



Patom Theory

White Paper

John Ball, Cognitive Scientist, B.Sc., M. Cog Sc, MBA

Founder, Pat, Inc.

Is it possible for a single paradigm to explain...the cognitive sciences?

One of the great challenges in neuroscience and related disciplines is to understand how the brain calculates and processes information^{1, 2, 3}. To the linguist, a great challenge is in understanding how the brain innately discovers language^{4,5} to produce unique sentences. Psychologists want to understand how we learn, how we can change the meaning of what we have learned⁶ and how we manage our emotions⁷. To the Artificial Intelligence (AI) researcher, the challenge is in processing visual and other sensory information to construct what you see⁸ and then placing this information in generalized information storage⁹.

Current literature argues for the idea that the human information processing system has a number of intelligences⁹ including linguistic, musical, and logical-mathematical. Books describing "emotional intelligence" and "visual intelligence" have been written, further diffusing the definition of intelligence.

Is it possible for a single paradigm to explain the majority of our observations in the cognitive sciences? Even with our massive knowledge base in the cognitive sciences, our best models today remain inadequate to explain fundamental human capabilities. Patom theory was developed to explain how a brain can do what it does so we can emulate it on a machine. It is a science and therefore it must first explain our observations, and second make useful predictions to validate it.

¹ McCrone, J., *Going Inside: A Tour Round a Single Moment of Consciousness*, 6 (Faber and Faber, London, 1999).

² Maddox, J., *What Remains to be Discovered*, 280, 294-301 (Papermac, London, 1999).

³ Pinker, S., *How the Mind Works*, 24-27 (Penguin Books, London, 1998).

⁴ Ibid. 356-357.

⁵ Pinker, S., *The Language Instinct: The New Science of Language and Mind*, 32-33 (Penguin Books, London, 1994).

⁶ Wessells, M. G., *Cognitive Psychology*, 35-36 (Harper & Row, New York, 1982).

⁷ Goleman, D., *Emotional Intelligence*, 3-29 (Bloomsbury, London, 1996).

⁸ Waldrop, M. M., *Machinations of Thought*. Science 85 Volume 6, Number 2, 41 (1985).

⁹ Gardner, H., *Frames of Mind: The Theory of Multiple Intelligences*, 73-278 (Fontana Press, London, 1993).



Pattern-matching, not processing

Many believe that our brain is the greatest processing machine in the universe. They think that when you see something, your brain recognizes it by rapidly processing its characteristics using incomprehensibly complex algorithms.

Cognitive science tells us that the brain is more of a *pattern-matching* machine than a *processing* machine. Imagine that your brain stores vision as photographs—we know it doesn't but the analogy helps to explain the idea. Your brain will compile a vast library of pictures over time, as each new experience adds multiple photos to your collection. This library stores, matches and uses patterns. To recognize a coffee pot, for example, your brain will go through its picture library to find the best match. You don't need intelligence to find the best fit. While picture matching appears of little value at this stage, given a large enough picture library, it isn't hard to imagine regularly finding matches.

It doesn't really matter that you find a similar picture in your library since the photo doesn't tell you anything other than you have seen it before. Looking through an album is not enough to explain your thoughts, is it? Let's say while your brain is taking the photos, it is also making a short tape recording of the sounds around you. It also records other senses—smell, touch, taste—whatever it perceives at that moment. Each little tape is stored in your library and thin wires linking them to their corresponding photos. This is the basic idea underlying Patom Theory—automatically linking together atoms based on experience.

Pattern Atoms

Patom Theory proposes the brain to be a collection of pattern-matching units¹⁰. Each unit is self-contained and matches specific patterns. The patom¹¹ (named to combine the word "pattern" with the word "atom") is the smallest defined unit that stores, matches and uses patterns. As the brain is comprised of cells called neurons, it follows that patoms are neural networks.

Patoms store patterns when conditioned by repetition, sequence and continuity. A pattern comes in 2 flavors—snapshot and sequential. Snapshots are an instant in time, like a photograph, and sequential is a list of snapshots. A snapshot can be one instance of sensory attributes or the set of multi-sensory instances of objects. Snapshots are stored in sufficient detail to match or repeat the same pattern again. Multiple snapshot patterns are necessary to represent physical objects.

Snapshots are bidirectional. Upon recognition of a previously stored snapshot, a unique signal is sent both forwards to the next layer and backwards to the previous layer. The signal is not an *encoded message* in the traditional sense. It simply identifies a match. Identifying a part of a snapshot is sufficient to identify the entire hierarchy. For example, a song can be recognized by a small snippet of the song because each part of the song is a sequence of snapshots.

¹⁰ Carter, R., *Mapping the Mind*, 8 (Phoenix, London, 1998).

¹¹ Ball, J. S., *Our Brain, the Patom-matcher, Ockham's Razor*, (ABC Radio National, 16 January, 2000; <http://www.abc.net.au/rn/science/ockham/stories/s73842.htm>).

A patom automatically stores and matches patterns based on received inputs. Matched patterns send a unique output, weighted by experience which directly link to other patoms. The location, dictated by its topology of connected links, determines whether the patterns are single-sense or multiple-sense. Input to the brain is at the edge. Vision edge is the eyes receptors. Touch edge is where receptors in the skin start. Balance edge is at the semicircular canals. Higher levels are initially within the one sense area. These are then combined together to form multi-sensory patterns.

While higher-level Patoms are essential to multi-sensory patterns over long timeframes, only sensory patterns in edge Patoms interact directly with our body. Snapshots stored in edge Patoms contain sufficient detail to recognize the transduced sensory pattern, while the rest of the system simply supplies links. The contents of the links from the edge Patoms are incomprehensible without referencing back. Our brain's architecture is structured with multiple layers of bi-directionally connected neural networks. Thus, patoms are hierarchically connected and operate bi-directionally and are based on concurrency and continuity. Figure 1 is a high-level diagram of interconnected patoms.

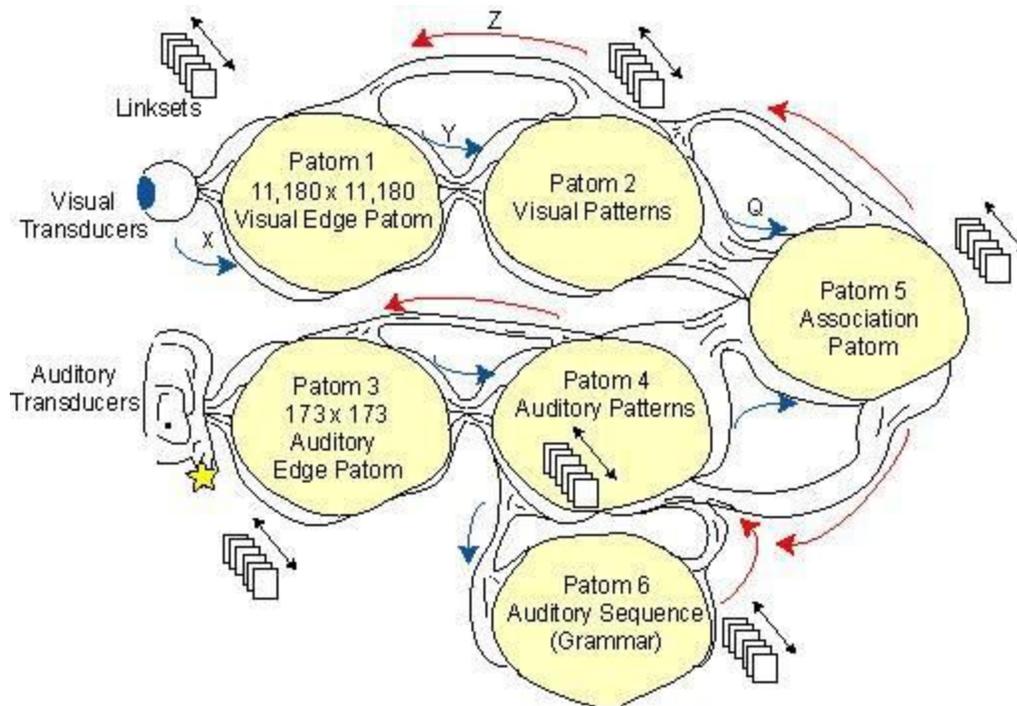


Figure 1. A high-level diagram of interconnected patoms. Here sensory patoms from the edge (sensors) combine within senses first, then in higher levels as associations. By thinking of each patom as a system to track objects, the need for processing is eliminated and replaced by connected patterns. This has implications to machine learning, which need only store and link the actual experience received by senses, and combining them through repetition to create the re-usable fragments seen by neuroscientists.

While it may appear confusing (Damasio, 1994) when compared with a model in which senses are connected directly to motor control areas, it aligns with the needs of a reasonably slow brain driven by



evolutionary necessity to act rapidly and the proposed layered approach to control complex multi-sensory patterns.

Linking patoms

A linked group of patterns form “linksets”. A linkset is a *specific* pattern, which is comprised of a number of patterns—snapshot or sequential. With language, linkset intersection finds the common features between two linksets. The word “eat” is in a linkset including its pronunciation (motion to move muscles in speech), its sound (auditory recognition), its function (semantic relations between objects and their properties) and so on.

If we go back to our original challenge of recognizing a coffee pot, we can see how these snapshots fit together. As your brain sees the new object, it selects the best fitting photograph from its record—in this case, a picture of another coffee pot. It then follows the strings from this stored picture to the audio tape, the smell of coffee and any other sense captured to identify the new object as a coffee pot. This is the concept of “linked sets.”

When activated, linksets activate related patterns and, due to the hierarchical nature, these connect to their original senses. This explains why the observation of a picture of your lost-long friend activates memories of where you last were and how you felt at the time. Those memories come from the activation of the linkset, stored by experience and then reactivated by pattern matching.

Sequences of actions, like eating in a restaurant, form sequential patterns through experience. As the formation of patoms ensures pattern atoms are the smallest pattern elements, indivisible, and unique: patoms are readily re-used. In sequential patterns, for example, the standard sequence for a meal may deviate by the waiter bringing along the bill before dessert. Our ability to detect the deviation need be nothing more complex than a pattern error being detected.

Patom theory expects that such deviations are commonplace, and as such, we learn continuously through life, although deviations reduce through experience.

Like snapshots, elements of a sequential pattern are bi-directional. By automatically linking together snapshot patterns based on continuity (sequence), logical objects are formed therefore eliminating the need to identify reduced or idealized object atoms. Successful object recognition relies on correctly adding snapshots to existing linksets. The pattern-atoms that enable object identification are therefore collections of snapshot patterns joined into linksets over time.

Equally, how these objects interact in the world are linked with other objects and actions through a higher-level patom. These higher level Patoms are learned tracking relationships that, in linguistics, are called selectional restrictions that help us identify the correct meanings for words in context.

Linkset patterns operate using the concept of *intersection*. By combining more than one linkset together, the weighted links will reduce the total active patterns to a subset in which the elements are fully consistent. There will typically be more than one consistent pattern, in which case the strongest links will determine the intersection’s best-fit.



The relationship between snapshots and sequences

Snapshot patterns include not only the sensor patterns, but also links to other sensor patterns when stored in higher layers. The difference between a snapshot and a sequence is that while snapshots are instantaneous patterns, sequences are the time-based collections of snapshot patterns. But as sequences are composed of lists of snapshots, any recognition of a part of the sequence identifies the whole as we experience as a small part of a song readily identifies the whole.

Snapshots rarely have much value without the time-based components. In language, for example, it is only at a specific point in time that a physical object is named. Linksets enable language by allowing a single association (a word) to be linked to numerous instances (snapshot patterns), effectively naming multiple instances in time with a single snapshot. While both the physical object and its name appear to be independent of time, the name may only be associated verbally for part of a second, despite hours of experience with the object. Yet our recall of the name for any of the object snapshots must be almost instantaneous. To stress the point, regardless of the shortness of object naming, immediate recall of this association is available through the linkset.

To illustrate the interaction, let's walk through the sequence of events necessary to name a person you meet. Snapshots are established based on concurrency and linksets are established based on continuity. When introduced to Robert and hearing the sound "robert", components of the images of Robert are stored as visual snapshots. These are grouped together in a linkset due to continuity. Elements of the sounds comprising the word "robert" are stored as snapshots, grouped together in an auditory linkset. Based on the concurrency of the two active linksets, at the layer above them, a snapshot connects the two sensory patterns. This connection effectively joins the two linksets. In other words, by association, the sounds comprising "robert" are joined with the visual images of Robert.

Linksets enable the concept of decision making through intersection. Intersection is the result of combining two or more linksets to establish the common ground between them. At a low level, intersection maintains consistency with the stored patterns involved. At a high level, intersection results in the most favorable or least unfavorable emotional outcome. The following example illustrates these concepts.

When viewing a known box, one of the previously stored snapshot patterns is matched. The snapshot signals the match to the layer above, which activates the linkset, or possibly many. The linkset signals forward and backward to its constituents including all the other snapshots comprising the box. These other snapshots are activated, in turn. One of the forward links may result in a link to the auditory linkset pattern for the word "box", effectively naming the object. This forward and backward linking is bidirectional pattern matching since either pattern identifies the other, albeit indirectly.

Our brain needs to do more than simply name objects, of course. More typically the brain's requirement is to direct motion based on numerous pattern inputs. Leaving aside the question of learning at this stage, let's consider two examples in which a brain needs to take action based on previously established patterns.

In the first example, let's contrast the use of relative patterns. When a decision must be made, the best choice is selected. If smell A, movement B, sight C, and touch D happen simultaneously, take action E. In isolation to all other stored patterns, this allows for A and B only to activate action E (Figure 2). As linksets

are bidirectional, the pattern E must identify A-D when receiving only one of the inputs. This bidirectional nature of linksets is invaluable to function in the world's complex sensory environment.

If the sensory pattern group above results in E, but a similar one with A, B, C plus F leads to G, how does the brain select which action to take when presented with just A? The concept of pattern *intersection* comes into play. Linksets connect using activating and inhibiting *weighted* links, not *binary* switches. Provided the A link is slightly stronger than E, in isolation A leads to E. In more complex intersections, the strongest links identified over a number of patterns dictate the action. If intersection fails to find a consistent pattern, no action is taken.

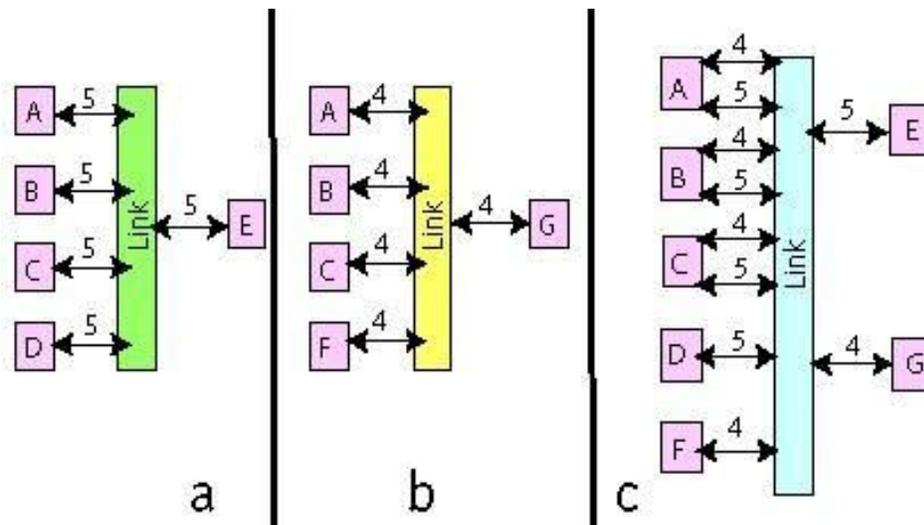


Figure 2. Examples of bi-directional "best fit" pattern matching. In example a, senses A, B, C, and D link to action E (strength 5). In example b, senses A, B, C and F link to action G (strength 4). In example c, imagine that A is the only sense active. This leads to E due to strength 5 ($5 > 4$). If A and F are active, G would be the action ($4 + 4 > 5$). In the second case, G would activate

all of A, B, C and F in reverse.

Now consider another example where ambiguity is obvious at one level, but not at another. Linguistics often makes the case easily.

"I am hear, no I mean here." If said out loud, the sentence would be perplexing as the spelling of hear in the first homonym, is wrong. But in written form, and with the benefit of the English writing system disambiguating words like these, the lack of meaning before the correction is evident. The patterns for "I am hear" are not valid English, because 'hear' is a verb and cannot be in that form in that position and equally, it makes no sense. The replacement "I am here" both makes sense as here is a location and it makes sense to identify my location.

If the patterns of language were simply, say, based on the parts of speech, simple sentences would make no sense, but remain valid. But when combinations of parts of speech, word agreement, tense, locations, times and other words and phrases related to the real world are experienced, they are communicated readily as the real world learned through experience eliminates invalid possibilities. In



languages, the combinations are usually exact, unlike in the decision-making example, with the real world in context able to determine the correct meaning.

Patom Theory and AI

The dominant AI paradigm considers the brain to be a parallel processor, encoding information in one area to transmit it to other areas of the brain for subsequent processing and construction. There has been little supporting evidence for this theory, despite its dominance since the 1950s. In the absence of a rival paradigm, the computational model has dominated even though examples of brain function are seemingly inexplicable by a processing model. Consider the following examples:

- The visual image of the word mind (almost completely obscured by an ink blot¹²) is readily identified despite the challenge to compute this. (The pattern is matched from our experience, despite insufficient visual information available to construct the image.)
- The false sighting of a glowing disk when a group of lines are drawn appropriately¹³. (A false pattern is detected because it matches a stored pattern.)
- The inability to recognize a face following a stroke, despite other functions being unimpaired¹⁴. (Visual patterns are lost.)
- The loss of the ability to reach correctly for a cup following parietal damage¹⁵. (Patterns are lost linking our sense of touch and visual experience.)
- The ability for children to learn a language, despite the lack of instruction to the child^{16,17,18}. (Language is learned due to the innate location of patoms providing common recognition of multi-sensory patterns to connect with auditory words.)
- The brain is comprised of similar materials, but each area of the brain processes information differently. (A patom's function includes using patterns. The ability to match words, recognize visual objects and recognize mathematical equations are pattern skills, rather than pre-programmed computational skills.)
- "The mind actively imposes organization on incoming information"¹⁹ is the Gestalt psychologist's view and the "whole is more than a collection of parts". (The brain's bi-directional patoms find and

¹² Pinker, S., *How the Mind Works*, 105 (Penguin Books, London, 1998).

¹³ Hoffman, D. D., *Visual Intelligence: How We Create What We See*, 111 (W. W. Norton & Company, New York, 1998).

¹⁴ Ibid. XI.

¹⁵ Gellatly A. & Zarate, O., *Mind & Brain for Beginners*, 114 (1998).

¹⁶ Maher, J. & Groves, J., *Chomsky for Beginners*, 47-51 (Penguin Books, London, 1996).

¹⁷ Chomsky, N., *Review of Verbal Behaviour by B.F. Skinner*. Language Volume 35, Number 1, 42 (1959).

¹⁸ Chomsky, N., *Recent Contributions to the Theory of Innate Ideas*. Synthese 17, 122-123 (1967).

¹⁹ Wessells, M. G., Op. Cit. 20-21.



combine matches in stored sensory information. This effectively results in a bias to use existing patterns as the basis of understanding - imposing organization on experience).

Patom Theory explains our brain's capability in cases like these that cannot be explained by processing models. There is a growing body of evidence supporting the existence of a Patom construct. The observable effects of brain damage variations resulting in cell loss suggests that localized brain areas are collectively capable of capturing specific patterns. Damasio (1994), for example, describes effects resulting from losing the six layers of cortical cells. Greenfield (1997), McCrone (1999) and others popularize this understanding through the visual presentation of graphic brain scan techniques. The question of whether the brain is modular has moved to the question of how many modules there are in the brain (Jacobs, Jordan and Barto, 1991).

Innate Language Acquisition

Patom Theory explains language acquisition as an automatic outcome if patoms are appropriately positioned. Braine's explanation of how to create syntax from only basic semantic representation is another way to put this²⁰. Critical patterns used in language include its symbols (letters/characters) plus its grammar, phonemes, intonation, syntax, pitch and intensity²¹.

Neuroscientists show that particular brain locations control specific parts of the body²². When lip reading, for example, the most active brain's areas include auditory patterns, visual face patterns, mouth movement and language²⁶. AI researchers may argue this is due to a clever processing approach, not yet understood, overseen by a controller, but there is no controller or programmer. Patom Theory suggests that lip readers use the learned linkages of patterns that connect the language areas. By learning a language (and watching the speaker's lips), Patoms automatically store the visual patterns used in lip-reading.

For a person to learn a language, patterns must be acquired in the language areas identified, starting of course with sensory patterns leading to Braine's semantic representation. As sensory patterns are recognized, they are joined together by the "name patom". Language is the total effect resulting from the collection of patterns. There are areas in the brain identified for multi-sensory integration, event storage, emotion and so forth. Patom Theory requires mere connections between appropriate patoms to link language to a complex sensory pattern or event in contrast to a computer's encoding or other general abstraction. When presented with incomplete experiences such as when lip reading, the brain uses the existing learned bi-directional links to complete the incomplete experience "by inference". Inference is the use of learned links representing the best fit available, in which available information is otherwise insufficient.

²⁰ Braine, MDS, *What sort of innate structure is needed to "bootstrap" into syntax*, Cognition 45 77-100, 1992

²¹ Fromkin, V. and Rodman, R., *An Introduction to Language*, 28-29 (CBS Publishing, Japan, 1983).

²² Hinton, G. E., *How Neural Networks Learn from Experience*. Scientific American Volume 267, Number 3, 104-109 (1992).



Generalized Information Storage

A curious aspect of brain damage can be the loss of specific categories of knowledge²³, such as a farmer being unable to recognize his sheep, while keeping his ability to recognize his family²⁴. Brain damage resulting in the specific loss of item categories is predictable by patom theory.

As in the visual example, AI researchers start with a general example, such as a face, and move to find a specific example, such as John Smith's face. AI concepts like *frames* attempt to specify the general information structure on which to base specific cases²⁵ by identifying the types of information included with a particular experience in advance.

Patom Theory uses the opposite approach. It moves from the specific examples experienced (a matched pattern), to link to the general case (typically using language, an auditory pattern). Of importance is that the “specific defining the general” eliminates the problems caused where (a) the general case cannot be defined and (b) the need for a programmer, while the “general defining the specific” has been the requirement for AI to date.

Learning and memory

Learning takes place as patterns are stored (in patoms). When learning mathematics, for example, children learn to recognize numbers visually and to link these patterns to other sensory patterns, such as auditory and tactile patterns. They also learn the visual symbols used to manipulate the numbers, and the entire visual pattern that provides an answer.

Addition, such as $2 + 2 = 4$ is learned, not computed. A child's education includes rote learning, storing patterns by repetition. Patom structure provides multi-level learning, storing sensory patterns first. Then these are linked at higher levels to identify additional patterns.

When learning higher level patterns, such as how to manipulate 2346×124 , they learn the steps to get to the result. These steps are memorized patterns. Mathematics may be a visual sequential pattern, similar in concept to language's grammar - linking visually recognized numbers and operators with meaning associated to other senses.

Emotions

The human brain's emotion center, located in the limbic system, is amazingly complex. The system can bias our attention (direct the patterns likely to be used) and change our body's chemistry through the release of hormones²⁶.

Psychologists are very interested in the patterns that link to these emotions, as meaning typically involves patterns linked to emotion. A new industry, with practitioners like Anthony Robbins, takes

²³ Carter, R., Op. Cit., 190.

²⁴ Greedfield, S. (editor), *The Human Mind Explained: An Owner's Guide to the Mysteries of the Mind*, 25 (Reader's Digest, Sydney, 1996).

²⁵ Stillings, N. A. et al, *Cognitive Science: An Introduction*, 129-131 (Bradford Book, Cambridge MA, 1987).

²⁶ Carter, R., Op. Cit., 129-131.



advantage of the emotional brain's bi-directional nature²⁷. By acting "as if they can" (imitating the behaviors of role models) students can be more effective in their activities. Patom Theory predicts the simulated experience is automatically stored in innate patoms for future use (the actions involved in pleasure and fear in humans is innate, not learned). The brain's bi-directional nature provides future benefit by activating these patterns when acting out.

Intelligence

Some experts in education refer to the word *intelligence* as including emotional, linguistic, musical, spatial, bodily-kinesthetic and visual proficiency. As each of these requires specific pattern storage, matching and use, intelligence is a measure of the effectiveness of patoms in their role regardless of which senses are involved.

AI researchers seeking to replicate human intelligence should consider Patom Theory. Intelligence is a measure of pattern-matching prowess in the areas identified.

Discussion

How a patom works at the physical level in biological brains is still to be understood, and yet the benefits in using it to model macro level brain functions are unmistakable. In evolutionary terms, the earliest parts of our brain control automatic functions, such as temperature control (simple pattern use skills). Patom Theory provides plausible explanations for the brain's function where none currently exists.

Opportunities exist to disprove the theory in most of the cognitive sciences and the ultimate success will depend on continued consistency in the many related disciplines.

One of the best targets for immediate research is in language. AI researchers have struggled for some time to program computers to use word sequences and grammar, but with limited success. Languages are patterns of sensory patterns (meaning), not simply patterns of sounds (words).

Without sensory pattern devices, the human basis for language is simply not available in machine form today, however patom theory allows for the replication of the patterns at any level. By removing the sensory input itself, and replacing it with higher level patterns, the same type of language understanding in humans can be achieved on computers. Computer-based sensory deprivation and the lack of artificial patoms and associated connection technology limits progress, but that shouldn't stop us producing useful language machines today.

Evidence from language manipulation simulations for conversation and translation support patom theory by dealing with existing open scientific problems. As brains deal with a rich world of ambiguity, the ability to disambiguate in languages is important. The open problems of disambiguation include word sense disambiguation (WSD), word-boundary identification (both in written and spoken forms) and the tracking of

²⁷ Robbins, A., *Awaken the Giant Within How to take immediate control of your mental, emotional, physical & financial destiny*, (Fireside Book, New York, 1991).



context in conversation with the use of pronouns. These all yield to the breakdown of language to smaller components for manipulation by linkset patterns, based on meaning coming from brain emulation.

WSD is addressed by connecting patterns based on experience, and then including only those that are valid after linkset intersection. Word boundary identification comes from testing all possible patterns allowed by the received phoneme sequences, but excluding all those that cannot form words, and then excluding those that cannot form valid phrases (because the resulting phrases are meaningless). Context tracking, once sentences have disambiguated words, is not complex as it uses disambiguated patterns as input.

The current focus on the brain as a processing machine has hindered our efforts to develop true machine intelligence. Patom theory, with its pattern-matching approach provides a different paradigm and opens up new avenues of development in AI—true machine intelligence.