The Atoms of Cognition: A Theory of Ground Epistemics

Seng-Beng Ho (hosengbeng@nus.edu.sg)
Temasek Laboratories, National University of Singapore
5A Engineering Drive 1, #09-02, Singapore 117411

Abstract

We propose a set of "atomic cognitive operational representations" on which higher level cognitive representations and processes can be built, thus providing fundamental building blocks for cognitive mechanisms necessary for intelligent actions. The fundamental concepts involved are elemental temporal changes of some quantities and in this representational scheme the temporal dimension is explicitly represented to fully characterize the meanings of the concepts involved at the epistemic ground level. This provides full grounding for all subsequent concepts that are built upon them, allowing cognitive systems embodying these concepts to have full and complete understanding and characterization of the concepts involved that it can use for various cognitive ends. This provides a firm theoretical foundation for the study of cognition and intelligence.

Keywords: representation; operational representation; spatiotemporal representation; conceptual grounding; fundamental building blocks of cognition; experiential memory

Introduction

Unlike the physical sciences, "standard" paradigms for scientific investigation of cognitive phenomena still do not exist for the sciences of cognition and intelligence (Arbib, 2002; Gazzaniga, 2008; Russell & Norvig 2010). Chief among the achievements of the physical sciences is the discovery of various fundamental particles and forces that provide the foundation for the understanding of physical reality. These fundamental particles and forces form the necessary and sufficient building blocks upon which the characterization of all other higher level physical phenomena can be constructed.

We show that a set of representations, which we refer to as "atomic cognitive operational representations" or *ACORs*, can perform an equivalent function of providing the fundamental building blocks for building cognition. Like elementary particles and forces, ACORs allow ground level semantics to be represented and all higher level semantics and cognitive operations to be based on them. As a dual of physical ontology, the atomic cognitive operational representations provide the *cognitive ontology* for building all cognitive phenomena.

In our theory, a set of correctly formulated ground level "atomic" building blocks of cognition will be necessary for allowing all cognitive processes to be constituted. Furthermore, the same fundamental building blocks apply at all epistemic scales to enable useful knowledge to be derived through cognition for intelligent functioning. Consider the following.

Before the advent of science, we dealt with the world as best we could, at a level of description provided by our natural sensory systems. Take for example the case of a tree as perceived by our human senses. From a distance, a tree can be sensed and conceptualized as consisting of a trunk, some branches and many leaves. However close up, our senses can tell the texture and the detailed shapes of the trunk, the branches and the leaves, and perceive their movements. The detailed perception of these subparts of a tree in turn provides us with the necessary knowledge to be able to use those subparts for various purposes, such as for decoration or other more functional ends. When the wind blows and the leaves on a tree move, the softness of the movement might allow us to "imagine" using the leaves as a broom for the purpose of sweeping, or as a fan to fan oneself; in due course, we might proceed to act on those ideas when the need for them arises. This would be the characterization of the tree at an epistemic ground level at which normal perception operates to provide useful information for a cognitive system to function intelligently.

As science improves our understanding of the natural world, our conceptualization of trees goes beyond just trunk, branches and leaves as they appear to our natural senses. We discovered chlorophyll and the photosynthetic processes, for example. These biological processes involve much smaller entities – various molecules – and attendant complicated interactions that our natural senses cannot detect directly. With this deeper understanding, we reach another epistemic ground level. It is the thesis of this paper that no matter which ground level we are looking at – i.e., whether it be the earlier one reached by our natural senses or the deeper ones reached through scientific means – the same fundamental building blocks of cognition are involved in subserving cognitive mechanisms from which intelligent actions emerge.

With the fundamental building blocks, we put the sciences of cognition and intelligence on a firm theoretical foundation, much as what the theory of fundamental particles and forces has done for the physical sciences.

The Atomic Operational Representations of Appearance and Existence

The vast canvas of space-time is the arena in which physical and mental processes take place. One fundamentally important concept that any mental process must deal with is appearance and existence (of objects, events and processes). We begin with the formulation of the atomic concept of appearance.

Figure 1a shows a one dimensional space with 5 discrete positions x0 – x4. At time t0 it contains nothing. At time t1, something appears at position x1. It could be an elemental bit of substance or a point of light. Now, assuming an intelligent system has perceptual detectors that can sense and signal something appearing at a certain spot at a certain time and has an explicit *experiential memory* that can store its perceptual experiences and lay them out in a spatial extent for simultaneous processing, Figure 1b shows what the explicit experiential memory looks like in the time interval t0 to t1. Figure 1b also shows a spatiotemporal "appear" template that can pick up this change in the physical situation in the one dimensional space. Basically this template encodes the meaning of "appear." It captures the situation in which in a moment in time just prior to the appearance of something, there was nothing in that point in space and in the next moment something came into existence. In the spatiotemporal template, time is spatialized and the temporal dimension is explicitly represented.

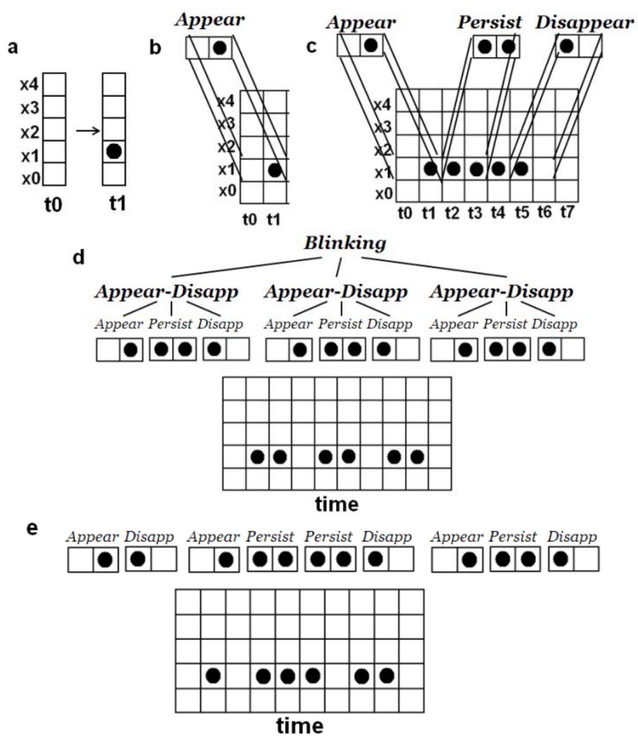

Figure 1: Atomic operational representations for "appearance" and related concepts. a. Appearance of an object at a specific time and location. b. The experiential memory and the "appear" operational representation. c. "Persist" and "disappear" operators. d. Concept of "blinking." e. Re-composing operators – cognitive manipulatability.

It should be stressed that the elemental "blob" that is picked up and represented in the experiential memory of Figure 1b may or may not correspond exactly to an elemental "blob" in the physical world. It is an "elemental" bit of occurrence as characterized by the sensory organ of the cognitive system relevant to the cognitive task at hand. Depending on the resolution of the sensory system and the cognitive task involved, the "blob" could correspond to an atom in the physical world, a vehicle or a person, or the recognizable points on a human body or on the leaves of a tree. It could also be something created by other internal mental processes in a "mental space."

We submit that this spatiotemporal *appear* template captures the full meaning of *appear* at the ground level. For the purpose of identifying an *appear* event, the template is used as follows: it is matched to the spatiotemporal patterns in the experiential memory as shown in Fig 1b. If there is a match, an *appear* event is identified to have taken place.

The *appear* template, other than being useful for the purpose of identifying or recognizing the occurrence of an *appear* event, is also usable in the opposite direction – that of *generating* the action or idea of *appear*. Assuming that an intervening system is available to translate an intelligent system's intention into physical realization, the intelligent system can act, under the direction of the *appear* template, to move a material substance into a specific point in space and thus making it "appear" at that location, to materialize a bit of substance from nowhere at a specific point in space, or to cause a point of light to appear at a specific point, etc. These are examples of "make-appear" – the *appear* template used in the generation direction. The *appear* template can also be used to generate an "idea" (i.e., a bit of "mental substance" – e.g., the mental analog of a corresponding physical quantity – or something more abstract) at a point in a mental space for the purpose of mental manipulation. An example of a mental space would be the spatiotemporal representational structures that hold the experiential memory for mental manipulation, such as the structure shown in Figure 1c that might be implemented using a computer memory. Generation of ideas in the mental space is also sometimes referred to as "imaging" or "imagining" (Kosslyn 1994). We will refer to this characteristic of these atomic operational representations as "operational bidirectionality" – i.e., they can be used for recognition as well as generation.

We thus term the *appear* spatiotemporal template an "operational representation" as it encodes the meaning of an operational concept – *appear* – directly in terms of the operations that it performs. Such an operational representation is *atomic* in that it captures knowledge at the ground level. An atomic operational representation is *elemental* in that it represents a smallest discrete change in the spatial and temporal dimensions relevant to the cognitive task at hand. A more succinct description of "atomic operational representations" is "atomic operators."

Figure 1c depicts two other concepts related to *appear* – "persist" and "disappear." Together with *appear*, they can be used to capture the process of something appearing at a point, persisting (existing) for some time, and then disappearing. Figure 1b and 1c also suggest that these templates or operators can be picked up from the environment, in the same manner as the “cookie-cutter” approach described by Uhr and Vossler (1981) – i.e., when a subpattern is perceived in the environment, it is simply “cut-out”/“picked-up” and stored as an elemental pattern to be used later to match with further pattern information coming in through the perceptual system to recognize future occurrences of the subpatterns. This is usually known as an unsupervised learning process (Duda, Hart & Stork, 2001; Fukushima, 1988; Malsburg, 1973). Whereas in Uhr and Vossler (1981), the “cookie-cutters” work on static images containing certain spatial patterns, thus they extract sub-features of the *spatial* patterns for the purpose of characterizing these patterns, the corresponding application here would be extracting *spatiotemporal* sub-patterns that serve to characterize activities or changes in the environment, external or internal to the cognitive system.

Figure 1d shows how higher level concepts can be built upon the atomic operators of *appear*, *persist* and *disappear* – if something alternately appears and disappears over time, the concept of “blinking” can be used to encapsulate the process. The higher level concepts are firmly grounded through the atomic operational representations. This process is termed “cognitive hierarchy construction.”

Figure 1e shows how the atomic operators can be recombined into novel sequences that may not have been encountered in the environment earlier in a generation process through the conceptual hierarchy, directed by reasoning and problem solving requirements. This characteristic is referred to as “cognitive manipulatability” – i.e., these representations can be directly manipulated in cognitive processes.

Base on the foregoing discussion, the critical characteristics of operational representations in general and atomic operational representations in particular are summarized as follows:

1. *Explicit Temporal Representation* - the temporal dimension is explicitly represented in operational representations – this requires the intelligent system utilizing the representation to have an explicit experiential memory – temporal changes are laid out in a spatial extent for simultaneous processing.
2. *Elemental Representation* – atomic operators are elemental in that they represent the smallest discrete changes in the respective dimensions.
3. *Cognitive Hierarchy Construction* – higher level concepts can be built directly upon the lower level as well as the atomic level operational representations.
4. *Operational Bidirectionality* - operational representations can be used for recognition as well as generation purposes. Recognition and generation can involve things, events and processes in both the physical environment and the mental space.
5. *Unsupervised Extraction* – atomic level as well as higher level operational representations can be extracted from the environment in an unsupervised learning process.
6. *Grounding of Representations* – atomic operational representations are grounded directly in the environment or the mental space, higher level representations are grounded on the atomic representations.
7. *Cognitive Manipulatability* – The operational representations, atomic or higher level, can be recomposed and cognitively manipulated in cognitive processes.

The atomic operational representations capture and encode “meaning” in the operations themselves. For example, *appear* is not defined in terms of other atomic or non-atomic concepts but is instead defined directly in the recognition and generation operations that it stipulates.

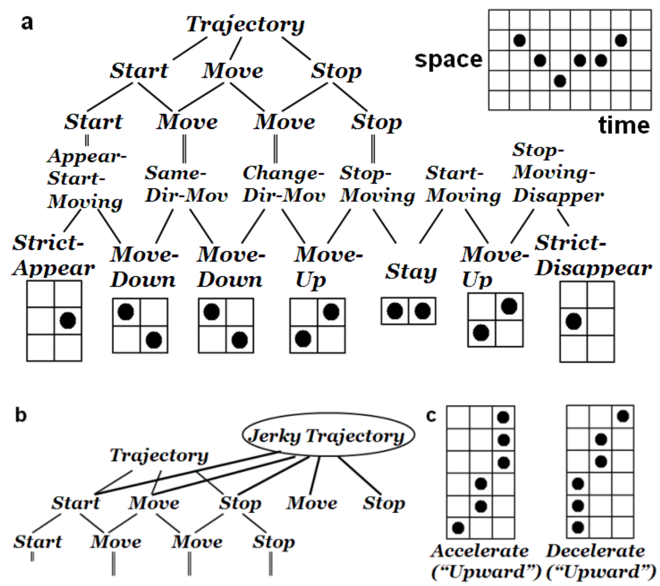

Figure 2: Operational representations of movements. a. Up and down movements in a one-dimensional space and the associated operational representations and conceptual hierarchy. b. A “jerky trajectory.” c. Accelerate and decelerate operators

Atomic Operational Representations of Movement

Again, using events in a discrete one dimensional space, we illustrate the atomic operational representations associated with movements in Figure 2. In Figure 2a, we show the operational representational characterizations of a point object appearing at a certain location, moving “up” and “down” in the one dimensional space, and then disappearing. There are basically two kinds of spatiotemporal templates associated with atomic movements – “move-up” and “move-down.” Here we also introduce a more specific kind of *appear* template called “strict-appear” – i.e., the appearance of something in a specific location is not caused by things moving into the location. If the earlier, more general *appear* in Figure 1 is used, then at every location where the point object moves into, the *appear* template will be triggered. Similarly for *disappear* – we introduce a "strict-disappear" template. The "stay" template is similar to the *persist* template in Figure 1. It is interesting to note that the concept of *stationarity* – "stay" – can only be represented adequately if there is an explicit temporal dimension that can capture and encode the meaning of "no change in position in time." Similarly for the concept of "persist" – "no change in the state of existence" – depicted in Figure 1c.

Figure 2a shows a cognitive hierarchy built upon the atomic operational representations of appearance and movement, reaching a level where the concept of "trajectory" is formed. Other than *trajectory*, concepts such as "jerky trajectory" can emerge from the conceptual hierarchy, as shown in Figure 2b.

Figure 2c shows the concepts of *acceleration* (an "upward" acceleration) and *deceleration* (an "upward" deceleration) captured in their corresponding operational representations. Again, here we can see that the *meanings* of *move-up*, *move-down* or *acceleration* and *deceleration* are captured directly in the corresponding associated operations stipulated by the operational representations.

Atomic Operational Representations of Scalar and General Parameters

At any given point in physical space, a physical quantity, such as the hardness of some substance or the brightness of a point of light, may change elementally. Similarly, in mental space, the strength of an idea or feeling may also change elementally. These non-spatial scalar parameters are represented in the same manner as the spatial parameters as atomic operational representations as shown in Figure 3a, using the examples of a physical parameter such as "brightness" or a mental parameter such as "pain" that increases and decreases over time.

Atomic movement through space (physical or mental), which involves an elemental change in spatial position, can be considered a special case of a change in something, so does a change in the state of existence in space – i.e., *appear* or *disappear*. Therefore, the most general characterization of all the atomic operators including those that capture non-spatial parameters is that they characterize an elemental change of these parameters, spatial or otherwise, across elemental time. Figure 3b shows a generalized atomic operational representation of parameter change – like a "God particle" of atomic operators – that subsumes all possible atomic operators.

The number and kind of atomic operators that can be found in a particular intelligent organism or system depends on the parameters that are necessary for its survival or intelligent operations. For example, if an organism needs salt for survival and has the necessary sensory apparatus to detect the presence of salt with varying intensities over time, then for the purpose of cognitive processing it would need "saltiness" atomic operators. In conscious perception such as that typically experienced and reported by human beings, the atomic level parameters typically appear as indivisible cognitive entities endowed with certain "qualia" such as pain, sadness, anxiousness, brightness, darkness, redness, saltiness, sweetness, spatiousness, etc. These atomic level cognitive parameters as appear to consciousness can vary in their intensities over time but otherwise have no further internal structures scrutinizable by consciousness.

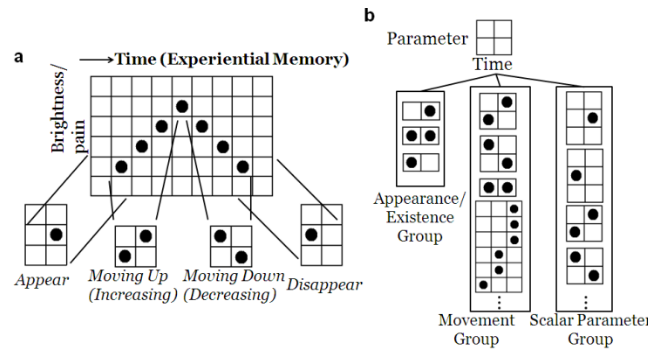

Figure 3: Other operational representations. a. Atomic operational representations of scalar parameters (brightness or pain) increasing and decreasing over time. b. Generalized atomic representation of parameter change.

Representing and Reasoning about Time

If a cognitive system provides for the explicit representation of time through an experiential memory such as that described above, it can lead to another interesting mental operation, which is the conceptualization of time as a "thing" and the "movement" of it mapped onto spatiotemporal dimensions in the cognitive system's *mental space-time* that allows the cognitive system to cogitate about time itself. Figure 4 shows the representations and operations involved for this.

In Figure 4 we re-label the horizontal axis as "cognitive time," to distinguish from the "real" physical time out there in the real world. (All previous labels of "time" in Figures 1 - 3 should rightfully be "cognitive time" as these operational representations are *mental* entities.) Here, time is characterized as a "thing" that can "move." Normally, we would conceptualize time as something that moves on inexorably with a "constant speed." In the physical universe, it is hard to ascribe "speed" to time as one needs a time reference to talk about speed, and such "super-time" is non-existent. However, in our mental processes, we can freely cogitate about these possibilities. Firstly, we often imagine real-world events going faster or slowing down. Therefore time can "accelerate" or "decelerate," and that would be applying the *acceleration* or *deceleration* operators as depicted in Figure 2c in the time representation here. We sometimes wonder if time has a beginning or end, and these would correspond to applying the *appear* and *disappear* operators as shown in Figure 4. If we cogitate whether time can come to a halt, that would correspond to applying the *persist/stay* operator. And in physics the concept of time reversal has been invoked to explain certain physical phenomena and that would correspond to applying a *move-down* operator as shown in Figure 4. The concept of time travel to the past or future would correspond to a sudden jump or accelerated movement to a different "point" in "time" from the present point. Our reasoning processes with time thus use similar mental processes as we use for space, once time is conceptualized as a "thing" and its "movement" mapped onto spatiotemporal dimensions in our mental space-time.

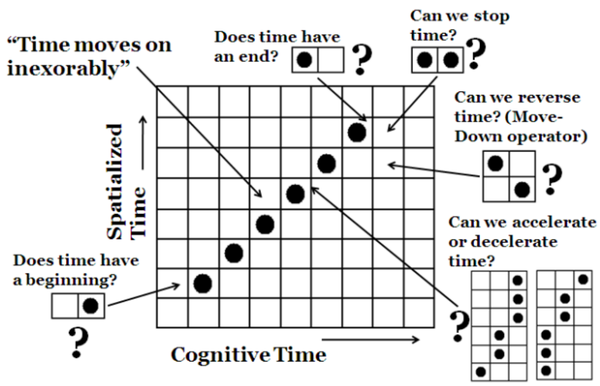

Figure 4: Cogitating about time. Representations capturing the movement, acceleration, deceleration, halting, beginning (appearing), ending (disappearing), and reversal of time for the purpose of cogitation about time.

Representing Interactions

The interactions between objects can be characterized using operational representations. Figure 5a shows a situation in a one dimensional world in which 2 objects hit each other and are reflected. The concepts of "meet", "part" and "reflect" can arise in this situation as shown in Figure 5a and be encoded in the form of operational representations. Figure 5b shows the situation with an alternative physics where a moving object penetrates a stationary object and no reflection takes place. The concept of "penetration" can likewise be characterized at the ground level through operational representations.

Extended Spatial Objects and Higher Dimensional Representations

The object to be represented by an operational representation can also be an extended object consisting of more than one point. Each point on the object can be represented by an atomic operational representation. If all the points on an object move in unison, it is a rigid object (the concept of "rigidity" can be characterized by some higher level operational representations building on top of the atomic operators in the same vein as the *blinking* and *trajectory* concepts shown in Figures 1d and 2a respectively). Otherwise, the object is deformable and the deformations can be characterized by how the points on the object move relative to each other likewise through the use of some higher level characterizations of the deformations built-up from the atomic operators.

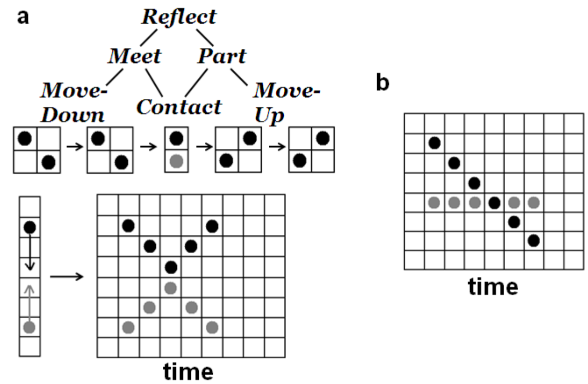

Figure 5: Interactions captured in operational representations. a. Two objects reflect off each other. b. Penetration of one object through another one.

An atomic operational representation that represents elemental movements in one of 8 directions in a 2-dimensional (2D), discrete, and Cartesian space would be a cube of 3X3X1 as shown in Figure 6a. If some liquid on a flat surface moves and changes shape more or less in a plane, its complex shape changes can be captured by the 3D (2D space + 1D time) atomic movement operator as shown in Figure 6a. Higher level characterizations of the liquid movement can be built from the 3D atomic operators.

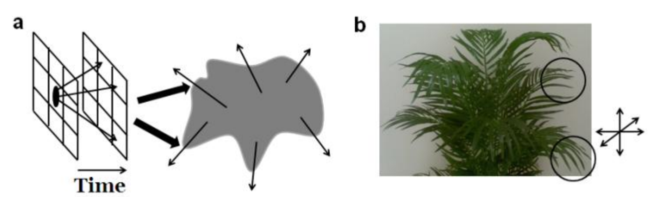

Figure 6: Complex object representations. a. Liquid moving and changing shape in 2D and the associated 3D atomic operator. b. Leaves and 3D micro trajectories of all their elemental movements.

Figure 6b shows a plant's leaves whose movements can be characterized using the atomic operators that can lead to the intelligent construction of a "broom" or "fan" made from leaves – i.e., having captured and understood a certain soft character of the potential elemental movements of the leaves on the tree, the cognitive system is then able to *imagine* the use of the leaves for the purpose of sweeping dirt or fanning air to cool oneself. For characterizing the leaves, because of 3D movements, a 4D atomic movement operator (3D space + 1D time) is needed. Bare branches will not function correctly as a broom or a fan because their elemental parts lack certain surfaces that can move in a certain "soft" manner. This understanding requires the high resolution ground level characterization of the potential movements of the leaves as captured by the 4D atomic movement operators and higher level characterizations of the movements built on the 4D atomic operators.

The atomic operational representations are the fundamental building blocks of cognition. All events and processes – including the operations of natural phenomena ranging from the micro to the cosmic scales (e.g., from the operations of bacteria to trees, weather, events and processes in the heavens, etc.), operations of human-made machineries (e.g., the operations of the finest mechanical contraptions to the operations of a gargantuan rocket, etc.), and mental operations (e.g., thinking processes in general, abstract mathematical manipulations in the mind, etc.) – are fundamentally cognitively describable based on the atomic operational representations.

Discussion

It is not a coincidence that the "mental" atomic operational representations described in this paper for the purpose of cognitive processing bear a strong resemblance to the usual "physical" spatiotemporal representations many a scientist and engineer employs to represent physical events and processes for the ease of understanding and manipulating them. This is because there is an intimate link between the physical and the mental world. Spatiotemporal representations are fundamental and lie at the very foundation of the descriptions of all events and processes that can take place in this reality. Cognitive processes are the dual of the physical processes in that they seek to represent the physical processes but they have additional and unique characteristics and operational requirements as described in this paper as they need to manipulate the representations in certain cognitively useful ways.

The grounding of concepts, and hence the learning of the atomic operators, would take place in the early days of a cognitive system's interaction with the environment (e.g., the infant stage in humans). Future studies of cognitive developmental processes could direct their investigative efforts to uncover whether or how these atomic operators are learned and used.

Our atomic operational representational scheme dictates that the temporal dimension needs to be represented explicitly for the adequate capture of the ground meanings of concepts. Hence, it follows that for a system to qualify as "cognitive," it must have the capability of representing and manipulating the temporal dimension explicitly. A critical mechanism that subserves explicit temporal representation is the experiential memory (Figure 1b). Neuroanatomically, the experiential memory can be identified with the hippocampus, which is supposedly the neural structure responsible for memory over time – episodic memory (Anderson et al, 2007). In this paper, for the sake of exploring the basic principles we describe the experiential memory as a structure that simply lays out the temporal information in an uncompressed manner. In reality, because of the amount of information in the temporal domain that a cognitive system needs to process, compression in the time dimension is probably necessary. We submit that the hippocampus performs both the experiential memory function as well as the compression function necessary for its practical functioning. We also submit that for a system to have "cognitive abilities" and be considered a "cognitive system", it must possess a structure that functions like the hippocampus or the experiential memory that processes and represents temporal information explicitly. This allows it to truly conceptualize and understand events, processes and causality in the world. Even the concept of *stationarity* and *changelessness* can be captured adequately only if there is an explicit temporal representational dimension as discussed above. Hence, artificial or natural neural networks systems that merely reproduce certain input-output mappings of some seemingly "intelligent" processes are "reflexive" and do not qualify to be characterized as "cognitive systems."

A last and interesting fact to note is that the spatialization of time in operational representations parallels the spatialization of time in relativity which brought us a deeper understanding of the reality in which we inhabit (Lorentz et al, 1923).

Acknowledgments

I thank Dr. Kenneth Kwok and the anonymous reviewers for their suggested changes to the manuscript.

References

Anderson, P., Morris, R., Amaral, D., Bliss, T., and O'Keefe, J. (2007). *The Hippocampus Book.* Oxford UK: Oxford University Press.

Arbib, M. A. (Ed.) (2002). The *Handbook of Brain Theory and Neural Networks.* Cambridge, Massachusetts: MIT Press.

Duda, R. O., Hart, P. E. & Stork, D. G. (2001). *Pattern Classification*. New York: John Wiley & Sons, Inc.

Fukushima, K. (1988). Neocognitron: A hierarchical neural network capable of visual pattern recognition. *Neural Networks*, 1,119-130.

Gazzaniga, M. S., Ivry, R. B. & Mangun, G. R. (2008). *Cognitive Neuroscience: The Biology of the Mind*. New York: W. W. Norton & Company.

Griffiths, D. J. (2008). *Introduction to Elementary Particle.* Weinheim: Wiley-VCH.

Kosslyn, S. M. (1994) *Image and Brain: The Resolution of the Imagery Debate*. Cambridge, Massachusetts: MIT Press.

Lorentz, H. A., Einstein, A., Minkowski, H., and Weyl H. (1923) *The Principle of Relativity: A Collection of Original Memoirs on the Special and General Relativity*, with notes by A. Sommerfield. Methuen and Company, Ltd.

Malsburg, C. V. D. (1973). Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* **14**, 85-100.

Russell, S. & Norvig, P. (2010). *Artificial Intelligence: A Modern Approach*. New Jersey: Pearson Education, Inc.

Uhr, L. & Vossler, C. (1981) A pattern-recognition program that generates, evaluates, and adusts its own operators. In Feigenbaum, E. A. & Feldman, J. (eds.) *Computers and Thought*. Malabar, Florida: Robert E. Krieger Publishing Company, Inc.