

Modeling sentence processing in ACT-R

Shravan Vasishth

Department of Computational Linguistics
Saarland University, PO Box 15 11 05
66041 Saarbrücken, Germany
vasishth@acm.org

Richard L. Lewis

Department of Psychology
University of Michigan
Ann Arbor, MI, USA
rickl@umich.edu

Abstract

We present a series of simulations of behavioral data by casting a simple parsing model in the cognitive architecture ACT-R. We show that constraints defined in ACT-R, specifically those relating to activation, can account for a range of facts about human sentence processing. In doing so, we argue that resource limitation in working memory is better defined as an artefact of very general and independently motivated principles of cognitive processing.

1 Introduction

Although language processing may be a specialized cognitive faculty, it is possible that it is nevertheless shaped by general constraints on the human cognitive architecture. This point has been addressed extensively in the connectionist literature, but we present a somewhat different approach to this problem by casting parsing within the cognitive architecture ACT-R (Anderson et al., 2002) and directly using the constraints provided in ACT-R to account for several interesting cross-linguistic facts: the well-known sentential complement/relative clause asymmetry (Gibson, 2000; Grodner and Gibson, 2003) and the subject/object relative clause asymmetry in English (Homes and O'Regan, 1981); and some recent results (Vasishth, 2003) involving Hindi center embeddings, including a principled account of individual variation in subject behavior.

In developing this approach, we argue that resource limitation in working memory is better defined as an artefact of very general constraints on information processing – specifically, rehearsal and activation – rather than as an inherent numerical bound on memory capacity (cf. (Gibson, 2000; Hawkins, 1994); also see Section 3.5).

In the rest of this paper, we first introduce the ACT-R architecture. Then we present the results of several simulations of experiments available in the psycholinguistic literature. The paper concludes with a discussion of the potential advantages and

shortcomings of this approach, and of the broader consequences of modeling parsing within a cognitive architecture.

2 A brief introduction to the cognitive architecture ACT-R

ACT-R is a theory of the human cognitive architecture. It allows the development of computational models that can closely simulate experimental methodologies such as eye-tracking and self-paced reading, and has been used to model a wide array of behavioral data from learning and memory, problem solving and decision making, language and communication, perception and attention, cognitive development, and individual differences (Anderson et al., 2002).

The ACT-R architecture is attractive as a modeling tool for three reasons. First, it is based on a wide array of empirical results in various domains of cognitive psychology. Second, it is flexible enough to permit the modeler to add their own assumptions and theories about the specific task to be modeled. Finally, ACT-R models yield dependent measures such as reading time in much the same way as humans performing the experiment; e.g., the system can easily be programmed to simulate key presses after it processes material presented on the screen.

As shown in Figure 1, the architecture consists of several MODULES such as Declarative, Visual, and Manual. Each module is associated with a BUFFER which temporarily stores information for a given action. For example, the visual buffer is used to store an item “seen” by the system in the environment before it is used in the service of some task.

The module that is especially important for the present paper is the Declarative (henceforth, DM). DM represents permanent memory: every fact that is assumed to be known is encoded as a CHUNK in declarative memory. A chunk is an attribute-value list structure with a special attribute, ISA, which defines its type. The attributes are also referred to as slots. The value of a chunk’s slot is also (by definition) a chunk, unless it is double-quoted or is the

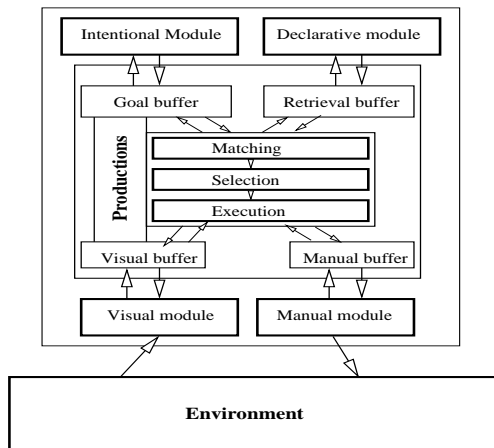


Figure 1: This is a schematic view of the ACT-R system. “Environment” is the outside world that ACT-R is programmed to interact with. The arrows show the possible flows of information. Productions and the central box with the boxes labeled “Matching”, “Selection”, and “Execution” are intended to represent a set of central executive mechanisms and processes.

lisp primitive “nil”.

Each DM chunk has an activation that determines its speed of retrieval, and the probability that it will be retrieved; the initial activation for a given chunk can be set manually.

There is a GOAL BUFFER that holds a current goal under consideration (there can be only one goal at one time); this goal is a chunk with a given type and possibly instantiated slots.

The control structure for modeling a sequence of events is a set of PRODUCTIONS; a production is simply an if-then statement of the following general form: for a given state of one or more buffers and/or DM, execute some actions. Examples of executing actions are retrieving something from DM; changing a value in one of the goal’s slots; repositioning the hand over a keyboard; a visual shift of attention; changing the goal to a new one, etc. If the goal is changed, then this new goal now occupies the goal buffer.

Building an ACT-R model is essentially a definition of possible sequences of actions for a given state of affairs. Events like retrievals from DM are triggered by looking at the contents of one or more buffers. For example, the ACT-R system “sees” an item/object on the screen and then encodes it as a visual chunk. This chunk can then be harvested from the visual buffer; it includes (as slot-value specifi-

cations) information about the content of the item seen, its x-y coordinates, etc. One can define an action based on this information, such as retrieving a chunk from DM.

3 Modeling sentence parsing in ACT-R

Previous research suggests that humans employ some variant of left-corner parsing (see, e.g., (Resnik, 1992)), which in essence involves a bottom-up and a top-down (predictive) step. We adopt this parsing strategy in the simulations. In order to model the prediction of syntactic structure based on incrementally appearing input, we assume that sentence structure templates are available in declarative memory as underspecified chunks. These chunks are retrieved every time a new word is integrated into the structure, as are prior arguments necessary for semantic integration.

We illustrate the parsing process with a simple example (Figure 2). Suppose that the sentence to be parsed is *The girl ran*, and suppose that we are simulating self-paced reading (Just et al., 1982). When the word *the* is seen, a bottom-up and top-down structure building step results in a sentence with an intransitive verb being predicted. This structure becomes the current goal. Then the word *girl* is seen and processed, i.e., its lexical entry is retrieved from declarative memory. The noun slot in the goal is then instantiated with that lexical entry. In the next step, if the word *ran* is seen the relevant lexical item for the verb is retrieved and instantiated with the verb slot of the goal; here, the verb’s argument is also retrieved and integrated with the subcategorization frame of the verb. If, instead of *ran* the word *that* appears, a new goal is created, with any previously instantiated slots of the preceding goal being passed on to the new goal, and parsing proceeds from there.

Each retrieval of a goal from memory results in a surge in its activation, so that repeated retrievals result in increased activation; and the higher the activation of an item the faster it is processed. At the same time, activation decays according to the power law of forgetting (Anderson et al., 2002). In the same way that the goals undergo decay and reactivation, so do the previously seen words. This means that the speed of retrieval of a previously seen argument at a verb will be determined by the activation level of that argument. Thus, the activation of both the goals (predicted structures) and the arguments affect processing.

In our simulations, for simplicity we code in the exact steps that ACT-R takes for particular sentences. Although it is feasible to build a very gen-




Figure 2: A simple illustration of parsing steps in the ACT-R simulations presented.

eral parser in pure ACT-R, before doing this we wanted to first establish whether ACT-R’s reactivation mechanisms can account for a reasonable array of facts from the sentence processing literature. In (Lewis and Vasishth, *An activation-based model of sentence processing as skilled memory retrieval*, (tentative title; in preparation)) we provide a detailed description of a model employing mechanisms similar to those described here, but one that behaves more like a standard parser.

3.1 English subject versus object relative clauses

It is well known (Homes and O’Regan, 1981) that English subject relatives are easier to process than object relatives (1). In the parsing model outlined above, we can model this result without changing any ACT-R parameters at all (i.e., we use the default settings for the parameters).

- (1) a. The reporter who sent the photographer to the editor hoped for a good story.
- b. The reporter who the photographer sent to the editor hoped for a good story.

The explanation comes from the decay of the arguments of the verb *sent*: in object relatives the argument *reporter* decays much more than in the subject relative by the time it is integrated with the verb’s subcategorization frame (Figure 3). This is because more time elapses between the argument being first seen and its retrieval at the verb.¹

¹A reviewer points out that several head-final languages such as German and Dutch also have a subject relative preference and in these languages the activation level cannot be the explanation. We do not claim that decay is the only constraint operating in parsing; frequency effects (greater preference for

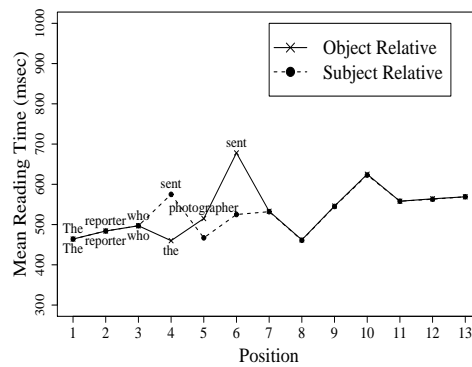


Figure 3: The reading times provided by the model. Retrieval of *reporter* at *sent* is harder in the object relative because of increased argument decay.

3.2 The SC/RC asymmetry in English

It is also well-known (Gibson, 2000) that a sentential complement (SC) followed by a relative clause (RC) is easier to process than an RC followed by an SC:

- (2) a. The fact that the employee who the manager hired stole office supplies worried the executive.
- b. #The executive who the fact that the employee stole office supplies worried hired the manager.

As in the previous discussion about relative clauses, in the harder case the decay of the argument *executive* at the verb *worried* is greater compared to the decay of the argument *employee* at *hired* in the easier-to-process sentence. In addition, the total reading time for the harder sentence is about 120 msec longer.²

3.3 Hindi center embeddings

Previous work (Hakes, 1972), (Konieczny, 2000) has shown that if argument-verb distance is increased, processing is easier at the verb. (Vasishth, more frequently occurring subject relatives) etc. could certainly dominate where the amount of decay is constant in subject and object relatives. It is an open empirical question whether frequency alone can account for the subject/object asymmetry in English, but given that we have independent empirical justification for decay (see Section 3.5), the above is a plausible explanation.

²As a reviewer points out, “the account in terms of activation decay suggests that the SC/RC asymmetry can be annihilated or even reversed by inserting longer or shorter NPs between the critical verbs (worried, hired) and their arguments (executive, employee). This seems unrealistic.” This is surely an empirical question that needs to be verified experimentally; we intend to pursue this very interesting issue in future work.

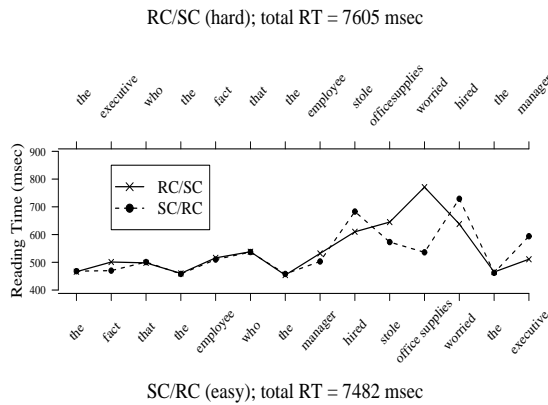


Figure 4: Model’s behavior in the complement-clause/relative-clause contrast.

2003) presented similar results in Hindi. The Hindi experiment manipulated distance by comparing the baseline condition (3a) with the case where an adverb intervened (3b), a verb-modifying PP intervened (3c), and relative clause intervened that modified the preceding NP (3d).

- (3) a. Siitaa-ne Hari-ko Ravi-ko [kitaab-ko Sita-erg Hari-dat Ravi-dat book-acc khariid-neko] bol-neko kahaa buy-inf tell-inf told
‘Sita told Hari to tell Ravi to buy the book.’
- b. Siitaa-ne Hari-ko Ravi-ko [kitaab-ko Sita-erg Hari-dat Ravi-dat book-acc **jitnii-jaldii-ho-sake** khariid-neko] as-soon-as-possible buy-inf bol-neko kahaa tell-inf told
‘Sita told Hari to tell Ravi to buy the book as soon as possible.’
- c. Siitaa-ne Hari-ko Ravi-ko [kitaab-ko Sita-erg Hari-dat Ravi-dat book-acc **ek barhiya dukaan se** khariid-neko] from-a-good-shop buy-inf bol-neko kahaa tell-inf told
‘Sita told Hari to tell Ravi to buy the book from a good shop.’
- d. Siitaa-ne Hari-ko Ravi-ko [kitaab-ko Sita-erg Hari-dat Ravi-dat book-acc **jo-mez-par-thii** khariid-neko] that-was-on-a-table buy-inf

bol-neko kahaa
tell-inf told
‘Sita told Hari to tell Ravi to buy the book that was lying on a/the table.’

In all the “insertion” cases a statistically significant speedup was observed at the verb, compared to the baseline condition.

This experiment’s results were replicated in the ACT-R system; the replication is based on the assumption that the goal (predicted syntactic structure) is reactivated each time it (i.e., the entire predicted structure) is modified. The intervening items result in an extra retrieval compared to the baseline, resulting in faster processing at the verb. In this model, one parameter was changed: the rate of decay of items. We justify this change in the next sub-section.

The modeling results are shown in Figure 5.

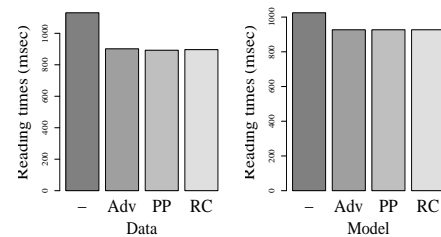


Figure 5: Reading times from data versus model, at the first verb.

3.4 Individual variation in Hindi center embedding data

In the Hindi experiment, there was a further variation in the data when individual subjects’ data were considered: only about 48% of subjects showed a speedup at the verb. About 21% showed a slowdown and there was only a few milliseconds difference (essentially no difference) in the reading times for about 31% of the subjects. The observed variation was a systematic trend in the sense that the 47% of the subjects who showed a speedup or slowdown in adverb-insertion case also showed the *same* trend in the PP- and RC-inserted cases – the probability of this happening is considerably below chance level.

The rate of decay defined in ACT-R’s rehearsal equation can systematically explain this variation. Consider the situation where a chunk i with an initial activation of B_i is retrieved. The activation is

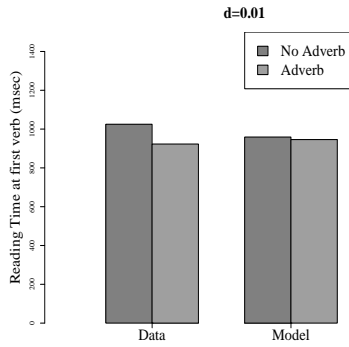


Figure 6: Modeling speedup.

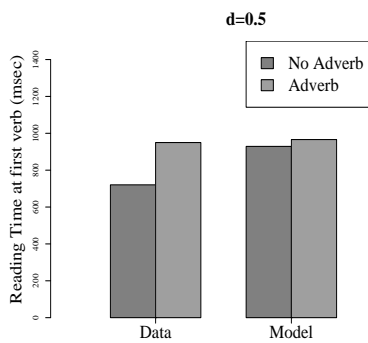


Figure 7: Modeling slowdown.

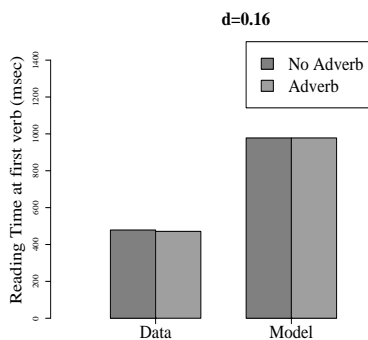


Figure 8: Modeling no difference in reading time.

recalculated each time a retrieval occurs, according to the following equation.

$$(4) \quad B_i = \ln\left(\sum_{j=1}^n t_j^{-d}\right)$$

Here, n is the number of times the chunk i was successfully retrieved, t_j is the time elapsed since the j -th retrieval, and d is a decay rate that defaults to 0.5 in ACT-R. This equation reflects the log odds

that a chunk would reoccur as function of how it has appeared in the past (Anderson et al., 2002, 17).

It turns out that the d parameter take us beyond boolean predictions: $d = .01$ results in a speedup; $d = .5$ results in a slowdown; and $d = .16$ results in no difference in RT at the verb; see Figures 6 to 8.³

3.5 Comparison with other models

The model presented here is very different in conception from existing models of sentence processing. For example, consider Early Immediate Constituents (Hawkins, 1994) and Discourse Locality Theory (Gibson, 2000), two theories with significant empirical coverage. Both theories propose variants of what we will call the distance hypothesis: increasing the distance between arguments and a subsequently appearing verb (head) that selects for them results in increased processing difficulty at the verb. Distance here is quantified in terms of the number of words in a constituent (EIC) or the number of new discourse referents introduced between the arguments and head (DLT).

The present model claims that distance effects are actually a result of argument decay. Preliminary evidence that it is really decay and not EIC- or DLT-defined distance comes from a recent self-paced listening experiment (Vasishth et al., 2004) in which two conditions were contrasted: arguments and verbs with (a) an adjunct intervening, (b) silence:

- (5) a. vo-kaagaz / jisko us-larke-ne / mez
that-paper which that-boy-erg table
ke-piiche gire-hue / dekhaa /
behind fallen saw
bahut-puraanaa thaa
very-old was
'That paper which that boy saw fallen
behind a/the table was very old.'
- b. vo-kaagaz / jisko us-larke-ne /
that-paper which that-boy-erg
SILENCE / dekhaa / bahut-puraanaa
saw very-old

thaa
was
'That paper which that boy saw was
very old.'

In (5), the arguments *kaagaz*, 'paper', and *larkaa*, 'boy' are separated from the verb *dekhaa*, 'saw' by

³Of course, modeling individual variation in terms of differing rates of decay assumes that subjects exhibit varying degrees of decay rates. An experiment is currently in progress that attempts to correlate varying verbal sentence span with subject behavior in the insertion cases.

an adjunct containing two⁴ discourse referents (5a); or by silence (5b). Subjects were allowed to interrupt the silence and continue listening to the rest of the sentence whenever they wanted to. Subjects interrupted the silence (on an average) after about 1.4 seconds.

Distance based theories predict that having an intervening adjunct that introduces discourse referents should result in greater processing difficulty at the verb *dekhaa*, ‘saw’, compared to when silence intervenes. If decay rather than distance is the critical factor here that affects processing, then there should be greater difficulty at the verb in the silence condition than when in the items intervene (see Section 3.3 for why intervening items may facilitate processing). The results support the activation account: introducing silence results in significantly longer reading times at the verb *dekhaa* than when intervening items occur.

4 Conclusion

These modeling efforts suggest that very general constraints on information processing can provide a principled account of parsing phenomena, and also brings human sentence parsing in closer contact with models of human working memory in cognitive psychology (Miyake and Shah, 1999).

There are of course certain potential limitations in the work presented here. Several alternative hypotheses remain to be explored, e.g., the role of competence grammar and its own (possibly theory-internal) operations on processing; the role of experience (Crocker and Brants, 2000), etc. However, the present research is a necessary first step since it provides a basis for such a comparison.

Secondly, there are specific assumptions in the model that may be controversial. For example, we assume that entire sentence structures are predicted as goal chunks, and not verb-types (cf. (Konieczny, 2000)). We are conducting further experiments to explore the predictions made by different assumptions.

Finally, we have used toy simulations to explore the ACT-R constraint-interaction space, the task of scaling up such a model to parse essentially any kind of input is necessary, but still in the future. However, we believe that the results presented are suggestive of the way in which a cognitively-oriented parser could be constructed.

⁴In DLT finite verbs also assumed to introduce a discourse referent.

5 Acknowledgements

We thank the two anonymous reviewers. This research was partly funded by the Sonderforschungsbereich 378 (EM6, NEGRA).

References

- J.R. Anderson, D. Bothell, M.D. Byrne, and C. Lebiere. 2002. An integrated theory of the mind. MS, available from <http://www.act-r.psy.cmu.edu/papers/403/IntegratedTheory.pdf>.
- M. W. Crocker and T. Brants. 2000. Wide-coverage probabilistic sentence processing. *Journal of Psycholinguistic Research*, 29(6):647–669.
- Edward Gibson. 2000. Dependency locality theory: A distance-based theory of linguistic complexity. In Alec Marantz, Yasushi Miyashita, and Wayne O’Neil, editors, *Image, Language, Brain: Papers from the First Mind Articulation Project Symposium*. MIT Press, Cambridge, MA.
- Daniel Grodner and Edward Gibson. 2003. Consequences of the serial nature of linguistic input. MS.
- David T. Hakes. 1972. On understanding sentences: In search of a theory of sentence comprehension. Microfilm, University of Texas, Austin.
- John A. Hawkins. 1994. *A Performance Theory of Order and Constituency*. Cambridge University Press, New York.
- V.M. Homes and J.K. O’Regan. 1981. Eye fixation patterns during the reading of relative clause sentences. *Journal of Verbal Learning and Verbal Behavior*, 20:417–430.
- M. A. Just, P. A. Carpenter, and J. D. Woolley. 1982. Paradigms and processes in reading comprehension. *Journal of Experimental Psychology: General*, 111(2):228–238.
- Lars Konieczny. 2000. Locality and parsing complexity. *Journal of Psycholinguistic Research*, 29(6):627–645.
- Akira Miyake and Priti Shah, editors. 1999. *Models of Working Memory*. Cambridge University Press, New York.
- Philip Resnik. 1992. Left-corner parsing and psychological plausibility. In *Proceedings of COLING*, pages 191–197.
- Shravan Vasishth, Richard L. Lewis, Rama Kant Agnihotri, and Hans Uszkoreit. 2004. Distinguishing distance and decay. Submitted.
- Shravan Vasishth. 2003. Quantifying processing difficulty in human sentence parsing: The role of decay, activation, and similarity-based interference. In *Proceedings of the EuroCogSci conference*, Osnabrueck, Germany.