of intelligence, but not the central characteristic or primary definition of being intelligent
- You can be intelligent just lying in the dark, thinking and understanding
- Auto-associative memories are built out of simple “neurons” that connected to each other and fired when they reached a certain threshold. They are interconnected differently than regular neural networks, using lots of feedback. Instead of only passing information forward, as in a back propagation network, auto-associative memories fed the output of each neuron back into the input. This feedback loop led

to some interesting features. When a pattern of activity was imposed on the artificial neurons, they formed a memory of this pattern

- The most important property is that you don't have to have the entire pattern you want to retrieve in order to retrieve it
- An auto-associative memory can be designed to store sequences of patterns, or temporal patterns
- Hawkins proposes the brain uses circuits similar to auto-associative memory
- We don't have machine that think, we have machines that do
- Intelligence is something that is happening in your head. Behavior is an optional ingredient
- According to functionalism, being intelligent or having a mind is purely a property of organization and has nothing to do with what you're organized out of
- By this principle, an artificial system that used the same functional architecture as an intelligent, living brain should be likewise intelligent - and not just contrivedly so, but actually, truly intelligent
- Some philosophers of mind have taken a shine to the metaphor of the "cognitive wheel," that is, an AI solution to some problem that although entirely different from how the brain does it is just as good
- To succeed (have intelligent machines), we will need to crib heavily from nature's engine of intelligence, the neocortex. We have to extract intelligence from within the brain

3.1 3 The Human Brain

- Being human and being intelligent are separate matters. An intelligent machine need not have sexual urges, hunger, a pulse, muscles, emotions, or a humanlike body
- The neocortex is about 2 millimeters thick and has six layers
- Stretched flat, the human neocortical sheet is roughly the size of a large dinner napkin
- Human are smarter because our cortex, relative to body size, covers a larger area, not because our layers are thicker or contain some special class of "smart" cells
- Some anatomists have estimated that the typical human neocortex contains around thirty billion neurons (30,000,000,000), but no one would be surprised if the figure was significantly higher or lower
- Each of these regions (of the cortex) is semi-independent and seems to be specialized for certain aspects of perception or thought. Physically they are arranged in an irregular patchwork quilt, which varies only a little from person to person. Rarely are the functions cleanly delineated. Functionally they are arranged in a branching hierarchy
- The lowest of the functional regions, the primary sensory areas, are where sensory information first arrives in the cortex
- V1 (the primary visual area) is concerned with low-level visual features such as tiny edge-segments, small-scale components of motion, binocular disparity (for stereo-vision), and basic color and contrast information
- Eventually, sensory information passes into "association areas," which is the name sometimes used for the regions of the cortex that receive inputs from more than one sense
- It is thanks to association regions that you are able to be aware that the sight of a fly crawling up your arm and the tickling sensation you feel there share the same cause
- Pyramidal neurons comprises 8 out of every 10 neuron cells of the neocortex
- Mountcastle suggests that since (all the brain) regions all look the same, perhaps they are actually performing the same basic operation. He proposed that the cortex uses the same computational tool to accomplish everything it does
- Mountcastle argues that the reason one region of cortex looks slightly different from another is because of what it is connected to, and not because its basic function is different. He concludes that there is a common function, a common algorithm, that is performed by all cortical regions
- The human brain has an incredible capacity to learn and adapt to thousands of environments that didn't exist until very recently. This argues for an extremely flexible system, not one with a thousand solutions for a thousand problems
- Apparently no area of cortex is content to represent nothing. The visual cortex, not receiving information from the eyes like it is "supposed" to, casts around for other input patterns to sift through - in this case, from other cortical regions
- Visual information from the outside world is sent to your brain via a million fibers in your optic nerve.

After a brief transit through the thalamus, they arrive at the primary visual cortex

- Sounds are carried via the thirty thousand fibers of your auditory nerve. They pass through some older parts of your brain and then arrive at your primary auditory cortex
- Your spinal cord carries information about touch and internal sensations to your brain via another million fibers. They are received by your primary somatosensory cortex
- In vision, the sense organ is your retina. An image enters your pupil, gets inverted by your lens, hits your retina, and creates a spatial pattern. This pattern gets relayed to your brain
- About three times every second, your eyes make a sudden movement called a saccade. They fixate on one point, and then suddenly jump to another point. Every time your eyes move, the image on your retina changes. This means that the patterns carried into your brain are also changing completely with each saccade
- Sounds are converted into action potentials through a coiled-up organ in each ear called the cochlea. High-frequency tones cause vibrations in the cochlea's stiff base. Low-frequency tones cause vibrations in the cochlea's floppier and wider outer portion. Mid-frequency tones vibrate intermediate segments. Each site on the cochlea is studded with neurons that fire as they are shaken
- Vision is more like three senses - motion, color and luminance (black-and-white contrast)
- Touch has pressure, temperature, pain, and vibration
- We also have an entire system of sensors that tell us about our joint angles and bodily position. It is called the proprioceptive system
- We also have the vestibular system in the inner ear, which gives us our sense of balance
- Your cortex doesn't really know or sense the world directly. The only thing the cortex knows is the pattern streaming in on the input axons

3.2 4 Memory

- A neuron collects inputs from its synapses, and combines these inputs together to decide when to output a spike to other neurons. A typical neuron can do this and reset itself in about five milliseconds, or around two hundred times per second
- The one hundred-step rule: In half a second, the information entering your brain can only traverse a chain one hundred neurons long. That is, the brain "computes" solutions to problems like this in one hundred steps or fewer, regardless of how many total neurons might be involved (comparison computer instructions with neurons "instructions", which is a faulty comparison)
- How can a brain perform difficult tasks in one hundred steps that the largest parallel computer imaginable can't solve in a million or a billion steps? The answer is the brain doesn't "compute" the answers to problems; it retrieves the answers from memory
- The entire cortex is a memory system. It isn't a computer at all
- There are four attributes of neocortical memory that are fundamentally different from computer memory:
 - The neocortex stores sequences of patterns
 - The neocortex recalls patterns auto-associatively
 - The neocortex stores patterns in an invariant form
 - The neocortex stores patterns in a hierarchy
- You can only recall memory the same way you learned it (in the same sequence)
- Texture can only be perceived through pattern sequences across the pressure-and vibration-sensing neurons in your skin (differences?)
- An auto-associative memory system is one that can recall complete patterns when given only partial or distorted inputs
- How can we perceive something as being the same or constant when the input patterns representing it are novel and changing?
- We use the term invariant representation to refer to the brain's internal representation (the same neuron cells are triggered while in presence of the same high level concept (e.g. a face))
- Memories are stored in a form that captures the essence of relationships, not the details of the moment
- When you see, feel, or hear something, the cortex takes the detailed, highly specific input and converts it to an invariant form. It is the invariant form that is stored in memory, and it is the invariant form of each new input pattern that it gets compared to

- The way you understand the world is by finding invariant structure in the constantly changing stream of input

3.3 5 A New Framework of Intelligence

- To notice that something is different, some neurons in my brain that weren't active before would have to become active. [...] When I look around the room, my brain is using memories to form predictions about what it expects to experience before I experience it. [...] But when some visual pattern comes in that I had not memorized in that context, a prediction is violated. And my attention is drawn to the error
- Prediction is not just one of the things your brain does. It is the primary function of the neocortex, and the foundation of intelligence
- The important point is that higher intelligence is not a different kind of process from perceptual intelligence
- Intelligence is measured by the capacity to remember and predict patterns in the world, including language, mathematics, physical properties of objects, and social situations
- Behavior came first, then intelligence
- Here then is the core of my argument on how to understand the neocortex, and why memory and prediction are the keys to unlocking the mystery of intelligence. We start with the reptilian brain with no cortex. Evolution discovers that if it tacks on a memory system (the neocortex) to the sensory path of the primitive brain, the animal gains an ability to predict the future. [...] By comparing the actual sensory input with recalled memory, the animal not only understands where it is but can see into the future.
- Intelligence and understanding started as a memory system that fed predictions into the sensory stream
- To know something means that you can make predictions about it
- The cortex has evolved in two directions. First it got larger and more sophisticated in the types of memories it could store. Second, it started interacting with the motor system of the old brain
- Prediction, not behavior, is the proof of intelligence

3.4 6 How the Cortex Works

- Using the “top-down” approach, you start with the image of what the solved puzzle should look like, and use this to decide which pieces to ignore and which pieces to search for
- The other approach is “bottom-up”, where you focus on the individual pieces themselves. You study them for unusual features and look for close matches with other puzzle pieces
- If you don't have a picture of the puzzle's solution, the bottom-up method is sometimes the only way to proceed
- To make predictions, your cortex needs a way to memorize and store knowledge about sequences of events
- To make predictions of novel events, the cortex must form invariant representations

3.4.1 Invariant Representations

- Each V1 neuron has a so-called receptive field that is highly specific to a minute part of your total field of vision
- Each V1 cell is also tuned for specific kinds of input patterns. For instance, a particular cell might fire vigorously when it sees a line or an edge slanted at thirty degrees within its receptive field
- In the course of spanning four cortical stages from retina to IT, cells have changed from being rapidly changing, spatially specific, tiny-feature recognition cells, to being constantly firing, spatially nonspecific, object recognition cells
- There are many, if not more feedback connections in visual cortex as there are feedforward connections
- With hearing and touch you can't recognize an object with a momentary input. Therefore the neural activity corresponding to the mental perception of objects, such as spoken words, must last longer in time than the individual input patterns

3.4.2 Integrating the Senses

- Information flows up and down sensory hierarchies to form predictions and create a unified sensory experience
- We interpret downward-flowing patterns as predictions. In the motor cortex we interpret them as motor commands

3.4.3 A New View of V1

- If all cortical regions perform the same function, why should IT be special?
- Experiments clearly show that nonadjacent patches of V1 are not directly connected; the left side of V1 can't know directly what the right side is seeing
- Subregions or clusters of V1 are physically disconnected but do the same thing
- Experiments show that all higher regions of cortex receive converging inputs from two or more sensory regions below themselves
- The job of any cortical region is to find out how its inputs are related, to memorize the sequence of correlations between them, and to use this memory to predict how the inputs will behave in the future
- Invariant representations are formed in every cortical region

3.4.4 A Model of the World

- One of the most important concepts in this book is that the cortex's hierarchical structure stores a model of the hierarchical structure of the real world. The real world's nested structure is mirrored by the nested structure of your cortex
- You can only experience a subset of the world at any moment in time
- A sequence is a set of patterns that generally accompany each other but not always in a fixed order. What is important is that patterns of a sequence follow one another in time even if not in a fixed order
- The fact that certain input patterns repeat time and again is what lets a cortical region know that those experiences are caused by a real object in the world
- "Whenever I see any of these events, I will refer to them by a common name. It is this group name, not the individual patterns, that I will pass on to higher regions of the cortex" - Cortical region

3.4.5 Sequences of Sequences

- By collapsing predictable sequences into "named objects" at each region in our hierarchy, we achieve more and more stability the higher we go. This creates invariant representations (named objects do not need to have a "name", but to be uniquely identified though an integer for instance)
- A hierarchy of nested sequences allows the sharing and reuse of lower-level objects
- The way you memorize sequences and represent them by name as information goes up and down your cortical hierarchy may remind you of the hierarchy of military command
- Recognizing any sequence would be impossible if you have not first classifying each item
- Classification and sequence formation are both necessary to create invariant representations, and each region of the cortex does them
- Context of known sequences can be used to resolve ambiguity
- An invariant representation in any region of the cortex can be turned into a detailed prediction of how it will appear on your senses by propagating the pattern down the hierarchy

3.4.6 What a Region of Cortex Looks Like

- (top) Layer 1 is the most distinct of the six layers. It has very few cells, consisting primarily of a mat of axons running parallel to the cortical surface
- Layer 2 and 3 look similar. They contain many tightly packed pyramidal cells
- Layer 4 has a type of star-shaped cell
- Layer 5 has regular pyramidal cells as well as a class of extra-big pyramid-shaped cells
- (bottom) Layer 6 also has several types of unique neurons
- Information flows horizontally in layer 1 and vertically in layers 2 through 6

- On close inspect, we see that at least 90 percent of the synapses on cells within each column come from places outside the column itself
- Upward flow
 - Converging inputs from lower regions always arrive at layer 4 - the main input layer. In passing, they also form a connection in layer 6
 - Layer 4 cells then send projections up to cells in layers 2 and 3 within their column
 - When a column projects information up, many layer 2 and layer 3 cells send axons to the input layer of that next higher region
- Downward flow
 - Layer 6 cells are the downward-projecting output cells from a cortical column and project to layer 1 in the regions hierarchically below
 - In layer 1, the axon spreads over long distances in the lower cortical region
 - Information flowing down the hierarchy from one column has the potential to activate many columns in the regions below it
 - Cells in layers 2, 3, and 5 have dendrites in layer 1, so these cells can be excited by the feedback running all across layer 1
 - The axons coming from layers 2 and 3 cells form synapses in layer 5 as they leave the cortex and are believed to excite cells in layers 5 and 6
- The thalamus is essential to normal living; a damaged thalamus leads to a persistent vegetative state

3.4.7 How a Region of Cortex Works: The Details

- How does a region of cortex classify its inputs?
- How does it learn sequences of patterns?
- How does it form a constant pattern or “name” for a sequence?
- How does it make specific predictions?
- Let’s assume that the columns in a region of cortex are like buckets
- Each column represent the label of a bucket
- The layer 4 cells in every column receive input fibers from several regions below it and will fire if they have the right combination of inputs
- When a layer 4 cell fires, it is “voting” that the input fits its label
- How does our cortical region store the sequence of classified patterns?
- Before learning, the column can only become active if driven by a layer 4 cell
- After learning, the column can become partially active via memory
- Layer 2 cells are used to “name” patterns, while Layer 3b cells represent unexpected patterns
- As a column learns, its layer 3b cell becomes quiet while its layer 2 cell is consistently active
- Layer 3a cell inhibits layer 3b cell from firing if it observes the expected input
- The method of combining partial prediction with partial input resolves ambiguous input, it fills in missing pieces of information, and it decides between alternative views

3.4.8 Flowing Up and Flowing Down

- Unexpected patterns are automatically passed to the next higher cortical region
- If recognition does not occur, an unexpected pattern will keep propagating up the cortical hierarchy until some higher region can interpret it as part of its normal sequence of events

3.4.9 Can Feedback Really Do That?

- As a general rule, information flowing up the cortical hierarchy is transferred via synapses close to cell bodies. Information going up the hierarchy is therefore more certain to pass from region to region
- Feedback flowing down the cortical hierarchy does so via synapses far from the cell body
- Massive feedback and massive numbers of synapses exist for a reason. Using this insight, we can say that a typical neuron has the ability to learn hundreds of precise coincidences on feedback fibers as they make synapses on thin dendrites

3.4.10 How the Cortex Learns

- When two neurons fire at the same time, the synapses between them get strengthened
- The two basic components of learning are forming the classifications of patterns and building sequences
- The basics of forming sequences is to group patterns together that are part of the same object
- During the early years of your life, your memories of the world first form in higher regions of cortex, but as you learn they are re-formed in lower and lower parts of the cortical hierarchy

3.4.11 The Hippocampus: on Top of it All

- With a very broad brush, we can say that the basal ganglia were the primitive motor system, the cerebellum learned precise timing relationships of events, and the hippocampus stored memories of specific events and places
- If you lose both halves of the hippocampus, you lose the ability to form most new memories
- The connection between the hippocampus and the neocortex suggest that the hippocampus is the top region of the neocortex, not a separate structure
- It is the unexplained and unanticipated remainders, the new stuff, that enter the hippocampus and are stored there
- This information won't be stored forever. Either it will be transferred down into the cortex below or it will eventually be lost
- We could say the more you know, the less you remember (the more you have experienced in the past, the less likely you are to remember a particular event)

3.4.12 An Alternate Path Up the Hierarchy

- Hawkins speculates that the alternative pathway through the thalamus is the mechanism by which we attend to details that normally we wouldn't notice
- Biologists have shown that the alternate pathway can be turned on in one of two ways
 - One is by a signal from the higher region of the cortex itself
 - The second method is a large, unexpected signal from below

3.4.13 Closing Thoughts

- There is a lesson here about the neocortex. It isn't made of superfast components and the rules under which it operates are not that complex. However, it does have a hierarchical structure that contains billions of neurons and trillions of synapses

3.5 7 Consciousness and Creativity

3.5.1 Are Animals Intelligent?

- Everything I have written so far about the neocortex and how it works depends on a very basic premise - that the world has structure and is therefore predictable
- All behavior is a means of exploiting the structure of the world for the benefit of reproduction

3.5.2 What's Different About Human Intelligence?

- Our neocortex is larger
- We have language
- Intelligence could be traced over three epochs, each using memory and prediction
 - The first would be when species used DNA as the medium for memory
 - The second epoch began when nature invented modifiable nervous systems that could quickly form memories
 - The third and final epoch is unique to humans. It begins with the invention of language and the expansion of our large neocortex

3.5.3 What Is Creativity?

- Creativity can be defined simply as making predictions by analogy, something that occurs everywhere in the cortex and something you do continually while awake

3.5.4 Are Some People More Creative Than Others?

- An expert is someone who through practice and repeated exposure can recognize patterns that are more subtle than can be recognized by a nonexpert
- Ultimately there is a physical limit to what we can learn constrained by the size of our cortex

3.5.5 Can You Train Yourself to Be More Creative?

- First, you need to assume up front that there is an answer to what you are trying to solve
- Second, you need to let your mind wander
- If you are stuck on a problem, the memory-prediction model suggests that you should find different ways to look at it to increase the likelihood of seeing an analogy with a past experience

3.5.6 What is Consciousness?

- Hawkins believes consciousness is simply what it feels to have a neocortex
- Consciousness can be broken down into two major categories:
 - One is similar to self-awareness - the everyday notion of being conscious
 - * This meaning of consciousness is not absolute. It can be changed after the fact by memory erasure
 - The second is qualia - the idea that feelings associated with sensation are somehow independent of sensory input
- If the cortex is the same everywhere, if it works with the same processes, if it is just dealing with patterns, if no sound or light enters the brain, just patterns, then why does vision seem so different from hearing?
 - Although hearing, touch and vision work under similar principles in the neocortex, they are handled differently below the cortex
 - The structure of inputs - differences in the patterns themselves - dictates how you experience qualitative aspects of the information
 - * The optic nerve has a million fibers and carries quite a bit of spatial information
 - * The auditory nerve has only thirty thousand fibers and carries more temporal information
- The brain is in a quiet and dark box
- It knows about the world only via patterns on the sensory nerve fibers
- The cortex has no ability to model the brain itself because there are no senses in the brain

3.6 8 The Future of Intelligence

3.6.1 Can We Build Intelligent Machines?

- Here is the recipe for building intelligent machines.
 - Start with a set of senses to extract patterns from the world
 - Attach to these senses a hierarchical memory system that works on the same principles as the cortex
 - Train the memory system much as we teach children
 - Once our intelligent machine has created a model of its world, it can then see analogies to past experiences, make predictions of future events, propose solutions to new problems, and make this knowledge available to us
- What makes an intelligent machine intelligent is that it can understand and interact with its world via a hierarchical memory model and can think about its world in a way analogous to how you and I think about our world

- Intelligence is measured by the predictive ability of a hierarchical memory, not by humanlike behavior
- The largest technical challenge we will face when building intelligent machines is creating the memory
- To build intelligent machines, we will need to construct large memory systems that are hierarchically organized and that work like the cortex
- We will confront challenges with capacity and connectivity
- Let's say the cortex has 32 trillion synapses
- If we represented each synapse using only two bits (giving us four possible values per synapse) and each byte has eight bits (so one byte could represent four synapses), then we would need roughly 8 trillion bytes of memory
- A hard drive on a personal computer today has 100 billion bytes (~100 GB), so we would need about eighty of today's hard drives to have the same amount of memory as a human cortex
- An individual cell in the cortex may connect to five or ten thousand other cells. This kind of massively parallel wiring is difficult or impossible to implement using traditional silicon manufacturing techniques
- Electrical wires send signals much more quickly than the axons of a neuron. A single wire on a chip can be shared, and therefore used for many different connections, whereas in the brain each axon belongs to just one neuron

3.6.2 Sensory Systems

- Our inability to tackle the issue (of protein folding) may be related, primarily, to a mismatch between the human senses and the physical phenomena we want to understand

3.7 Testable Predictions

- Prediction 1: We should find cells in all areas of cortex, including primary sensory cortex, that show enhanced activity in anticipation of a sensory event, as opposed to in reaction to a sensory event
- Prediction 2: The more spatially specific a prediction can be, the closer to primary sensory cortex we should find cells that become active in anticipation of an event
- Prediction 3: Cells that exhibit enhanced activity in anticipation of sensory input should be preferentially located in cortical layer 2, 3, and 6 and the prediction should stop moving down the hierarchy in layers 2 and 3.
- Prediction 4: One class of cells in layers 2 and 3 should preferentially receive input from layer 6 cells in higher cortical regions
- Prediction 5: A set of "name" cells described in prediction 4 should remain active during learned sequences
- Prediction 6: Another class of cells in layers 2 or 3 (different from the name cells referred to in predictions 4 and 5) should be active in response to an unanticipated input, but should be inactive in response to an anticipated input
- Prediction 7: Related to prediction 6, unanticipated events should propagate up the hierarchy. The more novel the event the higher the unanticipated input should flow. Completely novel events should reach the hippocampus
- Prediction 8: Sudden understanding should result in a precise cascading of predictive activity that flows down the cortical hierarchy
- Prediction 9: The memory-prediction framework requires that pyramidal neurons can detect precise coincidences of synaptic input on thin dendrites
- Prediction 10: Representation move down the hierarchy with training
- Prediction 11: Invariant representations should be found in all cortical areas

4 See also

5 References

- <https://www.amazon.com/Intelligence-Jeff-Hawkins/dp/0805078533>