Neural Networks 41 (2013) 3–14

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2013 Special Issue

Cognitive memory

Bernard Widrow*, Juan Carlos Aragon

ISL, Department of Electrical Engineering, Stanford University, CA, United States

ARTICLE INFO

Keywords:
Autoassociative neural networks
Cognitive memory
Pattern recognition
Human memory

ABSTRACT

Regarding the workings of the human mind, memory and pattern recognition seem to be intertwined. You generally do not have one without the other. Taking inspiration from life experience, a new form of computer memory has been devised. Certain conjectures about human memory are keys to the central idea. The design of a practical and useful "cognitive" memory system is contemplated, a memory system that may also serve as a model for many aspects of human memory. The new memory does not function like a computer memory where specific data is stored in specific numbered registers and retrieval is done by reading the contents of the specified memory register, or done by matching key words as with a document search. Incoming sensory data would be stored at the next available empty memory location, and indeed could be stored redundantly at several empty locations. The stored sensory data would neither have key words nor would it be located in known or specified memory locations. Sensory inputs concerning a single object or subject are stored together as patterns in a single "file folder" or "memory folder". When the contents of the folder are retrieved, sights, sounds, tactile feel, smell, etc., are obtained all at the same time. Retrieval would be initiated by a query or a prompt signal from a current set of sensory inputs or patterns. A search through the memory would be made to locate stored data that correlates with or relates to the prompt input. The search would be done by a retrieval system whose first stage makes use of autoassociative artificial neural networks and whose second stage relies on exhaustive search. Applications of cognitive memory systems have been made to visual aircraft identification, aircraft navigation, and human facial recognition.

Concerning human memory, reasons are given why it is unlikely that long-term memory is stored in the synapses of the brain's neural networks. Reasons are given suggesting that long-term memory is stored in DNA or RNA. Neural networks are an important component of the human memory system, and their purpose is for information retrieval, not for information storage. The brain's neural networks are analog devices, subject to drift and unplanned change. Only with constant training is reliable action possible. Good training time is during sleep and while awake and making use of one's memory.

A cognitive memory is a learning system. Learning involves storage of patterns or data in a cognitive memory. The learning process for cognitive memory is unsupervised, i.e. autonomous.

1. Introduction

During the summer of 1956 a seminar was held at Dartmouth College on the subject of artificial intelligence. This was the first AI meeting, and it was open to all. The founders of the field were there. It was an exciting time, with great anticipation in the air. What was floating around was the idea of developing an artificial brain by means of software, hardware, or a combination of both. Today, after more than fifty years, the goal is still elusive. The problem is too big, and how the brain works is for the most part still largely unknown.

In this paper, we are not trying to explain how the brain works. We are only trying to understand how human memory works, and how to build a human-like memory for computers. Human memory with all its complexity is not yet understood. This paper approaches the problem and is able to yield insight into some of memory's processes. An artificial memory can now be built taking advantage of what can be learned from nature. This kind of memory connected to a conventional digital computer facilitates solutions to problems in pattern recognition, control systems, and learning systems.

Fundamentally, human and animal pattern recognition involves matching an unknown incoming pattern with a pattern seen before and currently stored in memory. This does not fit all of the existing pattern recognition paradigms and not everyone will accept this, but this is what we believe. When a visitor appears and an interesting subject is being discussed, the "mental tape recorder" is recording the sights, sounds, etc. of the visit. In a half hour of discussion, perhaps 100,000 images of the visitor's face are recorded. These images are retinal views that capture

* Corresponding author.
E-mail address: widrow@stanford.edu (B. Widrow).

0893-6080/$ – see front matter
doi:10.1016/j.neunet.2013.01.016

the visitor's face with different translations, rotations, scale, light levels, perspectives, etc. These images are stored permanently in memory, in sequential locations wherever there is an empty place. The subject of the conversation is also recorded.

A digital computer has numbered memory registers. Data storage locations are program controlled, and data is retrieved when needed by calling for it by register number. Human memory has no numbered registers. Data is stored in human memory wherever there is an empty place and once stored, the memory has no idea where the data has been stored.

We contemplate a memory of enormous, almost unbelievable capacity, enough to hold many lifetimes of stored visual, auditory, tactile, olfactory, vestibular, etc. patterns of interest. Data and patterns are retrieved in response to an input query, an input pattern, whether visual, auditory, tactile, etc. or a combination of these. The input pattern serves as a prompt to initiate retrieval of related data patterns, if they are stored in the memory. If data patterns are retrieved and if they contain an identification, the input pattern is thereby identified. It is surprising that many aspects of human mental activity can be explained by such a simple idea of memory. Some of these aspects will be described below.

On the engineering side, we will introduce new approaches to computer memory and pattern recognition. Pattern matching is complicated by the fact that unknown incoming prompt patterns may be different from stored patterns, different in perspective, translation, rotation, scale, etc. Simple pattern matching may not be adequate, but this will be addressed.

The memory capacity must be enormous, and it should be implemented with a parallel architecture so that search time would be independent of memory size. There are many ways to structure a content-addressable memory. The "cognitive memory" proposed herein is content addressable and is of a unique design that could be physically built to give a computer a "human-like" memory, and furthermore, it is intended to serve as a behavioral model for some aspects of human and animal memory.

2. A design for a cognitive memory

The various components of the proposed cognitive memory system are represented by the diagrams of Figs. 1–4. Fig. 1 shows storing input patterns from various sensors such as eyes, ears, etc. These patterns are stored in memory folders. Autoassociative neural networks are used for memory retrieval. Fig. 2 illustrates the training method for these neural networks. Fig. 3 shows how the trained neural networks are used for retrieval when the memory system is presented with a prompt input. The memory system is composed of segments. Each segment operates independently, and retrieval is done in parallel over the segments. Fig. 4 shows how final searches are done in the segments, matching folder contents with prompt patterns. A more detailed description of the operations of the cognitive memory follows.

Returning to Fig. 1, input patterns are acquired from sensors and are then processed by cortexes. Visual patterns go through the visual cortex, auditory patterns go through the auditory cortex, tactile patterns through the somatosensory cortex, etc. In the visual cortex, for example, visual patterns are simultaneously rotated over many small angle increments, translated over many left–right up–down positions, scaled for many sizes, brightened for many brightness levels, etc. From a single snapshot, with all the combinations, hundreds to thousands of images can be generated by the visual cortex. How this can be physically done with neural networks is shown in a paper by Widrow and Winter (1988). All these images are transported by the memory input line to all the segments. They will be recorded in one of the folders of one of the segments, wherever there is an empty folder. A folder stores all of the sensory input patterns taken over a finite time interval, the time duration of an event, for example a visit to an office, etc. The storage is digital and the patterns could be stored redundantly in more than one folder. Recording time does not increase if the number of segments is increased. (Note that in a living brain, "folders", and "segments" may not exist. They are metaphors for functions that are needed.)

In Fig. 2, each memory segment is equipped with a unique neural network, a multi-layer perceptron indicated by "NN" in the diagram. Each neural network is trained independently to be "autoassociative". All the patterns stored in all the folders of the individual memory segment are used to train the neural network of that segment. A multiplexer in each segment scans through its stored patterns and presents these patterns in sequence to the input of the neural network. Training the network to autoassociate is accomplished by using the same input pattern as both the input pattern and the corresponding desired response pattern. Training can be done with the backpropagation algorithm (Werbos, 1974).

The backpropagation algorithm is a supervised training process that requires desired responses to be supplied along with the input patterns. The desired responses are usually supplied by a "supervisor", but in the case of autoassociative training as described above, the desired responses are the same as the input patterns. Thus, there is no need for a supervisor. The input patterns come from input sensors, and they are all that is required for training. Therefore the training of the neural networks is unsupervised, i.e. autonomous. (Note that in a living brain, the "backpropagation algorithm" may not exist, but is a metaphor for an essential training function.) Unsupervised learning is a good thing. Otherwise, the intelligence of a supervisor would be required. Who would be the supervisor?

Once an individual neural network is trained, it can be sensed by applying input prompt patterns and observing the corresponding output patterns. Refer to Fig. 3. If one of the prompt patterns, when presented to the neural network, causes an output pattern that closely matches the prompt pattern, there is a "hit", the difference between the input and output of the neural network was less than a predetermined threshold. The hit causes a switch to be closed and the hit pattern is then stored in a buffer memory.

If a pattern is presented as a prompt input and this is not one of the patterns trained into the neural network, i.e. not one of the patterns stored in the folders of the segment, there will be no hit. The output pattern will not match the input pattern and the difference between input and output patterns will be large. The neural network can be trained with thousands of patterns, and with high reliability will indicate very quickly when sensed if the presented pattern matches one of the patterns stored in the folders of the segment or not. How this property will be used will be explained below.

Thus far, folders in the various segments store input patterns or input data, and the neural networks have been trained to autoassociate. Fig. 3 illustrates how the neural networks are used in the memory retrieval process. Prompting patterns are fed in parallel to all of the segments. Prompting patterns could be thought of as input queries seeking contents of folders which may contain patterns that match the prompt. The prompting patterns originate with the input sensors: eyes, ears, etc. The sensed patterns are processed by the corresponding cortexes. For example, the visual cortex takes an input snapshot and creates hundreds of visual images from it by combinations of rotation, translation, scaling, brightness variations, etc. All of these prompt patterns are fed in sequence over the prompt input line. If any of the neural networks in any of the segments find an input–output match, there will be "hits". We call these primary hits. When a primary hit occurs, the switch closes and the hit prompt pattern is grabbed and held temporarily by the buffer memory.

In Fig. 4, the contents of a buffer memory that may be containing a hit prompt pattern is compared by exhaustive search with all

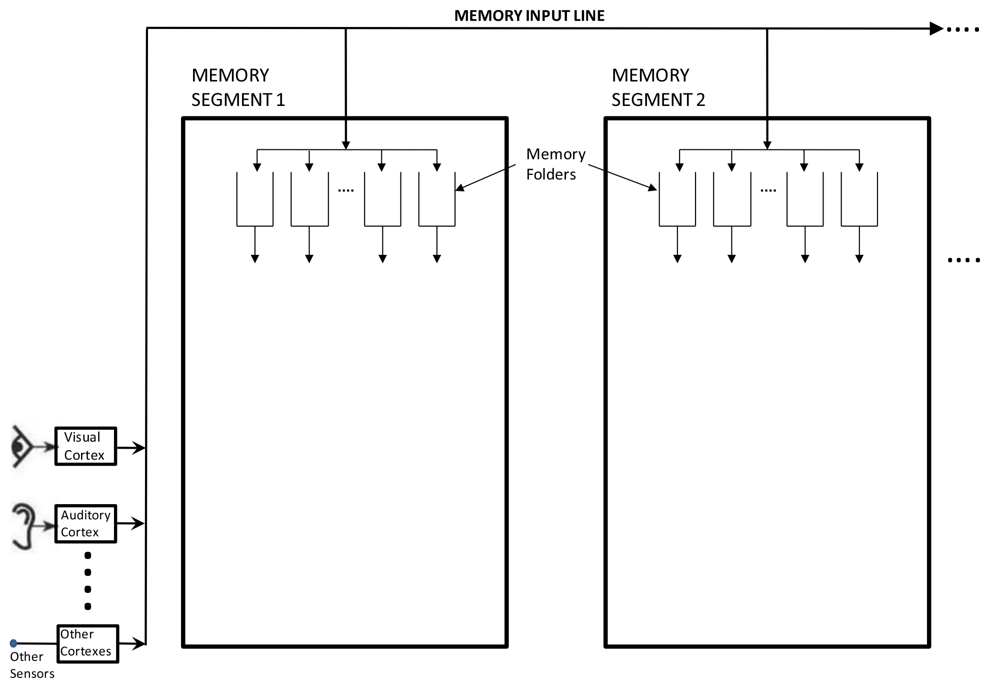

Fig. 1. Storing patterns and data into memory folders.

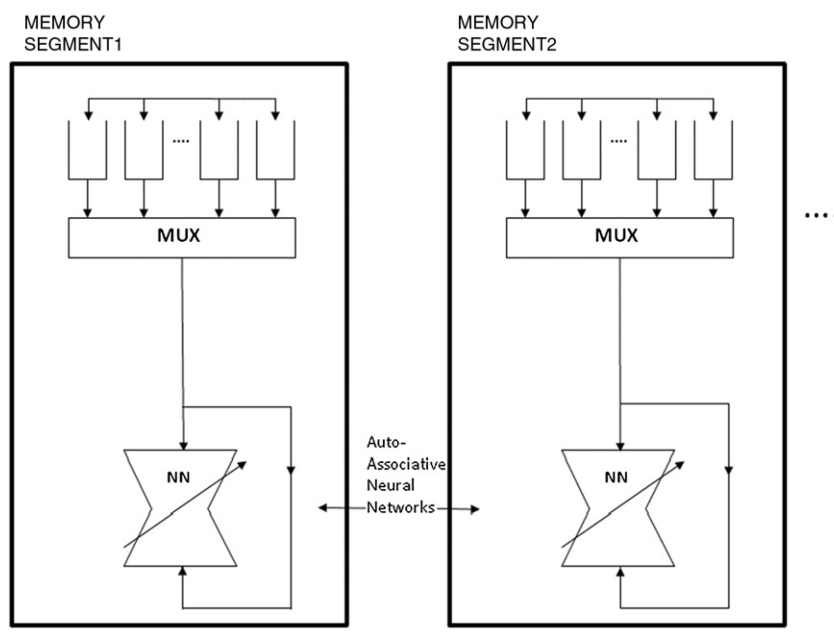

Fig. 2. Training autoassociative neural networks.

the patterns in all the folders of the corresponding segment. Since the buffer pattern is a hit pattern, a match will be found. The contents of the folders are scanned by the multiplexer. When a match is found, there is another "hit". We call this a secondary hit. When this takes place, the entire contents of the match folder are delivered as the memory output in response to the prompt. The prompt could be an audio pattern, a visual pattern, a tactile pattern, an olfactory pattern, etc., or a combination. The memory output consists of everything in the selected match folder, i.e., visual patterns, auditory patterns, tactile patterns, etc.

If one were to construct a cognitive memory system as above described, the storage of sensory input patterns in folders could be done with conventional digital memory devices. The actions of the cortexes could be implemented in software by standard computational techniques. If the neural networks were implemented in software by standard computational techniques, a massive amount of computation would be required. We are now implementing the neural networks with parallel hardware making use of graphical processing units (GPU's). Speed up of the whole system by a very large factor has been achieved. This is still a work in progress.

3. Pattern recognition by cognitive memory

One application of the cognitive memory is in the field of pattern recognition. We first record a database of a very large number of patterns and their variations (rotation, translations, etc.), plus their identifications, if known, and the autoassociative

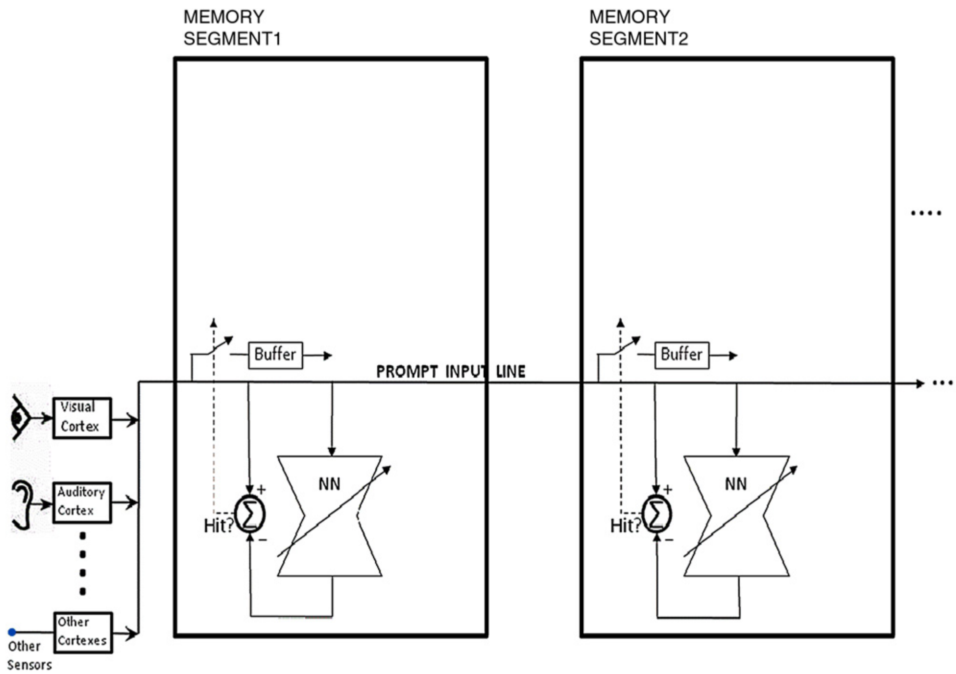

Fig. 3. Prompting the memory.

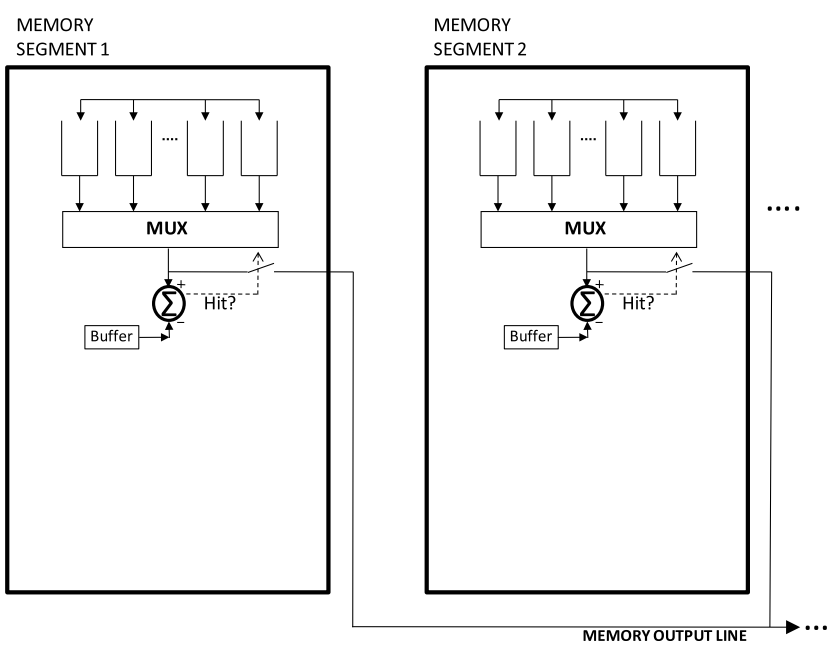

Fig. 4. Final search.

neural networks are trained with the data patterns. The training is done off line, when the cognitive memory is not performing data retrieval.

When we are given an unknown object to be recognized, its image serves as a prompt pattern. If a connection exists between the prompt pattern and one of the patterns in the database, a match will be found. This could have been done in each segment in parallel by comparing the input prompt pattern and all of its variations with every pattern in all of the folders of the segment. However, the number of combinations to be explored in each of the segments would be excessive. Instead, the input prompt pattern and all its variations is fed in sequence to all the neural networks in all the segments, and if there is a primary hit in one or more segments, we know this very quickly (each neural network is a parallel processor). We know that in each segment having a primary hit, one of the folders contains a hit pattern. Finding this folder (getting a secondary hit) requires that only one version of the prompt pattern, that stored in the segment's memory buffer, needs to be compared to the patterns in the segment's folders. This greatly reduces the number of combinations to be explored. Thus the search process is substantially sped up by using neural networks in the data retrieval process.

A practical application for the cognitive memory system is that of aircraft navigation. Once an aerial photo of an area of the earth's surface has been obtained, an airplane flying over this area could use the photo for navigation. A telescope with an attached video camera mounted on the bottom of the fuselage of the airplane could provide the necessary ground data. Navigation

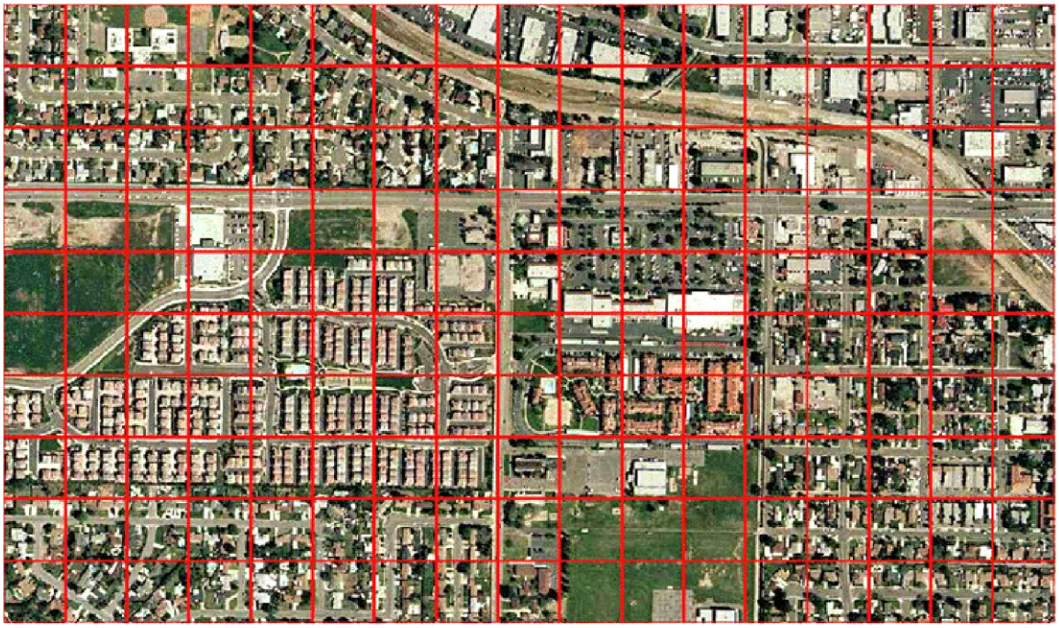

Fig. 5. 170 square images of Simi Valley, CA.

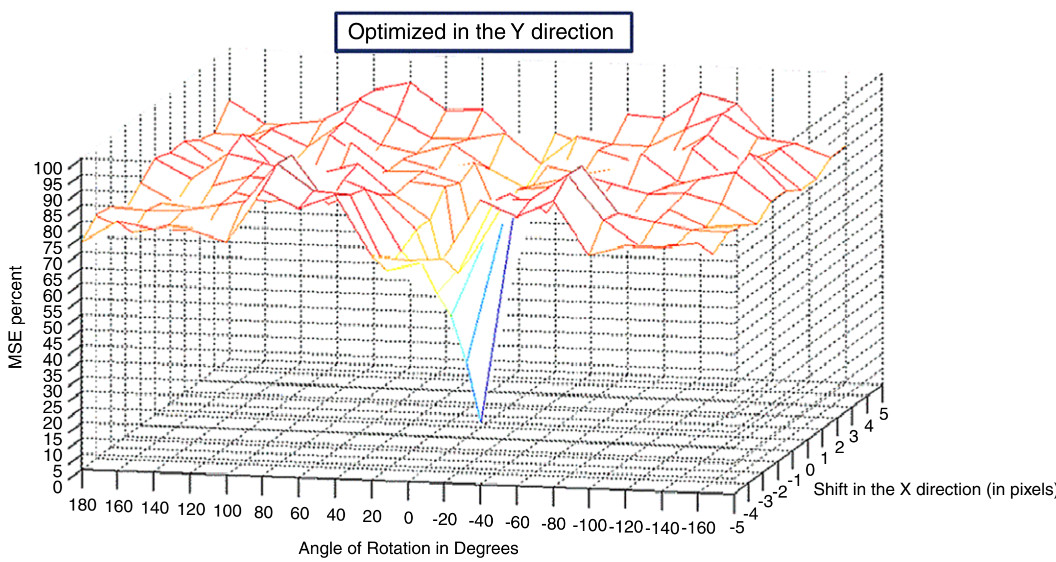

Fig. 6. Mean square error vs. rotation angle and translation.

is accomplished by correlating the telescopic images with the original aerial photo. Location and heading of the airplane can be found in this way by means of a cognitive memory system.

Fig. 5 shows an aerial photo of a portion of Simi Valley, near Los Angeles, California. The photo was downloaded from a website. The photo was divided into 170 squares. Each square image, downsampled to 15 pixels $\times$ 15 pixels, was stored in an individual folder along with the X–Y coordinates of the center of the square. The vertical Y-axis is north–south and the horizontal X-axis is east–west. An autoassociative neural network was trained with the 170 15 $\times$ 15 pixel square patterns. An airplane was assumed to be flying at the same altitude as that of the camera that took the original photo. The circular field of view of the airplane's telescope was large enough to encompass a number of the square images.

A square window within the telescope's field of view having the same size as the square images was positioned by rotation and translation to obtain an image of the ground, an image that was used as a prompt. Many prompt images were obtained by rotation and translation seeking a hit. Upon obtaining a hit, the hit folder was identified along with its X–Y coordinates. Taking into account the amount of rotation and translation that was needed to make the hit, the exact position and heading of the airplane was able to be determined. If the airplane were not flying at the same altitude as that of the original camera, then scaling would be necessary in addition to translation and rotation in order to make a hit. Fig. 6 shows a typical plot of mean square error (the difference between the input and the output patterns) versus both rotation angle and translation along the X-axis that was done to make a hit. The big dip in mean square error indicates the hit. For this plot, translation along the Y-axis was already optimized.

A simulated airplane was assumed to be flown over Simi Valley on an exactly circular track, and telescopic shots of the ground were made at uniform time intervals. The red arrows shown in Fig. 7 indicate the positions and headings of the airplane at the time of the shots of the ground. Each arrow was derived independently by the cognitive memory. They all came out equally spaced and tangent to the circular track, as they should have.

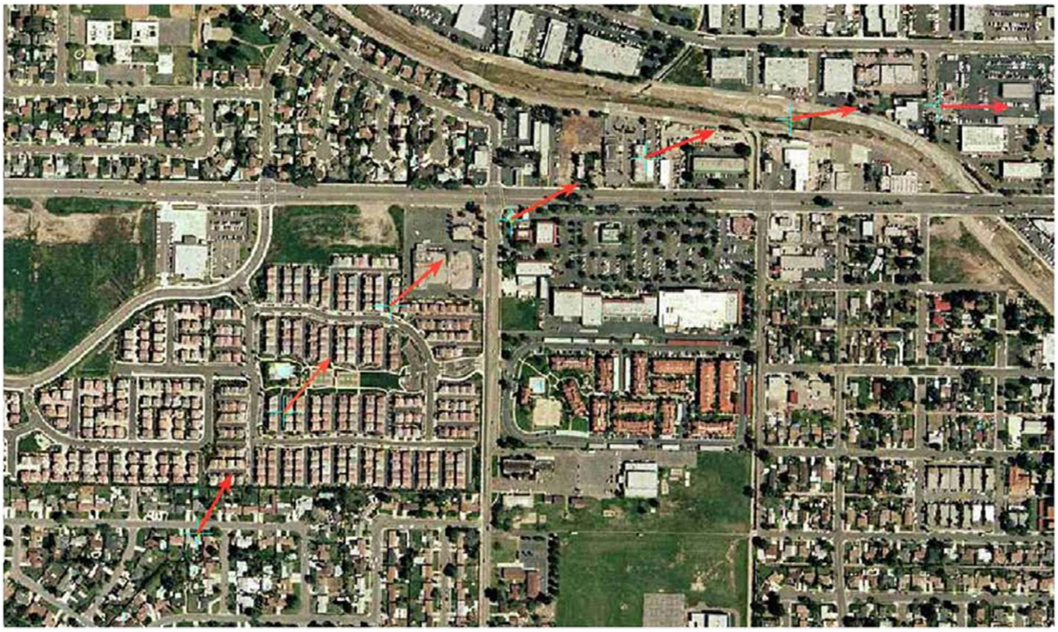

Fig. 7. Circular airplane track over Simi Valley, CA. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 8. Aircraft parked in an area near the main runway.

Another application for the cognitive memory system is that of aircraft identification and location from satellite photos. Fig. 8 is a photo taken from a satellite at about 200 miles in altitude over the British island Diego Garcia located in the Indian Ocean. There is a US air force base on this island. The aircraft in the photo are military planes of the US Air Force. The photo was downloaded from the Internet, and the website indicated that there were three types of aircraft present, B52's, KC135's, and B2's. A square window was placed over one of the B52's. The window had 25 pixels $\times$ 25 pixels and was large enough to more than cover the airplane. The original image of the B52 was translated left–right, up–down within the window making many digital images. All these images were placed in a single folder, labeled B52. The same procedure was followed for one of the KC135's, whose images were placed in another folder labeled KC135. There were only two B2 aircraft present. One of these was windowed, translated left–right, up–down, and the many images were stored in another folder labeled B2. All of the aircraft images were stored in folders of a single memory segment and were trained into a single autoassociative neural network. There were a total of 125 training patterns from the three aircraft. The neural network had three adaptive layers, with 250 neurons in the first layer with 625 synapses per neuron, 420 neurons in the second layer with 250 synapses per neuron, and 625 neurons in the third layer with 420 synapses per neuron.

Once the autoassociative neural network was trained, it was then possible to search for and identify aircraft. The search window was square, 25 pixels $\times$ 25 pixels. This window was used for scanning over the total satellite image.

For each scanned position, a variety of rotations were tried. Each resulting pattern served as a prompt. In scanning, a low mean square error of the autoassociative neural network indicated that an airplane was present in the scanning window. Once the airplane was detected, a more careful analysis was initiated that involved small translations of the window, left–right, up–down, and small rotations seeking the minimum mean square error, corresponding to a primary hit. The hit pattern was then compared

Fig. 9. This is a hit, the object is KC135.

with all the patterns in the folders looking for a secondary hit, and the identification of the folder having a pattern matching the hit pattern was taken as the identity of the unknown airplane, and the center of the optimized window was recorded as the position of the now identified airplane. Fig. 9 shows one case where an aircraft was detected, located, and identified.

4. Face recognition

Yet another application for the cognitive memory system is that of face detection and recognition. If a person's face approximately fills a square window having 20 × 20 pixels, displaying the face at this low resolution with only 400 pixels allows one to determine that the image is that of a person's face, but it is almost impossible to determine from this image who the person is. A higher resolution image of the face is necessary to do recognition. A 50 × 50 pixel image with 2500 pixels will allow recognition.

The first step is face detection, being able to detect people's faces in a scene containing a variety of background objects. This was done with a 400 pixel low-resolution system using a three-layer autoassociative neural network. A photographic image of a person's face was used for training. It did not matter who this person was, man or woman. A 20 × 20 pixel window was applied to the photograph, and the window was translated left–right, up–down with one-pixel increments, for a total of nine different positions. For each position, the window was rotated with 2 degree increments over seven different angles, and for each position and rotation, the image brightness was varied over five intensity levels. The square windows covered the face between the eyebrows and the upper lip. The total number of training patterns in the database was 315. These patterns were stored in a single folder, labeled human face. The first layer of the neural network had 400 neurons with 400 synapses or weights each. The second layer had 300 neurons with 400 synapses each. The third layer had 400 neurons with 300 synapses each. Once the network was trained, it was then able to be used for face detection.

Square windows were scanned over a test photograph to detect the faces, if they were present. Six different window sizes were applied, each scanned over the test image. Each of the windows was quantized to 20 × 20 pixels, and each 20 × 20 image was fed as an input to the trained low-resolution neural network. The mean square error was determined for each image, and when the error was below a pre-set threshold, there was a primary hit and a face was detected.

Once a face was detected, the system switched to high-resolution for face recognition. The high-resolution system had 50 × 50 pixels within its square window. The 50 × 50 images were fed to a separate high-resolution autoassociative neural network.

The high-resolution neural network had three layers. The first layer had 1800 neurons with 2500 synapses per neuron. The second layer had 1500 neurons with 1800 synapses per neuron. The third layer had 2500 neurons with 1500 synapses per neuron. How this neural network was trained can be explained by means of a simple example.

Fig. 10 shows three photos of Bernard Widrow that were used for training. A window was manually placed so that its top edge was above the eyebrows, and its bottom edge was between the nose and the lips. This was done independently for the three photos. The three windows were slightly increased in size and slightly shrunk in size generating two more window sizes per photo. For each photo, the three window sizes were rotated with 2 degree increments over seven angles, and translated left–right, up–down with 1 pixel increments over 25 positions. The total number of 50 × 50 training patterns was 1575. These patterns were recorded in a single folder that was labeled Dr. Widrow. Once the high-resolution network was trained, it was then possible to do recognition. Note that it was not necessary that the window be square. It could have been rectangular, circular, elliptical, etc., but a square window was actually used for this experiment.

Fig. 11 shows a photo of Juan Carlos Aragon, Victor Eliashberg, and Bernard Widrow that was used for test recognition. The first step was scanning with the low resolution 20 × 20 system, looking for the faces. The scanning window was moved from left–right, top–bottom, like an old style TV. The first face found was that of Juan Carlos Aragon. When the hit was made, the center of

Fig. 10. Three photos of Bernard Widrow used for training.

Fig. 11. A photo of Juan Carlos Aragon Victor Eliashberg, and Bernard Widrow use for sensing.

the corresponding 20×20 window was noted. The system was switched to high-resolution, searching with the center of the low-resolution hit window as a starting point. A new high-resolution window was scaled with six different sizes, translated with 2 pixel increments over 25 positions, and the brightness was varied over six levels of intensity. Within each window, a 50×50 pixel image was taken and input to the trained high-resolution network. No primary high-resolution hits were obtained when scanning the image of Juan Carlos Aragon.

The next step was to continue scanning with the low-resolution system. The next face detected was that of Bernard Widrow. The high-resolution system was then engaged and the image of Widrow's face was scanned by scaling, translation, rotation, and varying brightness. A primary high-resolution hit pattern was obtained and compared with the patterns in the folder. A secondary hit confirmed that the image was indeed that of Dr. Widrow, and so he was identified. Reverting back to the low-resolution system and continuing the scanning, the face of Victor Eliashberg was detected. Scanning his face with the high-resolution system generated no hit, so he was not able to be identified. Widrow was the only one identified, since his facial images were the only ones stored in the cognitive memory.

When a 50×50 hit pattern was obtained with Bernard Widrow, the mean square error had been minimized by optimization of the settings of x-position, y-position, rotation, scaling, and brightness. The window was then translated left–right in the x-direction. A plot of mean square error versus the amount of translation away from the optimal position is shown in Fig. 12. In this case the window was translated left–right in the X-direction. A plot of the mean square error versus the amount of translation along the Y-axis away from the optimal position is shown in Fig. 13. A plot of the mean square error versus the angle of rotation from the optimal orientation is shown in Fig. 14. Similar plots were obtained for scaling and variation of brightness. These mean square error plots are easy to search. These plots show that simple relaxation methods will find optimal solutions without hanging-up on relative optima. Similar easy to search plots are obtained with other faces and with other objects. These are empirical results, but they are consistent.

5. A model of human memory

The cognitive memory system described above is a mechanistic system. It is designed to function in a useful way, and it is also

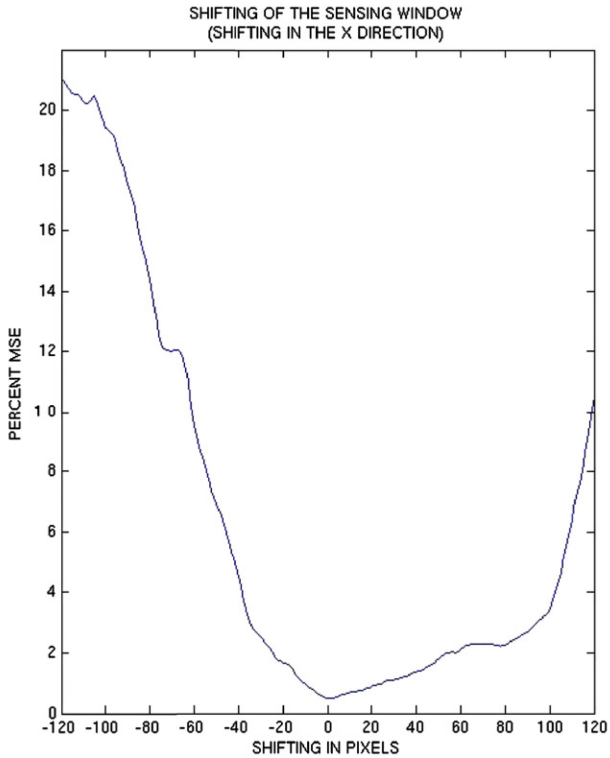

Fig. 12. Effects on mean square error due to left/right translation from the optimal position.

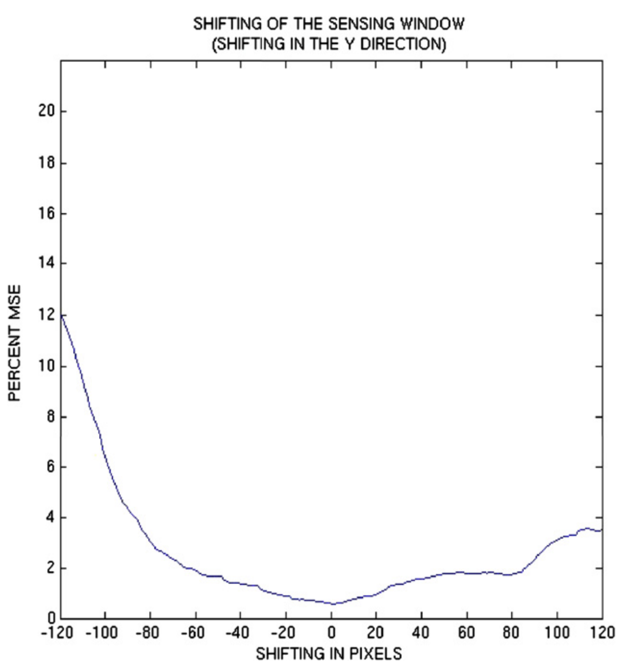

Fig. 13. Effects on mean square error due to up/down translation from the optimal position.

designed to behave in a way that resembles the behavior of human memory as well as we understand this. Between both of the authors of this paper, we have experienced more than 100 years of living with human memory. Observations from everyday life have given us insight in the workings and the behavior of human memory and have led to the cognitive memory model described above.

In accord with Fig. 1, sensory input patterns come from the eyes, ears, tactile, olfactory, vestibular, and other sensors. Incoming patterns are stored in empty folders, wherever they are located. Sequences of patterns, like videos, are stored in the same folder. Visual, auditory, and other sensory patterns that were received at the same time are stored in the same folder. The evidence of this is that upon hearing a song on the radio, one can often "see" an image of the singer's face in one's mind, or likewise, seeing a picture of the face of a singer, a favorite song by that singer comes to mind. Often when you hear a song that you have not heard for a long time, memories of that era come back to you. All this would have to have been recorded together, otherwise sights and sounds, etc., that occurred simultaneously could not be recalled as a unit. This implies that locations in the brain for visual storage are not distinct from locations in the brain for auditory storage. Therefore, visual patterns are not stored in the visual cortex and auditory patterns are not stored in the auditory cortex.

Only "interesting" input patterns are stored in the memory. They remain for the rest of one's life. Sensory input patterns first go to short term memory (of the order of a few seconds) and, if interesting, are transferred to long term memory and recorded for life. Short term memory is used in the determination of what is interesting. Patterns in short term memory are fed as prompts to the long term memory and if there is a hit in any of the segments, then the new input is related to previous inputs and is therefore interesting. How this could work will be discussed in a future paper.

With interesting inputs, sequences of new input patterns are recorded in memory and thereby are added to the stored database. Most of the sounds that we hear are not interesting, and they "go in one ear and out the other". Most of the sights that we see are not interesting and "go in one eye and out the other". When we see something or hear something that "rings a bell", we turn on "the tape recorder" and start adding to the database.

Arguments will be presented below in support of the idea that long term memory is probably not based on storage in the synapses, dendrites, and neurons, i.e. in the brain's neural networks. If so, what are neural networks for? It is known that destruction of the neural networks as with Alzheimer disease results in the destruction of memory. Therefore the neural networks have something to do with memory. We argue that the purpose of neural networks associated with memory is memory retrieval, not memory storage. We believe that autoassociative neural networks are used in the pattern retrieval process.

Fig. 2 shows the training of the autoassociative neural networks. This is done "off-line" during non-REM sleep when the neural networks are otherwise non-active. It is also done during REM sleep while dreaming, and is done while awake and active. How all this could work will be discussed in a future paper.

In each memory segment, the patterns stored in the folders are one-by-one applied to the autoassociative neural network for training. Once the neural networks are trained, they can then be used for memory retrieval.

Data, patterns, and information are retrieved in response to input prompts. Fig. 3 shows the source of the prompts, the input sensors such as the eyes, the ears, etc., whose patterns are preprocessed by the visual cortex, auditory cortex, etc. Visual prompts start with a visual pattern, a snapshot, that is applied to the visual cortex that in turn generates hundreds of images by rotation, translation, scaling, varying brightness, varying contrast, etc. These images are transported throughout the memory by the prompt input line. Each of these images is input to all of the autoassociative neural networks. If there is a hit or hits, the hit prompt patterns are stored temporarily in the associated buffer or buffers. A hit prompt pattern must correspond to at least one pattern in one of the folders of the segment having a hit. The search to find that folder or folders is illustrated by the diagram of Fig. 4. All the patterns in all the folders are one-at-a-time compared with

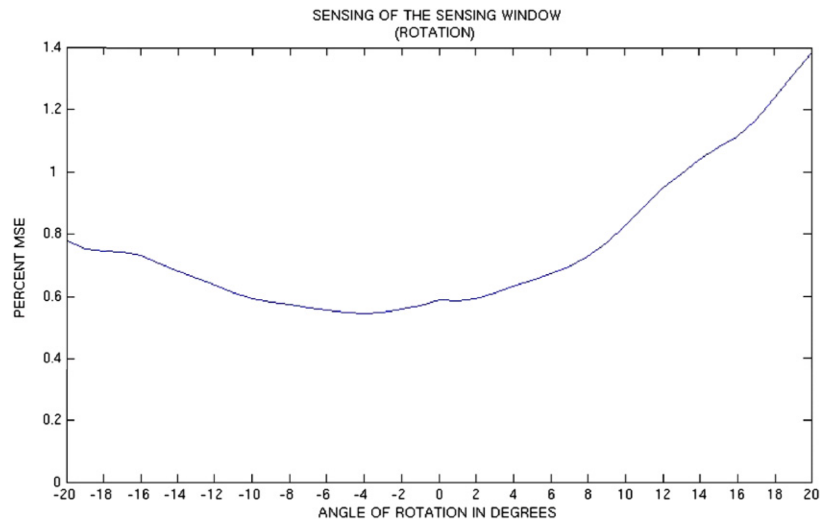

Fig. 14. Effects on mean square error due to rotation from the optimal position.

the hit prompt pattern, and if there is agreement, i.e. if there is a secondary hit, then the entire contents of the folder or folders containing secondary hit patterns are delivered as the memory output.

Subsequent visual snapshots would work in like manner as prompts. This could cause the contents of additional folders to be delivered as memory outputs. All of the memory output patterns could in turn serve as prompts, generating additional primary and secondary hits and delivering additional memory outputs. This is feedback that could work like a "chain reaction". Visual prompts could cause retrieval of auditory patterns, and auditory prompts could cause retrieval of visual patterns. Any sensory pattern or combination of sensory patterns could serve as prompts that could cause retrieval of whatever sensory patterns were originally recorded in the same folder.

In support of these ideas, an example from everyday life is the following. You are looking at a crowd of hundreds of people. Scanning the crowd, you spot a familiar face. You know this person. You come together and start a conversation. You have been having a nice conversation for ten minutes, and you are saying to yourself, what is his name? What is his name? What on earth is his name? Then finally his name pops into your head, it is Jonathan Jones. And you say, "well Jonathan,…". Jonathan thinks that all the while you knew his name. Actually, you did know his name, but you had difficulty in retrieving it.

When you first spotted Jonathan in the crowd, you had a primary hit. You then made a search of the folders of the memory segment that had the hit, and then you had a secondary hit. The folder that you retrieved with the secondary hit did not happen to have his name. You met with Jonathan previously many times, having many Jonathan Jones folders in memory. They do not all contain his name. Continuing the discussion (that you were able to do since you retrieved one or more folders containing previous discussions), you were retrieving the contents of more folders, prompted by the look of his face and sound of his voice. The contents of the retrieved folders also served as prompts. Finally you retrieve a folder containing his name, perhaps the folder containing the events when you were first introduced, where the new visitor introduced himself as "Jonathan Jones".

The model of cognitive memory described above has folders, and segments, and autoassociative neural networks, and multiplexers, and comparators. These are metaphors for functions that would be required for the proposed memory model. Perhaps these functions are performed in human memory. We do not know for sure, but we do know that many experiences that we have had living with human memory correspond to the functioning of the proposed cognitive memory model.

5.1. Memory patterns are probably not stored in synapses and in the brain's neural networks

All memory systems, whether living or electronic, must have some means for information storage and some other means for information retrieval. For example, regarding the hard drive of a computer, data is stored with magnetic particles on a spinning disc. Magnetic particles are the means of data storage. A solenoid floating near the surface of the disc generates voltage pulses whose timing and polarity allows recovery of the recorded digital data. The solenoid and associated electro-mechanical hardware is the means of data retrieval.

A common belief among biologists is that synapses and the associated neural networks are the means for storage in human memory. We will present arguments here that contradict this idea.

Synapses do have memory capability. There is evidence that they vary with experience (Kandel, 2006). However, synaptic memory is analog memory and this is not likely to be stable and consistent over time. Moreover, neural networks in the brain are continually changing. Neurons die out with age and disease. On the other hand, in young children, the brain grows at a very fast rate and new circuitry is added day and night.

To have consistent input–output responses from neural networks whose circuitry may change in unpredictable ways over time, it is imperative that these neural networks be retrained on a regular basis. It is likely that much of this retraining is done every night during REM and non-REM sleep. There is evidence that sleep is crucial for good memory function (Maquet et al., 2000).

In order to train the neural networks, training patterns are required. The training patterns must be stored permanently, somehow. This brings up a question: Why would nature store training patterns to train neural networks to store the training patterns? This does not make sense. It is a "catch 22". If nature has a way to store the training patterns, what is the purpose of the neural networks? We believe that the purpose of the neural networks is to assist in data retrieval, not to do data storage. Maintaining the neural networks by training them with permanently stored training patterns allows them to be a working part of the data retrieval system in spite of unpredictable circuit variations.

Experience with artificial neural networks implemented with electronic hardware has shown that consistent input–output responses can be obtained with partially defective circuits and with circuits which piece by piece gradually become defective. This works as long as the neural networks are continually trained, as long as there is enough left to be trained.

The training patterns are the stored database of the memory. These patterns need to be stored digitally and permanently for long term memory. The question is, how is this done in the human brain? No one knows for sure. We can only speculate.

A very rare person is one with a photographic memory. There are such people. They appear to be normal people except for the photographic memory. They can recall visual images with perfect accuracy over decades. This suggests digital memory. Although other people's memories cannot perform this way, there is probably not much difference in the memory systems of all people. We are probably all capable of remembering with photographic ability, we just do not do this. The question is, how does digital long-term memory work? What is the means of storage?

5.2. Innate knowledge

Humans and animals are born with a considerable amount of inborn knowledge. By observing the behavior of a few-weeks-old dog, one can see that this little animal has behavior, at least in how it relates to humans, that is very similar to that of a mature dog. It is likely that this behavior is not learned but is inborn, inherited from the baby dog's parents and they from their parents, on and on back in time.

The Widrow family adopted a 6-week-old West Highland Terrier puppy named Charlie. This type of dog requires a lot of exercise. Widrow was walking Charlie around and around a back yard swimming pool when suddenly his leash broke and he fell into the pool. Widrow got down on his hands and knees and tried to grab the dog, without success. The water was icy cold, as this was not the swimming season. Widrow was ready to jump in to save the dog's life, but looking up he noticed that Charlie was swimming in a big circle and returned to Widrow who grabbed him, wrapped him in a towel and held him for about an hour until he stopped shaking. He was chilled and terrified, but he knew how to swim! He had never been in water before. The knowledge of how to swim was inherited over thousands of generations, probably back to wolves.

The Widrows have friends, Dr. and Mrs. Jim Spilker. They live in a wooded area abounding with wildlife. One day the Spilkers looked through a window and they saw a deer outside standing several meters away. It was strange to see a deer just standing there. Then they noticed a baby deer emerging from the rear end of this deer.

The baby landed on the ground, the mother turned around to look at the baby, and started slowly to walk away. The baby followed the mother and also walked away. This was a remarkable sight. The baby was born with knowledge to walk, and indeed, with knowledge to follow its mother. All this being essential for survival.

A baby bird emerges from its shell and is hungry. It knows how to say "peep", how to open its mouth, and when its parents put a piece of worm in its mouth, it knows how to swallow. It grows feathers, exercises its wings, and when it comes time to fly, it flies. It could not have learned to fly from the time it jumped from the nest, before it hit the ground. Muscle control of the wings to enable flight is no doubt a very complicated process. The baby bird was born with knowledge of how to fly. All this is necessary for survival.

A human baby is born with knowledge of how to cry, how to pee, how to poop, how to suck, how to swallow, etc. A human baby is even born with knowledge of how to walk, although walking begins about one year after birth, after body development and muscle development make walking possible.

This delay is analogous to the delay for the baby bird before flying. Within a week or so of first attempting to walk, the human baby walks. It is unlikely that such a complicated control process could have been learned by the baby in such a short time. The baby was born with knowledge of how to walk. Dogs, deer, birds, and humans, all are born with substantial knowledge which is essential for survival.

5.3. Creation of a new living animal

At the moment of conception, a sperm and an egg fuse and a new cell is created. The nucleus of this cell contains DNA, some from the mother, some from the father. The DNA in the nucleus contains the information (the "blueprint") needed to construct the new living animal, to build all the organs, to determine how to interconnect them, and how to build the brain. Inborn knowledge inherited from the parents is also contained in the nucleus of the new cell, stored in DNA. Where else could the innate knowledge be stored? There are no neurons, synapses, or dendrites in the new cell to store this knowledge.

Knowledge of how to walk is stored in the developing brain. Originally, this knowledge was stored in the DNA of the new cell. It is most likely that this knowledge stays in the DNA or RNA of the cells of the developing brain and becomes part of its stored database. This does not result in a system for walking that is "hard wired" in the developing brain. The knowledge of how to walk must change as the animal grows, as the legs get longer and the body heavier. It is most likely that this knowledge is stored as software, data in the memory, not hard wired in brain circuits. Software is subject to change as required.

Data in the memory is stored in DNA or RNA, and this is digital memory. Believing that means of storage for long-term memory is DNA or RNA, a question arises about the role of the neural networks associated with long-term memory. What do the neurons, synapses, and dendrites have to do with memory? The answer proposed here is that the neural networks play a critical role, that they are part of the data retrieval mechanism, not the data storage mechanism.

5.4. Grand assumption

Innate knowledge is stored in DNA or RNA. Since nature has a tendency to use the same tricks over and over again, we will assume that the means of storage for knowledge gained after birth, during a normal lifetime, is the same as the means of storage for innate knowledge. We will also assume that the means of retrieval for knowledge gained during a lifetime is the same as the means of retrieval for innate knowledge. Under these assumptions, an individual should be able to seamlessly go from innate knowledge to knowledge gained after birth. The storage and retrieval mechanisms would be the same.

Based on these assumptions, knowledge throughout life is stored in DNA or RNA and retrieval is accomplished by neural networks. These have been the motivating ideas for the design of the human-like "cognitive memory" described above.

6. Conclusions

The design of a practical and useful cognitive memory system is described based on certain ideas of how long-term human memory works. The cognitive memory is content addressable. Sensory inputs concerning a single object or subject are stored together in a memory folder. When the contents of the folder are retrieved, sights, sounds, tactile feel, smell, etc., are all obtained at the same time. Retrieval would be initiated by a query or prompt signal from a current set of sensory input patterns. A parallel search through the memory would be made to retrieve stored data that correlates with or relates to the prompt input. The retrieval system uses autoassociative neural networks. Applications to pattern recognition and face recognition are demonstrated.

A cognitive memory is a learning system. Learning involves storage and retrieval of patterns and data in a cognitive memory. The learning process is unsupervised, i.e. autonomous. Human learning is also based on a cognitive memory, and is autonomous.

It is a common belief that long-term human memory involves storing patterns and data in the brain's neural networks, in the synapses, neurons, dendrites, etc. This is highly unlikely for several reasons. Storage for a lifetime requires a highly stable memory system. Living neural networks do not meet this requirement. In adults and in the elderly, neurons, dendrites, and synapses die out regularly. In newborn babies, neurons, synapses, etc., are grown and connected at a prodigious rate, 24 h per day, as the young brain grows. It is hard to imagine permanent data storage in analog circuits that are changing chaotically.

In spite of the chaotic changes to brain circuitry, the neural networks are able to function admirably as long as they are trained continuously. "ADALINE" and "MADALINE" were constructed by Widrow and his Ph.D. students at Stanford in the early 1960's. MADALINE, a hand-made six neuron system, was trained to perform a significant pattern classification task with, unbeknownst to its designers, twenty five percent of its analog synapses totally defective. Fixing the defects and then retraining, the improvements in performance were hardly noticed.

Neural circuits with defects and unpredictable circuit changes can function perfectly well, but they must be continually trained. To train the brain's neural networks, training patterns must be stored permanently, in some form of digital memory, not in neural networks with unstable analog memory. The training patterns comprise the contents of long-term memory. It is not likely that nature would permanently store the training patterns to train the neural networks to permanently store the training patterns. The neural networks must serve some purpose other than permanent data storage.

A new life is formed at the instant of conception when a sperm and egg fuse to form a new single cell whose DNA comes from both parents. The DNA stores the formula for building the new animal, constructing the body, the organs, all their interconnections, and the brain. Inborn knowledge, patterns and data, are stored in the long-term memory of the developing brain. Inborn knowledge is inherited from both parents and must have been initially stored in the nucleus of the original single cell, probably in its DNA. Two ideas come to mind. One, DNA is a permanent digital storage mechanism, and two, there are no neurons, synapses, dendrites, etc., in the nucleus of the single cell, yet inborn knowledge is stored there. Knowledge, patterns, and data are evidently stored there without neural networks, and are probably stored digitally in DNA. Since nature tends to use the same methods over and over again, we speculate that long-term memory is stored in DNA and/or RNA. Considering the number of cells in the brain, neurons and glia, the amount of DNA that could be available for storage is enormous.

Alzheimer's disease attacks and destroys neurons, synapses, dendrites, etc., and a result is memory loss. It is therefore clear that the brain's neural networks are involved with memory. However, if memory is not stored in the neural networks, what purpose do they serve to memory? We speculate that the purpose of neural networks in the long-term memory system is for information retrieval, not information storage.

Recent analyses about the progression of Alzheimer's disease indicates that the destructive process begins about ten to twenty years before the patient becomes symptomatic. Thinking about MADALINE, a machine that was trained with ease with 25% of its circuits defective, one can imagine a person with this disease functioning very well in spite of the damage. Constant training keeps the memory working, until the damage becomes excessive. Elderly people would be advised to keep working, keep thinking, and keep doing things. Use it or lose it, as long as you can.

Contrary to accepted theories about the workings of human memory but based on the above reasons, we believe that long-term memory is stored in DNA or RNA, and that the purpose of the memory's neural networks is memory retrieval, not memory storage.

References

Kandel, E. (2006). *In search of memory: the emergence of a new science of mind*. New York: Norton & Company.

Maquet, P., et al. (2000). Experience-dependent changes in cerebral activation during human rem sleep. *Nature Neuroscience*, *3*(8), 831–836.

Werbos, P. J. (1974). *Beyond regression: new tools for prediction and analysis in the behavioral sciences*. Ph.D. Thesis. Harvard University, Cambridge, MA.

Widrow, B., & Winter, R. (1988). Neural nets for adaptive filtering and adaptive pattern recognition. *Computer*, *21*(3), 25–39.

Further reading

Eichenbaum, H. (2002). *The cognitive neuroscience of memory: an introduction*. New York: Oxford University Press.

Fukushima, K., & Miyake, S. (1978). A self-organizing neural network with a function of associative memory: feedback-type cognitron. *Biological Cybernetics*, *28*(4), 201–208.

Fuster, J. M. (2002). *Cortex and mind: unifying cognition*. New York: Oxford University Press.

Haykin, S. (1999). *Neural networks: a comprehensive foundation*. New Jersey: Prentice Hall.

Hecht-Nielsen, R. (1989). *Neurocomputing*. Reading, MA: Addison Wesley.

Kohonen, T. (1989). *Self-organization and associative memory* (3rd ed.). Berlin: Springer-Verlag.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Phillips, P. J. et al. (2005). Overview of the face recognition grand challenge. In *IEEE Computer society conference on computer vision and pattern recognition 2005, Vol. 1* (pp. 947–954).