# Glocal Memory: A New Perspective on Knowledge Representation, Neurodynamics, Distributed Cognition, and the Nature of Mind

Ben Goertzel

Novamente LLC, 1405 Bernerd Place, Rockville MD 20851 {ben@goertzel.org}

**Abstract.** We describe a novel conceptual and formal model of memory structure that combines key aspects of global and local knowledge representation. Applications of this "glocal memory" model to artificial general intelligence are discussed, in the context of the "Novamente Cognition Engine" and OpenCog software systems, and simpler prototype systems constituting "glocal Hopfield nets." Building on recent results regarding visual memory derived from single-neuron recordings, an hypothesis is made regarding how some instances of glocal memory might be achieved in the human brain. It is also argued that glocal memory models the concept of "distributed cognition" according to which an individual human mind is both localized in a brain and distributed through a network of interactions with tools and other embodied minds.

## 1 Introduction

"Memory," considered in a general sense, is central to intelligence and to complex systems in general. Without the ability to reflect its past experiences in its present structure, a system cannot adapt, and hence cannot manifest either intelligence or complexity.

However, the conceptual tools currently available for modeling and analyzing memory structures and dynamics (on the conceptual or formal level) are disappointingly simplistic. In particular, popular approaches to understanding memory tend to focus either on local memory (in which a memory is stored in one place within a system) or global memory (in which a memory is stored as some sort of pattern of activation distributed across a system). My suggestion is that neither global nor local memory models are sufficiently subtle or flexible to capture the way memory works in complex intelligent systems like human brains, nor the way memory should work in AI systems if these systems are to demonstrate robust intelligence. Here I introduce a new concept of *glocal memory*, intended specifically to describe certain sorts of memory structures that combine local and global aspects. The basic notion of glocal memory is not entirely new; for instance, the hypothesis of (what is here called) glocal memory in the human cortex was described in detail in my1997 book *From Complexity to Creativity*. However, the glocal memory concept has not previously been named, explicitly formalized, nor discussed in such a general context as is done here.

The central idea of glocal memory is that declarative, episodic or procedural items may be stored in memory in the form of paired structures that are called (key, map) pairs. The key is a localized version of the item and records some significant aspects of the items in a simple and crisp way. The map is a dispersed, distributed version of the item which, represents the item as a (to some extent, dynamically shifting) combination of fragments of other items. The map includes the key as a subset; activation of the key generally (but not necessarily always) causes activation of the map; and changes in the memory item will generally involve complexly coordinated changes on the key and map level both.

After presenting a simple formalization of the glocal memory concept, applications of the idea to various aspects of human and artificial intelligence are discussed. The hypothesis of glocal memory in the human brain, from (Goertzel, 1997), is updated in the light of more recent neuroscience thinking, and some implications for the design of formal neural networks are considered. Next, the manifestation of glocal memory in AI is reviewed in the context of the Novamente Cognition Engine (Goertzel, 2006) and OpenCog (Goertzel, 2008) software systems. Finally, attention is then paid to the application of glocal memory to notions of "distributed cognition" (in which the mind of a human or other embodied organism is considered as extending beyond its body into those aspects of the world that regularly and intricately interact with its body).

## 2 A Simple Formalization of Glocal Memory

To explain the notion of glocal memory more precisely, we will introduce a simple formal model of a system S that uses a memory to record information relevant to the actions it carries out. The overall concept of glocal memory should not be considered as restricted to this particular formal model. The formal model is not intended for maximal generality, but is intended to encompass a variety of current AI system designs and formal neurological models.

In this model, we will consider S's memory subsystem as a set of objects we'll call "tokens," embedded in some metric space. The metric in the space, which we will call the "basic distance" of the memory, generally will not be defined in terms of the semantics of the items stored in the memory; though it may come to shape these dynamics through the specific architecture and evolution of the memory. Note that these tokens are not intended as generally being mapped one-to-one onto meaningful items stored in the memory. The "tokens" are the raw materials that the memory uses to store items.

We assume that each token, at each point in time, may meaningfully be assigned a certain quantitative "activation level." Also, tokens may have other numerical or discrete quantities associated with them, depending on the particular memory architecture. Finally, tokens may relate other tokens, so that optionally a token may come equipped with an (ordered or unordered) list of other tokens.

To understand the meaning of the activation levels, one should think about S's memory subsystem as being coupled with an action-selection subsystem, that dynamically chooses the actions to be taken by the overall system in which the two subsystems are embedded. Each combination of actions, in each particular type of context, will generally be associated with the activation of certain tokens in memory.

Then, as analysts of the system S, we may associate each token T with an "activation vector" v(T,t), whose value for time t consists of the activation of the token T at time t.

"Items stored in memory" over a certain period of time, may then be defined as clusters in the set of activation vectors associated with memory during that period of time. Note that the system S itself may explicitly recognize and remember patterns regarding what items are stored in its memory – but, from an external analyst's perspective, the set of items in S's memory is not restricted to the ones that S has explicitly recognized as memory items.

The "localization" of a memory item may be defined as the degree to which the various tokens involved in the item are close to each other according to the metric in the memory metric-space. This degree may be formalized in various ways, but choosing a particular quantitative measure is not important here. A highly localized item may be called "local" and a not-very-localized item may be called "global."

We may define the "activation distance" of two tokens as the distance between their activation vectors. We may then say that a memory is "well aligned" to the extent that there is a correlation between the activation distance of tokens, and the basic distance of the memory metric-space.

Given the above set-up, the basic notion of glocal memory can be enounced fairly simply. A glocal memory is one:

- That is reasonably well-aligned (i.e. the correlation between activation and basic distance is significantly greater than random)
- In which most memory items come in pairs, consisting of one local item and one global item, so that activation of the local item (the "key") frequently leads in the near future to activation of the global item (the "map")

Obviously, in the scope of all possible memory structures constructible within the above formalism, glocal memories are going to be very rare and special. But, I suggest that they are important, because they are generally going to be the most effective way for intelligent systems to structure their memories.

An example of a predominantly local memory structure, in which nearly all significant memory items are local according to the above definition, is the Cyc logical reasoning engine (Lenat and Guha, 1990). To cast the Cyc knowledge base in the present formal model, the tokens are logical predicates. Cyc does not have an in-built notion of activation, but one may conceive the activation of a logical formula in Cyc as the degree to which the formula is used in reasoning or query processing at a certain point in time. And one may define a basic metric for Cyc by associating a predicate with its extension, and definiting the similarity of two predicates as the symmetric distance of their extensions. Cyc is reasonably well-aligned, but according to the dynamics of its querying and reasoning engines, it is basically a local memory structure without significant global memory structure.

On the other hand, an example of a predominantly global memory structure, in which nearly all significant memory items are global according to the above definition, is the Hopfield associative memory network (Amit, 1989). Here memories are stored in the pattern of weights associated with synapses within a network of formal neurons, and each memory in general involves a large number of the neurons in the network. To cast the Hopfield net in the present formal model, the tokens are neurons and synapses; the activations are neural net activations; the basic distance between two neurons A and B may be defined as the percentage of the time that stimulating one of the neurons leads to the other one firing; and to calculate a basic distance involving a synapse, one may associate the synapse with its source and target neurons. With these definitions, a Hopfield network is a well-aligned memory, and (by intentional construction) a markedly global one. Local memory items will be very rare in a Hopfield net.

While predominantly local and predominantly global memories may have great value for particular applications, my suggestion is that they also have inherent limitations. If so, it means that the most useful memories are going to be those that involve both local and global memory items in central roles. However, this is a more general and less risky claim than the assertion that glocal memory structure as defined above is important. Because, "glocal" as defined above doesn't just mean "neither predominantly global nor predominantly local." Rather, it refers to a specific pattern of coordination between local and global memory items – what I have called the "keys and maps" pattern.

## 3 Hints of Glocal Memory in the Human Brain

Our understanding of human brain dynamics is still very primitive, one manifestation of which is the fact that we really don't understand how the brain represents knowledge, except in some very simple respects. So anything anyone says about knowledge representation in the brain, at this stage, has to be considered highly speculative. Existing neuroscience knowledge does imply constraints on how knowledge representation in the brain may work, but these are relatively loose constraints. These constraints do imply that, for instance, the brain is neither a relational database (in which information is stored in a wholly localized manner) nor a collection of "grandmother neurons" that respond individually to high-level percepts or concepts; nor a simple Hopfield type neural net (in which all memories are attractors globally distributed across the whole network). But they don't tell us nearly enough to, for instance, create a formal neural net model that can confidently be said to represent knowledge in the manner of the human brain.

As an initial example of the current state of knowledge, I'll discuss here a series of papers regarding the neural representation of visual stimuli (Quinoga et al, 2005; 2008), which deal with the fascinating discovery of a subset of neurons in the medial temporal lobe (MTL) that are selectively activated by strikingly different pictures of given individuals, landmarks or objects, and in some cases even by letter strings. For instance, in the 2005 paper, it is noted that

> in one case, a unit responded only to three completely different images of the

> ex-president Bill Clinton. Another unit (from a different patient) responded only to images of The Beatles, another one to cartoons from The Simpson's television series and another one to pictures of the basketball player Michael Jordan.

These empirical results seem quite clear and exciting, yet the authors' theoretical interpretation of the data has consistently been much less so. In the 2005 abstract, they note that their results "suggest that neurons might encode an abstract representation of an individual." And indeed, the title of the 2005 paper is the rather gutsily worded "Invariant visual representation by single neurons in the human brain." Yet in the paper's conclusion the authors tell a somewhat more conservative story:

> How neurons encode different percepts is one of the most intriguing questions in neuroscience. Two extreme hypotheses are schemes based on the explicit representations by highly selective (cardinal, gnostic or grandmother) neurons and schemes that rely on an implicit representation over a very broad and distributed population of neurons. In the latter case, recognition would require the simultaneous activation of a large number of cells and therefore we would expect each cell to respond to many pictures with similar basic features. This is in contrast to the sparse firing we observe, because most MTL cells do not respond to the great majority of images seen by the patient. Furthermore, cells signal a particular individual or object in an explicit manner[27], in the sense that the presence of the individual can, in principle, be reliably decoded from a very small number of neurons.We do not mean to imply the existence of single neurons coding uniquely for discrete percepts for several reasons: first, some of these units responded to pictures of more than one individual or object; second, given the limited duration of our recording sessions, we can only explore a tiny portion of stimulus space; and third, the fact that we can discover in this short time some images—such as photographs of Jennifer Aniston—that drive the cells suggests that each cell might represent more than one class of images. Yet, this subset of MTL cells is selectively activated by different views of individuals, landmarks, animals or objects. This is quite distinct from a completely distributed population code and suggests a sparse, explicit and invariant encoding of visual percepts in MTL.

It seems that the alternate title "Invariant visual representation by sparse neuronal subnetworks in the human brain," might have better captured their actual conclusions – and yet, this weaker title would not have communicated the exciting nature of some of their individual findings, such as the subject who apparently had "Bill Clinton" neurons which did not fire in response to other test images (though obviously this doesn't rule out that those neurons might have fired in a variety of other circumstances, which indeed I suspect would be the case).

The 2008 paper backed away from the more extreme interpretation in the title as well as the conclusion, with the title "Sparse but not "Grandmother-cell" coding in the medial temporal lobe." As the authors emphasize there,

> Given the very sparse and abstract representation of visual information by these neurons, they could in principle be considered as ‘grandmother cells’. However, we give several arguments that make such an extreme interpretation unlikely.
>
> ...
>
> MTL neurons are situated at the juncture of transformation of percepts into constructs that can be consciously recollected. These cells respond to percepts rather than to the detailed information falling on the retina. Thus, their activity reflects the full transformation that visual information undergoes through the ventral pathway. A crucial aspect of this transformation is the complementary development of both selectivity and invariance. The evidence presented here, obtained from recordings of single-neuron activity in humans, suggests that a subset of MTL neurons possesses a striking invariant representation for consciously perceived objects, responding to abstract concepts rather than more basic metric details. This representation is sparse, in the sense that responsive neurons fire only to very few stimuli (and are mostly silent except for theirpreferred stimuli), but it is far from a Grandmother-cell representation. The fact that the MTL represents conscious abstract information in such a sparse and invariant way is consistent with its prominent role in the consolidation of long-term semantic memories.

It’s interesting to note how inadequate the Quinoga et al data really is for exploring the notion of glocal memory in the brain. Suppose it’s the case that individual visual memories correspond to keys consisting of small neuronal subnetworks, and maps consisting of larger neuronal subnetworks. Then it would be not at all surprising if neurons in the “key” network corresponding to a visual concept like “Bill Clinton’s face” would be found to respond differentially to the presentation of appropriate images. Yet, it would also be wrong to overinterpret such data as implying that the key network somehow comprises the “representation” of Bill Clinton’s face in the individual’s brain. In fact this key network would comprise only one aspect of said representation.

In the glocal memory hypothesis, a visual memory like “Bill Clinton’s face” would be hypothesized to correspond to an attractor spanning a significant subnetwork of the individual’s brain – but this subnetwork still might occupy only a small fraction of the neurons in the brain (say, 1/100 or less), since there are very many neurons available. This attractor would constitute the map. But then, there would be a much smaller number of neurons serving as key to unlock this map: i.e. if a few of these key neurons were stimulated, then the overall attractor pattern in the map as a whole would unfold and come to play a significant role in the overall brain activity landscape. In prior publications (e.g. Goertzel, 1997) I have explored this hypothesis in more detail in terms of the known architecture of the cortex and the mathematics of complex dynamical attractors.

So, one possible interpretation of the Quinoga et al data is that the MTL neurons they’re measuring are part of key networks that correspond to broader map networks recording percepts. The map networks might then extend more broadly throughout the brain, beyond the MTL and into other perceptual and cognitive areas of cortex. Furthermore, in this case, if some MTL key neurons were removed, the maps might well regenerate the missing keys (as would happen e.g. in the glocal Hopfield model to be discussed in the following section).

Related, interesting evidence for glocal memory in the brain comes from a recent study of semantic memory (Patterson et al, 2007), which probed the architecture of semantic memory via comparing patients suffering from semantic dementia (SD) with patients suffering from three other neuropathologies, and found reasonably convincing evidence for what they call a “distributed-plus-hub” view of memory.

The SD patients they studied displayed highly distinctive symptomology; for instance, their vocabularies and knowledge of the properties of everyday objects were strongly impaired, whereas their memories of recent events and other cognitive capacities remain perfectly intact. And these patients also showed highly distinctive patterns of brain damage: focal brain lesions in their anterior temporal lobes (ATL), unlike the other patients who had either less severe or more widely distributed damage in their ATLs. This led Patterson et al to conclude that the ATL (which is adjacent to the amygdala and limbic systems, which process reward and emotion; and the anterior parts of the medial temporal lobe memory system, which processes episodic memory) is a “hub” for amodal semantic memory, drawing general semantic information from episodic memories based on emotional salience.

So, in this view, the memory of something like a “banana” would contain a distributed aspect, spanning multiple brain systems, and also a localized aspect, centralized in the ATL. The distributed aspect would likely contain information on various particular aspects of bananas, including their sights, smells, and touches, the emotions the evoke, and the goals and motivations they related to. The distributed and localized aspects would influence each other dynamically, but, the data Patterson et al gathered do not address dynamics and they don’t venture hypotheses in this direction.

There is a relationship between the “distributed-plus-hub” view and Damasio's better-known notion of a “convergence zone” (Damasio, 2000), defined roughly as a location where the brain binds features together. A convergence zone, in Damasio’s perspective, is not a "store" of information but an agent capable of decoding a signal (of reconstructing information). He also uses the metaphor that convergence zones behave like indexes drawing information from other areas of the brain – but they are not static but rather dynamic indices, containing the instructions needed to recognize and combine the features constituting the memory of something. The mechanism involved in the distributed-plus-hub model is similar to a convergence zone, but with the important difference that hubs are less local: Patterson et al’s semantic hub may be thought of a kind of “cluster of convergence zones” consisting of a network of convergence zones for various semantic memories.

What is missing in Patterson’s and Damasio’s perspective is a vision of distributed memories as attractors. The idea of localized memories serving as indices into distributed knowledge stores is important, but is only half the picture of glocal memory: the creative, constructive, dynamical-attractor aspect of the distributed representation is the other half. The closest thing to a clear depiction of this aspect of glocal memory that seems to exist in the neuroscience literature is a portion of William Calvin's theory of the "cerebral code" (Calvin, 1996). Calvin proposes a set of quite specific mechanisms by which knowledge may be represented in the brain using complexly-structured strange attractors, and by which these strange attractors may be propagated throughout the brain. Calvin explores in great detail how a distributed attractor may propagate from one part of the brain to another in pieces, with one portion of the attractor getting propagated first, and then seeding the formation in the destination brain region of a close approximation of the whole attractor.

Calvin's theory may be considered a genuinely glocal theory of memory. However, it also makes a large number of other specific commitments that are not part of the notion of glocality, such as his proposal of hexagonal meta-columns in the cortex, and his commitment to evolutionary learning as the primary driver of neural knowledge creation. We find these other hypotheses interesting and highly promising, yet feel it is also important to separate out the notion of glocal memory for separate consideration.

Regarding specifics, our suggestion is that Calvin's approach may overemphasize the distributed aspect of memory, not giving sufficient due to the relatively localized aspect as accounted for in the Quinoga et al results discussed above. In Calvin's glocal approach, global memories are attractors and local memories are parts of attractors. We suggest a possible alternative, in which global memories are attractors and local memories are particular neuronal subnetworks such as the specialized ones identified by Quinoga et al. However, this alternative does not seem contradictory to Calvin's overall conceptual approach, even though it is different from the particular proposals made in (Calvin, 1996).

The above paragraphs are far from a complete survey of the relevant neuroscience literature; there are literally dozens of studies one could survey pointing toward the glocality of various sorts of human memory. Yet experimental neuroscience tools are still relatively primitive, and every one of these studies could be interpreted in various other ways. In the next couple decades, as neuroscience tools improve in accuracy, our understanding of the role of glocality in human memory will doubtless improve tremendously.

## 4 Glocal Hopfield Nets

Following up on the ideas of the previous section, it is interesting to explore the notion of formal neural network models that embody the notion of glocal memory. My colleagues and I have recently run some interesting simulations with a variation on Hopfield netural nets that explicitly incorporates the notion of glocality. Our technical results will be reported elsewhere, but, a brief discussion of the main ideas would seem appropriate here.

Essentially, we augment the standard Hopfield net architecture by adding a set of "key neurons." These are a small percentage of the neurons in the network, and are intended to be roughly equinumerous to the number of memories the network is supposed to store. When the Hopfield net converges to an attractor A, then new links are created between the neurons that are active in A, and one of the key neurons. Which key neuron is chosen? The one that, when it is stimulated, gives rise to an attractor pattern maximally similar to A.

The ultimate result of this is that, in addition to the distributed memory of attractors in the Hopfield net, one has a set of key neurons that in effect index the attractors. Each attractor corresponds to a single key neuron. In the glocal memory model, the key neurons are the keys and the Hopfield net attractors are the maps.

This algorithm has been tested in sparse Hopfield nets, using both standard Hopfield net learning rules and Storkey's modified palimpsest learning rule (Storkey and Valabregue, 1999), which provides greater memory capacity in a continuous learning context. The use of key neurons turns out to slightly increase Hopfield net memory capacity, but this isn't the main point. The main point is that one now has a local representation of each global memory, so that if one wants to create a link between the memory and something else, it's extremely easy to do so – one just needs to link to the corresponding key neuron. Or, rather, one of the corresponding key neurons: depending on how many key neurons are allocated, one might end up with a number of key neurons corresponding to each memory, not just one.

In spite of their considerable theoretical power, Hopfield nets are not particularly useful for practical applications on von Neumann computer hardware (appropriately inexpensive massively parallel computer hardware would be another story, but that's not the direction the computer industry has taken), so the above-described experiments with glocal Hopfield nets were conducted with a view toward intellectual exploration – in order to understand the possible nature of glocal memory in the brain via a concrete computational model; and in order to provide a simple prototype domain for experimenting with related ideas in the more complex context of integrative AGI systems such as those discussed in the following section.

## 5 Glocal Memory in Integrative AGI Systems

One of the main motivations for the development of the glocal memory concept has been the design of artificial memories, which is a task different in many ways from the analysis of modeling of naturally occurring memories. In our work on the Novamente Cognition Engine (Goertzel, 1996) and OpenCog (Goertzel, 2008) AI systems, my colleagues and I have been motivated by the glocal memory concept to design memory approaches that are explicitly glocal in nature.

The glocality concept hits straight at the center of one of the biggest debates of theoretical AI: symbolic versus subsymbolic knowledge representation. This dichotomy is often discussed but rarely drawn in a formal and rigorous way, and I have argued elsewhere that it is actually a largely bogus dichotomy (Goertzel et al, 2008). Traditionally, logic-based AI systems are viewed as "symbolic", and neural net systems are viewed as "subsymbolic." But this distinction has gotten fuzzier and fuzzier in recent years, with developments such as

- logic-based systems being used to control embodied agents (hence using logical terms to deal with data that is apparently perception or actuation-oriented in nature, rather than being symbolic in the semiotic sense), see (Santore and Shapiro, 2003; Goertzel et al, 2008)
- hybrid systems combining neural net and logical parts, or using logical or neural net components interchangeably in the same role (Lebiere and Anderson, in preparation)
- neural net systems being used for strongly symbolic tasks such as automated grammar learning (Elman, 1991 plus a great deal of more recent work)

In my own AI systems referenced above, I have explicitly sought to span the symbolic/subsymbolic pseudo-dichotomy, via creating integrative systems that combine logic-based aspects with neural-net-like aspects, not in the manner of multimodular systems, but via attaching uncertain-logical truth values and neural-net-like weight and activation values to the same nodes and links in a knowledge-representation hypergraph. Furthermore, both the logical and neural-net-like features are used to handle all sorts of knowledge, from the most concrete perception and actuation related knowledge to the most abstract relationships. The concept of glocality lies at the heart of this combination, in a way that spans the pseudo-dichotomy:

- Local knowledge is represented in abstract logical relationships stored in explicit logical form, and also in Hebbian-type associations between nodes and links
- Global knowledge is represented in large-scale patterns of node and link weights, which lead to large-scale patterns of network activity, which

often take the form of attractors qualitatively similar to Hopfield net attractors

The data-store of nodes and links is acted on by a variety of cognitive processes, encapsulated in software objects called MindAgents. Some MindAgents work together to carry out probabilistic logical reasoning according to the mathematics given in (Goertzel et al, 2008); others spread neural-net-like weights and activation values (called importance values) according to equations based on artificial economics but somewhat similar in nature to neural net spreading activation equations. The attractors of this nonlinear activation spreading process constitute global memories; and there are then explicit MapEncapsulation MindAgents that identify these attractors and build "key nodes" corresponding to them, similarly to in the glocal Hopfield net described above. The logical inference and activation spreading processes feed off each other in particular ways, so that the formation and maintenance of the glocal memory is a result of the integrated behavior of the system's multiple cognitive dynamics.

The result of all this is that a concept like "cat" might be represented as a combination of

- a small number of logical relationships and strong associations, that constitute the "key" subnetwork for the "cat" concept
- a large network of weak associations, binding together various nodes and links of various types and various levels of abstraction, representing the "cat map"

The activation of the key will generally cause the activation of the map ... and the activation of a significant percentage of the map will cause the activation of the rest of the map, including the key. Furthermore, if the key were for some reason forgotten, then after a significant amount of effort, the system would likely to be able to reconstitute it (perhaps with various small changes) from the information in the map. I suspect that this particular kind of glocal memory will turn out to be very powerful for AI, due to its ability to combine the strengths of formal logical inference with those of self-organizing attractor neural networks.

## 6 Socio-culturo-technological Glocality and Distributed Cognition

So far I have talked about glocal memory within particular, physically localized brains, or within particular, delimited AI systems with coherent self/identity structures. Now I will extend the concept further, making use of the notion of "distributed cognition" in its broadest sense. A number of theorists (e.g. Hutchins, 1995; Perry, 2003) have argued that the human mind is not really contained in a single brain and body – but that, in fact, each individual person is best conceived as a pattern of activation across a sociocultural network, and across a subset of the physical world including e.g. the tools that the body associated with the mind habitually uses.

A paradigm case for distributed cognition would be a large boat like an aircraft carrier and the crew on it. It is not the isolated, encapsulated mind of any one single person or machine that is important for the successful operation of the boat It is the cognition that is distributed over the personnel, sensors, and machinery both on the boat and, to a lesser extent, in the various other machines interacting with the boat, such as airplanes or port-based communication centers.

As Mike Tintner pointed out to me in conversation, one interesting example of glocal memory in distributed cognition is military strategy, in which a single command by a single commander can trigger actions by a vast number of people with huge real-world consequences. This leads to the possibility of powerful glocal memory processes.

To consider the case of a war machine focused largely on global memory, consider the Russian army as depicted in War and Peace, where General Kutuzov proposes to essentially let the Russian army self-organize into its own context-appropriate battle patterns, rather than providing any kind of detailed top-down control. Loosely speaking this is a sort of large-scale guerilla warfare. The knowledge is in the whole – in the common sense of the common soldiers, not as individuals but as groups.

On the other hand, the opposite would be something like Operation Desert Storm, which was carefully orchestrated and planned (in spite of some errors e.g. deaths by friendly fire), so that the individual actors were largely doing what the software told them to do, based on the programming created by its programmers under the strategic guidance of the military leaders. Here the knowledge of strategy is localized, and from a strategic perspective, the individual humans carrying out the strategy are acting more like sensors and actuators rather than cognizers (though from their point of view as individual human actors, they are of course still carrying out cognition).

In Kutuzov's battles, the knowledge of the military plan was contained in the army as a whole; in Desert Storm, the knowledge was contained in the central planning software and the minds of the relevant military leaders (hence localized from the view of the whole army). In a glocal approach, on the other hand, both central planning and distributed, self-organized activity would be highly refined and productive; and they would be coordinated together dynamically and effectively. And this is in fact the sort of thing one hears the US military talking about these days: it realizes it needs to achieve in future in order to combine rapid, flexible adaptivity with global coordination.

A related example would be live performance of improvisatory music. In this case, there is an interaction between the performers, the audience and the instrument. A "tight band" accustomed to improvising together is a great example of distributed cognition and its intersection with glocal memory: the knowledge of each song lives within the individual minds of the performers, and also within the collective mind of the group. Replace one of the performers, and the global knowledge will pervade his mind somewhat, but his individuality will also have an effect, causing a shift in the nature of the music.

Consider first the case of an individual performer, improvising on the piano. The performer may have certain ideas, patterns and tastes in his mind – say, a habit of playing a melodic line that increasingly deviates from the scale implicit in the chords he's playing, until it eventually gets so far away from the original scale that the connection is impossible to detect ... at which point he brings the melody back home to the original scale. This habit may take a certain form in his mind, but it may take a quite different form when he sits down in front of the keyboard, because the feedback from actually hearing the music played makes him listen and play differently. And then, when he plays for an audience, the habit may shape itself quite differently than when he's playing for himself, because he gets feedback from the audience regarding when *they* think the melodic line has deviated too far from the original scale, based on the looks on their faces, their body language and so forth.

But next, introduce another musician. Suppose there is a saxophonist improvising along with the pianist. Then things get subtler, because the saxophonist may choose to intervene anywhere in the course of the pianist's improvisation, nudging the pianist back toward the original scale or further away from it. In this case, the knowledge of how to construct the melodic line is both local within the mind/body/tool combination of the pianist, and global within the system consisting of the pianist/piano and saxophonist/saxophone. This kind of dynamic is easy to hear, for example, in the interplay between John Coltrane and McCoy Tyner, especially in live recordings. This kind of dynamic exists both generically within a band, and specifically within a band's approach to an individual song, thus yielding a quite refined and detailed distributed multi-person glocal memory.

This sort of qualitative analysis doesn't prove anything scientifically, of course. And, rigorously demonstrating things about the relationship between locally embodied and environmentally distributed mind is going to be particularly difficult, since it requires accurately studying patterns both in the locally embodied mind, and in the body's interaction with the external physical and social environment. Designing experiments to test the hypothesis of the "socio-culturo-technologically glocal mind" is an important challenge.

## 7 Conclusion

The concept of glocal memory is simple but as I have argued here has quite broad scope. It may seem perplexing that such a simple and natural concept has not been extensively popularized and applied already, but I attribute this to the cognitive preference of human theorists for maximally simple models, rather than to any problem with the idea itself. Purely local and purely global models have a philosophical and analytic simplicity to them that glocality lacks. However, I have sought to show here that the glocal memory idea has the potential to cast clarity on a variety of issues on a variety of levels. The correct formulation of Occam's razor is Einstein's aphorism "As simple as possible, but no simpler," and I suggest that glocal memory fits this description. In considering human and AGI memory, I suggest that purely local and purely global models are "simpler than possible" whereas the glocal model is not.

Specifically, we can look at an individual embodied mind as consisting of two different glocal-memory feedback loops: one where the keys and maps both lie within the local knowledge base associated with the mind (the brain, in the case of human intelligence; or the RAM and disk of the host computers, in the case of software intelligence); and one where this whole local knowledge base is the key, and the map lies in the interactions of the individual's body with tools, other individuals, and other aspects of the world around it.

Each of the angles on glocal memory described in the preceding sections may be elaborated in its own, disciplinarily appropriate way. In the AI case, the challenge is to use glocal memory as the basis for a pragmatically useful AGI system. In the neuroscience case, the challenge is to create and then experimentally explore more fine-grained biological and biopsychological hypotheses along the lines roughly explored above. This is challenging in the experimental sense, because the current tools of experimental neuroscience tend to be either too low-level (single-neuron recordings) or too high-level (coarse whole-brain scans), but this problem will decrease as technology improves. In the socio-culturo-technological case, the challenge is to formulate the glocal memory hypothesis more crisply in some particular set of situations (e.g. a collaborative research team, a sports team, etc.) and turn it into a specific set of relatively easily testable hypotheses.

In conclusion, my suggestion is that keys, maps and glocality should become central parts of our working vocabulary for discussing, analyzing and engineering complex intelligent systems. In a scientific context: the feedback loop of interconstruction between semantically related keys and maps is a critical aspect of memory in complex systems such as human minds and social systems, and ignoring it will cause us ongoing confusion. In an AI engineering context: continuing to ignore glocality of memory will result in continuing to construct systems incapable of precise reasoning (due to lack of appropriate local manipulations) or incapable of intelligent contextual guidance of reasoning (due to lack of appropriate global dynamics corresponding to local manipulations).

## References


- Amit, Daniel (1989). Modeling Brain Function. Cambridge University Press.
- Calvin, William (1996). The Cerebral Code. MIT Press.
- Damasio, Antonio (2000). The Feeling of What Happens. Harvest Books
- Goertzel, Ben (1997). From Complexity to Creativity. Plenum Press.
- Elman, J. (1991). Distributed representations, simple recurrent networks, and grammatical structure. Machine Learning, 7(2/3):195–226.
- Goertzel, Ben (2006). The Hidden Pattern. Brown Walker Press.
- Goertzel, Ben (2006). Patterns, Hypergraphs and General Intelligence. Proceedings of IJCNN 2006, Vancouver Canada.
- Goertzel, Ben (2008). OpenCog Prime: Design for a Thinking Machine. Online book at http://opencog.org
- Goertzel, Ben, Matthew Ikle', Izabela Freire Goertzel and Ari Heljakka (2008). Probabilistic Logic Networks. Springer Verlag.
- Hutchins, E. (1995) Cognition in the Wild. MIT Press.
- Lebiere, C., & Anderson, J. R. (in preparation). The case for a hybrid architecture of cognition.

- Lenat, Douglas and R. V. Guha. (1990). Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project. Addison-Wesley
- Patterson, Karalny, Peter J. Nestor & Timothy T. Rogers (2007). Nature Reviews Neuroscience 8, 976-987 (December 2007)
- Perry, M. (2003) "Distributed Cognition". In J.M. Carroll (Ed.) "HCI Models, Theories, and Frameworks: Toward an Interdisciplinary Science" (Morgan Kaufmann) 193-223.
- Quiroga, R. Quian, L. Reddy, G. Kreiman, C. Koch and I. Fried (2005). Invariant visual representation by single-neurons in the human brain. Nature, 435: 1102-1107; 2005.
- Quiroga, R. Quian, Kreiman G, Koch C and Fried I. (2008). Sparse but not "Grandmother-cell" coding in the medial temporal lobe. Trends in Cognitive Sciences. 12: 87-91; 2008
- Santore, John F. and Stuart C. Shapiro (2003). Crystal Cassie: Use of a 3-D Gaming Environment for a Cognitive Agent. Papers of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions
- Storkey A.J. and R. Valabregue (1999) The basins of attraction of a new Hopfield learning rule, Neural Networks 12 869-876.