# **Abstraction: Nature, Costs, and Benefits**

Graeme S. Halford

University of Queensland, 4072, Australia

William H. Wilson

University of New South Wales

Steven Phillips

Electrotechnical Laboratory

Tsukuba Science City, Japan

Running head: Relational processing

This work was supported by grants from the Australian Research Council.

Author's address: Graeme Halford, Department of Psychology, University of Queensland, 4072, AUSTRALIA, AIRMAIL. Fax: 61 7 3365 4466.

Email: gsh@psy.uq.oz.au.

## Abstract

It is argued that abstract cognitive processes entail the processing of relations, which differ from more primitive cognitive processes in being more accessible, more flexible, and less content-specific. A relation is a binding between a relation-symbol or predicate, and one or more arguments. Each argument corresponds to a slot which can be filled in a number of ways, so the relation is independent of specific arguments. The binding of arguments to a relation-symbol means that the relation can be an argument to other relations, and is therefore accessible to other cognitive processes based on relations. It can be shown that the building blocks of cognitive processes, such as propositions, and trees can be expressed as relations. Each argument constitutes a dimension in the space represented by the relation, and the number of arguments provides a metric for conceptual complexity. Neural net modelling shows why relational representations impose a processing load which is a function of the number of dimensions.

# Abstraction: Nature, Costs, and Benefits

Though the problem of abstraction in Psychology has been never exactly been a major area of investigation in its own right, it has nevertheless been approached in many ways. I will briefly outline some of the major approaches, then present a perspective that has emerged from our research programme in the last decade. Our particular focus will be on the role of relational knowledge in abstraction. Relations are the core of higher cognitive processes, and the properties of relational processing mechanisms in humans seem a good place to start if we want to eventually uncover fundamental properties of higher cognition. I will quickly review the major approaches to abstraction, and indicate how our perspective can be related to them.

## Induction

Perhaps the most extensive treatment of induction in Psychology is that by Holland, Holyoak, Nisbett and Thagard (1986). In their model, induction is akin to a high level learning process in which condition-action rules are acquired that predict contingencies between events in the world. C-A rules are strengthened when the predictions they make are confirmed, provided the predictions are non-redundant. C-A rules are weakened when the predictions they make are not confirmed. In our laboratory we have applied this model to the induction of relational schemas from relational instances (Halford, Bain & Maybery, submitted). We start with items of the form:

state, operator $\rightarrow$ state

For example:

0,C $\rightarrow$ 1

1,C $\rightarrow$ 2

2,C $\rightarrow$ 3

3,C → 0

The actual items are coded in the form of meaningless materials.  For example, the state elements 0,1,2,3 might be first names of persons, while the operator C might be a geometric symbol, so an actual item might be: "Eileen, triangle → Jenny". The items are meaningless in themselves, but can be assembled into a relational schema. For example, the states can be thought of as distributed around the vertices of a square in the order 0,1,2,3, and the operator *C* can be thought of as a clockwise movement from one vertex to the next. Alternatively, we could think of the states as the integers they are and the operator *C* as "add 1" in modulus addition. Other interpretations are possible also, but the point for the moment is that a set of initially meaningless items can be assembled into a coherent relational schema. The acquisition of such a relational schema depends on its ability to predict future items. For example:

A,# → B

B,# → ?

C,# → ?

D,# → ?

The responses to the second, third and fourth items should be B,D,A respectively. This inference can be made by recognising the correspondence between this set of items and the previous set, so A,B,C,D map to 0,1,2,3 respectively, and # maps to *C*. It is a case of transfer between isomorphs. It is really a case of analogical reasoning, because it entails a structure-preserving map from the first set of elements to the second.

Induction of such a relational schema is itself a form of abstraction, because it effectively transforms content-specific knowledge into content-independent knowledge. The first set of items above, taken individually, are content-specific. Once assembled into a relational schema they entail a structure that is independent of the specific items, and can be transferred to isomorphs. The formation of the relational schema therefore creates knowledge that has a degree of content independence.

Analogy

Analogy has been conceptualised as a structure-preserving map from base to target (Gentner, 1983). Both base and target are coded as sets of relational instances, so relational knowledge is the essence of analogy. Analogical reasoning mechanisms are important to virtually every area of higher cognition, including language comprehension, reasoning and creativity (Holyoak & Thagard, 1995). Human reasoning appears to be based less on an application of formal laws of logic than on memory retrieval and analogy (Halford, 1992; 1995)

Learning sets

Learning sets are a form of abstraction that applies both to humans and other animals, but it has never been satisfactorily explained. It entails training participants on a series of isomorphic problems with different items. For example, participants such as rats, cats, monkeys or children, might be trained to discriminate between a square and a triangle, with triangle positive. Then they are transferred to a new pair of stimuli, such as a circle and square, with square positive, then to a new pair and so on, perhaps for hundreds of problems. The main phenomenon of interest is inter-problem learning, that is more rapid acquisition of later problems, sometimes called a "learning set" or set to learn. Monkeys, apes and humans can reach the point where they solve a new problem in minimum trials. If given two new stimuli A and B, on finding A unrewarded they choose B, and cease to make errors from that point on. This is a form of learning which transcends specific contents. Elsewhere we have suggested that it entails learning the structure of the task (Halford, 1993; Halford et al., submitted). For a two-object discrimination, this entails two elements with an exclusion relation between them, so if A is rewarded B is not and *vice verse*. Furthermore it entails relations between entities on different trials so that, if A is rewarded on trial 1 (or 2) it will be rewarded on subsequent trials, and if B is unrewarded on trial 1 (or 2) it will be unrewarded on subsequent trials. The transfer process that underlies inter-problem learning entails forming a schema that incorporates these relations, and transferring the schema to new problems by analogical mapping. Since learning set acquisition entails transfer between problems that have nothing in common except their structure, it is a form of analogical reasoning *par excellence*. It is also a form of abstraction that depends on relational knowledge.

Variables

The shift from constants to variables is a form of abstraction that has a multitude of applications. We want to draw attention to the fact that understanding variables entails representing relations, because variables are necessarily defined by their relations to other variables or constants. A constant such as "2" can be defined by set reference so that, for example, young children are taught about sets of two elements (2 ears, 2 hands etc.). A variable cannot be defined by set reference but must be defined by an expression relating it to other variables or constants. Thus in $a = 3x$ $a$ is defined by its relations to a constant and another variable. Children learning algebra sometimes have difficulty grasping this, because they think $a$ must refer to something like "apples".

Understanding the role of variables in algebra can be facilitated by using arithmetic as an analog (English & Halford, 1995). For example, the distributive law a(b+C) = ab + bc can be learned by analogy with arithmetic examples such as $3(2+4) = (3 \times 2) + (3\times4)$. The process requires many examples, as well as other procedures, but our point here is that acquisition of an abstract principle, such as the distributive law in algebra, entails learning relations between constants initially, then mapping that schema into an expression that includes variables.

Implicit-explicit transformation

The transformation of knowledge from implicit to explicit is a form of abstraction that has assumed considerable importance in the last decade due to the work of Karmiloff-Smith and her collaborators (Clark & Karmiloff-Smith, 1993). Implicit knowledge enables tasks to be performed, but does not enable a performance to be changed strategically. For example, young children with implicit knowledge can draw a house, but have difficulty drawing a fundamentally kind of house (a “funny house”) because they lack explicit knowledge of what they have done. Implicit knowledge is more accessible to other cognitive processes than implicit knowledge, and can be modified “on-line” without having to be relearned all over again. Elsewhere my colleagues and I have proposed that conversion from implicit to explicit knowledge entails the shift from associative to relational knowledge (Phillips, Halford, & Wilson, submitted).

Evolution and development of abstraction

Abstraction appears to increase through the evolutionary scale, and also with human development. Thus higher animals are considered capable of greater abstraction than lower animals. It might even be considered that abstraction is a uniquely human capability, but evidence of symbolic processing has been obtained from other primates, most notably from chimpanzees (Premack, 1983). Learning set acquisition results in knowledge that some degree of content-independence, as noted above. It occurs in mammals and even, though to a lesser extent, in lower vertebrates. It generally occurs more efficiently in higher animals. If we accept that learning set acquisition depends on acquisition of relational knowledge, the implication is that evolution toward more abstract thought reflects increasing capacity to process relations.

Abstraction is also thought to increase with age, older children being capable of more abstract thought than younger children. Piaget’s (1950) stages of cognitive development represent increasing levels of abstraction, though Piaget emphasised properties such as mobility and reversibility. The levels of relational knowledge to be defined below, correspond in approximate fashion to Piaget’s major stages, as we have indicated elsewhere (Halford, 1993). Therefore development of abstract thought also depends on increase capacity to process relations. We will define levels of processing of relations and indicate how they can be identified with categories of animals on the evolutionary scale, and also with levels of cognitive development in children.

## Abstraction and relational knowledge

We have briefly reviewed a number of abstraction processes in Psychology to illustrate that they all entail processing of relational knowledge in some form. Our work on neural net modelling of such processes as analogy (Halford et al., 1994; 1995), has also made us realise that relations are not easily represented in neural net architectures. Although there are numerous approaches to the problem (Halford et al., 1995; Hinton, 1990; Hummel & Holyoak, In press; Plate, 1995; Shastri & Ajjanagadde, 1993; Smolensky, 1990) it is still a source of controversy. Relational knowledge is incorporated in many kinds of models, so that for example, propositional networks entail labelled links that express relations between entities. It would appear to be useful if we had a set of criteria for relational knowledge, so we know what needs to be incorporated into a model.

It is important to recognise however that relations have not been the main explanatory construct in Psychology, or at least in that part of it that has been influenced mainly by the British Empiricist tradition. The concept of association has been very influential, at least since the time of Aristotle (Humphrey, 1951). Despite decades of controversy between associationists and proponents of various forms of relational knowledge, such as gestaltists and Piagetians, there does not appear to be a generally accepted set of criteria for this distinction. Therefore we will attempt to provide a set of properties that are designed to capture the essence of associative and relational knowledge.

The specific distinction we make between associative and relational processing was first proposed by Phillips, Halford and Wilson (1995). The suggestion is that associative processing is more primitive, less explicit, and less accessible than relational processing. By accessible we mean that it can be operated on by other cognitive processes in the manner which Clark & Karmiloff-Smith (1993) identify with explicit thought. The definitions we provide lead to a continuum of cognitive processes from elemental association, which we consider the most primitive, to the highest levels of relational knowledge that our research indicates humans can process in parallel. We will define a set of points along this continuum. We will show how, as organisms progress from associative to relational processing, and then to structures of higher dimensionality, cognition becomes more content-independent, accessible, and flexible. However while this progression to more complex relations increases the power and flexibility of thought, it incurs the cost of higher processing loads. We will summarize the properties of associative and relational knowledge which have been defined formally by Phillips et al. (1995).

An association is a link between a cue and associate. Activation of the cue evokes the associate, but the reverse is not necessarily the case. Although there is a possibility that backward conditioning may sometime occur, associations are not inherently bidirectional, and acquisition of the reverse association may require additional learning. One of the most important ways that associations differ from relational knowledge is that there is no symbol that represents the associative link. That is, a person can have an association between a cue and an associate without necessarily knowing that the link exists. They know something, but do not (necessarily) know that they know it. Another important property is that an associative link cannot be an element, a cue or an associate, in another association. Associations can of course be chained, so that the output of one association is the input to another. However in this case the associations share a common element, which is output to one and input to the other, but a link in one association cannot be an element in another. That is, an associative link *per se* cannot be a component of another association. There appear to be two distinct types of associations, elemental and configural.

Elemental associations (Rudy, 1991) in this context would comprise links between pairs of entities, of the form:

$entity_1 \rightarrow entity_2$

The best-known example in this century has probably been the association between a stimulus and a response, that is:

stimulus → response

Elemental associations do not require any representation other than input and output, and therefore cannot achieve any abstraction.

Configural associations entail two stimuli each of which modifies the link between the other stimulus and the response. In general, configural associations entail links of the form:

$E_1/E_2 \rightarrow E_3$

Conditional discrimination, in which choice of stimuli is dependent on background, is a well-known instance of configural learning (Rudy, 1991). For example, if the stimuli were triangle and square, on black or white backgrounds, configural discriminations might have the form:

Configuration triangle/black → R+

Configuration square/black → R-

Configuration triangle/white → R-

Configuration square/white → R+

Conditional discrimination is isomorphic to the exclusive-OR (XOR). Configural associations require a representation other than the input and output. This is essentially the same as neural nets that compute the XOR requiring a hidden unit layer. However the representation is content-specific and of limited generality. It cannot support transfer between problem isomorphs (Halford et al., submitted). Therefore configural association achieves only the most minimal abstraction.

A relational knowledge system consists of a set of bindings between a relation symbol or predicate and a set of arguments.[1] For example, the binary relation LARGER-THAN(whale,dolphin) is a binding between the predicate LARGER-THAN and the arguments,

[1] an N-ary relation $R^n(a,b,..,n)$ is a subset of the cartesian product $S_a \times S_b \times .. \times S_n$. It is a set of ordered n-tuples {. . (a,b, . . n)} such that R(a,b, . n) is true.

whale and dolphin. Some properties that distinguish a relational knowledge system from an association are:

Predication. The predicate provides an explicit symbol for the link between the arguments. For example, the predicate LARGER-THAN symbolizes a specific link or relation between whale and dolphin. By contrast, there is no symbol for the cue-associate link in an association. An important effect of this is that a relation can be an argument to another relation. That is, we can have relations between relations. Higher-order relations are relations whose arguments are relations. For example, CAUSE(kiss(John,Jane), hit(Tom,John)). The relation "kiss" between John and Jane is the cause of the "hit" relation between Tom and John. Cause is a higher-order relation, the arguments of which are "kiss(John,Jane)" and "hit(Tom,John)".

By contrast, with associations, a relation can be an argument to another relation. A higher-order relation has other relations as arguments. This means that relational systems can be recursive, whereas associative systems cannot (Phillips et al., submitted). A relation can be an argument to another relation, so hierarchies of relations can be created, comprising higher-order relations whose arguments are lower-order relations. The fact that a relation can be an argument to another relation also means that a relation can be accessible to other cognitive processes.

Omni-directionality. Relational knowledge entails omni-directional access. To take a simple example, it makes no sense to say that we know that a>b, but do not know that b<a. However it does make sense to say that someone has an a->b association but does not have the b->a association. Access to components of relations is even more flexible than this. If we know the arguments "whale" and "dolphin", we can at least retrieve a plausible relation; that is, we can recognize that at least a possible relation is LARGER-THAN. Alternatively, if we have the predicate LARGER-THAN and the argument "whale", we can retrieve a plausible second argument. This is equivalent to asking "name something that a whale is larger-than", to which possible answers are "dolphin", "cat", "person" etc.. Similarly, if we have the predicate LARGER-THAN and the argument "dolphin", we can retrieve a plausible first argument such as “whale” (equivalent to asking what is larger than a dolphin). These questions do not yield unique answers, but they yield sets of plausible answers; for example, the last question might yield answers such as “whale”, “battleship”, “mountain”, but it would not yield “mouse”. We therefore have omni-directional access to representations of relations, in the sense that given any N-1 components, the Nth component can be retrieved, at least in principle. Associations do not necessarily entail this omni-directional access, and may be uni-directional.

There is a well-known paradox concerning the asymmetry of associations. When Bryant and Trabasso (1971) attempted to train children in the relative lengths of sticks as preparation for their classical studies of transitivity, they were puzzled to find that they had to train them both that $a>b$ and $b<a$. We would not normally expect this, because if we know that $a>b$ we must know that $b<a$, so training in the inverse relation should be redundant. No explanation for this paradox appears to have been forthcoming. However from the present perspective it has a clear interpretation: It implies that the children were learning associations rather than relations. This is consistent with more recent work suggesting that quite a lot of empirical work on transitivity in both humans and (other) animals is really associative rather than relational. From our present perspective it is clear that extrapolations cannot be made from one to the other, and the literature on transitivity may need to be re-examined from this point of view.

The fact that relational knowledge is accessible from all directions has always been put to good, if unwitting, use in teaching mathematics. School mathematics depends on two distinct operations, addition and multiplication. Subtraction and division entail accessing addition and multiplication respectively from a different direction. For example, the sum 5 - 3 = ? simply requires us to access the addition operation by asking “what number, added to 3, gives 5?” Notice that children are not taught subtraction and division tables. They are only taught addition and multiplication tables because, once learned, those tables can be accessed to answer subtraction and division questions respectively. The reason is that addition and multiplication entail relational knowledge, both being ternary relations. Addition and multiplication knowledge therefore have the omni-directional access properties of relations.

Argumentation. Relational knowledge systems entail representation of argument roles or slots. Thus representation of the relation LARGER-THAN(-,-) includes representation of the two argument roles, one for a larger and one for a smaller entity. Representation of roles is inherent in relational knowledge. It would make no sense to say that someone knows what LARGER-THAN means but does not know that it relates a larger to a smaller entity. Knowing the roles is part of knowing the relation.

The representation of roles gives relations some degree of content-independence. The relation LARGER-THAN is the same irrespective of specific fillers in the roles. That is LARGER-THAN(whale,dolphin), LARGER-THAN(rat,mouse), LARGER-THAN(mountain,mole-hill) all entail the same relation. Associations do not entail explicit representation of roles, and therefore lack this content-independence.

Compositionality. Relations can be composed of simpler relations. Associations are essentially links between pairs of elements, a cue and an associate, whereas a relation can link an arbitrary number of elements. Associations can be chained, as noted above, and they can also be formed into converging sets, when cues $c_1$-$c_n$ are associated with a single associate a., or into diverging sets, as when a single cue c. is linked to associates $a_1$-$a_n$. However associative systems can be reduced to sets of cue-associate pairs which do not interact with, or modify, each other.

However relations can have arbitrary numbers of arguments. More complex relations can imply less complex relations. For example, the ternary relation MONOTONICALLY-LARGER(a,b,c) can be decomposed into >(a,b), >(b,c), >(a,c). However relations with more than two arguments are simply collections of simpler relations. Some relations with more than two arguments cannot be reduced to binary relations without remainder (Halford et al., submitted).

Associations, relations, and abstraction

The four properties of relational knowledge systems, predication, directionality, argumentation, and compositionality, indicate that relations can be considered to be abstract in ways that associations cannot. Relations have some degree of content-independence, which is one property that has been traditionally associated with abstract knowledge. Furthermore all the data structures that form the building blocks of higher cognition can be expressed as relations (Phillips et al., 1995).

A proposition is an instantiation of a relation. For example, the proposition LARGER-THAN(whale,dolphin) is an instance of the relation LARGER-THAN(-,-), in which the arguments are whale and dolphin. Other structures such as lists and trees can also be represented as relations (Codd, 1990; Phillips et al., 1995). Thus relational systems can be said to underlie higher cognitive processes generally, and are the basis of symbolic cognition.

From a computational perspective, relations are not more primitive than lists or trees (each can be implemented in terms of the other). Hence, the availability of both relation- and list-based programming languages (e.g., SQL and Lisp, respectively). From a psychological perspective, however, as is shown below, relation-based behaviour generalizes association-based behaviour in terms of predication, directionality, argumentation and compositionality. Thus, the distinction between associative and relational processing provides a more natural and parsimonious basis for abstraction in cognitive development.

Associations tend to be more content-specific. That is, an association is a link between two representations each of which has a specific content. There is no predicate symbolizing the link, as there is with a relation, and there is no representation of the roles, independently of the fillers. Associations are really a more primitive, more basic form of cognitive processing.

Our next step is to consider levels of processing within the associative and relational mode. Associative processing can be divided into two levels, elemental and configural, while human relational processing can be divided into four levels, unary to quaternary. Though more complex relations are theoretically possible of course, there is evidence that quaternary relations are the most complex than humans normally process in parallel (Halford, 1993; Halford et al., 1994) This entire classificatory system can be integrated into a single dimension, representational rank, which we will define in the next section.

Representational rank

Seven levels of representational rank, from 0 to 6, are shown in Figure 1. The ranks refer to the number of distinct entities that are related in a representation. We will consider each rank in turn.

Insert Figure 1 here

Rank 0: Elemental associations

At this level there is no representation, other than the output from the perceptual system and the input to the motor system.  Processing depends on links between input and output, without immediate representation. This level corresponds to elemental associations, between a content-specific input representation, such as a representation of a stimulus, and an output representation, such as a programme for a motor or verbal response. In a broad sense it includes reflexes, stimulus-response associations, and perceptuo-motor skills, insofar as they are not guided by a representation of the task. These processes can be captured in neural net models by pattern associators, or two-layered nets, as shown in Figure 1. That is, there is an association between input and output, without internal representation, and the association is either acquired incrementally through the operation of a learning algorithm, or is hard-wired. The acquisition is not influenced by other cognitive processes. This is the most primitive level of functioning, and presumably occurs in all animals that have at least the ganglion level of neural organization.

Rank 1: Configural associations

At this level there is internal representation, but it is undifferentiated, and does not have any of the properties of relational systems outlined above. Many features, and quite a large amount of information, can be represented, but no relations with semantic interpretations are defined between components of the representation. A prototypical example of Rank 1 processing would be configurational learning tasks such as conditional discrimination, transverse pattern and negative pattern discrimination (Halford, 1993; Rudy, 1991). In conditional discrimination, discussed earlier, each stimulus (circle, triangle, white, striped), taken singly, is equally associated with positive or negative responses. However the task could be performed by associations in which the cues are configurations, such as "circle/white", "circle/striped", etc. This performance requires internal representation of the configurations, and there is evidence that rats are capable of it, although it depends on maturation of the hippocampus (Rudy, 1991). Configurational associations are sufficient to permit individual tasks to be learned, but transfer between isomorphic task requires relational representations (Halford, et al., submitted; Halford, 1993)

Another example of this level of performance would be the representation of objects in space in the manner that has been observed in 5 month-old infants by Baillargeon (Baillargeon, 1987). Infants represent features of objects, such as height, and their distance from other objects, but there is no evidence that they can relate such representations to other representations. They can represent perceptible features, but presumably have no explicit representations of relations.

Performances at this level can be captured by three-layered nets, comprising input, output, and hidden layers, as shown in Figure 1. Configural learning tasks such as conditional discrimination can be performed by using units in the hidden layer which represent configurations of features such as "circle&white" (Schmajuk & DiCarlo, 1992), just as hidden units can be used to represent conjunctions in the exclusive-OR task. In spatial tasks, hidden units could represent various features of objects such as height and distance from other objects.

Rank 2: Unary relations

At this rank there is a representation comprising two components that can be related to one another. A prototypical case is the representation of a binding between a predicate and one argument, such as BIG(dog). Therefore a unary relation can be represented at Rank 2, and it is the lowest level at which propositions can be represented.

Psychologically, representation of a unary relation such as BIG(dog) is a binding between the unary relation BIG and one argument. The relation can be interpreted as expressing a state or an attribute. Other cases of propositions with one argument can be interpreted as class membership; e.g. DOG(fido). A binding between a variable and a constant can also be expressed as a unary relation; e.g. HEIGHT(1-metre). This proposition represents a binding between a variable, height, and an argument, which can take a range of values (Halford, Wiles, Humphreys, & Wilson, 1993; Smolensky, 1990)

Halford et al. (1994) have shown that predicate-argument bindings can be modelled using tensor product networks. We are not necessarily committed to this specific formalism, but we are committed to the properties of relational knowledge outlined earlier. We are also committed to representations in which the components remain distinct, though bound together. Thus when we represent a binary relation such as BIGGER(dog, mouse) the three components, BIGGER, dog, and mouse must be bound together, but retain their individual identity.

We will consider representations based on tensor products because they incorporate these properties, and thereby illustrate how they can be incorporated into neural nets. Each tensor product representation comprises a vector representing the predicate, and a vector representing each argument, as shown schematically in Figure 1. The binding units (whose values correspond to the computation of the tensor product) are interpreted as activations or as weights. Where they are interpreted as activations, bindings can be changed dynamically and the relation may be changed in all-or-none fashion without external input. For example we can change HAPPY(John) to SAD(John). One component of the representation (John) remains the same, but when the other component is changed a new binding is formed. This means representations can be changed in the course of reasoning tasks. They are not dependent on incremental change as a function of experience. On the other hand rank 0 and rank 1 neural nets depend on incremental adjustment of weights.

Several important new properties emerge at Rank 2 that are not possible with lower rank representations. These include ability to represent propositions, relations, and variables. This is a sizeable step towards explicit, conceptual thought. Further properties emerge at each of the higher ranks, as we will see.

Rank 3: Binary relations

Representations of this rank have three components, a predicate and two arguments. Binary relations, can be represented, with one component representing the predicate and two components representing the arguments. For example, BIGGER(dog, mouse) has two arguments, which can represent any pair of objects such that the first is bigger than the second. Univariate functions, f(a) = b, and unary operators, {(x, -x)} can also be represented at this level. At Rank 3 more complex variations between components can be represented. The binary relation R(x,y) represents the way x varies as a function of y, and vice versa, neither of which is possible with Rank 2 representations.

Rank 4: Ternary relations

At this rank representations have four components that can be used to represent a predicate and three arguments, so concepts based on ternary relations can be represented. An example of a concept based on ternary relations would be the "love-triangle", in which two persons, x and y, both love a third person, z. Concepts such as transitivity and class inclusion, to be discussed later, entail core representations that are ternary relations (Halford, 1993).

The number of possible relations between elements increases again with ternary relations, and the compositionality property of relations begins to become apparent: R(x,y,z) represents three binary relations, $R_1$(x,y) , $R_2$(y,z) and $R_3$(x,z) , as well as the three-way relation, R(x,y,z), which is not defined in a binary relation. This means that with a ternary relation, but not with unary or binary relations, it is possible to compare x with yz, or y with xz, or z with xy. It thus becomes possible to compute the effects on x of variations in yz, and so on. More complex interactions can be represented than at the lower levels, which increases the flexibility and power of thought.

Bivariate functions and binary operations may also be represented at this level. A binary operation is a special case of a bivariate function. A binary operation on a set S is a function from the set S x S of ordered pairs of elements of S into S; i.e. S x S -> S. For example, the binary operation of arithmetic addition consists of the set of ordered pairs of {. . (3,2,5), . . (5,3,8), , }.

At this level mappings can be made between structures on the basis of formal similarity, independent of content (Halford, 1993). Consider two premises such as “Tom is wiser than Mike” and “Mike is wiser than John”. Suppose we want to make a transitive inference, “Tom is wiser than John”. Based on well validated models of transitive inference reasoning, this entails mapping the premises into an ordering schema such as top, middle, bottom, as shown in Figure 2. We could do this by mapping Tom into top, Mike into middle and John into bottom, as I Figure 2A. However it would be equally valid to invert the map. That is, map Tom into the bottom position, Mike into middle, and John into top, as in Figure 2B. What is NOT valid is to (say) map Mike into top, Tom into middle, and John into bottom, as in Figure 2C. Why is this so? Applying the principles of analogical mapping, the problem is that relations in the top-middle-bottom ordering schema do not consistently correspond to those between the problem elements. As Figure 2C shows, “above” corresponds to “less wise than” on one occasion and to “wiser” on the other, which is inconsistent.

Figure 2 here

Thus mappings that entail more than one relation can be made solely on the basis of consistency criteria, independent of content. This cannot be done with lower level structures (Halford, 1993). This again illustrates the way relations of higher dimensionality are associated with greater abstraction. Psychological realistic mapping models, such as the LISA model of Hummel & Holyoak (in press) have shown that processing relations in parallel is necessary for reliable performance on such tasks.

Rank 5: Quaternary relations

At this rank representations have five components, and quaternary relations, R(w,x,y,z) may be represented, with one component representing the predicate and the other four the arguments. An example would be proportion; a/b = c/d expresses a relation between the four variables a,b,c, and d. It is possible to compute how any element will vary as a function of one or more of the others. With a quaternary relation all the comparisons that are possible with ternary relations can be made, as well as four-way comparisons; the effect on w of variations in x,y,z, the effects on x of variations in w,y,z, and so on.

Quaternary relations can also be interpreted as functions or as operations. A trivariate function is a special case of a quaternary relation. Quaternary relations may be interpreted as a composition of binary operations. For example (a + b) x c = d is a quaternary relation.

Rank 6: Quinary relations

At this level representations have six components, and it is possible to represent a predicate and five arguments. The psychological existence of this level is speculative, and if it exists, it is probably available only for a minority of adults, which makes evidence about it difficult to obtain. It would enable people to reason about relations between systems, each of which is composed of compositions of binary operations. If we assume that theories are composed of mathematical expressions each of which integrates a set of binary operations, then representation of the relation between theories amounts to representing the relation between structures defined by a composition of binary operations. This means that the ability to work within a theory would require rank 5 reasoning, whereas ability to deal with relations between theories would require rank 6. Thus the work of the theoretician, or theory builder, would entail Rank 6 processing.

Processing capacity

There is no mathematical limit to the number of arguments that a relation can have, but there may be psychological limits to the complexity of relations that can be computed. There is evidence that normal human cognition entails representations only up to Rank 5 (Halford, 1993; Halford, et al., 1993). This means that the most complex structure humans can process in parallel is equivalent to one quaternary relation, with a minority of people being able to process quinary relations.

Concepts more complex than Rank 5 are handled by *conceptual chunking* and *segmentation* (Halford, 1993). *Conceptual chunking* means that representations of high dimensionality, represented by tensor products of high rank, are recoded into fewer dimensions, with tensor products of lower rank. An example would be velocity, defined as $v=s/t$ (velocity = distance/time). This is 3-dimensional, but it can be recoded into one dimension (as occurs when we think of speed as the position of a pointer on a dial). However this does not mean that all processing loads can ultimately be reduced to that for a 1-D concept, because conceptual chunking results in loss of access to some relationships; e.g. when velocity is chunked as one dimension, the 3-way relation between *v*, *s*, and *t* is no longer represented, so changes in *v* as a function of *s* and/or *t* cannot be computed without returning to the 3-dimensional representation, which entails the higher processing load.

*Segmentation* means that complex tasks can be decomposed into smaller segments that can be processed serially. The development of serial processing strategies permits complex tasks to be performed without exceeding limits on the amount of information that can be processed in parallel. It might be asked why all tasks cannot be segmented into steps small enough to avoid capacity limitations. There are some relations that are indecomposable, meaning that they cannot be decomposed into simpler relations and then recomposed into the original relation, without loss. However, even when relations are decomposable in principle, it may not be a Psychologically realistic option, because people might not have the necessary strategies. Part of the difficulty is that, in order to develop strategies, a concept of the task is needed, which requires representation of the essential structure.

We can illustrate the problem with transitive inference. As mentioned earlier, there is extensive Psychological evidence that it entails mapping premises into some kind of ordering schema, and this entails a processing load. The problem is, as Halford (1993) has shown, that both premises have to be processed to determine the correct mapping. Returning to the transitive inference discussed earlier with respect to Figure 2, if we consider only the premise “Tom is wiser than Mike”, we know that Tom should go into top or middle position, but we do not know which. If we consider only the premise “Mike is wiser than John”, we only know that John should to into the middle or bottom position, but again we cannot decide which. Both premises are required in order to map the task correctly into the schema. Consistent with this, processing load in premise integration is higher than in other components of the task (Maybery, Bain & Halford, 1986).

On the other hand there is no task that cannot be decomposed into serial steps, and there are certainly ordering algorithms that could perform this task one step at a time. There is however no evidence that humans naturally employ those algorithms. They can work out serial steps if given appropriate experience, but it usually entails quite a lot of trial and error, and is guided by a concept of the task, which itself entails representing the relations. For example in the model of Halford et al. (1995) strategy development in transitive inference is guided by a concept of order that comprises an ordered set of at least three elements. Thus humans apparently need to process enough information in parallel to represent the essential structure of a task. The limitation here is that humans are largely self-programming organisms, and artificial intelligence models that assume a serial strategy is provided by an external agent or programmer are simply inappropriate.

Thought and rank of representation

Rank of representation is a measure of the power and flexibility of thought. Higher rank representations permit higher levels of relations to be represented, and permit more of the structure of a concept or situation to be processed in thought. Rank is analogous to the number of facets of a situation that can be viewed simultaneously.

Rank of representation has been associated with phylogenetic and ontogenetic development. We will briefly review evidence of the type of representation that is within the capacity of each level in the evolutionary scale, and at each age-range in human development.

Representational rank of animals and children

Levels of representation equivalent to those defined here have been related to species differences (Halford, Wilson & Phillips, submitted; Holyoak & Thagard, 1995). Rank 0, equivalent to simple associations, appears to be possible for all species of animals that have ganglia, albeit with varying degrees of efficiency. Mammals appear capable of Rank 1, as indicated by the fact that rats are capable of configural learning (Rudy, 1991). Rank 2 appears to be possible for monkeys, because they can recognize the binding between attributes and objects. For example they can learn to choose which of two objects is like a sample. If they are shown, say, an apple as sample object, and required to choose between an apple and (say) a hammer, they can learn to choose the apple. An individual task requires only content-specific information, and could be acquired by associative learning, which is Rank 0. However at least some monkeys can transfer the principle to a new task in which (say) the sample is a hammer, and the choices are a banana and a hammer. Transfer of the principle implies they can represent a binding between an attribute (hammer-like) and a specific object (Halford et al., submitted). This can be represented by a Rank 2 tensor product. They can represent a dynamic binding between an attribute and an object.

Chimpanzees appear capable of representing a relation between objects, which implies Rank 3 representations, although this ability has only been demonstrated for language-trained chimpanzees (Halford et al., submitted; Premack, 1983). For example they can learn that if the sample is XX, they should choose AA rather than BC. If the sample is XY, they should choose BC rather than AA. They learn to choose a pair of objects that has the same relation as the sample. The animals are sensitive to the fact that an attribute, or a relation, in the sample is the same as in the correct choice item. This implies they can represent binary relations, a Rank 3 representation. Such performances are really a form of analogical reasoning (Halford et al., submitted; Holyoak & Thagard, 1995).

Chimpanzees can also learn to recognize that a key represents the difference between a closed and an open lock, or that a knife represents the difference between a cut and an intact apple. This not only demonstrates understanding of binary relations, but exemplifies the flexibility of access noted earlier. That is, for chimpanzees a key symbolizes the relation between closed and open, indicating that the key has the role of predicate representing the relation. They can also recognize the state that will result from applying a key to (say) a closed lock, which is equivalent to retrieving the second argument, given the predicate and first argument.

This interpretation is borne out by studies of what has been called reversal learning (Bitterman, 1960). Suppose an animal is trained in a discrimination in which (say) circle is positive, and triangle is negative. Then, without any warning, the contingency is reversed, so triangle is now positive and circle negative. Then when that is learned the reversal process is repeated, so circle becomes positive again, and triangle negative, and so on. The question of interest is whether animals learn to make successive reversals more and more quickly, indicating that they "catch on" to the reversal principle. Species differ considerably in the rate at which they learn to do this. Apes will learn to make a reversal immediately the contingency is reversed, and monkeys also learn, but at a slower rate, while other mammals learn more slowly still. Birds, reptiles and amphibians show less benefit from training with successive reversals, while fish show no benefit at all. This performance clearly differentiates between species at different levels of the phylogenetic scale. As with other tasks that make this differentiation, it requires representation of the relations in the task, together with dynamic binding. That is, it requires recognition that there is always one positive and one negative stimulus, together with the ability to change the mapping of positive/negative into the stimuli dynamically. The higher animals can do this, but the lower animals can only relearn the discrimination incrementally. They are not capable of dynamic change.

It appears then that ability to represent relations is a factor which differentiates species at different points in the phylogenetic scale. Putting it another way, the capacity to represent relations is a recent evolutionary development, which is present only in the highest non-human animals, and then in rudimentary form.

Development of abstraction in children

Children's representational and reasoning abilities have been reviewed elsewhere (Halford, 1993), but will be summarized here. Neonates are capable of Rank 0 associative learning, but there is evidence that Rank 1 representations are possible at 5 months. This is indicated by infants' awareness of properties of recently vanished objects (Baillargeon, 1987). Rank 2 representations are possible at approximately one year, when infants become capable of treating hiding place of an object as a variable, and recognize category memberships explicitly, which in turn leads to their understanding referents of words; for example "Doggy gone" can be represented as GONE(doggy), which is a predicate with one argument.

Rank 3 representations appear to be possible at age two, because children of this age can make discriminations based on explicit representation of binary relations. One consequence is that they can perform proportional analogies of the form A:B::C:D, provided the content is familiar to them.

Rank 4 representations become possible at approximately age 5, and they open up a wide range of new performances, including transitive inference, hierarchical classification and inclusion, certain kinds of hypothesis testing, and many others. Transitivity entails a ternary relation, because premises R(a,b) and R(b,c) are integrated into the ordered triple, monotonic-R(a,b,c). Class inclusion and the part-whole hierarchy are essentially ternary relations. A class inclusion hierarchy has three components, a superordinate class, a subclass and a complementary class; e.g. fruit, divided into apples and non-apples. Part-whole hierarchies are similar, and comprise a whole divided into a part and a complementary part. Though the age of attainment of these concepts has been controversial, it appears that children have difficulty processing tasks that entail ternary relations until approximately five years. This issue has been discussed in greater detail elsewhere (Halford, 1992; 1993).

We have developed a computational model (Halford, et al., 1995) which simulates the acquisition of transitive inference strategies in children and adults. The development of adequate strategies depends on ability to represent an ordered set of three elements. Without this ability, inadequate strategies result. To acquire such a strategy spontaneously, as most children do, is undoubtedly creative, but it depends on ability to represent ternary relations.

Representations of Rank 5 are typically understood at about age 11, as evidenced by understanding of proportion and a number of other concepts, including understanding of the balance scale. This entails representing the interaction of four factors, weight and distance on the left and weight and distance on the right.

Cost of abstraction

Capacity to represent relational concepts increases with phylogenetic level, and also with age in childhood. The reason is that the magnitude of the processing load imposed by relational processes depends on the complexity of the relations processed in parallel. We have shown elsewhere (Halford, 1993; Halford, et al., submitted) that the number of entities related, or number of arguments of a relation, makes a good metric for conceptual complexity. Each argument of a relation is a dimension of variation. That is, a relation with N arguments corresponds to a set of points in N-dimensional space, so the dimensionality of a concept increases with the number of arguments. Therefore quaternary relations impose higher processing loads than ternary relations, which impose higher loads than binary and so on.

The reason for the relation between processing load and number of entities related becomes apparent from distributed representation models. In our model (Halford, et al., 1994; 1995), it is due to the number of binding units being an exponential function of the number of arguments. Other models may yield different functions relating relational complexity to processing load, but there does not appear to be any disagreement that load is a function of relational complexity (Hinton, 1990; Hummel & Holyoak, In press; Plate, 1995; Shastri & Ajjanagadde, 1993; Smolensky, 1990). This theoretical relationship is supported by empirical evidence that processing load is a function of the complexity of relations being processed (Halford, Maybery, & Bain, 1986; Maybery, Bain, & Halford, 1986; Posner & Boies, 1971). Therefore one of the conclusions to which our experiments and attempts at theory construction have led us is that associative processing is not noticeably capacity limited, but there are rather serious capacity limitations on relational processing.

Associative processing is rapid, effortless and intuitive, but it is not accessible to other cognitive processes, and cannot be modified without external input. It is therefore like implicit processing (Clark & Karmiloff-Smith, 1993) in that it represents performance which cannot be modified by strategic cognitive processes. Relational processing, by contrast, imposes high processing loads, the magnitude of which is a function of the complexity of relations that have to be processed in parallel. A model of implicit and explicit processing based on the distinction between associative and relational knowledge is being developed (Phillips, Halford & Wilson, submitted). Because of limitations in the amount of relational information that can be processed in parallel, tasks must be segmented into components small enough not to exceed processing capacity. This means that complex relational tasks necessarily entail a lot of serial processing, which explains why they are slower than associative tasks.

Relational cognition has all the power, flexibility, and abstractness of human symbolic processing. Because of this, relational processes can be modified strategically, whereas associative processes can only be modified through external input. However these benefits come at the cost of higher processing loads, which necessitate the use of procedures such as conceptual chunking and segmentation, which in turn force us into serial processing. This in turn forces us into the labour of acquiring serial processing strategies, and explains the enormous dependence on expertise that is characteristic of higher cognitive processes.

References

Baillargeon, R. (1987). Young infants' reasoning about the physical and spatial properties of a hidden object. Cognitive Development, 2, 179-200.

Bitterman, M. E. (1960). Toward a comparative psychology of learning. American Psychologist, 15, 704-712.

Bryant, P. E., & Trabasso, T. (1971). Transitive inferences and memory in young children. Nature, 232, 456-458.

Clark, A., & Karmiloff-Smith, A. (1993). The cognizer's innards: A psychological and philosophical perspective on the development of thought. Mind and Language, 8(4), 487-519.

Codd, E. F. (1990). The Relational Model for Database Management: Version 2. Addison-Wesley.

English, L. D., & Halford, G. S. (1995). Mathematics education: Models and processes. Hillsdale, NJ: Erlbaum.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. Cognitive Science, 7, 155-170.

Halford, G. S. (1992). Analogical reasoning and conceptual complexity in cognitive development. Human Development, 35, 193-217.

Halford, G. S. (1993). Children's understanding: the development of mental models. Hillsdale, N.J.: Erlbaum.

Halford, G. S. (1995). Commentary on Moshman (1995). Human Development, 38, 65-70.

Halford, G. S., Bain, J. D., & Maybery, M. T. (submitted). Schema induction from isomorphic problems: Common processes in analogy and learning set acquisition.

Halford, G. S., Maybery, M. T., & Bain, J. D. (1986). Capacity limitations in children's reasoning: A dual task approach. Child Development, 57, 616-627.

Halford, G. S., Smith, S. B., Dickson, J. C., Maybery, M. T., Kelly, M. E., Bain, J. D., & Stewart, J. E. M. (1995). Modelling the development of reasoning strategies: The roles of analogy, knowledge, and capacity. In T. Simon & G. S. Halford (Eds.), Developing Cognitive Competence: New Approaches to Cognitive Modelling Hillsdale, NJ: Erlbaum.

Halford, G. S., Wiles, J., Humphreys, M. S., & Wilson, W. H. (1993). Parallel distributed processing approaches to creative reasoning: Tensor models of memory and analogy. In T. Dartnall, S. Kim, & F. Sudweeks (Ed.), AI and creativity: Proceedings of the AAAI Spring Symposium.

Halford, G. S., Wilson, W. H., Guo, J., Gayler, R. W., Wiles, J., & Stewart, J. E. M. (1994). Connectionist implications for processing capacity limitations in analogies. In K. J. Holyoak & J. Barnden (Eds.), Advances in connnectionist and neural computation theory, Vol. 2: Analogical connections (pp. 363-415). Norwood, NJ: Ablex.

Halford, G. S., Wilson, W. H., & Phillips, S. (submitted). Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology.

Hinton, G. E. (1990). Mapping part-whole hierarchies into connectionist networks. Artificial Intelligence, 46, 47-75.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. (1986). Induction: Processes of inference, learning and discovery. Cambridge, MA: Bradford Books/MIT Press.

Holyoak, K. J., & Thagard, P. (1995). Mental leaps. Cambridge, MA: MIT Press.

Hummel, J. E., & Holyoak, K. J. (In press). From analogy to schema induction in a structure-sensitive connectionist model. In T. Dartnall & D. Peterson (Eds.), Creativity and computation : MIT Press.

Humphrey, G. (1951). Thinking: An introduction to its experimental psychology. London: Methuen.

Humphreys, M. S., Bain, J. D., & Pike, R. (1989). Different ways to cue a coherent memory system: A theory for episodic, semantic and procedural tasks. Psychological Review, 96(2), 208-233.

Maybery, M. T., Bain, J. D., & Halford, G. S. (1986). Information processing demands of transitive inference. Journal of Experimental Psychology: Learning, Memory and Cognition, 12, 600-613.

Phillips, S., Halford, G. S., & Wilson, W. H. (1995). The processing of associations versus the processing of relations and symbols: A systematic comparison. In J. D. Moore & J. F. Lehman (Ed.), Annual Conference of the Cognitive Science Society, (pp. 688-691). Pittsburgh, Pennsylvania.

Phillips, S., Halford, G. S., & Wilson, W. H. (submitted). Representational redescription: From associative to relational systems.

Piaget, J. (1950). The psychology of intelligence. (M. Piercy & D. E. Berlyne, Trans.) London: Routledge & Kegan Paul, (Original work published 1947).

Plate, T. A. (1995). Holographic reduced representations. IEEE Transactions on Neural Networks, 6(3), 623-641.

Posner, M. I., & Boies, S. J. (1971). Components of attention. Psychological Review, 78, 391-408.

Premack, D. (1983). The codes of man and beasts. The Behavioral and Brain Sciences, 6, 125-167.

Rudy, J. W. (1991). Elemental and configural associations, the hippocampus and development. Developmental Psychobiology, 24(4), 221-236.

Schmajuk, N. A., & DiCarlo, J. J. (1992). Stimulus configuration, classical conditioning, and hippocampal function. Psychological Review, 99, 268-305.

Shastri, L., & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. Behavioral and Brain Sciences, 16(3), 417-494.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. Artificial Intelligence, 46(1-2), 159-216.

## Figure captions

Figure 1. Representational rank, with typical cognitive process, and schematic neural net implementation.

Figure 2. A. Consistent mapping of premises into ordering schema. B. Alternate, consistent mapping of premises into ordering schema. C. Inconsistent mapping of premises into ordering schema.