Entangled Associative Structures and Context

Douglas L. Nelson and Cathy L. McEvoy

University of South Florida

Department of Psychology and School of Aging Studies

Tampa, FL 33620

dnelson@cas.usf.edu

Abstract

Scientists are engaged in mapping the universe, the land, the genome, and word knowledge. In this paper we describe results of a 30-year effort to map and understand how preexisting word knowledge affects memory for a recently experienced word. We show how some of these findings are inconsistent with widely held views in psychological science and support the incorporation of the quantum formalism in our attempts to understand how prior knowledge interacts with recent experience and context. We relate our work to the state context property (SCOP) quantum formalism.

Introduction

There exists a large body of evidence showing that seeing or hearing a word activates words related to it through prior language experience. For example, seeing PLANET activates the associates *earth*, *moon*, and so on, because *planet-earth, planet-moon, moon-space* and other associations have been acquired in the past. This activation aids comprehension, is implicit, and provides rapid, probably synchronous access to what is associated to a word. Understanding how such activation affects memory requires a map of links among known words, and nearly 30 years ago we set about this task by using free association to construct an associative map of word knowledge.

In free association, we present individual words to samples of 150 participants asked to produce the first associated word to come to mind. We started using this task to index the strengths of preexisting links between pairs of words, with strength computed by dividing the production frequency of a response word by the sample size (e.g., the probabilities that the word planet produces *earth* and *mars* are .61 and .10, respectively). We soon discovered that some words had relatively small sets of associates whereas others had much larger sets. Although no theory at the time predicted that a word's preexisting associative set size would affect memory, we set about exploring its effects (e.g., Nelson and McEvoy 1979; Nelson, Schreiber, and McEvoy 1992). By the late 1980s, we realized that link strength and set size were capturing important dimensions of associative structure but were ignoring links between the associates (e.g., *moon-space)* and between the associates and the initiating stimulus (*earth-planet*). We then set about collecting free association norms for each word's *associates*. Norms were collected through the 1990s until over 5,000 words were normed using 6,000 participants (Nelson, McEvoy, and Schreiber 2004).

We discovered that the associative structures of individual words differ in both size and connectivity. Each word can be represented in an *nxn* matrix that describes the number and strength of three types of links: target-to-associate, associate-to-associate, and associate-to-target (see Nelson et al. 2004, for a database of 4,000 examples). Figure 1 shows the word PLANET in a network format to illustrate the three types of links.

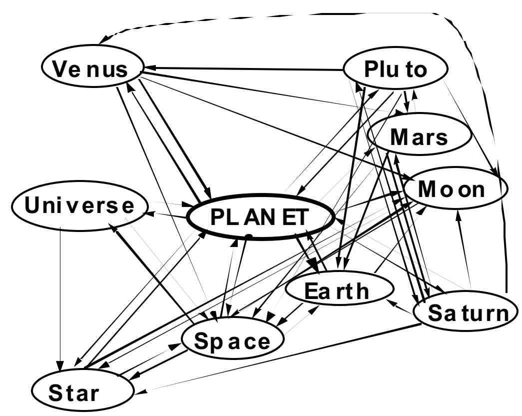

Figure 1. Planet's associative structure

Planet produces a relatively small set of 9 target-to-associate links (e.g., *earth*, *moon*) having many associate-to-associate links (e.g., *mars-to-earth*) and many associate-to-target links (e.g., *earth*-to-*planet*). Other words have relatively small sets and few associate-to-associate and associate-to-target links. All combinations of set size and connectivity are represented in the database. Because of the size of the database and because these indices are not highly correlated (Nelson and Zhang 2000), we were able to select words that systematically varied on these links in order to investigate how they affect recall and recognition.

Initially we were interested in determining whether the nature of a word's preexisting associative structure had any effects at all on memory, and later in the research effort we began to question prevailing notions about the nature of activation and fixed representations. Before turning to this work, three points about the nature of associative structure indexed via free association are relevant. The first is that free association provides a reliable index of response probability. Different samples of participants produce similar probabilities. The second point is that the collection of the norms has always been as context free as possible. Our interest has been on indexing what words people generally think of when there is little or nothing in the immediate context to bias responding. The third point is that free association provides a probabilistic index of link strength. The procedure allows us to expect that an association is likely to be present at the measured probability level for a new sample of participants but the state of that association for any individual participant in that sample cannot be determined. We model free association probability on the assumption that it represents mean relative associative strength from the normed word to the response word (Nelson, McEvoy, and Dennis 2000).

The Learning Task

We have investigated the effects of preexisting structure using a variety of different methods, but here we focus on the extralist cuing task. In this task, participants process a list of 24 words shown on the computer for 3 seconds each. Typically, the study instructions ask them to read each word aloud when shown and to remember as many as possible. They are told how they will be tested only after the last study word has been presented. The test instructions indicate that a new set of words, the test cues, will be shown and that each cue word is related to one of the target words just studied. These cues were not present during study (hence, the name extralist cuing). As each cue is shown, participants attempt to recall its associatively related word from the study list, and this test is self-paced.

This simple task allows for many variations in the learning and testing conditions and in the associative characteristics of the studied words and their test cues. For example, the studied words can be systematically selected from the norms based on their individual associative structures. With other variables controlled, half of the studied words might have many and half might have few associate-to-associate links. Similarly, half could have many and half could have few associate-to-target links, or the split could reflect differences in network size. The effects of factorial combinations of these variables can be investigated as well. The characteristics of the test cues can be held constant, or they can vary in strength in one or more of four different ways. Figure 2 shows Planet as a studied target, with Universe as the tes t cue.

Preexisting cue-to-target strength as indexed in free association is .18, and preexisting target-to-cue strength is .02. These two links directly connect the cue and target together. The stronger such links are the higher the probability of correct recall (e.g., Nelson, el al. 1998; Nelson and Goodmon 2003). Recall also varies with indirect links (Nelson, Bennett, and Leibert 1997). Recall is higher when mediated links (Universe-to-*space*-to-Planet) and shared associate links are present (both Universe and Planet produce *star* as an associate). Finally, Figure 2 also shows two associates, one linked to the cue (eternity) and one linked to the target (mars). Such associates do not link the cue and target together and they increase in number with a word's set size.

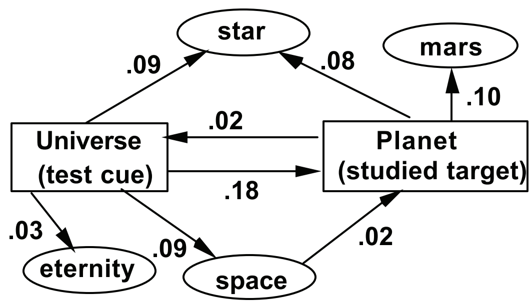

Figure 2. Links that join the test cue and target and competing associates that do not. Adapted from Nelson and McEvoy (2005).

An extralist cuing experiment involves manipulating or controlling the nine variables mentioned here in addition to other word characteristics such as target frequency. The results of hundreds of comparisons have shown that each of these variables affects recall under a variety of learning and testing conditions (for reviews see Nelson et al. 1992, 1998; Nels on and McEvoy 2005; Nelson and Zhang 2000). In terms of our initial purpose, the findings indicate that memory for a recently studied target word is strongly influenced by its preexisting associative structure. First, recall probability *increases* as the number of associate-to-associate and associate-to-target links increase (e.g., Nelson et al., 1992, 1993). Similar effects occur in recognition (Nelson, Zhang, and McKinney 2001). Words having more densely connected associates are more accurately recognized. Second, recall probability *decreases* as set size increases because associates like *mars* fail to link the cue and target. They compete with the target for recall. The goal is to recall the studied word linked to a particular cue, and recall decreases as the number of target competitors increases (Nelson et al. 1992).

Hence, target recall benefits from the presence of links between its associates and it declines when its associates are not connected to the test cue. Both the positive effects of the target's associative connectivity and negative effects of the target's set size occur despite the fact that attention is never drawn to the associates at any time. Furthermore, the effects of target connectivity and set size are not a fluke related to confounded word attributes, nor are they found only with particular types of participants or conditions (Gee, Nelson, and Krawczyk 1999; Nelson et al. 1993; Nelson and Goodmon 2002; Nelson et al. 1992, 2003). Both effects are evident regardless of target frequency, concreteness, and number of target meanings. The effects are found for young and old participants, under very fast and very slow presentation rates, as well as under incidental and intentional learning and testing conditions. Associative connectivity and set size have robust effects on probability of recall in the extralist cuing task. In trying to understand why associative structure has such robust effects on memory we learned that standard psychological explanations failed, and that the quantum formalism offers a promising alternative (e.g., Bruza and Cole 2005; Gabora and Aerts 2002).

Spooky Activation At A Distance

Nelson, McEvoy, and Pointer (2003) evaluated two explanations for why words having more associate-to-associate links are more likely to be recalled. Figure 3 shows a hypothetical target with two pre-existing target-to-associate links (indicated by arrows). There is also an associate-to-associate link between *Associates 1* and *2*, and an associate-to-target link from *Associate 2* to the Target. The values on the links indicate relative strengths measured independently via free association.

Two explanations for why associate-to-associate links benefit recall are shown below in the figure. The Spreading Activation equation is based on the classic idea that activation spreads through a fixed associative network, weakening with conceptual distance (e.g., Collins and Loftus 1975). Multiplying link strengths captures this weakening effect. Activation ostensibly travels from the target to and among its associates and back to the target in a continuous chain. The target is strengthened by activation that returns to it from pre-existing connections in long-term memory involving two- and three-step loops. In Figure 3, both two- and three-step loops add to the level of target activation. More associate-to-associate links create more three-step loops and theoretically benefit recall and recognition by increasing target activation strength. Importantly, in this rule the effects of associate-to-associate links are *contingent* on the number and strength of associate-to-target links because they allow activation to return to the target. If associate-to-target links were absent for a given target, even the maximum possible number of associate-to-associate links would have no effect on recall because activation could not return to the target (i.e., the number of 3-step loops would equal zero).

In the Spooky Activation at a Distance equation, the target activates its representation and the associates that comprise its network in synchrony. This equation assumes that each link in the associative set contributes additively to net strength. The beneficial effects of associate-to-associate links are not contingent on associate-to-target links. Stronger target activation is predicted when there are many associate-to-associate links even when associate-to-target links are absent. Target activation strength is solely determined by the sum of the link strengths within the target's associative set, regardless of origin or direction. The activation-at-a-distance rule assumes that the target is, in quantum terms, entangled with its associates because of learning and practicing language in the world. Associative entanglement causes the studied target word to simultaneously activate its associative structure.

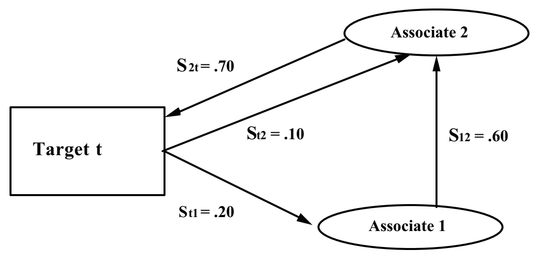

Spreading Activation Equation:

$$S(T) = \sum_{i=1}^{n} S_{ti}S_{it} + \sum_{i=1}^{n}\sum_{j=1}^{n} S_{ti}S_{ij}S_{jt} = (.10)(.70) + (.20)(.60)(.70) = .154$$

Activation at a Distance Equation:

$$S(T) = \sum_{i=1}^{n} S_{ti} + \sum_{i=1}^{n} S_{it} + \sum_{i=1}^{n}\sum_{j=1}^{n} S_{ij} = .20 + .10 + .70 + .60 = 1.60$$

Sti = target-to-associate i strength
Sit = associate i-to-target strength
Sij = associate i-to-associate j strength
n = number of associates
$i \neq j$

Figure 3. A hypothetical target with two associates and single associate-to-target and associate-to-associate links. From Nelson, McEvoy, and Pointer (2003).

Figure 4 shows the results of an extralist cuing experiment with cue-to-target strength set at an average of .17. This probability estimates the likelihood that a cue will produce its target in the absence of a study trial and serves as a lower boundary on expected recall. As can be seen, recall was well above this boundary. Recall is more likely when target words have more associate-to-associate links and more associate-to-target links. Most importantly, the effects of associate-to-associate links do not depend on the number of associate-to-target links. The interaction between the two variables is unreliable.

These results indicate that targets having more links *between* their associates and *from* their associates are more likely to be recalled. Memory for a word appearing in a new environmental context is probabilistically influenced by its past associative history involving related words. These results are also important because the effects of associate-to-associate links are not contingent on the presence of links from the associates to the target. This finding is inconsistent with spreading activation theory. Spreading activation that diminishes in strength with conceptual distance predicts that associate-to-associate links will have reduced effects compared to associate-to-target links because they are more distant. Furthermore, this theory predicts that associate-to-associate links will affect recall only when activation can return to the target through an associate-to-target link. In contrast, the findings of this experiment are inconsistent with both predictions indicating that this spreading activation theory is incorrect. Similar results are obtained in recognition (Nelson et al. 2001).

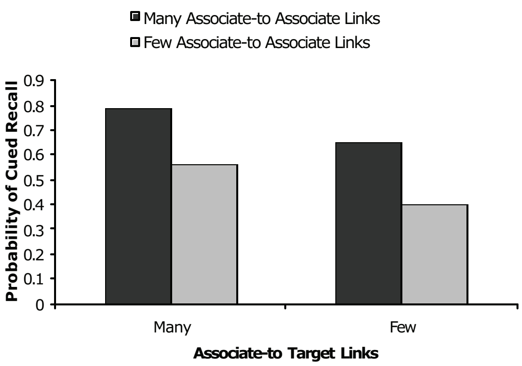

Figure 4. Probability of cued recall as a function of the numbers of associate-to-target and associate-to-associate links.

The additive effects of associate-to-associate and associate-to-target links on correct recall and recognition suggest that the target and its associates act as superposed entities under the present conditions. Because of their history of entanglement, the target and its associates act as correlated instead of separated entities (Aerts 1985a, 1985b; Bruza and Cole 2005; Gabora and Aerts 2002). This history of entanglement can be implemented by conceptualizing the mental lexicon as an *nxnxn* global semantic space (S). Words, their associates, and the associates of these associates are represented in S as *x*, *y*, and *z* vectors, respectively. We link this 3-coordinate conceptualization to the quantum formalism by treating S as a real-valued Hilbert space using free association values as the inner products. A word's local representation is described as a superposition of vectors that reflect its interactive history with other words appearing in the same language context. Figure 5 illustrates a portion of the interactive history of the word Planet that is captured by the activation-at-a-distance rule by adding the connection strengths defined in S. The local associative space of Planet equals the sum of its target-to-associate (.61 + .10 +…+ .02), associate-to-associate (.02 + .14 +…+ .09), and associate-to-target strengths (.16 + .02 +…+ .18). The latter two link types arise from the *associates* of the target's associates (coordinate *z*). In the activation-at-a-distance rule, the associative meaning of Planet is determined by both its associates and the associates of these associates falling within the target's set. In S, the *moon-space* association (.02) contributes as much to Planet's activation strength as the *planet-moon* association (.02).

	P	E	M	ST	SP	MO	PL	U
P		.61	.10	.08	.04	.02	.02	.02
E	.16				.02	.14		
M		.03			.06		.05	
ST	.02				.02	.12		
SP	.04	.01		.09		.03		.06
MO				.12	.02			
PL	.29	.01	.11		.03	.01		
U	.18			.09	.09			

Figure 5. The interactive history of Planet (P) with seven of its 9 associates, Earth (E), Mars (M), Star (ST), Space (SP), Moon (MO), Pluto (PL), and Universe (U). Numerical values reflect free association probabilities, so reading left to right along the rows Planet produces earth, moon, and so forth, with probabilities of .61, .10…; Earth produces moon, star, space, and so on, with probabilities of 0.0,0.0, .02…

In the quantum formalism, a quantum state encodes the potentialities of its measurable properties. By analogy, each link in Planet's local associative space represents a potentiality, and the sum of these links represents a superposition of potentialities. A superposition state is relevant whenever Planet, or any word, is experienced in isolation or in the absence of words that bias their meaning. A word's activation state, its accessible associative meaning, is context driven and when environmental context approaches "zero," meaning is uncertain and context is supplied internally by activating a set of potential meanings. When the context fails to specify a word's meaning, memory activates its associative history and this history provides an internally generated context for the word, as in the case of the extralist cuing experiments just described. The associative structure of *planet* changes from an inactive to an active state of potentialities, and this change is caused by experiencing it in an ambiguous context in which all of its associative meanings are potentially relevant. Context, whether absent or present, plays a critical role in determining what a word activates in memory.

The Role of Context

In addition to internally generated context, work in cognitive science indicates that more general context cues associated with the environment of a word experience also play an important role in determining the likelihood of recall. In the extralist cuing task, target words are studied in a general environment of physical, temporal, social, and emotional cues. Words are presented one at a time on a screen in a particular room at a particular time, and in the presence of an experimenter. The influence that general context has on recall becomes evident when study and testing contexts are different (Nelson and Goodmon 2002), when testing is delayed by problem solving or list learning tasks (Nelson et al. 1993), and when testing is incidental or implicit (Nelson and Goodmon 2002). For each of these manipulations we know that general context is important because variables that block or disrupt its retrieval reduce the influence that associative structure has on recall compared to when testing is immediate and disruptions are minimized. We model the context by associative structure interaction by cuing target recall with both general context cues and the extralist cue (Nelson et al. in press). Theoretically, test delays and other disruptions ostensibly reduce the effectiveness or the use of general context as a cue during retrieval while leaving the activation state of the target's associative structure intact. This solution allows the model to predict recall effectively on both immediate and disrupted tests. However, the activation state of the target may in fact change (e.g., decay), but as of the present time, we have no way of determining whether it has. Although the findings are clear and replicable, the potential mutability of the target's associative structure, the difficulties of measuring general context information, and the nature of the context-structure interaction remain uncertain under the conditions of these experiments. Given our state of ignorance about these uncertainties, the generalized quantum formalism may prove helpful in reconceptualizing the findings associated with manipulations of both general as well as specific context cues (Gabora 2007, this issue).

In contrast to the uncertainties associated with general context cues, specific contexts clearly change the target's activation state (e.g., Nelson et al. 1993). The most dramatic effects of context occur when the target is studied in the presence of a related word that biases its meaning. For example, if instead of studying Planet in isolation, the pair Universe-Planet is studied, the effects of target competitors and associate-to-associate links are reduced, dramatically. We compared extralist cued recall where the context cue is absent during study to intralist cued recall where the cue is present during study. Targets or Context-Target pairs were studied for 3 seconds, followed by an immediate cued recall test administered in the usual manner. As shown in Figure 6 and as would be expected, target set size effects produced by target competitors are apparent in extralist cuing because targets are studied in isolation and meaningful context is absent during study. When target meaning is uncertain, the lexical system fills it in by activating the target's entangled associative history. In contrast, when meaningful context is present during study in intralist cuing, recall is higher and the effects of target competitors are eliminated.

Figure 6. Target set size effects as a function of cuing task (adapted from Nelson, McEvoy, et al. 1993).

Similarly, as shown in Figure 7, the effects of associate-to-associate connectivity between the target's associates are apparent when the semantic context is absent during study, but not when it is present.

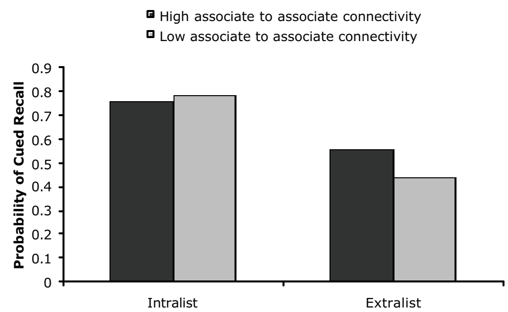

Figure 7. Effects of associate-to-associate connectivity as a function of cuing task (adapted from Nelson, McEvoy, et al. 1993).

In intralist cuing, effects of the target's associative structure are no longer apparent when it is studied in the presence of an associatively related word. Given the presence of a semantically biasing context, other studies have shown that the elimination of target competitor effects persists even when the context cue present during study is not used to prompt recall during testing (Nelson et al. 1992). Target competitor effects are essentially eliminated when word pairs are studied, regardless of whether recall is prompted by the meaningfully related context word the target was studied with, or by switching to extralist cues (e.g., study Universe Planet, cue recall with Earth or Janet). The recall effects of Planet's associative structure are present when it is studied in isolation, but not when it is studied in the presence of an associatively related context. Both the presence and absence of the effects of associative structure are apparent regardless of how recall is cued during testing. Such findings indicate that the activation of the target's associative structure and its elimination in the presence of a meaningfully related context represent encoding effects, not retrieval effects. The activation of the target's associative history represents a search for meaning that occurs whenever the context is uninformative or ambiguous.

Interestingly, the timing of the context word and target is critical in determining the influence of the target's associative structure. The elimination of target competitor effects occurs only when the word pairs appear simultaneously (Nelson, Gee, and Schreiber 1992). Regardless of whether the delay is varied by presenting word pairs alone or in the context of a sentence, if the context word appears slightly before or after the target, the effects of a target's associative structure are readily apparent though somewhat smaller compared to the context absent condition. Both the presence of context and its timing relative to the target influence the magnitude of the effects of associative structure.

We once assumed that such findings could be explained by adopting the post hoc assumption that the target's associates are activated and then inhibited when the target is studied is the presence of a meaningfully related context (Kintsch 1988; Nelson et al. 1992). However, direct evidence for such inhibition in the memory literature is scant and not compelling. We never published much of our own work on this issue because of difficulties associated with cuing the recall and recognition of ostensibly inhibited associates. Such associates generally failed to be less accessible as we had expected they would be if they had been inhibited. The experimental rational required using retrieval cues that were related to the ostensibly inhibited associates but not to the target, and we eventually abandoned this line of work because such cues proved difficult to find.

The generalized quantum formalism offers an alternative explanation for the context-generated elimination of associative structure effects that avoids the notion of inhibition altogether (Aerts and Gabora 2005a, 2005b; Gabora and Aerts 2002). In the SCOP model, sets of contexts and the associates of a studied word are embedded in complex Hilbert space to model how context affects the activation state of its associates. Activation states are unit vectors or density operators, with contexts and a word's associates treated as orthogonal projections. The tensor product provides a means for determining how context affects the activation state of a word's associates. SCOP provides an innovative solution for what is known as the *pet fish* problem. *Guppy* is rated as an atypical exemplar of both *pet* and *fish*, but as a typical exemplar of the conjunction *pet fish*. Existing theories of concepts cannot explain this finding. Similar to the SCOP solution for the *pet fish* problem, we assume that the associative structures for the activation states of Planet and Universe-Planet differ. When Planet (or Universe) is studied in relative isolation, its normatively defined associative structure is activated in a search after context-generated meaning. In quantum terms, Planet is in its superposition state in which its past interactions with other words are probable to varying degrees. In contrast, when Planet is studied in conjunction with the context provided by the word Universe, the superposition can be described as collapsing onto a specific meaning, the Universe-Planet meaning of Planet. This meaning represents one of Planet's potentialities. Other normally stronger potentialities such as Planet's *earth, mars and star* meanings become irrelevant because they are not activated. This interpretation assumes that the lexical system can be represented as a quantum system and that processing a target word in the context of another word is tantamount to measuring an observable in physics. Experiencing a word in context is presumptively equivalent to measuring the momentum or location of a particle. Such experience changes its activation state in memory, collapsing its potentialities to a more precise meaning. When the context that caused the collapse is re-presented as the retrieval cue, our findings indicate the semantic collapse of a studied target to a more precise meaning increases the probability that it will be recalled. Putatively stronger associates that compete with the target and interfere with its recall when the context is missing at study no longer interfere because their potentialities have been eliminated by the associative collapse during encoding.

This description of the interaction between context and collapse differs somewhat from that of Aerts and Gabora (2005). In their view, a state of a concept that is *not* an eigenstate of a context represents a potentiality state with respect to that context (e.g., *guppy* is not an eigenstate of *pet* but a potentiality state). An eigenstate of a concept can only be an eigenstate with respect to a particular context and, similarly, a potentiality state is always with respect to a particular context. Additional context changes the potentiality state of a context, and this change is described as collapse (e.g., the simultaneous presentation of Pet Fish alters the meaning of Pet so that *guppy* is now rated as a typical exemplar). We agree that potentiality states are always defined with respect to context, but we doubt that any operation or measurement can differentiate a concept into discrete states. The free association operation, "produce the first associated word to come to mind to the word Pet (or Planet)" produces a continuous, not a discrete, spectrum. This operator produces a continuum of possible values or probabilities that vary from 0.0-1.0 and does not admit eigenvectors (Hughes, 1989). We have generalized the meaning of collapse in assuming that isolated words such as *pet* or *planet* activate their potential meanings, and that context produces an associative equivalent of collapse by changing the state of these potentialities. In our view, attempts to determine that a concept is not an eigenstate of some context are likely to fail. For example, free association measures the very strongest associates of a word. If a word is not produced reliably in a sample of 150 participants, we might be tempted to conclude that it is not an eigenstate of Planet but is a potentiality (Aerts and Gabora 2005). However, we cannot say that this word is *not* an eigenstate because it is not a measured exemplar in the set. For example, Jupiter and Uranus are not produced as associates of Planet, but the conclusion that these words are not associates of Planet are unwarranted. We can only say that these items were not strong enough to overcome the interference produced in the free association task by Planet's stronger associates. A sample of 500 participants may show that they belong in the set with the implication that these words too can be eigenstates of Planet (by our reasoning).

Similarly, given a large enough sample, *guppy* is likely to be a member of the *Pet* category whether that determination is made via free association or typicality ratings. We cannot say with assurance when something is *not* an eigenstate of some context. In our view, words are intensely entangled as a result of processing in thousands if not millions of contexts. Potential set size is enormous when a word appears isolated from context, and even the weakest context appears to place strong constraints on these potentialities. So, it seems to us that it is reasonable to assume that context causes "collapse" to a more specific meaning. Compared to when experienced in isolation, a word's potential meaning set is likely to be greatly reduced by even the weakest of contexts. For example, Fence and Planet are not directly associated, but they are indirectly associated through other words. Among other associates, Fence produces *barrier, guard, keep out*, and Planet produces *earth*, which produces *ground*. Seen together, our understanding of the utterance, "The US is fencing the planet to define its financial self-interest" is entirely fluent. As with *guppy*, Fence approaches the Aerts and Gabora criterion as a potentiality state for Planet because it is not an associate of Planet. However, Fence is still associated with Planet albeit indirectly. The difference between the pair Universe-Planet and Fence-Planet is that the potentiality state is direct for the former pair and indirect for the latter. As an example of just how extensively entangled words are, Steyvers and Tannenbaum (2005) have shown that it takes only an average of three associative steps to get from any one word in our 5,000 word associative database to any other.

Associative entanglement may mean that collapse occurring because of context may never produce a "pure" state. Bruza and Cole (2005) show that the meaning of "Reagan" in the context of "Iran" appears to reflect two senses of "Reagan," one related to the Iran-contra scandal and one relating to his interactions with Iraq during the Iran-Iraq war. The context "Iran" activates a mixed state representing a mixture of meanings. This mixture, however, still represents a substantially reduced set of meanings compared to when "Reagan" is experienced as an isolated word. The collapse is partial instead of pure because context reduces the potentiality set, but not necessarily to a single meaning. Bruza and Cole propose an alternative interpretation of collapse analogous to a quantum system having many particles. With each particle corresponding to a word, the lexical system is represented by a global density state that incorporates an entire vocabulary. When a word is experienced in isolation, it provides its own context by changing the global density state to a density state that is unique to this word. The new density state is a subset of the global state. In generalizing this situation, they suggest that context can be represented as a projection of a density matrix onto a subspace represented by another density matrix.

This interpretation fits well with our findings and addresses the issue of associative entanglement because it allows for the complexity of associative meaning. The contiguous conjunction of the two words changes the meaning of each. The associative meaning of Planet is altered by the word Universe, and the associative meaning of Universe is altered by Planet. New concepts associated with exploring planets and extraterrestrial life can emerge. The density states of Universe and of Planet in isolation change when they are simultaneously experienced together and this interaction in turn can activate new potentialities. The concept "explore" is embedded in Explore-Universe-Planet, which in turn is embedded in Universe-Planet, which in turn is embedded in Planet. Increasingly more precise meaning is in a sense "curled up" in less precise meaning. Given entanglement, mixed density states are likely to be more common than rare when a single word or phrase is providing the only contextual constraints. In contrast, when reading extended discourse such as a story about discovering new planets, the context presumably can be mutually reinforcing so the meaning of any single word has the potential to approach purity.

The usefulness of the quantum formalism for memory research remains to be determined. Nevertheless, it provides a new way to think about how a word sometimes acts as if it were a collection of associated words that is instantly available and other times acts more or less as a precise entity. Given the nondeterministic nature of associative structures, context, and their interaction, the quantum formalism seems to provide an ideal means for conceptualizing such entities. Understanding the influence of associative structures on memory for recent experience is likely to require a deeper appreciation of uncertainty, superposition, and collapse.

Acknowledgments

This research was supported by grant MH16360 from the National Institute of Mental Health to D. L. Nelson and by grant 13973 from the National Institute on Aging to C. L. McEvoy.

References

Aerts, D. 1985a. The Physical Origin of the Einstein Podolsky Rosen Paradox. In G. Tarozzi and A van der Merwe (Eds.), *Open Questions in Quantum Physics: Invited Papers on the Foundations of Microphysics,* 33-50. Dordrecht: Kluwer Academic.

Aerts, D. 1985b. The Physical Origin of the ERP Paradox and How to Violate Bell Inequalities by Macroscopic Systems. In P. Lathi and P. Mittelstaedt (Eds.), *Symposium on the Foundations of Modern Physics: 50 years of the Einstein-Podolsky-Rosen Gedankenexperiment,* 305-320. Singapore: World Scientific.

Aerts, D. and Gabora, L. 2005a. A State-Context-Property Model of Concepts and their Combinations I: The Structure of the Sets of Contexts and Properties. *Kybernetes* 34(1&2): 167-191.

Aerts, D. and Gabora, L. 2005b. A State-Context-Property Model of Concepts and their Combinations II: A Hilbert Space Representation. *Kybernetes* 34(1&2): 192-221.

Bruza, P. D., and Cole, R. J. 2005. Quantum Logic of Semantic Space: An Exploratory Investigation of Context Effects in Practical Reasoning. In S. Artemov, H. Barringer, A. S. d'Avila Garcez, L. C. Lamb and J. Woods (Eds.) *We Will Show Them: Essays in Honor of Dov Gabbay*, 339-361. College Publications.

Collins, A. M., and Loftus, E. F. 1975. A Spreading-Activation Theory of Semantic Processing. *Psychological Review* 82(6): 407-428.

Gabora, L. and Aerts, D. 2002. Contextualizing Concepts Using a Mathematical Generalization of the Quantum Formalism. *Journal of Experimental and Theoretical Artificial Intelligence* 14(4): 327-358.

Gee, N. R., Nelson, D. L., and Krawczyk, D. 1999. Is the Concreteness Effect a Result of Underlying Network Interconnectivity? *Journal of Memory and Language* 40(4): 479-497.

Hughes, R. I. G. 1989. *The Structure and Interpretation of Quantum Mechanics*. Cambridge, MA: Harvard University Press.

Kintsch, W. 1988. The Role of Knowledge in Discourse Comprehension: A Construction-Integration Model. *Psychological Review* 95(2): 163-182.

Nelson, D.L., Bennett, D.J., Gee, N.R., Schreiber, T.A., and McKinney, V.M. 1993. Implicit Memory: Effects of Network Size and Interconnectivity on Cued Recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19(4): 747-764.

Nelson, D.L., Bennett, D.J., and Leibert, T. 1997. One Step is Not Enough: Making Better Use of Association Norms to Predict Cued Recall. *Memory & Cognition* 25(6): 785-796.

Nelson, D. L., Gee, N. R., and Schreiber, T. A. 1992. Sentence Encoding and Implicitly Activated Memories. *Memory & Cognition* 20(6): 643-654.

Nelson, D. L., and Goodmon, L.B. 2002. Experiencing a Word can Prime its Accessibility and its Associative Connections to Related Words. *Memory & Cognition* 30(3): 380-398.

Nelson, D. L., and Goodmon, L.B. 2003. Disrupting Attention: The Need for Retrieval Cues in Working Memory Theories. *Memory & Cognition* 31(1): 65-76.

Nelson, D. L., Goodmon. L. B., and Ceo, D. in press. How Does Delayed Testing Reduce Effects of Implicit Memory: Context Infusion or Cuing with Context? *Memory & Cognition*.

Nelson, D.L., and McEvoy, C.L. l979. Encoding Context and Set Size. *Journal of Experimental Psychology: Human Learning and Memory* 5(3): 292-314.

Nelson, D. L., and McEvoy, C. L. 2005. Implicitly Activated Memories: The Missing Links of Remembering. In C. Izawa and N. Ohta (Eds.) *Learning and Memory: Advances in Theory and Applications,* 177-198. New Jersey: Erlbaum.

Nelson, D. L., McEvoy, C. L., and Dennis, S. 2000. What is Free Association and What Does it Measure? *Memory & Cognition* 28(6): 887-899.

Nelson, D. L., McEvoy, C. L., and Pointer, L. 2003. Spreading Activation or Spooky Action at a Distance? *Journal of Experimental Psychology: Learning, Memory and Cognition* 29(1): 42-52.

Nelson, D.L., McEvoy, C.L., Janczura, G.A., and Xu, J. 1993. Implicit Memory and Inhibition. *Journal of Memory and Language* 32(5): 667-691.

Nelson, D.L., McEvoy, C.L., and Schreiber, T.A. 2004. The University of South Florida Word Association, Rhyme and Word Fragment Norms. *Behavior Research Methods, Instruments, & Computers* 36: 408-420.

Nelson, D.L., McKinney, V.M., Gee, N.R., and Janczura, G.A. 1998. Interpreting the Influence of Implicitly Activated Memories on Recall and Recognition. *Psychological Review* 105(2): 299-324.

Nelson, D.L, Schreiber, T.A., and McEvoy, C.L. 1992. Processing Implicit and Explicit Representations. *Psychological Review* 99(2): 322-348.

Nelson, D. L., and Zhang, N. 2000. The Ties that Bind What is Known to the Recall of What is New. *Psychonomic Bulletin & Review* 7(4): 604-617.

Nelson, D. L., Zhang, N., and McKinney, V. M. 2001. The Ties that Bind What is Known to the Recognition of What is New. *Journal of Experimental Psychology: Learning, Memory and Cognition* 27(5): 1147-1159.

Steyvers, M., and Tannenbaum, J. B. 2005. Graph Theoretic Analyses of Semantic Networks: Small Worlds in Semantic Networks. *Cognitive Science* 29(1):41-78.