Quantum Collapse in Semantic Space: Interpreting Natural Language Argumentation

P.D. Bruza
Queensland University of Technology
Australia
p.bruza@qut.edu.au

J.H. Woods
University of British Columbia
Canada
jhwoods@interchange.ubc.ca

Abstract

The interpretation of natural language utterances in argumentation is largely a tacit procedure, that is to say, a procedure transacted for the most part sublinguistically, inattentively, automatically and involuntarily. In the particular case of the interpretation of argumentative texts, the tacitness thesis provides that interpreters are able to discern intended messages without forming – and usually without being able to – propositional representations that wholly contain their contents. How is this done? In this paper, we propose the following theses: (1) Interpretation can be represented as collapse of meaning in a semantic space model. (2) Semantic collapse, in turn, can be likened to the quantum collapse of superpositional states of word meaning. (3) Non-trivial structural similarities with quantum collapse are discernible in matrix models of memory. The advantage of (3) is that it provides independent reason to suppose that quantum structures admit of psychological construal. The paper concludes with suggestions for further research.

Cognitive Agency

During the past forty or so years, the study of argumentation has achieved a considerable interdisciplinary momentum, impelled by important work in dialectical and dialogue logic (Barth & Krabbe 1982), argumentation theory – which is a loose assemblage of work by informal logicians, speech communication theorists, critical thinking (Van Eemeren *et al.* 1996) and artificial intelligence theorists (Norman & Reed 2004). A principal departure from mainstream mathematical logic is the emphasis given by these accounts to factors of agency and context. Unlike what we find in the mainstream of mathematical logic – set theory, model theory, proof theory or recursion theory – argumentation theory gives to *people* a load-bearing theoretical role.

It would be wrong to leave the impression that 20th century logic has no truck at all with agents and their circumstances. In a number of "non-classical" systems – notably, epistemic, temporal and deontic logics – agents have a place both notationally and semantically, although sometimes their semantic treatment is rather slight. Agent-relative systems that attempt to retain the basic methodologies of mathematical modeling are best seen as theories of *formal pragmatics*. Thus a major change in logic in its broader sense is its drift from an emphasis on formal *semantics* to a focus on an encompassing formal *pragmatics*.

Curiously enough, in most of these agent-oriented approaches not much is to be learned of the actual constitution of agents. In such systems a behaviour-type is analyzed by way of a set of rules or behavioural constraints, and agents are simply read off as those beings or devices that implement these constraints. So, for example, if a theory seeks to investigate the principles of human reasoning, it proceeds to lay down the norms that (it says) govern this practice, independently of any stand-alone consideration of how the reasoner is actually constituted, independently of what he is interested in and able to do. Similarly, if a theory has as its target the norms of correct argumentation, these are laid down with scant attention at best to how human arguers are actually put together. In virtually all such cases, the nature of the agent is inferred from the nature of the norms purported to govern the practice in question.

In the opinion of the present authors, it is better to reverse this order of precedence; that is, it is likely that one will have a better theory of reasoning (or arguing, or deciding, etc), only after determining the nature and wherewithal of actual human agents operating in real-time in the here and now. Given this shift in analytical priority, it is easy to see that the individual human is dominantly a being with cognitive interests. He wants to *know* what to believe and he wants to *know* what to do. He makes his way in life and owes his survival and prosperity to using his head, and he owes these uses to the way in which he is built for them. It is also apparent that individuals transact their cognitive agendas with comparatively few cognitive resources – information, time, and storage and computational capacity[1]. This being so, individuals tend to be proportionate in the setting of their goals. They favour targets whose attainment lies in principle within their reach. A related feature of cognitive agency is that the resource-boundedness of human agency places the individual in a cognitive economy in which, nearly always, resource-usage carries a cost. Overall there is an abiding necessity to be a *cognitive economizer*.

[1]Resource-modesty often runs to outright scarcity, but in the general case it is more a comparative matter. Contrast an individual's command of information, time, storage and computational capacity with that of an institutional agent such as Nato or the International Monetary Fund.

A further advantage of this agents-before-actions this priority-reversal is that it opens up the study of cognitive behaviour, including reasoning and argumentative behaviour, to a rich and complex body of findings from cognitive psychology. Logics that take seriously input from cognitive science have been called "practical logics of cognitive systems" (Gabbay & Woods 2005). A major difference between, on the one hand, the approaches to reasoning of standard logic and those agent-relative systems which fail to make independent provision for what agents are actually like and, on the other, theories to which agents are admitted as they actually are, warts and all, is psychologism. Psychologism is anathema to mainstream logicians and to those whose analysis of agency is wholly implicit in what count as the theory's norms. For those favouring a more robust notion of agency, psychologism is as welcome as it is unavoidable.

As the empirical record amply attests, there are a great many "shortcuts" taken by individual agents, often exhibiting the same levels of payoff had more expensive strategies been employed (Gigerenzer 2000). Of these, we wish in this paper to comment on two. One is the considerable savings achieved by cognitive processings that occur sublinguistically and unconsciously or, as we might say "down below". Another savings, and a related one, is a discourse-economy involving the suppression of articulation in favour of contextual cuing. This last is well-grounded in argumentational practice, both in the interpretation of an interlocutor's utterances and in the content-selection of one's own contributions.

Consider a simple-looking case. Peter and Rupert pass in the hallway of an IT research organization. Peter, a research scientist utters to Rupert, the business development manager, "How is it going with John?" This utterance is the tip of an ice-berg rich in implicit associations. Due to their shared context, Peter and Rupert both know that "John" refers to "John Smith" of "ACME Corp", who is negotiating a commercial license for "Guidebeam", a next generation web-based search technology. An interaction ensues. Peter argues that a research license should be offered to John should he reply negatively to the commercial license. Rupert argues to the contrary. We see in this modest exchange that a theory of argument must take note of two aspects of argumentational practice. One is argument-interpretation, and the other is argument-assessment, with the former taking net analytical precedence over the latter.[2]

In the not so distant future our information environment will feature all sorts of devices and displays. Even now, technologies loom in the background which process the above argument, draw appropriate context sensitive associations in order to flesh it out, and thereafter uses the result to query for emails, license documents, podcasts of relevant conversations etc., and tacitly retrieves these to prime Rupert and Peter's immediate information environment. For example, the licence document and associated emails could be brought up on the wall display should they be needed for further reference in Peter and Rupert's spontaneous hallway interaction.

[2]'Net' is a prudent qualification, owing to the presence in many accounts – and, as we believe, in actual practice – of a "charity" principle, according to which an interpretation of a person's argument which, if generalized to his general behaviour, would make him normatively subpar, is (defeasibly) a flawed interpretation.

It is worth noting that the historic role of logic was to lay bare the logical structure of human reasoning. Aristotle is clear on this point. The logic of syllogisms would serve as the theoretical core of a wholly general theory of real-life, two-party argumentation. Therefore, even at its historical inception, natural language is central to the logic of argumentation. This poses non-trivial challenges for technical solutions. In the day-to-day cut and thrust argumentation on the ground, sentences may not conform to grammatical norms, thereby compromising the precision of parsing technologies. What is more, the challenging problem of semantics becomes even more vexing as important elements of the semantics are sensitive to the shared, and often quite specific context of the interlocutors. The meanings of concepts and words are dynamic and have evolved in a community of practice. Consequently, such meanings are far more feral than the stylized examples portrayed in linguistics text books. This, in turn, impacts on the inferences drawn by the interlocutors. The role of inference is important, since common knowledge is part and parcel of shared context. Accordingly, as remarked above, for reasons of cognitive economy, things remain unsaid, because they are assumed known. Correct linguistic interpretation of the argument therefore relies on drawing appropriate inferences. Contrary to the syllogisms proposed by Aristotle, these often have the character of forming tentative hypotheses about the context and intentions of other speakers. In other words, a mode of inference at play is not deduction, but rather abduction (Gabbay & Woods 2005).

Interpreting Understatement

As noted just above, a dominant feature of communication, including the utterances that constitute *n*-person argumentional exchanges, is *understatement*. Understatements leave parts of content unexpressed; they leave those parts unarticulated. By and large, the message conveyed by an utterance or inscription is not wholly contained in it. In the general case no composite of the semantic interpretations of the lexical items of an utterance, as ordered by the utterance's or inscription's syntax, is sufficient to identify what, in making that utterance or inscribing that text, the utterer or inscriber has actually said. We have in this the long-recognized distinction between utterance-meaning and utter-meaning and, relatedly, the distinction between explicit meaning and tacit meaning. An important fact about understatements is that, while they do not themselves say what their utters intend to say in uttering them, this omission, typically does not produce communicational breakdown. Unlike misstatements, understatements are not in general impediments to the conveyance of intended meaning. Why is this so? How does it come to be the case that human agents are so adept at mapping utterances that don't contain the intended message to the intended message?

It is widely recognized that central to his success in effecting utterance-message mappings is the interpreter's ability to bring to bear considerations of common knowledge, contextual particularities, procedural conventions and empathy, the ability to put himself in the utterer's or inscriber's shoes (Gabbay & Woods 2005, chapter 9). What is often overlooked is the relationship in which an utterer stands to his own choice of utterance in the light of what he intends to say. On some tellings, the speaker has a privileged position in the construction of his own utterance-message mappings. Intuitively, this might strike us as right. For doesn't the speaker know what he intends to communicate prior to utterance-selection? Doesn't the interlocutor have to wait for the utterance before he takes a crack at fathoming the intended message? On this view, an utter's use of understatement is discretionary, abetted by his natural interest in economizing. Unlike his interlocutor, the utterer himself always possesses the wherewithal to embed the totality of his message in an utterance, albeit a longer one than he would normally have occasion to make in actual practice. But, again, there is little in the empirical record to sustain this opinion. (How often do we hear: "I didn't know what I wanted to say until the words were out of my mouth"?) For utterer and interlocutor alike, for both an utterer's utterance-message mappings as well as the hearer's decoding of them, a good deal of what happens happens down below – unconsciously inattentively, involuntarily and sublinguistically. It would appear, then, that the apparent asymmetry of access to intended meaning is largely an illusion. In attaching a speaker's meaning to his utterance or inscription, it is true that an interpreter is *responsive* to factors of context, common knowledge, procedural conventions, as well as approaching his task as an empathist. But in the general case, these factors operate in ways that outreach the agent's command and awareness. As for the speaker himself, he seeks for an utterance or inscription which, models those very same factors of context, common knowledge and so on, will convey the very message that a successful interpreter will be able to discern. Whereas the interpreter moves from utterance or inscription + these other features to message, the speaker moves from message + those other factors to utterance or inscription. Speakers and interpreters perform the converses of one another's tasks. This being so, in the interest of space we can now confine our remarks to follow on the interpreter's role in construing an utterance.

As briefly mentioned above, it has recently been proposed that the structure of inscription-interpretation is usefully modelled as a form of abductive inference. Informally speaking, what makes interpretations abductive is that they are inferred on the basis of their contribution to the overall coherence and efficacy of the discourse goals presently in play. A simple example will illustrate this point. When presented with the utterance, "How can you say that?", it bears on its interpretation that it arises in the context of a dispute rather than in the context of a language-learning class. Whatever may be said for the thesis of interpretation-as-abduction, it remains the case that the dynamics of speech interpretation are very complex, especially given the hiddenness of the various factors that operate down below.

It is well-known among AI theorists that there are no easy ways in which to automatically unpack the messages embedded in utterance and render them into propositional form. This is a considerable problem for the mechanization of textual interpretation. Since argument interpretation carries difficulties of its own, it remains a particularly serious problem in that context as well. What we now wish to do is to turn our attention to the dynamics of down below as they bear upon the difficult problem of the interpretation of texts, for which we shall employ semantic space models which are computational models of word meaning from the field of cognitive science.

The problem that motivates the semantic space approach is that when AI researchers build systems that are able to reason over substantial texts, the deployment of techniques for the *propositional* representation of the message embedded in the text fails to achieve the objective in a satisfactory way. Even so, on the assumption that texts embed messages of some sort, and on the further assumption that textual interpretation and inference somehow "gets at" those messages, it is reasonable that a message representation capability of a non-propositional kind. Accordingly, a semantic space analysis of textual interpretation may be viewed as a contribution to a "logic of down below" (Bruza, Widdows, & Woods 2008). As we shall see two sections hence, a distinctive feature of our semantic space approach is the tie it postulates between the interpretation of natural language argumentation and what physicists call *quantum collapse*. It is our further conjecture that the creative inarticulacies involved in both the understated transmission and interpretation of messages has a natural explanation in a semantic space model subject to these same quantum constraints.

Semantic Space

Cognitive scientists have produced an ensemble of models which have an encouraging, and at times impressive, track record of replicating human information processing, such as word associations norms. These models are generally referred to as *semantic space*. As used here, the term "semantic" derives from the intuition that the meaning of a word derives from the "company it keeps", as the linguist J.R. Firth (1890-1960) famously remarked.

Although the details of the various semantic space models differ, they all process a corpus of text and "learn" representations of words in high dimensional space. Semantic space models are interesting in light of the scenario presented above, since they open the door to gaining operational command of socio-cognitive "meanings" in a community of practice together with mechanisms to replicate our ability to draw context-sensitive associations within the scheme of an argument, or dialogue.

To illustrate how the gap between socio-cognitive semantics and actual computational representations may be bridged, the Hyperspace Analogue to Language (HAL) semantic space model is employed (Burgess, Livesay, & Lund 1998). HAL constructs a matrix whereby the columns and rows correspond to words and the values of a cell represents the strength of co-occurrence of one word seen in the context of another. The strengths are computed by sliding a context window of fixed size over the text 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. Remembering Firth's quotation above, all words in a given context window are the "company" of the last word in the window. Intuitively, if the window size is set too large, spurious co-occurrence associations are represented in the matrix. Conversely, if the window size is too small, relevant associations may be missed. In most studies, a window size of between eight and ten has proved optimal. Consider the trace "President Reagan ignorant of the arms scandal". Table depicts the HAL matrix of the example trace with a window size of five. Each row i in a HAL matrix represents accumulated weighted associations of word i with respect to other words which preceded i in a context window. Conversely, column i represents accumulated weighted associations with words that appeared after i in a window. If the word order is not of interest, the matrix can be added to its transpose giving rise to a symmetric matrix. A given column vector then represents the "meaning" of the word associated with the column. The semantic association between two words a and b can then be computed by measuring the cosine of the angle between the corresponding vector representations - the smaller the angle the stronger the semantic association. An alternative measure of semantic association used in the literature is the Euclidean distance between the vector representations of a and b. Irrespective of the means employed to compute semantic association, semantic space is *clumpy* – words with similar meaning will tend to cluster. A semantic space like a HAL matrix can be dimensionally reduced via for example, singular value decomposition, a theorem from linear algebra. Replication of a variety of human information processing tasks relies on dimension reduction as it apparently picks up higher order associations between words which are not captured by straight co-occurrence.

	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: A simple semantic space computed by HAL

Semantic space models like HAL are normally run over a large corpus of text, sometimes millions of words. In this way, a global semantic space can be computed, which has been of most interest to those cognitive scientists who deploy semantic space models to replicate aspects of human information processing. It is important to bear in mind the word "meanings" are relative to the corpus. Consider a community of practice. A corpus quite naturally develops around it in the form of electronic documents, emails, on-line postings to a forum, etc. This corpus can be used to compute a semantic space. In this case, the "meanings" of the words are relative to the community - they are computational approximations of socio-cognitive meanings harboured by the human agents in the community. What is more, these meanings are not fixed, they will evolve as the corpus evolves. Bear in mind, these meanings have been computed without grammatical processing and by their very nature are rich in associations. This reflects cognition itself whereby it is purported that the upper symbolic level level of cognition is where higher order linguistic structures are manipulated, such as propositional representation. In the level below, concepts are purported to a have a geometric a representation whereby notions such as context-senstive similarity are naturally expressed (Gärdenfors 2000). In other words, semantic space would seem to be an appropriate computational approximation of the sublinguistic issues involved in interpretation mentioned above.

Interpretation and Collapse of Meaning

A big question lurking in the scenario-argument previously sketched is how to effectively model the interplay between meaning, context and human abductive inference in relation to interpretation. Surprisingly, quantum mechanics (QM) gives rise to some innovative and possibly ground breaking answers in relation to this challenging question. In recent times there have been number of speculative attempts to explore the connection between density matrices and semantic space models (Aerts & Czachor 2004; Bruza & Cole 2005; Widdows & Bruza 2007). In these works, a word in semantic space may be likened to a quantum particle in the following sense. In the absence of context it is in a superposed state - it is a collection of all the possible meanings of the word. Seeing the word in context, however, gives rise to a "collapse" of potential meanings onto an actual one. As mentioned above, this collapse of meaning happens "down below" – it is a prevalent, unnoticed facet of interpretation during argumentation.

Suppose that an argument begins with the affirmation "Reagan got it wrong with Iran". The word "Reagan" has several possible senses, or basis states. For example, there is the run of mill Presidential Reagan dealing with congress, Reagan's trade war with Japan, the Iran-Contra scandal, and even the sense pertaining to the aircraft carrier U.S.S Ronald Reagan. When encountering "Reagan" in the context of "Iran", a collapse of meaning occurs onto the basis state corresponding to the Iran-Contra scandal. As mentioned previously, a density matrix representing a basis state has a single eigenvector. This matters, since the eigenvector is a source of relevant context-sensitive associations to the Iran-Contra scandal such as "arms", "scandal", "illegal" etc.

The density matrix corresponding to the word "Reagan" can be constructed as a linear combination of density matrices, each corresponding to a given sense. Consider the traces "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". A symmetric HAL matrix can be computed from each of these. As more traces are encountered, it becomes clear that they start to cluster and a sense begins to emerge – for example, "Reagan" in the Iran-Contra sense, the missile treaty with the Soviets, etc. Each sense can be represented as a density matrix which is quite easily derived from summing the HAL matrices of the associated traces. In addition, a probability can be ascribed the to a given sense. For example, the density matrix ρ_r for the meaning of the word "Reagan" can be formalized at the following linear combination:

$$\rho_r = p_1\rho_1 + \dots p_m\rho_m$$

where each ρ_i is a basis state representing one of the m senses of the word "Reagan" and the probabilities p_i sum to unity. This is fully in accord with QM whereby a density matrix can be expressed as a weighted combination of density matrices corresponding to basis states (senses). There is no requirement that the eigenvectors of the basis states be orthogonal to one another.

This is a very important point. Intuitively, it is unrealistic to require the senses of a word meaning be mutually orthogonal. The above equation is depicted in figure 1.

A striking aspect of figure 1 is its similarity to a matrix model of human memory published in the psychological literature (Humphreys, Bain, & Pike 1989). Seen this way, the density matrix comprises vertical slices, each slice corresponding to a sense (basis). The analog in the matrix model of memory is a tensor of rank 3, which has a three dimensional structure. Rank-three memory is a superposition of memory traces. For accessing episodic memories (memories of specific events), a context vector is required. This is relevant to our account, as the Iran-Contra scandal can be considered as a specific event. In the density matrix representation, each vertical slice corresponds to a context, and in the terminology of the matrix model of human memory, the eigenvector of the slice, can be considered a context vector.

The collapse of meaning

The intuition we will attempt to develop is the collapse of word meaning due to context is akin to a cued-recall retrieval operation from human memory. More specifically, context is modelled a projection operator which is applied to a given density matrix corresponding to the state of a word meaning resulting in its "collapse". The probability of collapse p is a function of the scalar quantity resulting from matching. The analogy with orthodox QM is the following - a projection operator models a measurement on a quantum particle resulting in a collapse onto a basis state.

In the matrix model of memory (Humphreys, Bain, & Pike 1989) , memory representations can include items, contexts or, combinations of items and contexts (associations). Items can comprise stimuli, words, or concepts. Each item is modelled as a vector of feature weights. Feature weights are used to specify the degree to which certain features form part of an item. There are two possible levels of vector representation for items. These include:

- modality specific peripheral representations (e.g., graphemic or phonemic representations of words)
- modality independent central representations (e.g., semantic representations of words)

Our discussion will focus on the latter, given the assumption that semantic spaces deliver semantic representations of words.

Memories are associative by nature, and unique representations are created by combining features of items and contexts. Several different types of associations are possible. The association of interest in our running example is a two way association between a word "Reagan" and a context word "Iran". In the matrix model of memory, an association between context and a word is represented by an outer product of the vectors corresponding the meanings of "Reagan" and "Iran". Seeing a given word (a target) in the context of other words (cue) forms an association which probes memory. The object being probed is a density matrix corresponding to a word or concept, the meaning of which, as we have seen, is a superposition. As in orthodox QM, the probe can be formalized as a projection operator which essentially retrieves a particular slice from figure 1. In this case, the collapse of meaning is total. In other words, no ambiguity remains about the state of the word's meaning.

Motivating the collapse of meaning using the matrix model of memory introduces a deviation from orthodox QM. The measurement devices in QM are precise in the sense that a measurement collapses a quantum particle onto a basis state. Context, however, can be imprecise. For example, there are actually two senses of "Reagan" in the context of "Iran". The first, and more probable sense, associates with the Iran-Contra scandal. The other is the freeing of the hostages from the American embassy in Teheran at the beginning of Reagan's presidency. In terms of the figure above, what is retrieved from memory by the probe is a mixture of two slices, one slice (basis state) corresponding to each of the two senses just mentioned. In quantum terminology, the result of the measurement is a superposition, not a basis state, albeit that the resulting superposition involves far fewer possible senses that the original. In other words, the collapse of meaning is not total.

Conceptualizing interpretation in this way aligns well with the issues of cognitive economy and abduction mentioned above. If the the collapse is not complete, it is economic to "clean up" the resulting superposition by inattentively hypothesizing ("abducing") the most probable sense, especially if one sense has a significantly higher probability than others. If this condition is not met, the interpreter may then resort to the more cognitively demanding mode of interrogation characterized by "What do you mean by X?". It is not hard to imagine that such heuristics could be deployed in a computational argumentation system based on density matrices derived from semantic space around a community of practice.

More generally, the study of faulty projection operators on human memory may generate some insights into the errors humans make in inference. If the interpreter does not resort to interrogation, then they may "clean up" the resulting superposition by "abducing" the wrong sense, which in turn leads to inappropriate context-sensitive associations being produced "down below". It may well be that these feed into higher level inference processes at the symbolic level of cognition and thus ultimately lead to the drawing of erroneous

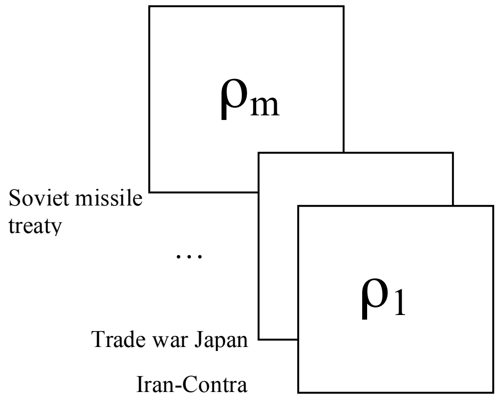

Figure 1: ρ_r: the density matrix representation of the meaning of "Reagan"

conclusions within the scheme of the argument. Viewed this way, many so-called errors of inference may have little to do with inference at all, but rather faulty collapse of meaning in memory. In is interesting to note that some studies of analogical mapping show analogical reasoning is confounded with processes in human memory. The same may also be said in regard to default reasoning. When confronted with the utterance, "Tweety is a bird", the interpreter will likely hold the hypothesis that "Tweety flies". However, when confronted with the subsequent utterance, "No, Tweety is a penguin", the interpreter will drop the hypothesis. The subfield of Artificial Intelligence known has non-montonic reasoning proliferates with studies of the reasoning processes underlying such examples. It is debatable, however, as to whether any reasoning is actually going on. The example can be explained in terms of two probes to memory resulting in two collapses of meaning of "Tweety".

Summary and Outlook

It is widely accepted that much of a human agent's cognitive behaviour involves appreciations of and responses to contexts, meaninings, memories and intentions – as well as of targets and their standards of attainmant – that are implicit. Owing to the resource-limitations characteristic of individual agency, the present authors conceive of the processes of cognition "down below" as strategies that abet the agent's interest in using his scant resources with requisite economy. Such strategies are, first and foremost, *fast*. Since operations down below are dominantly sublinguistic, it is necessary to postulate cognitive devices and procedures that are put into play without the necessity of propositional representation. In the particular case of the interpretation of an argumentative text, these same factors are also dominantly present. This being so, the human cognizer manages to grasp the message conveyed by such a text without contriving a – and in, in general, without being able to – a propositional expression of it,

It is hardly surprising that the mechanics of cognitive behaviour down below present the experimental psychologist with challenges so robust that the entire field is still presently much *terra incognita*. Ethical considerations alone are an experimental deterrent. Accordingly until these impediments are more effectively subdued, the logic of down below – and its role in a comprehensive theory of argument-interpretation – must be largely conjectural, that is to say, an exercise in theoretical *abduction*. What is needed for the task at hand is a model of textual interpretation that takes these attributes and constraints seriously into account. The semantic space model certainly fits this bill, and it has to our mind the further advantage of giving to associationist assumptions a credible theoretical role.

On the approach developed here, interpretation is semantic collapse. Taken on its own, "quantum collapse" of meaning is a metaphor, albeit an attractive one. Without further explication, it is a metaphor in no fit state to bear theoretical burdens that we intend for it. Accordingly, our further abduction is that semantic collapse is a form of the quantum collapse of superpositions, except that the orthogonality requirement of the latter is given a somewhat relaxed providence as regards the former. If the quantum explication of semantic collapse holds water, then a good deal can be inferred about how textual interpretations are achieved. This is advantage enough to discourage dismissal of the quantum hypothesis out of hand. Another advantage of the quantum collapse hypothesis is its apparent conformity with the matrix model of memory developed in the psychological literature. In any case, it is easy to see why this is a similarity from which the present authors would derive some comfort. The packaging, storage and retrieval of memories is a paradigm of cognitive processing down below. The structural similarities between the matrix model of memory and the quantum collapse treatment of the semantic space model is especially welcome, if for no other reason than that it provides independent reason for supposing quantum structures capable of bearing non-trivial psychological readings.

We offer the hypotheses developed here without categorical assurance, and certainly not as the final world. Their intended function is the stimulation of further enquiry. On our own agenda is the investigation of Nelson & McEvoys's (Nelson & McEvoy 2007; Bruza *et al.* 2008) linking of word association and quantum entanglement, as well as the possibility of embedding the down below aspects of error in a general account of that subject. The value of the former project speaks for itself. The interest of the second requires a word or two of explanation. In (Gigerenzer 2000)and (Gabbay & Woods 2007), it is observed that certain classes of error have the net advantage of facilitating speedy and accurate learning. This leaves the question as to whether these constructive errors are random occurrences – luck of the draw – or whether somehow the human agent is adept at "selecting" them. If the latter obtains, it is plainly a further datum for the logic of down below.

Acknowledgment This research was supported under Australian Research Council's Discovery Projects funding scheme (project number DP0773341).

References

Aerts, D., and Czachor, M. 2004. Quantum Aspects of Semantic Analysis and Symbolic Artificial Intelligence. *Journal of Physics A-Mathematical and General* 37:L123–L132. http://uk.arxiv.org/abs/quant-ph/0309022.

Barth, E., and Krabbe, E. 1982. *From Axiom to Dialogue: A Philosophical Study of Logic and Argumentation*. Berlin and New York: De Gruyter.

Bruza, P., and Cole, R. 2005. Quantum Logic of Semantic Space: An Exploratory Investigation of Context Effects in Practical Reasoning . In Artemov, S.; Barringer, H.; d'Avila Garcez, A. S.; Lamb, L. C.; and Woods, J., eds., *We Will Show Them: Essays in Honour of Dov Gabbay*, volume 1. London: College Publications. 339–361.

Bruza, P.; Kitto, K.; Nelson, D.; and McEvoy, C. 2008. Entangling words and meaning. In *Proceedings of the Second Quantum Interaction Symposium (QI-2008)*. College Publications.

Bruza, P.; Widdows, D.; and Woods, J. 2008. A Quantum Logic of Down Below. In Engesser, K.; Gabbay, D.; and Lehmann, D., eds., *Handbook of Quantum Logic and Quantum Structures*, volume 2. Elsevier. (In press).

Burgess, C.; Livesay, K.; and Lund, K. 1998. Explorations in context space: words, sentences, discourse. *Discourse Processes* 25(2&3):211–257.

Gabbay, D., and Woods, J. 2005. *The Reach of Abduction: Insight and Trial*, volume 2 of *A Practical Logic of Cognitive Systems*. Elsevier.

Gabbay, D., and Woods, J. 2007. Error. *Journal of Logic and Computation*. (In Press).

Gärdenfors, P. 2000. *Conceptual Spaces: The Geometry of Thought*. MIT Press.

Gigerenzer, G. 2000. *Adaptive Thinking: Rationality in the Real World*. Oxford University Press.

Humphreys, M.; Bain, J.; and Pike, R. 1989. Different ways to cue a coherent memory system: A theory for episodic, semantic and procedural tasks. *Psychological Review* 96:208–233.

Nelson, D., and McEvoy, C. 2007. Entangled associative structures and context. In Bruza, P.; Lawless, W.; Rijsbergen, C. v.; and Sofge, D., eds., *Proceedings of the AAAI Spring Symposium on Quantum Interaction*. AAAI Press.

Norman, T., and Reed, C., eds. 2004. *Argumentation Machines. New Frontiers in Argument and Computation*. Kluwer.

Van Eemeren, F.; Grootendorst, R.; Henekemans, F.; Blair, J.; Johnson, R.; Krabbe, E.; Plantin, C.; Walton, D.; Willard, C.; Woods, J.; and Zarevsky, D., eds. 1996. *Fundamentals of Argumentation Theory*. Lawrence Erlbaum Associates.

Widdows, D., and Bruza, P. 2007. Quantum information dynamics and open world science. In Bruza, P.; Lawless, W.; Rijsbergen, C. v.; and Sofge, D., eds., *Quantum Interaction*, AAAI Spring Symposium Series. AAAI Press.