Ambiguity in sentence processing

Gerry T.M. Altmann

As listeners and readers, we rarely notice the ambiguities that pervade our everyday language. When we hear the proverb 'Time flies like an arrow' we might ponder its meaning, but not the fact that there are almost 100 grammatically permissible interpretations of this short sentence. On occasion, however, we do notice sentential ambiguity: headlines, such as 'Two Sisters Reunited After 18 Years in Checkout Counter', are amusing because they so consistently lead to the unintended interpretation (presumably, the sisters did not spend 18 years at the checkout). It is this consistent preference for one interpretation – and one grammatical structure – rather than another that has fuelled research into sentence processing for more than 20 years. Until relatively recently, the dominant belief had been that these preferences arise from general principles that underlie our use of grammar, with certain grammatical constructions being preferred over others. There has now accrued, however, a considerable body of evidence demonstrating that these preferences are not absolute, but can change in particular circumstances. With this evidence have come new theories of sentence processing, some of which, at first glance, radically question the standard notions of linguistic representation, grammar and understanding.

We rarely stop to notice that 'eating pizza with a friend' could (but does not) mean that the friend fulfilled the same function as the knife and fork in 'eating pizza with a knife and fork', or as the glass of wine in 'with a glass of wine'. The pervasiveness of ambiguity during sentence processing does not always pass us by, however. Occasionally we choose an interpretation that is not the one intended by the speaker or, in printed texts, the writer. A sentence such as (1) below is grammatical but, more often than not, interpreted as ungrammatical:

(1)   *The horse raced past the crowd stumbled.*

Commonly, people reach the word 'stumbled' in (1) and believe either that the sentence is ungrammatical, or that the sentence was supposed to be either: 'The horse raced past the crowd <u>and</u> stumbled' or 'The horse raced past, (<u>and</u>) the crowd stumbled'. The problem in (1) is that the verb 'raced' tends to be interpreted as a main verb, which subsequently proves incompatible with the final word 'stumbled'. In fact, the verb 'raced' is ambiguous between a main verb in the past tense and a **past participle** (see Glossary), and only the past participle version is compatible with the final continuation of this sentence. In (2) below, the verb 'ridden' can only be a past participle, and this sentence is therefore relatively unproblematic:

(2)   *The horse ridden past the crowd stumbled.*

Other problematic cases include sentences such as the following:

(3)   *He will read the paper that he received tomorrow.*

We tend to interpret the final **adverb** 'tomorrow' as saying something about the immediately preceding verb ('received') – yet the past tense of the verb is incompatible with an adverb that refers to the future. Instead, we should interpret the adverb as saying something about the earlier verb (that is, as saying something about when the reading will be done).

In some respects, sentences such as these might at first glance be considered oddities and, at best, unrepresentative of the language we normally encounter. But we can nonetheless ask: Why do readers of such sentences consistently misinterpret them? What is it about the grammatical structure of the language, or the structure and organization of our language faculty, that causes us to mis-analyse consistently the grammatical structure of these sentences?

Syntactic explanations

The earliest attempts to explain these consistent misinterpretations were based on an analysis of the underlying **syntactic structure** of such sentences. John Kimball[1] and subsequently Lyn Frazier[2] pointed out that the interpretation of these ambiguities seemed to be determined by consistent differences in the syntactic structures that accompanied the alternative readings of each sentence. Frazier developed two principles to explain the behaviour of the 'human sentence processing mechanism' when faced with

G.T.M. Altmann is at the Department of Psychology, University of York, York, UK YO1 5DD.

tel: +44 1904 434362
fax: +44 1904 433181
e-mail: g.altmann@ psych.york.ac.uk

Glossary

Adverb: a word referring to some property of an event, such as when the event happened ('yesterday') or how it happened.

Agent: the active instigator of an event (usually also the subject, in English).

Argument: a participant (animate or otherwise) in an event defined by a verb. The verb 'put' takes three 'arguments': he put the ball in the box.

Argument structure: the information associated with a verb regarding the syntactic context in which it should occur.

Connectionism: a computational approach to modelling various phenomena that assumes networks of interconnecting units whose activation levels can, by virtue of those interconnections, influence the activation levels of their neighbours. Typically, information is encoded in the pattern of connectivity between the units, and the pattern of activation across those units in response to an input reflects properties of the input (as defined by the information already encoded within the network).

Event-related brain potential: electrical activity recorded at the scalp that changes in response to some (experimentally induced) event.

Interactive activation: a term used to refer to a class of model in which representations are activated to varying degrees (depending on the input), and in which the activation of one representation can influence the activation of another (hence the interaction).

Intransitive: an intransitive verb requires only a subject, as in 'he slept'.

Noun phrase: a phrase that refers to a thing or things, such as 'the horse', 'the horse that lost the race', etc.

Object: the grammatical object of a sentence is, typically, the patient or theme in the event described by the sentence. In 'he enjoyed the food', the object is 'the food'.

Past participle: the form of a verb used to indicate a passive in the past tense, such as 'driven', in 'the car driven by the world champion is a …'.

Patient: a participant who is changed or directly affected by an instigated event.

Referential context: a term used to refer to that aspect of the context that relates to the different entities, which may, or may not, be referred to subsequently.

Relative clause: a phrase that can be used to provide additional information about whatever is being referred to by a noun phrase, as in 'the horse that lost the race'.

Subject: the grammatical subject of a sentence is, typically, the agent of the event described by the sentence. In 'he enjoyed the food', the subject is 'he'.

Syntactic structure: the organization of the words in a sentence into constituent phrases, and the relationships between them. Such structures are often notated (see Figs 1 and 2).

Theme: a participant who is spatially displaced as a result of an instigated event (cf. patient).

Transitive: a transitive verb requires both a subject and an object, as in 'she ate the apple'.

---

syntactic ambiguities. The first of these, called 'minimal attachment', was based on the observation that, for at least one class of ambiguity, the structure associated with the least preferred interpretation is more complex – it contains more branching points ('nodes') than the structure associated with the preferred interpretation (see Fig. 1). This principle dictates that the sentence processor adopts the analysis that requires the simpler structure. In example (1) given above, the unpreferred (but correct) reading requires more nodes in the corresponding syntactic structure.

The second principle applies to example (3) above, in which each interpretation contains the same number of nodes (see Fig. 2). Called 'late closure', this principle dictates that the clause currently being processed is kept 'open' as long as possible (with the consequence that incoming material will be incorporated within the current, or more recent, clause). A wide range of sentential ambiguities can be analysed in terms of one or other of these principles. Both were motivated, originally, in terms of reducing memory load, although Frazier subsequently suggested, following Fodor[3], that the human sentence processing mechanism consists of informationally encapsulated modules, with the syntactic module being 'informationally encapsulated' from the others (i.e. there is no non-syntactic influence on syntactic decision-making)[4,5] – in effect, decisions regarding syntactic structure are made solely on the basis of syntactic information because no other knowledge is available to the subsystem responsible for syntactic decision-making.

Multiple constraints in sentence processing

More recently, an alternative view of human sentence processing has emerged that directly challenges this notion of informational encapsulation. It assumes instead that multiple sources of information interact, with each source of information constraining the interpretation in a particular way[6,7]. An important aspect of this approach is that the different constraints (see below) are applied in parallel and are probabilistic. This idea borrows much from findings in the word-recognition literature. In particular, there are three findings within that literature that motivate the constraint-satisfaction approach to sentence processing:

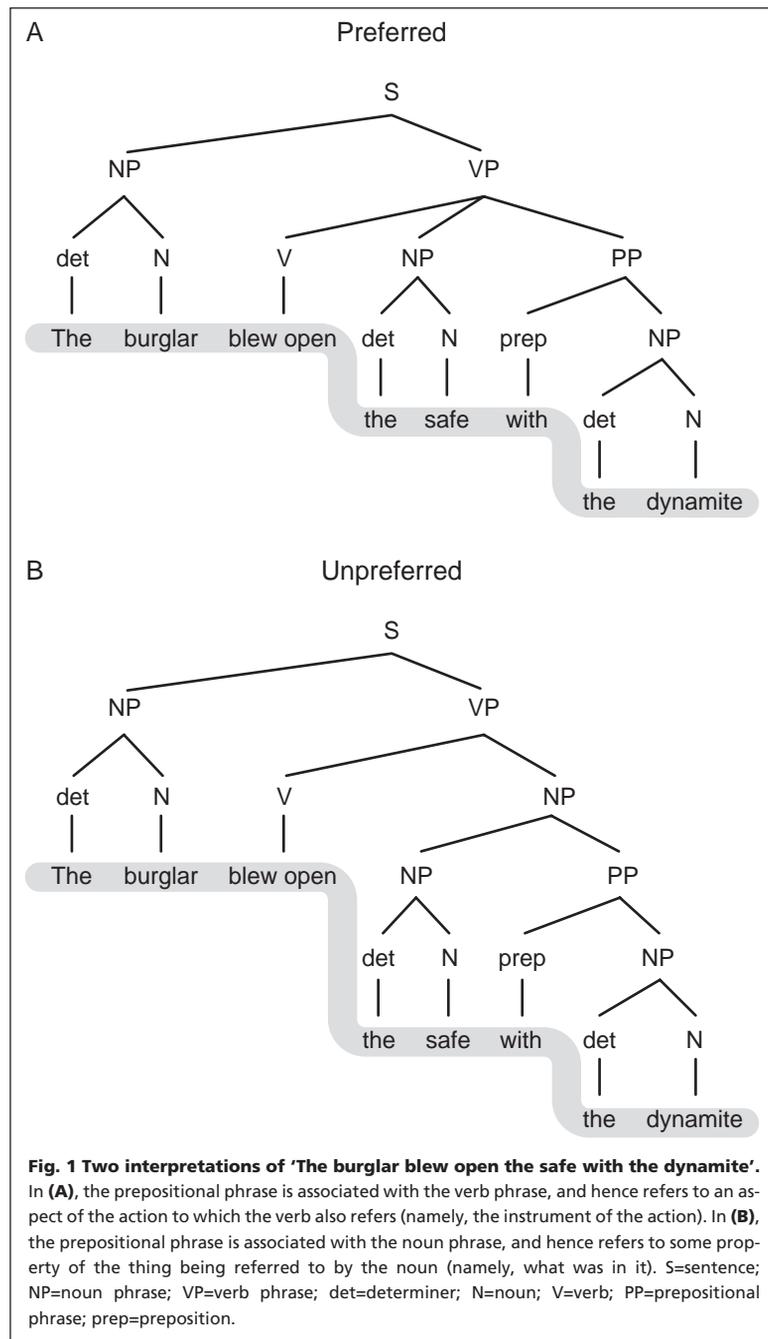(1) Multiple meanings of a single ambiguous word are briefly activated[8,9].

(2) The frequency (or probability) of occurrence of each meaning in the language at large determines the relative degree of activation of the alternatives[10].

(3) Biasing contexts can increase the activation of one or other of the alternatives (so that, for example, a less frequent meaning can become activated to the same degree as a more frequent meaning)[11].

How do these findings relate to ambiguity in sentence processing? Consider again the first example:

(1)　*The horse raced past the crowd stumbled.*

Because the verb 'raced' is ambiguous, multiple representations, corresponding to the meanings of the main verb and past participle versions, will be activated. These representations will include information not simply about what 'race' means, but also about the syntactic (or argument) structures it can occur in (it can be **intransitive** – taking just one **argument**, as in 'he raced', or **transitive** – taking two arguments, as in 'he raced the horse'). Whichever is the most frequent will be the more active[6]. However, the frequency with which 'raced' occurs as a past participle or as

Fig. 1 Two interpretations of 'The burglar blew open the safe with the dynamite'. In **(A)**, the prepositional phrase is associated with the verb phrase, and hence refers to an aspect of the action to which the verb also refers (namely, the instrument of the action). In **(B)**, the prepositional phrase is associated with the noun phrase, and hence refers to some property of the thing being referred to by the noun (namely, what was in it). S=sentence; NP=noun phrase; VP=verb phrase; det=determiner; N=noun; V=verb; PP=prepositional phrase; prep=preposition.

'**agent**', but a good '**patient**' or '**theme**'). In 'the cop arrested…', the cop is likely to be doing the arresting, and is therefore a good agent (although, legitimately, the cop could be the person being arrested – perhaps with less likelihood). Thus, the plausibility with which the surface **subject** can fill one role (either as the thing doing the action referred to by the verb, or as the thing that the action is being done to) is another source of constraint[13].

Each of the factors discussed so far has been shown, by a variety of means (for example, see Box 1), to have an independent influence on the initial choice of interpretation in cases like the 'raced' example. And because the influences are independent, they can on occasion conflict. The first sentence in the pair shown below should be entirely straightforward:

(4)   *The sofa scratched by the cat was badly damaged.*
(5)   *The sofa that was scratched by the cat was badly damaged.*

Sofas are more likely to be scratched than to scratch, yet 'scratched' is rarely used as a past participle. Indeed, analysis of the time taken to read 'by the cat' in sentence (4) compared with the unambiguous version in (5) reveals that 'by the cat' is unexpected in version (4) – it engenders longer reading times, which suggests that despite the implausibility of sofas scratching, the main verb interpretation is initially preferred[13].

An important feature of this constraint-based approach to syntactic ambiguity is that representations corresponding to alternative interpretations are activated, but that this activation is both graded and dynamically changing as the sentence unfolds, and as constraints continue to apply. Thus, it is not the case, according to such theories, that all possible interpretations are held in working memory, each equally accessible, until some decision process decides between them.

The role of extra-sentential context
Thus far, the only information we have considered that can influence the interpretation of sentential ambiguity has been information derived from within the sentence itself. However, a further source of information that can influence sentence interpretation is the context, linguistic or otherwise, in which the ambiguous sentence occurs. The following is another example of a sentence that is often initially perceived as ungrammatical:
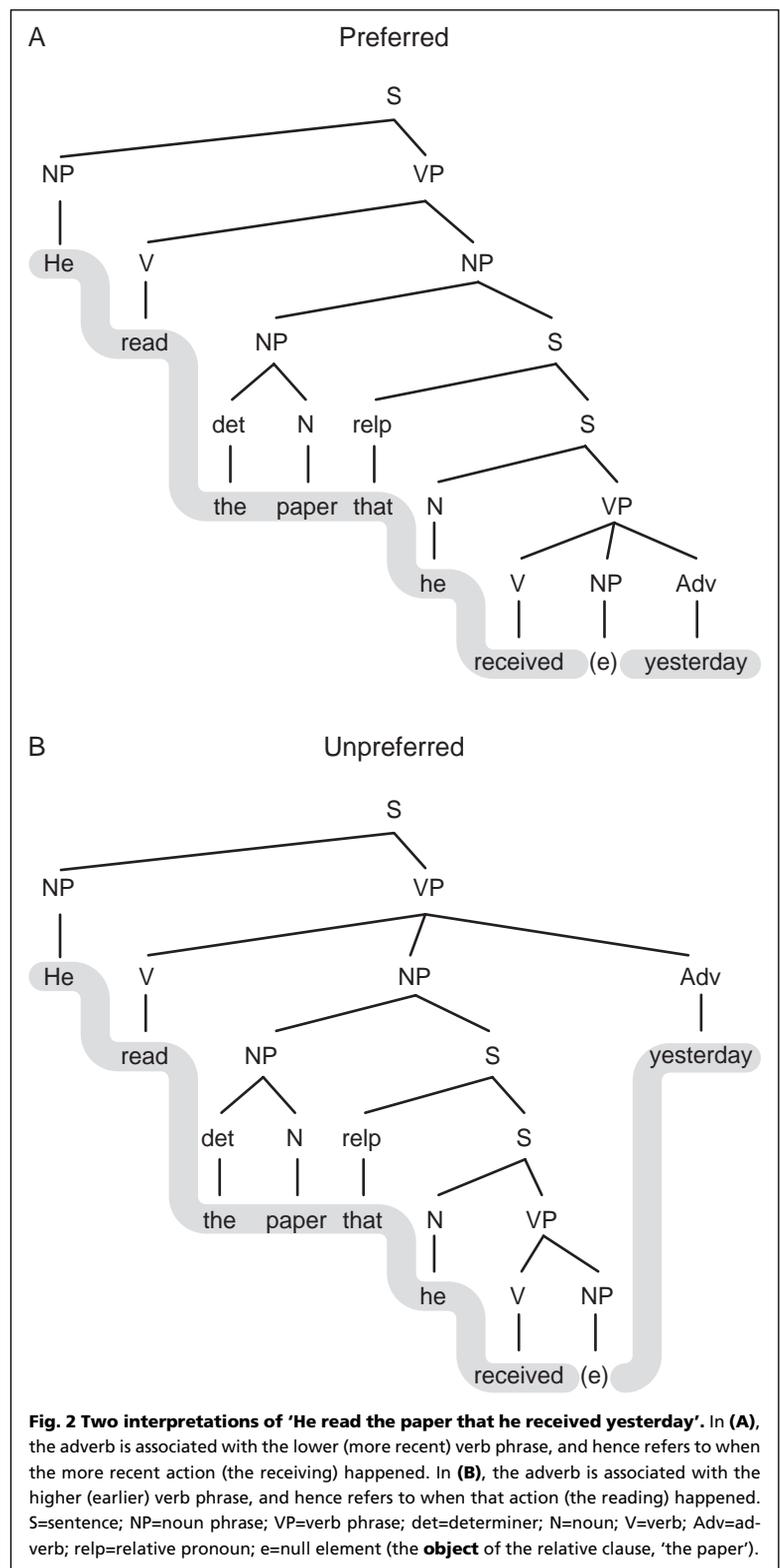
(6)   *He told the journalist that he had recently written to to hurry up and submit his piece.*

Commonly, people interpret the sequence starting 'that he had…' as referring to what was told to the journalist, in which case the second 'to' appears ungrammatical. Instead, this sequence should be interpreted as a **relative clause** – that is, as telling us something about which journalist is being referred to, as in 'He told the one whom he had recently written to to hurry up with his piece'. Stephen Crain, working with children, has observed that children are often better able to process relative clauses when they are presented in a context that makes them necessary[14]. For this last example, if the fragment 'He told the journalist' were

a main verb is only one factor that will influence which meaning of 'raced' is the more activated. Another factor might be the frequency with which 'raced' is used in its intransitive form or its transitive form[12,13]. These two factors interact inasmuch as the passive past participle form is necessarily transitive, whereas the active main verb form can be either. Different verbs will have different properties with respect to these factors (e.g. for 'enjoyed', the past tense form is more frequent than the past participle form; for 'examined' the two forms are equiprobable; for 'received' the past participle form is the more frequent – each case is obligatorily transitive).

The context within which each verb occurs can also play a part. Take a sentence fragment, such as 'the burglar arrested…'. Burglars tend to *be* arrested, rather than to arrest, and so the burglar here is unlikely to be doing the arresting (it is, in the context of the verb 'arrest', a bad

presented in a context that introduced two or more journalists, a relative clause would be required in order to specify which journalist was intended (that is, which journalist was being referred to): 'the journalist that he had recently written to'. Crain and Steedman presented pilot data from adult studies suggesting that the problems normally encountered with sentences such as the one above could indeed be overcome in appropriate '**referential contexts**'[15]. Several studies followed that supported this view[16–21]. One recent study involved sitting people in front of a table on which there were various objects, and monitoring their eye movements as they acted out various instructions, such as 'Put the apple on the towel in the box'. In this case, the visual world, rather than a mental representation of previous linguistic material, constituted the context. Ordinarily, 'on the towel' is initially interpreted as where the apple should be put (as predicted, in fact, by those earlier syntactic accounts of ambiguity resolution). However, the data demonstrate that if people are given this instruction when faced with two apples in front of them, just one of which is on a towel, they *immediately* interpret 'on the towel' as indicating which apple is being referred to, and not as indicating where it should be put[22]. It was found, in this study, that the eyes would move, on hearing 'on the…', to the apple that was on the towel, as opposed to either the other apple or a second (empty) towel that could be a potential target location.

How do such contextual factors interact with the more local constraints described earlier? According to the constraint-satisfaction approach to sentence processing, the availability of alternative structures will be graded (according to the probabilistic effects of frequency, plausibility and so on), and this differential availability determines the extent to which any single constraint can influence the interpretive process. Thus, the stronger one kind of constraint (for example, the high frequency of occurrence of one **argument structure** rather than another), the weaker the potential influence of others. We saw an example of this with the 'sofa scratched' example earlier – there, the frequency of occurrence of 'scratched' as an active main verb is more influential than the implausibility (or untypicality) of sofas scratching. Exactly the same kind of trade-off has been observed with context; the stronger the local constraint (frequency of occurrence of one structure over another), the weaker the influence of contextual constraints[21].

The accumulation of evidence in favour of a constraint-based approach to sentence processing has led proponents of the earlier 'syntactic' accounts to modify their theories so as to incorporate the possibility of non-syntactic influences on the initial processing of ambiguities. One such account, termed 'construal'[23], maintains nonetheless that there are restricted cases where non-syntactic influences are ruled out. For example, this account maintains that in a sentence such as 'he put the apple on the towel…', syntactic principles dictate that the phrase 'on the towel' must be interpreted as an argument of the verb 'put' – indicating the location of the putting. Such an account is clearly at odds with the earlier finding that, in appropriate contexts, 'on the towel' is *not* interpreted as an argument of 'put'[22].

Fig. 2 Two interpretations of 'He read the paper that he received yesterday'. In **(A)**, the adverb is associated with the lower (more recent) verb phrase, and hence refers to when the more recent action (the receiving) happened. In **(B)**, the adverb is associated with the higher (earlier) verb phrase, and hence refers to when that action (the reading) happened. S=sentence; NP=noun phrase; VP=verb phrase; det=determiner; N=noun; V=verb; Adv=adverb; relp=relative pronoun; e=null element (the **object** of the relative clause, 'the paper').

Activation and decay during sentence processing

The constraint-satisfaction view of sentence processing does more than simply propose that there exist constraints on sentence processing: it proposes also that these constraints are applied within a framework of '**interactive activation**', as exemplified by the **connectionist** approach to cognitive modelling[24]. In that approach, interconnecting units increase and decrease in activation as a function of the inputs they receive across their connections. Recently, a number of computational models have been developed, using such

Box 1. Monitoring the resolution process with event-related brain potentials

When people read something that becomes ungrammatical, a particular pattern of electrical activity, an '**event-related brain potential**'[a], can be recorded that peaks at around 600 ms after the point at which the sentence ceases to be grammatical[b]. This phenomenon can be used to study how different constraints influence the resolution of syntactic ambiguity. In the first of the following sentences, this peak (labelled 'P600') occurs in response to the word 'was':

(1)   *The captain liked the crew <u>was</u> unhappy*
(2)   *The captain heard the crew <u>was</u> unhappy*
(3)   *The captain believed the crew was unhappy*

The word 'was' is ungrammatical in (1) because the sentence should end after 'crew'. Sentence (2) contains the verb 'heard' which is ambiguous insofar as it can be followed either (as in this case) by a phrase describing what the captain heard (that the crew was unhappy), or simply by a **noun phrase** describing the thing or person that he heard (compare with 'The captain heard them'). However, this second alternative is very much more common than the first, and this is reflected in the same P600 indicator of ungrammaticality as in the first case[c] –

people assume that 'the crew' is what the captain heard (not that he heard something about them), and consequently the continuation 'was unhappy' is grammatically anomalous. Of course, the anomaly is not so severe that we cannot easily recover from it and re-analyse the sentence accordingly, but evidently the recording technique is sensitive to small local anomalies of this kind. Finally, sentence (3) contains the verb 'believed', which, like 'heard', permits the same two-argument structures. In this case, however, the actual sentence-continuation structure is the more common, and now no P600 is observed.

References

a Osterhout, L., McLaughlin, J. and Bersick, M. (1997) Event-related brain potentials and human language *Trends Cognit. Sci.* 1, 203–209
b Osterhout, L. and Holcomb, P.J. (1992) Event-related brain potentials elicited by syntactic anomaly *J. Mem. Lang.* 31, 785–806
c Osterhout, L., Holcomb, P.J. and Swinney, D.A. (1994) Brain potentials elicited by garden-path sentences: evidence of the application of verb information during parsing *J. Exp. Psychol. Learn. Mem. Cognit.* 20, 786–803

systems, that successfully capture the interplay of the different constraints[25,26]. Within the original formulation of the constraint-satisfaction approach it has been suggested that certain preferences of interpretation are due not to probabilistic constraints of the kind outlined thus far, but to the way in which a pattern of activity across those interconnecting units will gradually decay with time[6]. Thus, whereas the constraints we have seen so far are due to the encoding of probabilistic information across those interconnecting units (encoded in the pattern of 'connectivity' between the units), other constraints are due, in effect, to the physical properties of the system (its 'architecture'). An example of this latter kind of constraint is embodied in sentence (7) below [similar to sentence (3) above], in which the final adverb ('yesterday') can be associated, in principle, with either of the two verbs that precede it:

(7)   *He read the paper that he received yesterday.*

People tend to associate the word 'yesterday' with the more recent verb. If one assumes that the representations corresponding to the argument structures of each verb become activated when that verb is encountered, but then gradually decay over time, it follows that the representation associated with the more recent verb will have decayed less than the representation associated with the earlier one. Consequently, if incoming words are incorporated within the most active representation, there will be a preference to associate the final adverb, in examples such as this last one, with the more recent verb (in essence, this is simply a re-statement of the 'late closure' preference, but cast in activation terms).

Recently, it has been shown that even this preference can be reversed in appropriate contexts – if the context includes, for example, a phrase such as 'they wondered when he read the paper that he received', it is as easy in (7) to as-

sociate the final adverb with the earlier verb as it is to associate it with the more recent one[27]. This is because the context serves to keep the representation associated with 'read' active – the question focuses attention on that verb. Interestingly, if some additional words are inserted between 'received' and 'yesterday', it becomes easier to associate the adverb with the earlier verb:

(8)   *He read the paper that he received <u>from the journal</u> yesterday.*

The reason that the preference reverses in this case, but is merely eliminated in version (7), is that the addition of the extra material ('from the journal') allows time for the representations associated with 'received' to decay, leaving those associated with 'read' (which receive contextual support) as the most active. When that extra material is absent, as in the first case, the representations associated with the more recent verb have not had time to decay, and are therefore available for the adverb to associate with.

The finding that contextual information can influence the interpretation of sentential ambiguity is clearly relevant insofar as it represents another source of constraint that can be applied during sentence processing. However, we need also to explain what it means to say that a representation 'receives contextual support'? What is the mechanism of that support?

Sentence interpretation as a predictive process

To explain the contextual effects observed with the last two examples, it has been proposed that the context ('they wondered when he would read the paper…') sets up predictive expectations regarding where, within the subsequent sentence, the relevant information may be found[27]. As that sentence unfolds, the reader will expect the adverb to occur only in very specific locations. In representational terms,

this corresponds to the activation of a representation of a following adverb at each of the positions in which an adverb like 'yesterday' could occur. This is reminiscent of research with small-scale connectionist networks that can learn, in effect, to predict the syntactic categories that could occur in particular sentential positions[28,29]. Effectively, then, the context sets up an expectation regarding the location of the relevant information in the target sentence, and this is manifested as the 'predictive activation' of structures that support the subsequent integration of that information. Much the same principles can be applied to the cases, described above, in which the processing of fragments, such as 'He told the journalist that…', was shown to be dependent on whether the context introduced one or more than one subject (journalist). According to the theory, these principles would apply so long as the human cognitive system was sensitive to the predictive contingencies that hold between relative clauses and the contexts (introducing multiple entities) within which those relative clauses are used.

Of course, the claim that structure can be predictively activated in this way, and that this activation supports the subsequent integration of the information corresponding to the structure that is activated, is just the same as the claim that representations corresponding to argument structures are activated in response to each verb that is encountered. An argument structure is nothing more than the set of predictive contingencies that hold between a verb and the sentential structures within which that verb can occur. (The nature of such contingencies necessarily makes them sensitive to factors like frequency of occurrence, and necessarily encodes factors like typicality or plausibility.) These contingencies 'support integration' insofar as the meaning of a verb, or indeed of any word, is, broadly speaking, little more than the encoding of the predictive contingencies that hold between that word and the contexts within which it can be experienced[30]. Simple artificial neural networks can be built that are able to encode such contingencies, and in so doing can induce internal representations that are analogous to linguistic structures[31]. However, it remains an empirical and (to an extent) theoretical issue whether human neural structures encode, and acquire, linguistic structures in equivalent ways. Connectionist systems are not fully implemented models of sentence processing – rather, they are small-scale implementations of fragments of the language system that are designed to capture principles believed to be shared with human language processing. Thus, they are better considered as a framework within which to cast theories of such processing.

Ambiguity in language is part of the legacy we inherit through having to *acquire* a language. The relationship between the sounds of a language and the world to which those sounds refer is inherently ambiguous – which of the different sounds refer, in combination with which others, to which things in the infant's environment? This acquisition problem is solved, in part at least, by the extraordinary sensitivity to predictive contingencies that we possess. This is not to say that language is nothing more than complex statistics. The challenge is to explore what else is required, and why the complexities required are apparently beyond the sensitivities of other species. Research into ambiguity

resolution during sentence processing is just one of many fields of study that should inform this exploration.

....

References

 1 Kimball, J. (1973) Seven principles of surface structure parsing in natural language *Cognition* 2, 15–47
 2 Frazier, L. (1979) *On Comprehending Sentences: Syntactic Parsing Strategies*, Indiana University Linguistics Club, Bloomington
 3 Fodor, J.A. (1983) *Modularity of Mind*, MIT Press
 4 Frazier, L. (1987) Sentence processing: a tutorial review, in *Attention and Performance XII: The Psychology of Reading* (Coltheart, M., ed.), pp. 559–586, Erlbaum
 5 Frazier, L. (1990) Exploring the architecture of the language-processing system, in *Cognitive Models of Speech Processing: Psycholinguistic and Computational Perspectives* (Altmann, G.T.M., ed.), pp. 409–433, MIT Press
 6 MacDonald, M.C., Pearlmutter, N.J. and Seidenberg, M.S. (1994) The lexical nature of syntactic ambiguity resolution *Psychol. Rev.* 101, 676–703
 7 Trueswell, J.C. and Tanenhaus, M.K. (1994) Towards a lexicalist framework of constraint-based syntactic ambiguity resolution, in *Perspectives on Sentence Processing* (Clifton, C., Frazier, L. and Rayner, K., eds), pp. 155–179, Erlbaum
 8 Tanenhaus, M.K., Leiman, J.M. and Seidenberg, M.S. (1979) Evidence for multiple stages in the processing of ambiguous words in syntactic context *J. Verbal Learn. Verbal Behav.* 18, 427–441
 9 Swinney, D.A. (1979) Lexical access during sentence comprehension: (re)considerations of context effects, *J. Verbal Learn. Verbal Behav.* 18, 645–659
 10 Simpson, G. (1981) Lexical ambiguity and its role in models of word recognition *Psychol. Bull.* 96, 316–340
 11 Duffy, S., Morris, R. and Rayner, K. (1988) Lexical ambiguity and fixation times in reading *J. Mem. Lang.* 27, 429–446
 12 MacDonald, M. (1994) Probabilistic constraints and syntactic ambiguity resolution *Lang. Cognit. Processes* 9, 157–202
 13 Trueswell, J.C. (1996) The role of lexical frequency in syntactic ambiguity resolution *J. Mem. Lang.* 35, 566–585
 14 Hamburger, H. and Crain, S. (1982) Relative acquisition, in *Language Development* (Kuczaj, S.A., ed.), pp. 245–274, Erlbaum
 15 Crain, S. and Steedman, M.J. (1985) On not being led up the garden path: the use of context by the psychological parser, in *Natural Language Parsing: Psychological, Computational, and Theoretical Perspectives* (Dowty, D., Karttunen, L. and Zwicky, A., eds), pp. 320–358, Cambridge University Press
 16 Altmann, G.T.M. and Steedman, M.J. (1988) Interaction with context during human sentence processing *Cognition* 30, 191–238
 17 Altmann, G.T.M. (1988) Ambiguity, parsing strategies, and computational models *Lang. Cognit. Processes* 3, 73–97
 18 Altmann, G.T.M., Garnham, A. and Dennis, Y. (1992) Avoiding the garden path: eye movements in context *J. Mem. Lang.* 31, 685–712
 19 Altmann, G.T.M., Garnham, A. and Henstra, J.A. (1994) Effects of

syntax in human sentence parsing: evidence against a structure-based proposal mechanism *J. Exp. Psychol. Learn. Mem. Cognit.* 20, 209–216

20 Britt, M.A. (1994) The interaction of referential ambiguity and argument structure *J. Mem. Lang.* 33, 251–283

21 Spivey-Knowlton, M. and Tanenhaus, M. (1994) Referential context and syntactic ambiguity resolution, in *Perspectives on Sentence Processing* (Clifton, C., Frazier, L. and Rayner, K., eds), pp. 415–439, Erlbaum

22 Tanenhaus, M.K. *et al.* (1995) Integration of visual and linguistic information in spoken language comprehension *Science* 268, 1632–1634

23 Frazier, L. and Clifton, C. (1995) *Construal*, MIT Press

24 McLeod, P., Plunkett, K. and Rolls, E. (1998). *Introduction to Connectionist Modelling of Cognitive Processes*, Oxford University Press

25 McRae, K., Spivey-Knowlton, M.J. and Tanenhaus, M.K. Modelling the influence of thematic fit (and other constraints) in on-line sentence comprehension *J. Mem. Lang.* (in press)

26 Tabor, W., Juliano, C. and Tanenhaus, M.K. (1997) Parsing in a dynamical system: an attractor-based account of the interaction of lexical and structural constraints in sentence processing *Lang. Cognit. Processes* 12, 211–271

27 Altmann, G.T.M., Garnham, A., van Nice, K. and Henstra, J.A. Late closure in context *J. Mem. Lang.* (in press)

28 Elman, J.L. *et al.* (1996) *Rethinking Innateness: A Connectionist Perspective on Development*, MIT Press/Bradford Books

29 Elman, J.L. (1993) Learning and development in neural networks: the importance of starting small *Cognition* 48, 71–99

30 Altmann, G.T.M. (1997) *The Ascent of Babel: An Exploration of Language, Mind, and Understanding*, Oxford University Press

31 Elman, J.L. (1990) Finding structure in time *Cognit. Sci.* 14, 179–211

32 Cuetos, F. and Mitchell, D.C. (1988) Cross-linguistic differences in parsing: restrictions on the use of the late closure strategy in Spanish *Cognition* 30, 73–105

33 Pearlmutter, N.J. and MacDonald, M.C. (1995) Individual differences and probabilistic constraints in syntactic ambiguity resolution *J. Mem. Lang.* 34, 521–542

Cerebral asymmetry: motoring on

Michael C. Corballis

I argue that the phylogenetic and neurobiological bases for cerebral asymmetry in humans are likely to be found in motor systems rather than in perceptual systems. Current genetic models of human laterality suggest that a 'dextral' allele might be responsible for right-handedness and left-cerebral dominance for speech in the majority of humans. The linking of handedness with language lateralization might reflect the early evolution of language as predominantly a system of manual gestures, perhaps switching to a mainly vocal system only with the emergence of *Homo sapiens*. So-called 'mirror neurons' in the prefrontal cortex of the monkey, which fire both when the animal makes a grasping response and when it sees the same response performed by others, might be part of a circuit that is the precursor to language circuits in the brain. This circuit appears to be bilateral in monkeys, but left-hemispheric in humans.

M.C. Corballis is at the Research Centre for Cognitive Neuroscience, University of Auckland, Private Bag 92019, Auckland, New Zealand.

tel: +64 9 373 7599
fax: +64 9 373 7450
e-mail: m.corballis@
auckland.ac.nz

Perhaps the major question confronting research on cerebral asymmetry is whether it will survive into the new millennium. The early wave that began with Broca's discoveries in the 1860s faded after the turn of the century[1]. Interest was revived in the 1960s following Sperry's work on the split brain, and has continued more or less unabated into the 1990s. However, as in the late 19th century, much of the theorizing about the dual brain is fanciful, even exploitative, and one wonders if it will once again fade through want of credibility, or perhaps from a sheer surfeit of information, much of which is confusing and contradictory. Indeed its death has already been announced[2]. Yet, there are signs that we might be moving towards a more realistic and biologically based understanding of what cerebral asymmetry is really about.

Empirical research has overwhelmingly stressed asymmetry at the perceptual level. This is largely a result of the development of perceptual techniques, such as divided visual-field studies or dichotic listening, that appear to be sensitive to cerebral asymmetry in normal people, although to a lesser extent than in split-brained patients. Reviewing the evidence, Hellige[3] concluded that there is no single principle that can account for variety of cerebral asymmetries even within the visual domain. Efron[2] is more damning, suggesting that the 'cottage industry' of perceptual laterality studies actually has little to say about hemispheric specialization. Yet, perception theorists remain undaunted, and the effort to discover general principles continues[4,5].

Nevertheless, there are compelling reasons to look to motor systems rather than perceptual ones for the origins of cerebral asymmetry. First, we must suppose that there would be disadvantages to a perceptual system that was asymmetrical to any substantial degree. The perceptual world is largely indifferent with respect to left and right, and a deficit on one side would leave an animal vulnerable to attack from that