A Cognitive Model of Coherence-Driven Story Comprehension

Elliot Smith
School of Computer Science, University of Birmingham,
Edgbaston, Birmingham B15 2TT. United Kingdom.
email: e.smith@cs.bham.ac.uk

Abstract

Current models of story comprehension have three major deficiencies: (1) lack of experimental support for the inference processes they involve (e.g. reliance on prediction); (2) indifference to 'kinds' of coherence (e.g. local and global); and (3) inability to find interpretations at variable depths. I propose that comprehension is driven by the need to find a representation that reaches a 'coherence threshold'. Variable inference processes are a reflection of different thresholds, and the skepticism of an individual inference process determines how thresholds are reached.

1 Introduction

Recent research in psychology maintains that comprehension is 'explanation-driven' (Graesser et al., 1994) and guided by the 'need for coherence' (van den Broek et al., 1995). The comprehender's goal is construction of a more-or-less coherent representation which includes explanations for and relations between the story's eventualities. This representation is generated via inferences, which enrich the representation until it reaches the threshold specified by the comprehender's *coherence need* (van den Broek et al., 1995).

By contrast, early models of comprehension emphasised its expectation-driven nature: prediction of future eventualities, followed by substantiation of these predictions (DeJong, 1979). The inference processes described in these early models are still implemented in many contemporary systems.

One problem with these models is their failure to account for experimental evidence about inferences: predictive inferences are not generated at point x in the story, unless strongly supported by the story up to point x (Trabasso and Magliano, 1996); in addition, predictive inferences not immediately confirmed by the story after point x are not incorporated into the representation (Murray et al., 1993). While it is difficult to define 'strong support' or 'confirmation', it is clear that an overly-assumptive model does not reflect mundane comprehension.

A second problem is the failure of these models to account for differential establishment of local and global coherence. Local coherence holds between 'short sequences of clauses', while global coherence is measured in terms of 'overarching themes' (Graesser et al., 1994). McKoon and Ratcliff (1992) maintain that only local coherence is normally established during comprehension (the *minimalist* hypothesis). Others state that readers 'attempt to construct a meaning representation that is coherent at both local and global levels' (the *constructionist* hypothesis) (Graesser et al., 1994). Script-based models allow globally-coherent structures to be constructed automatically, contradicting the minimalist hypothesis; the inclusion of promiscuous predictive inferences also contradicts the constructionist hypothesis.

A third problem is that previous models deny comprehension's flexibility. This issue is sometimes side-stepped by assuming that comprehension concludes with the instantiation of one or more 'primitive' or 'top-level' patterns. Another approach is to apply lower-level patterns which account for smaller subsets of the input, but the aim is still to connect a story's first eventuality to its last (van den Broek et al., 1995).

This paper describes a model which treats inferences as *coherence generators*, where an inference's occurrence depends on its coherence contribution. Unusual inference-making, establishment of local and global coherence, and variable-precision comprehension can be described within this framework.

2 Coherence and Satisficing

A *schema* is any function which maps inputs onto mental representations. It contains slots which can be instantiated using explicit input statements, or implicit statements derived via proof or assumption. Instantiated schemas form the building blocks of the comprehender's representation. A comprehender has available both 'weak' schemas, which locally link small amounts of input (e.g. causal schemas); and 'strong' schemas, which globally link larger sections of input (e.g. scripts).

All schemas generate 'connections of intelligibility' which affect the coherence of a representation (Harman, 1986). Coherence is a common 'currency' with which to measure the benefit of applying a schema. Instead of requiring that a top-level structure be instantiated, the system instead applies schemas to produce a representation of sufficient 'value'. This process can be naturally described as *abduction*, or 'inference to the best explanation' (Ng and Mooney, 1990).

Previous natural-language abduction systems *can* form more-or-less coherent representations: for example, by halting comprehension when assumptions start to reduce coherence (ibid.). However, these systems still have a fixed 'cut-off' point: there is no way to change the criteria for a good representation, for example, by requiring high coherence, even if this means making poorly-supported assumptions. By treating coherence as the currency of comprehension, the emphasis shifts from creating a 'complete' representation, to creating a *satisficing* one. (A satisficing representation is not necessarily optimal, but one which satisfies some minimal constraint: in this case, a *coherence threshold.*)

3 Coherence-Driven Comprehension

In this section, I outline some general principles which may attenuate the performance of a comprehension system. I begin with the general definition of a schema:

$c_1, ..., c_n \rightarrow I.$

where $c_1, ..., c_n$ are the elements connected by I. The left-hand side of a schema is its *condition set*, and the right-hand side represents the *interpretation* of those conditions in terms of other concepts (e.g. a temporal relation, or a compound event sequence). During each processing cycle, condition sets are matched against the set of *observations.*

At present, I am developing a metric which measures *coherence contribution* with respect to a schema and a set of observations:

$C = (V \times U) - (P \times S)$

where C = coherence contribution; V = Coverage; U = Utility; P = Completion; and S = Skepticism. This metric is based on work in categorisation and diagnosis, and measures the similarity between the observations and a condition set (Tversky, 1977).

3.1 Coverage and Completion

Coverage captures the principle of conflict resolution in production systems. The more elements matched by a schema, the more coherence that schema imparts on the representation, and the higher the Coverage. By contrast, *Completion* represents the percentage of the schema that is matched by the input (i.e. the completeness of the match). Coverage and Completion thus measure different aspects of the applicability of a schema. A schema with high Coverage may match all of the observations; however, there may be schema conditions that are unmatched. In this case, a schema with lower Coverage but higher Completion may generate more coherence.

3.2 Utility

The more observations a schema can explain, the greater its coherence contribution. *Utility* measures this inherent usefulness: schemas with many conditions are considered to contribute more coherence than schemas with few. Utility is independent of the number of observations matched, and reflects the structure of the knowledge base (KB). In previous comprehension models, the importance of schema size is often ignored: for example, an explanation requiring a long chain of small steps may be less costly than a proof requiring a single large step. To alleviate this problem, I have made a commitment to schema 'size', in line with the notion of 'chunking' (Laird et al., 1987). Chunked schemas are more efficient as they require fewer processing cycles to arrive at explanations.

3.3 Skepticism

This parameter represents the unwillingness of the comprehender to 'jump to conclusions'. For example, a credulous comprehender (with low Skepticism) may make a thematic inference that a trip to a restaurant is being described, when the observations lend only scant support to this inference. By raising the Skepticism parameter, the system may be forced to prove that such an inference is valid, as missing evidence now decreases coherence more drastically.[1]

4 Example

Skepticism can have a significant impact on the coherence contribution of a schema. Let the set of observations consist of two statements:

enter(john,restaurant), order(john,burger)

Let the KB consist of the schema (with Utility of 1, as it is the longest schema in the KB):

$$enter(Per, Rest), order(Per, Meal), leave(Per, Rest) \rightarrow restaurantvisit(Per, Meal, Rest).$$

In this case, $C = (V \times U) - (P \times S)$, where:

$$Coverage(V) = \frac{ObservationsCovered}{NumberOfObservations} = \frac{2}{2}$$

$$Utility(U) = 1$$

$$Completion(P) = \frac{ConditionsUnmatched}{NumberOfConditions} = \frac{1}{3}$$

$$Skepticism(S) = \frac{1}{2}$$

Therefore, $C = \frac{5}{6}$, with *leave(john, restaurant)* being the assumption. If S is raised to 1, C now equals $\frac{2}{3}$, with the same assumption. Raising S makes the system more skeptical, and may prevent hasty thematic inferences.

5 Future Work

Previous models of comprehension have relied on an 'all-or-nothing' approach which denies partial representations. I believe that changing the goal of comprehension from top-level-pattern instantiation to coherence-need satisfaction may produce models capable of producing partial representations.

One issue to be addressed is how coherence is incrementally derived. The current metric, and many previous ones, derive coherence from a static set of observations. This seems implausible, as interpretations are available at any point during comprehension. A second issue is the cost of assuming various conditions. Some models use weighted conditions, which differentially impact on the quality of the representation (Hobbs et al., 1993). A problem with these schemes is the sometimes ad hoc character of weight assignment: as an antidote to this, I am currently constructing a method for deriving weights from condition distributions over the KB. This moves the onus from subjective decisions to structural criteria.

[1]Skepticism is a global parameter which 'weights' all schema applications. Local weights could also be attached to individual conditions (see section 5).

References

G.F. DeJong. 1979. Prediction and substantiation: A new approach to natural language processing. *Cognitive Science*, 3:251–273.

A.C. Graesser, M. Singer, and T. Trabasso. 1994. Constructing inferences during narrative text comprehension. *Psychological Review*, 101(3):371–395.

G. Harman. 1986. *Change in View*. MIT Press, Cambridge, MA.

J.R. Hobbs, M.E. Stickel, D.E. Appelt, and P. Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142.

J.E. Laird, A. Newell, and P.S. Rosenbloom. 1987. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33:1–64.

G. McKoon and R. Ratcliff. 1992. Inference during reading. *Psychological Review*, 99(3):440–466.

J.D. Murray, C.M. Klin, and J.L. Myers. 1993. Forward inferences in narrative text. *Journal of Memory and Language*, 32:464–473.

H.T. Ng and R.J. Mooney. 1990. On the role of coherence in abductive explanation. In *Proceedings of the 8th AAAI*, pages 337–342, Boston, MA, July-August.

T. Trabasso and J.P. Magliano. 1996. Conscious understanding during comprehension. *Discourse Processes*, 21:255–287.

A. Tversky. 1977. Features of similarity. *Psychological Review*, 84:327–352.

P. van den Broek, K. Risden, and E. Husebye-Hartmann. 1995. The role of readers' standards for coherence in the generation of inferences during reading. In R.F. Lorch, Jr., and E.J. O'Brien, editors, *Sources of Coherence in Reading*, pages 353–373. Lawrence Erlbaum, Hillsdale, NJ.