

Complexity of Structure Mapping in Human Analogical Reasoning: A PDP Model

Graeme S. Halford

Department of Psychology
University of Queensland
Brisbane 4072 Australia
gsh@psych.psy.uq.oz.au

William H. Wilson

School of Computer Science & Engineering
University of New South Wales
Sydney 2052 Australia
billw@cse.unsw.edu.au

Matthew McDonald

Department of Psychology
Murdoch University
Perth 6150 Australia
maf@cs.uwa.oz.au

Abstract

A PDP model of human analogical reasoning is presented which is designed to incorporate psychologically realistic processing capacity limitations. Capacity is defined in terms of the complexity of relations that can be processed in parallel. Relations are represented in the model by computing the tensor product of vectors representing predicates and arguments. Relations in base and target are superimposed. Based on empirical evidence of capacity limitations, the model is limited to processing one quaternary relation in parallel (rank 5 tensor product). More complex relations are processed by conceptual chunking (recoding to fewer arguments, but with loss of access to some relations) or segmentation (processing components of the structure serially). The model processes complex analogies, such as heat flow-water flow, and atom-solar system, while remaining within capacity limitations.

Introduction

The growth of parallel processing models of human analogical reasoning raises the question of the complexity of structures which humans can map in parallel. A good case can be made for parallel processing in analogies (Holyoak & Thagard, 1989) but it is implausible that the most complex analogies can be processed entirely in parallel. The Rutherford analogy between the hydrogen atom and the solar system (Gentner, 1983; Holyoak & Thagard, 1995) entails a very complex set of relationships, and it is far from clear that humans process the entire structure in parallel. If computational models are to be psychologically realistic, a means must be found for quantifying the complexity of structures that can be processed in parallel. It is also necessary to explain how problems that exceed this capacity are processed. The Structured Tensor Analogical Reasoning (STAR) model was designed to incorporate realistic human information processing capacities into a PDP model of analogy. An earlier version of the model (Halford, Maybery, O'Hare, & Grant, 1994) did not handle problems that exceeded human capacity to process structures in parallel. In this paper we present extensions of the model which introduce a number of new processes, and handle more complex tasks.

An analogy is a structure-preserving map from a base or source to a target (Gentner, 1983). The structure of base and target are coded in the form of one or more relations. In the simple proportional analogy cat:kitten::horse:foal the base would be coded as MOTHER-OF(cat,kitten) and the target

as MOTHER-OF(horse,foal). More complex analogies might comprise ternary and higher rank relations, or they might be coded as a higher-order relation which provides an overarching structure, and which has first-order relations as arguments. In general, a first-order relation has objects as arguments, and a second-order relation has first-order relations as arguments, and so on.

Capacity and complexity

The complexity of structures in base or target can be quantified by the complexity of relations, which may be defined by the number of arguments. A binary relation (e.g. BIGGER-THAN) has two arguments, and a ternary relation has three arguments (e.g. LOVE-TRIANGLE is a ternary relation, and has arguments comprising three people, two of whom love a third). Each argument can be instantiated in more than one way. For example, each argument of BIGGER-THAN can be instantiated in an arbitrarily large number of ways (e.g. BIGGER-THAN(horse,mouse), BIGGER-THAN(whale,dolphin) etc.). Consequently, each argument provides a source of variation, or dimension, and thereby makes a contribution to the complexity of the relation. In general, an N-ary relation $R_n(a_1, a_2, \dots, a_n)$ is a subset of the cartesian product $S_1 \times S_2 \times \dots \times S_n$. It is the set of ordered n-tuples $\{(a_1, a_2, \dots, a_n) \mid R(a_1, a_2, \dots, a_n) \text{ is true}\}$. An N-ary relation can be thought of as a set of points in N-dimensional space. Relations of higher dimensionality (more arguments) impose higher processing loads.

We have proposed (Halford et al., 1994; Halford & Wilson, submitted) that processing capacity of higher cognitive processes can be quantified in terms of this *dimensionality* concept. Assessment of the working memory literature, plus some specific experimentation, has led to the conclusion that adult humans can process a maximum of four dimensions in parallel, equivalent to one quaternary relation (Halford, 1993; Halford, et al., 1994; Halford & Wilson, submitted). Structures more complex than this must be processed by either conceptual chunking or segmentation.

Conceptual chunking is recoding a concept into fewer dimensions. For example the ternary relation $R(a,b,c)$ can be chunked to a binary relation $R'(a,b/c)$ by combining b,c into a single argument. The relation between a and b/c can be computed, but the relation between b and c cannot, because they are processed as a single argument. $R(a,b,c)$ can also be chunked to a unary relation, $R''(a/b/c)$ in which a,b,c constitute a single dimension, and relations between them cannot be computed. A relation can also be chunked

to a single entity, in which predicate and argument(s) are not distinguished. In our tensor product representations, this is represented by a single vector. Chunked representations can also be unpacked so as to return to the original relation.

The chunking principles are: (1) a chunk functions as a single entity, predicate or argument, in a relation. (2) no relations can be accessed between items within a chunk. (3) relations between the chunk and other items, or other chunks, can be represented.

Segmentation is decomposing tasks into steps small enough not to exceed processing capacity, as in serial processing strategies.

The Model

An N-ary relation, $R(a_1, a_2, \dots, a_n)$ is a binding between a relation symbol or predicate, R , and the arguments a_1, a_2, \dots, a_n . In the STAR model, relations are represented as the tensor product of vectors representing the predicate and each argument. Thus, given a set of unary relations R_1 on a set A , the relations are represented in a vector space V_{R_1} and the members of A in a vector space V_A , a relational instance $r_1(a)$ (i.e. "r1 is true of a") where $r_1 \in R_1$ and $a \in A$ is represented by a tensor $v_{r_1} \otimes v_a$ in the tensor product space $V_{R_1} \otimes V_A$. Similarly, if R_2 is a set of binary relations on $A \times B$, then the representation space for these binary relations would be denoted $V_{R_2} \otimes V_A \otimes V_B$, and a particular relational instance $r_2(a, b)$ by the tensor $v_{r_2} \otimes v_a \otimes v_b$. This notation extends naturally to any number of arguments – for example, the quaternary relational instance $r_4(a, b, c, d)$ would be represented by $v_{r_4} \otimes v_a \otimes v_b \otimes v_c \otimes v_d$. Conceptual chunking is implemented by convolution of the vectors in the tensor product, in the limiting case to a single vector. Segmentation is implemented by processing one tensor product representation at a time.

The relations in base and target are superimposed on the same tensor product, as shown in Figure 1A, and the mathematical treatment is given in Halford et al. (1994).

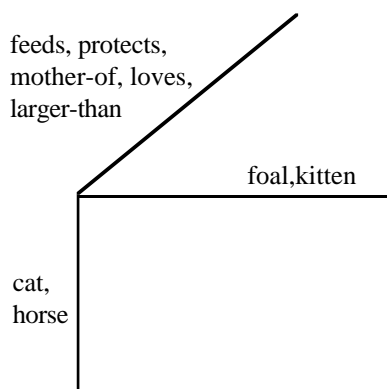


Figure 1A: Tensor product representation

Simple proportional analogy

The representation of base and target in the analogy cat:kitten::horse:foal is shown in Figure 1A. There is a vector representing the predicate MOTHER-OF, and other predicates such as LOVES, FEEDS, PROTECTS, LARGER-THAN. The first arguments of both base and target are superimposed on one vector, and the second arguments on another, as shown. Predicate-argument bindings other than those essential to the problem are represented (e.g. LARGER-THAN(mare,rabbit)) to demonstrate that the model can select the appropriate solution and avoid irrelevancies.

The solution of the problem cat:kitten::mare:? is presented schematically in Figure 1B. In the first step, the input is cat:kitten and the output is all the predicates that have "cat" and "kitten" as arguments (a "predicate bundle"). That is, the output represents the set {MOTHER-OF, LOVES, FEEDS, PROTECTS, LARGER-THAN}.

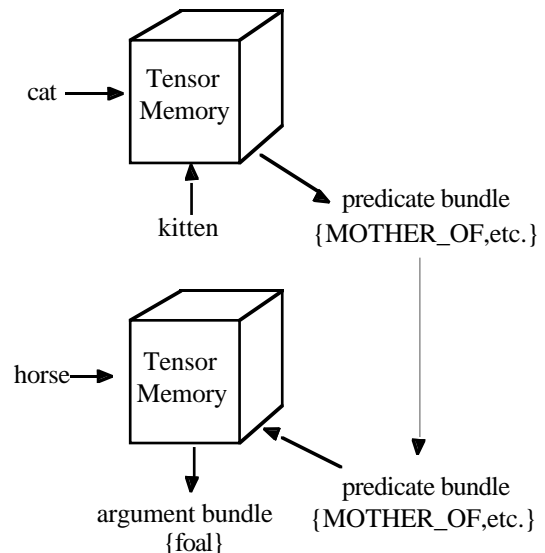


Figure 1B: Processing of a simple analogy.

In the second step, this predicate bundle is used as input, together with "mare". The output is an "argument bundle" representing all possible solutions. The possible solutions can be recognized by computing the inner product of vectors representing each possible solution with the output vector. Alternatively, it can be done by an auto-association technique (Chappell & Humphreys, 1994) which selects the most appropriate solution.

Notice that "foal" is not the only valid solution. For example, "rabbit" is a syntactically correct solution because cat:kitten::mare:rabbit is a valid analogy. The base can be represented as LARGER-THAN(cat,kitten) and the target as LARGER-THAN(mare,rabbit). The preferred solution "foal" can be justified on a number of grounds however. One is the salience of predicates. MOTHER-OF is a more salient predicate relating "cat" and "kitten" than is LARGER-THAN, because of the stronger associations between terms such as "kitten" and the mother-infant relation. Second, the

solution "foal" fits more predicates than does "rabbit". The solution "foal" is consistent with all the predicates MOTHER-OF(mare, foal), LOVES(mare, foal), FEEDS(mare, foal), PROTECTS(mare, foal), LARGER-THAN(mare, foal). However "rabbit" fits only one predicate: LARGER-THAN(mare,rabbit). The model can use both salience and the number of predicates consistent with a solution to produce an analogy corresponding to the one which we would find most satisfying. Therefore the model acknowledges that more than one solution is syntactically consistent, but can distinguish between solutions according to their plausibility.

Analogies with higher rank relations

More complex analogies can be processed. Analogies based on ternary relations can be processed using Rank 4 tensor products, $V_R \otimes V_A \otimes V_B \otimes V_C$. An example would be when premises representing two asymmetric binary relations are mapped into a conventional ordering schema, such as top to bottom. The premises Tom is happier than John, John is happier than Mark might be mapped into a top-down schema in which Tom is in the top position, John in the middle, and Mark in the bottom. The base is MONOTONICALLY-HIGHER(top,middle,bottom) and the target is MONOTONICALLY-HAPPIER(Tom,John,Mark). This can be represented by a Rank 4 tensor product, as shown in Figure 2. Concepts based on quaternary relations (e.g. proportion $a/b = c/d$) can be represented in Rank 5 tensor products in analogous fashion. The representation and processing of other complex concepts in the architecture of this model are discussed elsewhere (Halford, 1993; Halford & Wilson, submitted).

These representations have two important properties, which we call flexibility and derivation. **Flexibility** means that there must be no fixed input or output. To represent the relation $R_n(a_1, a_2, \dots, a_n)$, it must be possible to use the predicate and any $n-1$ arguments as inputs, and compute the remaining argument as output. Similarly, it must be possible to retrieve the predicate, given the arguments as input. Both these functions were illustrated with the simple proportional analogy discussed above, and they are important to analogical reasoning generally. **Derivation** means that it must be possible to derive lower dimensional relations. Given a representation of the n -ary relation $R_n(a_1, a_2, \dots, a_n)$, it must be possible to derive all the $(n-1)$ -ary induced relations $R_{i, n-1}^{n-1}(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$, all the $(n-2)$ -ary induced relations $R_{ij}^{n-2}(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_{j-1}, a_{j+1}, \dots, a_n)$, and so on. This property is also used in analogical reasoning.


Complex Analogies - Serial and Parallel Processing

The main focus of this paper is analogies which are too complex to be completely represented in parallel. We will consider two examples, the analogy between water-flow and heat-flow, and the Rutherford analogy between the structure of the solar system and the structure of the hydrogen atom (Falkenhainer, Forbus, & Gentner, 1989).

A Premises:

Tom is happier than John.
John is happier than Mark.

base schema



target
(problem representation) Tom — happier than — John — happier than — Mark

B




Figure 2. Mapping of ternary relation into ordering schema (A) with tensor product representation (B).

The water flow-heat flow analogy is shown in Figure 3. The figure shows that water-flow is caused by pressure-difference. The component relations are GREATER(PRESSURE(vesselA),PRESSURE(vesselB)) and FLOW(vesselA,vesselB,water,pipe). There is a higher-order relation CAUSE which has pressure-difference and water-flow as arguments. The way this complex structure can be represented in the model without exceeding processing capacity limitations is shown in Figure 3. Pressure-difference and water-flow are each chunked to a single vector, by convolution. Cause is then represented as a binary relation, with the chunked representations of pressure-difference and water-flow as arguments. The model can actively represent the causal relation between pressure-difference and water-flow, but cannot simultaneously represent the structure of the pressure-difference and water-flow concepts, because these are chunked into single entities. These must be unpacked in order for their constituent structure to be represented. However while the constituent structure of pressure-difference and water-flow are being actively represented, the overarching causal relation between them cannot be.

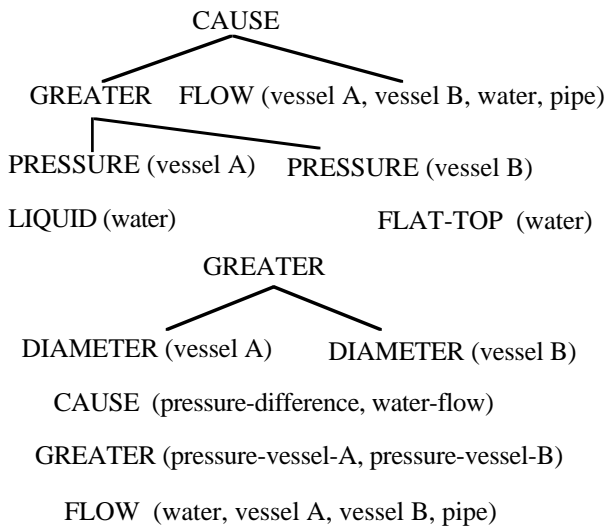


Figure 3. Representation of water-flow.

The Rutherford analogy is represented, as in the model of Falkenhainer et al. (1989) in Figure 4. The orbital motion of a planet around the sun is caused by the difference in mass between planet and sun, and by the gravitational attraction between planet and sun. In analogous fashion, the orbital motion of an electron around the nucleus of an atom is caused by the difference in mass and the electrostatic attraction of the bodies. As with water-flow and heat-flow, the structures are represented in a hierarchy in which higher-order predicates have chunked lower-order relations as arguments. The hierarchy has more levels however than in the case of water-flow and heat-flow. Distracting irrelevant relations, such as the temperature difference between sun and planet, are also represented.

The model works by matching relations in base to relations in target one at a time, staying within the limitation of not representing more than one quaternary relation in parallel. The model can move either up or down the tree looking for matches, which are accumulated in a match matrix, and checked to ensure that the uniqueness and correspondence properties of structure mappings are maintained. The *uniqueness property* means that mappings are one-to-one. The *correspondence property* means that if a predicate P in structure 1 is mapped to a predicate P' in structure 2, the arguments of P are mapped to the arguments of P', and vice versa. Matches which violate these properties are rejected. Each match that is made adds an increment to a structural evaluation score, which is designed to assess the consistency of the matches made.

The operation of the model was also assessed by presenting it with the following base, and testing to see whether a mapping to target 1 or target 2 was preferred.

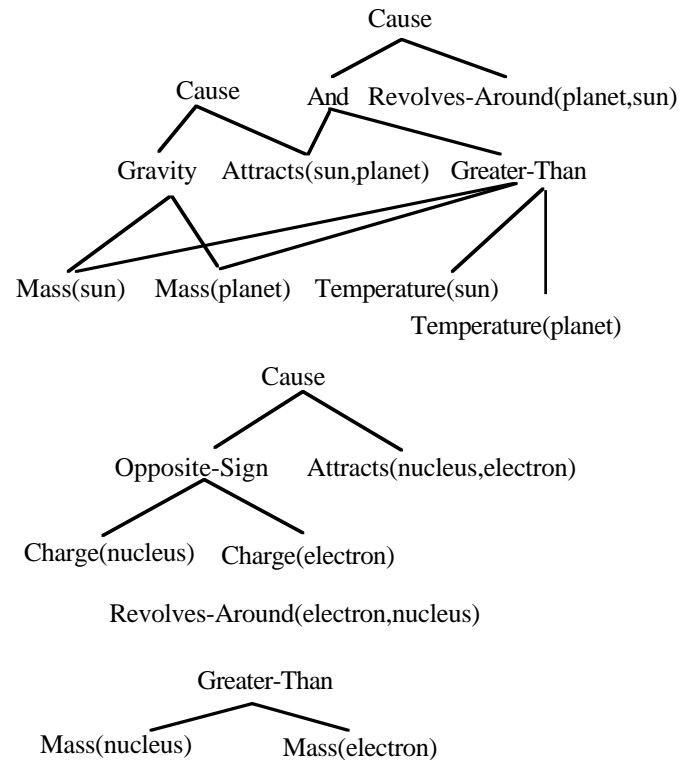


Figure 4: Representation of atom-solar system analogy.

Base:

John is anxious
 John's anxiety is caused by a thesis
 John's anxiety affects his thesis

Target 1:

Joan is anxious
 Joan's anxiety is caused by an exam
 Joan's anxiety affects her sleep

Target 2:

Martin is anxious
 Martin's anxiety is caused by his obesity
 Martin's anxiety affects his obesity

Target 2 is more structurally similar to the base than is target 1. This is reflected in the output of the model, which was able to perform both mappings, but gave a higher structural evaluation score to the mapping between base and target 2.

The representation of each relation in the hierarchy is based on tensor products, and is distributed. However in the present version of the model the movement up and down the tree, from one relation to another, and the compilation of the structural evaluation score, is implemented by a conventional C-program. Means of implementing this aspect in a PDP architecture are being investigated now. However the logic of the model will remain essentially as at present.

The model implies that complex analogies must entail a combination of parallel and serial processing because parallel processing of the entire structure would exceed capacity limitations. The task is segmented into relations, with parallel processing within each relation, but serial processing between relations. This is implemented by coding each relation as a tensor product which binds predicate and argument vectors, and permits predicates and arguments to be recovered. These vectors can be convolved into a single vector, which can be an argument to a higher-order predicate, enabling a hierarchy of relations to be represented. The computational cost of high-rank tensor products provides a natural explanation for the processing load imposed by complex relations. The model can handle complex analogies, such as water flow-heat flow and atom-solar system, while remaining within psychologically realistic capacity limitations.

References

- Chappell, M., & Humphreys, M.S. (1994). An auto-associative neural network for sparse representations: Analysis and application to models of recognition and cued recall. *Psychological Review*, 101(1) 103-128.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41 1-63.
- Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7 155-170.
- Halford, G.S. (1993). *Children's understanding: the development of mental models*. Hillsdale, N. J.: Erlbaum.
- Halford, G.S., Maybery, M.T., O'Hare, A.W., & Grant, P. (1994). The development of memory and processing capacity. *Child Development* .
- Halford, G.S., & Wilson, W.H. (submitted). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology.
- Holyoak, K.J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3) 295-355.
- Holyoak, K.J., & Thagard, P. (1995). *Mental leaps*. Cambridge, MA: MIT Press.