

A Quantum Information Retrieval Approach to Memory

Kirsty Kitto and Peter Bruza

School of Information Systems

Queensland University of Technology, Brisbane, Australia

Email: [kirsty.kitto,p.bruza]@qut.edu.au

Liane Gabora

Psychology and Computer Science

University of British Columbia, Kelowna, Canada

Email: liane.gabora@ubc.ca

Abstract—As computers approach the physical limits of information storable in memory, new methods will be needed to further improve information storage and retrieval. We propose a quantum inspired vector based approach, which offers a contextually dependent mapping from the subsymbolic to the symbolic representations of information. If implemented computationally, this approach would provide exceptionally high density of information storage, without the traditionally required physical increase in storage capacity. The approach is inspired by the structure of human memory and incorporates elements of Gärdenfors’ Conceptual Space approach and Humphreys et al.’s matrix model of memory.

I. MEMORY STRUCTURE AND INFORMATION DENSITY

The age of density driven computer memory increase is fast approaching its conclusion. With Moore’s law suggesting that we are nearing the physical limits of information density storable in standard computational memory, it is time to investigate new paradigms of information storage and retrieval. This paper proposes that a recently developed class of cognitive models provide a highly promising avenue, that can be used to shift the current information storage paradigm from a density dependent model to a more structural methodology. Our approach is inspired by insights from neuroscientific studies of memory, and focuses upon the context in which information is encoded and subsequently recalled. Mathematically, it is grounded in a vector-based formalism that utilizes the probability structure of quantum theory which draws upon two related lines of research. One derives from modern approaches to information retrieval which attempt to incorporate a sophisticated notion of context into the classification of information as relevant to a query ([1], [2], [3], [4]). The second approach is more squarely based in cognitive science, and uses a quantum approach to model concepts and their combinations ([5], [6], [7], [8]).

In summary, the key purpose of this paper is to suggest a new paradigm for information storage and retrieval *in context* that allows for a marked increase in the amount of information storable by a given resource. This will require the identification of a mechanism by which stored information can be retrieved, which somehow links that information to relevant storage and retrieval contexts. We provide tentative solutions for both of these problems.

We begin with a brief summary of how a subsymbolic encoding in human memory can still give rise to a symbolic

capacity. This is followed by a review of the Conceptual Space approach advocated by Gärdenfors [9], which proposes a framework of three tiers for understanding human memory. We then discuss the Matrix Model of Memory [10], which shows how a memory can be encoded along with information about the context in which it occurred. This will lead us to consider the treatment of context in that model and finally to extend it through reference to a quantum information retrieval framework which combines the desirable features of each approach. We propose that our approach not only allows for an exceptionally high density memory storage but also provides a memory architecture that can process information in a way that is flexible, adaptive, and possibly even creative.

II. SYMBOLIC AND SUBSYMBOLIC LEVELS OF HUMAN MEMORY

Let us begin by examining the architecture of human memory (summarized in [11]). This will serve as a starting point to build a computer memory that uses similar basic mechanisms to human memory.

A. The Subsymbolic Level

We take as a starting point some fairly well established characteristics of memory. Human memories are encoded in neurons that are sensitive to ranges (or values) of what has been called subsymbolic microfeatures [12], [13]. For example, one might respond to lines of a particular orientation, or the quality of honesty, or quite possibly something that does not exactly match an established term [14]. Note that sometimes the word concept is used by non-neuroscientists (*e.g.* [15]) to refer to subsymbolic microfeatures. In this paper, the word microfeatures is used to refer to stimuli responded to by single cells, which may or may not be meaningful in daily life, and the word concepts to refer to things like DOG or BEAUTY that are generally comprised of many microfeatures, and refer collectively to a class of instances or exemplars that are meaningful in daily life.

Another characteristic of memory is that although each neuron responds maximally to a particular microfeature, it responds to a lesser extent to related microfeatures, an organizational structure referred to as coarse coding [16]. For example, neuron A may respond preferentially to sounds of a certain frequency, while its neighbor B responds preferentially

to sounds of a slightly different frequency, and so forth. However, although A responds maximally to sounds of one frequency, it responds to a lesser degree to sounds of a similar frequency. The upshot is that an item in memory is stored in a distributed manner across a cell assembly that contains many neurons, and likewise, each neuron participates in the storage of many items [17]. A given experience activates not just one neuron, nor every neuron to an equal degree, but activation is spread across members of an assembly. This means that the same neurons get used and re-used in different capacities, a phenomenon referred to as neural re-entrance [18].

The final key attribute of memory is its content addressability, meaning that there is a systematic relationship between the content of a representation, and the neurons where it gets encoded. This emerges naturally as a consequence of the fact that representations activate neurons that are tuned to respond to particular features, so representations that get encoded in overlapping regions of memory share features. As a result, they can thereafter be evoked by stimuli that are similar or *resonant* in some (perhaps context-specific) way [17], [19]. Note that even if a brain does not possess a neuron that is maximally tuned to a particular microfeature, the brain is still able to encode stimuli in which that microfeature predominates, because representations are distributed across many neurons.

Note that on the basis of the discovery of single cells in the human brain that have highly selective, abstract and invariant responses to complex, natural stimuli, which have unfortunately been called concept cells, some neuroscientists have questioned the idea that representations are distributed [15]. This is not inconsistent with the variety of distributed representation discussed here. If you artificially activate one neuron, it gives an invariant response. It is because real-world stimuli and experiences activate not just one neuron but many that actual representations in memory are distributed.

B. The Symbolic Level

Consciously experienced symbolic meanings emerge in response to the set of subsymbolic microfeatures responded to by the entire constellation of activated neurons. Sometimes these neurons have been activated as a unit many times before, at others the constellation consists of neurons that have never been activated simultaneously as a whole. In this case an emergent meaning may simultaneously incorporate elements of different symbolic representations.

The distributed, content addressable architecture of memory is critical to the adaptive, flexible, and creative manner in which the information it stores is not just retrieved when required, but frequently reconstructed in contextually appropriate and sometimes even creative ways. If this memory were not distributed then there would be no overlap between items that share microfeatures, and thus no means of forging associations between them. If it were not content-addressable then these associations would not be meaningful. The upshot is that representations which share features are encoded in overlapping distributions of neurons, and therefore activation can spread

from one to another. Thus representations are encoded in memory in a way that takes into account how they are related, even if this relationship has never been consciously noticed [20], [11], [21]. This is not earth shattering; indeed it seems fairly obvious with respect to the hierarchical structure of knowledge. We may never have explicitly learned that a white hamster is a mammal, but we know it is one nonetheless. In this sense it is reasonable to claim that people implicitly know more than they have ever explicitly learnt. This architecture has implications that extend far beyond issues related to the hierarchical structure of knowledge.

It should be pointed out how different this is from the typical structure of computer memory. In a computer memory, each possible input is stored in a unique address. Retrieval is thus a matter of looking up the address in a register and then fetching the corresponding item at the specified location. Since there is no overlap of representations, there is no means of creatively forging new associations based on newly perceived similarities. The exceptions are computer architectures that are designed to mimic, or are inspired by, the distributed, content-addressable nature of human memory, but these are difficult to discuss formally. In this paper we shall propose a theoretical structure that can be used to map subsymbolic architectures to symbolic representations, and so potentially provide a more flexible, adaptable and creative approach to computer memory.

C. Forging Unusual Associations through Reconstructive Interference of Memories

A fascinating finding to come out of the early connectionist literature is that in a distributed, content addressable memory, not only do representations that share features activate each other, they sometimes interact in a way that is creative. Even a simple neural network is able to abstract a prototype, fill in missing features of a noisy or incomplete pattern, or create a new pattern on the fly that is more appropriate to the situation than anything it has ever been fed as input [22]. In fact, similar representations can interfere with one another [23], [24], [25], and these same papers provide numerous names for this phenomenon: crosstalk; false memories; spurious memories; ghosts; and superposition catastrophe. These phenomena are suggestive of a form of thought that, if not outright creative, involves a departure from known reality. Findings from neuroscience are also highly consistent with this phenomenon; as Edelman puts it, one does not retrieve a stored item from memory so much as reconstruct it [26]. That is, an item in memory is never re-experienced in exactly the form it was first experienced, but colored, however subtly, by what has been experienced in the meantime, re-assembled spontaneously in a way that relates to the task at hand (one reason eye-witness accounts cannot always be trusted [27], [28], [29]).

Because information is encoded in a distributed manner across ensembles of neurons interacting by way of synapses, the meaning of a representation is in part derived from the meanings of other representations that excite similar constellations of neurons; that is why memory is sometimes referred

to as associative. Content addressability ensures that the brain naturally brings to mind items that are similar in some perhaps unexpected or indefinable but useful or appealing way to what is currently being experienced. Recall that if the regions in memory where two distributed representations are encoded overlap then they share one or more microfeatures. They may have been encoded at different times, under different circumstances, and the correlation between them never explicitly noticed. But the fact that their distributions overlap means that some context could come along for which this overlap would be relevant, causing one to evoke the other. There are as many routes by which an association between two representations can be forged as there are microfeatures by which they overlap; i.e., there is room for typical as well as atypical connections. Therefore what gets evoked in a given situation is relevant, and that happens for free; no search is necessary at all because memory is content-addressable. The *like attracts like* principle is embedded in our neural architecture.

Moreover, because memory is distributed and subject to crosstalk, if a situation does come along that is relevant to multiple representations, they merge together, a phenomenon that has been termed *reconstructive interference* [30]. The multiple items may be so similar to each other that you never detect that the recollection is actually a blend of many items, and in this case the distributions of neurons activated overlap substantially. Alternatively, they may differ in mundane ways, as in everyday mind-wandering. They may be superficially different but related in a way you never noticed before, in which case the distributions of neurons activated overlaps only with respect to only a few features that happen to be relevant or important in the present context. Finally, the present experience may infuse recall of a previous experience that is relevant or important with respect to only a few key features.

We now turn to some of the models that have been proposed to describe this merging of the subsymbolic and the symbolic levels of human memory. We shall find a number which share a key set of features.

III. MEMORY MODELS INSPIRED BY THE MULTI-LEVELLED ARCHITECTURE OF HUMAN MEMORY

Gärdenfors [9] has proposed a three level model of cognition in which the representation of information varies greatly at each level. Within the lowest level, information is pre-conceptual or subsymbolic, and is carried by a connectionist representation.

At the uppermost level information is represented in terms of higher order symbolic structures such as sentences. Grammars specify the parts of a sentence, and the manner in which they fit together. It is at this upper symbolic level of cognition where a significant portion of the computational literature resides. Indeed, the very storage of information in a standard computer architecture could be understood as belonging to this level.

While the need for at least these two levels seems intuitively plausible, the gap between the upper, logical level and the lowest connectionist level is difficult to bridge. How are we to

connect the symbolic approaches with the structural? There is a possibility for some approach that shares logical *and* structural characteristics *between* the symbolic and the structural levels of cognition, and this is precisely where Gärdenfors' intermediate, conceptual level, or *conceptual space*, is introduced. Rather than relying upon a connectionist structure, an intermediate geometric representation is used which provides an expressive theoretical framework capable of linking the 'hardware' of a 'neuronal' level with the more commonly described, and theoretically understood, logical level.

IV. ENCODING INFORMATION IN A CONCEPTUAL STRUCTURE

Within a conceptual space, knowledge has a dimensional structure. For example, the property COLOR can be represented in terms of three dimensions: hue, chromaticity, and brightness, which can be mapped into a convex region in a geometric space. Thus, the property RED is a convex region within the tri-dimensional space made up of hue, chromaticity and brightness, and the property BLUE would occupy a different region of this same space. A domain is a set of integral dimensions in the sense that values in particular dimensions can determine (or affect) the values possible in others. For example, the three dimensions defining the color space are integral since the brightness of a color will affect both its saturation (chromaticity) and hue. Gärdenfors extends the notion of properties into concepts, which are likewise based on domains.

For example, the concept APPLE may have domains taste, shape, color, etc. Context is modeled as a weighting function on the domains, for example, when eating an apple, the taste domain will be prominent, but when playing with it, the shape domain (i.e. its roundness) will be heavily weighted.

Observe the distinction between representations at the symbolic and conceptual levels. At the symbolic level, the concept APPLE can be represented as the atomic proposition $\text{apple}(x)$. However, within a conceptual space (conceptual level), it has a representation involving multiple inter-related dimensions and domains. Colloquially speaking, the token "apple" (which might be spoken, written *etc.*) is the tip of an iceberg with a rich underlying representation at the conceptual level. Gärdenfors points out that the symbolic and conceptual representations of information are not in conflict with each other, but are to be seen as "different perspectives on how information is described".

However, an implementation problem arises in that both the representation and the generation of a conceptual space from its underlying content has generally been discussed only for simple examples such as those above. It is not clear how more complex examples could be implemented. A more comprehensive and systematic approach to the representation of conceptual spaces is required.

Vector space based models (VSBM) provide a viable first avenue here. These can be traced back to the seminal paper of Salton et al. [31] who were searching for an appropriate mathematical space to represent documents for the purposes

of Information Retrieval. Starting from a few basic desiderata, they settled upon a vector in a high dimensional vector space as an appropriate representation of a document. Within this framework, a query is treated like a small (pseudo) document that is also converted to vector form. The documents in the corpus are then ranked according to their distance to the query; closer documents are considered more relevant than ones that are further away. One of the main drawbacks of this system was that it had trouble returning documents that would have been highly relevant if one of the words in the query was replaced by a synonym. The next advance came from representing concepts latently in a so-called *semantic space* where they are not formally represented or labeled. Semantic spaces are instances of vector spaces, and represent words in a basis created from other words, concepts, documents, or topics. They are generally built from the observation of co-occurrences in large text corpora. In word spaces such as the Hyperspace Analogue to Language (HAL) [32], the basis consists of every word in the vocabulary. Thus, the vector for a given word w is calculated by summing the number of occurrences of word $w(i)$ in a given context window around each occurrence of w and writing that number at the position i in the vector that represents w . This number can be adjusted using the distance (defined in terms of the number of words) or mutual information measures such as Point-Wise Mutual Information, which allows for a weighting of the importance of the word at that position. It is also possible to take word order into account [33], [34]. Later models derived a more fundamental semantic value through a reduction of the initial word space using mathematical tools such as Singular Value Decomposition [35], Non Negative Matrix factorization [36], or random projection [37], all of which generate a new smaller basis which can under certain conditions be naturally related to certain topics, objects and concepts [36].

Semantic space models, however, do not make provision for integral dimensions (a notion related to that of ‘core properties’ of a concept). This leaves them too situation dependent, and relevant primarily for the text collection from which they were constructed. For the purposes of next generation information storage, we will require a more objective information storage mechanism that can function satisfactorily at the conceptual level. Thus, while learning from the semantic space approaches, this paper will propose that we extend them from a text based, and corpus dependent information representation, to a concept and property inspired approach. However, the vector based analytical contributions of the semantic space approaches will play a key inspirational role as we start with this extension.

V. INFORMATION RETRIEVAL IN A MATRIX MODEL OF MEMORY

We now turn to a promising treatment of retrieval, that is capable of providing a map between the neural and conceptual levels of information storage. The matrix model of memory [10] is a well known cognitive model. It stores and encodes memories as patterns of interconnections between

the elements that define items in memory (*i.e.* the neurons for a subsymbolic structure). All memories are superimposed (summed) in this representation so that, without appropriate cuing, their individual identities are lost. Thus, the model provides a natural link between the lower and mid levels of information that Gärdenfors proposes. For example, when a set of interconnected neurons fires, this can be represented in the matrix model as a set of entries in a matrix, with the entries in the matrix corresponding to the probability that a particular pairing of nodes will concurrently fire (although this is not a necessary interpretation of the model [10]).

Humphreys et al. [10] take the position that there are two fundamental but pervasive memory access operations; matching and retrieval. Matching involves the comparison of test cue(s) with the information stored in memory, and gives the strength of the match as output. In contrast, retrieval involves the recovery of an associate of a cue (*i.e.* the return of actual information), and so is the concept in which we are currently interested.

The matrix model takes an item A_i , occurring in a context X , to retrieve a list associate B_i . This assumes that a three-way association (between the context, the cue and the desired item) must have been stored. This is represented mathematically as the three-dimensional array $\mathbf{x}\mathbf{a}'_i\mathbf{b}''_i$, where \mathbf{x} is a column vector, \mathbf{a}'_i is a row vector, $\mathbf{x}\mathbf{a}'_i$ is a $n \times n$ matrix, and \mathbf{b}''_i is another vector in an orthogonal dimension to \mathbf{x} and \mathbf{a}'_i . (Primes are used to indicate this set of orthogonality relationships.) The matrix $\mathbf{x}\mathbf{a}'_i$ represents the *association* between the context \mathbf{x} and the cue \mathbf{a}'_i . If a list of items $A_1B_1, A_2B_2 \dots A_kB_k$ is learned in a context X , then Humphreys et al. define the memory for the list as a simple sum over all the three-dimensional arrays that were formed:

$$\mathbf{E} = \sum_{i=1}^k \mathbf{x}\mathbf{a}'_i\mathbf{b}''_i. \quad (1)$$

This list memory \mathbf{E} is added onto any pre-existing memories in a process that we leave to the original article [10].

Retrieval from this list memory is defined by Humphreys et al. [10] to work as follows. First, a test cue $\mathbf{x}\mathbf{a}'_j$ is applied to the list memory:

$$(\mathbf{x}\mathbf{a}'_j) \cdot \mathbf{E} = (\mathbf{x}\mathbf{a}'_j) \cdot \left(\sum_{i=1}^k \mathbf{x}\mathbf{a}'_i\mathbf{b}''_i \right) \quad (2)$$

$$= \sum_{i=1}^k [(\mathbf{x}\mathbf{a}'_j) \cdot (\mathbf{x}\mathbf{a}'_i)] \mathbf{b}''_i \quad (3)$$

$$= \sum_{i=1}^k [(\mathbf{x} \cdot \mathbf{x})(\mathbf{a}'_j \cdot \mathbf{a}'_i)] \mathbf{b}''_i \quad (4)$$

$$= [(\mathbf{x} \cdot \mathbf{x})(\mathbf{a}'_j \cdot \mathbf{a}'_j)] \mathbf{b}''_j + \sum_{i \neq j} [(\mathbf{x} \cdot \mathbf{x})(\mathbf{a}'_j \cdot \mathbf{a}'_i)] \mathbf{b}''_i. \quad (5)$$

We can learn a little about this model through a consideration of the two terms in (5). The first term represents the desired vector \mathbf{b}''_j , weighted by a scalar term that results from taking

the dot products of two vectors with themselves. The second term is effectively an error term; if the similarity between the cue \mathbf{a}'_j and the other cues that were used to store the memory (\mathbf{a}'_i) is too great then this error term will become large and the chances of the correct item being recalled will decrease. In short, the other stored memories ($\mathbf{b}''_i, (i \neq j)$) will interfere with the desired term. It is also worth noting that the explicit representation of the context vectors using a dot product (i.e. $\mathbf{x} \cdot \mathbf{x}$) in (5) suggests that the authors were open to the idea of a different context being present during the recall process, even if this was not explicitly discussed [10]. We shall return to this point in section VII, however, a brief foreshadowing of that argument runs as follows: We consider the use of context in this model to be unsatisfactory. Firstly, while the role of context is fundamental to this model, it must be explicitly and fully recorded at the time of storage. A slightly different context, or even a more detailed specification of the same context could result in the retrieval of a very different piece of information, even though a very similar cue was utilized. The static treatment of context that is provided by this model therefore leaves us with what is an interesting retrieval paradigm, which is however perhaps unnecessarily limited. We are left wondering if there is scope for a more adaptive treatment of context, one provided by the geometric models of conceptual space that were introduced in section IV.

In the remainder of this article we shall endeavor to connect these two interesting approaches into an integrated cognitive memory model which could be used to form the basis of a future physical implementation of computational memory. This approach takes its inspiration from a set of models that consider information retrieval in context, utilizing the powerful formalism of quantum theory [1]–[4], [38], [39].

VI. INCORPORATING CONTEXT INTO INFORMATION ENCODING AND RETRIEVAL

The seminal book by van Rijsbergen [1] provides a novel approach to the modeling of semantic spaces, inspired by Quantum Theory (QT). This approach models a word w as a vector

$$|w\rangle = \begin{pmatrix} w_1 \\ \vdots \\ w_n \end{pmatrix} \quad (6)$$

where $|w\rangle$ is termed a *ket*, in contrast to the row vector obtained by taking the transpose: $\langle w| = |w\rangle^T = (w_1, \dots, w_n)$. In this case, we shall take the subcomponents $w_1 \dots w_n$ to be the weights allocated to each of the available *senses* that the word might take in a n -dimensional vector space.

We can quickly see the connection to both vector space based approaches and the Matrix Model of Memory. In both cases a vector was obtained, (although in each case this was via a different process) and formed the basis of further analysis. However, the formalism of quantum theory provides an extra level of structure that implicitly incorporates a more adaptive notion of context into information recall.

This is done by seriously considering what it means to define a *basis* for a conceptual space. Thus, the vector of (6) must be considered *in its context*; it is a representation of information within a high dimensional vector space, with the vector entries determining the extent of the vector in each of the relevant dimensions.

This geometric model provides predictions about the likely recall of an item from memory within a given context. This is achieved via an application of Pythagoras' theorem. Thus, simplifying equation (6) down to a vector occurring in a two dimensional space, we might find that it could be drawn as shown in figure 1, where

$$|w\rangle = a_0 \begin{pmatrix} 0 \\ 1 \end{pmatrix} + a_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (7)$$

$$= a_0|0_p\rangle + a_1|1_p\rangle, \text{ and } |a_0|^2 + |a_1|^2 = 1. \quad (8)$$

Here, $\{|0_p\rangle \equiv (0, 1)^T, |1_p\rangle \equiv (1, 0)^T\}$ define an orthonormal

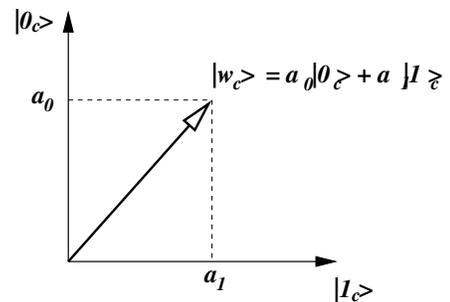


Fig. 1. A concept w , for example *red*, is represented in some contextual probe c which takes the form of a choice of basis. This low dimensional representation shows a case of some object being classified by a probe as “red” ($|1\rangle$) or “not red” ($|0\rangle$), with the probability of that a “red” judgment being given by $|a_1|^2$. In a higher number of dimensions the property of redness will be enclosed by a convex region, and the probabilities lie in a range of values specified by the extent of that region.

basis, and so the inner product of these basis vectors returns 0 or 1: $\langle 0_p|0_p\rangle = \langle 1_p|1_p\rangle = 1$ and $\langle 1_p|0_p\rangle = \langle 0_p|1_p\rangle = 0$. The state (6), can be re-written using an extension of this formalism, giving

$$|w\rangle = w_1 \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + w_2 \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix} + \dots + w_n \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix} \quad (9)$$

which allows us to capture high dimensional vector representations of information. Here, the w_j 's, or weights, represent the extent to which a piece of information falls into each of the dimensions of the vector space, and thus how much it overlaps with the individual concepts represented by each axis in the basis. This means that the convex region representing a property in a conceptual space can be mapped out by a collection of vectors covering that region, with each of the weights mapping how much a given property is represented by that dimension. A piece of stored information, (e.g. a concept w) is thus represented in this framework as a state, or a vector in a high dimensional space. For a low dimensional example,

consider the concept of “redness” that might be stored about two different objects, such as an apple, and some wine. In a two dimensional, or q -bit representation, each object will be classified as either “red” or “not red” but this classification will depend upon the context. Figure 1 depicts a possible state which one of these objects might have, within a particular concept space where $|1\rangle$ represents “red” and $|0\rangle$ “not red”. Within this specific context, we might find that an apple is more likely to be returned as a response to a query that asks for a “red object” than is red wine, although this might change were the information to be sought in a different context. We shall return to this point shortly, showing how this formalism can very naturally capture this behavior.

We propose that once information is stored in this complex multidimensional space, it can be recovered through use of a probe which enacts a *quantum measurement* of the state (7). This is defined with respect to a projection operator V , where, for the two dimensional case outlined above

$$V = |0_p\rangle\langle 0_p| + |1_p\rangle\langle 1_p| = V_0 + V_1. \quad (10)$$

According to the quantum formalism, the probability of a probe represented by the p basis returning the desired value is given by

$$P(|1\rangle) = \langle w|V_1|w\rangle \quad (11)$$

$$= \langle w|1_p\rangle\langle 1_p|w\rangle \quad (12)$$

$$= (a_0^*\langle 0_p|1_p\rangle + a_1^*\langle 1_p|1_p\rangle) \times (a_0\langle 0_p|1_p\rangle + a_1\langle 1_p|1_p\rangle) \quad (13)$$

$$= |a_1|^2. \quad (14)$$

However, in the context represented by the shifted basis in figure 2 the probe returns the desired information with a probability given by $P(|1\rangle) = |b_1|^2$. Thus, a search for a

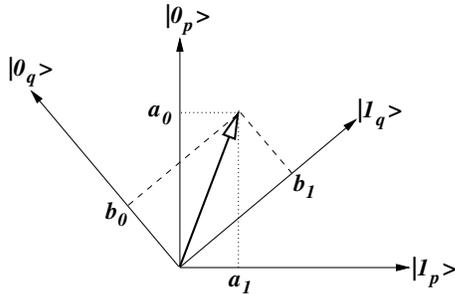


Fig. 2. Changing the context of a probe can significantly affect the chances of recall. Once this effect is incorporated into a Matrix-like memory model, a structural information storage system becomes more viable.

“red object” in the context represented by p may yield a very different result to that which searches in the context q .

The assumption in (7) that the squared coefficients of the basis vectors sum to 1 allows for the treatment of these values as probabilities since $0 \leq P \leq 1$ etc. This approach makes use of a geometrical notion of probability [40], which contrasts with standard probability theory, where probabilistic outcomes arise from our lack of knowledge as to what has actually occurred. Quantum probabilities are profoundly different, arising

from a genuine state of uncertainty; the context in which the information is to be represented must be defined before the recall can start to make sense.

We shall now show how this more sophisticated treatment of context can be utilized in an extension of the Matrix Memory Model which, while retaining the representation of items and cues as vectors, embeds them within a context that is spatial rather than vectorial.

VII. REMEMBERING AS A PROCESS OF INFORMATION RETRIEVAL

In line with the proposal by Wiles et. al [41], we take the position that the recall of information from a memory structure can be well represented by a contextual probe to an underlying network structure. The construction of such a probe has been a difficult problem for neural modelers, as it is difficult to correlate the activation of neural connections with a logical, or even conceptual, structure. However, with the three tier model advocated by Gärdenfors we can begin to see how a probe that has both a logical (symbolic) structure, and a connection to the lower (subsymbolic) level neural model can be created. This section will sketch out the key details of this construction.

We start with a reference to the result shown in van Rijsbergen that projection operators such as the one in (10) can be used to define a conditional logic [1], meaning that the link between the quantum conceptual space that we discussed in section VI and higher order logic has already been found.

This leaves the connection between the subsymbolic neural level and the conceptual levels to be made. Returning to the consideration of the Matrix Model that was started in section V we recall its use of a somewhat unsatisfactory *explicit context*. The representation of context in this model as a vector (i.e. \mathbf{x}) means that it acquires an ontological status equivalent to that of the items that are used as cues, or stored to be retrieved by those cues, however, we believe that this identification is incorrect. Rather than behaving as a thing, or absolute entity, context appears to be more of a *relationship* between the thing currently under consideration (i.e. the memory for this scenario) and a *perspective* from which it will be viewed. This is a very new approach to the treatment of context in computational representations, which most commonly take context to be a thing, or a parameter [42], [43] with a similar ontological status as very system which is being considered *within that context*. This is unlikely to be a satisfactory approach, but the lack of alternative formal models has hindered the adoption of a more sophisticated understanding. However, the quantum inspired model presented above makes use of a very different conceptualization, that we shall refer to here as an *implicit context*, which frames the system under consideration rather than being of a similar form to it. In what follows, we shall make use of this implicit notion of context in an extension to the Matrix Model of Memory which treats context as a space rather than a vector.

This will be achieved by utilizing projection operators rather than vectors to represent the context in which storage and recall takes place. Thus, in place of the context vector \mathbf{x} in

(5), we propose to utilize a projection operator that arises in the same space as the memory itself

$$\mathbf{V}_x = \sum_{h=1}^n |x_h\rangle\langle x_h|. \quad (15)$$

This equation takes the vector notion of context utilized in equation (5) and translates it into a set of projection operators defined using basis vectors, each of which could have been a context vector in the Matrix Model. Returning to equation (5) we rewrite it with this extended understanding of the context of a memory:

$$\mathbf{V}_y \mathbf{a}_j \mathbf{E} = \mathbf{V}_y \mathbf{a}_j \sum_{i=1}^k \mathbf{V}_x \mathbf{a}_i \mathbf{b}_i \quad (16)$$

where \mathbf{V}_y is a second cueing context defined with respect to the vector \mathbf{y} which could be specified with a different set of basis vectors to \mathbf{x} . Expanding the projection operators in this equation starts to give us some indication of how this model can be expected to behave:

$$\mathbf{V}_y \mathbf{a}_j \mathbf{E} = \sum_{h=1}^n \sum_{i=1}^k |y_h\rangle\langle y_h|a_j\rangle\langle x_i|a_i\rangle|b_i\rangle \quad (17)$$

$$= \sum_{h=1}^n \sum_{i=1}^k u_{h,j} v_i |y_h\rangle\langle x_i|b_i\rangle \quad (18)$$

$$= \sum_{h,i} u_{h,j} v_i \mathbf{y}_h \mathbf{x}_i \mathbf{b}_i, \quad (19)$$

where $u_{h,j} = \langle y_h|a_j\rangle$ and $v_i = \langle x_i|a_i\rangle$ are scalars, obtained by taking the associated dot products of the corresponding vectors. These scalars weight the contribution of the individual cue, context and stored item vectors. If the context of recall is the same as the context of recording (*i.e.* $\mathbf{y} = \mathbf{x}$) then we can say a little more about the recall process using a standard property of projection operators: $\mathbf{V}_x \mathbf{V}_x = \mathbf{V}_x$ [40].

$$\mathbf{V}_x \mathbf{a}_j \mathbf{E} = \mathbf{V}_x \mathbf{a}_j \sum_{i=1}^k \mathbf{V}_x \mathbf{a}_i \mathbf{b}_i \quad (20)$$

$$= \mathbf{V}_x \mathbf{a}_j \sum_{i=1}^k \mathbf{a}_i \mathbf{b}_i \quad (21)$$

$$= \sum_{i=1}^k v_i \mathbf{x}_i \mathbf{a}_i \mathbf{b}_i. \quad (22)$$

Finally, breaking (22) into the two components utilized in (5) we find that

$$\mathbf{V}_y \mathbf{a}_j \mathbf{E} = v_j \mathbf{x}_j \mathbf{a}_j \mathbf{b}_j + \sum_{i \neq j} v_i \mathbf{x}_i \mathbf{a}_i \mathbf{b}_i \quad (23)$$

which has the same item to be retrieved + error terms of (5) but in more complex space that contains all cue, contexts and items stored in the memory. We finish by noting that this equation suggests that a context which maximizes v_j will increase the probability of a correct retrieval result, but many different contexts could have been used. Indeed, we need

merely shift the basis in equation (15) in order to obtain a very different set of representations for the items in memory, and these would have a very different set of probabilities for recall. Thus, with a shift to a geometric space we see a way in which information might be stored and retrieved from a system based upon traditional storage mechanisms than is currently the case, all through the use of a sophisticated notion of context.

VIII. CONCLUSION

A strength of this approach lies in the density of information that it is likely to be able to store. The choice of a structural approach to information storage, with a potentially infinite set of contexts, allows for memory to move from a density driven approach, where the quantity of information stored is inversely proportional to the size of the components used to store it, towards one where storage capacity is limited only by how many sensible contexts can be used to retrieve the required information. Even with a very small storage space, a wide range of representations can still be obtained from a conceptual space that takes the underlying subsymbolic structure and complexifies it according to the context in which it is accessed. Such an ‘‘actualization of potential’’ [30] provides both creative ability and extra storage capability. Indeed, with this approach, we can start to see how the lowest, or neural level of cognition can be made redundant despite its strong dependence upon a specific structure.

While we have emphasized the background of this quantum inspired model in the field of Information Retrieval, a related line of work [5]–[8], [45], [46] makes use of a quantum approach to model concepts and their combinations. Thus, a growing body of literature points to the utility of the quantum formalism in modeling Information in context from both the cognitive and the computational sides of Information storage and retrieval. This approach has also been utilized in a preliminary approach illustrating how context might be incorporated into vector spaces described with reference to a point of view [47]. A solution which that paper shows might circumvent the apparent incompatibility between metricity and the similarity judgments that humans actually make [48].

While the proposed approach is in its very early days, we feel that its incorporation of a wide range of both cognitive and computational insights makes it a highly interesting avenue to pursue as we search for new paradigms of computational memory and information storage. Future work will investigate the manner in which different contexts might interfere with specified cues to produce different probabilities of recall and hence different items from memory. It will also seek to further clarify the role of the projection operators in specifying a context space, and to extend the formalism proposed at the end of the previous section. Finally, we intend to take seriously the notion of creativity as it arises in human memory, and to investigate the manner in which a similar notion might arise in a system such as this. Such a result would bring us one step closer towards a system capable of exhibiting a true form of computational intelligence.

ACKNOWLEDGMENTS

This project was supported in part by the Australian Research Council Discovery grant DP1094974, the Social Sciences and Humanities Research Council of Canada, and the Fund for Scientific Research of Flanders, Belgium. Welcome support was also provided by the Marie Curie International Research Staff Exchange Scheme: Project 247590, “QONTEXT - Quantum Contextual Information Access and Retrieval”).

REFERENCES

- [1] C. Van Rijsbergen, *The Geometry of Information Retrieval*. Cambridge University Press, 2004.
- [2] D. Widdows, *Geometry and Meaning*. CSLI Publications, 2004.
- [3] D. Song and P. Bruza, “Towards context sensitive information inference,” *Journal of the American Society for Information Science and Technology*, vol. 54, no. 4, pp. 321–334, 2003.
- [4] P. Bruza, D. Widdows, and J. A. Woods, “Quantum Logic of Down Below,” in *Handbook of Quantum Logic and Quantum Structures*, K. Engesser, D. Gabbay, and D. Lehmann, Eds. Elsevier, 2007, vol. 2.
- [5] D. Aerts and L. Gabora, “A theory of concepts and their combinations I: the structure of the sets of contexts and properties,” *Kybernetes*, vol. 34, pp. 151–175, 2005.
- [6] —, “A theory of concepts and their combinations II: a Hilbert space representation,” *Kybernetes*, vol. 34, pp. 192–221, 2005.
- [7] K. Kitto, B. Ramm, P. D. Bruza, and L. Sitbon, “Testing for the non-separability of bi-ambiguous words,” in *Proceedings of the AAAI Fall Symposium on Quantum Informatics for Cognitive, Social, and Semantic Processes (QI 2010)*. AAAI Press, 2010.
- [8] K. Kitto, B. Ramm, L. Sitbon, and P. D. Bruza, “Quantum theory beyond the physical: information in context,” *Axiomathes*, vol. 21, no. 2, pp. 331–345, 2011.
- [9] P. Gärdenfors, *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000.
- [10] M. Humphreys, J. Bain, and R. Pike, “Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks,” *Psychological Review*, vol. 96, no. 2, pp. 208–233, 1989.
- [11] L. Gabora, “Revenge of the ‘neurds’: Characterizing creative thought in terms of the structure and dynamics of human memory,” *Creativity Research Journal*, vol. 22, no. 1, pp. 1–13, 2010.
- [12] P. Smolensky, “On the proper treatment of connectionism,” *Behavioral and Brain Sciences*, vol. 11, pp. 1–43, 1988.
- [13] P. S. Churchland and T. Sejnowski, *The computational brain*. MIT Press, 1992.
- [14] Mikkulainen, “Neural network perspectives on cognition and adaptive robotics,” in *Natural language processing with subsymbolic neural networks*, A. Brown, Ed. Institute of Physics Press, 1997.
- [15] A. Roy, “Discovery of concept cells in the human brain — could it change our science?” *Natural Intelligence*, vol. 1, no. 1, pp. 23–29, 2011.
- [16] D. Hubel and T. N. Wiesel, “Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat,” *Journal of Neurophysiology*, vol. 28, pp. 229–289, 1965.
- [17] D. Hebb, *The organization of behavior*. Wiley, 1949.
- [18] G. M. Edelman, *Neural Darwinism: the theory of neuronal group selection*. Oxford University Press, 1989.
- [19] D. A. Marr, “Theory of the cerebellar cortex,” *Journal of Physiology*, vol. 202, pp. 437–470, 1969.
- [20] L. Gabora, “Cognitive mechanisms underlying the origin and evolution of culture,” Ph.D. dissertation, Center Leo Apostel for Interdisciplinary Studies, Vrije Universiteit Brussel, 2001.
- [21] L. Gabora and A. Ranjan, “How insight emerges in distributed, content-addressable memory,” in *The Neuroscience of Creativity*, A. Bristol, O. Vartanian, and J. Kaufman, Eds. Oxford University Press, In Press.
- [22] J. L. McClelland and D. E. Rumelhart, “A distributed model of memory,” in *Parallel distributed processing: Explorations in the microstructure of cognition*, D. E. Rumelhart, J. L. McClelland, and the PDP research group, Eds. MIT Press, Cambridge, 1986, vol. II.
- [23] J. Feldman and D. Ballard, “Connectionist models and their properties,” *Cognitive Science*, vol. 6, pp. 204–254, 1982.
- [24] J. Hopfield, “Neural Networks and Physical Systems with Emergent Collective Computational Abilities,” *Proceedings of the National Academy of Sciences*, vol. 79, pp. 2554–2558, 1982.
- [25] J. Hopfield, D. I. Feinstein, and R. D. Palmer, “‘unlearning’ has a stabilizing effect in collective memories,” *Nature*, vol. 304, pp. 159–160, 1983.
- [26] G. M. Edelman, *Bright Air, Brilliant Fire: On the Matter of Mind*. Basic Books, 1993.
- [27] H. M. Paterson, R. I. Kemp, and J. P. Forgas, “Co-witnesses, confederates, and conformity: The effects of discussion and delay on eyewitness memory,” *Psychiatry, Psychology and Law*, vol. 16, no. 1, pp. S112–S124, 2009.
- [28] E. F. Loftus, *Memory: Surprising new insights Into how we remember and why we forget*. Addison-Wesley Pub. Co., 1980.
- [29] D. L. Schacter, *The seven sins of memory: How the mind forgets and remembers*. Houghton Mifflin Co., 2001.
- [30] L. Gabora and A. Saab, “Creative interference and states of potentiality in analogy problem solving,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*. Boston MA.: Cognitive Science Society, 2011, pp. 3506–3511.
- [31] G. Salton, A. Wong, and C. Yang, “A vector space model for automatic indexing,” *Communications of the ACM*, vol. 18, no. 11, pp. 613–620, 1975.
- [32] H. Schütze, “Automatic word sense discrimination,” *Computational linguistics*, vol. 24, no. 1, pp. 97–123, 1998.
- [33] M. N. Jones and D. J. K. Mewhort, “Representing Word Meaning and Order Information in a Composite Holographic Lexicon,” *Psychological Review*, vol. 114, no. 1, pp. 1–37, 2007.
- [34] M. Sahlgren, A. Holst, and P. Kanerva, “Permutations as a means to encode order in word space,” in *Proceedings of the 30th Annual Meeting of the Cognitive Science Society*, 2008, pp. 11–18.
- [35] T. Landauer and S. Dumais, “A solution to plato’s problem: the latent semantic analysis theory of acquisition, induction and representation of knowledge,” *Psychological review*, vol. 104, no. 2, pp. 211–240, 1997.
- [36] D. D. Lee and H. S. Seung, “Learning the parts of objects by non-negative matrix factorization,” *Nature*, vol. 401, no. 6755, pp. 788–791, 1999.
- [37] M. Sahlgren, “An introduction to random indexing,” in *Proceedings of Methods and Applications of Semantic Indexing Workshop at the 7th International Conference on Terminology and Knowledge Engineering*, 2005.
- [38] P. Bruza and R. Cole, “Quantum Logic of Semantic Space: An Exploratory Investigation of Context Effects in Practical Reasoning,” in *We Will Show Them: Essays in Honour of Dov Gabbay*, S. Artemov, H. Barringer, A. d’Avila Garcez, L. Lamb, and J. Woods, Eds. College Publications, 2005, vol. 1, pp. 339–361.
- [39] M. Melucci, “A basis for information retrieval in context,” *ACM Trans. Inf. Syst.*, vol. 26, pp. 14:1–14:41, June 2008.
- [40] C. J. Isham, *Lectures on Quantum Theory*. London: Imperial College Press, 1995.
- [41] J. Wiles, G. Halford, J. Stewart, M. Humphreys, J. Bain, and W. Wilson, “Tensor Models: A creative basis for memory and analogical mapping,” in *Artificial Intelligence and Creativity*, T. Dartnall, Ed. Kluwer Academic Publishers, 1994, pp. 145–159.
- [42] P. Brézillon, “Context in problem solving: a survey,” *Knowledge Engineering Review*, vol. 14, pp. 47–80, May 1999.
- [43] R. Guha and J. McCarthy, “Varieties of contexts,” in *Modeling and Using Context*, ser. Lecture Notes in Computer Science, P. Blackburn, C. Ghidini, R. Turner, and F. Giunchiglia, Eds. Springer Berlin / Heidelberg, 2003, vol. 2680, pp. 164–177.
- [44] K. Kitto and P. Bruza, “Tests and models of non-compositional concepts,” in *Proceedings: Cognitive Science 2012*, Japan, 2012, in Press.
- [45] P. Bruza, K. Kitto, B. Ramm, and L. Sitbon, “The non-decomposability of concept combinations,” under review.
- [46] D. Aerts, J. Broekaert, and L. Gabora, “A case for applying an abstracted quantum formalism to cognition,” *New Ideas in Psychology*, vol. 29, no. 1, pp. 136–146, 2011.
- [47] S. Aerts, K. Kitto, and L. Sitbon, “Similarity metrics within a point of view,” in *Quantum Interaction 5th International Symposium, QI 2011, Aberdeen, UK, June 26-29, 2011, Revised Selected Papers*, ser. LNCS, D. Song, M. Melucci, and et al., Eds., vol. 7052. Springer, 2011, pp. 13–24.
- [48] A. Tversky and I. Gati, “Similarity, separability, and the triangle inequality,” *Psychological Review*, vol. 89, no. 2, pp. 123–54, 1982.