A production system model of memory for spatial descriptions

Gareth E. Miles (MilesGE@cf.ac.uk)

Stephen J. Payne (PayneS@cf.ac.uk)
School of Psychology, University of Wales, Cardiff, P.O. Box 901, CF1 3YG, UK.

Thom Baguley (T.S.Baguley@lboro.ac.uk)
Department of Human Sciences, University of Loughborough Loughborough, Leicestershire, LE11 3TU, UK.

Abstract

When people read spatial descriptions they construct a mental model. When they attempt to remember the spatial description they may rely on memory for the description itself, memory for the constructed model, and/or memory for the operations used to construct the mental model an episodic construction trace (Payne, 1993). This paper reports an ACT-R simulation of this multiple-representation account of memory for spatial descriptions. The simulation shows that the idea of a remembered construction trace can arise naturally from ACT-R's treatment of goals as declarative memory elements. The simulation captures the most important experimental data in favour of the construction trace hypothesis.

Introduction

Many authors have suggested that comprehending linguistic descriptions involves the construction of a *mental model* ; a representation of the described situation (Bower & Morrow, 1990; Bransford, Barclay, & Franks, 1972; Johnson-Laird, 1983; Tversky, 1991). In Johnson-Laird's theory (which provides the starting point of the current research), a mental model is an analog representation, which shares the structure of the situation it represents.

The mental model theory of comprehension has had a widespread influence in text comprehension and reasoning research. One of the clearest illustrations of the concept of mental models is provided by considering simple spatial descriptions, such as:

> The vest is to the left of the shawl
> The blouse is to the left of the vest
> The coat is below the blouse
> The kilt is to the right of the coat

Mani and Johnson-Laird (1982) asked subjects to read descriptions of this general kind, and then judge whether they were consistent with a diagram of the layout of objects. Later, the same subjects were given recognition tests for the descriptions. Mani and Johnson-Laird found that subjects' gist memory (their ability to distinguish a description of the same layout of objects from a description of a different layout) was better if the original spatial description was determinate (i.e. was consistent with only a single configuration of objects, given assumptions about scale and symmetry, and therefore with a single mental model). However, subjects' verbatim memory (their ability to distinguish the original description from a configuration-preserving paraphrase) was better if the description was indeterminate (i.e. consistent with two different configurations, and therefore two alternative mental models).

Johnson-Laird (1983) gives the above pattern of findings pride of place in his exposition of the theory of mental models. He explains the findings by assuming that memory for a spatial description can involve the propositions of the description or a mental model that was constructed in order to comprehend the description. When the description is determinate, it is easier to construct a mental model, which later supports gist memory. When the description is indeterminate it is harder to construct a mental model, so the participant tends to remember the propositions themselves, and can thus exhibit verbatim memory. Mental models of these simple spatial descriptions are assumed to be structural analogues of the described configurations. Johnson-Laird (1983) describes a computer program for constructing such models from the spatial propositions; the program encodes mental models using a simple 2-dimensional array.

Payne (1993) reported difficulty in replicating Mani & Johnson-Laird's phenomenon, and offered a new account of the long-term memory for spatial descriptions. He accepted that participants construct mental models in order to comprehend the descriptions, but proposed that they remembered not the model itself, but the process of constructing the model - an "episodic construction trace."

Thus, in the case of the description above, a participant might remember beginning the construction of a model by "placing" a vest to the left of a shawl, then adding the blouse to the model, to the left of the vest, then adding the coat below the blouse and finally adding the kilt to the right of the coat. Payne extended Johnson-Laird's array-building program to make this proposal explicit: his program stored a list of propositions encoding each array-building step. In the current example, the construction trace would be represented as follows:

> [start vest shawl left]
> [blouse vest left]
> [coat blouse below]
> [kilt coat right]

The key idea is that the construction trace preserves which objects are used to start building a new model, and, thereafter, it distinguishes the objects that are already in the model from those that are new. (For indeterminate descriptions the trace also encodes the nature of the indeterminacy, see Payne, 1993, for details.)

One of the predictions from the construction trace account is that participants should find descriptions difficult to recognise if the sentences are reordered between presentation and test, even if the participants are told to ignore sentence order.

To see this, consider a re-ordering of the above description:

> The coat is below the blouse
> The kilt is to the right of the coat
> The blouse is to the left of the vest
> The vest is to the left of the shawl

When the participant tries to build a model of this description they must begin by placing the coat and the blouse. The construction trace is now as follows:

> [start coat blouse below]
> [kilt coat right]
> [vest blouse right]
> [shawl vest right]

This trace shares only a single element with the original trace, because the order in which objects enter the model has been disrupted. If, to recognise descriptions, participants construct a model and compare the new construction trace with a remembered trace, then recognition of re-ordered descriptions will be affected.

Payne (1993) reported that re-ordering did indeed suppress recognition memory. Note that the reordered descriptions contain exactly the same propositions, and lead to exactly the same mental model, so that this result is not readily explained by propositional memory or by memory for a mental model. Recently Baguley and Payne (submitted) have provided further support for the episodic construction trace, and extended the re-ordering phenomenon to temporal descriptions. In addition, they report new arguments that the mental model itself is also remembered. They conclude that memory for a spatial description may contain three separate primary levels of representation the verbal propositions, the mental model, and the episodic construction trace.

The aim of this article is to present a production system model of this multiple-representation account of memory for descriptions. The model performs the tasks of comprehending and later recognising verbal descriptions, and it captures the re-ordering phenomenon.

An ACT-R model of the Episodic Construction Trace (ACT-ECT)

ACT-R (Anderson, 1993) is the latest version of Anderson's ACT family of cognitive architectures. In keeping with the others it proposes that a production system operates on a declarative memory. To allow modelling of the time course and stochastic nature of human behaviour ACT-R embodies several assumptions about the activation of elements in declarative memory, how these change over time and how they affect and are affected by processing. These assumptions are justified by rational analysis, so that they work to produce a system that is tuned to the statistics of its environment. In addition to activation-based processing, ACT-R assumes that productions can be selected for use on the basis of the history of their successes and failures and their associated costs.

The explicit division of memory into procedural and declarative memories in ACT-R, and the earlier ACTs, has proven particularly useful when memory tasks are modelled (Anderson, 1983, 1993). In the case of ACT-ECT it allows a clear distinction between over-learned procedures (i.e. those that create and interpret mental models) and transient storage (i.e. the propositions, models and construction traces). ACT-ECT contains productions that create mental models, cope with indeterminate descriptions, discriminate true statements using a remembered model and recognise previously comprehended descriptions by combining evidence from the remembered propositions, mental model and construction trace.

The key decisions that must be made to implement the Episodic Construction Trace (ECT) hypothesis in ACT-R are as follows: How should mental models be represented? How should the ECT be represented as a by-product of the productions that construct such mental models? How should recognition be achieved so as to exploit remembered propositions, mental models and the ECT?

Representing mental models in ACT-ECT

A mental model is treated as a distributed data structure: it is represented by having a working memory element (WME) corresponding to each object in the model. Each such WME has slots for a spatial x and y co-ordinate. This captures the source independence of mental models in that it allows inferences about the relative positions of objects in the mental model irrespective of whether the relations in question have been explicitly noted or encoded in propositional form. Thus a mental model might look like this (sentences leading to construction are shown on right):

Object1 (x = 0, y = 0)	Object2 is to the left of Object1
Object2(x = -1, y = 0)	
Object3(x = 0, y = -1)	Object3 is below Object1
Object4(x = -2, y = 0)	Object4 is to the left of Object2
Object5(x = -1, y = -1)	Object5 is to the left of Object3

If the question 'Is Object5 to the left of Object1?' is asked; then the answer is computed using the co-ordinates in the model representation. Here is an example of the representation of an object-in-a-model:

imag_object8

ISA	imag_object
Type	coat
X_coord	0
Y_coord	-1
Model	model2

The use of generated identifiers facilitates both multiple mental models from the same description and mental models generated from different descriptions that use the same objects. Each Imag_object also contains a slot with a reference to a WME of type model; each model WME is generated at run-time and is used to indicate which imag_objects are in the same model. In theory ACT-ECT could maintain an unrealistically large number of mental models; however the descriptions used in experiments that have investigated spatial mental models do not contain the multiple indeterminacy necessary for this to occur.

Constructing mental models and remembering the construction process

To construct a model from a propositional representation is relatively straightforward. Essentially, we encoded Johnson-Laird's algorithm with a set of productions. There are productions for testing whether either of the objects is already in a model; for creating a new model if neither object is; for adding new objects to models (and, critically, for preparing for this by transforming propositions as necessary so that they express the relation between the new object and the one that is already in the model, instead of the other way around.). Other productions resolve indeterminacy by creating multiple mental models.

The more interesting modelling problem is how to code these construction productions so that they leave a trace of their behaviour (the episodic construction trace) that can later by used in recognition. One rather obvious solution might be for the model-construction productions to not only perform their main function, but to include right-hand side clauses (i.e. actions) that create a special type of representation to record what the production has just done (and that is later used by other productions to recognise this process). This has some similarity to the way the analogy mechanism works in recent versions of ACT-R; where a special 'dependency' representation created in the same way encodes information about what the system is doing for use in later analogy making. However, in the present case such an addition was felt to be arbitrary and unparsimonious. One can use ACT-R as a programming language in this way, to implement the ECT hypothesis, but a preferred solution would be to use ACT-R more as a cognitive theory. After all, why should these particular productions have this special effect on memory?

Our preferred solution relies on the properties of goal handling in ACT-R. Goals are represented as working memory elements (WMEs) that share common properties with other WMEs (e.g. the mental models). One of the properties of ACT-R's WMEs is that they decay, so every separate WME created in the system's past will have some level of activation (and hence retrievability) however minuscule. Hence, all previous goals of the system are maintained in long term memory (eventually decaying to oblivion if they are not reused) and may be treated as data by other productions. A record of the different goals created by the production system therefore provides some record of the processing done. When a new mental model needs to be created, a goal must be set to do this; similarly a goal must be set to rearrange and reprocess a sentence. The memory for these goals can be used to implement the construction trace. ACT-ECT recognises processing steps that have been traversed before by comparing the record of previous goals with the current goals, removing the need for a dedicated memory type.

An important feature of ACT-ECT that enables this scheme to be used is its somewhat unusual representation of goals. Many ACT-R models use distinct memory types for different types of goal. Such a representation implies no special similarity between different types of goal; for instance the goal to read a sentence might be no more related to a goal to write that sentence than to the representation of that sentence:

```
Read_the_sentence
        ISA     Read
        Object  The_sentence

Write_the_sentence
        ISA     Write
        Object  The_sentence

The_sentence
        ISA     Sentence
        Noun1   Tea
        Relation above
        Noun2   Coffee
```

Fig 1: Examples of common notation for two goals and a related memory element in ACT-R

For example in Figure 1 we see three different working memory types (Read, Write and Sentence respectively), and all have one over-lapping element (i.e. the label The_sentence). However in ACT-ECT the relationship between steps involving processing needs to be made explicit, so these elements would instead be represented as shown in Figure 2.

```
Read_the_sentence
        ISA     Goal
        Action  Read
        Object  The_sentence

Write_the_sentence
        ISA Goal
        Action  Write
        Object  The_sentence

The_sentence
        ISA     Sentence
        Noun1   Tea
        Relation Above
        Noun2   Coffee
```

Fig 2: Alternative notation for elements in fig 1 as used by ACT-ECT

Hence in ACT-ECT all working memory elements that act as a processing goal, are of working memory type 'Goal'; their nature is indicated by the slot 'Action', with usually only one other slot needed; i.e. the 'Object' the action is to applied to.

Recognising processing steps

The shared type of all processing goals is important to the recognition productions in ACT-ECT. Only three productions are required; one that recognises propositions seen before (recognise_text), one that recognises mental models constructed before (recognise_model), and the final novel production that recognises the processing steps engaged in previously (recognise_construction_trace). The system agent itself, as implemented in ACT-ECT, does not explicitly set a goal to recognise a model, it merely sets goals to read a test-phase description and create a new model from this description. The afore-mentioned productions match to these processing goals in what amounts to incidental recognition (no explicit goal to recognise a description is necessary). The recognition productions are sufficiently preferred in the conflict resolution process that they will always match first. When any of the three recognition productions fire one of three recognition variables are incremented; *processing* when recognise_construction_trace fires, *gist* when recognise_model fires, and *text_recog* when recognise_text fires. The latter two productions are restricted to firing once for every proposition processed; whilst recognise_construction_trace can fire once per new system goal.

Three different algorithms for calculating the correspondence between two mental models were explored in ACT-ECT. They are essentially very similar and calculate correspondence by seeing if objects are in the same relative positions in the two models; the algorithms only vary in how many objects need to be in the same positions for recognition to occur, with the value in brackets indicating that number. The gist(5) algorithm only increments gist once all five objects are in place; gist(4) is less specific incrementing gist when any 4 objects correspond (not caring about the fifth), and the gist(3) algorithm is the most tolerant. These different algorithms allow a crude exploration of how specific correspondence needs be between two mental models for gist recognition to occur.

The *use* of recognition information is not modelled in ACT-ECT. Currently, ACT-ECT is solely concerned with the collection of recognition evidence. An issue that is therefore left unresolved here is how information from the three different recognition variables should be combined. In the examples below a simple additive algorithm is used to derive the overall recognition score (*Recog_level*). Detailing the use of recognition information , and the weighting of the three identified components of recognition, would require the addition of further assumptions to ACT-ECT.

Simulation runs of the model

Here we provide two examples of the model processing an original description, then attempting to recognise target descriptions. The first of these examples will use a determinate description, the second will extend this to an indeterminate description.

Example 1: Determinate description

The description originally processed by the model was:

[s1] The vest is to the left of the shawl
[s2] The blouse is to the left of the vest
[s3] The coat is below the blouse
[s4] The kilt is to the right of the coat

After creating the mental model associated with this description, ACT-ECT was then given three target descriptions to recognise; original, re-ordered and inferable. The former was a repetition with the sentences in the same order; the re-ordered description was a repetition but in order s3 , s4 , s2 , s1 ; and the inferable description featured none of the original sentences but was consistent with the same mental model. The values of the recognition variables for these three 'conditions' are shown in table 1.

The *recog_level* values in table 1 are the sum of the recognition scores computed from the three kinds of representation (using the gist(3) version of the model-recognition process). Note that a) the re-ordered description suffers relative to the description in the original order (processing recognition value of 2 versus 5, respectively); and b) the inferable description shows better recognition than an unrelated description (notional recog_level = 0 for an unrelated description). c) The three gist algorithms all effectively give the same result in this example, showing no variation between conditions.

Example 2: Indeterminate description

For example 2 the original indeterminate description was as follows:

[s1] The Vest is below the Blouse
[s2] The Blouse is to the left of the Shawl
[s3] The Vest is below the Kilt
[s4] The Coat is to the right of the Vest

Table 1: Recognition values recorded by ACT-ECT for example 1

	Recog level	Processing	Text recog	Gist (3)	Gist (4)	Gist (5)
Original order	12	5	4	3	2	1
Re-ordered	9	2	4	3	2	1
Inferable	3	0	0	3	2	1

Table 2: Recognition values recorded by ACT-ECT for example 2, when both models are abandoned after encountering indeterminacy

	Recog level	Processing	Text recog	Gist (3)	Gist (4)	Gist (5)
Original order	14	6	4	4	2	0
Re-ordered	7	1	4	2	0	0
Inferable	7	3	1	3	2	0
Foil (unrelated)	3	0	1	2	0	0

Table 3: Recognition values recorded by ACT-ECT for example 2, when only one model is maintained after encountering indeterminacy

	Recog level	Processing	Text recog	Gist (3)	Gist (4)	Gist (5)
Original order	16	7	4	5	3	1
Re-ordered	11	3	4	4	2	0
Inferable	9	3	1	5	4	2
Foil (unrelated)	3	0	1	2	0	0

Table 4: Recognition values recorded by ACT-ECT for example 2, when all models are maintained after encountering indeterminacy

	Recog level	Processing	Text recog	Gist (3)	Gist (4)	Gist (5)
Original order	16	7	4	5	3	1
Re-ordered	12	3	4	5	3	1
Inferable	9	3	1	5	4	2
Foil (unrelated)	3	0	1	2	0	0

This description can result in two possible models: either:

Blouse	Shawl		Kilt	
Kilt		or	Blouse	Shawl
Vest	Coat		Vest	Coat

The target descriptions used for this simulation run included a repetition condition, a re-ordered condition (s4, s3, s1, s2), an inferable condition and an unrelated foil condition. The inferable and the foil descriptions shared a single sentence with the original description (s2).

It was necessary to consider the different possible strategies adopted by participants when indeterminacy was encountered (i.e. at s3 above). Participants can either choose to give up mental model creation altogether, continue with only one model (the first computed, we assume), or maintain both possible models; all of these possibilities were simulated, results dependent on the differing assumptions are presented in tables 2, 3 and 4 respectively. As with the determinate descriptions we find ACT-ECT predicting better recognition for the target descriptions presented in their original order; if anything the differences are more pronounced than for example 1 due to the extra processing involved in resolving the impasse caused by an indeterminacy (and therefore the more distinctive construction trace). Note that when both models are abandoned then the sentences following the indeterminacy are not fully processed, so that whilst in the case of the originally ordered description the same sentences will be processed; the re-ordered description will cause the processing of sentences not originally processed and the ignoring of ones that were; hence the large difference in the processing value (see table 2.1).

The superior recognition for the Inferable condition over the foil condition (each of which have one common sentence with the original) is also illustrated under all three indeterminacy coping strategies. An important difference from example 1 is variation in the results from the different gist recognition algorithms. When both models are abandoned then no models ever have 5 objects in, so the gist(5) algorithm will not register (see table 2.1). The gist(5) algorithm is also very sensitive to indeterminacies that occur once all objects are in the model; i.e. when the indeterminacy is caused by the sentence presented last, as in the gist condition above. Another point to note is that the gist(3) algorithm throws up spurious occurrences of gist recognition; notably in the case of the foil above. The foil is designed so the mental model generated does not correspond to that derived from the original description (i.e.

gist recognition should be zero); but the same objects are used and coincidental arrangements of three objects are bound to occur when dealing with indeterminacies (as two models generated from the foil are compared to the two models generated from the original).

Discussion

The episodic construction trace has been proposed as an account of memory for spatial descriptions (Payne, 1993; Baguley and Payne, submitted). The account draws on Johnson-Laird's theory of mental models, and is consistent with a processing view of memory (Crowder, 1993; Kolers, 1973; Morris, Bransford, & Franks, 1977; Weldon & Roediger, 1987).

One of our goals in this project was to better integrate the episodic construction trace with other aspects of cognitive processing and memory. An attractive way to attempt such an integration is to embed the theory in a cognitive architecture, and to this end we have implemented the theory in ACT-R (Anderson, 1993).

Interestingly, the episodic construction trace does not itself require any special encoding, instead it arises naturally from ACT-R's treatment of goals. Because old goals are automatically available as elements of declarative memory, all that is required is that the productions that control recognition attempt to match previous processing goals with current processing goals. This very neat dovetailing between the episodic construction trace theory and ACT-R is to our minds a novel source of evidence in support of both theories.

The ACT-R simulation that uses these principles successfully predicts the most salient experimental data. In future work we plan to compare the simulation's predictions to other aspects of the experimental data. We hope the model will allow us to further explore the interaction of the episodic construction trace theory with pre-existing quantitative theories of recognition memory as embodied in ACT-R.

References

Anderson, J. R. (1983). *The architecture of cognition.* Cambridge, MA: Harvard University Press

Anderson, J. R. (1993). *Rules of the mind.* Hillsdale, NJ: Erlbaum.

Baguley, T. & Payne, S. J. (1998). Long-term memory for spatial and temporal mental models includes construction processes and model structure. *Manuscript submitted for publication.*

Bower, G. H., & Morrow, D. G. (1990). Mental models in narrative comprehension. *Science, 247*, 44-48.

Bransford, J. D., Barclay, J. R., & Franks, J. J. (1972). Sentence memory: a constructive versus interpretive approach. *Cognitive Psychology, 3*, 193-209.

Crowder, R. G. (1993). Systems and principles in memory theory: another critique of pure memory. In A. Collins, M. A. Conway, S. E. Gathercole, & P. E. Morris (Eds.), *Theories of memory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Johnson-Laird, P. N. (1983). *Mental models: towards a cognitive science of language, inference and consciousness.* Cambridge: Cambridge University Press.

Kolers, P. A. (1973). Remembering operations. *Memory and Cognition, 1*, 347-355.

Mani, K., & Johnson-Laird, P. N. (1982). The mental representation of spatial descriptions. *Memory and Cognition, 10*, 181-187.

Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning and Verbal Behaviour, 16*, 519-533.

Payne, S. J. (1993). Memory for mental models of spatial descriptions: an episodic construction trace account. *Memory and Cognition, 21*, 591-603.

Tversky, B. (1991). Spatial mental models. In G. H. Bower (Ed.), *The psychology of learning and motivation: advances in research and theory* (Vol. 27). San Diego: Academic Press.

Weldon, M. S., & Roediger, H. L. (1987). Altering retrieval demands reverses the picture superiority effect. *Memory and Cognition, 15*, 269-280.

Acknowledgments

This research was supported by the United Kingdom Economic and Social Research Council (project grant number R000235641) at the University of Wales, Cardiff.