Computational Processes that Appear to Model Human Long-term Memory

John L. Pfaltz
Dept. of Computer Science, University of Virginia

July 13, 2017

Abstract

This paper presents two computable functions, ω and ε, that map networks into networks. If all cognition occurs as an active neural network, then it is thought that ω models long-term memory consolidation and ε models memory recall. A derived, intermediate network form, consisting of chordless cycles, could be the structural substrate of long-term memory; just as the double helix is the necessary substrate for genomic memory.

1 Introduction

There seems to be consensus that our sensations, ideas, and memories are really just active networks of neurons in our brains [18, 40]. And we have a good idea where in the brain specific kinds of mental activity occur, *e.g.* [26, 37] But, to our knowledge, no one has any idea as to what kinds of networks correspond to any specific sensation, concept or memory.

We know that neurons can stimulate other neurons by means of electric (or chemical) charges proceding along an axon to one, or more, synapses [39]. That would suggest that a directed, asymmetric network is a reasonable model. However, such an asymmetric network may best model neuronal behavior, but not neuronal state. Many neurons are interconnected by dendrites. These are thought to be bi-directional, thus implementing symmetric relationships that may recognize a state necessary to activate a neuron.

Given this state of uncertainty, we have chosen to explore symmetric relationships, or graphs or networks, in this paper. Some of the mathematical results we present may be true as well for asymmetric (directed) networks; some would require minor rewording; and some will no longer be true at all.

Regardless of whether our neural networks are essentially symmetric or asymmetric, it would appear that a mathematical treatment of networks, or graphs, or relationships is a fruitful way to approach them. That we will do in this paper which is an expanded version of [36].

In Section 2, we clarify our interpretation of relationships and their visual representation as graphs or networks. We also introduce the concept of "closure". In Section 3, we describe a computational process, ω, which reduces any network to its unique, irreducible "trace". We will claim that this procedure appears to model the process of long-term memory "consolidation".

The ω process is a well-defined function over the space of all finite networks in that for any network $\mathcal{N}$, ω yields a unique irreducible trace $\mathcal{T}$. Thus the inverse set, ω^{-1}, defines the abstract set of all networks that reduce to the same specific irreducible trace. In Section 4, we present a computational process which generates specific members within ω^{-1}. We will argue that this can model "memory reconstruction".

In Section 5 we present additional evidence to support our claims to model long-term memory consolidation and recall. Certain mathematical details are developed in an appendix, Section 6.

2 Sets, Relations, and Closure

Our computations are set based. The nature of the elements comprising the sets play no part, and can be quite arbitrary. So unlike most computational systems in which the variables will be `int` or `float`, our variables have type `setid`. We program using a set manipulation package in C++ with operators such as `is_contained_in` and `union_of`. Sets themselves are represented as extensible bit strings, so that the operators above are effectively of order O(1). There is no theoretical upper bound of these sets, but we have not tested it with sets of cardinality exceeding 50,000. A somewhat fuller description is given in [28]. All the following set-based operators and procedures have been implemented, and fully tested, using this system.

We use a standard set notation. A **set** S is comprised of elements $\{a, b, \ldots, y, z\}$ of unspecified type. The curly braces $\{ \}$ indicate that these elements are regarded as a "set". Sets are denoted by upper case letters, $e.g.\ X, Y$; elements are always lower case, $e.g.\ x, y$. Sometimes we elide the commas, as in $Y = \{abc\}$.

If an element x is a **member of** the set X, we write $x \in X$. If a set X is **contained in** another set Y, that is, $x \in X$ implies $x \in Y$ (here x is a variable running over all elements of X), we write $X \subseteq Y$. If the containment is **strict**, that is there exists $y \in Y$, $y \notin X$, we write $X \subset Y$. By $X \cup Y$ and $X \cap Y$ we mean the union and intersection (meet) of X and Y respectively.

One may have a "set of sets", which we call a **collection**, and denote with a caligraphic letter. Thus we may have $X \in \mathcal{C}$.

2.1 Relationships

Let S be any set. A **relation**, η, on S is a function, which given any subset $Y = \{y_1, y_2, \ldots, y_k\} \subseteq S$ returns the related set $Y.\eta = \{z_1, z_2, \ldots, z_n\} \subseteq S$. This is a bit unusual. It is more common to think of relations as links, or edges, between elements, such as illustrated by the undirected graph, or network, of Figure 1, which we will use as a running example. η is sometimes regarded as a set of "edges" in a graph theoretic approach. But we prefer to define relations in terms of sets and functional operators. It provides an additional measure of generality which can be of value. We emphasize this set-based definition by using suffix notation, such as $Y.\eta$ to mean the set of elements $\{z\}$ that are related to Y by η. We call $Y.\eta$ the **neighborhood** of Y. In Figure 1, $\{a\}.\eta = \{a, d, f\}$, $\{d\}.\eta = \{a, b, d, f, g\}$, and $\{a, e\}.\eta = \{a, b, c, d, e, f, h\}$.

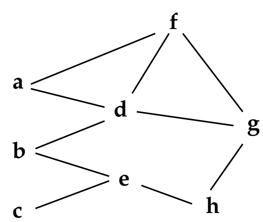

Figure 1: A small network illustrating neighborhood properties

If η has the property that

P1: $Y.\eta = \bigcup_{y \in Y} \{y\}.\eta$ (extensibility)

that is, $Y.\eta$ is the union of all the subsets $\{y\}.\eta$ for all $y \in Y$, we say that η is **extensible**, or *graphically representable*, so that Figure 1 is an accurate representation of η.

If the relationship is not extensible, then it constitutes a "hypergraph" [4, 12]. To more easily illustrate the concepts of this paper with graphs, we will assume that η is extensible; but unless explicitly noted none of the mathematical assertions require it. Moreover, we observe that for large, sparse relationships, matrix representations and operations are quite impractical [44].

In addition to extensibility, P1, a relationship η may also have any of the following 3 properties: that for all $X, Y \subseteq S$,

P2: $Y \subseteq Y.\eta$ (expansive or reflexive)[1]
P3: $X \subseteq Y$ implies $X.\eta \subseteq Y.\eta$ (monotone)[2]
P4: $X.\eta = Y$ implies $Y.\eta = X$ (symmetric)[3]

The relation of Figure 1 is symmetric; its graph is **undirected**. By a **network**, $\mathcal{N} = (N, \eta)$, we mean a set N of nodes or elements, together with any relationship η. For this paper, we require that η satisfy the functional properties P2, P3 and P4.

2.2 Closure

The mathematical concept of "closure" plays a key role in our approach. In a discrete world, the interpretation of closed sets is somewhat different from the more traditional concepts encountered in classical point-set topology. Our view is that a closure operator, φ, is a set-valued function whose domains are also sets. If Y denotes any set, $Y.\varphi$ denotes its closure; that is it is the smallest closed set containing Y. Thus, like η, φ is a well-defined function mapping subsets, $X, Y \subseteq N$ of a given network into other subsets of N. More formally, φ is a **closure operator** that satisfies the following 3 closure axioms, C1: expansive ($Y \subseteq Y.\varphi$), C2: monotone ($X \subseteq Y$ implies $X.\varphi \subseteq Y.\varphi$), and C3: idempotent ($Y.\varphi.\varphi = Y.\varphi$). Readily, any relationship operator, η, satisfying properties P2 and P3 is *almost* a closure operator. It has only to satisfy the idempotency axiom. But normally, $Y.\eta \subset Y.\eta.\eta$ since neighborhoods tend to grow.

An alternative definition of closure asserts that a collection $\mathcal{C} = \{C_1, \ldots, C_n\}$ can be regarded as the closed sets of a superset S if and only if C4: the intersection $C_i \cap C_k$ of any these closed sets is itself closed (in $\mathcal{C}$). It is not difficult to prove that C4 implies C1, C2, and C3, and conversely.

We normally think of closure in terms of its operator definition. Because φ is expansive, C1, the superset S must be closed; by C4, if any two closed sets are disjoint, the empty set Ø must also be closed.

2.3 Neighborhood Closure

One important closure operator φ, called **neighborhood closure**, can be defined with respect to network relationships. We let

$$Y.\varphi = \bigcup_{z \in Y.\eta} \{\{z\}.\eta \subseteq Y.\eta\} \tag{1}$$

[1] This is primarily for mathematical convenience.
[2] Probably essential. If η is not monotone, we can prove very few mathematical results of interest.
[3] Unnecessary, relaxed in other papers such as [34, 35].

That is, if $z \in Y.\varphi$ then z is not related to any elements that are not already related to Y. Readily, since $\{a\}.\eta = \{adf\} \subseteq \{abdfg\} = \{d\}.\eta$ in Figure 1, $a \in \{d\}.\varphi$. Convince yourself that $\{d\}.\varphi = \{adf\}$ and $\{h\}.\varphi = \{h\}$. It is not hard to show that φ, so defined with respect to η satisfies the closure axioms C1, C2 and C3 and that for all $\{y\}$, $\{y\} \subseteq \{y\}.\varphi \subseteq \{y\}.\eta$.

3 Irreducible Networks

We are fairly sure that all mental activity, that is sensory apprehension, cognition, and ideation are creatures of our brain's neural system and we know what parts of the brain this activity is located [17, 20, 40]; but what configurations of activated neurons might correspond to a particular experience or idea is totally unknown. The 206 interconnected nodes of Figure 2, some of which have been labeled with

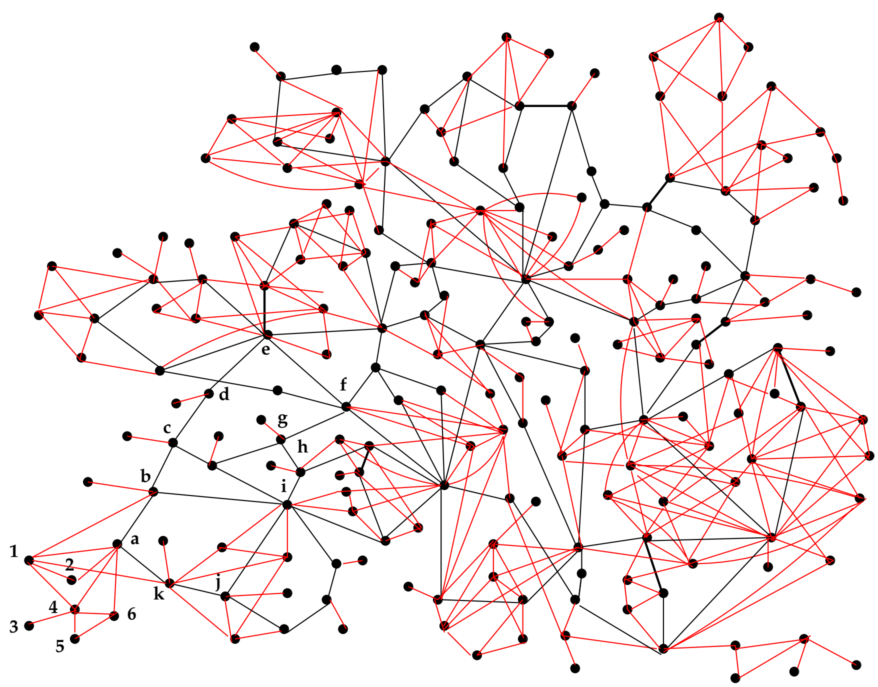

Figure 2: A moderately complex network whose nodes and links might model a neural configuration [5, 6].

letters and integers, may be thought to schematically represent a momentary configuration of neurons in a mental process. Assuming this, we ask how could this mental experience be "remembered"?

If a singleton set $\{y\}$ is not closed, say $z \in \{y\}.\varphi$, then $\{z\}.\eta \subseteq \{y\}.\eta$, so z contributes little to understanding the structure of η in terms of closure. In Figure 2, $\{4\}.\varphi = \{3, 4, 5, 6\}$. Removing the nodes 3, 5, and 6, and their connections to 4, results in minimal information loss with respect to η as a whole. However, with these nodes gone, we now have $\{4\}.\eta = \{a, 1, 4\} \subseteq \{a, 1, 2, 4\} = \{a\}.\eta$. So the two elements, 1, 4, and their connections can be removed as well. In fact the entire pendant substructure $\{1, 2, 3, 4, 5, 6\}$ on $\{a\}$ can be removed with no loss of global information. If it is removed $\{a\}.\varphi = \{a\}$, so $\{a\}$ is closed.

We say a network $\mathcal{N} = (N, \eta)$ is **irreducible** if every singleton set, $\{y\}$, is closed. That is, if for all $y \in N$, $\{y\}.\varphi = \{y\}$. In Figure 2, $\{f\}.\varphi = \{f\}$, but from observations above, the entire network is not irreducible.

If $\{y\}$ is not closed, only elements z in $\{y\}.\eta$ could possibly be in $\{y\}.\varphi$ so only those need be considered. If $\{z\}.\eta \subseteq \{y\}.\eta$ so that $\{z\}.\varphi \subseteq \{y\}.\varphi$, we say z is **subsumed** by y, or z **belongs** to y. We can remove z from N, together with all its connections, and add z to $\{y\}.\beta$, the set of all nodes belonging to $\{y\}$ which we call its β**-set**. Since, $y \in \{y\}.\beta$, its cardinality, or β**-count**, $|\{y\}.\beta| \geq 1$, a value we will use in the next section. We use the pseudocode of Figure 3 to implement the process ω that reduces any network $\mathcal{N}$ to its irreducible core, which is called its **trace**, $\mathcal{T}$.[4] This version of ω only records β-counts, not entire β-sets.

```
while there exist reduceable nodes
   {
   for_each {y} in N
      {
      get {y}.nbhd;
      for_each {z} in {y}.nbhd - {y}
         {
         if ({z}.nbhd contained_in {y}.nbhd
            {         //  z is subsumed by y
            remove z from network;
            |{y}.beta| = |{y}.beta| + |{z}.beta|;
            }
         }
      }
   }
```

Figure 3: Reduction code, implementing ω

The irreducible trace, $\mathcal{T}$, of Figure 2 is shown in Figure 4. The trace is the dark network on the 83 elements with 234 bolder connections. Sets of subsumed nodes, or β-sets, have been encircled with dashed lines. The largest β-set, comprised of 12

[4]In [33], this was called the "spine" of $\mathcal{N}$.

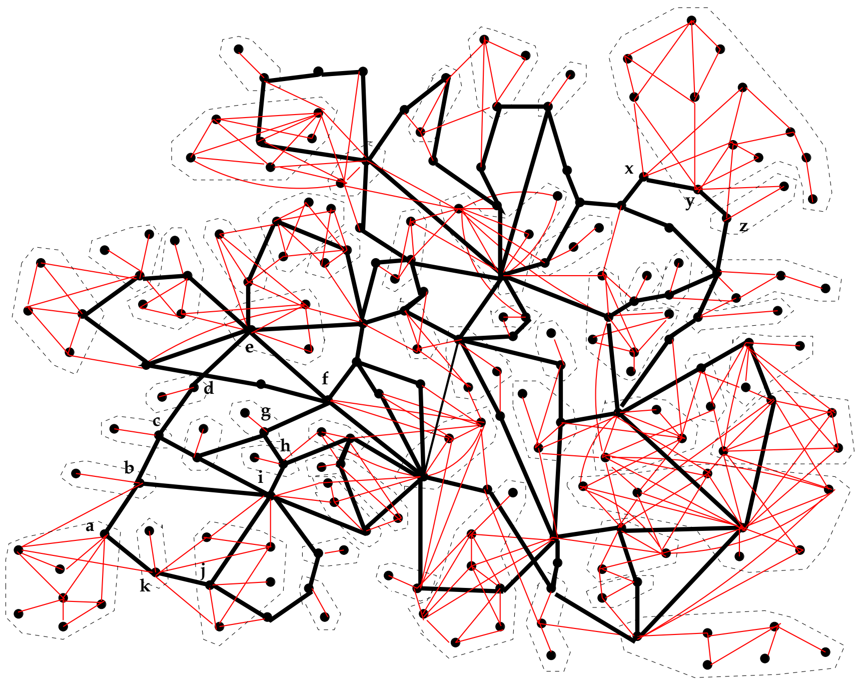

Figure 4: The irreducible trace $\mathcal{T}$ of Figure 2

nodes, is $\{y\}.\beta$ in the upper right corner. The β-set $\{a\}.\beta$ in the lower left corner consists of 7 nodes (including a itself). A total of 123 nodes were subsumed and eliminated.

In this particular network the outer loop of Figure 3 was executed 4 times with a 5th pass to verify that there remained no more reduceable nodes. The order in which individual nodes y are examined is arbitrary. One can create networks that require $n = |N|$ iterations of this outer loop. So this process has a theoretical complexity of $O(n^2)$. However, in tests with rather complex networks of several thousand nodes, the maximal number of iterations has never exceeded 7. Its effective complexity appears to be quite reasonable. Moreover, because of its local nature, the inner loop could be easily implemented in parallel.

We keep speaking of the *function* ω. It can be shown (see Section 6), that for any network $\mathcal{N}$, its irreducible trace, $\mathcal{T}$, is *unique* (up to isomorphism). Therefore, the pseudocode of Figure 3 does indeed embody a well-defined computational function which we denote by ω. Not only is ω a function, we can actually characterize its image, $\mathcal{T} = \mathcal{N}.\omega$. When η is symmetric, if y is a node in $\mathcal{T} = \mathcal{N}.\omega$, then y will be either: (a) an isolated node; (b) an element of a chordless cycle of length ≥ 4; or (c) an element in a path between two chordless cycles of length ≥ 4 (again see Section 6).

3.1 Chordless Cycles

A **chordless cycle** is most easily visualized as a necklace of pearls (or beads). More formally, it is a sequence $< y_1, y_2, \ldots, y_n, y_1 >$ where $y_{i\pm 1} \in \{y_i\}.\eta$, $y_1 \in \{y_n\}.\eta$ and $y_{i\pm k} \notin \{y_i\}.\eta$ if $k > 1$. In Figure 1 the 5-cycle $< b, d, g, h, e, b >$ is chordless. In Figures 2 and 4 the 5-cycle $< a, b, i, j, k, a >$ is also chordless, as is the 8-cycle $< b, c, d, e, f, g, h, i, b >$. But the combined 11-cycle $< a, b, c, d, e, f, g, h, i, j, k, a >$ is not. The link (b, i) is a chord.

Granovetter [19] called chordless cycles the "weak connections" of a social network. He felt they were the key to understanding the network structure as a whole. However, chordless cycle structures have been relatively unstudied, while "chordal graphs" (with no chordless cycles) have an exhaustive literature [25]. Even when η is not symmetric, chordless cycles are basic to the characterization of an irreducible trace [35].

We can explore the mathematical properties of chordless networks a bit further. If some collection, $\mathcal{S}$, of subsets C_i of elements in S has the property that no set C_i is completely contained in another, then $\mathcal{S}^n$, where $n = |S|$, constitutes a Sperner system.[5] Such systems are so called after Emanuel Sperner who first described them [11]. With a little thought, we see that if C_i and C_k are chordless cycles, then C_i cannot contain C_k. Consider C_1 and C_2 in Figure 2, C_2 is a subset of C_1 so C_1 cannot be chordless, and indeed it is not. Consequently, a system of chordless cycles $\{C_i\}$ constitutes a Sperner system, $\mathcal{S}^n$, with each cycle being a unique set in this system. This allows us to treat chordless cycles as sets of elements, without considering links or edges.

The primary interest in Sperner systems has been combinatoric, that is counting how many distinct Sperner systems on n elements can exist. That number is exceedingly large. For example, if $n = 7$, there are $S^7 = 7,581$ distinct configurations where no subset is contained in another. This combinatoric result assures us that if memories are encoded as chordless cycles, it will be very rich in coding possibilities.

This trace, $\mathcal{T}$, of chordless cycles preserves a number of important properties found in the original network, $\mathcal{N}$. First, it preserves the shortest path structure between retained nodes. Consequently, connectivity and the distances between nodes (as usually defined) are preserved. Further, "network centers", [3, 13, 14], whether with respect to distance or "betweenness", are preserved in the trace.

When η is not symmetric, the reduction process ω still yields a unique irre-

[5] A collection of subsets $\{C_i\}$ is a Sperner system if $C_i \not\subset C_k$ for all $i \neq k$.

ducible network $\mathcal{T}$; but the characterization preceeding Section 3.1 above no longer holds. Instead, when η is not symmetric, then for all $y \in \mathcal{T} = \mathcal{N}.\omega$, if there exists $z \in \{y\}.\eta$ then there exists a directed path from y through z that terminates in a chordless cycle of length ≥ 4. That path may itself be the chordless cycle.

3.2 Distribution of Chordless Cycles

The combinatorics of Sperner sets provides one mechinism for encoding information. The distribution of cycle lengths in a single irreducible network provides another. Counting the numbers of cycles of length k in a specific chordless network, $\mathcal{N}$, is not easy. Using the Sperner set property, the author has employed a brute force counting process that is limited by the size of $\mathcal{N}$. Diane Castonguay, Elisangela Silva Dias *et al.* have developed more effective ways of counting [9, 21]. In [33], this distribution is called the *signature* of the network. The average cycle

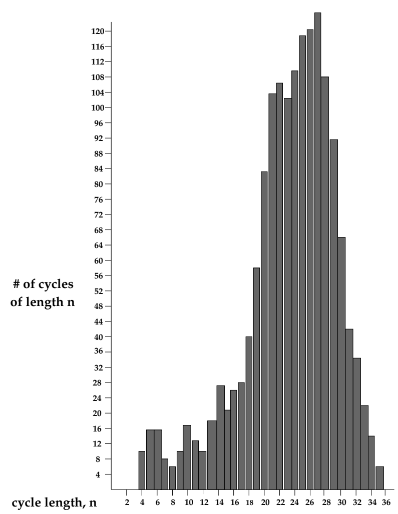

Figure 5: Distribution of chordless cycles in the irreducible trace of Figure 4.

length is 23.4 in this network and the 6 longest chordless cycles have length 35.

3.3 Consolidation

The physical nature of human long-term memory is not at all a settled matter. We are fairly certain that the hippocampus of the brain is heavily involved [1, 16, 37]; but just how is not completely understood. One school of thought posits that long-term memories are recorded in some form of "memory trace" [7, 42, 46, 38]. But, because no trace of these supposed "memory traces" has ever been physically detected (pun intended), others disbelieve this theory [8, 27]. We explore the possibility of biological "memory traces" in Section 5.2.

There is more consensus that some form of processing which distinguishes long-term memory from short-term memory does occur. This process is commonly called **consolidation** [2, 23, 27]. We believe that ω is analogous to consolidation, and that chordless cycles, in some form, are analogous to the elusive "memory trace", whence our terminology.

4 Computing Similar Networks

Since ω is a well-defined function mapping the space of all finite, symmetric networks into itself, one can consider $\mathcal{N}.\omega^{-1}$, which is the collection of all networks $\mathcal{N}_i$ such that $\mathcal{N}_i.\omega = \mathcal{T} = \mathcal{N}.\omega$. Two such networks, $\mathcal{N}_i$ and $\mathcal{N}_k$, that have the same irreducible trace are said to be **structurally similar**. Readily, structural similiarity is an equivalence relation. Even though $\mathcal{N}_k$ may be similar to $\mathcal{N}_i$, they may have very different cardinalities. A network $\mathcal{N}_k(N_k, \eta_k)$ is said to be **strongly similar** to $\mathcal{N}_i(N_i, \eta_i)$ if $\mathcal{N}_k.\omega = \mathcal{N}_i.\omega$ and $|N_k| = |N_i|$.

The pseudocode below in Figure 6 describes a computational process ε that, given the trace $\mathcal{T}$ of a network $\mathcal{N}$ together with β-counts, randomly expands it to a strongly similar network $\mathcal{N}' = \mathcal{N}.\omega.\varepsilon$. The process **choose_random_in** returns

```
for all {y} in N
    {
    while (|{y}.beta| > 1)
        {
        create new node z;
        S = choose_random_in ({y}.nbhd);
        {z}.nbhd = S;
        k = random_int(1, |{y}.beta|-1);
        |{y}.beta| = |{y}.beta| - k;
        |{z}.beta| = k;
        add {z} to N;
        }
    }
```

Figure 6: Pseudocode for ε which generates strongly similar networks.

a random subset of its argument. Since $\{z\}.\eta = S \subseteq \{y\}.\eta$, the node z will be subsumed by (or belong to) y if reduced again ensuring that $\mathcal{N}.\omega.\varepsilon.\omega = \mathcal{N}.\omega$. When a node $\{y\}$ is expanded, its β-count is decremented, and if > 1, part of the remainder may be added to the β-count of $\{z\}$. Consequently, by creating just as many new nodes as had belonged to any node $\{y\}$, we ensure that $|N'| = |N|$. This kind of ε process has been called an “expansion grammar” in [31]. The construction of ε, where $\{z\}.\eta \subseteq \{y\}.\eta$, assures us that $\mathcal{T}.\varepsilon.\omega$ will be $\mathcal{T}$ again. Consequently, for any network $\mathcal{N}' = \mathcal{T}.\varepsilon$, $\mathcal{N}' \in \mathcal{N}.\omega^{-1}$, so $\mathcal{N}'$ and $\mathcal{N}$ are structurally similar.

Let $\mathcal{N}$ be the network of Figure 2. The following Figure 7 shows a network $\mathcal{N}'$ that was randomly expanded by ε, given the irreducible trace $\mathcal{T}$ of Figure 4. The numbered

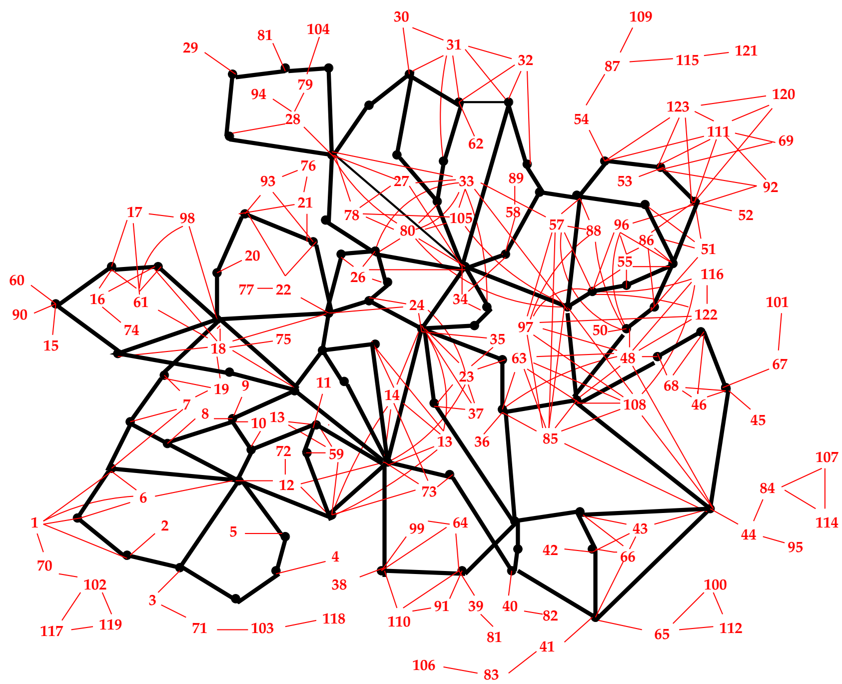

Figure 7: A reconstructed network $\mathcal{N}' = \mathcal{T}.\varepsilon$ in $\mathcal{N}.\omega^{-1}$ that is strongly similar to $\mathcal{N}$ of Figure 2

nodes were randomly appended to the trace and roughly correspond to the 123 subsumed nodes. They are numbered in the order that they were attached to the darker irreducible trace. $\mathcal{N}'$ is strongly similar to the network $\mathcal{N}$ of Figure 2 because $\mathcal{N}'.\omega = \mathcal{T} = \mathcal{N}.\omega$.

Such a semi-random “retrieval” process may be inappropriate in computer applications [35], but it seems to model biological recall rather well. It has been observed that the recall and reconstruction of our long-term memories is seldom exact [23]. Our memories often are confused with respect to detail, even when they are generally correct. Reconstruction of a network trace by ε has these very properties.

Given that for all networks $\mathcal{N}$, $\mathcal{N}.\omega.\varepsilon.\omega = \mathcal{N}.\omega = \mathcal{T}$, it also supports the notion of “re-consolidation” which asserts than long-term memories are repeatedly recalled and re-written with no change, unless deliberately distorted in our (semi)conscious mind [27, 45].

5 Biological Memory

A computational model need not actually explain the behavior that it models. For example, the path of a thrown projectile has an excellent parabolic model. However, further study of this conic formulation contributes little to the understanding of either gravity or air resistence. By the same token, there need not be closure operators or chordless cycles involved in the performance of human memory, for the model to be valid. But, it would be a powerful verification of this model if we could demonstrate the existence of chordless cycle structures in a memory representation. We can't. Neither, to our knowledge, does anyone else know the the structural format of our long-term memory.

Throughout this paper we have suggested parallels found in various memory studies. But, do these computational processes, ω and ε, really model biological memory? We just don't know. Are long-term memories really encoded as chordless cycles? In this section we offer a few more tantalizing clues which may, or may not, be significant.

5.1 Role of Closure

We employed "closure" as the basic mathematical concept in the preceeding development. But, are instances of closure actually found in biological organisms? We offer two suggestive examples.

First, The visual pathway consists of layers of cells, beginning with the rods and cones of the retina passing stimuli toward the primary visual cortex. The neurological structure of this visual pathway is reasonably well understood, *c.f.* [18, 39, 43]. The individual functions of its layers are less well so.

Imagine that Figure 8 depicts a cross section of the retinal region. Dark cells denote

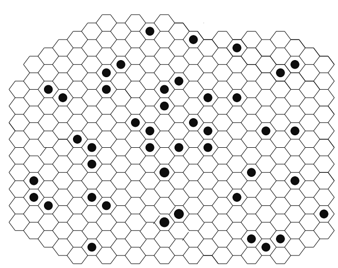

Figure 8: Excited cells in a cross section of the visual cortex.

visually excited cells. Although tightly packed, the actual neuronal structure is not as regular as this hexagonal grid; but this regularity plays no part in the process.

Let α be an existential operator defined as $Y.\alpha = black$ (excited), if and only if $\exists z \in Y.\eta$ where z is $black$ (excited). Let β be the existential operator defined by $Y.\beta = white$ (quiesent), if and only if $\exists z \in Y.\eta$ such that z is $white$ (quiesent). Figure 9(a) illustrates the excited (small $\times$) neighbors of Figure 8. Figure 9(b) illustrates $Y.\alpha.\beta$ in which all excited cells of Figure 9(a) that have at least one quiesent (white) neighbor become quiesent (white). The resulting central figure becomes evident; it is a closed object, because the pair

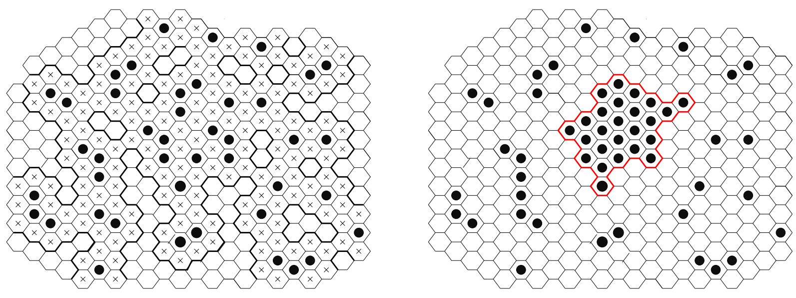

Figure 9: (a) $Y.\alpha$, excited cells, (b) $Y.\alpha.\beta$, remaining excited cells.

of operators $(\alpha.\beta)$ is a closure operator. The pair $(\alpha.\beta)$ is idempotent because iterating them, as in $Y.(\alpha.\beta).(\alpha.\beta)$ yields no new black (excited) cells.

This two step operation can occur at the neural firing rate. It is an effective parallel process that was first proposed to eliminate salt and pepper noise in computer imagery [41]. Readily, such a "blob detection" capability would have evolutionary value. Does such a capacity exist? We don't know for sure. But, it is thought that the visual pathway is organized in an alternating manner to facilitate precisely this kind of two-step processing [43]. In any case, this example illustrates that this kind of all, or nothing, logic in which "for all" $(\forall x)$ can be interpreted as "there does not exist" $(\neg\exists\neg x)$ needed to implement the closure operator of (1) can be rendered in neural circuitry.

The second example is also "cognitive". In the development of "Knowledge Spaces" [10], Doignon and Falmagne call a coherent collection of facts or skills a "knowledge state". These are closed sets which are partially ordered by containment to form a lattice structure [30], which they call a "knowledge space". There is a considerable literature concerning closed knowledge "states" and knowledge "spaces".[6] A somewhat similar approach to cognitive closure was presented in [34]. Closure operators can be an important aspect of cognitive behavior.

5.2 Role of Chordless Cycles

Also central to our paper is the concept of "chordless cycles" which constitute the structure of an irreducible trace. Surprisingly, chordless cycles abound in all biological organisms as protein polymers.

One example, found in every cell of our bodies, is a 154 node phenylalaninic-glycine-repeat (nuclear pore protein), $\mathcal{N}$, which is shown in Figure 10.[7] One can easily see the chordless loops, with various linear tendrils attached to them. When these are removed by

[6] Cord Hockemeyer, http://www.uni-graz.at/cord.hockemeyer/KST_Bibliographie/kst-bib.html, maintains a bibliography of over 400 related references.

[7] This network, $\mathcal{N}$, that we received from a lab at Johns Hopkins Univ. was only identified as $GrN2$. We believe it is an natively unfolded phenylalanine-glycine (FG)-repeat [24].

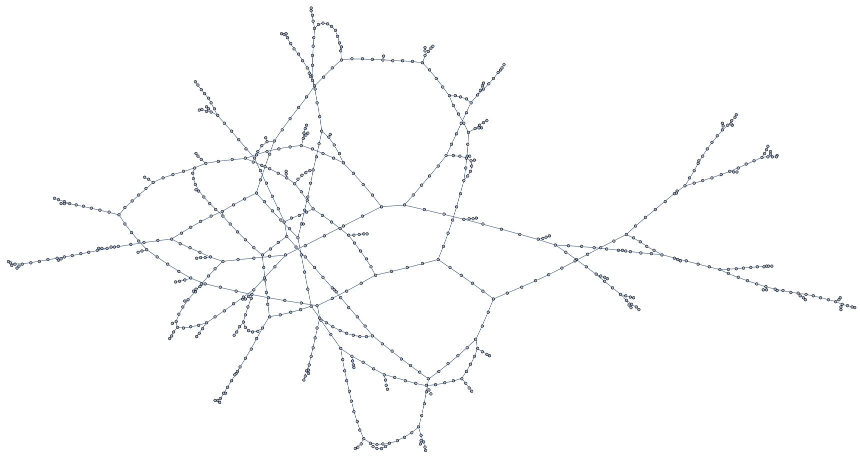

Figure 10: A 154 node protein polymer

ω, there are 107 remaining elements involved in the chordless cycle structure. These are thought to regulate transport of other proteins across the nuculear membrane [15, 29, 47].

Readily, organisms with any form of memory, *e.g.* "movement toward light yields food", have survival benefit. Nature appears to reuse successful structures. If chordless cycles can successfully regulate one form of transport, it would not be surprising if evolutionary pressure led to their use in other control mechanisms.

Moreover, modification of protein polymers by means of phosphorylation [48] is thought to be involved in short-term memory [23]. For long-term memories, chordless cycles within the dendritic connections between pyramidal neurons seem more likely [22, 38].

But is there reason to suspect that memory has any "structural" properties at all?

Perhaps the most important biological memory mechanism is our genetic memory which records the nature of our species. It is known to have a double helix structure which facilitates a near perfect recall. These coded sequences are subsequently "expressed" during development by an expansion process which might be similar to a *non-random* ε.

While the double helix facilitates a reliable read-only memory (ROM); chordless cycles appear facilitate the encoding of eposidic information in a dynamic memory via a consolidation process such as ω.

Much of this section is speculation. But, both "closure" [34] and "chordless cycles" [35] would appear to have biological significance. The assertions of this paper have a solid mathematical base. As such, ω and ε provide useful examples within a category of *networks* that can formally model dynamic biological networks. If in addition, they actually model memory consolidation and recall as we suspect, that would be an additional bonus.

6 Appendix

Too much formal mathematics makes a paper hard to read. Yet, it is important to be able to check some of the statements made in the body of the paper. In this appendix we provide a few propositions to formally prove some of our assertions.

The order in which nodes, or more accurately the singleton subsets, of $\mathcal{N}$ are encountered can alter which points are subsumed and subsequently deleted. Nevertheless, we show below that the reduced trace $\mathcal{T} = \mathcal{N}.\omega$ will be unique, up to isomorphism.

Proposition 6.1 *Let $\mathcal{T} = \mathcal{N}.\omega$ and $\mathcal{T}' = \mathcal{N}.\omega'$ be irreducible subsets of a finite network $\mathcal{N}$, then $\mathcal{T} \cong \mathcal{T}'$.*

Proof: Let $y_0 \in \mathcal{T}$, $y_0 \notin \mathcal{T}'$. Then y_0 can be subsumed by some point y_1 in $\mathcal{T}'$ and $y_1 \notin \mathcal{T}$ else because $y_0.\eta \subseteq y_1.\eta$ implies $y_0 \in \{y_1\}.\varphi$ and $\mathcal{T}$ would not be irreducible.
Similarly, since $y_1 \in \mathcal{T}'$ and $y_1 \notin \mathcal{T}$, there exists $y_2 \in \mathcal{T}$ such that y_1 is subsumed by y_2. So, $y_1.\eta \subseteq y_2.\eta$.
Now we have two possible cases; either $y_2 = y_0$, or not.
Suppose $y_2 = y_0$ (which is often the case), then $y_0.\eta \subseteq y_1.\eta$ and $y_1.\eta \subseteq y_2.\eta$ or $y_0.\eta = y_1.\eta$. Hence $i(y_0) = y_1$ is part of the desired isometry, i.
Now suppose $y_2 \neq y_0$. There exists $y_3 \neq y_1 \in \mathcal{T}'$ such that $y_2.\eta \subseteq y_3.\eta$, and so forth. Since $\mathcal{T}$ is finite this construction must halt with some y_n. The points $\{y_0, y_1, y_2, \ldots y_n\}$ constitute a complete graph Y_n with $\{y_i\}.\eta = Y_n.\eta$, for $i \in [0, n]$. In any reduction all $y_i \in Y_n$ reduce to a single point. All possibilities lead to mutually isomorphic maps. □

In addition to $\mathcal{N}.\omega$ being unique, we may observe that the transformation ω is functional because ω maps all subsets of N onto N_ω. So we can have $\{z\}.\omega = \emptyset$, thus "deleting" z. Similarly, ε is functional because $\emptyset.\varepsilon = \{y\}$ provides for the inclusion of new elements. Both ω and ε are monotone, if we only modify its definition to be $X \subseteq Y$ implies $X.\varepsilon \subseteq Y.\varepsilon$, *provided* $X \neq \emptyset$.

The following proposition characterizes the structure of irreducible traces.

Proposition 6.2 *Let $\mathcal{N}$ be a finite symmetric network with $\mathcal{T} = \mathcal{N}.\omega$ being its irreducible trace. If $y \in \mathcal{T}$ is not an isolated point then either*

(1) there exists a chordless k-cycle C, $k \geq 4$ such that $y \in C$, or

(2) there exist chordless k-cycles C_1, C_2 each of length ≥ 4 with $x \in C_1$ $z \in C_2$ and y lies on a path from x to z.

Proof: (1) Let $y \in N_{\mathcal{T}}$. Since y is not isolated, we let $y = y_0$ with $y_1 \in y_0.\eta$, so $(y_0, y_1) \in E$. Since y_1 is not subsumed by y_0, $\exists y_2 \in y_1.\eta, y_2 \notin y_0.\eta$, and since y_2 is not subsumed by y_1, $\exists y_3 \in y_2.\eta$, $y_3 \notin y_1.\eta$. Since $y_2 \notin y_0.\eta$, $y_3 \neq y_0$.
Suppose $y_3 \in y_0.\eta$, then $< y_0, y_1, y_2, y_3, y_0 >$ constitutes a k-cycle $k \geq 4$, and we are done.
Suppose $y_3 \notin y_0.\eta$. We repeat the same path extension. $y_3.\eta \not\subseteq y_2.\eta$ implies $\exists y_4 \in y_3.\eta$, $y_4 \notin y_2.\eta$. If $y_4 \in y_0.\eta$ or $y_4 \in y_1.\eta$, we have the desired cycle. If not $\exists\ y_5, \ldots$ and so forth. Because $\mathcal{N}$ is finite, this path extension must terminate with $y_k \in y_i.\eta$, where $0 \leq i \leq n-3$, $n = |N|$. Let $x = y_0, z = y_k$.
(2) follows naturally. □

Finally, we show that ω preserves the shortest paths between all elements of the trace, $\mathcal{T}$.

Proposition 6.3 *Let $\sigma(x, z)$ denote a shortest path between x and z in $\mathcal{N}$. Then for all $y \neq x, z, \in \sigma(x, z)$, if y can be subsumed by y', then there exists a shortest path $\sigma'(x, z)$ through y'.*

Proof: We may assume without loss of generality that y is adjacent to z in $\sigma(x, z)$.
Let $< x, \ldots, x_n, y, z >$ constitute $\sigma(x, z)$. If y is subsumed by y', then $y.\eta = \{x_n, y, z\} \subseteq y'.\eta$. So we have $\sigma'(x, z) =< x. \ldots, x_n, y', z >$ of equal length. (Also proven in [32].) $\Box$

In other words, z can be removed from $\mathcal{N}$ with the certainty that if there was a path from some node x to z through y, there will still exist a path of equal length from x to z after y's removal.

Figure 11 visually illustrates the situation described in Proposition 6.3, which we call a **diamond**. There may, or may not, be a connection between y and y' as indicated by

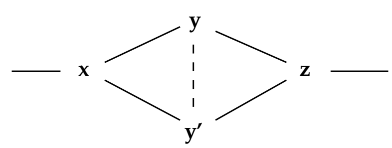

Figure 11: A network diamond

the dashed line. If there is, as assumed in Proposition 6.3, then either y' subsumes y or *vice versa*, depending on the order in which y and y' are encountered by ω. This provides one example of the isomorphism described in Proposition 6.1. If there is no connection between y and y' then we have two distinct paths between x and z of the same length.

References

[1] Gareth R. I. Barker, Paul J. Banks, Hannah Scott, G. Scott Ralph, and et.al. Separate elements of episodic memory subserved by distinct hippocampal-prefrontal connections. *Nature Neuroscience*, pages 1–28, 2017. DOI: 10.1038/nn.4472.

[2] Clive R. Braham and Eloucine Messaoudi. BDNF function in adult synaptic plasticity; The synaptic consolidation hypothesis. *Progress in Neurobiology*, 76(2):99–125, 2005.

[3] Ulrik Brandes. A Faster Algorithm for Betweeness Centrality. *J.Mathematical Sociology*, 25(2):163–177, 2001.

[4] Alain Bretto. *Hypergraph Theory: an Introduction*. Springer, 2013.

[5] Ed Bullmore and Olaf Sporns. Complex brain networks: graph theoretical analysis and functional systems. *Nature Reviews: Neuroscience*, 10:186–198, March 2009.

[6] Edward T. Bullmore and Danielle S. Bassett. Brain Graphs: Graphical Models of the Human Brain Connectome. *Annual Review of Clinical Psychology*, 7:113–140, 2011.

[7] Robert G. Crowder. *Principles of Learning and Memory*. Psychology Press, New York, 2015.

[8] Lila Davachi. Encoding: The proof is still required . In H.L. Roediger III, Y. Dudai, and S.M. Fitzpatrick, editors, *Science of Memory: Concepts*, pages 129–135. 2007.

[9] Elisângela Silva Dias, Diane Castonguay, Humberto Longo, and Walid A. R. Jradi. Efficient enumeration of chordless cycles. *arXiv:1309.1051v4*, 24 Nov. 2014.

[10] Jean-Paul Doignon and Jean-Claude Falmagne. *Knowledge Spaces*. Springer, Berlin, 1999.

[11] Konrad Engle. Sperner Theory . In Michiel Hazewinkle, editor, *Encyclopedia of Mathematics*. Springer, Berln, 2001.

[12] Martin Farber and Robert E. Jamison. Convexity in Graphs and Hypergraphs. *SIAM J. Algebra and Discrete Methods*, 7(3):433–444, July 1986.

[13] Linton C. Freeman. Centrality in Social Networks, Conceptual Clarification. *Social Networks*, 1:215–239, 1978/79.

[14] Linton C. Freeman. Going the Wrong Way on a One-Way Street: Centrality in Physics and Biology. *J. of Social Structure*, 9:1–15, 2008.

[15] Steffen Frey and Dirk Görlich. A Saturated FG-Repeat Hydrogel Can Reproduce the Permeability of Nuclear Pore Complexes. *Cell*, 130:512–523, Aug. 2007.

[16] J.D.E. Gabrieli. Cognitive Neuroscience of Human Memory. *Annual Review of Psychology*, 49:87–115, 1998.

[17] John D.E. Gabrieli, John E. Desmond, Jonathan B. Demb, Anthony D. Wagner, Maria V. Stone, Chandan J. Vaidya, and Gary H. Glover. Functional Magnetic Resonance Imaging of Semantic Memory Processes in the Frontal Lobes. *Psychological Science*, 7(5):278–283, Seot, 1996.

[18] Michael S. Gazzaniga, Richard B. Ivry, George R. Mangun, and Megan S. Steven. *Cognitive Neuroscience, The Biology of the Mind*. W.W. Norton, New York, 2009.

[19] Mark S. Granovetter. The Strength of Weak Ties. *Amer. J. of Sociology*, 78(6):1360–1380, 1973.

[20] Annika Hultén, Hannu Laaksonen, Matti Laine, and Riitta Salmelin. Modulation of Brain Activity after Learning Predicts Long-Term Memory for Words. *J. of Neuroscience*, 30(45):15160–15165, Nov. 2010.

[21] Walid A. R. Jradi, Elisângela S. Dias, Diane Castonguay, Humberto Longo, and Hugo A. D. do Nascimento. A GPU-based parallel algorithm for enumerating all chordless cycles in graphs. *arXiv:submit*, 1288731:1–15, June 2015.

[22] David LaBerge and Ray Kasevich. The apical dendrite theory of consciousness. *Science Direct*, 20:1004–1020, 2007.

[23] Joseph E. LeDoux. Consolidation: Challenging the traditional view. In H.L. Roediger III, Y. Dudai, and S.M. Fitzpatrick, editors, *Science of Memory: Concepts*, pages 171–175. 2007.

[24] Roderick Y. H. Lim, Ning-Ping Huang, Joachim Köser, and Jie Deng *et al.* Flexible phenylalanine-glycine nucleoporins as entropic barriers to nucleocytoplasmic transport. *Proc. Nat. Acad. Science (PNAS)*, 103(25):9512–9517, June 20 2006. DOI:10.1073/pnas.0603521103.

[25] Terry A. McKee. How Chordal Graphs Work. *Bulletin of the ICA*, 9:27–39, 1993.

[26] Yashar Moshfeghi, Peter Triantafillou, and Frank E. Pollick. Understanding Information Need: an fMRI Study. In *SIGIR '16*, pages 335–344. Pisa, Italy, July 2016. DOI: http://dx.doi.org/10.1145/2911451.2911534.

[27] Lynn Nadel. Consolidation: The demise of the fixed trace . In H.L. Roediger III, Y. Dudai, and S.M. Fitzpatrick, editors, *Science of Memory: Concepts*, pages 177–181. 2007.

[28] Ratko Orlandic, John L. Pfaltz, and Christopher Takylor. A Functional Database Representation of Large Sets of Objects. In Hua Wang and Mohamed A. Saraf, editors, *25th Australasian Database Conference, ADC 2014*, pages 189–196, Brisbane, QLD, 2014.

[29] Samir S. Patel, Brian J. Belmont, Joshua M. Sante, and Michael F Rexach. Natively Unfolded Nucleoporins Gate Protein Diffusion across the Nuclear Pore Complex. *Cell*, 129:83–96, April 6 2007.

[30] John L. Pfaltz. Closure Lattices. *Discrete Mathematics*, 154:217–236, 1996.

[31] John L. Pfaltz. Neighborhood Expansion Grammars. In *Intern'l Workshop on Theory and Application of Graph Transformation, TAGT'98*, Lecture Notes in Computer Science, #1764, pages 30–44, Paderborn, Germany, Nov. 1998. Springer-Verlag.

[32] John L. Pfaltz. Finding the Mule in the Network. In Reda Alhajj and Bob Werner, editors, *Intern. Conf on Advances in Social Network Analysis and Mining, ASONAM 2012*, pages 667–672, Istanbul, Turkey, August 2012.

[33] John L. Pfaltz. The Irreducible Spine(s) of Discrete Networks. In Xuemin Li, Yannis Manolopoulos, Divesh Srivastava, and Guangyan Huang, editors, *Web Information Systems Engineering - WISE 2013*, volume LNCS # 6984, Part 2, pages 104–117), Nanjing, PRC, Oct. 2013.

[34] John L. Pfaltz. Using Closed Sets to Model Cognitive Behavior. In Tapabrata Ray, Ruhul Sarker, and Xiaodong Li, editors, *Proc. Australian Conf. on Artificial Life and Computational Intelligence (ACALCI 2016)*, volume LNCS 9592, pages 13–26, Canberra, ACT, 2016.

[35] John L. Pfaltz. A Role for Chordless Cycles in the Retrieval and Representation of Information. In Giorgio Orsi, editor, *Proc.of 6th Inter'l Workshop on Querying Graph Structured Data (GraphQ 2017) at EDBT/ICDT 2017*, pages 10–15, Venice, IT, March 2017.

[36] John L. Pfaltz. Computational Processes that Appear to Model Human Memory. In Daniel Figueiredo, Carlos Martín-Vide, Diogo Gratas, and Miguel A.Vega-Rodríguez, editors, *Proc. 4th Intern. Conf. on Algorithms for Computational Biology (AlCoB 2017)*, pages 85–99, Aveiro Portugal, 2017.

[37] A. L. Phinney, M. E. Calhoun, D. P. Wolfer, H. P. Lipp, H. Zheng, and M. Jucker. No hippocampal neuron or synaptic bouton loss in learning-impaired aged β-amyloid precursor protein-null mice. *Neuroscience*, 90(4):1207–1216, 1999.

[38] Mu-ming Poo, Michele Pignatelli, Tomás J. Ryan, Susumu Tonegawa, and et.al.'. What is memory? The present state of the engram. *BioMedCentral Biology*, pages 1–18, 2016. DOI 10.1186/s12915-016-0261-6.

[39] Dale Purves, George J. Augustine, David Fitzpatrick, and *et al. Neuroscience*. Sinauer Assoc., Sunderland, MA, 2008.

[40] Edmund T. Rolls and Alessandro Treves. *Neural Networks and Brain Function*. Oxford Univ. Press, Oxford, 1998.

[41] Azriel Rosenfeld and John L. Pfaltz. Sequential Operations in Digital Picture Processing. *J. of the ACM*, 13(4):471–494, Oct. 1966.

[42] Tomás J. Ryan, Dheeraj S. Roy, Michele Pignatelli, Autumn Arons, and Susumu Tonegawa. Engram cells retain memory under retrograde amnesia. *Science*, 348(6238):1007–1013, May 2015.

[43] Allessandro Sarti, Giovanna Citti, and Jean Petitot. Functional geometry of the horizontal connectivity in the primary visual cortex. *Journal of Physiology — Paris*, 103(1-2):37–45, Jan.-Mar. 2009.

[44] Zhao Sun, Hongzhi Wang, Haixun Wang, Bin Shao, and Jianzhong Li. Efficient Subgraph Matching on Billion Node Graphs. *Proceedings of the VLDB Endowment*, 5(9):788–799, 2012. Originally presented at VLDB Conf., Istanbul, Turkey, 2012.

[45] Natalie C. Tronson and Jane R. Taylor. Molecular mechanisms of memory reconsolidation. *Nature Review Neuroscience*, 8(4):262–275, 2007.

[46] Endel Tulving. Coding and representation: Searching for a home in the brain. In H.L. Roediger III, Y. Dudai, and S.M. Fitzpatrick, editors, *Science of Memory: Concepts*, pages 65–68. 2007.

[47] Karstan Weis. The Nuclear Pore Complex: Oily Spaghetti or Gummy Bear? *Cell*, 130:405–407, August 10 2007.

[48] Fan Zhu and Yuanfang Guan. Predicting dynamic signaling network response under unseen perturbations. *Bioinformatics*, 30(19):2772–2778, 2014. doi:10.1093/bioinformatics/btu382.