
Quantum Logic of Semantic Space: an exploratory investigation of context effects in practical reasoning

P. D. BRUZA AND R. J. COLE

1 Introduction

The field of non-monotonic reasoning (NMR) has successfully provided an impressive symbolic account of human practical reasoning over the last two and half decades. There remains, however, a disappointment - the dearth of large-scale operational NMR systems on the ground. During Lora Morgenstern's keynote address at the International Joint Conference on Artificial Intelligence (IJCAI-97) with the title "Inheritance Comes of Age: Applying Non-monotonic Techniques to Problems in Industry" she warned researchers that NMR needs to go beyond the examination of toy examples and to tackle serious, large scale problems, or run the risk of NMR becoming a backwater at artificial intelligence conferences. That is getting on ten years ago. Since then NMR has largely crystallized and is well understood from a stratum of theoretical perspectives. Morgenstern's warning still lingers, in our opinion. Theoretical insight without corresponding reasoning systems on the ground belies NMR's promise of embodying human practical reasoning.

We feel that the symbolic characterization of practical reasoning is only part of the picture. Gärdenfors ([?], p127) argues that one must go under the symbolic level of cognition. In this vein, he states, "...information about an object may be of two kinds: *propositional* and *conceptual*. When the new information is propositional, one learns new *facts* about the object, for example, that x is a penguin. When the new information is conceptual, one *categorizes* the object in a new way, for example, x is *seen as* a penguin instead of as just a bird". Gärdenfors' mention of "conceptual" refers to the conceptual level of a three level model of cognition [?]. How information is represented varies greatly across the different levels. The sub-conceptual level is the lowest level within which information is carried by a connectionist representation. Within the uppermost level information is represented symbolically. It is the intermediate, *conceptual level*, or *conceptual space*, which is of particular relevance to this account. Here properties and concepts have

a geometric representation in a dimensional space. For example, the property of “redness” is represented as a convex region in a tri-dimensional space determined by the dimensions hue, chromaticity and brightness. The point left dangling for the moment is that representation at the conceptual level is rich in associations, both explicit and implicit. We speculate that the dynamics of associations are primordial stimuli for practical inferences drawn at the symbolic level of cognition. For example, it seems that associations and analogies generated within conceptual space play an important role in hypothesis generation. Gärdenfors ([?], p48) alludes to this point when he states, “most of scientific theorizing takes place within the conceptual level.” His conjecture is aligned with Gabbay and Woods’ insights regarding the cognitive economic basis of abduction [?]. Put crudely, it is cheaper to “guess” than to pursue a deductive agenda in relation to a problem at hand. Gabbay and Woods’ notion of cognitive economy rests on compensation strategies employed by a practical agent to alleviate the consequences of key cognitive resources such as *information*, *time*, and *computational capacity*. Practical reasoning is reasoning performed by practical agents, and is therefore subject to cognitive economy. In this connection, we put forward the following conjecture: It may well be that because such associations are formed below the symbolic level of cognition, significant cognitive economy results. This is not only interesting from a cognitive point of view, but also opens the door to providing a computationally tractable practical reasoning systems, for example, operational abduction to drive scientific discovery in biomedical literature [?, ?]

The appeal of Gärdenfors’ cognitive model is that it allows inference to be considered not only at the symbolic level, but also at the conceptual (geometric) level. Inference at the symbolic level is typically a linear, deductive process. Within a conceptual space, inference takes on a decidedly associational character because associations are often based on similarity (e.g., semantic or analogical similarity), and notions of similarity are naturally expressed within a dimensional space. For example, Gärdenfors’ states that a more natural interpretation of “defaults” is to view them as “relations between concepts”. This is a view which flows into the account which follows: the strength of associations between concepts change dynamically under the influence of context. This, in turn, influences the defaults harboured within the symbolic level of cognition.

It is important to note the paucity of representation at the symbolic level and reflect how symbolic reasoning systems are hamstrung as a result. In this connection, Gärdenfors ([?], p127) states, “..information about categorization can be quite naturally transfered to propositional information: categorizing x as an emu, for example, can be expressed by the proposition

“ x is an emu”. This transformation into the propositional form, however, tends to suppress the internal *structure* of concepts. Once one formalizes categorizations of objects by *predicates* in a first-order language, there is a strong tendency to view the predicates as primitive, atomic notions and to forget that there are rich relations among concepts that disappear when put into standard logical formalism.”

The above contrast between the conceptual and symbolic levels raises the question as to what are the implications for providing an account of practical reasoning. Gärdenfors states that concepts generate “expectations that result in different forms of *non-monotonic reasoning*”, which are summarized as follows:

Change from a general category to a subordinate

When shifting from a basic category, e.g., “bird” to a subordinate category, e.g., “penguin”, certain default associations are given up (e.g., “Tweety flies”), and new default properties may arise (e.g., “Tweety lives in Antarctica”).

Context effects

The context of a concept triggers different associations that “lead to non-monotonic inferences”. For example, *Reagan* has default associations “Reagan is a president”, “Reagan is a republican” etc., but *Reagan* seen in the context of *Iran* triggers associations of “Reagan” with “arms scandal”, etc.

The effect of contrast classes

Properties can be relative, for example, “a tall Chihuahua is not a tall dog” ([?], p119). In the first contrast class “tall” is applied to Chihuahuas and the second instance it is applied to dogs in general. Contrast classes generate conceptual subspaces, for example, skin colours form a subspace of the space generated by colours in general. Embedding into a subspace produces non-monotonic effects. For example, from the fact that x is a white wine and also an object, one cannot conclude that x is a white object (as it is yellow).

Concept combination

Combining concepts results in non-monotonic effects. For example, *metaphors* ([?], p130) Knowing that something is a lion usually leads to inferences of the form that it is alive, that it has fur, and so forth. In the combination, *stone lion*, however, the only aspect of the object that is lion-like is its shape. One cannot conclude that a stone lion has the other usual properties of a lion, and thus we see the non-monotonicity of the combined concept.

An example of the non-monotonic effects of concept combination not involving metaphor is the following: *A guppy is not a typical pet, nor is guppy is a typical fish, but a guppy is a typical pet fish.*

In short, concept combination leads to conceptual change. These correspond to revisions of the concept and parallel belief revisions modelled at the symbolic level, the latter having received thorough examination in the artificial intelligence literature.

The preceding brief characterization of the dynamics of concepts and associated non-monotonic effects is intended to leave the impression that a lot of what is happening in relation with practical reasoning is taking place within a conceptual (geometric) space. In addition, this impression may provide a foothold towards realizing genuine operational systems. This would require at least three issues to be addressed. The first is that a computational variant of the conceptual level of cognition is required. Secondly, the non-monotonic effects surrounding concepts would need to be formalized and implemented. Thirdly, the connection between these effects and NMR at the symbolic level needs to be specified. This account will attempt to address the first two of these questions. Computational approximations of conceptual space will be furnished by semantic space models which are emerging from the fields of cognition and computational linguistics. Semantic space models not only provide a cognitively motivated basis to underpin human practical reasoning, but from a mathematical perspective, they are real-valued Hilbert spaces. This introduces the tantalizing and highly speculative prospect of formalizing aspects of human practical reasoning via quantum mechanics. In this account will focus on a treatment of how to formalize context effects as well as keeping an eye on operational issues.

2 Semantic space: computational approximations of conceptual space

To illustrate how the gap between cognitive knowledge representation and actual computational representations, the Hyperspace Analogue to Language (HAL) model is employed [?, ?]. HAL produces representations of words in a high dimensional space that seem to correlate with the equivalent human representations. For example, “...simulations using HAL accounted for a variety of semantic and associative word priming effects that can be found in the literature...and shed light on the nature of the word relations found in human word-association norm data”[?]. Given an n -word vocabulary, a HAL space is an $n \times n$ matrix constructed by moving a window of length l over the corpus by one word increment ignoring punctuation, sentence and paragraph boundaries. All words within the window are considered as co-occurring with the last word in the window with a strength inversely proportional to the distance between the words. Each row i in the matrix represents accumulated weighted associations of word i with respect to other words which preceded i in a context window. Conversely, column

```

def calculate_hal(documents, n)
  HAL = 2DArray.new()
  for d in documents {
    for i in 1 .. d.len {
      for j in max(1,i-n) .. i-1 {
        HAL[d.word(i),d.word(j)] += n+1-(i-j)
      }
    }
  }
  return HAL
end

```

Figure 1.1. Algorithm to compute the HAL matrix for a collection of documents. It is assumed that the documents have been pruned of stop words and punctuation.

i represents accumulated weighted associations with words that appeared after i in a window. For example, consider the text “President Reagan ignorant of the arms scandal”, with $l = 5$, the resulting HAL matrix H would be:

	arms	ig	of	pres	reag	scand	the
arms	0	3	4	1	2	0	5
ig	0	0	0	4	5	0	0
of	0	5	0	3	4	0	0
pres	0	0	0	0	0	0	0
reag	0	0	0	5	0	0	0
scand	5	2	3	0	1	0	4
the	0	4	5	2	3	0	0

Table 1.1. A simple semantic space computed by HAL

If word precedence information is considered unimportant the matrix $S = H + H^T$ denotes a symmetric matrix in which $S[i, j]$ reflects the strength of association of word i seen in the context of word j , irrespective of whether word i appeared before or after word j in the context window. The column vector S_j represents the strengths of association between j and other words seen in the context of the sliding window: the higher the weight of a word, the more it has lexically co-occurred with j in the same context(s). For example, table 1.2 illustrates the vector representation for “Reagan” taken from a matrix S computed from a corpus of 21578 Reuters news feeds taken from the year 1988.

president (5259), administration (2859), trade (1451), house (1426), budget (1023), congress (991), bill (889), tax (795), veto (786), white (779), japan (767), senate (726), iran (687), billion (666), dlrs (615), japanese (597), officials (554), arms (547), tariffs (536) ...
--

Table 1.2. Example representation of the word “Reagan”

HAL is an exemplar of a growing ensemble of computational models emerging from cognitive science, which are generally referred to as *semantic spaces* [?, ?, ?, ?, ?, ?, ?, ?, ?]. Even though there is ongoing debate about specific details of the respective models, they all feature a remarkable level of compatibility with a variety of human information processing tasks such as word association. Semantic spaces provide a geometric, rather than propositional, representation of knowledge. They can be considered to be approximations of conceptual space proposed by Gärdenfors [?].

Within a conceptual space, knowledge has a dimensional structure. For example, the property colour can be represented in terms of three dimensions: hue, chromaticity, and brightness. Gärdenfors argues that a property is represented as a convex region in a geometric space. In terms of the example, the property “red” is a convex region within the tri-dimensional space made up of hue, chromaticity and brightness. The property “blue” would occupy a different region of this space. A domain is a set of integral dimensions in the sense that a value in one dimension(s) determines or affects the value in another dimension(s). For example, the three dimensions defining the colour space are integral since the brightness of a colour will affect both its saturation (chromaticity) and hue. Gärdenfors extends the notion of properties into concepts, which are based on domains. The concept “apple” may have domains taste, shape, colour, etc. Context is modelled 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 will be heavily weighted (i.e., it’s roundness). One of the goals of this article is to provide both a formal and operational account of this weighting function.

Observe the distinction between representations at the symbolic and conceptual levels. At the symbolic level “apple” can be represented as the atomic proposition *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” (symbolic level) 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”.

Barwise and Seligman [?] also propose a geometric foundation to their account of inferential information content via the use of real-valued state spaces. In a state space, the colour “red” would be represented as a point in a tri-dimensional real-valued space. For example, brightness can be modelled as a real-value between white (0) and black (1). Integral dimensions are modelled by so called observation functions defining how the value(s) in dimension(s) determine the value in another dimension. Observe that this is a similar proposal, albeit more primitive, to that of Gärdenfors as the representations correspond to points rather than regions in the space.

Semantic space models are an approximation of Barwise and Seligman state spaces whereby the dimensions of the space correspond to words. A word j is a point in the space. This point represents the “state” in the context of the associated text collection from which the semantic space was computed. If the collection changes, the state of the word may also change. Semantic space models, however, do not make provision for integral dimensions. An important intuition for the following is the state of a word in semantic space is tied very much with its “meaning”, and this meaning is context-sensitive. Further, context-sensitivity will be realized by state changes of a word.

In short, HAL, and more generally semantic spaces, are a promising, pragmatic means for knowledge representation based on text. They are computational approximations, albeit rather primitive, of Gärdenfors’ conceptual space. Moreover, due to their cognitive track record, semantic spaces would seem to be a fitting foundation for considering realizing computational variants of human reasoning. Finally, a semantic space is a real-valued Hilbert space which opens the door to connections with quantum mechanics.

3 Context effects in Semantic Space

Human beings are adept at producing context-sensitive inferences. Shifts in context effect the inferences made, even to a dramatic degree. The well known “Tweety” problem exemplifies this. A rough account of this example in terms of Gärdenfors’ model of cognition is as follows: When given “Tweety is a bird”, a prototypical concept of bird is activated within conceptual space and default inferences at the symbolic level such as “Tweety flies” arise as a strong association is primed between “Tweety” and “flies” at the conceptual level. The prototypical “Tweety” would be a point in the centre of a convex region in conceptual space representing birds (see [?],[?], p139). Learning “Tweety is a penguin”, shifts the representation of “Tweety” towards the edge of the region representing birds as penguins differ signif-

icantly to the prototypical bird. As a consequence the association with “flies” diminishes radically and new associations arise, e.g., with “Antarctica”. Even though Gärdenfors characterizes this type of NMR as being driven by a change from a general category to a subordinate, we feel that the associations can be more generally considered as a product of a shift of context — in this case the context is being *refined* from broader to the more specific. Initially the object “Tweety” is placed in the context of the concept “bird”. The context is then refined to “penguin” leading to a change in the associations being primed, and consequently a change in the inferences being drawn.

The “Reagan” example exhibits similar characteristics. The vector representation given in table 1.2 is almost the prototypical representation of “Reagan” in the context of the underlying corpus. This is because HAL accumulates the association weights as it goes along. In fact, the weights in table 1.2 need only be divided by the frequency of the term “Reagan” in the underlying corpus to produce the vector representing prototypical “Reagan”. Highly weighted associations in the representation have the character of being default like - “Reagan was a president”, “Reagan had an administration” etc. Such default associations reflect the run of the mill presidential Reagan dealing with trade, budgets, congress etc.

The above two examples exhibit very common, or “garden” variety of practical inference. In this section, we attempt to provide a formal account in terms of quantum mechanics (QM).

3.1 Bridging semantic space and QM

A semantic space is a vector space and these can be expressed in the notation of quantum mechanics (The following draws heavily from [?]).

A semantic space S is a $m \times n$ matrix where the columns $\{1, \dots, n\}$ correspond to a vocabulary V of n words. A typical method for deriving the vocabulary is to tokenize the associated corpus and remove non information bearing words such as “the”, “a”, etc. The letters u, v, w will be used to identify individual words.

The interpretation of the rows $\{1 \dots m\}$ depends of the type of semantic space in question. For example, table 2 illustrates that HAL produces a square matrix in which the rows are also interpreted as representations of terms. In contrast, a row in the semantic space models produced by Latent Semantic Analysis [?] corresponds to a text item, for example, a whole document, a paragraph, or even a fixed window of text, as above. The value $S[t, w] = x$ denotes the salience x of word w in text t . Information-theoretic approaches are sometimes use to compute salience. Alternatively, the frequency of word w in context t can be used.

For reasons of a more straightforward embedding of semantic space into QM, we will focus on square, symmetric semantic spaces ($m = n$). A word w is represented as a column vector in S :

$$(1.1) \quad |w\rangle = \begin{pmatrix} w_1 \\ \vdots \\ w_n \end{pmatrix}$$

The notation on the LHS is called a *ket*, and originates from quantum physicist Paul Dirac. Conversely, a row vector $v = (v_1, \dots, v_n)$ is denoted by the *bra* $\langle v|$.

Multiplying a ket by a scalar α is as would be expected:

$$(1.2) \quad \alpha|w\rangle = \begin{pmatrix} \alpha w_1 \\ \vdots \\ \alpha w_n \end{pmatrix}$$

Addition of vectors $|u\rangle + |v\rangle$ is also as one would expect. In Dirac notation, the scalar product of two n -dimensional real¹ valued vectors u and v produces a real number:

$$(1.3) \quad \langle u|v\rangle = \sum_{i=1}^n u_i v_i$$

The product $|u\rangle\langle u|$ produces a symmetric matrix. Vectors u and v are *orthogonal* iff $\langle u|v\rangle = 0$. Scalar product allows the length of a vector to be defined: $\|u\| = \sqrt{\langle u|u\rangle}$. A vector $|u\rangle$ can be normalized to unit length ($\|u\| = 1$) by dividing each of its components by the vector's length: $\frac{1}{\|u\|}|u\rangle$.

A Hilbert space is a complete² inner product space. In the formalization to be presented in ensuing sections, a semantic space S is an n -dimensional real-valued Hilbert space using Euclidean scalar product as the inner product.

A Hilbert spaces allows the state of a quantum system to be represented. It is important to note that a Hilbert space is an *abstract* state space meaning QM does not prescribe the the state space of specific systems such as electrons. This is the responsibility of a physical theory such as quantum electrodynamics. Accordingly, it is the responsibility of semantic space theory to offer the specifics: In a nutshell, a ket $|w\rangle$ describes the state of a

¹QM is founded on complex vector spaces. We restrict our attention to finite vector spaces of real numbers.

²The notion of a “complete” vector space should not be confused with “completeness” in logic. The definition of a completeness in a vector space is rather technical, the details of which are not relevant to this account.

word w . It is akin to a particle in QM. The state of a word changes due to context effects in a process somewhat akin to quantum collapse. This in turn bears on practical inferences drawn due to context effects of word seen together with other words as described above.

In QM, the state can represent a superposition of potentialities. By way of illustration consider the state σ of a quantum bit, or *qubit* as:

$$(1.4) \quad |\sigma\rangle = \alpha|0\rangle + \beta|1\rangle$$

where $\alpha^2 + \beta^2 = 1$. The vectors $|0\rangle$ and $|1\rangle$ represent the potentialities, or *eigenstates* of “off” and “on”. Eigenstates are sometimes referred to as *pure* states. They can be pictured as defining orthogonal axes in a 2-D plane:

$$(1.5) \quad \alpha|0\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

and

$$(1.6) \quad \alpha|1\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

The state σ is a linear combination of eigenstates. Hard though it is to conceptualize, the linear combination allows the state of the qubit to be a mixture of the potentialities of being “off” and “on” at the same time.

In summary, a quantum state encodes the probabilities of its measurable properties, or eigenstates. The probability of observing the qubit being off (i.e., $|0\rangle$ is α^2). Similarly, β^2 is the probability of observing it being “on”.

The above detour into QM raises questions in relation to semantic space. What does it mean that a word is a superposition - a “mixture of potentialities”? What are the eigenstates of a word?

3.2 Mixed and eigenstates of a word

Consider the following traces of text from the Reuters-21578 collection:

- *President Reagan was ignorant about much of the Iran arms scandal*
- *Reagan says U.S to offer missile treaty*
- *Reagan seeks more aid for Central America*
- *Kemp urges Reagan to oppose stock tax.*

Each of these is a window which HAL will process accumulating weighted word associations in relation to the word “Reagan”, say. In other words, included in the HAL vector for “Reagan” are associations dealing with the

Iran-contra scandal, missile treaty negotiations with the Soviets, stock tax etc. The point is the HAL vector for “Reagan” represents a mixture of potentialities.

Let us now generalize the situation somewhat. Consider once again the HAL matrix H computed from the text “President Reagan ignorant of the arms scandal”. As mentioned before, $S = H + H^T$ is a symmetric matrix. Technically, S is a *Hermitian linear operator*. Consider a set of text windows of length l which are centred around a word w . Associated with each such text window j , $1 \leq j \leq m$, is a semantic space S_j . It is assumed that the semantic space is n -dimensional, whereby the n dimensions correspond to a fixed vocabulary V as above. The semantic space around word w , denoted by S_w , can be calculated by the sum:

$$(1.7) \quad S_w = \sum_{j=1}^k S_j$$

In other words, the semantic space around the word “Reagan” is a summation of n -dimensional Hermitian linear operators computed from text windows centred around “Reagan”.

In turn, the semantic space of the associated corpus, termed the *global semantic space*, denoted \mathcal{S} can be considered as a mixture of the semantic space of the words in the associated vocabulary V :

$$(1.8) \quad \mathcal{S} = \sum_{w \in V} S_w$$

An important intuition drawn from QM is that a word meaning equates with a state. The state may be *mixed*, that is the state embodies different potentialities corresponding to different “senses” of the word Reagan. Here we use the word “sense” with some poetic licence, but we do so deliberately because the “Reagan” example is similar to the case of an ambiguous word. Consider the word “suit” in isolation. Is it an item of clothing or a legal procedure? We put forward the intuition that, both “Reagan” or “suit” are states involving mixtures of senses, which parallels the superposition of eigenstates in the qubit given above. More formally, let $|r\rangle$ be the vector representing the state of “Reagan” in a semantic space S , and $\{e_1, \dots, e_k\}$ represent the eigenstates of S , then

$$(1.9) \quad |r\rangle = \alpha_1|e_1\rangle + \dots + \alpha_n|e_k\rangle$$

where $\alpha_1^2 + \dots + \alpha_n^2 = 1$. The preceding intuition connecting word “meanings” in semantic space to QM seems to have independently arisen. (See [?, ?, ?]). The eigenstates define the different senses of the word in question. In QM terms, these correspond to the eigenstates of “Reagan”.

In QM, the interpretation of the eigenstates are clearly grounded, e.g., the “on”, “off” states of a qubit, or the momentum eigenstate of a particle. The eigenstates of a word are more subtle. This subtlety is not due to subjective interpretations of word meanings. By using a semantic space constructed from a corpus of text, the “meanings” ultimately are derived from this corpus. It could be argued that such meanings are inter-subjective due to the track record of semantic space model in replicating human word association norms. The subtlety derives more from the range of potential eigenstates. We shall see as we go along, however, the state of a word is nevertheless amenable to a formal treatment.

Computing eigenstates of a word by Singular Value Decomposition

Singular value decomposition, a theorem from linear algebra, allows a matrix to be decomposed. In the following, many of the technical details of SVD will be skipped, and only the essential elements will be presented. See [?] for a comprehensive account. As S_w is a symmetric matrix, SVD decomposes it as follows:

$$(1.10) \quad S_w = UDU^T$$

where U is a $n \times n$ unitary matrix, the columns of which are the orthonormal basis of S_w . This means, the columns of U are pairwise orthogonal. To remain consistent with our notation, the i -th column vector of U will be denoted by $|e_i\rangle$. Matrix D is a positive $n \times n$ diagonal matrix, the values of which are the eigen-values of S_w . The value $D[i, i]$ will be denoted d_i .

The spectral decomposition of SVD allows S_w to be reconstructed, where $k \leq n$:

$$S_w = \sum_{i=1}^k |e_i\rangle d_i \langle e_i| \quad (1.11)$$

$$= \sum_{i=1}^k d_i |e_i\rangle \langle e_i| \quad (1.12)$$

$$= d_1 |e_1\rangle \langle e_1| + \dots + d_k |e_k\rangle \langle e_k| \quad (1.13)$$

This shows once again how a word w is a mixture of eigenstates $|e_i\rangle$. The eigenvalues are related to the probabilities of the eigenstates occurring after a quantum measurement. In the semantic space interpretation, the eigenstates $|e_i\rangle$ of S_w correspond to the senses of word w .

The spectral decomposition of S_w parallels the decomposition of a density state ρ [?]:

$$(1.14) \quad \rho = a_1 P_1 + \dots + a_k P_k$$

where P_i is a projection operator and $a_1 + \dots + a_k = 1$. Projection operators, in real valued state spaces, are idempotent, symmetric matrices. (Note that $P_i = |e_i\rangle\langle e_i|$ is a projection operator)

A *density state*, or *density operator*, or *density matrix* expresses the distribution of quantum states in an ensemble of particles. The intuition is that we will run with is that words are particles, and that context acts like a “measurement”, e.g., on a particle. A density matrix is a Hermitian operator with all its eigenvalues between 0 and 1.

Aerts and Czachor [?] have shown how to render a semantic space computed by Latent Semantic Analysis into a density matrix. For the purposes of this article, however, equation 1.14 allows the notion of density matrix to be directly equated with semantic space when the weights in the space are normalized. This ensures the eigenvalues lie between 0 and 1.

3.3 An analysis of the eigenstates of “Reagan”

Table 1.3 contains the first four eigenstates, or eigenvectors, of S_{reagan} . The first eigenstate contains all positive values and can be seen as a kind of average of the space. The subsequent eigenstate has a single positive component and a collection of negative components. Eigenstates having both positive and negative components represent two contrasting aspects of the space. Individual word vectors may project onto either the positive or the negative portion of each eigenstate. If they were to project onto both then their projection into the subspace would be small, and since SVD maximises the average length of the projected vectors, the negative and positive parts of the eigenstates tend to be in opposition. The third eigenstate, for example, seems to indicate that reports about Reagan concerning exports, tariffs and Japan are in opposition to reports about the senate, vetos and the budget.

It is important to recognize that the eigenstates do not neatly partition the meanings of Reagan into distinct clusters but rather span a subspace describing the topics in which Reagan is involved. The space can be thought of as being lumpy but continuous rather than being due to a small number of discrete and largely disjoint topics.

In short, the eigenstates computed by SVD do not seem to correspond well with the intuitively expected eigenstates of the word “Reagan”.

3.4 Summary

As this section has proceeded a sort of duality has emerged. Initially, the state of a word w was presented as a ket $|w\rangle$ in a n -dimensional semantic (Hilbert) space S . The ket $|w\rangle$ may represent a mixture of the senses of

1: reagan (0.62), president (0.48), administration (0.22), house (0.17), trade (0.15), congress (0.11), budget (0.11), bill (0.10), veto (0.10), white (0.09), tax (0.09), japan (0.08), senate (0.08), billion (0.08), iran (0.07),
2: reagan (0.74), ... bill (−0.04), congress (−0.05), trade (−0.07), house (−0.08), administration (−0.23), president (−0.55)
3: japan (0.25), trade (0.25), japanese (0.24), tariffs (0.21), administration (0.13), united (0.11), sanctions (0.11), exports (0.11) ... tax (−0.11), senate (−0.13), veto (−0.14), budget (−0.19), white (−0.31), house (−0.38)
4: billion (0.44), dlrs (0.37), dlr (0.21), budget (0.18), veto (0.18), deficit (0.17), bill (0.14), highway (0.13), mln (0.10), ... conference (−0.07), house (−0.08), baker (−0.09), scandal (−0.12), white (−0.14), arms (−0.24), iran (−0.25)

Table 1.3. First four eigenstates of S_{reagan} . Components are listed in order. Only the largest components (by magnitude) are for each eigenstate are shown.

the word w . Later, a connection was made between the semantic space S_w constructed around a word w and a density matrix, a notion from QM. From a technical point of view, this is not a problem. A density matrix can represent both eigenstates and mixed states: If $|w\rangle$ represents an eigenstate, then $|w\rangle\langle w|$ represents the corresponding density matrix. As a consequence, the state of a word, whether pure or mixed, can be represented as density matrix.

However, this technical resolution, does not seem to fully resolve the perceived duality. For example Widdows [?] has proposed a quantum of word meanings drawn from semantic space. The meanings are represented as kets with no recourse to density matrices. Aerts and Gabora [?], on the other hand employ kets for the pure states of a concept, and a density matrix for a mixed state of a concept. It would seem that more research is needed to resolve this duality.

Quantum Collapse and Context Effects in Semantic Space

A quantum system is usually not in an eigenstate of whatever observable (e.g., momentum) is intended to be measured. However, if the observable is measured, the state of the system will immediately become an eigenstate of that observable. This process is known as *quantum collapse*.

A parallel can be drawn with respect to words in semantic space. When a word is seen in context, the superposition (mixed) state of the word collapses onto one of its senses. The senses of a word are the observables. For example, when “Reagan” is seen in the context of “Iran”, the mixture of potentialities of “Reagan” collapses onto the eigenstate representing the sense dealing with the Iran-Contra scandal. After collapse, weights of associations to words such as “Contra”, “illegal”, “arms”, “scandal”, “sale” will be high, whereas before collapse the weights of such associations may have been weak. The highly weighted associations may, for example, “bubble up” and give rise to defaults at the symbolic level of cognition. This intuition gives rise to the tantalizing possibility that context effects within the conceptual level of cognition may be formalized by quantum collapse. This change in weighting can be dramatic and thus produce non-monotonic effects in relation to the weights of associations. For the moment, the observables can be conceived of as the different senses of a word. Seeing a word in the context of other word(s) acts like a “measurement”. This measurement collapses the word meaning into one of its potential senses.

The description above of the interaction between context and collapse is essentially the same as that of Aerts and Gabora [?]. They state: “A state [of a concept] that is not an eigenstate of the context is called a potentiality state with respect to this context. The effect of a context is to change a potentiality state of this context, and this change will be referred to as collapse”.

The context effects that are being considered here are similar to Aerts and Gabora. For simplicity, the case of a word v seen in the context of word u will be considered, the prototype of the running example: “Reagan” in the context of “Iran”.

3.5 Formalizing context effects by quantum collapse

Aerts and Gabora [?] state a measurement in QM is described by a Hermitian operator M . For the context word u , there is an associated operator M_u . It is assumed that the state of word v is represented by the $|v\rangle$ drawn from some density matrix ρ . The parallel with QM is the following — a particle (word) v is drawn from a quantum system, the state of which is ρ . It is subjected to a “measurement”, which is a product of word v being seen in the context of word u . The state of v collapses as a result. This intuition is formalized as follows, where $|v_u\rangle$ denotes the state of word v after collapse:

$$(1.15) \quad |v_u\rangle = \frac{M_u|v\rangle}{\sqrt{\langle v|M_u|v\rangle}}$$

The value $\sqrt{\langle v|M_u|v\rangle}$ is a normalizing factor. One way of inspecting non-monotonic effects in relation to associations is simply to compare $|v_u\rangle$ with $|v\rangle$. Recall the ket representation of a word is a vector whose components correspond to words. The value x of the component i represents the strength of association of $|v_u\rangle$ with the i -th word of vocabulary V . Examples will follow shortly.

The above equation is a more liberal interpretation of that proposed by Aerts and Gabora’s equation 11 in [?]. Their equation requires $|v\rangle$ to be a pure state. Our more liberal proposal arises from the following intuition: Collapse due to context may not necessarily result in a pure state. For example, Reagan’s involvement with Iran included the U.S. embassy hostage crisis as well as the Iran-contra scandal. Intuitively this phenomena corresponds to a *partial* collapse of “Reagan”, whereby the resultant state is less mixed than originally. In other words, the context “Iran” has not fully led to a collapse of the “meaning” of “Reagan” onto an unambiguous sense. This phenomenon shows the embedding of semantic space into QM is not always straightforward.

3.6 Example: “Reagan” in the context of “Iran”

To illustrate the effect of equation 1.15, $|v\rangle$ is primed to be the state of the word “Reagan” extracted from the density matrix ρ_{Reagan} computed from the Reuters collection. This ket represents the prototypical presidential Reagan, and is illustrated in table 1.2. The measurement operator M_u is primed as the Hermitian operator S_w with w equal to the term “Iran”. One interpretation of the resulting quantum collapse is that it promotes words occurring in the vicinity of “Iran” based on how similar their meaning in the context of “Iran” is to the meaning of the prototypical Reagan.

iran (59), reagan (27), arms (21), iraq (12), gulf (12), scandal (10), war (7), oil (7), iranian (7), sales (7), house (6), president (5), attack (5), contra (5), united (5), states (4), white (4), missiles (4), profits (4), action (3), military (3), officials (3), senate (3), new (3), tehran (3), shipping (3), news (3), offensive (3), sale (3), rebels (2), speech (2), secret (2), warned (2), iraqi (2), policy (2), fighting (2), commission (2), response (2), hussein (2), diversion (2), major (2), official (2), tower (2), ship (2), denied (2), foreign (2), deal (2), affair (2), administration (2), saddam (2), ...

Table 1.4. “Reagan” in the context of “Iran”.

Compare the above weighted associations with those of table 1.2. Observe how the above no longer represent the prototypical “Reagan”, but

where associations relevant to the Iran-contra scandal are apparent, e.g., “scandal”, “arms”, “sales”, “contra”. Therefore, the Iran-contra sense of Reagan is coming through. Also there are prominent associations to “Iraq” and “oil”. These may be related to the sense of “Reagan” reflecting President Reagan’s dealings with Iraq during the Iran-Iraq war. Therefore, table 1.4 seems to reflect two senses of “Reagan”. In other words, the resultant state after collapse is mixed. The reason for this is that the context word “Iran” is also a mixture of senses.

Perhaps, for this reason, Aerts and Gabora [?] do not directly employ measurement operator M as a whole, but its spectral decomposition:

$$M = d_1|e_1\rangle\langle e_1| + \dots + d_k|e_k\rangle\langle e_k|$$

where the projector $P_j = |e_j\rangle\langle e_j|$. They refer to projector P_j as a “piece of context”. Take for example, $u = \{\text{Iran}\}$. This context is a mixture of senses involving oil trade, Iran-Iraq war, the US embassy siege etc. The intuition of the projector P_j is that it represents one of these senses, and this in turn is a “piece of context” which can be substituted in equation 1.15 instead of M_u .

Recourse to “pieces of context” does not satisfactorily remove an incongruence. Why is it that when presented with “Reagan” in the context of “Iran”, most will readily assume the Iran-contra sense, which we argued earlier, is an eigenstate of “Reagan”. This stands in contrast to the above mixed state of “Reagan” after collapse illustrated in table 1.4. The progression is as follows. Initially the state of “Reagan” is mixed as reflected by the following ket $|r\rangle$ drawn from the Reagan density matrix ρ_r . Assume that the eigenstate of e_i corresponds to the Iran-contra sense of “Reagan”:

$$|r\rangle = \alpha_1|e_1\rangle + \dots + \alpha_i|e_i\rangle + \dots + \alpha_k|e_k\rangle$$

After collapse due to context “Iran”, the state of “Reagan” is still mixed but less mixed than before. The result computed above suggests two senses, denoted e_i and e_j :

$$|r\rangle = \beta_i|e_i\rangle + \beta_j|e_j\rangle$$

The eigenvalues β_i are related to probabilities. For the sake of argument, let us assume that $\beta_i > \beta_j$. We speculate that the reason that the eigenstate e_i is assumed by most, is because it is the more probable sense left after collapse. Bear in mind, these probabilities are furnished by the geometry of the space and not by a frequentist approach which dominates statistical language processing.

3.7 Another way to view collapse of word meanings

Let us assume that before any words are seen or uttered there is a global density state represented by ρ_S , where S signifies the global semantic space. This is akin to a quantum system with many particles, each particle corresponding to a word in the vocabulary V , which may number in the hundreds of thousands. Consider what happens when a word v is expressed in isolation of other words. This is transforming a situation without context into one where the context is simply given by the word v . We contend that this changes the density state from ρ_S to ρ_v , which is a subspace of ρ_S . Generalizing from this intuition leads to the hypothesis that context can be represented as a projection of a density matrix ρ onto a subspace represented by another density matrix:

$$\rho_X = P_X \rho$$

P_X is a projection operator constructed from one or more context words represented by X . A word v collapses from $|v\rangle = \rho|e_i\rangle$ to $|v_X\rangle = \rho_X|e_i\rangle = P_X\rho|e_i\rangle = P_X|v\rangle$, where $|e_i\rangle$ selects the column from ρ that corresponds to word v . It is curious to note that the v column of ρ , namely $|v\rangle$, is invariant under the transform P_v , because the components of $|v\rangle$ are all drawn from contexts containing the word v , so restricting the context to those containing the word v , i.e. applying P_v , doesn't change $|v\rangle$.

The full technical details of P_X still need to be worked out in relation to semantic space, however the effect of one context word $X = \{u\}$ can nevertheless be illustrated as ρ_w can be constructed directly from the Reuters collection via equation 1.7.

Table 1.5 depicts the state of "Reagan" in ρ_{Iran} and table 1.6 depicts the state of "Iran" in ρ_{Reagan} . Both tables represent unnormalized kets with the strength of association to other words represented as values in brackets.

This ket shows an collapse of the prototypical "Reagan" onto a state where associations relevant to the Iran-contra sense are prominently weighted. The ket depicted in table 1.6, however, reflects "Iran" in the context of "Reagan". One can clearly discern by comparing both kets that context effects are not symmetric.

4 Summary and Outlook

This article began with speculation that important aspects of human practical reasoning are manifest within the conceptual level of cognition referred to as conceptual space. Within conceptual space, information is represented in a geometric space, and inference has an associational, rather than a deductive, linear character. Our investigation focused on providing a formal

iran (827), president (538), arms (430), scandal (208), administration (118), sales (101), contra (97), speech (91), senate (83), house (80), profits (77), tower (77), contras (76), commission (73), deal (69), approved (60), conference (58), policy (58), aid (57), diversion (53), approval (52), fighting (51), poindexter (50), ronald (50), new (49), rating (48), decision (48), sale (47), funds (47), mistake (47), rebels (46), wrong (46), investigating (46), denied (45), knew (45), nation (44), news (42), secret (42), initiative (39), role (39), response (39), pct (38), defense (38), recollection (38), money (37), congress (37), televised (37), sell (35), adviser (35), gave (35), ...

Table 1.5. “Reagan” in the context of “Iran”.

arms (1522), iraq (1494), gulf (1432), war (939), oil (864), reagan (827), scandal (639), missiles (620), iranian (594), president (540), attack (528), offensive (504), sales (463), new (424), shipping (399), united (396), military (395), states (379), house (370), iraqi (364), contra (355), action (327), silkworm (291), news (285), hormuz (280), launched (270), diplomats (268), warned (258), southern (248), sale (247), major (244), attacked (243), tehran (239), strait (239), officials (236), kuwait (233), fighting (232), profits (230), north (225), senate (216), forces (213), foreign (212), washington (203), shipments (197), soviet (197), strike (196), attacks (193), american (191), crude (188), mln (185), ...

Table 1.6. “Iran” in the context of “Reagan”.

account of the non-monotonic effects on conceptual associations due to context. Our aim is to provide the foundations for operational practical reasoning systems. To this end, the conceptual space was approximated by a semantic space model which can be automatically derived from a corpus of text. Within semantic space, words, or concepts, are represented as vectors in a high dimensional space. Semantic space models have emerged from cognitive science and computational linguistics. They have an encouraging, and at times impressive, track record of cognitive compatibility with humans across a number of information processing tasks. Due to their cognitive credentials semantic space models would seem to be a fitting foundation for realizing computational variants of human practical reasoning. The particular focus was formalizing the non-monotonic dynamics of associations within semantic space due to context effects. Context is a notoriously slippery notion to pin down. Yet context effects seem to trigger many garden variety non-monotonic inferences.

It has recently been pointed out in a letter to the editor of a journal in physics and mathematics that semantic space models bear some interesting similarities with the framework of quantum mechanics (QM) [?]. We have explored the connection between the two in the light of human practical reasoning and our intention has been more to provoke thought than provide concrete answers. It was shown that there is a very close parallel between semantic space and the notion of a density operator in QM. In a nutshell, the non-monotonic dynamics of word associations due to context are formalized by means of the quantum collapse of the state of a word in semantic space onto a sense which is determined by context words. A product of the collapse is a change of state, or “meaning” of the word. As a consequence, word associations also change. QM is one of the few frameworks in which context is neatly integrated. Essentially, context is something akin to a quantum measurement which brings about collapse. We speculate these changes in word association are the primordial beginnings of non-monotonic inferences at the symbolic level of cognition.

The embedding of semantic space into QM is not perfect. A summary of the major problem areas is given as follows:

- In QM, eigenstates are orthogonal, whereas the senses of a word need not be.
- In QM, collapse results in an eigenstate, whereas the collapse of word meaning in semantic space may be partial.

Neither of these problems would seem to fatally undermine further research. Aerts , Broekaert and Gabaora [?] go so far to state “..generalizations of the mathematical formalisms of quantum mechanics are transferable to the modeling of the creative, contextual manner in which concepts are formed, evoked, and often merged together in cognition”. The theory developed in this article is complemented by realistic illustrations in an operational setting. The non-monotonic effects witnessed in the illustrations allow for cautious optimism.

The title of this account includes the phrase “quantum logic”. Where is the logic? The phrase “quantum logic” is promissory. It reflects our belief that quantum logic, or something akin to it, can be employed to provide an account of logics of “down below” meaning practical reasoning as it is transacted below the symbolic level of cognition. It is important to stress that the view of reasoning presented in this account does not rest on traditional conception of logic. Gabbay and Woods ([?], p63) speculate that a logic of “down below” could be “a logic of semantic processing without rules”. We feel that collapse of word meanings in semantic space falls very much within the ambit of such speculation and actually reinforces it.

There are yet many stones that need be laid to provide an adequate bridge between semantic space and quantum logic. In this regard, Widdows [?, ?, ?] have provided an important contribution with his quantum logic of word meanings and initial explorations into the lattice structure of vector subspaces. Such lattices provide the meeting point for Gabbay and Engesser’s pioneering investigation into the connection between non-monotonic logic and quantum logic [?].

Finally, there is the bigger picture. QM is emerging out of physics and permeating into other areas, for example, information retrieval [?], human language [?] and cognition [?]. This offers tantalizing possibilities and bizarre implications. (See Malin [?] for a wonderfully daring view of the philosophical implications of QM). In terms of semantic space, intriguing questions arise in relation to QM notions such as entanglement. For example, Aerts and Gabora [?] contend that the pet fish example mentioned in the introductions arises because the concepts “pet” and “fish” are entangled. If so, does entanglement manifest in semantic space and can it be exploited in an operational setting? Certainly we agree with Aerts and Czachor [?] that the emedding of semantic space models into QM is mostly unexplored. This article documents a tiny exploratory step.

Acknowledgments

The work reported in this paper has been funded in part by the Co-operative Research Centre for Enterprise Distributed Systems Technology (DSTC) through the Australian Federal Government’s CRC Programme (Department of Education, Science, and Training).

In the year 2000, the first author attended a course at ESSLLI on abduction given by Dov Gabbay and John Woods. Since then he has been following their bold and inspiring trail into the “New Logic”. The first author thanks Professor Shimon Malin and Dominic Widdows for their thought provoking discussion and deep insight.