Context, Latent Semantic Analysis (LSA), and atypical cognition.

John R. Skoyles
j.skoyles@ucl.ac.uk
http://www.human-existence.com/

Latent semantic analysis (LSA), a tested model of cognition, suggests that context processing depends upon the information contained in higher-order associations (rather than first-order or direct ones). LSA also indicates that higher-order associations underlie the nonreferential “sense” properties related to context of meaningfulness such as semantic closeness and synonymity. This review focuses first upon the potential of LSA to understand the processing of context in everyday cognition such as language, harm awareness, sociability, sense of self and other, affect, motor control and perceptual gating. Second, it focuses upon LSA and atypical cognitive development in these areas. It argues that initially context extraction depends upon processes that are modularized, and that only following this extraction, is context created when this information is consolidated into "upstream" cognitions. Early failure in such a module or such consolidation thus can cause atypical information processing by selectively depriving such upstream cognitions of extracted context.

Keywords: context, context cognition, co-occurrence information, language development, latent semantic analysis, LSA, meaning, semantics, social psychology, atypical cognition.

1. Introduction

1.1. What is context?

Context is a descriptive word that is often employed in cognitive (Sperber & Wilson, 1986), A.I. (Ekbia & Magnitman, 2001), pragmatics (Van Dijk, 1977), social theory (Rohlfing, Rehm, & Goecke, 2003), neurocognitive (Koechlin, Ody, & Kouneiher, 2003) and social neuroscience models (Todorov, Harris, & Fiske, 2006). Broadly, it concerns the associations that exist between entities and the episodes in which they appear. Unfortunately, context as a descriptive word is ambiguous and can refer to at least two kinds of information. (1) Situational associations between episode constituents due them sharing the same circumstance or situation of being together in that episode (the information contained in first-order [or direct] associations). And (2), associations due to substitutability or near substitutability of constituents in regard to appearing together in such episodes (the information for this relates predominately to that contained in higher-order associations). This distinction is rarely made in models about context.

To better understand context as involving the information found in substitutability associations, the nature of an episode needs to be somewhat clarified. An episode is an information processing chunking in regard to some descriptive or functional relevance. Associations in regard to such relevance can exist both between constituents in a particular episode and between constituents of different ones. Substitutability between such entities is a kind of meta-association in that it concerns not a real association but a potential one –that of having the same associative relationships as possessed by another constituent entity. What follows explains this in terms of linguistic entities (words) and linguistic episodes (sentences), but it is universal to other kinds of entities that can exist in the same kind of relationship with each other within and across episodes.

1.2. Substitutability

Substitutability associations between entities occurs most familiarly with synonyms, such as “big” or “large”. While synonyms hardly ever appear near each other in the same word situation (they are not linked by a direct association), if one word appears with a particular episode set of words, so could its synonyms – it does not matter whether someone says or writes, “the large stone broke the weighing-machine”, or “the big stone broke the weighing-machine”. They both occupy the same context since they share much the same first-order associations with “broke”, “stone”, and “weighing-machine” (and indeed all other English words). They are intersubstitutability "associated" in this way even though as two words they may never appear together in the same sentences (see Fig. 2). Another example are semantically close words such as “admiral” and “general” which while not synonyms partially associate with many other words in the same substitutable manner in regard to words such as “command”, “orders” “military” and “rank”, though they are not in terms of other ones, such as land and sea, or ships and tanks.

It is the core claim of this discussion that in many cognitive, linguistic, social and neuroscience models involving “context”, that what is being modeled at an information processing level would be better approached as involving this latter kind of substitutability "context" association than first-order "situational" ones. Further that the information that links entities in this substitutability context manner in episodes is extracted from that contained in their indirect, latent, or higher-order associations. This latter claim rests on the observation that latent class type models such as LSA (Landauer, 2002; Landauer & Dumais, 1997; Landauer, Foltz, & Laham, 1998; Landauer, McNamara, Dennis, & Kintsch, in press) link synonyms (and other semantically close words) routinely using the information extracted from higher-order associations. The mathematics underlying LSA does not constrain such modeling to the linguistic domain and so is potentially as relevant (or available) in all areas where episodes constructed of entities are processed by the brain and issues of substitutability occur. As noted later, this includes areas as diverse as sociability, awareness of self and others, and sensorimotor processing.

1.3. Context cognition

The third claim of this discussion is that context substitutability information when incorporated into cognition plays a central and crucial role in a broad range of domains. For example, context enriched cognitions will be competent at (1) the detection of the properties of entities that might be missing, hidden, or unfamiliar in an episode (episodic inference), (2) recognizing novel episodes that are equivalent (or near substitutes to past ones) in their properties from different sets of elements (episodic equivalence), or (3) identifying the constituents of episodes containing redundant members that because of this can be elliptically removed (episodic reduction). These are key processes for learning and behavioral adaptability in everyday awareness, interaction and accomplishment. Without episodic inference, it is difficult to learn about the properties of new entities such as unfamiliar words. Without episodic equivalence, it is difficult to see unfamiliar situations as being the same as familiar ones even though they are superficially made of different constituent elements. Without episodic reduction, it is difficult to learn more efficient or economical ways of being aware, attending, communicating, achieving goals or performing tasks.

1.4. Atypical cognition

The fourth claim of this discussion is that reasons exist to suppose (see below, section 4.1.) that the two kinds of association information (first-order and higher-order) linking entities and episodes are processed differently in the brain with the latter’s computations being specialized, modularized, and then transfer and consolidated ("uploaded") into everyday cognition. Following this, it is observed that as a result, such specialized and modularized processes could be developmentally impaired without this impairment directly happening to the cognitive processes that actually use this context information. Such impairment, partial or near complete, in these modules (or the processes that transfer or consolidate their processed information) would result in characteristic forms of atypical cognition. Anomalously these cognitions would be intact at the physical level of their neural implementation (for example, in terms of their functioning as neural networks) but functionally impaired due to lacking uploaded context information which makes their information processing entirely based on the information contained in first-order associations (not requiring such uploading). Moreover, since the extraction of such information is modularized, it could leave cognitions across diverse domains in this atypical state.

This paper in its fifth claim theoretically explores the plausibility, nature and possible occurrence of such atypical cognition. Whether such atypical cognition actually exists is not a primary concern of this discussion since it is unlikely that a theoretically postulated form of cognition could be sufficiently specified as to immediately enable its identification without more sophisticated investigation. However, it could be that such an atypical form of cognition while not characteristic of any recognizing condition, could exist in the altered cognitions in those identified as “unusual” or labeled as retarded. Atypical cognitions of an ill-defined nature, however, do occur in one condition – autism – and this raises the possibility that some aspects of the cognition in some individuals with this condition might arise from the information processing atypicallity raised in this paper. However, the theory proposed is more broad than just this condition and might be relevant to understanding other forms of impaired or atypical condition. Also while it is possibly relevant to understanding some aspects of autism, it should be strongly emphasized that autism is characterized by many other factors for which this theory is not relevant.

1.5. LSA and higher-order associations

Core to this discussion, is developing the cognitive implications for everyday cognition of latent semantic analysis (LSA) (Landauer & Dumais, 1997; Landauer et al., 1998, see also Burgess and Lund, 1997; Landauer et al., in press). This offers a novel approach to context that provides a computer-based and a substantially tested account of how context and intersubstitutability might be extracted from episodes constituent associations, and how they might contribute to cognitive development and normal cognitive functioning in the domain of language and word meaning --though its authors recognize it might have a potentially wider application.

LSA can be viewed as having two parts. First, it offers a mathematical account of the presence and extraction (by the singular value decomposition of a matrix and various transformation) of a previously unstudied form of the information contained in the past word usage: their higher-order associations from which context and substitutability can be extracted. This it does by deriving the "latent" higher-order associations that exist between words (elements) that appear in thousands of word strings (episodes). This computational reduction by LSA of such higher-order associations to context provides cognitive science with the opportunity to develop potentially precise definition of the nature of context and its role in cognition. Importantly, the mathematical framework of LSA, though applied to linguistic meaning, is applicable, in theory, for modeling the role of context in the information processing of a wide range of other types of cognition (Landauer & Dumais, 1997, p.228), and thus also in nonlinguistic domains.

Second, LSA and its modeling of context in terms of higher-order associations has had considerable empirical success in simulating the measured performance of many aspects of language such as word sorting and category judgments, estimations of passage coherence, and the quality and quantity of knowledge contained in student psychology essays (Landauer, 2002; Landauer & Dumais, 1997; Landauer et al., 1998; Landauer et al., in press). Moreover, LSA has notably been effective in modeling one aspect of child development that has until now been resistant to explanation: how children learn the meaning of unknown words

(Landauer & Dumais, 1997). Formerly seen as enigmatic, LSA shows that given the context information provided by surrounding known words, children possess sufficient information (contained in their knowledge of past word usage) to guess what an unknown word might mean. This success of LSA in modeling the different contributions of context to language processing strongly argues that it captures some nontrivial aspect of information processing that has been ignored until now in the scientific attempt to model human cognition. Since context nor the type of information contained in higher-order associations are specific to language, this raises the question as to whether LSA-type information processing might not also be involved in as yet unappreciated ways in the typical functioning and the typical development of other types of cognition, an issue which so far has gone unexplored. Further, since context plays an essential role in normal cognition, processes that extract it, could cause, when defective, forms of cognitive handicap that have yet to be recognized and identified as a cause of impairment or difference in retarded or atypical people.

LSA is a particular example of latent class analysis (Hofmann, 1999; McCutcheon, 1987). So far it is the only model at present using this approach that has been widely and successfully tested against real data. In the following discussion, because of this, the focus is upon LSA. Moreover, it should be appreciated that the LSA approach itself is undergoing active development (Blei, Ng, & Jordan, 2003; Hofmann, 1999; Landauer et al., in press; Steyvers & Griffiths, in press). At present, simplifications in its implementation result in it ignoring much of the information available to word learning such as syntax, word order, and real world associations (Landauer, 2002). However, ways have been proposed of incorporating them such as those that involve generating "probabilistic topics" (Hofmann, 1999; Steyvers & Griffiths, in press). The idea that the information contained in higher-order associations requires the stage of dimensional reduction as carried out in LSA has also been questioned and alternatives have been suggested (Hofmann, 1999; Kwantes, 2005; Steyvers & Griffiths, in press).

The implications of LSA for cognition are discussed here without reference to such issues of its implementation or further possibility of implementation. Moreover, the observations made in this theoretical analysis, it should be noted, could also apply more generally to other models created in the future that are based upon or related to latent class analysis. The theoretical discussion here is concerned with the functional details of LSA and not details that are not pertinent to this such as its particular method of data reduction and data transformation.

2. Latent Semantic Analysis (LSA)

2.1. Origins of LSA

Historically, LSA originated in the need of computer scientists to find an automatic means to retrieve documents by keywords (Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990). (For this reason it is sometimes called latent semantic indexing, LSI.) Computer scientists faced the problem that the link between keywords and the words in sought-after-documents depend upon context. If a searcher types in, say, the keywords, "film" and "Marilyn Monroe", they want to retrieve not only documents that mention "film", but also related synonyms found in the same context (such as "movie", "Hollywood" and "motion picture"). Further, they want to retrieve only documents that contain the word "film" that fits in with the context of "Marilyn Monroe", and not ones containing "film" when it means "thin coating". Computationally, the context sensitivity needed to identify such synonyms and homograph meanings cannot be reduced to the information contained its direct co-occurrence associations with adjacent ones (Haugeland 1989). Computer scientists to overcome this developed methods to determine the synonyms and homographs of keywords by extracting from texts a type of information that had been not previously investigated: the higher-order (or indirect) associations that words have with each other. (Previous work upon the information in texts [due to limits upon computer power] had been confined only to first-order or direct co-occurrence associations.) Computer scientists using information extracted from higher-order associations have been able to detect synonyms and homographs efficiently (Deerwester et al., 1990; Landauer & Dumais, 1997; Landauer et al., 1998). These techniques have been used to create powerful keyword document retrieval programs (Deerwester et al., 1990). Though commercial confidentially prevents the open publication of details of contemporary search engines such as Google, readily available web newsletters (such as http://www.free-seo-news.com/newsletter147.htm) report they incorporate, at least in part, LSE-related techniques.

2.2. 'Behavioral Properties' and Meaning

The importance of LSA and its conversion of higher-order associations into context and substitutability information is that it provides the first empirically tested account of the phenomena of "meaningfulness".

Meaning until the advent of LSA had been quixotic to scientific investigation. After all, there is the problem of what does it "mean" for a word or sentence to mean something? Indeed, philosophers since Socrates have studied this question for two and half thousands years without providing an answer. One tentative approach for understanding its nature from a neurocognitive perspective is to explore what might be called its "behavioral properties" (in contrast to the "semantic" ones studied traditionally by philosophers that concern reference or sense (B. Hale & Wright, 1997)). "Behavioral" is used here descriptively to refer to how the experience of meaningfulness by the brain effects judgments about words in terms of their similarity, their intersubstitutability in sentence contexts, or their predictability in certain sentence contexts.

The similar-dissimilarity property: All semantic words have various degrees of similarity and distance of meaning. "Vast" and "big" have a great deal of similarity and closeness, "vast" and "small", a smaller amount, while "vast" and "kiss" are very distant, for example. What might underlie this property of semantic similarity, however, is uncertain: the resemblance or lack of meaning between words cannot be inferred from their visual or sound identity nor their close association with other words (the level of first-order associations). Nothing, for example, about the perceptual nature of "vast" and "big" – its letters and phonemes – tells us that they are synonyms. Nor is this information provided by their immediate associations with surrounding words in the sentences containing them (Haugeland, 1989). In spite of this, humans (or rather human brains, since human competences in this arise from those of the human brain) have an intuitive grasp of which meanings are similar and which are not. This suggests that our experience of words might in some unknown way provide our brains with the information with which to make such judgments

The synonymy intersubstitutability context property: Meaning has the property that words that have roughly the same meaning – synonyms – are context intersubstitutable (Sparck Jones, 1986). Consider "vast" and its synonyms, "big", "large", or even "gargantuan". Each of these three words can be substituted in most sentence contexts without significant change of sentence meaning. "The big rock broke the wagon", for example, means roughly the same as "The vast rock broke the wagon", and even, "The gargantuan rock broke the wagon" (see Fig 5.). Like the distance-similarity property, while the human brain can understand intuitively which words can be swapped, it has proved impossible by philosophical or other analysis to specify the nature of the processes involved.

The predictability usage property: Linked to context is the property that words do not appear randomly in normally encountered sentences. Expectations and constraints exist – word usage patterns – about which individual words tend to appear with which other ones

(Gough, Alford, & Holley-Wilcox, 1981). This can be experimentally shown with cloze sentences in which words are omitted and human brains face multi-choice options to pick the missing word (Taylor, 1953). If human cognition is presented with the incomplete sentence, “The X stone broke the wagon”, and asked to predict the most likely word out of “big”, “forgetful”, and “sweet”, the word picked would be “big”. “The big stone broke the wagon” has an expected pattern of usage, that is absent – even though they are readily comprehendible – in the sentences, “the forgetful stone broke the wagon”, and “the sweet stone broke the wagon” (understandable, for example, if they were to appear in a “fairy story”). Indeed, even without suggested word choices, if single words are cut from writing, the human brain usually can guess using the remaining before and after words (depending upon sentence position) in up to half of them (Miller & Coleman, 1967, fig 1). Further, if the next word in a piece of ordinary writing is covered, and an individual has not read ahead, they will be able to guess it in about one in four times (Gough et al., 1981, pp. 91-92). The ability to fill in missing words from sentences is used in the testing of the comprehension progress of language learners (Storey, 1997), and the assessment of language competence, for example, in TOEFL (Test of English as a Foreign Language) certification. This argues that the way words are used contains, in predictable ways, information that is central to the brain’s comprehension of word meaning.

2.3. LSA and higher-order associations

LSA is based on extracting the information contained in the complex patterns of association contained in word usage found in the many tens of thousand of sample sequences of words. When they are analyzed, the information they provide is sufficient to enable LSA to make synonym identification with comparable competence to that of humans. For example, applicants with English as a second language to US colleges if given a word and four possible synonyms get 64.5% correct; the LSA model, 64.% (Landauer & Dumais, 1997, p.220).

The information that enables LSA to do this comes from various kinds of word associations present within word usage. These include ones that have traditionally been extensively studied such as word frequency (how common a word is), and word frequency co-occurrence associations (how commonly a word is found next to another one) (Dagan, Marcus, & Markovitch, 1995; Haugeland, 1989). LSA, however, depends (and it turns out, much more so) upon the less studied information that exists at the higher level by which words associate by indirect (or higher-order) associations.

2.4. Types of associations amongst words

First-order associations: words occur with different frequencies next to each other in sentences very often such as soap and wash, or very rarely such as soap and aardvark. Since it is statistically easy to count different words, and find how often they appear near to each other, these frequency associations have been studied in the past (Dagan et al., 1995; Haugeland, 1989). While such an association might seem to provide a ‘context’, it has not been possible to capture with the information contained in such direct associations, the type of context that underlies semantic distance or synonym intersubstitutability. Notions of context based upon first-order associations thus fail to articulate an important aspect of information processing done by the human brain.

Second-order associations: Any two words have many more indirect associations via further words than the direct one between them. Consider the synonyms “big” and “vast”: the human brain judges that there is an informational relation between them (they are synonyms), yet they rarely appear directly together in the same sentences. Normally, people do not say things like “the big and vast stone broke the wagon”; if one word is used, it will not be repeated as it is redundant. Instead of the mutual co-occurrence of first-order associations, what links “big” and “vast” is occurring in the same kind of sentences, and so with the same co-occurrence associations to other words. After all, for example, it is often a matter of indifference as to whether “big” or “vast” gets picked for use in any particular sentence context to indicate large size. Thus, a large number of words in such sentence contexts associate similarly with both these two words. If “big” appears frequently with another word, say, “stone”, so predictably will “vast”; if “big” does not with another word, then neither will “vast”. “Big” and “vast” as a result gain shared covariance associations with other words. These common associations create at a higher (or indirect) level, a strong link between “big” and “vast”. Such higher-order associations exist between “big” and “vast” in regard to every one of the tens of thousands of words used in vocabulary. Such high-order associations exist not just between “big” and “vast”, but also (with varying strength) between every other pair of words.

Higher-order associations: Higher and even more indirect levels of co-occurrence word association also exist. Some of the words with which “big” associate, co-associate with it more or less strongly in the presence of yet further words. Thus, if the word “big” appears in a sentence with “stone”, then it is likely that “heavy” will also appear. As a result, near innumerable associations of an even higher level exist. Together these higher-order associations in subtle and complex ways shape a web of information that permeates the entire vocabulary that gets used by a person and a language community. As in a spider’s web, where a pencil pushed against it distorts touched silk strains that in turn distort other threads communicating deformation throughout the entire web, the actual information of word usage that defines the context of a word in a sentence therefore arises from indirectly links that “communicate” from such distant associations.

2.5. LSA and higher-order association information

LSA is a method that extracts the word constraint information contained in those higher-order associations and it does so by analyzing tens of thousands of past episodes of word strings.

The extracted higher-order association information by LSA enables the making of meaning similarity and the distance judgments that match those provided by human subjects (Landauer & Dumais, 1997; Landauer et al., 1998). As a consequence, LSA provides strong support to suggest that higher-order associations play a key role in the brain’s generation of the context that underlies the meaning of words (as reflected in the brain’s making of similarity and dissimilarity judgments). This would be a very unlikely finding, if the higher-order associations extracted by LSA were a mere epiphenomenon to the brain’s processing of meaning.

The LSA computer model further tells us that most of the information needed to do this by the brain comes from higher-order associations not first-order ones (Landauer & Dumais, 1997, p.226). Theoretically, this might seem surprising since compared to first-order ones, most higher-order associations provide only relatively weak information about word usage. However, higher-order associations are by many orders of magnitude much more numerous than first-order ones, and so (when added up) in total contain much more information. Indeed, as Landauer and Dumais note, “About three quarters of LSA’s word knowledge [when tested] is the result of indirect induction, the effect of exposure to text not containing words in the tests” (1997, p. 226).

LSA investigates the information contained in the co-occurrence associations between elements (words) in episodes (strings of words). LSA does not ask where this information comes from: it merely extracts it and shows that it provides the basis of context processing. The source of that information, however, is almost certainly associations in the real world that arise when we describe its elements (words) in terms of episodes (as depicted, for example in sentences). This is not to suggest that these co-occurrence associations have an independent existence: episodes, after all, do not exist in the real world separate from human cognition as they arise from the brain's information processing need to process the world exterior to it in terms of episodic "chunks".

2.6. Semantic space

The raw information contained in higher-order associations is hidden and inaccessible to cognitive processes that might seek to use such information. Such information therefore needs to be transformed into a form that might be useful to the cognitive faculties that seek, for example, to identify unknown, hidden or missing entities from those that surround them (episodic inference). Thus, the "substitutability" of elements and so context needs to be made explicit. This is done in LSA by converting association information into a multidimensional vector "substitutability" space. In this space, synonyms occupy the same locality (they are little different in their intersubstitutability in any context), while the different meanings possessed by heteronyms (homographs and homophones) are widely separated (they produce different meanings when put in different contexts). More generally, the closer the meaning of words, the nearer they are in this space. In this multidimensional space, locations are not confined only to words: places can also be given to complete and incomplete sentences. A word can be synonymous, after all, not only with another word, but a group of words – for example, the definition – "being of extreme size" is synonymous with "vast".

This multidimensional space enables keywords to be linked to the documents in which they appear. By checking the location of keywords provided by a searcher, a program can find words with similar meaning (since they are found in similar contexts). Moreover, the different homograph meanings of the keywords such as "film" can be separated so that the locations of accompanying keywords ("Marilyn Monroe", "paint") can be used to select which of their meanings (movie, thin coating) the searcher intends.

It is noteworthy, as observed above, that humans relate to words in terms of them having a distance between each other in meaning space (see the above section upon the behavioral properties of meaning). This suggests that the brain could also treat words as existing in episodes, and convert associations between them into something like a context space, and that this could provide an information substrate upon which the neural processes of cognition might operate in everyday life.

2.7. Nonsemantic spaces

Things identified in nonlinguistic domains can also have cognitive attributes akin to "synonymy" and "heteronymy". A car, a donkey, and a sedan-chair may look very different, but they provide similar intersubstitutable means when on a holiday to get from one place to another. They – at least if our desire is holiday transport – can substitute for each other (much as the words, "vast", "big" or "being of extensive size" can in the same sentence substitute for each other). Here, however, what they share is not the same meaning but the same functionality as required by our needs and our cognition of those needs. "The tourist went in a car to the shops", or "The tourist went on a donkey to the shops", or "The tourist went in a sedan-chair to the shops" while different as physical activities are similar in terms of their being solutions to a need to travel. (And, reflecting this, the sentences describing them mean roughly the same). Likewise, in terms of function, an entity can be "heteronymous" in regard to different cognitive relevancies: a car in one context can be a means of transport, in another a place to escape rain, and in yet a further one, a status symbol.

Such episodic synonymy and heteronymy is important since our brains mostly seek similarities and dissimilarities between the constituent elements (such as events, processes and entities) in episodes, not in regard to their perceptual commonalties but in terms of their commonality in regard to our cognitive concerns and relevancies. Such relevancies may concern how we can use things, or how they might affect us; for example, that a particular means of transport in a given context is functionally the same for getting from A to B as another.

2.8. Locating the unknown from the known

Sentences and other groups of words, as noted above, can gain a location by mathematically adding together all the vectors defining the various locations of their individual words (Landauer & Dumais, 1997; Landauer et al., 1998). This provides a new vector location that is a kind of mathematical "center" of their individual vectors. LSA shows that once such a multidimensional space has been created, these new vectors can be used to identify unknown words from their word context. This is because when the individual locations of the surrounding known words are combined together, the new location produced is near in the multidimensional space of the location of the unknown word's meaning. Thus, the words surrounding an unknown word can be exploited to identity that word.

This method of exploiting contexts seems to be used by the human brain. Children learn on average 10 to 15 new word meanings each day, but only one of these words can be accounted for by direct instruction (Landauer & Dumais, 1997). The other nine to 14 word meanings need to be picked up in some other way. Landauer and Dumais (1997) have showed that when children meet an unfamiliar word, its context contains sufficient information to enable them to correctly guess its meaning. Since such word learning depends upon previously learnt words (via the contribution they make to creating the multidimensional space), the more words children learn, the easier it is for them to learn new ones. Landauer and Dumais (1997) argue this ability to infer unknown elements from the context of known ones in the high-dimensional space generated from higher-order associations plays a key role in language cognition and development.

2.8 Word meaning atypicallity

The findings about normal context cognition by LSA raises a question: what happens to word learning when the processes described by LSA are impaired in the developing brain, and children cannot use context either to learn words or comprehend them? Do individuals, for example, exist with abnormalities in how they understand words because their brains are unable to use context as explained by LSA? The clinical literature contains many conditions where retardation causes children to be delayed in learning the meaning of words, but actual atypicallity in what is learnt are rare. Interestingly, the one condition in which such atypicallity does occur is autism. This raises the possibility that LSA and its mathematics of word context might provide possible insights into their language and other cognitive atypicallities.

3. Atypical information processing and autism

The atypical cognitive information processing of autism (or ASD, autism spectrum disorders) is scientifically terra incognita. Francesca Happé has observed "the one thing that all those who study autism agree on is that no one really understands autism." (Happé, 2006, p. 633). Indeed, doubts have even been expressed as to whether a causal factor might exist (Gillberg & Coleman, 1992, p. 283; Happé, Ronald, & Plomin, 2006; Ozonoff, 1997, p.875), and the argument has even been made that autism might be a composite condition of unrelated cognitive defects (Waterhouse, Fein, & Modahl, 1996). However, it is possible that the present lack of understanding for autism could be due to a continuing incompleteness in the knowledge as to the full range of the information processes that underlie and enable everyday human cognition and its development. Completeness for cognitive science if it did exist would be unprecedented in the history of science -- all major areas of scientific investigation even after they have developed to a mature stage of apparently full development continue to discover major and unexpected additions to their topic phenomena (for example, the Kuiper pelt for astronomy in 1992; gases for chemistry in the late eighteenth century and the noble elements in 1894; plate tectonics for geology in the 1960s; and for biology, viruses in 1898 and archaea in 1977). It is thus possible that the failure to understand some aspects of the atypicalness of cognition in autism is due to a gap in the cognitive science inventory of the elementary information processing tasks carried out by the brain.

3.1. Possible context impairment

Within autism research, the existence of limitations in context processing is a considered to be a candidate possibility for its cognitive atypicallity. Indeed, one prominent etiology of autism, "the weak central coherence theory" (Frith, 1989; Happé, 1994; Happé & Frith, 2006), seeks to explain the heterogeneous aspects of autism in terms of a core defect in the information processing of a context-like process. For example, this theory describes as central to normal cognition, the 'built-in propensity to form coherence over as wide a range of stimuli as possible, and to generalize over as wide a range of contexts as possible' (Frith, 1989, p.100). This theory also claims that this process is central to many distinct types of cognition since it 'forces together complex information from totally disparate sources into a pattern which has *meaning*' (Frith, 1989, p. 174). The popularity of this theory has arisen in part because individuals with autism have been experimentally shown to have limited abilities to exploit and use context appropriately (Beversdorf et al., 2000; Brosnan, Scott, Fox, & Pye, 2004; Frith & Snowling, 1983; Happé, 1997; Jolliffe & Baron-Cohen, 1999, 2000); and also, in part, because, at an intuitive level, various features of autism, seem to relate to the dependence and independence of cognitions upon context (Frith, 1989; Happé, 1994). Context is also employed in the skill (that is impaired in autism) of generalizing from structured settings to naturalistic ones which contain unpredictable and context-dependent interpretation factors (Volkmar, Lord, Bailey, Schultz, & Klin, 2004a). A difference in experiencing things in context was also implied in the first description of autism by Kanner (1943, p. 246), for example: "A situation, a performance, a sentence is not regarded as complete if it is not made up of exactly the same elements that were present at the time the child was first confronted with it. If the slightest ingredient is altered or removed, the total situation is no longer the same and therefore is not accepted as such .."

Unfortunately, further progress in understanding the role of context as a possible factor in the atypicallity of autistic cognition is limited due to the intuitive nature of research definitions about "context". At present, no general theory of context has been detailed (at least as a theory of context) by those working upon autism reflecting the wider failure in cognitive science that shows how it might be mathematically reduced to more basic kinds of information, particularly that which might be developmentally extracted from everyday experience. Context in the absence of such analysis has come to be used to refer rather imprecisely to nearly any global, holistic, or 'higher level' aspect of a situation, circumstance or spatial configuration.

The problems caused by the lack of a cognitive science model of context can be seen in the attempts to test the central coherence theory of autism. Its proposed role for context has been analyzed with tasks measuring low-level analytic processes compared to high-level global ones (Jarrold & Russell, 1997; Mottron, Burack, Stauder, & Robaey, 1999). However, since the central coherence theory describes context in an intuitive way, no objective way exists of knowing if such tests of low-level analytic and high-level global processes do in fact test its suggested disruption to context processing. Another approach (Lopez & Leekam, 2003) has been to use a semantically related visual context priming task (Palmer, 1975), but the role of context in this task is intuitive not computational.

3.2. Methodology limitations and etiological analysis

Etiological theories of biomedical conditions can be proposed into the mechanisms of defective processes empirically (bottom-up), or theoretically (top-down) from an understanding of normal processes and the likely results that arise after they are impaired.

Top-down approaches could be particularly critical to furthering our understanding of the atypicallity of autism due to the unfortunately limited etiological import of much present empirical research. As Volkmar and colleagues in a recent major review have noted, “Typical research designs involve a simple comparison between a group with autism and a control on the measure of interest. Significant results in such comparisons have generated sweeping theories of autism in the past, regardless of effect sizes obtained or the extent to which the variable of interest could be directly related to social adaptive functioning”, and “small effect sizes are unlikely to account for the very large gaps in social abilities and disabilities separating individuals with autism from their peers” (2004b, p. 143). Consider also the limitations of the standard autistic-group matched control-group type experiment. Usually a few individuals in the control group exist that perform at the same level as the autistic group but the existence of these nonautistic “outliners” is statistically removed when the means of the two groups are compared. However, the existence of these nonautistic outliners can, and often is, of potential interpretational importance. For example, recent estimates of the prevalence of autism range from 34 to 60 per 10,000 (Fombonne, 2003). Taking the higher estimate, a population will need to be the size of 5,000 to contain 30 individuals with autism. Thus, a research group of 30 individuals with autism will match in terms of sampled *prevalence* with one of nearly 5000 nonautistic individuals. If one individual (and usually there are more) exists in a typical control group of 30 implicitly sampled from this 5000 population, then by arithmetic the vast majority individuals that perform like those in the autistic group will not in fact be labeled as autistic. Indeed that for every individual that has a autistic-like performance that in the control group, there will be for those 30 individuals with autism, 200 individuals in the prevalence population from which the control is implicitly sampled.

Therefore few of the "standard" findings in autism research so far sharply identify what is specific to the atypicalism of autism rather than what is predominant but etiologically nonspecific (see also (Mervis & Robinson, 1999, p. 123)) leaving the actual nature of the atypically of autism very unclear. This suggests a continuing need to draw in a top-down manner from developments in the cognitive science of normal cognition potential sources of atypical cognition that might potentially characterize its nature. Such approaches (as are all approaches to autism) candidates that subject to further refinement and articulation. However, they might provide beginnings for possible directions for research that might allow better understanding as to its actual nature.

3.3. Word meaning atypicallity in autism

From Leo Kanner (Kanner, 1943, 1946) to the *Diagnostic and Statistical Manual of Mental Disorders. Fourth edition*, (American Psychiatric Association, 1994), language and meaning atypicalness has been central to the diagnosis of autism. Their nature includes mutism, echolalia, pedanticism, and atypicallity in the understanding and use of word meaning. The latter will be of concern here. Kanner (1946, p. 243) noted that, “the autistic child has his own private, original, individualized references, the semantics of which are transferable only to the extent to which any listener can, through his own efforts, trace the source”, and that their word meanings are “rooted in concrete, specific, personal experiences of the child who used them” (1946, p. 243). Similar comments have been made more recently by Uta Frith and Francesca Happé (Frith & Happé, 1994). They note that children with autism, “may use single words in a simple, associative way, so that “Apple” always means, “Give me apple”. The single words acquired are often esoteric (e.g. “Beethoven”) and not like the first words of a normally-developing pre-schooler. Neologisms (e.g. “bawcet” for bossy), or familiar words with special meanings “yes” meaning “carry me on your shoulders”), also reflect the very concrete context of word and object”. These comments suggest that individuals with autism when they first hear unfamiliar words might be attempting to understand what they mean using immediate (and often misleading) associations in the physical world. This, indeed, was described by Kanner (1946, p. 242) in regard to a child called Paul G who said "Peter eater" when ever he saw anything resembling a saucepan. According to his mother, when Paul was two years old, while busy in the kitchen, she was reciting to him the nursery rhyme about "Peter, Peter, pumpkin eater", when she dropped a saucepan. Ever since Paul has understood "Peter eater" atypically to mean "saucepan". He had made an association between "Peter eater" and the most salient event at the time of its been said – the dropped saucepan.

3.4. Context impairments

This reliance upon direct co-occurrence associations would suggest a particular kind of failure in those with autism in word learning. First, that they fail to employ the context-based cognition processes by which nonautistic individuals acquire word meaning. Second, that in its absence, they rely upon the accidental nature of word and event co-occurrences to guess what words might mean. Third, that when they have done this, they are impaired in using context to appreciate the inappropriateness with which they have understood a word. Such an impairment has two results: (1) it delays them learning words, since it removes the main source of information by which normal individuals acquire word meaning. This could be the reason why some individuals with autism are mute or echo words and phrases – a word usage that does not require that their meaning is understood. (2) That when they do acquire language, this could explain why they tend to pick up bizarre and restricted ideas about the meaning of words: without an ability to use context to judge what an unknown word might mean, they are limited when doing this to its direct associations with events and situations.

3.5. Homographs

In support of a deficiency in the processing of context as modeled by LSA is their apparent miscomprehension of homographs. Four sets of experiments (Frith & Snowling, 1983; Happé, 1997; Jolliffe & Baron-Cohen, 1999, experiment 1; Lopez & Leekam, 2003, experiment 4) (but see experiment 2 in Hoy, Hatton, & Hare, 2004) upon those with autism that can read have found suggestive evidence that they ignore sentence context when asked to pronounce a homograph. Unlike those without autism, they tend to pronounce the most common occurrence of homographs, rather than the one that fits its sentence context. Thus, they will tend to pronounce the word "tear" in the sentence, "The lady had a TEAR on her dress" in the sense of a cry droplet rather than of a rip or cut. This suggests that they might not access the context used by nonautistic people in word comprehension. LSA has shown that the context needed to disambiguate the different meanings of homographs is extracted from higher-order associations. This raises the possibility that the processes needed to extract higher-order associations and consolidate their information into comprehension as modeled by LSA are potentially atypical in someway in autism.

3.6. Coherence

LSA has also shown that such context information is used by people to judge the degree of coherence present in texts (Foltz, in press; Foltz, Kintsch, & Landauer, 1998). Thus if LSA-like processes were damaged in autism, the perception and use of such coherence would be expected to be atypical. Evidence suggests that this is the case: individuals with autism tend to fail to comprehend the normal coherence that allows inferences between sentences (Jolliffe & Baron-Cohen, 1999; 2000, experiment 2). In further support of this, DSM-IV and others (C. M. Hale & Tager-Flusberg, 2005; Kanner, 1943) note that those with autism seem unable to create continuity between sentences as they fail to continue the topic of conversations, and instead, regardless of meaning, are repetitive of words or phrases.

3.7. LSA and autism

The language impairments of those with autism are consistent with an impaired ability to extract and use the higher-order associations modeled by LSA. Interestingly, Landauer and Dumais note, that the information processing principles behind LSA are not domain specific.

> There is no reason why much more complex structures, with mental (or neural) events at varying temporal scales and various degrees of repeatability could not exploit the dimensionally-matching mechanism to produce similarities and generation among and between psychological entities of many kinds (1997, p. 228).

There is thus no reason why these "psychological entities of many kinds" should not occur in across the diverse variety of cognitive faculties that are impaired in autism.

That other cognitions use similar processes to those described by LSA for language is not implausible as the problem of identification from context is not confined to words (as noted above in section, 2.7). A cognitive problem faced by many faculties is identifying something – a hidden object, an unfamiliar response, or an unidentified aspect of a situation – given the information provided by surrounding known events and entities. As discussed later, this problem is encountered particularly by the faculties that provide us with our sense of security, social interaction and understanding of mental states. Such faculties develop using the information contained in the hundreds, if not thousands, of episodes encountered in daily experience across the many years in which cognition matures. Thus, such faculties have available the higher-order association information contained across such episodes to generate a multidimensional space that can be used to identify unknowns from their surrounding context. This raises the possibility that the extraction of higher-order associations and their use to provide context might be widely exploited by many faculties in development and function. From this it follows that impairments to such processes might possibly produce context processing impairments in these faculties akin to the atypicallity of those with autism.

This suggestion is obviously preliminary given the general *terra incognita* of the role of context in nonlinguistic domains, such as security perception and sociability, and our still limited understanding of the nature of the atypically in these areas of those individuals with such atypical cognition. However, though preliminary given the limits of knowledge about atypicallity in these areas, it might provide insights worth investigating and discussing as their atypicalness.

One question is whether given the limits of present knowledge, there is a general consistency of this approach with the facts as far as they are known in regard to the atypical cognition syndrome as it is found in autism. Such consistency is required in two areas: first, in the general etiology profile of autism as a atypical condition, and second, in regard to autism's specific core symptoms. The former is discussed in section 4, the latter in section 5.

4. Atypical context-based cognition and autism

The etiology underlying autism is characterized by four overall characteristics: (i) it effects diverse cognitions, (ii) it is specific to development with no close parallel in acquired psychopathology, (iii) it is heterogeneous, and (iv) it preserves isolated nonretarded or even supercompetent abilities. Any provisional theory of the atypicallity of cognition found in autism must be consistent with these four characteristics.

4.1 Development and specialized modularized processes

Certain factors can be reasonably inferred about the processes that extract higher-order associations and consolidate them as context into cognitions, notably that they will in many cases involves specialized neural operations that are separate from the cognitions that they enhance with context information. This is because these computations must process all of the potential informational links that exist between entities and do so across the tens of thousands of episodes previously experienced in a domain. Such computations need also to be specialized to calculate the "percolation" of information since as Landauer and Dumais (1997) note when a new element “is encountered, the distance between its representation and that of every other stimulus that occurs in close proximity to it is adjusted ... the adjustment is then allowed to percolate through the whole previously constructed structure of relations” (1997, p. 217). While this in theory could be done by a single network (Gorrell & Webb, 2005), the information that links entities making up episodes is highly distributed across cerebral cortex areas making this unlikely. For example, the “noun” and “verb” qualities of entities (Vigliocco et al., 2006), and tools and animals (Martin, Wiggs, Ungerleider, & Haxby, 1996) are processed in different parts of the brain. This suggests the need for a process that is specialized in terms of its physical ability to gather higher-order information across the various entities processed in different areas of the cerebral cortex, as well as its ability to integrate the associations it finds hidden between them as a totality.

The physical separation of neurocomputational abilities (connected by white fiber tracts) is a prominent feature of the brain suggesting that evolution has selected that cognition should depend for its effective processing upon the splitting of information processing among different specialized neural networks. This consolidation/transfer process is known to occur for the cerebellum (Maquert, Schwartz, Passingham, & Frith, 2003), hippocampus (Wiltgen, Brown, Talton, & Silva, 2004) and basal ganglia (Weber, Wermter, & Elshaw, 2006) all in regard to consolidation in the cerebral cortex, and the hippocampus in regard to the basal ganglia (Orban et al., 2006). Such separation could provide the advantage that it allows specialized processes to gather and analyze information from across the brain to be processed by specialized neural networks that can then be consolidated in the cerebral cortex to extend its cognitive functionality. The computing requirements raised by LSA in the extraction of context make it a candidate for one of these cross-brain consolidation processes.

The consolidation of context if this was the case would create the situation whereby such processing could be developmentally impaired in isolation from the upstream cognitions that use its extracted context – even though "upstream" of such impairment they remain themselves in neurophysiological terms otherwise “intact”. One situation in which this might happen is if internal brain communication is disrupted, or the upstream cognitive faculties are unable to incorporate transferred information as might arise if there was developmental failure of one part of the brain but not the others with which it consolidated processed information. Deprived of context information in these ways, this would force upstream cognitive faculties to use only information available locally that did not need specialized extraction and consolidation . One of the characteristics of autism is that cognition is not necessarily retarded (at least in terms of IQ measured intelligence) with different cognitions retaining functioning albeit in an atypical way suggesting that they are intact. Moreover, such functioning, as noted, is suggestive of a context related defect with unimpaired processing of immediate associations, and so consistent with this scenario.

Another factor for their modularization is that many faculties draw upon information contained in the same past episodic experiences, though this happens in regard to different concerns. For example, as discussed, below our ability to sense security and danger in appropriate circumstances derives from our past experience of safe and dangerous situations. But the information that describes these situations also overlaps with that which provides the past episodic experiences from which context is drawn to enable language and social interactions. Thus, the processes extracting higher-order associations for different functions will draw upon the same stock of episodic experiences. Such extraction, since it concerns much the same higher-order associations hidden in the same episodes, is therefore unlikely to be done separately for each faculty, but rather as a whole by one process, that then passes the relevant extracted context information to different cognitions for their varying purposes. In consequence, a single impairment to such processes could effect, as found in autism, a wide range of cognitions.

4.2. Development and once only information capture

The extraction and the use of information can happen at different times. Search engines such as Google web crawl to create an entirely new index every few weeks, though the resulting index is searchable within a fraction of a second. Likewise once higher-order associations are extracted and converted into a multidimensional vector space, this process does not necessarily need to be repeated as the information extracted will be in most cases usable for new entities. Neurobiologically, the extraction of higher-order context is likely to be continuously updated from the constant stream of episodes that arise in everyday life. But, if for any reason this process is interrupted or stopped, the brain will not lose the context information that it has already acquired, employed and consolidated into mature cognition. As a result, any defect later in life in the processes extracting higher-order context need not necessarily produce noticeable cognitive defects since the cognition of such individuals would still possess functioning context information from the period when such processes had been intact. Impairment will thus be critical only for young individuals before they have acquired a reserve of such extracted context information. As a result, this would limit – as in autism – the effects of impairments upon such processes to when they happen early but not to when they are impaired later in life.

4.3. Heterogeneous development

The integrity of the various stages involved in the processing of higher-order association information, and its consolidation into upstream cognitive faculties can be affected in diverse ways due to the number of stages involved. Moreover, upon any such developmental sources of variability, further ones and be imposed related to the different environmental and educational support that ameliorates or exaggerates them. In addition, the neurobiological factors that cause impairment upon neurocognitive processes might independently of this also have their own separate effects that impair cognitive integrity (such as causing general mental retardation and epilepsy). There is thus no reason to assume (except in broad features) that the behavioral consequences of an impairment to context extraction processes will be in detail necessarily alike in different individuals –or even alike at different times in the same individual.

4.4. Isolated abilities and residual first-order association based cognition

The contribution of context to the normal development and function cognition varies widely. While some competences are critically dependent, others can function with only first-order associations. Thus, following an impairment to context-based cognition, many faculties that can function using mainly first-order co-occurrence associations will be intact or only mildly impaired. In the absence of higher-order associations, however, some faculties that would have otherwise employed context information might be only able to compensate -- atypically by substituting (or having greater dependence upon) first order co-occurrence associations. This might result in them developing in different ways – and even in rare cases acquiring greater competence – compared to that present in neurotypical individuals. One of the noted peculiarities of autism is the preservation (although atypical developed) of cognitions not dependent on context (see section 5.4.).

5. Atypical cognition

If, autism is the result of atypical context-dependent everyday cognition (and so the processing of higher-order associations) then parallels can be predicted to exist between (1) the dependence of faculties upon context derived from higher-order associations and the areas of autistic atypicallity; and (2) between the independence of faculties from needing such information and areas of autistic normality.

Two main traits were seen by Leo Kanner (1943) to be present in individuals with autism: an “anxiously obsessive desire for the preservation of sameness” (p. 245), and an “extreme autistic aloneness” (p. 242). The American Psychiatric Association’s DSM-IV criteria for Autistic Disorder (1994) likewise gives its two essential features as “a markedly restricted repertoire of activity and interests” (see section 5.1. below in this paper), and “the presence of markedly abnormal or impaired development in social interaction and communication” (section 5.2.). These together with several further aspects of autism: impairments in theory-of-mind (section 5.3.), the preservation of islets of ability and supernormal skills (section 5.4.) are discussed below. (Brief comments are also made upon the possible role of impaired context processing in sensory gating and motor control in section 5.5).

5.1. Harm perception atypicalness

One of our instinctive concerns is to avoid what might harm us. Harm avoidance depends upon the cognitive capacity to anticipate and then activate specially evolved neural circuits in our limbic system, and our hypothalamic-pituitary-adrenal (HPA) axis so that the physiological reactions of fright, fear and ”flee or fight” that protect us from harm are induced (Panksepp, 1990). The normality of these harm protection reactions therefore critically depends upon the integrity of the cognitive processes that detect threats of danger. If we falsely expect harm, we may panic even when we are safe (and so waste our energies or unnecessarily limit our actions), or if we fail to foresee harm, we may not panic when this might enable us to avoid later suffering and injury. Thus, experiences such as panic and anxiety are intimately linked to how successfully the brain’s information processing can predict and identify actual threats. This makes the efficacy of such cognitive processes central to our everyday ability to live without suffering constant and undue distress.

5.1.1. Context and harm awareness

Information used to anticipate danger can depend upon first-order cues or context. Anticipatory information provided by first-order clues concern such things as the appearance of a threat (the shape of a snake, say) or the unfamiliarity of a situation. However these danger cues are not always trustworthy. For example, the shape of a snake – a first-order perceptual cue – can warn us of a dangerous animal, but many entities from a distance look like snakes (large worms, water hoses, plastic snake imitations), and many snakes are not poisonous, (and poisonous ones are not dangerous if dead or in glass cages). Familiarity is also unreliable: many familiar places and things can also be dangerous – for example, a house on a beach does not change in appearance even though a forecast hurricane requires its immediate evacuation. Further, strange places can be safer than familiar ones (the sports center into which people seek refuge from an advancing hurricane). The brain thus is advantaged if it can use context to extend and refine its sense of security so it minimizes false positives (anticipations that are false), and false negatives (failure to anticipate dangers that are real). Unimpaired and experienced cognition has the information to do this since it can incorporate the information contained in the many tens of thousands of varied and diverse situations that have happened either to ourselves, or which we have witnessed happen to others, in which harmful or nonharmful events occur or not. Thus, cognition has the opportunity to extract higher-order associations for the contextual anticipation of danger by means of a multidimensional space in which situations can be found to be similar or dissimilar in regard to their risk and security. Such a multidimensional space, as with words, would allow cognitive processes to use the context of the known to identify the unknown (in this case whether a situation is likely to be safe or dangerous).

Our limbic and our hypothalamic-pituitary-adrenal (HPA) axis can use such information to supplement unfamiliarity and other unreliable co-occurrence cues. Thus, as a result of this, the brain can feel safe even when unfamiliarity and first-order cues suggest it should feel danger (for example, when it first travels in an aircraft), provided that context reassures us that there is no threat. The brain, however, will still remain sensitive to dangers if this context changes: for example, if the warning lights above the passenger seats light up and the captain announces that we must prepare for an emergency landing. These clues while they are not in themselves unfamiliar events (at other occasions the lights come on and the captain speaks), however, they create a new context (we do not think the captain is fooling or mistaken) that warns that we face actual risk.

This analysis of the role of context to anticipate risk and safety raises the question of what happens when it is impaired? In this situation, the limbic system and the hypothalamic-pituitary-adrenal (HPA) axis will have no option but to be limited in its threat activation to those provided by first-order perceptual clues. As a result, individuals with an impairment in extracting context, will be hypersensitive to changes in what is recognizable and well-known. Such individuals exist: those with autism are notable for panic and distress over small and irrelevant alterations in their routines and circumstances. As the DSM-IV (American Psychiatric Association, 1994) criteria explain, “they may insist on sameness and show resistance to or distress over trivial changes (e.g. a younger child may have a catastrophic reaction to a minor change in the environment such as a new set of curtains or a change in place at the dinner table)” (p. 69).

5.2. Atypical sociability

5.2.1. Sociability and context

Humans are a social primate. Due to this, the human brain generates emotions that bind us with our kin, friends, and other individuals (Panksepp, 1998). This interpersonal involvement with each other is moreover highly contextual (Clark, 1992; Halberstadt, Denham, & Dunsmore, 2001; Lewin, 1951; Rohlfing et al., 2003; Todorov et al., 2006). For example, when we meet a friend, how we interact will happen within various contexts such as our past relationship (there might be past obligations to repay); our and their relationship to others (they might know a mutual friend whose impressions we value); what they feel (if our friend is sad, we try and cheer them up); social recognition (our friend wants to be valued as a special person), and social mores, rules, manners, and morals (there are contextually defined boundaries upon behavior). All these concerns enrich social interaction with multiple levels of context; nothing analogous occurs in our relationship with physical objects.

Social psychology, moreover, in its attempts to understand our social existence gives a central position to phenomena such as ‘self’, social roles, sense of personal identity, self-presentation, pragmatics, and awareness of in-group and out-group boundaries (Cicourel, 1987; Ellemers, Spears, & Doosje, 2002; Lewin, 1951; Prentice & Banaji, 1994; Rose-Krasnor, 1997; Todorov et al., 2006; Turner, Oakes, Haslam, & McGarty, 1994). These depend for their functioning within social psychology upon the ability of people to understand or experience social events and entities as happening within an interpretative context. Unfortunately, social psychologists have not provided an information processing framework in which to understand at a basic information processing level such phenomena. As a result, processes such as “self”, “personal identity” and “social group perception”, from a computational perspective, are nebulous lacking an underlying cognitive mechanisms that could support their generation or perception by the brain. Obviously, due to space limits such an account cannot be provided here that grounds social psychology into more basic information processing operations. However, it is striking that social psychologists describe such phenomena as being carried out by processes that are interpretative, sensitive to context and concerned with what is loosely referred to as ‘meaning’.

Another link is that the anticipation of unknown elements from known ones in a social context is central to when we interact. Take the example of a greeting: though it has some usual first-order co-occurrence elements such as saying, “hello”, smiling, and hand shakes, these depend upon an interpreted context. A greeting done indifferently can cause offence. One done in another way, might cause us to ask what is troubling a friend. The elements, moreover, involved might be much the same, but small subtleties can change radically their social meaning. And people can improvise. For example, an individual carrying shopping bags sees a neighbor across the road. That individual cannot make the usual first-order signs of greeting such as, “hello” (they cannot hear them above the traffic), and they cannot wave (their hands are tied). They, however, can and will adapt and make a “social synonym”. They might, for example make an exaggerated smile while mouthing words of greeting, or they might slightly raise one of their shopping bags, and wave it and their body. Stopped from making a normal greeting, they adaptively generate social behavior which they know others due to its context can interpret as a greeting rather than annoyance or merely an odd action. An individual is not concerned with creating a particular “Hello” behavior, but concerned with creating behavior that given its context, will be understood as having that social outcome.

Sociability, moreover, builds upon context like the sentences of language by building upon earlier ones and contextually inter-chaining them. In a text, this is reflected in the fact that when a sentence follows another, it expands and develops the meaning of the earlier sentence. Because of this, when the sentences of a text are scrambled up, they do not make sense (their meaning depends upon the sequential flow of context). Likewise, sociability involves an interdependent flow of context: imagine the episodes of a social situation mixed up; they would cease to be recognizable as a social interaction. One action at one moment means one thing, at another moment, something very different. Due to this context dependence, if we miss one element in a social interaction, we can guess it from its surrounding ones, much as in a cloze sentence, an omitted word can be guessed from the context provided by surrounding ones. While physical entities and our relation with them to have some aspects of context, they lack this rich embeddedness in context that shapes social behavior and social interpretation.

Every human brain encounters from its earliest years tens of thousands of episodes of social interaction. Thus, it has considerable opportunity to extract social context information from the higher-order associations hidden in them and generate a multidimensional "social intercourse" space. Such a multidimensional space contains context information which can be used to generate social phenomena such as ‘self’, social roles and personal identity.

Given that such information is available (people are constantly engage in social interactions), that such information is known to be used by one faculty already (language), and that there exists a need for such a information processing system (we live in communities, and in our affects, we are social primates), parsimony suggests that we should assume (unless we have good reason to think otherwise), that this information processing system has an important role in enabling and enhancing sociability. If we do not, we would have to explain why we have a skill (sociability), (1) that has an information processing similarity to that used in another domain (language), (2) that has the information needed to be processed like it, and yet (3) uses some different and distinct process. Ockham in his razor requires that we do not to multiply entities beyond necessity. Thus, even though we lack precise models, it is reasonable to suggest, at least provisionally, that the cognitive processes under sociability depend upon LSA-like functions.

5.2.2. Social emotions and context

Social psychology tells us that emotions link deeply to these above discussed context aspects of sociability and self (Halberstadt et al., 2001). Again, limits of space allows only the most preliminary comments. However, one needs to note minimally that the interaction between emotions and social context are likely to be important. For example, consider this description of our need for social acknowledgement by William James:

> We have an innate propensity to get ourselves noticed, and noticed favorably, by our kind. No more fiendish punishment could be devised, were such a thing physically possible, than that one should be turned loose in society and remain absolutely unnoticed by all the members thereof. If no one turned round when we entered, answered when we spoke or minded what we did, but if every person we met "cut us dead," and acted as if we were non-existing things, a kind of rage and impotent despair would ere long well up in us. (1891, pp. 292-293).

People not only exist in a social context, but as James suggests here that our emotions make us concerned that others experience us as part of the context of their social world.

5.2.3. Atypical social interests

These comments – preliminary though they are – suggest that without the ability to extract higher-order associations from social interactions that sociability and a sense of self would not arise. First, individuals would not be able to acquire the context sensitivity needed to engage in social behaviors such as greeting and responding to greetings. Second, individuals would not be able to interchain our social interactions so that each social interaction followed on appropriately from earlier ones. Third, since roles, self and a sense of social presence derive from the multidimensional "social" space generated from such interactions, together with emotions (we care much to be noticed), individuals will not be motivated to interact. As a result, people lacking the ability to process higher-order association will still have emotions, but such emotions will enable them to relate only through perceptual-based types of cognition: their affect will link to specific behaviors, and resemble their behavior to objects. They will, moreover, lack the ability to experience the many emotions that arise out of context such as the "rage and impotent despair" that James mentions above. Indeed, rather than such social engagement emotions, they will be content if they are socially ignored. They will show poor or inappropriately developed capacity for social roles, personal identity, self-presentation, and social boundaries. In consequence, such individuals will seem to others socially remote, passive, odd and lacking in friends. This description, of course, characterizes the core deficits of the social impairment of those with autism. As put by the DSM-IV, in the autistic disorder:

> Younger individuals may have little or no interest in establishing friendships. Older individuals may have an interest in friendship but lack understanding of the conventions of social interaction. There may be a lack of spontaneous seeking to share enjoyment. interests, or achievements with other people (e.g. not showing, bringing or pointing out objects they find interesting) (American Psychiatric Association, 1994, p.68).

Further, that "Lack of social or emotional reciprocity may be present (e.g.: not actively participating in simple social play or games, preferring solitary activities, or involving others only as tools or mechanical aids)" (American Psychiatric Association, 1994, p.68).

5.3. Theory-of-mind atypicallity

One particularly important context is that carried within each person in regard to their goals, desires, beliefs and knowledge. In spite of these mental processes being hidden, individuals appreciate that they strongly shape the behavior of other individuals. For example, we can predict if a child believes a toy is in a box, that they will act in one way, and that if they believe it has been moved, that they will act in a different one. Skill in guessing how mental states shape behavior is referred to as theory-of-mind.

Mental states, however, cannot be inferred from first-order co-occurrence associations of the episodes in which they occur since they are not only hidden but rarely correlate to the immediate aspects of a situation. In information processing terms, mental state inference involves inferring the presence of unknown elements from a context of known ones – a cognitive task parallel to episodic inference one of the identification of the meaning of unknown words from the known ones that surround them. Here, however, the unknown element guessed from observed behavior and its situation is an aspect of mental activity.

Such a skill in mind reading will depend upon an ability of the brain to create a multidimensional space in which to locate the observed elements of behavior that can be used to infer such hidden mental states. The higher-order associations needed to create this space exist since all episodes of behavior arise from people having varied intents, beliefs and desires. As a result, over many years, the brain encounters tens of thousands of situations molded by mental phenomena. The higher-order associations contained in them allow for the existence of a multidimensional "mind" space in which episodes in which people do things can be redescribed in terms of contexts and hidden intentional states. From this space, the known elements of a behavioral situation can provide contextual information about the unperceivable ones of intents, beliefs and desires. (Though, as noted below, another source of information in creating this space are mentalistic language terms for intentional and mental states: it is conceivable that both sources of information have a role in developing theory-of-mind competence).

Due to the ability to guess unknown elements from the context of known ones (where these depend upon higher-order associations), we would expect impairments in extracting such co-occurrence associations to impair the cognitive ability to infer mental states. This fits in with the atypical deficits of autism. Awareness that others have desires, beliefs and intents (theory-of-mind) is impaired in those with autism (Baron, Leslie, & Frith. 1985). Some individuals with autism admittedly pass tasks that test theory-of-mind (Happé 1995). But they are different: ordinary people cannot explain in detail the answers they give to such tasks (much as they cannot when they use context for identifying homographs that fit sentence context) (Happé 1995), in contrast, those with autism that succeed on these tasks provide articulate justifications. This suggests that they have tackled them as puzzles, and extrapolate their answers from superficially similar examples. Such “example-based” puzzle solving is of limited use in replacing those based upon context: “mind-aware” people with autism fail harder tasks that test more difficult forms of mind awareness, moreover, they fail to use such skills in real life (Happé, 1995).

5.4. Islets of preserved cognition

A problem may require either the use of first-order co-occurrence associations or the context provided by extracted higher-order associations. But which kind of information processing is needed may not be obvious from the initial perception of a task: habits and superficial details could therefore cause a problem to be tackled inappropriately making it harder and more difficult to solve. The block design (Shah & Frith, 1993), and the embedded figures task (Shah & Frith, 1983), for example, embed a problem in information that invites a contextual approach but which in fact is irrelevant and distracting to its solution. Another instance is the learning of lists of meaningless words – since words are habitually heard as parts of sentence contexts, an individual will tend to spontaneously treat them as having potential meaning as a group – even though the memorizing of them is easier if they are treated as lists of unconnected elements. In consequence, such tasks will be more straightforwardly done if a deficit exists that suppresses the extraction or consolidation of higher-order information when faced with a problem. (Assuming, of course, that first-order information processing is intact.) If those with autism fail to process context information, they should solve such tasks better than people without autism. This indeed is the case. Kanner (1943) noted that people with autism could memorize lists of items better than ordinary people (see also Beversdorf et al., 2000; Tager-Flusberg, 1991), and people with autism have been reported to be superior to normal people upon tasks such as the block design (Shah & Frith, 1993), and the embedded ones (Shah & Frith, 1983).

As noted, in the absence of higher-order associations, the development of diverse cognitions might be maladaptive since they might use in the place of context the noncontextual information provided by available first-order co-occurrence associations. This abnormal development need not always be inferior to that normally acquired. It is possible that in the absence of context that cognition might become highly sophisticated in doing certain tasks (albeit with limited usefulness). This even might produce in some individuals novel cognitive skills that show superiorities. Consistent with this, a few individuals with autism show unusual abilities in such activities as drawing, music, calculation and memory. Their exceptionality, as this theory requires, seems to derive from replacing a sensitivity to context with elaborate and expert use of noncontextual information (Mottron & Belleville, 1993).

5.5. Sensory and motor processing

Some apparently noncontextual aspects of autism might hide the processing of context. The following comments attempt to show that this might be the case for sensory gating and motor related problems. The reader should note these are preliminary observations: while there is suggestive evidence context plays a key role in perception and motor control, our understanding of the basic processes is still limited.

Many hundreds of sensory and biomechanical interactions exist (by the nature of our articulate bodies being possessed of diverse and separated sensory organs) that bind together what the brain experiences, and governs what the body can do. These would appear in many cases to require for their efficient processing sensory gating, anticipatory motor adjustment and other modulation that depends upon the detection of an appropriate context. Contractions of agonist muscles, for example, must be initiated contextually to occur with time appropriate relaxation, not just of antagonist muscles but also muscles in parts of the muscular that are distant, (and then only in certain movements contexts) (Crone & Nielsen, 1989; Karst & Hasan, 1987). Many movements substitute for each other depending upon context (for example, underhand and overhand throwing). Sensation is actively gated in a similar way dependent upon context: for example, input from our joint and skin receptors is temporally "damped" so we do not to "feel" our own movements (Asanuma et al., 2003), and hear our own words when we speak (Houde, Nagarajan, Sekihara, & Merzenich, 2002). Touch input also is contextually heightened, (for example, when our finger tips contact an object) (Gardner, Ro, Debowy, & Ghosh, 1999). Theories of schizophrenia (Grace, 2000), further argue, that context gating underlies the normal control of behavior.

It is not unreasonable to suggest that such context gating and adjustment depends upon the extraction of the higher-order associations gained from the analysis of tens of thousands of past episodes of movement and sensation in which motor outputs and sensory inputs interact together. That such information is crucial to effective sensory-motor function could explain why motor skills take so long to perfect, for example, the ability to grab objects continue to develop in sophistication into the first decade of life (Kuhtz-Buschbeck, Stolze, John, Boczek-Funcke, & Illert, 1998). Due to this dependence of motor and sensory gating upon context, any deficits in extracting the context interactions between motor outputs, sensory input and the surrounding sensory-motor circumstances could produce impaired sensory gating and motor abnormalities. Whether this explains the motor problems (Baranek, 1999; Nayate, Bradshaw, & Rinehart, 2005; Teitelbaum, Teitelbaum, Nye, Fryman, & Maurer, 1998) and sensory disturbances (Harrison & Hare, 2004) faced by those with autism is unknown: we have, after all, only limited understanding of the role of context sensitivity in normal sensory-motor functioning. However, it suggests the possibility that general impairments in context processing could underlie aspects of autism that upon first impressions do not seem be context dependent. In support of this, Happé and Frith (Happé & Frith, 2006, p. 6) have observed that "perceptual abnormalities such as hypersensitivity, clinically/anecdotally reported but little studied in research to date, may relate to context-free processing as expectations and context-based interpretation are known to modulate experience of sensory stimuli",

6. Discussion

The LSA model provides a tested computer model of the normal context cognition involved in comprehending the meaning of words, though so far the consequences of its impairment are unexploited. "Lesioned" or compromised computer models of normal processes, however, have been used elsewhere to understand cognitive impairments, such as dyslexia, Williams syndrome, brain injury, specific language impairment use (Brown, 1997; Joanisse & Seidenberg, 1999, 2003; Thomas & Karmiloff-Smith, 2003). If the general hypothesis of this paper is correct, then LSA computer simulations of normal word learning and comprehension when lesioned should produce both known and some unknown meaning related atypical use and understanding of words. These upon further investigation should be found in those with autism and other forms of atypical cognition.

The reverse might also apply: atypicalness in autism might suggest areas in which the LSA could model aspects of language which it has not yet investigated. For example, the multidimensional space that LSA extracts represents, in a manner, conceptual knowledge – an aspect that has resulted in LSA models being successfully used to automatically mark the content of student essays in regard to their knowledge of psychology textbooks (Landauer et al., 1998). The acquisition of theory-of-mind skills is known to link to language (Astington & Jenkins, 1999), and the exposure to mental state words (Peterson & Siegal, 1999). This suggests it is likely to be extracted by LSA from the usage of mentalistic words. Theory-of-mind skills have not as yet been modeled by computer simulation, nor has the LSA model been tested as to whether it can simulate them. According to the theory proposed here, their lack of theory-of-mind skills in those with autism will derive, like their other impairments, at least in part, from a defect in extracting and using higher-order associations contained in mentalistic words. Given the dependence of theory-of-mind on language, and its proposed dependence upon higher-order context, this should be modelable by LSA.

Extrapolating the mathematics of the context processing underlying word meaning to the context processing of other faculties is unavoidably conjectural, particularly given so little is known experimentally or theoretically about them (for example, we do not know much about how children develop such a basic skill as a “street sense” of harm avoidance). One of the reasons for this, of course, has been past lack of appropriate scientific tools with which to model and formulate empirical investigation into cognition and context.

The recent novelty of the LSA approach, moreover, has provided insufficient time for the creation of general models of the role of context in cognition and development. The work of Landauer and Dumais, as a result, has not been followed by related analysis in other areas such as how human brain’s think, feel and socialize, and how such processes are acquired. The comments made here in regard to them are thus unavoidably preliminary and aimed to suggest the future plausibility and merit of making such analysis. However, without such theoretical development by cognitive and child developmental psychologists of the role of context, the ability to link impairment in LSA-type processes to the problems found to autism will be limited.

The impairment in the processes described by LSA, further, has been discussed as a single and uniform event. However, there are many stages involved and these could be affected in different ways. Thus, considerable variety could exist in the way context-based cognition and its extraction from higher-order associations becomes atypical. For example, varying the size of the window of association elements from which such information is extracted, and the number dimensions into which such extracted information is rerepresented might produce slightly different kinds of atypicalness. The multidimensional vector space created from higher-order association information might also be created but only accessible in part to normal cognition. As yet, it is not possible to theorize fully about the potential of this for heterogeneous forms of atypicallity.

In conclusion, the current lack of insight into forms of atypical cognition could be due to the incompleteness with which normal cognition and its development is modeled. LSA shows that until recently cognitive science missed in the domain of linguistic meaning the key role of context-based cognition (derived from extracted higher-order associations). This omission is understandable given that its investigation requires computer power that has only become recently available to the scientific investigation of cognition. LSA shows that in the domain of language the performance of many tasks can be best accounted for by simulations which employ such information. There is, of course, a theoretical jump as to whether cognition uses similar processes in nonlinguistic domains. However, in many of these it has been shown both that the circumstances are present for such information, and that the cognition processing in such domains shows characteristics – such as the use of context – that strongly suggests that they indeed use (or have a need to use) of such information. While it is possible that context information might be available to a cognition and go unused, it would be rather unlikely that cognition in one domain would employ such information, but not in another where it might be also useful. Parsimoniously, therefore, it is a good hypothesis that such information is used by cognition in areas outside that of language. Thus, the information provided by higher-order associations is likely to play a key role – albeit still an unknown one – in some form -- in many cognitive, social, affective, perceptual and motor faculties. Further, that following failure to consolidate such information would be produce atypicalness across diverse cognitive domains.

References

American Psychiatric Association. (1994). *Diagnostic and statistical manual of mental disorders* (4th ed.). Washington, DC: American Psychiatric Association.

Asanuma, K., Urushihara, R., Nakamura, K., Kitaoka, K., Sei, H., Morita, Y., et al. (2003). Premovement gating of somatosensory evoked potentials after tibial nerve stimulation. *Neuroreport, 14,* 375-379.

Astington, J. W., & Jenkins, J. M. (1999). A longitudinal study of the relation between language and theory of mind development. *Developmental Psychology, 35,* 1311-1320.

Baranek, G. T. (1999). Autism during infancy: a retrospective video analysis of sensory-motor and social behaviors at 9-12 months of age. *Journal of Autism and Developmental Disorders, 29,* 213-224.

Beversdorf, D. Q., Smith, B. W., Crucian, G. P., Anderson, J. M., Keillor, J. M., Barrett, A. M., et al. (2000). Increased discrimination of "false memories" in autism spectrum disorder. *Proceedings of the National Academy of Sciences of the United States of America, 27,* 8734-8737.

Blei, D. M., Ng, A. Y., & Jordan, M. J. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research, 3,* 993-1022.

Brosnan, M. J., Scott, F. J., Fox, S., & Pye, J. (2004). Gestalt processing in autism: failure to process perceptual relationships and the implications for contextual understanding. *Journal of Child Psychology and Psychiatry, 45,* 459-469.

Brown, G. D. (1997). Connectionism, phonology, reading, and regularity in developmental dyslexia. *Brain and Language, 59,* 207-235.

Cicourel, A. (1987). The interpretation of communicative contexts : Examples from medical encounters. *Social Psychology Quarterly, 50,* 217–226.

Clark, H. H. (1992). *Arenas of language use*. Chicago: University of Chicago Press.

Crone, C., & Nielsen, J. (1989). Spinal mechanisms in man contributing to reciprocal inhibition during voluntary dorsiflexion of the foot. *Journal of Physiology, 416,* 255-272.

Dagan, I., Marcus, S., & Markovitch, S. (1995). Contextual word similarity and
estimation from sparse data. *Computer Speech and Language, 9,* 123-152.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing By Latent Semantic Analysis. *Journal of the American Society For Information Science, 41,* 391-407.

Ekbia, H. R., & Magnitman, A. G. (2001). Context and relevance: A pragmatic approach. In V. Akman, P. Bouquet, R. H. Thomason & R. A. Young (Eds.), *Modeling and using Context, Third International and Interdisciplinary Conference* (pp. 156-169). Berlin: Springer.

Ellemers, N., Spears, R., & Doosje, B. (2002). Self and social identity. *Annual Review in Psychology, 53,* 161–186.

Foltz, P. W. (in press). Discourse coherence and LSA. In T. Landauer, D. McNamara, S. Dennis & W. Kintsch (Eds.), *Latent Semantic Analysis: A Road to Meaning.* Mahwah, NJ: Laurence Erlbaum.

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textural coherence with Latent Semantic Analysis. *Discourse Processes, 25,* 285-307.

Fombonne, E. (2003). The prevalence of autism. *JAMA, 289,* 87-89.

Frith, U. (1989). *Autism: explaining the enigma*. Oxford.: Blackwell.

Frith, U., & Happé, F. (1994). Language and communication in autistic disorders. *Philosophical Transactions of the Royal Society of London. Series B: Biological*

Sciences, 346, 97-104.

Frith, U., & Snowling, M. (1983). Reading for meaning and reading for sound in autistic and dyslexic children. *British Journal of Developmental Psychology, 1,* 329-342.

Gardner, E. P., Ro, J. Y., Debowy, D., & Ghosh, S. (1999). Facilitation of neuronal activity in somatosensory and posterior parietal cortex during prehension. *Experimental Brain Research, 127,* 329-354.

Gillberg, G., & Coleman, M. (1992). *The biology of the autistic syndrome* (2nd ed.). New York: Cambridge University Press.

Gorrell, G., & Webb, B. (2005). Generalized hebbian algorithm for incremental latent semantic analysis. *INTERSPEECH-2005*.

Gough, P. B., Alford, J. A., & Holley-Wilcox, P. (1981). Words and context. In O. J. Tzeng & H. Singer (Eds.), *Perception of print: Reading research in experimental psychology* (pp. 85-102). Hillsdale, NJ: Erlbaum.

Grace, A. A. (2000). Gating of information flow within the limbic system and the pathophysiology of schizophrenia. *Brain Research. Brain Research Reviews, 31,* 330-341.

Halberstadt, A. G., Denham, S. A., & Dunsmore, J. C. (2001). Affective Social Competence. *Social Development, 10,* 79-119.

Hale, B., & Wright, C. (1997). *A companion to the philosophy of language*. Oxford, UK ; Malden, Mass.: Blackwell Pub.

Hale, C. M., & Tager-Flusberg, H. (2005). The relationship between discourse deficits and autism symptomatology. *Journal of Autism and Developmental Disorders, 35,* 519-524.

Happé, F. G. (1994). *Autism*. London: UCL Press.

Happé, F. G. (1995). The role of age and verbal ability in the theory of mind task performance of subjects with autism. *Child Development, 66,* 843-855.

Happé, F. G. (1997). Central coherence and theory-of-mind in autism: Reading homographs in context. *British Journal of Developmental Psychology, 15,* 1-12.

Happé, F. G. (2006). Book review: Making sense of autism. *Nature, 442,* 632 - 633.

Happé, F. G., & Frith, U. (2006). The weak coherence account: Detail-focused cognitive style in autism spectrum disorders. *Journal of Autism and Developmental Disorders, 36,* 5-25.

Happé, F. G., Ronald, A., & Plomin, R. (2006). Time to give up on a single explanation for autism. *Nature Neuroscience, 9,* 1218-1220.

Harrison, J., & Hare, D. J. (2004). Assessment of sensory abnormalities in people with autistic disorders. *Journal of Autism and Developmental Disorders, 34,* 727-730.

Haugeland, J. (1989). *Artificial Intelligence: The very idea*. Mass, Cambridge: MIT Press.

Hofmann, T. (1999). Probabilistic latent semantic indexing. In *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 50–57). New York: ACM Press.

Houde, J. F., Nagarajan, S. S., Sekihara, K., & Merzenich, M. M. (2002). Modulation of the auditory cortex during speech: an MEG study. *Journal of Cognitive Neuroscience, 14,* 1125-1138.

Hoy, J. A., Hatton, C., & Hare, D. (2004). Weak central coherence: A cross-domain phenomenon specific to autism? *Autism, 8,* 267-281.

James, W. (1891). *Principles of psychology* (Vol. 1). London: Macmillan.

Jarrold, C., & Russell, J. (1997). Counting abilities in autism: Possible implications for central coherence theory. *Journal of Autism and Developmental Disorders, 27,* 25-37.

Joanisse, M. F., & Seidenberg, M. S. (1999). Impairments in verb morphology after brain

injury: a connectionist model. *Proceedings of the National Academy Sciences United States of America, 96,* 7592-7597.
Joanisse, M. F., & Seidenberg, M. S. (2003). Phonology and syntax in specific language impairment: evidence from a connectionist model. *Brain and Language, 86,* 40-56.
Jolliffe, T., & Baron-Cohen, S. (1999). A test of central coherence theory: linguistic processing in high-functioning adults with autism or Asperger's syndrome: Is local coherence impaired? *Cognition, 71,* 149-185.
Jolliffe, T., & Baron-Cohen, S. (2000). Linguistic processing in high-functioning adults with autism or Asperger's syndrome. Is global coherence impaired? *Psychological Medicine, 30,* 1169-1187.
Kanner, L. (1943). Autistic disturbances of affective contact. *Nervous Child, 2,* 217-250.
Kanner, L. (1946). Irrelevant and metaphorical language in early infantile autism. *American Journal of Psychiatry, 103,* 242-246.
Karst, G. M., & Hasan, Z. (1987). Antagonist muscle activity during human forearm movements under varying kinematic and loading conditions. *Experimental Brain Research, 67,* 391-401.
Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science, 302,* 1181-1185.
Kuhtz-Buschbeck, J. P., Stolze, H., John, K., Boczek-Funcke, A., & Illert, M. D. (1998). Development of prehension movements in children: a kinematic study. *Experimental Brain Research, 122,* 424-432.
Kwantes, P. J. (2005). Using context to build semantics. *Psychonomic Bulletin & Review, 12,* 703-710.
Landauer, T. K. (2002). On the computational basis oflearning and cognition: Arguments from LSA. In N. Ross (Ed.), *The psychology of learning and motivation* (Vol. 41, pp. 43-84).
Landauer, T. K., & Dumais, S. (1997). A Solution to Plato's Problem: The latent semantic analysis theory of acquisition, induction and representation of knowledge. *Psychological Review, 104,* 211-240.
Landauer, T. K., Foltz, P., & Laham, D. (1998). Introduction to Latent Semantic Analysis. *Discourse Processes, 25,* 259-284.
Landauer, T. K., McNamara, D., Dennis, S., & Kintsch, W. (in press). *Latent Semantic Analysis: A Road to Meaning*. Mahwah, NJ: Laurence Erlbaum.
Lewin, K. (1951). *Field Theory in Social Science*. New York: Harper and Row.
Lopez, B., & Leekam, S. R. (2003). Do children with autism fail to process information in context? *Journal of Child Psychology and Psychiatry and Allied Disciplines, 44*.
Maquert, P., Schwartz, S., Passingham, R., & Frith, C. (2003). Sleep-related consolidation of a visuomotor skill: Brain mechanisms as assessed by functinal magnetic resonance imaging. *Journal of Neuroscience, 23,* 1432-1440.
Martin, A., Wiggs, C. L., Ungerleider, L. G., & Haxby, J. V. (1996). Neural correlates of category-specific knowledge. *Nature, 379,* 649-652.
McCutcheon, A. C. (1987). *Latent Class Analysis*. Beverly Hills: Sage Publications.
Mervis, C. B., & Robinson, B. F. (1999). Commentary: Methodological issues in cross-syndrome comparisons: matching procedures, sensitivity (Se), and specify (Sp). *Monographs of the Society for Research in Child Development, 61,* 115-130.
Miller, G. R., & Coleman, E. B. (1967). A set of thirty-six prose passages calibrated for complexity. *Journal of Verbal Learning and Verbal Behavior, 6,* 851-854.
Mottron, L., & Belleville, S. (1993). A study of perceptual analysis in a high-level autistic subject with exceptional graphic abilities. *Brain and Cognition, 23,* 279-309.

Mottron, L., Burack, J. A., Stauder, J. E., & Robaey, P. (1999). Perceptual processing among high-functioning persons with autism. *Journal of Child Psychology and Psychiatry and Allied Disciplines, 40,* 203-211.

Nayate, A., Bradshaw, J. L., & Rinehart, N. J. (2005). Autism and Asperger's disorder: are they movement disorders involving the cerebellum and/or basal ganglia? *Brain Research Bulletin, 67,* 327-334.

Orban, P., Rauchs, G., Balteau, E., Degueldre, C., Luxen, A., Maquert, P., et al. (2006). Sleep after spatial learning promotes covert reorganization of brain activity. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 7124-7129.

Ozonoff, S. (1997). Casual mechanisms of autism, A unified perspective from an information-processing framework. In D. J. Cohen & F. R. Volkmar (Eds.), *Handbook of autism and pervasive developmental disorders* (2nd ed., pp. 868-879). New York: John Wiley.

Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition, 3,* 519-526.

Panksepp, J. (1990). The psychoneurology of fear: evolutionary perspectives and the role of animal models in understanding human anxiety. In R. Burrows (Ed.), *Handbook of Anxiety* (pp. 41–57). Amsterdam: Elsevier/North-Holland Biomedical.

Panksepp, J. (1998). *Affective Neuroscience: The Foundations of Human and Animal Emotions*. New York: Oxford University Press.

Peterson, C. C., & Siegal, M. (1999). Representing inner worlds: Theory-of-mind in autistic, deaf and normal hearing children. *Psychological Science, 10,* 126-129.

Prentice, D. A., & Banaji, M. R. (1994). The self in social contexts. *Annual Reviews in Psychology, 45,* 297-332.

Rohlfing, K. J., Rehm, M., & Goecke, K. U. (2003). Situatedness: The interplay between context(s) and situation. *Journal of Cognition and Culture, 3,* 132-156.

Rose-Krasnor, L. (1997). The nature of social competence. A Theoretical review. *Social Development, 6,* 111-135.

Shah, A., & Frith, U. (1983). An islet of ability in autistic children. *Journal of Child Psychology and Psychiatry and Allied Disciplines, 24,* 613-620.

Shah, A., & Frith, U. (1993). Why do autistic individuals show superior performance on the Block Design task? *Journal of Child Psychology and Psychiatry and Allied Disciplines, 34,* 1351-1364.

Sparck Jones, K. (1986). *Synonymy and Semantic Classification*. Edinburgh: Edinburgh University Press.

Sperber, D., & Wilson, D. (1986). *Relevance : communication and cognition*. Cambridge, Mass.: Harvard University Press.

Steyvers, M., & Griffiths, T. L. (in press). *Probabilistic topic models*. Mahwah, NJ: Laurence Erlbaum.

Storey, P. (1997). Examining the test-taking process: a cognitive perspective on the discourse cloze test. *Language Testing, 14,* 214-231.

Tager-Flusberg, H. (1991). Semantic processing in the free recall of autistic children: further evidence for a cognitive deficit. *British Journal of Developmental Psychology, 9,* 417-430.

Taylor, W. (1953). Cloze procedure: a new tool for measuring readability. *Journalism Quarterly, 30,* 415-433.

Teitelbaum, P., Teitelbaum, O., Nye, J., Fryman, J., & Maurer, R. G. (1998). Movement analysis in infancy may be useful for early diagnosis of autism. *Proceedings of the National Academy of Sciences of the United States of America, 95,* 13982-13987.

Thomas, M. S., & Karmiloff-Smith, A. (2003). Modeling language acquisition in atypical phenotypes. *Psychological Review, 110,* 647-682.

Todorov, A., Harris, L. T., & Fiske, S. T. (2006). Toward socially inspired social neuroscience. *Brain Research, 1079,* 76-85.

Turner, J. C., Oakes, P. J., Haslam, S. A., & McGarty, C. (1994). Self and collective: Cognition and social context. *Personality and Social Psychology Bulletin, 20,* 454-463.

Van Dijk, T. A. (1977). Context and cognition: Knowledge frames and speech act comprehension. *Journal of Pragmatics, 1,* 211–232.

Vigliocco, G., Warren, J., Siri, S., Arciuli, J., Scott, S., & Wise, R. (2006). The Role of Semantics and Grammatical Class in the Neural Representation of Words. *Cerebral Cortex*.

Volkmar, F. R., Lord, C., Bailey, A., Schultz, R. T., & Klin, A. (2004a). Autism and pervasive developmental disorders. *Journal of Child Psychology and Psychiatry, 45,* 135-170.

Volkmar, F. R., Lord, C., Bailey, A., Schultz, R. T., & Klin, A. (2004b). Autism and pervasive developmental disorders. *J Child Psychol Psychiatry, 45,* 135-170.

Waterhouse, L. W., Fein, D., & Modahl, C. (1996). Neurofunctional mechanisms in autism. *Psychological Review, 103,* 457-489.

Weber, C., Wermter, S., & Elshaw, M. (2006). A hybrid generative and predictive model of the motor cortex. *Neural Networks, 19,* 339-353.

Wiltgen, B. J., Brown, R. A., Talton, L. E., & Silva, A. J. (2004). New circuits for old memories: the role of the neocortex in consolidation. *Neuron, 44,* 101-108.