

Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology

Graeme S. Halford

*Department of Psychology, University of Queensland, Brisbane,
Queensland 4072, Australia*

gsh@psy.uq.edu.au

www.psy.uq.edu.au/people/Department/gsh/index.html

William H. Wilson

*Department of Computer Science & Engineering, University of New South
Wales, Sydney, New South Wales, Australia*

billw@cse.unsw.edu.au

www.cse.unsw.edu.au/~billw

Steven Phillips

*Information Science Division, Electrotechnical Laboratory,
Tsukuba 305, Japan*

stevep@etl.go.jp

www.etl.go.jp/etl/ninchi/stevep@etl.go.jp/welcome.html

Abstract: Working memory limits are best defined in terms of the complexity of the relations that can be processed in parallel. Complexity is defined as the number of related dimensions or sources of variation. A unary relation has one argument and one source of variation; its argument can be instantiated in only one way at a time. A binary relation has two arguments, two sources of variation, and two instantiations, and so on. Dimensionality is related to the number of chunks, because both attributes on dimensions and chunks are independent units of information of arbitrary size. Studies of working memory limits suggest that there is a soft limit corresponding to the parallel processing of one quaternary relation. More complex concepts are processed by “segmentation” or “conceptual chunking.” In segmentation, tasks are broken into components that do not exceed processing capacity and can be processed serially. In conceptual chunking, representations are “collapsed” to reduce their dimensionality and hence their processing load, but at the cost of making some relational information inaccessible. Neural net models of relational representations show that relations with more arguments have a higher computational cost that coincides with experimental findings on higher processing loads in humans. Relational complexity is related to processing load in reasoning and sentence comprehension and can distinguish between the capacities of higher species. The complexity of relations processed by children increases with age. Implications for neural net models and theories of cognition and cognitive development are discussed.

Keywords: capacity limits; chunking; cognitive development; comparative psychology; complexity; neural nets; relations; representations; resources; working memory

1. Introduction

We propose that information processing capacity limits in humans and higher animals should be defined not in terms of the number of items but in terms of the complexity of relations that can be processed in parallel. We will try to show that relational complexity covaries systematically with processing load in human problem solving, with age differences in children’s understanding of concepts and capacity differences between higher animal species. Relational complexity is not the only factor that influences difficulty; domain expertise, skill with problem-solving heuristics, memory availability, and perceptuomotor factors are also

important, but relational complexity can help explain some established findings in a wide range of literature. We will also explore processing limits in neural net models of cognition. We begin with a brief review of previous approaches to working memory and then present our own formulation.

1.1. Processing capacity and working memory

Processing capacity has often been treated as “working memory capacity,” which is defined as information stored for later processing (Hitch 1980). Capacity limits have been defined in terms of the number of items or units of information (Miller 1956). The storage and processing functions

of working memory are partly distinct, however, because short-term storage of information does not necessarily interfere with concurrent processing (Baddeley 1986; 1990; Baddeley & Hitch 1974; Halford 1993; Halford et al. 1984; 1994; Klapp et al. 1983). For this reason, Baddeley (1986) postulated three systems, a visuospatial scratchpad, a phonological loop, and a central executive. The first system is concerned with the storage of visual – imaginal information, the second with short-term serial recall or “memory span” tasks; actual processing is the function of the central executive.

The distinction between information stored for later processing and information that is actually being processed can be illustrated by the task of mentally adding 79 and 86. The storage of a partial result for later processing in this task is legitimately called “working memory,” but it is distinct from the actual processing, in which information constrains a decision. For example, when we ask “what number results from adding 9 to 6?” we are not merely storing the addends but are using them to carry out a computation; they constrain our decision.

Working memory capacity has also been defined in terms of limits on activation (Anderson et al. 1996; Just et al. 1996), but this does not provide a general metric for pro-

cessing complexity. Just et al. (1996) assess complexity in terms of the number of new goals generated, and Case (1985; 1992) uses a similar metric based on the number of embedded goals in a control structure. As will be explained in section 6.1.3, the relational complexity metric can subsume the levels of embedding metric and applies more widely. Working memory capacity in children has been assessed using tests that combine processing and storage (Case 1985; 1992; Pascual-Leone 1970), but this makes it difficult to determine whether successful prediction is due to the processing or the storage components of working memory. Anderson et al. (1996) assess processing complexity in terms of the number of symbols in an equation, but this may not reflect the difficulty of the underlying processes and does not provide a metric applicable to other types of tasks.

Storage complexity can be measured relatively directly as the number of items or chunks stored in memory span tests. The computational complexity of algorithms can also be measured (Garey & Johnson 1979; Tsotsos 1990), but (as will be discussed in sect. 5) it is not clear that this can be translated into a metric for processing complexity in human cognition. The problem can, however, be approached in the following way. Any cognitive process can be expressed as a function that transforms input(s) to output(s). The capacity to perform the transformation corresponds to the capacity to compute this function. A function is a special case of a relation (see sect. 2.3.2), so relational complexity is a potential measure of processing demands. With processing capacity defined in this way, the limiting factor is not merely the number of items or the amount of information but the relations between entities. This was first realised from a review of the empirical cognition literature (discussed in sects. 3 and 6). However, two approaches to modeling higher cognitive processes in neural nets (Halford et al. 1994; Shastri & Ajjanagadde 1993a) have independently identified a limitation consistent with this one.

We are concerned with capacity in higher cognitive processes, including reasoning, memory operations, and language comprehension. Processes such as vision, which are performed by modules, have higher processing capacity but are specialised for particular types of input, such as retinal or cochlear input (Fodor 1983; Fodor & Pylyshyn 1988). Modular processes do not impose measurable processing demands (as defined in sect. 2.1), and are not associated with individual differences in intelligence (Anderson 1992). Modular processes cannot be greatly modified by higher cognitive processes, so thought cannot be used to “reprogram” the visual system. By contrast, higher cognitive processes can be modified “online,” and the way in which a task is performed can be influenced strategically without relearning (Clark & Karmiloff-Smith 1993). We have attempted to provide a principled account of higher cognitive processes by identifying them with the properties of relational knowledge, defined in section 2.2, and we propose that the processing limitations we are considering apply to cognitive functions having these properties.

2. Relations and processing demand

The way in which relational complexity influences processing load can be illustrated by the difficulty of the following sentence:

“The boy the girl the man saw met slept.” (1)



GRAEME HALFORD, Professor of Psychology at the University of Queensland since 1989, does research on cognitive development and adult cognition, especially the role of learning and information processing limitations in shaping children’s cognitions. His 80 published technical works include *The Development of Thought* (Erlbaum 1982), and *Children’s Understanding: The Development of Mental Models* (Erlbaum 1993), *Mathematics Education: Models and Processes* (with L. English, Erlbaum 1995) and *Developing Cognitive Competence: New Approaches to Process Modeling* (Edited with Tony Simon, Erlbaum 1995) as well as articles in *Cognitive Psychology* and *Child Development*. He is fellow of the American Psychological Society, the Australian Psychological Society and the Academy of the Social Sciences in Australia and consulting editor of the *Psychological Review*. His main recreation is yachting, but he also enjoys theatre and ballet and is interested in the graphic arts.



WILLIAM WILSON is a Senior Lecturer in the School of Computer Science and Engineering at the University of New South Wales, Sydney, Australia. His research interests include cognitive modelling, natural language processing, and neural networks, including tensor product networks, feedforward networks, and generalizations of simple recurrent networks and their application to cognitive modelling and to pattern recognition problems including handwritten character recognition. He is the author of over 50 scientific publications.



STEVEN PHILLIPS (PhD., Computer Science, University of Queensland), a researcher in the Information Science Division of the Electrotechnical Laboratory (Japan), works on systematicity, schema induction, and relational knowledge, and their implications for connectionist models.

The problem here does not reside solely with storage of either the original sentence or the results of partial processing; it also reflects the amount of information about which decisions must be made. This sentence entails an integrated structure (discussed in sect. 6.1.4) that requires boy, girl, and man to be assigned to roles corresponding to subjects of three verbs and objects of two verbs. Strategies are not generally available to English speakers for serial processing of centre-embedded clauses, and the relative lack of semantic cues or syntactic case markers means people receive little help in deciding which nouns fill which case roles. The result is that we have to decide who saw, who met, and who slept as well as to identify the objects of “saw” and “met” all together, because we cannot positively identify subjects or objects of any of the verbs until we understand the whole sentence. Sentences with this type of reduced relative clause are known to impose high resource demands (Just & Carpenter 1992; Kimball 1973).

Another example is provided by Sweller (1993), who analysed the following problem: *Suppose five days after the day before yesterday is Friday. What day of the week is tomorrow?* Despite our expertise in reasoning about days of the week, this problem is frustratingly difficult. The reason is that, especially in the first sentence, numerous elements are related to each other and cannot be considered meaningfully in isolation. These relations have to be at least partially processed in order to segment the statement into sub-problems that can be processed serially. The processing load is felt most keenly when we try to plan this procedure.

The processing load imposed by interacting components of a task can be captured with the concept of *relational complexity*. We will begin by considering relations between different numbers of factors. At a low level of complexity, consider a case in which a cognitive process is constrained by a single factor; for example, our choice of restaurant might depend on the amount of money we have. We can express this as a binary relation, between money and restaurant choice, that is, a set of ordered pairs, in which each amount of money is associated with a particular restaurant (or with a subset of restaurants). The money–restaurant relation could be modified by another factor, such as importance: The more important the occasion, the more expensive our choice of restaurant, though importance might have more influence when we have plenty of money than when we have little. Here we have an interaction between two determining factors. This situation can be represented as a ternary relation, comprising a set of ordered triples in which each amount of money and each level of importance is associated with a restaurant. These variables could in turn interact with a third factor, such as hunger, which might make us profligate, but only when we are not really poor. We now have an interaction between three determining factors. This can be expressed as a quaternary relation, comprising a set of ordered quadruples, in which each amount of money, level of importance, and state of hunger is associated with a restaurant choice.

It is clear that the problem becomes more complex as the number of interacting factors increases. This complexity can be measured by the dimensionality of the relation, or the number of variables related. Problems that entail a binary relation are simpler than those that entail a ternary relation, which are simpler than those that entail a quaternary relation, and so on. The idea of relational complexity is analogous to the number of factors in an experimental design.

An experimental design can be thought of as a set of relations between independent and dependent variables. A one-way experimental design is equivalent to a binary relation between one independent and one dependent variable. A two-way experimental design is equivalent to a ternary relation between one dependent and two independent variables. Experimental designs with more factors permit more complex interactions, but at the cost of more observations (participants). This is analogous to the processing load imposed by problems of high relational complexity. [See also multiple book review of Chow: *Statistical Significance* BBS 21(2) 1998.]

The complexity of a relation may be defined by the number of arguments. A binary relation has two arguments: for example, the binary relation “bigger-than” has two arguments, a larger and a smaller entity. A ternary relation has three arguments: for example, “love-triangle” is a ternary relation and has arguments comprising three people, two of whom love a third. Quaternary relations have four arguments, and so on. Each argument can be instantiated in more than one way. For example, each argument of “bigger-than” can be instantiated in an infinity of ways, such as bigger-than(horse, mouse) . . . bigger-than(whale, dolphin) . . . and so on. Consequently, each argument provides a source of variation, or a dimension, and thereby makes a contribution to the complexity of the relation. Our next step is to link relational complexity to the concepts of demand (or load) and resources, which commonly have been used to account for performance limitations in cognitive psychology.

2.1. Processing complexity

Processing complexity might depend on the strategy used by the person in a particular set of circumstances; different people use different strategies, and even the same person may use different strategies at different times. The optimal cognitive strategy for human subjects may not correspond to the one that is best theoretically, or to any algorithm that would be used in artificial intelligence, because human cognition operates in ways that are very different in some respects from theoretically optimal algorithms. Strategies can be constrained to some extent by experimental procedures, but, if this cannot be done with confidence, then the strategy used must be determined by empirical investigation (a point to which we will return in sect. 6). Measures of processing complexity accordingly must be specific to the actual cognitive process used.

The *complexity of a cognitive process* is the number of interacting variables that must be represented in parallel to implement that process. Processing complexity also can vary over time within one task, hence the critical value is the complexity of the most complex step. Tasks can vary in the number of steps they require, but this does not necessarily affect processing load because a task with many steps might impose only a low demand for resources at any one time (e.g., counting peas in a box). Processing demand can be high when steps are embedded in a hierarchical structure, but this entails higher order relations and relational complexity is also high (discussed in sect. 2.2.5).

The *processing complexity of a task* is the number of interacting variables that must be represented in parallel to perform the most complex process involved in the task, using the least demanding strategy available to humans for that task.

Processing demand is the effect exerted by task complexity on a performer and it reflects the requisite cognitive resources. The core proposal of this target article is that demand is a function of relational complexity. That is, the more interacting variables to be processed in parallel, the higher the demand. “Demand” is synonymous with “load” and “effort,” and the three terms tend to be used interchangeably in the psychological literature. This is the psychological counterpart of computational cost, which will be discussed in section 5. Demand can be manipulated experimentally, with other aspects of the task controlled; a number of examples will be considered in section 6.

The *resources* allocated to a task vary as a function of demand and performance. More resources must be allocated to higher demand tasks to maintain performance. The methodology for dealing with demand and resources is now highly developed, and has been reviewed by Gopher (1994) and Halford (1993, Ch. 3). The resources used can be measured by physiological arousal indicators (Kahneman 1973), by brain imaging techniques (Carpenter & Just 1996), by the subjective feeling of effort assessed through self-report, and by decrements in competing tasks (Navon & Gopher 1980). The resources invested by an individual in a given task will vary over time as a function of the conditions of performance.

Processing capacity is the limit of the available resources. It will vary across individuals and may change over the life span (discussed in sect. 6.3). Within a short time frame, processing capacity is essentially constant, but it can be influenced by factors such as physiological state, diurnal rhythms, and drugs.

To put this argument on a more formal basis, we will consider the nature of relational knowledge.

2.2. Relational knowledge

Given that processing complexity can be related to the number of arguments in a relation, characterizing the nature of relational knowledge becomes important for a theory of processing capacity. Phillips, Halford, and Wilson (1995; submitted) have argued, on the basis of the work of Codd (1990), that essential properties of higher cognitive processes, including data structures such as lists and trees, can be captured by a model based on processing relations. Our explanation for capacity limitations in higher cognitive processes depends on the complexity of neural net models of relational knowledge (which are considered in sect. 4). Hence, we must specify the properties that relational knowledge must have for a neural net model to be considered adequate. For our present purposes, it will be appropriate to say that relational knowledge consists of relational schemas, which we will now define.

2.2.1. Relational schemas. Relational schemas are cognitive representations that include elements and relations between elements, and that represent situations or activities in the world. In general, an n -ary relation R is a subset of the Cartesian product of n sets: $S_1 \times S_2 \times \dots \times S_n$. Thus if $(a_1, a_2, \dots, a_n) \in R$ we say that $r(a_1, a_2, \dots, a_n)$ holds; for example, $(cat, mouse) \in larger-than$ signifies *larger-than* (cat, mouse). An n -ary relation comprises a set of n -tuples, where each “tuple” is a relational instance. We shall refer to R and “larger-than” as relation-symbols. Tuples such as (a_1, a_2, \dots, a_n) and (cat, mouse) we call “relational instances.” It may not be clear, whether, for example, (cat, mouse) is being considered as an instance of the relation “larger-than” or the relation “chases.” For this reason, we frequently use the term “relational instance” for an

expression of the form $r(a_1, a_2, \dots, a_n)$ or *larger-than*(cat, mouse). Strictly speaking, $r(a_1, a_2, \dots, a_n)$, *larger-than*(cat, mouse), and *larger-than*(mouse, cat) are propositions (see sect. 2.2.2). Thus, we use the term “relational instance” to refer both to the tuple and to the tuple labeled with the relation symbol. This is not an uncommon practice in cognitive science. When it would be unclear whether $r(a_1, a_2, \dots, a_n)$ is being considered as a proposition or a relational instance, we shall refer, for example, to “the relational instance $r(a_1, a_2, \dots, a_n)$.”

2.2.1.1. Representation of a relation. The representation of a relation requires a symbol to specify the relation R , a representation of the arguments a_1, a_2, \dots, a_n , and a set of bindings between symbol and arguments that maintain the truth of the relation. The binding must constrain the fillers for each argument role, so that appropriate members of the Cartesian product are bound. For example, in *bigger-than*(-, -), it is conventional for the entity in the first argument position to be bigger than the argument in the second argument position. Thus, *bigger-than*(cat, mouse) and *bigger-than*(mountain, molehill) should be bound, but not *bigger-than*(mouse, cat). The symbol and arguments must retain their identity in the binding. That is, they must not be fused into a whole in which the components cannot be identified.

Representations of relations are valid when they conform to the *structural correspondence principle*, according to which the relations in the representation must correspond to relations in the world. Formally, a relational schema comprising elements E^s and relations R^s corresponds to an aspect of the world with elements E^w and relations R^w if there exists a function f that assigns each member of E^s in the schema to a member of E^w in the world, in such a way that, for any $r_s(e_1^s, e_2^s, \dots, e_n^s)$ in the schema ($r^s \in R^s, e_i^s \in E^s$), there is an $r^w(e_1^w, e_2^w, \dots, e_n^w)$ in the world [$r^w \in R^w, e_i^w = f(e_i^s)$]. These criteria have been specified in more detail by Halford and Wilson (1980) and by Holland et al. (1986), and their neural net implementation has been specified by Halford et al. (1994). Structural correspondence is a soft constraint, and in some cognitive models it can be overridden by other constraints if they are of sufficient strength (see, e.g., Hummel & Holyoak 1997). Learning and induction processes, however, will tend to bring relational schemas into correspondence with the aspect of the world they represent (Holland et al. 1986). This can be done by reducing the strength of representations with inappropriate bindings.

2.2.1.2. Representation of a relational instance such as $r(a_1, a_2, \dots, a_n)$, which requires a binding between the relation symbol r and the fillers a_i , each bound to one argument role. Thus *bigger-than*(whale, dolphin) is a relational instance, whereas the *bigger-than* relation includes this instance plus appropriate others, such as *bigger-than*(cat, mouse) or *bigger-than*(mountain, molehill).

A relational instance in isolation can be represented by specifying the relation symbol r plus each of the role-filler bindings. The relational instance *loves*(John, Mary) requires a representation of “loves” plus a binding of John to the lover role and Mary to the loved role. We can write this as follows.

loves + lover:John + loved:Mary

Here the period signifies the role-filler binding and the plus sign serves to concatenate the bindings to the relation symbol. This is not sufficient to represent a relation (as distinct from relational instances), because ambiguity occurs as soon as more than one relational instance is represented. Suppose we now add the instance *loves*(Tom, Wendy), represented as follows.

loves + lover:Tom + loved:Wendy

When we put the representations of both relational instances together we have the following.

loves + lover:John + loved:Mary + loves + lover:Tom + loved:Wendy

This represents the fact that John and Tom are lovers and that Mary and Wendy are loved, but it does not distinguish between

John loving Mary and John loving Wendy and is similarly ambiguous with respect to Tom.

The ambiguity can be removed by indicating that John and Mary belong to one instance of “loves” and Tom and Wendy to another. One possibility is to index each instance with a unique identifier. This in effect separates (or brackets) the “location” of the representation of each instance. For example, the loves relation could be represented as follows.

- 1.(loves + lover.John + loved.Mary) +
- 2.(loves + lover.Tom + loved.Wendy)

One implication of this method is that, in general, the index does not indicate its contents (the index “1” does not indicate that John and Mary are involved). Thus, potentially all instances must be processed to determine the filler for a given role (in the worst case, one may have to search all instances to find the one in which John is the lover of Mary).

An alternative approach is to define each relational instance as a unique n -tuple. This can be done by representing bindings between the relation symbol and the fillers for each argument, such as the following.

loves.John.Mary + loves.Tom.Wendy

This representation is based on *symbol–argument–argument* bindings, each of which comprises an intact relational instance together with the symbol for the relation (“loves” in this case). The binding between a symbol and its arguments represents the link specified by the relational instance. Thus, representing the tuple loves.John.Mary identifies this relational instance unambiguously, and directly represents the link, identified as “loves,” between John and Mary.

2.2.1.3. Relations and working memory. Working memory has often been associated with storage of items of information, as in span tasks, which entail storing strings of items. Therefore, it might appear that working memory would store relational instances, that is, interconnected sets of items, rather than relations. In some cases this would be sufficient. For example, to understand “John loves Mary” we need only represent relational instance loves(John, Mary). In some cases, however, working memory calls for variables. Consider, for example, reasoning about velocity, defined as $V = s/t$ (where s is distance and t is time). We know that if we cover the same distance in half the time, velocity is doubled, even though we know no specific values. Thus we are reasoning about the interaction of three variables, velocity, distance, and time, and we are dealing with a ternary relation, rather than a relational instance.

2.2.2. Proposition. A proposition is defined as the smallest unit of knowledge that can have a truth value. If the proposition is true, then there will be a corresponding relational instance. For example, the proposition bigger-than(cat, mouse) is an instance of the bigger-than relation. However a proposition need not be true, and a proposition that is false is not a relational instance: bigger-than(mouse, cat) is false and is not an instance of bigger-than as defined above.

True propositions and false propositions correspond to different subsets of the Cartesian product. The relation $>$ is a subset of $S_1 \times S_2$: the subset $\{(a, b) \mid a \in S_1, b \in S_2 \text{ and } a > b\}$. Consider the false proposition $>$ (mouse, cat). The pair (mouse, cat) is not a part of the relation $>$ (there are no cases of mice being bigger than cats), but the proposition $>$ (mouse, cat) is representable, because (mouse, cat) $\in S_1 \times S_2$. Hence, false propositions can be represented and correspond to subsets of the Cartesian product. Learning and induction will tend to weaken representations of false propositions and will tend to incorporate semantic constraints. Thus a proposition such as owns(car, Tom) will tend to be excluded. However, false propositions do occur in real cognitive processes, and provision must be made for them to be represented.

2.2.3. Truth value. The truth value of a proposition can be assessed by matching against semantic memory using a mechanism

described in section 4.2.1. Truth value can be represented as a higher order relational instance, for example, false[bark(cats)], meaning that it is false that cats bark. There is a psychological bias to represent true propositions. For example, in mental models theory (Johnson-Laird et al. 1992) only true contingencies are represented to reduce load on working memory. [See also multiple book review of Johnson-Laird and Byrne: *Deduction*. *BBS* 16(2) 1993.]

Quantifiers are not explicitly represented in relational schemas or in mental models, but the closest psychological property is strength. A strong proposition is one that has a high probability of being true (propositions such as “dogs are bigger than cats,” which have no definite truth value, in the sense that they are not universally either true or false¹).

2.2.4. Symbolization. “Symbolization” means that a link between the arguments of a relation or relational instance is explicitly symbolized (e.g., one link between “whale” and “dolphin” is explicitly symbolized by “bigger-than”). The link is one of the requirements for identifying a relational instance, which is in turn necessary for a relational instance to be an argument in another relation, thereby forming higher order relations. It also allows us to distinguish between a number of different links between the same argument sequence.

2.2.5. Higher order relations and hierarchical structures.

Higher order relations have relational instances as arguments, whereas first-order relations have entities as arguments. Higher order relations can represent connectives, for example, implies[$>$ (a,b), $<$ (b,a)] or and[dog(Fido),pet (Fido)]. Higher order relations can be used to represent hierarchical structures, for example, cause[shout-at(John, Tom), hit(Tom,John)].

The *repeated variable constraint* operates with hierarchical structures. In cause[shout-at(x,y),hit(y,x)], x and y must be bound to the same entities in both cases. For example, cause[shout-at(John,Tom), hit(Tom,John)] has the required structure, but cause[shout-at(John,Tom), hit(John, Tom)] does not. The repeated variable constraint is implemented by ensuring that the first relational instance is in structural correspondence with cause[shout-at(x,y),hit(y,x)]. Notice that a mapping that conforms to the structural correspondence principle (in sect. 2.2.1.1) can be formed between them, whereas for the second relational instance such a mapping would be inconsistent (John and Tom must each be mapped to both x and y).

2.2.6. Omnidirectional access. “Omnidirectional access” means that, given all but one of the components of a relational instance, we can access (i.e., retrieve) the remaining component. For example, given the relational instance mother-of(woman, child), and given mother-of(woman, ?) we can access “child,” whereas, given mother-of(?, child) we can access “woman,” and, given ?(woman, child) we can access “mother-of.”

Omnidirectional access is also potentially true of relations. A relational instance $r(a_1, a_2, \dots, a_n)$ contains $n + 1$ objects, the relation symbol and the n arguments. Given any n components, we can retrieve one or more candidates for the $n + 1$ st component. When more than one relational instance is represented, however, the answers obtained will not always be unique. For example, given the relation “mother-of,” the query mother-of(woman, ?) may yield child, toddler, infant, baby, teenager, and so on. Access might not be equally efficient in each direction. For example, arithmetic addition corresponds to the ternary relation $+\{ \dots (3, 2, 5), \dots, (5, 4, 9) \dots \}$. It might be easier to access a sum given two addends, that is, $+(3, 2, ?)$, than to access an addend given the sum and the other addend, that is, $+(3, ?, 5)$, but access in both directions is possible. Having learned our addition tables we can perform subtraction, but perhaps not as efficiently as addition. (Another example from the psychological literature is discussed in sect. 6.2.5.1.)

2.2.7. Role representation. “Role representation” means that relational knowledge requires representing argument roles or slots, independent of specific instances. Thus bigger-than(—, —)

requires the representation of roles for a larger and a smaller entity. Roles must be distinguishable, but they need not be represented explicitly; the role an argument fills can be identified by its position relative to the other arguments. Thus, in “loves.John.Mary” (discussed in sect. 2.2.1.2), we know that John is in the lover role by his position in the binding.² The argument positions in a proposition correspond to the sets indicated by their position in the Cartesian product space. Given the relational instance $r(a_1, a_2, \dots, a_n)$ for a relational R that is a subset of $S_1 \times S_2 \times \dots \times S_n$, each a_i corresponds to a given S_i .

2.2.8. Operations on relations. Operations on relations are adapted from those defined in the theory of relational databases (Codd 1990) and they include select, project, and join, plus the usual set operators, intersection, union, and difference (Phillips et al. 1995). These operations permit information stored in relational knowledge structures to be accessed and manipulated in flexible and powerful ways. The select and project operations together permit access to any element within a relation. Informally, if one thinks of a relation as a table, in which rows correspond to relational instances and column names correspond to role names, then select and project return rows and columns of a table, respectively. The join operation corresponds to combining two tables at a specified pair(s) of columns.

The operators are best described by example; suppose the following relation: Taller = {(John, Mary), (Mary, Tom)}, where Taller is a subset of Person \times Person. Formally, a select operation parametrized with condition C takes a relation R and returns a new relation R' such that all instances of R' satisfy C . For example, $\text{select}_{\langle \text{Person1, John} \rangle}$ Taller \rightarrow {(John, Mary)}, returns a single-instance binary relation with John at role Person1. A project operation parametrized with attribute (role) name(s) A takes a relation R and returns only arguments at attribute A for each instance in R , for example, $\text{project}_{\langle \text{Person2} \rangle}$ Taller \rightarrow {(Mary), (Tom)} (i.e., a unary relation with two instances). Taken together, select and project provide omnidirectional access to all relational elements. For example, the query “who is taller than Mary” is realised as: $\text{project}_{\langle \text{Person1} \rangle}(\text{select}_{\langle \text{Person2} = \text{Mary} \rangle}$ Taller) \rightarrow {(John)}, which we write as Taller(-, Mary) \rightarrow John, for short.

There are a number of different *join* operators that take two relations and return a new relation. The outer (or natural) *join* is analogous to the Cartesian product. It returns a relation containing every unique pairwise combination of instances from the argument relations. In this way it permits the construction of higher “-arity” relations; the outer join of two unary relations is a binary relation. Of more interest for our purposes is the equijoin operator, which joins only instances that have the same arguments at the specified roles. It provides a way of implementing transitive inference.

For example, $\text{equijoin}_{\langle \text{Person2, Person1} \rangle}$ (Taller, Taller) \rightarrow {(John, Mary, Tom)} is an equijoin of the relation Taller with itself along roles Person2 and Person1. Projecting onto the first and third positions of the resulting relation results in {(John, Tom)} (i.e., the inference “John is taller than Tom”). Finally, insofar as relations are sets, the standard set operators *intersect*, *union*, and *difference* apply in the usual way.

2.2.9. Decomposability of relations. “Decomposability of relations” means that relations can be composed of simpler relations. A *decomposable* relation is one that can be written as a conjunct of instances of relations of lower “-arities.” For example, the ternary relation monotonically-larger(a, b, c) can be decomposed into $\succ(a, b) \ \& \ \succ(b, c) \ \& \ \succ(a, c)$. This is discussed in more detail in Appendix A. Relations can be decomposed using operators *select* and *project*, defined in section 2.2.8.

2.2.10. Relational systematicity. “Relational systematicity” means that certain relations imply other relations. For example, $\succ(a, b) \rightarrow (b, a)$, whereas $\text{sells}(\text{seller}, \text{buyer}, \text{object}) \text{ buys}(\text{buyer}, \text{seller}, \text{object})$. Implication can be handled as a higher order relation as noted in section 2.2.5.

2.2.11. Attribute. An attribute is a relation with one argument. An attribute value is an instance of a unary relation [e.g., $\text{ripe}(\text{apple23})$ indicates that apple23 satisfies the unary relation ripe].

2.2.12. Analogy, planning, and modifiability. Analogy is a structure-preserving map between a base or source and a target, and representation of relations is at the core of analogies (Gentner 1983; Gick & Holyoak 1983; Holyoak & Thagard 1989). Planning is the main process in strategy development and entails the organization of a sequence of actions to achieve a goal. Planning has been explicitly modeled by VanLehn and Brown (1980), Greeno et al. (1984), and Halford et al. (1995). The development of the strategy is guided by a concept or mental model of the task, which entails representing relations between components of the task.

Higher cognitive processes can be modified “online” without the necessity of learning all over again. A performance that distinguishes between higher and lower animal species is the ability to acquire the reversal learning set; that is, having learned to choose A in preference to B, the animal must switch to B without relearning (Bitterman 1960; 1975). If the animal learns the exclusion relation between A and B (i.e., one and only one of A and B is correct), the reversal can be effected by changing the mapping of the stimuli into the relational schema, so the stimulus that was formerly mapped to the positive element of the schema is remapped to the negative element, and vice versa (Halford 1993). Clark and Karmiloff-Smith (1993) have pointed out that modifiability is a criterial attribute of human cognition. Relational representations can achieve this by switching between relations, because when a relation is changed mappings between input and output change. For example, the binary operation of arithmetical addition, a ternary relation, entails a set of mappings between addends and sum { . . . (3, 2 \rightarrow 5) . . . (4, 7 \rightarrow 11) . . . }. Multiplication entails a different set of mappings { . . . (3, 2 \rightarrow 6) . . . (4, 7 \rightarrow 28) . . . }. A switch to a different relation activates a different set of mappings and modifies the performance.

Relational knowledge is symbolic, content-independent, flexible, and modifiable and can serve the functions of higher cognitive processes. Our next step is to consider the psychological properties associated with different levels of relational complexity.

2.3. Psychological interpretations of orders of relational complexity

Each level of relational complexity, from unary to quaternary, corresponds to a distinct category of cognitive tasks. Empirical indicators for each level will be considered in section 6.

A *unary relation* has relational instances $r(x)$ that can be interpreted as propositions with one argument, as variable-constant bindings, or as zero-variate functions. A proposition with one argument can represent a state, such as happy(John), an action, such as ran(Tom), an attribute, such as big(dog), or class membership, such as dog(Fido).

Binary relations have relational instances $r(x, y)$ that can sometimes be interpreted as univariate functions; $f(a) = b$ is a special case of a binary relation, in which the mappings are unique; it is a set of ordered pairs, (a, b) , such that for each a there is precisely one b such that $(a, b) \in f$. A unary operator is a special case of a univariate function; for example, the unary operator change-sign comprises the set of ordered pairs $\{(x, -x)\}$.

Ternary relations have relational instances $r(x, y, z)$ that can be bivariate functions and binary operations. A bivariate function is a special case of a ternary relation. It is a set of ordered triples (a, b, c) such that for each (a, b) there is precisely one c such that $(a, b, c) \in f$. A binary operation is a special case of a bivariate function; a binary operation on a set S is a function from the set $S \times S$ of ordered pairs of elements of S into S , that is, $S \times S \rightarrow S$. For example, the binary operation of arithmetic addition consists of the set of ordered pairs of { . . . (3, 2, 5) . . . (5, 3, 8) . . . }; that is, $\{(x, y, z) \mid x + y = z, x, y, z \text{ (say) natural numbers}\}$.

With a ternary relation, it is possible to seek x such that $r(x, y, z)$ is a relational instance given y, z , and similarly for y given x, z ,

or z given x, y . It thus becomes possible to compute the effects on x of variations in y, z and so on. This emergence of three-way comparisons in ternary relations is analogous to the emergence of interactions in two-way experimental designs.

Quaternary relations have relational instances of the form $r(w, x, y, z)$. Proportion, $a/b = c/d$, is a quaternary relation and entails relations between the four terms, a, b, c, d . Given any three terms, plus the knowledge that proportion is entailed, we can predict the remaining term, or, given all of a, b, c, d , we can decide whether proportion is entailed (a case of omnidirectional access; see sect. 2.2.6). 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 also encompass trivariate functions and ternary operations. A trivariate function is a special case of a quaternary relation. It is a set of ordered quadruples (a, b, c, d) such that for each (a, b, c) there is precisely one D such that $(a, b, c, d) \in f$. Compositions of binary operations, such as $a(b + c) = d$, correspond to quaternary relations.

Regarding *dimensionality of relations*, each argument x_i of $R(x_1, \dots, x_n)$ for an n -ary relation R can be instantiated in more than one way and therefore represents a source of variation, or dimension. An n -ary relation may be thought of as a set of points in n -dimensional space and can represent interaction between n variables. The number of arguments, n , corresponds to the number of dimensions in the space defined by the relation. This is the basis for our proposed complexity metric. The relation symbol can be predicted in principle from the arguments; for example, given $?(3, 2, 5)$, where $?$ is known to be an arithmetic operation, we know the operation must be addition, whereas given $*(?, 2, 6)$ we know the first multiplicand must be 3, and so on. Prediction of the relation symbol may depend on constraints, such as knowing the relevant domain, as in this example, in which it depended on knowing that the domain was arithmetic operations. A suitable set of relational instances must be available, and in the worst case all relational instances must be known. Because the relation symbol can be predicted, at least in principle, from the arguments, there are only n independent sources of variation in an n -ary relation, and the number of dimensions equals the number of arguments.

Algorithms that embody these properties will be discussed in section 4. However, our next step is to define processing capacity in terms of relational complexity.

3. Processing capacity

The amount of information that can be processed in parallel has long been recognized as a critically important datum in cognitive psychology. The most notable attempt to estimate this parameter was made by Miller (1956), who suggested that human capacity was limited to a small number of chunks.

3.1. Chunks

Miller's (1956) concept of a chunk may be defined as a unit of information of arbitrary size, so a digit, an alphabetic character, and an English word may all constitute one chunk, although they vary in information content. The paradox is that the limitation seems to be not in the amount of information but in the number of independent units.

3.2. Chunks and dimensions

There is some correspondence between the properties of chunks and the properties of dimensions. Each chunk is a separate signal and fills a separate slot in a message. Attributes or values on different dimensions are at least partly in-

dependent of each other, by definition (even nonorthogonal dimensions must convey some independent information, or they are redundant).³ Thus chunks, as with dimensions, represent units of information that are at least partly independent. Their similarity is illustrated by the proposition played(John, cricket, oval, Sunday), which has four roles or slots corresponding to player, game, location, and day. It seems equally appropriate to regard each filler of these roles as a different chunk or as a value on a different dimension.

The amount of information (in the information theoretic sense; Attneave 1959) conveyed by a chunk depends on the number of alternatives for that slot. For example, "cat" conveys $\text{Log}_2 2 = 1$ bit of information if there are two equally likely alternatives (e.g., cat or dog). If there are 32 (equally likely) alternatives, however, then the chunk "cat" conveys $\text{Log}_2 32 = 5$ bits. An attribute on a dimension also represents varying amounts of information, depending on the number of alternative values on that dimension. Hence the amount of information conveyed by a chunk or dimension is arbitrary; both chunks and dimensions are independent units of information of arbitrary size.

3.3. Number of dimensions processed in parallel

Given the link between dimensions and chunks, the number of dimensions that can be processed in parallel can be estimated by determining the number of chunks that can be processed in parallel. Miller (1956) proposed that approximately seven chunks were processed in parallel, but difficulties have arisen with his empirical database (Baddeley 1986; Halford et al. 1988; Schweickert & Boruff 1986).

More recent research has revised Miller's estimate downward. Broadbent (1975) examined temporal patterns in recall from semantic memory and found that items tended to be recalled in groups of approximately three. He suggested that the "magical number seven" proposed by Miller might have reflected the combined output of two systems, each with a capacity of three or four items. Fisher (1984) studied visual scanning and found a modal value of four items processed in parallel, with a range of three to five. Halford et al. (1988) assessed the capacity of primary memory, or the information that is currently active (James 1890) at four or five items. A number of other studies reviewed by Schneider and Detweiler (1987) also indicate that approximately four chunks are processed in parallel. Given the identification of chunks with dimensions in section 3.2, this implies that approximately four dimensions can be processed in parallel and that humans should be limited to processing quaternary relations in parallel. Most studies show a range of values, indicating that this should be a soft limit. Neural net models of relations agree in predicting a soft limit, as will be discussed in section 5.

An attempt has been made to assess the number of dimensions that can be processed in parallel using interpretation of interactions, because the factors in an interaction cannot be interpreted meaningfully alone; hence there is a constraint to process them in parallel (Halford et al. 1994). An interaction between N factors corresponds to a relation between N independent variables and the dependent variable, as discussed in section 2, so the ability to process four dimensions implies, *prima facie*, an understanding of three-way interactions. Academic staff and graduate students who were experienced in interpreting statistical interactions were asked what was the most complex interaction they

could interpret unambiguously and with confidence, ignoring scale effects and nonlinearity. Ten answered two-way, fourteen three-way, and six four-way. The variations in estimates probably reflect errors resulting from imprecision of the test, but the mode is three-way, suggesting that four dimensions are processed. Although no single study is definitive, there is a degree of consensus across studies employing a wide range of methodologies that approximately four dimensions can be processed in parallel.

If it is possible to process two relations in parallel but independently, then the sum of their processing demands would be less than for a single relation with the same number of arguments. If we consider a k -ary relation R on a set S with s members, then each component x_i in a tuple $(x_1, x_2, \dots, x_k) \in R$ might be filled in s ways: The number of possible tuples is s^k . Therefore, the number of tuples is $2s^2$ for two binary relations but s^4 for a quaternary relation. When we consider neural net representation of relations based on symbol – argument – argument bindings (in sect. 4.1.1.2), the number of binding units is of n^{k+1} , where n is the number of elements of each vector, and a similar argument applies. The limit will accordingly be reached more quickly with a single relation than with two relations having the same total number of arguments. Notice also that more links are defined in a quaternary relation than in two binary relations. Thus $R(a,b,c,d)$ defines links between six pairs, $a-b, a-c, \dots, b-c, \dots, c-d$, whereas $R(a, b)$ and $R(c, d)$ collectively define links between only two pairs, $a-b$ and $c-d$.

It is unlikely, however, that two or more relations can be processed in parallel and independently in the central executive, or within any one system, because they would have to be coordinated to avoid conflict. They could be coordinated by integration into a higher order relation, but this implies that they are effectively being processed as a single relation. Another method would be to superimpose two or more relations; this can be done in the neural net models (to be discussed in sect. 4.1.1.2). For example, the relational instances mother-of(mare, foal), loves(mare, foal), feeds(mare, foal) can all be superimposed. Notice, however, that superimposing relational instances in this way does not increase the number of dimensions being processed (there are only two arguments in this example). When we consider neural net implementations in section 4.1.1, it will be apparent that such superposition adds little computational cost. Superimposed relational instances can be treated as a whole, but they can also be processed separately, using the retrieval process described in section 4.2.1. On the other hand, the relation between them cannot be defined by superposition (e.g., “mother-of,” “loves,” and “feeds” can be fused into a whole equivalent to some kind of composite motherhood concept, or they can be processed as separate relational instances, but no relation is defined between the relational instances).

If four dimensions can be processed in parallel, the next question concerns how more complex concepts are processed. Many concepts are more complex than quaternary relations, so we must have some means of dealing with these concepts without exceeding our processing capacity.

3.4. Using capacity efficiently

We propose two mechanisms for reducing processing loads imposed by complex concepts. These are conceptual chunking and segmentation.

3.4.1. Conceptual chunking. Conceptual chunking is the recoding of concepts into fewer dimensions. In the limiting, and most typical, case, concepts are recoded into a single dimension. In a mnemonic chunk, items function as a unit (e.g., c, a, t becomes a chunk if the three letters form a single word “cat”). Similarly, elements that are formed into a conceptual chunk function as a whole in a relational structure, and relations between items within the chunk cannot be accessed.

We can illustrate conceptual chunking using the concept of velocity, defined as $V = s/t$ (where s is distance and t is time). The relation between velocity, distance, and time is three dimensional, but velocity can be expressed as a single-dimension, such as the position of a pointer on a dial; velocity(60 km/h). In this single-dimensional representation, no relation is defined between velocity, distance, and time. If we want to compute, say, the way velocity varies as distance increases and time decreases, we must return to the three-dimensional representation. Thus conceptual chunks save processing capacity, but the cost is that some relations become temporarily inaccessible. There is also a psychological factor that limits chunking, because experience is required in which there is a constant mapping of elements into chunks (Logan 1979). Chunking is a form of learning, which takes place over time.

Chunked concepts can be combined with further dimensions to represent higher level concepts, so acceleration can be defined as $A = (V_2 - V_1)/t$. Acceleration also can be chunked, and then force (F) can be defined as $F = MA$ (where $M = \text{mass}$). In this way we can bootstrap our way to higher and higher level concepts without ever exceeding four dimensions processed in parallel. A major function of expertise is to provide ways of chunking that permit the important features of concepts to be represented without imposing excessive processing demands.

When a role has only one possible filler, it can be chunked without loss, and the number of dimensions is reduced accordingly. Consider, for example, the relation mother-of{(Jenny, Tom),(Jenny, Mary),(Jenny, Jill)}. The mother role is always filled by Jenny, so the representation can be collapsed to the unary relation Jenny-mother-of{(Tom),(Mary), (Jill)}.

The general principles of chunking are: (1) a chunk functions as a single entity, relation-symbol or argument, in a relation; (2) no relations can be represented between items within a chunk; and (3) relations between the chunk and other items, or other chunks, can be represented. In order to assess processing capacity, chunking can be inhibited by using novel structures for which chunks have not been learned. This does not preclude using familiar domains. For example, in testing transitivity, size relations between unknown persons can be used (e.g., John > Tom, Tom > Peter). Size relations are a familiar domain, but the specific orderings (John, Tom, Peter, etc.) will not have been prelearned as chunks.

3.4.2. Segmentation. “Segmentation” is breaking tasks into steps that do not exceed processing capacity and that are processed serially. Examples include algorithms for arithmetic operations, counting, and ordering tasks. Arithmetic algorithms such as multidigit addition generally have to be taught, but there is some degree of autonomy in acquisition of counting and ordering algorithms. It is not possible to determine the precise number of elements in a large set solely by parallel processing, so we use the serial procedure of counting objects one (or, at most, a few) at a

time. Children have some understanding of the principles of counting, and this guides the development of their strategies (Greeno et al. 1984). People's concept of an ordered set can provide a mental model for an ordering algorithm (Halford et al. 1995). [See also Schusterman: "Language and Counting in Animals," *BBS* 11(4) 1988; Geary: "Sexual Selection and Sex Differences in Mathematical Abilities," *BBS* 19(2) 1996.]

The development of counting and ordering strategies illustrates the principle that autonomous development of serial processing strategies requires planning, which depends on representing relations (VanLehn & Brown 1980). If processing limitations prevent the structure of the task being represented, then a strategy cannot be developed without didactic help, which can be available for only a small subset of the cognitive tasks we perform. Thus the "self-programming" property of higher cognitive processes depends on the ability to represent relations.

3.4.3. Effective complexity determined by reduction technique. Because complexity can be reduced by conceptual chunking and segmentation, the number of arguments of a relation does not immediately translate into effective complexity. Also, simply increasing the number of arguments by conjunction does not necessarily contribute to the complexity of the resulting relation. The important point regarding relational complexity concerns the nature of the interaction between the relational elements. Effective relational complexity can be determined using a reduction technique.

More specifically, the effective complexity of a relation is the minimal dimensionality to which a relation can be reduced without loss of information. Thus, if a ternary relation can be reduced to two binary relations without loss of information, then effective relational complexity is binary, not ternary. One can determine whether a relation can be reduced to a combination of lower order relations by a procedure of decomposing and recombining. If the resulting relation is the same as before, then the relation can be decomposed into lower order relations without loss of information. Psychologically, *effective relational complexity* is the minimal dimensionality to which a relation can be reduced using decomposition and recombination procedures available to human performers.

For example, suppose the following three facts:

- (1) John played tennis at the school.
- (2) John played soccer at the park.
- (3) Mark played soccer at the park.

Intuitively, it would appear that this domain consists of a ternary relation over person, game, and location. That is, $\text{played}(\text{person}, \text{game}, \text{location}) = \{(\text{John}, \text{tennis}, \text{school}), (\text{John}, \text{soccer}, \text{park}), (\text{Mark}, \text{soccer}, \text{park})\}$. Thus, at first we may claim that the relational complexity of this domain is ternary. However, this ternary relation can be decomposed into two binary relations by splitting the ternary relation along the game attribute. The resulting relations are: $\text{played}(\text{person}, \text{game}) = \{(\text{John}, \text{tennis}), (\text{John}, \text{soccer}), (\text{Mark}, \text{soccer})\}$; and $\text{Is-played-at}(\text{game}, \text{location}) = \{(\text{tennis}, \text{school}), (\text{soccer}, \text{park})\}$. Now, recombining these two binary relations by joining them along the common attribute Game results in the original ternary relation $\text{played}(\text{person}, \text{game}, \text{location})$ containing exactly the same elements. Hence the ternary relation is decomposable into two binary relations, and effective relational complexity is binary.

Now suppose that the domain has changed to include a new fact: (4) Mark played soccer at the school. We will see that including this fact changes the relational complexity of the domain. The ternary relation played consists of elements (John, tennis, school), (John, soccer, park), (Mark, soccer, park), and (Mark, soccer, school). Splitting this relation along the game attribute results in the two binary relations: $\text{played}(\text{person}, \text{game}) = \{(\text{John}, \text{tennis}), (\text{John}, \text{soccer}), (\text{Mark}, \text{soccer})\}$ and $\text{is-played-at}(\text{game}, \text{location}) = \{(\text{tennis}, \text{school}), (\text{soccer}, \text{park}), (\text{soccer}, \text{school})\}$. Recombining these two relations results in the new triple: (John, soccer, school), formed by joining pairs (John, soccer) and (soccer, school). This triple, however, is not an element of the ternary relation $\text{played}(\text{person}, \text{game}, \text{location})$ before decomposing and is not recorded in any of the four facts for the new domain. Therefore, decomposing and recombining along the Game attribute have not recovered the original relation; hence there has been a loss of information.

The same procedure applied to the Person and Location attributes also results in triples, (John, tennis, park) and (John, soccer, school), respectively, that are not elements of the ternary relation $\text{played}(\text{person}, \text{game}, \text{location})$. Therefore, this new domain cannot be decomposed into two binary relations, so its effective relational complexity is ternary. One possible rejoinder to this sort of analysis is to claim that all relations can be decomposed into binary relations simply by creating a unique symbol for each element (tuple) of the higher order relation. Under this scheme, the above-mentioned domain could be decomposed into the single binary relation: $\text{involved}(\text{event}, \text{participant}) = \{(\text{JTS}, \text{John}), (\text{JTS}, \text{tennis}), (\text{JTS}, \text{school}), (\text{JSP}, \text{John}) \dots\}$, where JTS is a unique symbol for the event "John played tennis at the school," etc. However, a general encoding scheme requires processing the original ternary relation [e.g., (John, tennis, school) \rightarrow JST]. Decomposition becomes effective only once the creation process is complete. Less general encoding schemes are possible using only binary relations [e.g., (John, tennis) \rightarrow JT], but such schemes are inadequate for relations containing both (John, tennis, school) and (John, tennis, park). Furthermore, it is implausible that appropriate encoding strategies are immediately available for novel cognitive tasks.

To relate this example to the analysis of variance analogy, notice that, with facts 1 – 3, location is predicted solely by game (tennis at school, soccer in the park) independently of person. When fact 4 is added, person and game jointly predict location, and there are three interacting variables.

The reduction checking technique can also be applied to higher order relations. For example, intuitively one might expect that the relational complexity of transitive inferences (e.g., John is taller than Mary, and Mary is taller than Sue, so John is also taller than Sue) is binary, because such inferences operate over binary relations. However, transitive inference is not simply a collection of binary relations. Transitive inferences have the structure "(A R B) and (B R C); therefore (A R C)," where R is some binary relation and A, B, and C are variables ranging over the arguments of R. Transitive inference involves a constraint between two premises and a conclusion, and it is an example of systematicity, as defined in section 2.2.10. Reduction analysis shows that the structure of transitive inference is ternary; the structure cannot be reduced to a collection of binary relations without loss of information.

Transitive inference can be schematized as a higher or-

der ternary relation over binary relational instances. That is, transitive inference $(P1, P2, C) = \{(aRb, bRc, aRc), (aRb, bRd, aRd), (aRc, cRd, aRd), (bRc, cRd, bRd)\}$, where P1, P2, and C are the first and second premises, and consequent attribute names, respectively, and *a*, *b*, *c*, and *d* are symbols (place holders) to which elements of specific relational instances are aligned. Using the reduction checking technique, we show that transitive inference cannot be decomposed into binary relations.

Suppose we choose to split the relation along the P2 attribute, which results in the two binary relations $\text{and}(P1, P2) = \{(aRb, bRc), (aRb, bRd), (aRc, cRd), (bRc, cRd)\}$ and $\text{implies}(P2, C) = \{(bRc, aRc), (bRd, aRd), (cRd, aRd), (cRd, bRd)\}$. Rejoining these two relations along attribute P2 results in two ternary relational instances, (aRc, cRd, bRd) and (bRc, cRd, aRd) , which were not present in the original relation. (That is, it is not logically valid to conclude that, for example, Tom is taller than Mark given that John is taller than Bob, and Bob is taller than Mark, if we do not know the relationship between Tom and John or Bob.) Similarly, splitting and rejoining on attributes P1 and C results in additional relational instances not present in the original relation. Thus, transitive inference is not, in general, decomposable into binary relations.

Indecomposable relations are ultimately significant because if a relation is indecomposable then subjects cannot recode the problem by decomposing (ternary) relations into simpler (binary) relations. A case of an indecomposable relation is given in section 6.2.4.3. Even when decomposition is theoretically possible, participants might lack the required strategies (algorithms), and the need to cope with more relational instances (of a lower -arity) might impose loads of its own (e.g., if higher order relations are involved; see sect. 6.1.3).

3.5. Effects of processing overload

A participant who cannot construct a representation of the dimensionality required for a task has three options.

(1) The concept can be chunked to a lower dimensional representation. This will be possible only if appropriate chunks have been learned or can be constructed, and it results in loss of access to relations between chunked entities.

(2) The task can be segmented into smaller steps that are performed serially, but this requires a strategy, autonomous planning of which depends on the participant's ability to represent relations in the task.

(3) The participant can default to a lower level representation. This is analogous to performing an experiment with, say, a three-way design, then analysing the data by a series of two-way ANOVAs. Just as the analysis would lead to recovery of most of the relevant data in the experiment (all main effects and two-way interactions would be recovered⁴), the performance would probably be correct in most respects. However, just as the hypothetical experimenter would miss the three-way interactions, our hypothetical performer could not reason about high-level relations in the task.

3.6. Capacity and content

It is important to consider whether processing capacity is a function of content. As was indicated in section 1.1, the capacity limitations we have defined do not apply to modular

processes such as vision; they apply to higher cognitive functions, which entail processing explicit relational knowledge, defined in section 2.2. However, it is still reasonable to ask whether complexity might be influenced by content. We suggest that complexity effects of content variations can be attributed to processes such as conceptual chunking, segmentation, and the use of higher order relations. Relations in a familiar domain can be more readily chunked, or higher order relations may be known that allow the structure to be represented hierarchically, as illustrated in section 2.2.5. It can then be segmented by processing one level of the hierarchy at a time, as described in section 4.2.5.

An example of content effects is discussed in section 6.1.4, but here we will consider two examples of the way relational complexity can be applied to different task contents and formats. Andrews and Halford (in press) tested young children's ability to order colored blocks using premises such as "red above green" and "green above blue." In the construction condition, the children simply built towers with green above blue, then red above green, and so on. In the prediction condition, children had to say in advance which of two blocks, red or blue, would be higher in the tower. The construction condition was easier, apparently because of its concrete, "hands-on" nature. Notice, however, that in the construction task relations can be processed one at a time: Children can first place green on blue, then red on green, etc. By contrast, in prediction they must mentally integrate two relations "red above green" and "green above blue" to yield "red above blue." When the number of relations that had to be considered in a single decision was manipulated, with format controlled, it was found that this factor accounted for most of the variance in the task, and there was no significant residual effect of construction–prediction. Thus, format was completely subsumed under relational complexity.

The relational complexity metric has been applied successfully to a wide variety of tasks, of which those discussed in section 6 are samples. It was applied successfully to children's mathematics by English and Halford (1995). Consider, for example, the following relations between rational numbers: $1/2 = 3/6$, $1/2 < 4/6$, $1/3 < 3/6$, and $5/7 > 5/8$.

Proportion is a notoriously difficult concept for children learning mathematics, and there seems to be some uncertainty regarding this issue. However, as this example illustrates, proportions entail a quaternary relation, the variables being the two numerators and two denominators. The task is hence likely to be difficult because four dimensions must be processed. This simply illustrates that relational complexity has proved a very serviceable metric for conceptual complexity, in tasks as varied as proportion and ordering blocks.

4. Algorithmic design

The essence of the model is defined at the mathematical (computational) level and is designed to account for observed capacity limitations of higher cognitive processes. Research on neural net representation of relations, however, has uncovered limitations at least broadly consistent with those observed from psychological data. Integrating the psychological and neural net research on this question has the potential to deepen our understanding of the issue and to produce more refined questions for future research. This section considers how relations can be represented in neural nets; section 5 develops the argument that computational cost is a

function of relational complexity. Thus the underlying reason for limitations in processing capacity can be found in requirements for processing relations in neural nets. (If the reader does not wish to consider neural net modeling, however, at least in a first reading, then it is still possible to follow the paper by skipping to section 6.) The representation of relations in neural nets is currently the subject of extensive research, but even models that differ in architecture are in reasonable agreement about the nature of capacity limitations.

4.1. Neural net models.

Neural net models of relational knowledge can be categorized in two ways, namely, type of binding and type of architecture. The models of Hummel and Holyoak (1997), Plate (1995), Shastri and Ajanagadde (1993a), and Smolensky (1990) use role-filler bindings, whereas the model of Halford et al. (1994) uses symbol–argument–argument bindings (defined in sect. 2.2.1.2). Architectures can be divided mainly into models based on a product operation, either tensor product (Halford et al. 1994; Smolensky 1990) or circular convolution (Plate 1995) and models based on synchronous oscillation (Hummel & Holyoak 1997; Shastri & Ajanagadde 1993a). Other networks exist that learn to represent relations, for example, the recursive autoassociative memory (Pollack 1988) and BoltzCONS (Touretsky 1990). However, the large number of training examples needed to learn appropriate representations makes them unsuitable for models of working memory, so they are not considered here.

4.1.1. Tensor and convolution models. Tensor and convolution models represent bindings by performing some type of product operation on vectors representing bound entities. Tensor and convolution models can use either role-filler bindings or symbol–argument–argument bindings.

4.1.1.1. Role-filler bindings. In the model of Smolensky (1990) the role-filler binding is represented by the outer product of role and filler vectors, whereas in the model of Plate (1995) it is represented by the circular convolution of the vectors. Thus loves(John, Mary) can be represented in essence by $t_r = v_{role1} \otimes v_{John} + v_{role2} \otimes v_{Mary}$ or by $v_r = v_{role1} \circ v_{John} + v_{role2} \circ v_{Mary}$ where \otimes and \circ represent tensor product and circular convolution, respectively, and t_r and v_r represent the relation symbol r in the tensor and convolution models, respectively. A tensor product net that can represent a role-filler binding is shown in Figure 1A, with an arithmetic example given in Figure 1B. In these models all vectors representing roles are superimposed on a single set of units, and vectors representing fillers are superimposed on another set of units. A circular convolution of the vectors in Figure 1B is shown in Figure 1F. A circular convolution is like a compression (technically a projection) of the tensor product matrix, computed by adding along the curved lines as shown. (For a lucid explanation of circular convolution see Plate 1995.) The elements within the matrix are the *binding units*, and their activations are computed in one shot, rather than by incremental adjustment over trials as occurs in learning algorithms. Hence the matrix represents a dynamic binding in the sense that it represents the currently activated representation, rather than a product of past learning.

With circular convolution the number of elements remains constant (as illustrated in Fig. 1F); for example, the number of elements in each of v_{role1} and v_{John} is the same as the number in $v_{role1} \circ v_{John}$, whereas the tensor product of vectors with n and m elements contains nm elements (as illustrated in Fig. 1A,B). The implications of this for computational cost will be discussed in section 5.2.3.

Retrieval from circular convolution representations is noisy (Plate, in press) and requires a clean-up memory, whereas tensor product representations produce an unambiguous output provided that all the vectors form an orthonormal basis (that is, they form a basis for the vector space they span, their lengths are 1, and

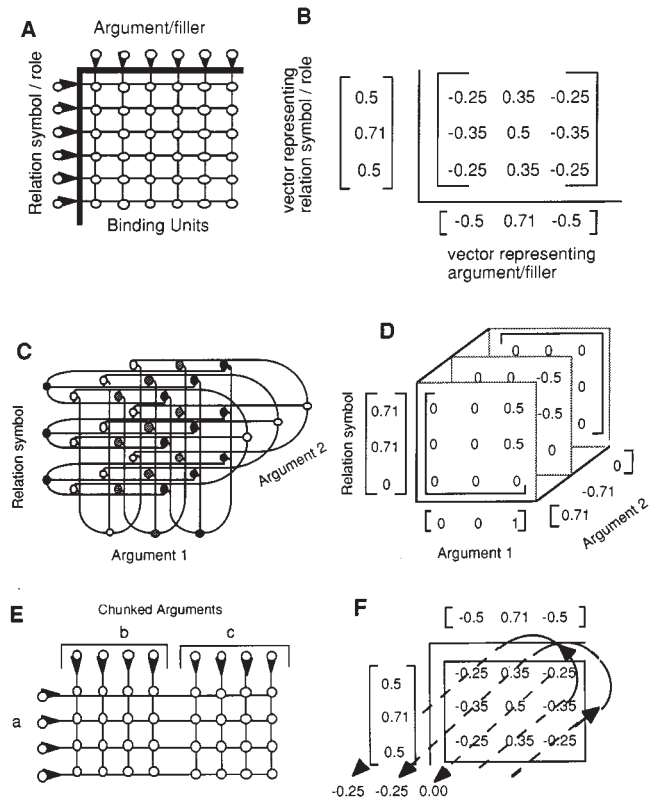


Figure 1. Neural net capable of tensor product representation of role-filler binding or unary relation (A) and as arithmetic example (B). Binary relation shown as a tensor product net (C) and as an arithmetic example (D). Ternary relation is chunked to a binary relation $R(a, b/c)$ in E (with symbol vector omitted for simplicity). A circular convolution calculated from the tensor product in B is shown in F. The circular convolution is computed by adding along the curved lines and is $[0.50 \ 0.71 \ 0.50] \cdot [-0.5 \ 0.71 \ -0.50] = [-0.25 \ -0.25 \ 0.00]$. The shadings in C are to make the spatial layout clear and do not represent levels of activation.

the inner products of distinct vectors are all 0). Though orthonormal vectors are convenient, they do not allow cross talk (interference between similar representations, or similar tasks) to be modeled. This can be done using sparse random vectors, in which similar entities share some units (Wilson et al. 1995). Then the possibility of confusion and cross talk arises.

4.1.1.2. Symbol–argument–argument bindings. Symbol–argument–argument bindings were illustrated in section 2.2.1.2. With this type of model, a relational instance is effectively represented by computing the outer product of symbol and argument vectors (Halford et al. 1994). A collection of relational instances can be superimposed on the same representation by adding the outer products. Thus, representations of loves(John, Mary) can be represented as $V_{loves} \otimes V_{John} \otimes V_{Mary}$, and loves(Tom, Wendy) can be represented as $V_{loves} \otimes V_{Tom} \otimes V_{Wendy}$. These representations can be superimposed by summing the outer products, yielding $V_{loves} \otimes V_{John} \otimes V_{Mary} + V_{loves} \otimes V_{Tom} \otimes V_{Wendy}$.

The resulting sum of outer products is referred to as a tensor. Thus the relational instance $r(a_1, a_2, \dots, a_n)$ would be represented in a tensor product space as $V_R \otimes V_1 \otimes V_2 \otimes \dots \otimes V_n$, where V_R represents alternative relation symbols including r , and v_i ($i > 0$) represents concepts appropriate to the i th argument position. A unary relational instance $r(a)$ can be represented in a rank two tensor product space $V_R \otimes V_1$. A binary relational instance $r(a_1, a_2)$ can be represented in a rank three tensor product $V_R \otimes V_1 \otimes V_2$. The net shown in Figure 1A represents a unary relation by this

method if one vector represents the relation symbol and the other vector represents the argument, and similarly for the arithmetical example shown in Figure 1B. Binary relations are illustrated by this method in Figure 1C,D. The arithmetic examples of outer products shown in Figures 1B,D show that each element in the matrix (rank two outer product) is the product of a component from each of the symbol and argument vectors. Arguments to a relation may also be regarded as role-fillers, and “argument” and “filler” are used interchangeably in this context depending on whether symbol–argument–argument or role–filler models are being considered.

Tensor product implementations of relations from unary to quaternary are shown schematically in Figure 2. In each case there is a vector representing the relation symbol and a vector representing each argument. A ternary relational instance $r(a_1, a_2, a_3)$ can be represented in a rank four tensor product space $V_R \otimes V_1 \otimes V_2 \otimes V_3$. Notice that the example in Figure 2 represents arithmetic addition and multiplication, which, as was noted earlier, are ternary relations, superimposed on the same tensor product. If two addends (multiplicands) are entered in the argument units (using the retrieval procedure described in sect. 4.2.1), and the vector representing addition (multiplication) is entered in the symbol units, the output represents the sum (product). Changing the symbol vector changes the relation that is implemented and is an example of modifiability as discussed in section 2.2.12. Our simulations have shown that addition and multiplication can be su-

perimposed in this way on a rank 4 tensor product without interference. A quaternary relational instance $r(a_1, a_2, a_3, a_4)$ is represented in a rank five tensor product $V_R \otimes V_1 \otimes V_2 \otimes V_3 \otimes V_4$, the example in Figure 2 being a composition of two binary operations.

In symbol–argument–argument models, roles are determined positionally, a type of coding also used in language. This implies that roles do not have to be explicitly represented and that role-filler bindings are unnecessary. The role to which an argument is assigned is defined by its position in the representation. In the tensor product implementation, roles are not represented by vectors, but separate sets of units are used for each argument vector. A procedure is required to ensure structural correspondence, so that arguments are represented on the correct set of units. The criteria for valid representation mentioned in section 2.2.1.1 are sufficient to ensure this.

To illustrate how roles need not be explicit if arguments are defined relative to each other, consider an instance of the ternary relation arithmetic addition, $+(3, 5, 8)$. Now, suppose we want to superimpose another instance $+(2, 4, 6)$. This can be done using tensor product representations of symbol–argument–argument bindings as mentioned above. Were we to misalign the representations by superimposing $+(2, 6, 4)$ on $+(3, 5, 8)$, that is, interchanging the second and third arguments, the error would be detected by the tests for structural correspondence (in sect. 2.2.1.1). Thus a valid relational representation can be established without roles being explicitly represented. The arguments are assigned to the correct role position by ensuring that they are correctly related to each other (in the current example this means that they are in the correct order).

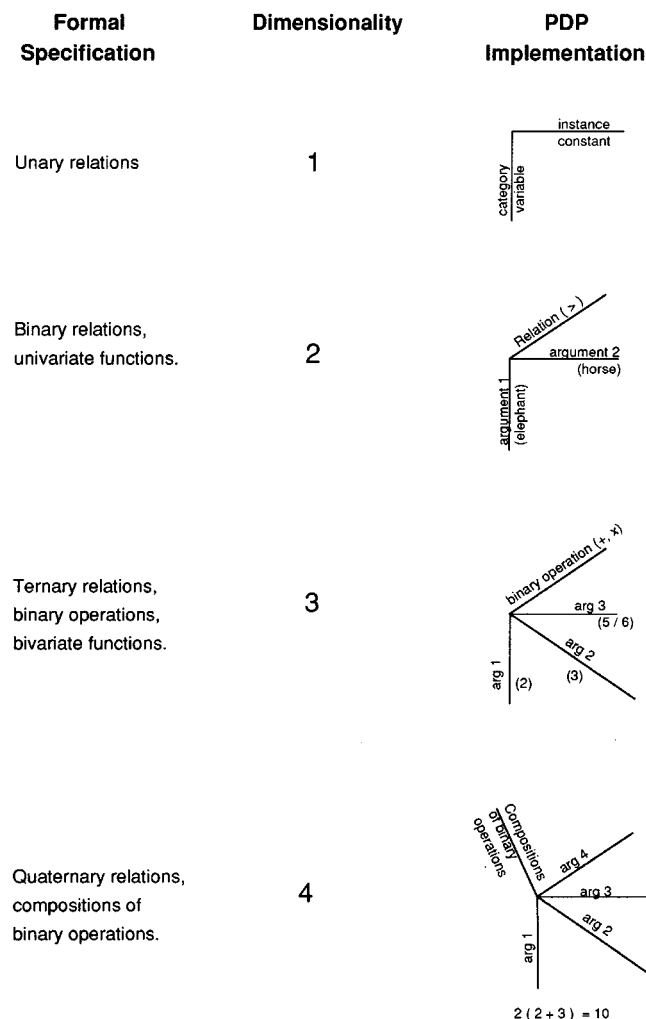


Figure 2. Four levels of relational complexity, with dimensionality and schematic outer product representations.

4.1.2. Synchronous oscillation models. In synchronous oscillation models, units representing a role oscillate in phase with units representing the filler bound to that role and out of phase with units representing other roles and fillers. The relational instance loves(John, Mary) would be represented by units representing the agent role of loves oscillating in synchrony with units representing John, while units representing the patient role of loves oscillated in synchrony with units representing Mary (see Fig. 3). However, units representing the agent role (filler) would oscillate out of synchrony with units representing the patient role (filler). The model of Shastri and Ajjanagadde (1993a) utilises synchronous oscillation for role–filler binding, but much of its power comes from additional nodes and connections. The analogy model of Hummel and Holyoak (1997) uses synchronous oscillation to perform mappings between analogs, but much of its power also comes from other systems, including a distributed semantic memory representation.

Relational instances can be superimposed on the representation, as illustrated in Figure 3. Thus kisses(John, Mary) and marry(John, Mary) can be superimposed on the representation of loves(John, Mary), by having corresponding roles of all relational instances oscillate in synchrony. Notice that no additional phases are required for the superimposed instances, and this is analogous to tensor product representations of symbol–argument–argument bindings discussed in section 4.1.1.2, where relational instances can be superimposed on the same set of vector spaces. Foreshadowing a point to be made in section 5.1, Shastri and Ajjanagadde (1993a) have shown that the major limitation is in the number of distinct phases, rather than the amount of information represented in each phase. This corresponds to the limitation in human processing capacity, which is defined by the number of arguments a relation has, rather than by total information processed.

4.1.3. Comparison of models. Tensor product and synchronous oscillation models appear to be equally capable of representing higher cognitive processes, and the similarity of their properties is at first sight somewhat surprising. However, Tesar and Smolensky (1994) have proposed that the architectures are formally reducible to one another, the primary difference being that tensor product




Figure 3. Synchronous oscillation representations of relational instances love(John, Mary), kiss(John, Mary), and marry(John, Mary).

models use spatial role vectors whereas synchronous oscillation models use temporal role vectors. Another possible explanation is that their similar properties are due to additional features designed to give them the power to simulate higher cognitive processes. As we will see, the similarity of their properties extends to the processing capacity limitations inherent in them.

It is important that, in order to account for working memory, a model must deal with relations. As was noted in section 2.2.1.2, the representation of relations by role-filler bindings requires that each relational instance be stored separately or be uniquely identified. We will now develop this point further by considering a ternary relation, the binary operation of addition. There are three roles, corresponding to the two addends and the sum, which we will represent as a_1 , a_2 , and s . Using the role-filler approach, we could bind numbers to each role, thus:

$a_1.2$ $a_2.3$ $s.5$
 $a_1.4$ $a_2.5$ $s.9$
 $a_1.3$ $a_2.4$ $s.7$
 $a_1.5$ $a_2.2$ $s.7$, and so on.

Were we to represent all addition facts in this way, then every number would be bound to every role because any number can serve as first or second addend or as sum. If all these role-filler bindings are entered into the same representation (e.g., by adding the resulting vectors, as in the models discussed above), and without specific identification of the tuples, then we cannot recover role-filler bindings or relational instances. If we ask, "What number is bound to the first addend role?" the answer is, "Every number," and the same is true for the other two roles. Furthermore, we have stored only role-filler bindings rather than relational instances, so there are no links between addends and sum. Thus even the fact that $a_1.2$ and $a_2.3$ are associated with $s.5$ is not represented, so it is not possible to access a component of the instance, given the remaining component. Thus we cannot ask, "If the addends are 2 and 3, what is the sum?" Suppose, for example, we were to identify first addend roles that are bound to 2. We can-

not then determine which of these cases have 3 bound to the second addend role, because no link has been stored between first and second addends. We cannot retrieve the sum, given the first and second addends, for the same reason.

The solution of identifying each relational instance also has its problems, first because individual identification of every relational instance is implausible when the number of instances is very large. A relational instance such as $2 + 3 = 5$ is identified by its content (e.g., $2 + 3 = ?$; $? + 3 = 5$) rather than by an index, such as a context vector, that identifies the relational instance.⁵ It is implausible that every addition fact we know is individually identified. Second, notice also that, even were instance identification to be used, every role-filler binding in a relational instance would require the same identifier. Thus, if we identify this instance as $add(2, 3)$, we must bind the identifier to all three role-filler bindings, thus:

$add(2, 3): (add2, 3).a_1.2 (add2, 3).a_2.3 (add2, 3).s.5$

In other words, the context in which the three role-filler bindings are learned/memorized must be sufficiently stable to result in the same identification vector across all three role-filler pairs. Notice also that this representation bears a close resemblance to symbol-argument-argument bindings. Furthermore, the identifier increases the computational cost of the representation, and appears to require an additional rank in the outer product (i.e., rank three rather than rank two, as in Smolensky's 1990 model).

Contrast this with representations based on symbol-argument-argument bindings. Omitting the relation symbol (as with role-filler bindings), we would represent the same facts as follows:

2.3.5
 4.5.9
 3.5.8, and so on.

The computational cost is high for one relational instance (a ternary relation requires a rank 4 tensor, including the relation symbol vector), but there is no increase in cost for further relational instances because they can be superimposed, the tuples are inherently identified by the bindings, the links between addends and sums are represented, and the other properties of relational knowledge are implemented, as explained in section 4.2. Thus, the initial cost of symbol-argument-argument bindings is high, but their power is considerable.

Although the synchronous oscillation models of Shastri and Ajanagadde (1993) and the tensor product symbol-argument-argument binding model of Halford et al. (1994) are very different, they have a common property that is important to capacity limitations. They both map dimensions of the relations (as defined in sect. 2.3) to separate dimensions of the representation. In the synchronous oscillation model each argument is assigned to a separate phase in the oscillation (as illustrated in Fig. 3). In the symbol-argument-argument binding model each argument is assigned to a separate vector space (illustrated in Fig. 1C,D). This means that the dimensions of the relation are mapped directly into phases of oscillation or into vector spaces. The role-filler models based on tensor products⁶ (Smolensky 1990) or circular convolution (Plate, in press) do not have this property, because all roles are superimposed, as are all fillers (see sect. 4.1.1.1). As we will see in section 5, models that map dimensions of the relation to dimensions of the representation imply similar capacity limitations.⁷

4.2. Modeling relational knowledge

The manner in which these models implement the properties of relational knowledge defined in section 2.2 will now be considered. The role-filler model of Smolensky (1990), based on tensor products, handles the storage and retrieval of relational instances, but in its original form it does not appear to incorporate the other features of relational knowledge. The symbol-argument binding model (Halford et al. 1994) was based on Smolensky's tensor prod-

uct formalism, but with modifications and extensions to handle all features of relational knowledge. The circular convolution model of Plate (in press) handles storage and retrieval of relational instances and gives a good account of similarity, but it does not appear to handle conceptual chunking (see sect. 4.2.4), nor does it provide a general solution to systematicity (see sect. 4.2.9). Role-filler binding models based on synchronous oscillation (Hummel & Holyoak, in press; Shastri & Ajjanagadde 1993a; 1993b) appear to have been designed to incorporate the properties of relational knowledge discussed in section 2.2. We will emphasise those models that have been designed to implement the properties of relational knowledge.

4.2.1. Retrieval of information. Information stored in a tensor memory can be retrieved by representing a question as a tensor product and computing the inner product (dot product) of the question and memory tensors. The query is a partial relational instance and can be expressed as an outer product with one entity deleted. For example, given $r(a_1, a_2, a_3)$ stored as part of the rank four tensor T_{pqrs} in the tensor product $V_R \otimes V_1 \otimes V_2 \otimes V_3$, then, say, a_3 can be retrieved by computing the generalized inner product $v_r \otimes v_{a_1} \otimes v_{a_2} \otimes \dots \cdot T$, where v_r is the vector representing the relation symbol r , v_{a_1} is the vector representing the argument a_1 , and v_{a_2} is the vector representing the argument a_2 . Generalized inner products are described in Appendix B: the underscore signifies the component of the tensor that is "retrieved" by computing this generalized inner product. Effectively, the query $r(a_1, a_2, ?)$ has been used as input to the tensor memory, and a_3 has been obtained as output. The details of how this generalized inner product may be computed are contained in Appendix B, where it is listed as operation (3). An analogous procedure is specified for synchronous oscillation models by Shastri and Ajjanagadde (1993a, sect. 4.3).

4.2.2. Truth value. The truth value of a proposition can be assessed by matching against memory, in what is essentially a recognition process. The proposition bark(cats) can be represented by a rank two tensor product, which can be matched against semantic memory by computing a generalized inner product (dot product) of the tensor with the representations in semantic memory (Humphreys et al. 1989). The relational instance bark(cats) is treated as a query by representing it as the tensor $v_{\text{bark}} \otimes v_{\text{cats}}$ as shown above, and the dot product of this tensor and tensor representations in memory is computed, as described in section 4.2.1 (the procedure corresponds to operation (0) in Appendix B). This can be done in parallel for superimposed memories. If the product is nonzero, the proposition is recognized. Thus bark(cats) and bark(dogs) would produce zero and nonzero dot products respectively, so the latter is recognized, whereas the former is not. Note the following example.

$$\begin{aligned} v_{\text{bark}} \otimes v_{\text{cats}} \cdot (v_{\text{bark}} \otimes v_{\text{dogs}} + \dots + v_{\text{sing}} \otimes v_{\text{birds}}) &= 0 \text{ but} \\ v_{\text{bark}} \otimes v_{\text{dogs}} \cdot (v_{\text{bark}} \otimes v_{\text{dogs}} + \dots + v_{\text{sing}} \otimes v_{\text{birds}}) &> 0 \end{aligned}$$

The procedure defined by Shastri and Ajjanagadde (1993a, sect. 4.4) provides a way of assessing the truth value of a proposition in synchronous oscillation models.

4.2.3. Relation symbols. Relation symbols are represented as separate vectors in the vector space V_R . In synchronous oscillation models, the symbol can be represented as a unit firing in a separate phase in the oscillation or by an additional node connected to role and filler nodes.

4.2.4. Conceptual chunking. Conceptual chunking serves to reduce the rank of a tensor product representation of a relation. It can be implemented by convolution, concatenation (illustrated in Fig. 1E), or superposition (in which vectors representing arguments are added), or by defining a special function that associates an outer product to a single vector. The outer product representation of $r(a, b, c)$ can be reduced to $r(a, b/c)$, by concatenating or

convolving vectors b and c into a single vector. Features of b and c can still influence the computation of the relation with a because activation can be propagated from units in b and c to a (Fig. 1E), but the representation functions as a binary relation, and neither the relation between b and c nor the three-way relation between a, b , and c is directly accessible.

Unchunking can be achieved by differentiating vectors into other vectors. Algorithms for this have been defined in the STAR analogical reasoning model (Halford et al. 1994; 1995; 1996). In general, lower rank representations can be differentiated, yielding more complex relations. For example, in Figure 2E, if the vector representing b/c were differentiated into separate vectors representing b and c , and if all four vectors including the relation-symbol vector (not shown) were then appropriately interconnected, as for a rank four tensor product, a ternary relation could be represented.

A chunked representation is wholistic in that features are represented but are not differentiated into dimensions. Many concepts are wholistic initially and progress to dimensional representation (Smith 1989). This is like unchunking in that it entails differentiation of a vector into two or more vectors and representation of the relation between them.

It is unclear how Plate's (1995) circular convolution model would handle conceptual chunking, at least without significant additions. Chunking involves a compression of a relational instance into an unstructured whole, so the relations between components become inaccessible. The circular convolution, however, is already a compression (a projection of the tensor product), and it is not clear how a further compression that incorporates the psychological properties of conceptual chunking could be achieved. A further problem is that circular convolution relies on component vectors randomly generated from a Gaussian or uniform distribution. This has the effect that there is no similarity (as measured by the dot product) between chunked and unchunked representations. Thus, for example, features from b and c would not, and in general could not, influence computation of the relation with a in $R(a, b/c)$.

Hummel and Holyoak (1997) represent the equivalent of a chunk in a synchronous oscillation model by having units that represent part or all of a proposition. For example loves(John, Mary) can be represented as features, as roles and fillers, or as an intact proposition. In the latter case, it can be an argument to a proposition such as knows[Sam, loves(John-Mary)].

4.2.5. Higher order relations and hierarchical structures. Higher order relations and hierarchical structures can be modeled by representing higher order relational instances with chunked lower order relational instances as arguments. Consider this relational instance.

cause[shout-at(John, Tom), hit(Tom, John)].

The relational instance shout-at(John, Tom), normally represented as a rank three outer product in our model, is chunked into a single vector shout-at₁, as described in section 4.2.4, and hit(Tom, John) is chunked similarly as hit₁. The higher order relation cause(shout-at₁, hit₁) is then represented as a rank three outer product.

The *repeated variable constraint* requires that fillers be bound to the correct roles, as was pointed out in section 2.2.5. The STAR analogy model (Halford et al. 1996) can achieve this by ensuring that hierarchical structures are in correspondence. Consider the following relational instances.

cause[shout-at(John, Tom), hit(Tom, John)]

cause[shout-at(Mary, Wendy), hit(Wendy, Mary)]

These would be represented as chunked relational instances, as described above. The model maps one level of the hierarchy at a time, then moves to another, usually lower, level and recursively matches corresponding arguments of source and target. Thus the model would first map cause(shout-at₁, hit₁) to cause(shout-at₂, hit₂). It would then unchunk shout-at₁ and shout-at₂ and map

shout-at(John, Tom) to shout-at(Mary, Wendy). The model has a bias to maintain the mappings of John to Mary and Tom to Wendy when processing other parts of the structure. It would map hit(Tom, John) to hit(Wendy, Mary), in keeping with previous mappings, thereby maintaining structural consistency. It would also compute a goodness-of-mapping score that reflects degree of structural correspondence. The score would be higher for this mapping than for the inconsistent mapping below.

cause[shout-at(John, Tom), hit(Tom, John)]
cause[shout-at(Mary, Wendy), hit(Mary, Wendy)]

The model enforces the repeated variable constraint as a consequence of maintaining structural consistency. Because the person bound to the agent role of “shout-at” is bound to the object role of “hit” in the source, this constraint is maintained in the target because of biases in the algorithm to ensure structural correspondence between base and target. Shastri and Ajjanagadde (1993a, sect. 4.5) provide a synchronous activation-based mechanism that enforces the repeated variable constraint.

4.2.6. Omnidirectional access. Omnidirectional access is implemented by the retrieval process described in section 4.2.1 because a query can be composed of a relational instance with any component missing. Thus a ternary relation represented as $v_R \otimes v_1 \otimes v_2 \otimes v_3$ can be queried by any of the following means.

$$\begin{aligned} v_R \otimes v_1 \otimes v_2 \otimes \dots \cdot v_R \otimes v_1 \otimes v_2 \otimes v_3 &= v_3 \\ v_R \otimes v_1 \otimes \dots \cdot v_R \otimes v_1 \otimes v_2 \otimes v_3 &= v_2 \\ v_R \otimes \dots \otimes v_2 \otimes v_3 \cdot v_R \otimes v_1 \otimes v_2 \otimes v_3 &= v_1 \\ \dots \otimes v_1 \otimes v_2 \otimes v_3 \cdot v_R \otimes v_1 \otimes v_2 \otimes v_3 &= v_R \end{aligned}$$

The procedure for answering wh-queries specified by Shastri and Ajjanagadde (1993a, sect. 4.7) essentially embodies the omnidirectional access property.

4.2.7. Role representation. The role that an argument fills can be indicated by its position relative to other arguments, as discussed in section 2.2.7, and its implementation in symbol–argument–argument bindings is described in section 4.1.1.2. The synchronous oscillation model of Shastri and Ajjanagadde (1993a) uses role-filler bindings as described in section 4.1.2.

4.2.8. Decomposability of relations. The relation represented can be decomposed into the derived relations by replacing the vector in any role with a special vector, namely, the sum of all the basis vectors used to represent fillers on that axis of the tensor. Thus a representation of the ternary relation $R(x, y, z)$ can be reduced to representation of $R_3 = (x, y)$ by entering this special vector on the units representing z . Such “collapsing” of a representation to a lower rank has been used in models of memory (Humphreys et al. 1989) and of analogical reasoning (Halford et al. 1994).

In the tensor product representation of an n -ary relation with instances $r(a_1, a_2, \dots, a_n)$, the effect of variations in any proper subset of $\{a_1, a_2, \dots, a_n\}$ on the remaining argument(s) can be computed. For example, suppose that one wishes to use the fixed value b in the final role (the a_n role) and consider the induced $(n - 1)$ -ary relation $Ra_n = b = \{(a_1, a_2, \dots, a_{n-1}) \mid (a_1, a_2, \dots, b) \in R\}$. This effect can be achieved by clamping the value in the a_n role of the tensor network to be b . This can, of course, be done with any role, not just the a_n role, and can be iterated so that, eventually, any desired set of roles is fixed in this way. For synchronous oscillation models, the representation of partially instantiated relations given by Shastri and Ajjanagadde (1993a, sect. 3.1) effectively decomposes relations in an analogous manner.

4.2.9. Relational systematicity. Relational systematicity can be handled by using higher order relations, as described in section 4.2.5. For example, implies[>(a, b), <(b, a)] can be represented by the tensor product of vectors representing implies and chunked representations of >(a, b) and <(b, a). Systematicity is achieved in the synchronous oscillation model of Shastri and Ajjanagadde

(1993a, sect. 4.2) by connections that ensure that corresponding arguments oscillate in synchrony, for example, that the first role representation in >(a, b) oscillates in synchrony with the second role-representation in <(b, a).

The circular convolution model of Plate (1995) incorporates systematicity, but there is some doubt regarding the generality of the procedure used. To allow the model to recognize the structural similarity between “Spot bit Jane, causing Jane to flee from Spot” and “Felix bit Mort, causing Mort to flee from Felix” (by contrast with the superficially similar, but structurally dissimilar, “Rover bit Fred, causing Rover to flee from Fred”), Plate used contextualized representations. These entailed adding the property “flee-from” to the representation of Spot, Felix, and Rover and the property “bite-object” to the representation of Jane, Mort, and Fred. This handles some structurally similar higher order relational instances but depends on representing dogs as entities people flee from and people as entities dogs bite. This approach lacks plausibility in relational instances such as “Jane smiled at John, causing John to like Jane” because it is implausible that smiling should be part of the representation of Jane (she may not always smile, even at John) or that liking should be part of the representation of John (he may not always like people). The circular convolution model appears to require additional means of representing structure in order to handle systematicity, and the computational cost of these additions is unknown.

4.2.10. Dimensionality of relations. The dimensionality of a relation was defined in section 2.3 as the number of arguments. In symbol–argument–argument tensor product representations, a separate vector is used for the relation symbol and for each argument, so rank is one more than dimensionality. We have used the convention of specifying number of arguments by n , and we will use the convention of specifying ranks by k . Therefore, in this type of model, $k = n + 1$. The *components of a representation* are the relation symbol and the argument representations, so the number of components = k . Furthermore, if $k - 1$ components are known, there is at least some potential to predict the k th component (illustrated in sect. 2.3), so dimensionality = $k - 1 = n$. Even if relations are represented by formalisms other than tensor products, symbol and arguments must be represented independently of each other, or be individually identified, so that they retain their identity when linked (bound) to other components. Note also that, in the context of neural net models, the dimensionality of a relational concept should not be confused with the number of elements in a vector, which is also sometimes referred to using the term “dimension.” As was noted in section 4.1.3, the models of Halford et al. (1994) and Shastri and Ajjanagadde (1993a) map dimensions of relations directly into dimensions of representations, whereas other models do not.

4.2.11. Analogy, planning, and modifiability. Analogy can be successfully modeled using the tensor product representations of relations outlined in section 4.1 (Halford et al. 1994; 1995; 1996). A sophisticated model of analogy based on synchronous oscillation has been presented by Hummel and Holyoak (1997).

With symbol–argument–argument models based on tensor products, relations can be modified online by changing the relation symbol, which selects a new set of relational instances. Relational instances are stored as outer products of symbol and argument vectors, and outer products are summed to form a tensor. We will illustrate with arithmetic addition and multiplication. Addition would be stored as $v_{\text{add}} \otimes v_2 \otimes v_3 \otimes v_5 + v_{\text{add}} \otimes v_3 \otimes v_6 \otimes v_9 + \dots$, and multiplication would be stored as $v_{\text{mult}} \otimes v_2 \otimes v_3 \otimes v_6 + v_{\text{mult}} \otimes v_3 \otimes v_6 \otimes v_{15} + \dots$. Changing the symbol vector from v_{add} to v_{mult} selects a new set of relational instances and changes the mappings between addends and sum or product.

4.2.12. Strength. Strength can be represented by multiplying the outer product representing the relational instance by a scalar, before adding it to the tensor. Typically the scalar would have a value

between 0 and 1 indicating how frequently the relational instance is found to be true; for example, *bigger-than*(dog, cat) would have a scalar a little bit less than 1 to take account of the small minority of dogs (e.g., chihuahuas) that are smaller than cats.

4.2.13. Operations on relations. Operations on relations can be implemented using tensors. For completeness we provide one tensor implementation for each of the relational operators. The simplest implementation assumes local (with all unit values 0 except for a single unit with value 1) argument vector representations. Relaxing this assumption introduces other properties, such as cross talk, but at the expense of not implementing exact analogues of relational operators. Under the local vector assumption, then, the set operators *union*, *intersection*, and *difference* are implemented by pairwise addition, multiplication, and subtraction (respectively) of binding units with the same index. The upper limit on activation eliminates multiple occurrences of the same element (consistent with set *union*), and the lower limit prevents subtraction of nonexisting elements (consistent with *difference*).

The *select* operator retrieves relational instances satisfying condition C. Insofar as C is an arbitrary Boolean expression, the *select* operator is very general and powerful. For our purposes, however, we consider a restricted version, in which C has the form of a conjunction of filler-role pairs: $(a_1, A_1) \wedge \dots \wedge (a_m, A_m)$ (i.e., *select* with filler a_1 at role A_1 and filler a_2 at role A_2 , etc.). The corresponding tensor implementation is to compute the outer product of the fillers a_1 to a_m at the specified tensor axes A_1 to A_m , respectively. Axes with unspecified fillers use a special filler vector $I = (1 \dots 1)$. Thus, the rank of the tensor (T_C) representing the condition C is the same as the rank of the tensor (T_R) representing relation R. Next, we perform a pairwise multiplication $T_C \cdot T_R$, resulting in a tensor (T_S) representing the selected relational instances.

The *project* operator returns the relation between components at the specified roles. The equivalent tensor operation is summation onto the corresponding axes. Formally, given a relation R with attributes (roles) $A_1 \dots A_k$ and a corresponding tensor T with axes labeled $A_1 \dots A_k$, then $\text{project}_A R$, is implemented as $\sum_{\langle A_1 \dots A_k \rangle - A}$, where A is the list of projected attributes (or tensor axes) and $\langle A_1, \dots, A_k \rangle - A$ is the difference of the two lists (i.e., sum onto the axes not in the list of projected attributes). The rank of the resulting tensor is the same as the *-arity* of the projected relation.

Often one wants to cue a *k*-ary relational memory with *k* - 1 components to retrieve the target at the *k*th role. At the relational level, the target is retrieved by successive application of the *select* and *project* operators. The *select* operator retrieves an instance(s) containing all *k* - 1 components at the specified roles, and the retrieved instance(s) is applied to the *project* operator, which returns the target at the *n*th role. At the tensor level, this combination of *select* and *project* is realised by taking the inner product of the *k* - 1 components with the tensor, resulting in a vector representing the retrieved target(s). This tensor operation is specified in section 4.2.1.

The outer join of two relations is simply the outer product of the corresponding tensors. The equijoin, however, is more complicated; it requires joining only those instances that share a common argument at the specified roles. Suppose the relation *Taller* = {(John, Mary), (John, Bob), (Mary, Tom)} and its corresponding tensor $T = v_{\text{John}} \otimes v_{\text{Mary}} + v_{\text{John}} \otimes v_{\text{Bob}} + v_{\text{Mary}} \otimes v_{\text{Tom}}$ (the relation symbol is not part of the join operation). One way of implementing the equijoin of the *Taller* relation with itself is to, first, cue the tensor T, with a possible argument at the Person1 role (e.g., $v_{\text{Mary}} \cdot T$). This results in the vector v_{Tom} . Then, cue the tensor again, but at the Person2 role (i.e., $T \cdot v_{\text{Mary}}$). This results in the vector v_{John} . Provided one maintains the cue vector v_{Mary} and the two retrieved vectors v_{John} and v_{Tom} , one can construct the tensor representing the equijoin as $v_{\text{John}} \otimes v_{\text{Mary}} \otimes v_{\text{Tom}}$. Other instances are retrieved in the same manner by cueing with different fillers at the joined roles and adding the result to the tensor representing the equijoin.

To summarise section 4, the properties of relational knowledge can be incorporated in neural net models based either on symbol-argument-argument bindings using tensor products or on role-filler bindings using synchronous oscillation. Representing these properties effectively depends on a number of additional features, but the basic properties of the representations are important. Despite their differences, both types of model have the property that the dimensions of a relation, the symbol and arguments (fillers), are mapped into dimensions of the representation, either vector spaces or phases of oscillation. This means that components of the relation are represented as intact entities, which retain their identity in the binding, and this is a major factor in the computational cost of representing relations.

5. Relational complexity and processing load

So far we have been considering properties of human cognitive functioning, with the aim of accounting for processing capacity limitations observed from psychological data. We have defined *cognitive complexity* intuitively in terms of the number of interacting variables represented in parallel and have conceptualized it in terms of the number of arguments in a relation. However, we wish to explain processing capacity limitations, and two approaches to neural net modeling of relational knowledge have independently identified possible explanations (Halford et al. 1994; Shastri & Ajjanagadde 1993a). To explain how processing loads are imposed by relations, we must consider *computational complexity*, which refers to the amount of computation required to perform a task.

Complexity analysis at the computational level is a very general, potentially algorithm-independent method of determining the inherent difficulty of a particular problem. The classic results from computer science have been to identify two very broad classes of problems, called P and NP. In the most general terms, the complexity of the former is a polynomial function of some measure of the input, whereas for the latter it is an exponential (or worse) function. Intuitively, an NP-complete problem is intractable, in the sense that the time required by any known algorithm to solve the problem grows explosively with the size of the problem (the *n*). NP-complete problems can be approached in a number of ways, including using an approximate/heuristic algorithm, avoiding large instances of the problem, or considering only subclasses of the problem. An algorithm-independent analysis is performed by showing that the problem can be transformed in polynomial time to another problem of known complexity. The power of this method was demonstrated by Tsotsos (1990) with respect to vision, by showing that certain problems in vision transform to NP-complete problems.

Whereas analysis at this level has been successful with vision, it does not seem to capture processing capacity limitations in cognitive tasks such as reasoning and language. The paradox is that, although aspects of vision are intractable by this analysis (Tsotsos 1990), vision tasks do not appear to impose the kinds of demands defined in section 2.1, which have been observed in higher cognition processes. The computationally complex tasks of vision appear to be performed without measurable processing demands of the kind discussed in sect. 2.1, the standard explanation being that the visual system is a module with high capacity for specialised input. By contrast, many computationally simple tasks in higher cognition impose high processing demands. Ordinary arithmetic, for example, requires relatively little computation, but imposes a high cognitive demand on the human performer; even such computationally simple tasks as transitive inference problems, in which for example $>(a, b)$ and $>(b, c)$ have to be integrated into monotonically-larger(a, b, c) impose a measurable processing load on adult humans. Tsotsos (1990) has shown that visual search is inherently complex because of the combinatorial explosion that occurs as the number of elements to be searched increases. No such combinatorial explosion occurs in ordinary arithmetic operations or transitive

inferences. We must seek the explanation for observed processing demands of such tasks in the architectures employed in higher cognitive processes, for which algorithmic complexity is more relevant.

Algorithmic complexity analysis considers how many steps are required, or how much space is required to compute a given problem for a given algorithm. In the former case, complexity depends on the function that links the number of computational steps required to the size (or length) of the input. A linear time algorithm will complete in $\Theta(n)$ steps (i.e., of the order of n steps), where n is size of problem (e.g., number of inputs)⁸. A polynomial-time algorithm will complete in $\Theta[p(n)]$ steps for some polynomial $p(n)$.⁹ An exponential-time algorithm will complete in $\Theta[c^{p(n)}]$ steps, so the number of steps grows explosively with size of input. Problem complexity is the least of the complexities of all algorithms to solve the problem.

Both synchronous oscillation and tensor product/convolution models predict limitations on the complexity of relational schemas that can be activated in parallel, although the bases of the limitations are somewhat different. We will examine both types of model in an attempt to find explanations for the limitations observed in the psychological data reviewed in section 3.3.

5.1. Synchronous oscillation models

Synchronous oscillation models are limited by the number of distinct oscillations. This is determined by the ratio of the period of oscillation to the window of synchrony (which is related to the duration of peak and is approximately the maximal temporal spacing between peaks that are recognized as in phase). This ratio is estimated by Shastri and Ajjanagadde (1993a) to be about five, and by Hummel and Holyoak (1997) to be four to six (to illustrate, notice that in Fig. 3 approximately five distinct oscillations would be possible). Given the criteria for relational knowledge set forth in section 2.2, five distinct entities would permit quaternary relations (a relation symbol and four arguments) to be represented without cross talk. However, Shastri and Ajjanagadde suggest that up to 10 entities could be related with cross talk. Psychological data discussed in section 3.3 appear to correspond to Shastri and Ajjanagadde’s prediction of capacity without cross talk, possibly because performance criteria used in experiments (e.g., low error rates to facilitate analysis of latencies) would tend to preclude cross talk. However, as was noted in section 4.1.2, the power of models by Shastri and Ajjanagadde and by Hummel and Holyoak depends on additional features, and there does not appear to be any way of calculating the cost of these using computational complexity theory.

5.2. Tensor product models

Tensor product models entail a computational cost in space and time. We will consider tensor product representations of relations using symbol–argument–argument bindings as proposed in section 4.1.1.2, focusing on the process of accessing the k th component of a relation, given the $k - 1$ other components, as described in sections 2.2.6 and 4.2.6. Then we will consider role-filler bindings, first based on tensor products to facilitate comparison and because existing circular convolution models do not appear to incorporate all properties of relational knowledge (as was noted in sect. 4.2). Next we will consider circular convolution models insofar as they can be directly compared with tensor product models. Computational cost can be considered from either a parallel or a sequential processor point of view, but the former is more appropriate to emphasise here.

5.2.1. Complexity for symbol–argument–argument bindings.

We will consider time complexity and then space complexity for the symbol–argument–argument binding models.

5.2.1.1. Time complexity. In the parallel processing model, one assumes that there is a processor for each unit of the vectors repre-




Figure 4. A three-level binary cascade adder.

sented symbol and argument(s), a processor for each binding unit, and some addition units, to be described below. To access the k th component of a relational instance, given arguments $a_1 \dots a_{k-1}$, it is necessary to propagate the component values to all the relevant binding units (one step); each binding unit then multiplies its binding memory contents with the values propagated for the $k - 1$ arguments; this requires $k - 1$ multiplications. Then, it is necessary to add all these products. How long this takes depends on the rank of the tensor and on the length of the vectors; let us suppose all vectors are of length n (so that there are n^k binding units in the tensor network). It will be necessary to add up n^{k-1} products to form each component of the symbol output. This is done by the addition units referred to above. The most rapid way to add many items with many processors is to cascade the additions (see Fig. 4 for a binary cascade adder).

This arrangement adds 2^3 items in three steps. In general, m items can be added together in $\text{ceiling}[\log_b(m)]$ steps, where b is the fan-in of the adders (two in the diagram). If we assume enough processors in the addition unit pool, then the addition step requires $\text{ceiling}[\log_b(n^{k-1})] = \text{ceiling}[(k-1)\log_b(n)]$ steps, for a total of

$$k + \text{ceiling} [(k - 1)\log_b(n)] \tag{2}$$

steps. If b is made large enough, then the second term can be made small, though it is always at least one.

Neurons may have on the order of 10,000 input connections, but, because not all connections may be appropriate to the computation, 10,000 should probably be regarded as the upper limit. A two-level cascade of 10,000-to-1 adders permits addition of 100 million binding units, so we could approximate the second term by the constant 2. Thus at most $k + 2$ steps are required to access a missing component of a relational instance.

Similar computations for a sequential implementation to access the k th component of a relational instance, given arguments $a_1 \dots a_{k-1}$ yields the expression $(2k - 1)n^k$ for the number of steps. The important finding, however, is that with full parallel implementation the tensor product representation is not expensive in terms of time (number of steps), but its spatial complexity is quite large, as shown below.

5.2.1.2. Space complexity. The basic requirement is for the n^k binding units of the rank k tensor and the kn input/output units. In addition to this, in a parallel implementation there would be a need for cascade adders for each side of the tensor network. Assuming that a two-stage cascade is adequate [i.e., that $b^2 \geq n^{k-1}$, so that $b \geq n^{(k-1)/2}$], each cascade would use at most $b + 1$ adders, and there would be n cascades per side (one for each component) and k sides: at most $nk(b + 1)$ adders in all. The n^k binding units dominate the space complexity, providing of course that b is not larger than necessary, that is, not significantly larger than $n^{(k-1)/2}$. Therefore, the limiting factor with this representation is the number of binding units, which increases exponentially with dimensionality.

The representation of a single relational instance is quite expensive in terms of space, requiring n^k units to represent that in-

stance, but all combinatorially possible other relational instances can then be represented in the same tensor. (Here n is the length of the vectors, and k is the number of vectors, i.e., one more than the number of arguments). Thus superposition does not incur additional computational cost.

To achieve dynamic binding, the binding units must be interpreted as activations, as was noted in section 4.1.1.1, and activations demand processing resources (Just & Carpenter 1992). Hence the rank of relations will be limited by resources available, that is, by capacity as defined in section 2.1.

This is a soft limit, because tensor product representations have the property of graceful degradation (Wilson & Halford 1994). More recent simulations in our laboratory have extended this finding to tensor product networks of ranks up to seven. For example, a rank five tensor of side 16 (i.e., $n = 16$, $k = 5$), with up to 93.75% of the binding units deleted, reliably distinguished stored facts (relational instances) from nonfacts. Such a tensor has the same number of active binding units as an intact rank four tensor with side 16. Our results suggest that the robustness depends on the ratio of number of facts stored to number of binding units: the lower the ratio, the more robust the network. Provided that the value of n , the number of components in each representation vector, is reasonably large (32 was typically adequate in our simulations, for tensors of rank three and up, and up to 4,000 facts stored), it appears to be possible to simulate a rank $k + 1$ tensor with the number of binding units available to an (intact) rank k tensor, for $k = 2-6$ (at least). Another way of looking at this is to say that the apparently very regimented architecture of a tensor product network is not necessary in order to achieve acceptable memory performance: 85% or more of the binding units can be removed (i.e., caused to have zero output) with impunity.

5.2.2. Complexity for the role-filler method. Relations can be represented using role-filler bindings, as explained in section 4.1.1.1, provided that relational instances are identified. We will assume this is done using a separate set of units, because of the implausibility of an identifying code. Smolensky (1990) used tensor products, and Plate (1994, Appendix I) used circular convolution. For the purposes of direct comparison with section 5.2.1, we calculate the time and space complexity of accessing relations using the tensor product. Again, we consider the case of accessing the k th argument of a k -ary relation given $k - 1$ arguments.

5.2.2.1. Time complexity. Given that the complex cue has already been composed, then the role-filler method has two major steps: (1) determining the tensor (representing the relational instance) with the highest similarity (dot product) to the complex cue and (2) retrieving the target from that tensor.

Assuming k roles and n fillers, then each relational instance requires nk binding units. Furthermore, because each relational instance is represented by a separate set of units, the dot product of the cue with each instance can be performed in parallel. Therefore, the time to compute the highest match is: one step to propagate the activations of each cue unit to each tensor unit plus one step to multiply cue and tensor elements (pairwise multiplication step of the dot product) and at most two steps to sum the activations of each multiplication (summation step of the dot product), assuming a fan-in of at most 10,000 units as for the symbol-argument method. Theoretically, a winner-take-all network can compute the highest match in one step assuming exponential functions and complete interconnection between competing units (Yuille & Geiger 1995). However, with limited fan-in, at most 10,000 relational instances can be compared in parallel. Once the winning tensor is determined, it must be reinstated into working memory so that the target component can be retrieved. We will assume that this takes one step. In all, at most six steps are required, for relations of less than 10,000 instances.

The time required to access the target component is: one step to propagate the target's role activations to the binding units, one step to multiply the role and tensor activations, plus one step to

sum activations (since we can assume that the number of roles will be small). The total time for the role filler is nine steps.

Although the time complexity is independent of the number of roles, we have not considered the time to compose the cue, which is necessarily $\Theta(k)$ steps, because each of the $k - 1$ arguments in the cue must be presented serially over the same group of binding units. The time complexity for composing the cue in the symbol-argument method depends on whether the cue arguments are presented to each separate group of units in parallel or in series. The total time to compose the cue and retrieve the target is k steps (parallel cues) or $2k - 2$ steps (serial cues). Either way, the time for both methods is low (i.e., linear in k).

5.2.2.2 Space complexity. Each relation is represented as the sum of outer products of role and filler vectors. This requires a rank two tensor with n possible fillers by k possible roles. Hence the number of units required to represent any one relational instance is nk . Under the assumption that each relational instance is stored on a separate group of units, the total number of units needed to store the entire relation is $|R|nk$, where $|R|$ is the number of instances in the relation. In addition, we require nk units to compute the complex cue, and $|R|$ winner-take-all units to compute the winning instance. Finally, we need nk units to store the retrieved instance, and n units for the retrieved argument. In total, there are $|R|nk + nk + |R| + nk + n = (|R| + 2)nk + n$ units [i.e., $\Theta(|R|nk)$]. Clearly, the space complexity depends on the growth in number of instances as a function of n and k . In the worst case ($|R| = n^k$ - all possible instances), the space complexity is $\Theta(kn^{k+1})$. In the best case ($|R| = 1$ - all one-instance relations), the space complexity is $\Theta(nk)$. Average case depends on knowing how many instances are likely to be in any one relation. Under the condition that working memory can store at most Θ , then tasks requiring storage of more instances would force a serial strategy. Under this scenario, the savings in space are traded for an increase in time.

5.2.3. Role-filler models using circular convolution. As was noted in section 4.2.9, Plate's (in press) circular convolution model does not appear at present to incorporate all properties of relational knowledge, and the computational cost of the additional features required cannot be estimated. Nevertheless, we will examine time and space complexity of circular convolution models for the insights that can be obtained.

The time complexity for circular convolution is the same as for the role-filler tensor method discussed in section 5.2.2.1. Assuming appropriately connected units for implementing circular convolution, each role-filler convolution requires one step to propagate activation for each argument, with the remaining processing requiring only a constant number of steps. Thus, the time complexity is still $\Theta(k)$.

Space complexity for circular convolution models is more difficult to determine. Plate (1994) used 840 unique role-filler combinations superimposed over 512 units. Thus, circular convolution permits more role-filler pairs than there are units. However, this method assumes a clean-up memory that does not appear to have been implemented as a neural net. To avoid ambiguity, each relational instance must be represented on a different set of units (unless we assume an implausible label attached to each relational instance). The required clean-up memory has essentially the same form as the tensor product implementation of role-filler bindings, the complexity of which is discussed in section 5.2.2. In the worst case, when most relational instances must be stored, complexity of the role-filler representation is worse than that for symbol-argument representations. Plate's (1994) circular convolution model has achieved interesting results, and the approach has much potential, but its restricted ability to implement the properties of relational knowledge and the requirements of the clean-up memory mean that in the context of working memory theory the savings in computational cost might be more apparent than real.

5.2.4. Neural net limits on relational complexity. Although there are still some unresolved issues in neural net representation of relations, it is strongly indicated that the limit is in the number of dimensions, or number of entities related, rather than in the total amount of information. Shastri and Ajanagadde (1993a) showed that synchronous oscillation models can represent quite large amounts of information in parallel, but only a small number of distinct entities can be related (the authors present many complex cases of reasoning without more than three distinct entities being represented in parallel). Tensor product symbol–argument–argument models also imply a limit in the number of distinct entities, rather than in the amount of information. In these models computational cost is polynomial in vector size but exponential in the number of dimensions, so the amount of information that can be represented by a single vector is not significantly limited, but the number of vectors that can be bound in one representation of a relation is limited. Both types of neural net models are consistent with psychological data in implying that the limit is in the number of distinct entities that can be related in parallel.

It is natural to ask why these two models share this limitation. The probable reason is that the synchronous oscillation models and tensor product symbol–argument–argument models have been designed to model higher cognitive processes comprehensively and consequently incorporate the properties of relational knowledge defined in section 2.2. The computational cost is attributable to dimensions of relations being mapped to dimensions of representations, either vector spaces or phases of oscillation, as noted in sections 4.1.3 and 4.2.13. Thus the computational costs that we have observed are not inherent in specific architectures but are inherent in processing relational knowledge. The more adequately a model incorporates the features of relational knowledge, the more clearly it entails these costs.

Tensor product role-filler binding models (Smolensky 1990) incorporate some but not all the properties of relational knowledge, and their computational cost depends on both number of arguments and number of instances stored. Importantly, in general these models do not map dimensions of relations to dimensions of representations, so their computational cost is less clearly related to dimensionality of relations. These models are efficient with few relational instances, but their cost relative to symbol–argument–argument binding models increases when many instances are stored. Circular convolution models (Plate 1995) appear at first sight to avoid computational costs in space because vector size is constant in the number of entities related. It is not clear, however, that circular convolution models incorporate all properties of relational knowledge, and the cost of the additional features required is unknown. Furthermore, they produce ambiguous output and depend on clean-up memories that store every relational instance. This incurs a major computational cost that depends on the number of arguments and on the number of instances stored but in the worst case can exceed the computational cost of symbol–argument–argument bindings.

Synchronous oscillation models suggest that about five entities can be processed in parallel, and this would permit one quaternary relation to be represented (symbol and four arguments). The empirical literature reviewed in section 3.3 indicates that quaternary relations are processed in parallel, so there does not appear to be a major disagreement here. Shastri and Ajanagadde (1993a) suggest that up to 10 entities can be related, with cross talk. Our model agrees to the extent that it implies a soft limit on processing capacity, with performance degrading gracefully as processing load increases.

The ability to process relations more complex than quaternary, though with increased error, may be important in creativity, where early ideas are often imprecise and difficult to communicate. Creative thought also probably requires processing of complex relations, because it entails integrating known relations and producing new relations that we do not yet know how to chunk or segment. It is possible that processing of relations of high dimensionality is important in creativity, but the increased risk of error would make confirmation and explicitness essential.

6. Empirical evidence

Processing load should be a function of relational complexity, which should limit the complexity of tasks that can be performed when chunking or segmentation are inhibited, either by task structure (e.g., nondecomposable relations) or by experimental manipulation. Performance predictions depend on analyses of the relations processed, and these analyses in turn have to be confirmed empirically. Therefore, testing the theory entails three steps.

(1) Develop a process model of the task, and empirically verify the model. This can entail an extensive program of developing and testing models.

(2) Analyse relations that must be processed. When possible, apply the reduction technique outlined in section 3.4.3 to determine effective relational complexity. If the task entails steps performed serially, the relevant step is the one in which the largest number of relations is processed (i.e., the peak load). Chunking and segmentation must be controlled (see sect. 3.4.1). When analysing tasks, it is useful to think of the number of interacting variables that are processed in parallel in a given step. Features that remain constant do not contribute to complexity, because they can be readily chunked.

(3) Test predictions derived from step 2 by manipulating relational complexity, with other factors controlled. It is necessary to manipulate the information that must be processed in parallel in order to make a decision. This sometimes entails preventing serial processing (an example is given in sect. 6.1.4).

6.1. Complexity and processing load

Our purpose in this section is to show how relational complexity analysis can be applied to tasks that are already well understood and for which reasonably well-validated process models exist. We adopt a “breadth first” approach, with the aim of showing that relational complexity is applicable to a wide range of phenomena in higher cognition and, therefore, offers worthwhile generality.

6.1.1. Transitive inference. Transitive inference has been shown to be a ternary relation by the reduction technique described in section 3.4.3. This is consistent with a number of well-substantiated models (Sternberg 1980; Trabasso 1975) showing that transitive inferences are made by integrating the premise elements into an ordered triple. For example, the premises “Tom is smarter than John, John is smarter than Stan” can be integrated into the ternary relational instance monotonically-smarter(Tom, John, Stan). Maybery et al. (1986) showed that premise integration, which entails a ternary relational instance, should impose a higher processing load than premise coding, which entails binary relational instances, such as smarter-than(Tom, Stan). The middle term can be ignored once integration has occurred, so generating a conclusion entails only a binary relational instance smarter(Tom, Stan). Previous models predicted that processing of negatives (e.g., “John is not as tall as Tom”) would increase processing load, but these models had not predicted the processing load of premise integration because it occurs in every form of the task, and the models were oriented toward accounting for task differences.

Maybery et al. (1986) tested the prediction that premise integration would impose a high processing load by using a sentence verification format for transitive inference with adult participants. Segmented presentation was used, so the second premise did not appear until the participant indicated that the first premise had been encoded. When the participant indicated that the second premise had been encoded and the premises had been integrated, a target appeared to which the participant responded by pressing one of two buttons indicating whether the target was consistent with the premises. Probe reaction time, saying “beep” to a tone, was used to assess information processing loads. The control task required verification of separate premises without integration, for example, “Tom is smarter than John, Peter is smarter than Stan.”

Experimental and control tasks were very closely matched in other respects.

Relational complexity theory predicts that there should be significantly longer probe reaction time when the probe occurs while the second premise is being processed than during processing of the first premise or the target but that there should be no such effect with the matched nonintegration control task. Maybery et al. (1986) found a significant probe position by integration/nonintegration interaction of this form and showed that alternative explanations based on response interference and similar processes could not account for the effects. Processing negatives, previously thought to impose high demands, imposed less load than premise integration.

6.1.2. Verifying relations. The prediction that binary relations impose higher processing loads than unary relations is supported by letter-match data. Posner and Boies (1971) showed that in the letter-match task processing load, as indicated by a probe reaction time secondary task, was greater when the second letter was presented. Coding of one letter is equivalent to a unary relation; for example, letter(c), representing that the stimulus is the letter “c.” When the second letter is presented it also must be encoded, but then a binary relation such as same(c, c), or different(c, k), must be represented. The theory accounts for the higher load observed in the comparison task, because coding requires a unary relation, whereas comparison requires a binary relation. Dimensionality does not preclude other factors, such as memory retrieval, contributing to difficulty, as suggested by the finding that name match is harder than physical match (Posner & Boies 1971).

6.1.3. The Tower of Hanoi. The Tower of Hanoi (TOH) puzzle is another task for which there are well-validated process models, and it is a good example of a task that entails planning, which depends on relational knowledge. TOH comprises three pegs and a variable number of discs. The discs are placed initially on peg A with the largest on the bottom, the next largest above it, and so on. The goal is to move all discs from peg A to peg C, without moving more than one disc at a time, or placing a larger on a smaller disc.

Complexity in TOH depends on the levels of embedding of the goal hierarchy, a metric that has been commonly used to assess complexity (Just et al. 1996). The more difficult moves require more levels of embedding of subgoals in the goal hierarchy. However, the goal hierarchy metric can be subsumed under the relational complexity metric because, as is shown in Table 1, moves with more subgoals entail relations with more dimensions of complexity. The first and every fourth move thereafter are shown, because it is only these that require planning (VanLehn 1991).

Consider a two-disc puzzle. To shift disc 2 from A to C, it is necessary first to shift disc 1 from A to B. The main goal is to shift 2 to C (2C), and the subgoal is to shift 1 to B (1B). The goal hierarchy therefore is 2C 1B, and has two levels (see Table 1). However, the task can be expressed as a relation:

$$\text{Prior}[\text{shift}(2, C), \text{shift}(1, B)]$$

Shift is a relation, so shifting 2 to C can be expressed as shift(2, C). Similarly for shifting 1 to B. The essence of the goal hierarchy is to perform a set of moves in order to perform another move. This can be expressed as the higher order relation “Prior”, the arguments of which are shift; that is $\text{Prior}[\text{shift}_2(\text{—}, \text{—}), \text{shift}_1(\text{—}, \text{—})]$. As with other relations, complexity is a function of the number of dimensions or roles to be filled (four in this example), so the task is prima facie four dimensional.

Now consider the more complex three-disc puzzle. To shift 3 to C, it is first necessary to shift 2 to B, in order to do which it is necessary to shift 1 to C (Table 1). There are now three levels of goals, and the corresponding relations are also more complex.

$$\text{Prior}[\text{shift}(3, C), \text{Prior}[\text{shift}(2, B), \text{shift}(1, C)]]$$

There are now six roles, so the task is six dimensional. By similar argument, the first move on the four-disc puzzle entails four levels of goals, and can be expressed by the relation:

$$\text{Prior}(\text{shift}(4, C), \text{Prior}(\text{shift}(3, B), \{\text{Prior}[\text{shift}(2, C), \text{shift}(1, B)]\}))$$

This is eight dimensional. Thus number of embedded subgoals corresponds to relational complexity, as Table 1 shows. Conceptual chunking and segmentation can be used to reduce complex-

Table 1. *Tower of Hanoi moves which require planning, showing goal(s) and dimensions*

Problem/move	Current state	Move	New state	Goal(s) ^a	Dimensions	
					All	New
2 disc						
1	12,_,_	1 to B	2,1,_	<u>2C</u> 1B	4	4
3 disc						
1	123,_,_	1 to C	23,_,1	<u>3C</u> <u>2B</u> 1C	6	6
5	_,12, 3	1 to A	1,2,3	<u>2C</u> 1A		
4 disc						
1	1234,_,_	1 to B	234,1,_	<u>4C</u> <u>3B</u> <u>2C</u> 1B	8	8
5	4, 3,12	1 to A	14,3,2	<u>4C</u> <u>2B</u> 1A	6	4
9	_,123,4	1 to C	_,23,14	<u>3C</u> <u>2A</u> 1C	6	6
13	12,_, 34	1 to B	2,1, 34	<u>2C</u> 1B	4	4
5 disc						
1	12345,_,_	1 to C	2345,_,1	<u>5C</u> <u>4B</u> <u>3C</u> <u>2B</u> 1C	8	8
5	45,12,3	1 to A	145,2,3	<u>5C</u> <u>4B</u> <u>2C</u> 1A	8	4
9	5,4,123	1 to B	5,14,23	<u>5C</u> <u>3B</u> <u>2A</u> 1B	8	6
13	125,34, _	1 to C	25,34,1	<u>5C</u> <u>2B</u> 1C	6	4
17	_,1234,5	1 to A	1,234,5	<u>4C</u> <u>3A</u> <u>2C</u> 1A	8	6
21	3,4,125	1 to B	3,14,25	<u>4C</u> <u>2A</u> 1B	6	4
25	123,_,45	1 to C	23,_,145	<u>3C</u> <u>2B</u> 1C	6	4
29	_,12,345	1 to A	1,2,345	<u>2C</u> 1A	4	2

^aNew goals are underscored.

ity, as with other relational tasks. The first move of the three-disc puzzle can be simplified by chunking discs 1 and 2.

Prior[shift(3, C), shift (1/2, B/C)]

Shift(1/2, B/C) can be unchunked thus.

Prior[shift(2, B), shift(1, C)]

Hence conceptual chunking and segmentation allow the task to be divided into two four-dimensional subtasks. This captures the recursive subgoal strategy that underlies successful performance (VanLehn 1991), and a conceptual chunk of this kind is called a “pyramid.”

Just et al. (1996) have shown that processing resources are related to the number of new goals that have to be generated for a move. Planning requires only representation of new goals, so the relations that correspond to new goals provide a more realistic estimate of the dimensionality of a move. These are shown in Table 1 by underscoring the new goals at each step. When number of goals is reduced, relational complexity is reduced correspondingly.

Our estimate that humans are limited to processing approximately four dimensions in parallel implies that humans would normally process no more than one goal and one subgoal in a single move. That is, they would process one relation of the form Prior[shift(2, C), shift(1, A)] or Prior[shift(3, C), shift (1/2, B/C)]. This is consistent with protocol information (VanLehn 1991, Appendix, pp. 42–47). A number of predictions based on this analysis have been tested, with positive results (Loveday 1995).

The relational complexity metric subsumes the metric based on levels of embedding of a goal hierarchy, because number of levels of embedding can be mapped directly into relational complexity. However, relational complexity also applies to tasks that do not entail subgoals, including tasks for which decisions can be made in a single step; it therefore has greater generality. Just as important, relational analysis gives insights into the kind of decisions that are made to construct the goal hierarchy. For example, it enables us to determine how much information is likely to be processed in one step when a decision is made that, in order to move 3 to C, 1 and 2 must be moved to B. Notice that the TOH can be performed without processing steps more complex than a quaternary relation. Relational complexity also has the advantage that there is extensive developmental data, to be reviewed in section 6.2, indicating ages at which children can typically process each level of complexity. This allows predictions to be made about typical successes of children on specific decisions within the TOH task.

6.1.4. Sentence comprehension. Sentences with reduced relative clauses (i.e., without syntactic markers), with a centre-embedded structure, and without semantic cues, make it difficult for most English speakers to identify cases without parsing the whole sentence. We assume that participants normally segment sentences into constituents, which are processed serially as far as possible. In all our modeling, in this and other contexts, we have found it a fruitful assumption that participants tend to minimize processing demand, implying that they never process more information in parallel than necessary. This type of structure tends to preclude serial processing, however, thereby preventing the processing load from being reduced by segmentation. This logic has been used by Just and Carpenter (1992) and by Henderson (1994) to test processing load predictions from the theory of Shastri and Ajjana-gadde (1993a). An example of such a sentence was mentioned in section 2.

The boy the girl the man saw met slept. (1)

The subjects and objects of the verbs cannot be identified individually (Kimball 1973), and such sentences are associated with high processing loads (Just & Carpenter 1992). Because such sentences tend to inhibit serial processing, they can be used to explore our capacity to process relations in parallel. The meaning of the sentence can be expressed in the following propositions.

slept(boy)
met(girl, boy)
saw(man, girl)

There are five roles to be filled, corresponding to subject and object of the verbs.

slept(Subject1)
met(Subject2, Object2)
saw(Subject3, Object3)

The sentence can be parsed by applying a set of rules (not necessarily the only possible set) shown in Appendix C, which collectively constrain a unique parsing, shown as P10. Given that serial processing is effectively inhibited, parsing the sentence amounts to finding an assignment of noun phrases to roles that fits a set of constraints that correspond to the rules. There are five roles that must be filled, so the task is five-dimensional, and beyond the capacity of most adults, with even exceptional individuals finding it at the limit of their powers.

Andrews and Halford (1994) tested these predictions using center-embedded and right-branching sentences, with reversible content (e.g., “The cow followed the horse”) or nonreversible content (e.g., “The boy patted the puppy”). Nonreversible sentences reduce the need for parallel processing (e.g., boy can be assigned directly to the subject role and puppy to the object role), whereas in reversible sentences there are no semantic constraints to assist with identification of subject and object. This again illustrates the point made in section 3.6 that effects of content can operate through the complexity of relations that have to be processed in parallel.

Examples of centre-embedded sentences of each dimensionality, together with the corresponding propositions, follow (centre-embedded and right-branching structures are not distinguishable with one-dimensional sentences):

One-dimensional:

“The dog ran.” Ran (dog)

Two-dimensional (ignoring initial clause in parentheses, which is used to make the center-embedded structure meaningful, was the same for all two-dimensional sentences, and can be processed before the remainder of the sentence):

“(The boy saw) the dog that the cat chased.” Chase (cat, dog)

Three-dimensional:

“The emu that the kangaroo passed slept.”
Sleep(emu)
Passed (kangaroo, emu)

Four-dimensional:

“The baker that the fireman introduced the doctor to died.”
Die(baker)
Introduce-to(fireman, doctor, baker)

Five-dimensional:

“The clown that the teacher that the actor liked watched laughed.”
Like(actor, teacher)
Watch (teacher, clown)
Laugh(clown)

Participants rated sentences for ease of comprehension, with content controlled. Rated difficulty increased monotonically with dimensionality, but this was modified by an interaction with surface form, so dimensionality had a significantly stronger effect with centre-embedded structure. Dimensionality was also modified by reversibility, but not as strongly as by center-embedded/right-branching structure. The difficulty ratings strongly reflect the number of bindings that had to be processed in parallel. Participants also indicated whether they found a sentence to be incom-

prehensible. Only four- and five-dimensional sentences were judged to be incomprehensible, and 88% of such judgments applied to five-dimensional sentences. Among these, 96% were applied to center-embedded sentences. Reversibility did not affect comprehensibility.

These results suggest that, when serial processing is inhibited by a center-embedded structure, people have difficulty assigning words to more than four case roles. The fact that the same limitation does not occur with right-branching sentences supports the dimensionality interpretation in preference to an alternative explanation in terms of the repeated variable constraint, because this would apply equally to right-branching sentences.

6.2. Relational complexity and cognitive development

The theory also predicts that children's performance will be poorer when relational complexity is greater. Furthermore, if children's processing capacity, or the efficiency with which they use their available capacity, develops (Case 1985; Halford 1993), they should be able to represent concepts of higher relational complexity with increasing age.

6.2.1. Infancy. Content-specific representations appear to be possible in infancy. Baillargeon (1987a; 1987b) has shown that 4–5-month-olds can represent attributes and position of vanished objects, at least within the immediate spatiotemporal frame. They dishabituate when a rotating drawbridge moves through the space that should have been occupied by a hidden object, suggesting that they can represent its position in the apparatus in front of them. They are also sensitive to attributes such as height and compressibility of the vanished object. Such performances are consistent with representation of the object in the immediate spatiotemporal frame. There is no evidence that semantically interpretable relations are represented, however, or that inferences go beyond the perceptible properties of objects. For example, there is no evidence that infants infer that an object must have been removed if a drawbridge moves through the space that it should have occupied.

6.2.2. Unary relations. Unary relations appear to be processed at one year of age, as indicated by category membership (Sugarman 1982; Younger 1993) and by disappearance of the A not-B error. This paradoxical phenomenon in infant object constancy research (Wellman et al. 1986) can be interpreted as inability to treat hiding place as a variable, reflecting lack of variable-constant binding (a unary relation; see sect. 2.3). That is, when an infant has repeatedly retrieved an object from hiding place A, then continues to search for it at A despite having just seen it hidden at B, the infant is treating hiding place as a constant. To treat hiding place as a variable requires representation of the binding between the variable location and the object, that is, location(object₁).

Performance on this task deteriorates as a function of delay between hiding and retrieval, and this effect is greater for younger children (Wellman et al. 1986). This would follow from a representation with the form of a rank two outer product of vectors representing a variable and a constant, if we make the reasonable assumption that the representation degrades with delay but becomes clearer with age, so older children can tolerate more degradation before the representation becomes uninterpretable.

6.2.3. Binary relations. Binary relations such as “larger” and “more” appear to be well understood by 2 years of age, even though there may be some confusion as to which relation is referred to by a particular term (Halford 1982; 1993). Proportional analogies of the form $a:b::c:d$ are frequently based on binary relations, and there is evidence that young children can make such analogies in familiar domains (Goswami 1992). Proportional analogies do not require processing quaternary relations; they require processing two binary relational instances belonging to the

same relation. Because the relation is constant to both sides of the analogy, only a single binary relation need be considered. For example, in the analogy “mother is to baby as horse is to what?,” (mother, baby) identifies the relation mother-of, and mother-of and mother-of(horse, ?) identifies foal.

This point can be illustrated by comparing a proportion with an analogy. A proportion $a/b = c/d$ is a quaternary relation in that the relation between each of a, b, c, d , and the other three entities is defined, whereas this is not true for an analogy. Thus $8:4::27:2$ is a valid analogy in that the same relation “>” holds between $8:4$ and $27:2$, but it is obviously not a proportion. To illustrate this in another way, we can define proportion as the quaternary relation proportional(a, b, c, d). Now, knowing four of the elements (i.e., any four out of “proportional/nonproportional,” a, b, c, d), we can determine the fifth, as discussed in section 2.3 (e.g., given $4,8,3,6$ we know that this is proportional, but $4,8,3,5$ is not proportional). With the analogy $8:4::27:2$, however, this is clearly not possible. Thus a so-called proportional analogy bears only a superficial resemblance to a proportion and they differ markedly in relational complexity.

6.2.4. Ternary relations. A number of concepts based on ternary relations have been associated with persistent difficulties for young children. Transitivity and class inclusion are the best known examples. Attempts at explanation based on stages of development or on flawed methodology leading to false negatives have provided many insights and have yielded improved assessments, but they still leave important sources of difficulty unexplained (Halford 1989; 1992; 1993).

6.2.4.1. Transitivity. Transitivity has been a source of difficulty for young children, the reasons for which not been wholly explained (Breslow 1981; Bryant & Trabasso 1971; Halford 1982; 1989; 1992; 1993; Thayer & Collyer 1978; Trabasso 1977). We suggest that the unrecognized factor is the processing load imposed by premise integration, which also affects adults (see sect. 6.1.1) but has a greater effect at younger ages (Halford et al. 1986).

Piaget's (1950) contention that transitivity is a concrete operational task was challenged by Bryant and Trabasso (1971), who trained children in the premises and found above-chance performance in 3–4 year-olds. Subsequent work, however, suggested that the children might have been given undue assistance in ordering the premise elements (i.e., given premises $a < b, b < c$, etc., both children and adults typically integrate the premises into the ordered set $\{a, b, c \dots\}$). The elimination of children who failed to learn the premises might have biased the results, because premises would be difficult to learn if they could not be integrated. When these factors were controlled, children under 5 years no longer succeeded (Halford & Kelly 1984; Kallio 1982). Evidence for transitive inference, in children, adults, or other animals, provides evidence of processing ternary relations only if participants are not assisted in ordering the premise elements.

More recent work (Pears & Bryant 1990) has shown that if 4-year-olds are given premises in the form of pairs of colored blocks stacked one above the other (e.g., red above green, green above blue, etc.), they can infer the order of blocks in a tower (e.g., red above blue). Andrews and Halford (in press), however, showed that 4-year-olds' performance is marginal at best and is influenced by relational complexity. This work is consistent with the present theory and with the empirical work of Halford (1984) in showing that when children under 5 years are required to consider two binary relations in a single decision, they have very little success, but they virtually always succeed when they can process one relation at a time. Children older than 5 years succeed on both tasks. We again find evidence that relational complexity affects performance when other factors are controlled, and the effect is greater with younger children.

6.2.4.2. Class inclusion. Class inclusion entails understanding that a and a' are included in b (e.g., apples and nonapples are included

in fruit), and therefore $b > a$ (assuming a and a' to be nonempty). As with transitivity, class inclusion was regarded by Piaget (1950) as concrete operational, and unattainable before 7–8 years of age. There have been alternative explanations, including misapplication of the rule that a set is counted only once (Klahr & Wallace 1976; Trabasso et al. 1978; Wilkinson 1976), and misinterpretation of the question as requiring subclass comparison (Grieve & Garton 1981; Markman & Seibert 1976; McGarrigle et al. 1978; Shipley 1979). These issues have been discussed elsewhere and, although some sources of false negatives have been discovered, class inclusion presents a source of difficulty for children under age 5 years that has not been fully explained (Halford 1992; 1993; Halford & Leitch 1989; Hodkin 1987).

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 nonapples). More extended hierarchies are obviously possible, but the concept of inclusion necessarily entails a class and its complement being included in a superordinate class. More extended hierarchies can be handled by conceptual chunking or segmentation. For example, the inclusion of apples, bananas, pears, and so on, in “fruit” can be handled by chunking bananas, pears, and so on, in “nonapples.” When there are more than two levels, they can be handled by segmenting the hierarchy into subhierarchies and processing two levels at a time. The models of both Hummel and Holyoak (1997) and Halford et al. (1996; 1997) entail this process. Chunking apples and nonapples, however, would make the concept of inclusion inaccessible. Part–whole hierarchies, which cause difficulties for young children in arithmetic word problems (Cummins et al. 1988; Halford 1993; Kintsch & Greeno 1985) are similar and comprise a whole divided into two complementary parts.

The difficulty children have with these problems supports the hypothesis that children under 4–5 years have difficulty with ternary relations. Furthermore, as with transitivity, relational complexity has been shown to interact with age; children under 5 years succeed when the task requires them to consider only one binary relation, but not when they have to integrate binary relations, whereas older children succeed in both cases (Halford & Leitch 1989). The same finding has been made with matrix classification (Halford 1980).

6.2.4.3. Concept of mind. Two reviews of children’s concept of mind (Astington 1993; Halford 1993) have noted phenomena that, although at first appearing anomalous, can be interpreted in terms of relational complexity. Very young children seem to have difficulty understanding that a person can have two representations of an object. For example, the perceived colour of an object can be modified by a coloured filter (the appearance–reality and perspective-taking tasks), or a person’s knowledge of an object’s whereabouts can depend on whether he or she knows that it has been moved since it was last seen (the false-belief task). Based on an extensive assessment of the literature, Flavell et al. (1990) have proposed that young children cannot handle two ways of representing an object (e.g., as both blue and white), but the authors have not explained why young children should be limited in this way. Some recent analyses (Frye et al. 1995; Halford 1993; 1996a) converge on the same explanation.

The essence of the problem is that children can represent the relation between a person’s knowledge (which we call percept) and the properties of an object or situation. This is a binary relation between percept and object attribute. However, they cannot represent the case when this is conditional on a third variable. Consider, for example, a person who sees an object as white (without filter) and as blue (with filter). This entails representing a ternary relation between viewing condition, object, and percept.

Seen-object2(<condition>, <object-color>, <percept>)

Instances of this relation would be as follows:

Seen-object2(no-filter, object-white, percept-white)

Seen-object2(blue-filter, object-white, percept-blue)

This is a ternary relation between the condition¹⁰, the object and the person’s representation of the object. Young children seem unable to do this and seem to represent the relation between an object and one percept only. That is, they represent either Seen-object1(object-white, percept-white) or Seen-object1(object-blue, percept-blue).

These are both instances of a binary relation, and either could be represented alone by a person who could not represent ternary relations. To represent the fact that a person can see an object in either of two ways, however, entails conditionalising these relational instances on a third variable, which is equivalent to integrating the binary relations into a ternary relation. Notice that the ternary relation is not decomposable in the sense that it is not reducible without remainder to instances of the binary relation Seen-object1, because this relation does not represent the fact that alternative ways of viewing the same object are conditional on the viewing condition.

The same limitation would occur with false belief, which entails representing the relation between an object and two different representations of its location, one based on knowledge of where it is and the other based on a false belief about its location. For example, a person sees an object placed in a box, then leaves, and the object is shifted to a basket. Young children have difficulty understanding that the person will believe the object to be in the box though it is really in the basket (Wimmer & Perner 1983). This can be expressed as the ternary relation

Find-object(<known-event>, <actual-location>, <believed-location>)

instances of which are as follows:

Find-object(<saw-moved>, <obj-in-basket>, <believe-obj-in-basket>)

Find-object(<not-seen-moved>, <obj-in-basket>, <believe-obj-in-box>).

The somewhat paradoxical difficulty that young children have with appearance–reality, perspective-taking, and false belief can be interpreted in these terms: they readily understand any of the component binary relations, yet they cannot “put the situation together” and integrate two object–percept relations into a single representation (note that this is analogous to transitivity, which entails integrating two binary-relational premises into a ternary relation; see sect. 6.1.1). Young children’s apparently anomalous performance in the concept-of-mind tasks is perfectly consistent with their performance on other tasks that entail ternary relations. There is evidence that processing capacity is a factor in children’s concept of mind (Davis & Pratt 1995; Frye et al. 1995), but the predictions offer scope for more empirical work.

6.2.4.4. Contrary evidence. The most recent, and probably strongest, challenge to the proposition that children under 5 years have difficulty with ternary relations comes from Goswami (1995). In experiment 1 she presented 3- and 4-year-old children with two sets of three stacking cups. The experimenter indicated a cup (smallest, middle, or largest) in one set, and the child had to identify the corresponding cup in the other set (smallest, middle, or largest, respectively). Performance was high and appears to provide impressive evidence that 3–4-year-olds can process ternary relations, but there was little attempt to analyze the processes by which children made their discriminations.

We can perform at least a first analysis using the reported sizes of the cups, as shown in Table 2. As with Goswami’s report, the 12 cup sizes are shown as the values 1–12, the four sets used being: 1, 5, 9; 2, 6, 10, and so on. Table 2 indicates three types of cases. The first is where the corresponding cup can be selected on the basis of *absolute size* (indicated by a “1” in Table 2). For example, if the experimenter’s set is 1, 5, 9, and the child’s set is 2, 6, 10, and

Table 2. Mapping ordered triples, based on absolute size^a

			Experimenter's set											
			1	5	9	2	6	10	3	7	11	4	8	12
Children's set														
1	5	9	1	1	1	1	1	1	.5	.5	1	0	0	1
2	6	10	1	1	1	1	1	1	1	1	1	.5	.5	1
3	7	11	1	.5	.5	1	1	1	1	1	1	1	1	1
4	8	12	1	0	0	1	.5	.5	1	1	1	1	1	1

^a1 = absolute size gives correct answer; .5 = absolute size gives two answers, one of which is correct; 0 = absolute size gives incorrect answer.

if the experimenter indicates cup 5, the correct choice for the child is cup 6, because it is in the corresponding ordinal position (middle). However, cup 6 is also closest to cup 5 in absolute size. Hence a child who paid no attention to binary or ternary relations, and judged on the basis of absolute size (a unary relation), would be likely to choose cup 6, which is the correct answer. This is true in every case where a "1" is entered in Table 2. The critical case is Goswami's different-size cup group in the different spatial position condition, because it is only here that cups were neither identical nor in the same spatial position in the two sets. This case corresponds to the off-diagonal entries in Table 2. There are 24 of 36 cases, or 66.67%, in which the correct cup can be determined on the basis of absolute size. There are a further eight cases, or 22.22%, in which absolute size gives two equally likely answers, one of which is correct (indicated by ".5" in Table 2). These would be expected to yield a further 11.11% correct answers. Thus the expected performance, if children attend only to absolute size without processing any relations of higher rank than unary, is 78% correct. Furthermore, some additional correct answers could be obtained by processing binary relations (to be discussed below). Goswami reports 86% success for 4-year-olds and 70% success for 3-year-olds (averaging over the analogy/no-analogy treatment, which is not relevant to our present argument). Thus these superficially impressive performances provide no evidence for the processing of ternary relations by 3-year-olds and doubtful evidence for 4-year-olds.

In experiment 2 Goswami had 3- and 4-year-old children map fractions from one set to another. For example, given three glasses of lemonade, one-quarter full, one-half full, and full, the child would be asked to see the correspondence between one-half glass of lemonade and one-half box of chocolates. In the different spatial position condition, the performance of 4-year-olds was 87% correct and that of 3-year-olds 54%. Absolute size matching is not possible in this experiment, but it is necessary to ask what success can be achieved using binary relations. Two elements in an ordered set of three can be chunked, reducing the task to the binary relation larger/smaller; for example, one-quarter might be labeled "smaller," whereas one-half and full are chunked as "larger" (this is an example of the kind of chunking presented in Fig. 1E). If the correct item is one-quarter, ($p = .33$) this leads to 100% correct. If the correct item is either one-half or full ($p = .67$), it leads to 50%. The overall expected percentage correct is therefore $.33 + .67 \times .50 = .67$. The three-year olds are clearly not above the binary-relation baseline, and no test was made to see whether 4-year-olds were significantly above this baseline.

Experiment 3 entailed mapping between levels of loudness, pitch, hardness, height, and so on, which did not permit use of absolute cues. There is the possibility of chunking to binary relations as discussed above, but we will not pursue that issue here. The 4-year-olds were again successful but, in both experiments 2 and 3, *the mean age of the 4-year-olds was 4 years 11 months, range 4.8–5.1*. No data for 3-year-olds are reported on this task, so the

only evidence for mapping ternary relations is obtained from children who are at, or very close to, 5 years old. The data actually support the proposition that ternary relations are first processed at a median age of 5 years.

6.2.4.5. Summary of ternary relations evidence. The fundamental problem here is that cognitive development research must take account of actual cognitive processes to be theoretically meaningful (Halford 1982; 1989; 1993; Siegler 1981). The more important evidence however comes from those studies in which relational complexity has been varied while other factors are held constant. It appears that tasks entailing ternary relations are consistently found to cause difficulty for young children, even though the children readily process binary relations. This suggests that relational complexity is an important factor in children's cognitive performance, and it offers a solution to the mystery of why these tasks have seemed unaccountably difficult.

6.2.5. Quaternary relations. Proportion (see sect. 2.3) and the balance scale entail quaternary relations, because they entail relations between four variables, and both have been found to be difficult for young children, but there is more extensive research on the balance scale.

6.2.5.1. The balance scale. The balance scale entails the quaternary relation balance-state($W_1 D_1, W_r D_r$). The task is difficult for children under about 11 years old, and they tend to use lower rank rules, but even adults rarely use the cross-products rule without specific instruction (Siegler 1981; Surber & Gzesh 1984). The task can be segmented by, for example, computing the moments on each side and then comparing them to see which side will go down or whether the beam will balance: $W_1 \times D_1 = M_l$; $W_r \times D_r = M_r$. Each of these steps requires processing a ternary relation. The comparisons entail the rules $M_l = M_r \rightarrow$ balance, $M_l > M_r \rightarrow$ left side down, and $M_l < M_r \rightarrow$ right side down. Each is a comparison of two values and is a binary relation. The task can hence be performed by processing ternary relations one at a time. However, planning this strategy means being able to represent the fact that the moments are determined by weight and distance on each side (see sect. 2.2.12). This entails representing the ternary relation relation balance-state($W_1 D_1, W_r D_r$).

6.2.5.2. Neural net model of balance scale. McClelland (1995) has shown that a three-layered net (with input, hidden, and output layers of units) can be trained to indicate whether a beam will balance, given weight and distance on left and right as input. The model accounts for a number of important empirical observations that challenge earlier theories of children's balance scale performance. However this model does not fully represent the principle of the balance scale, and the particular way in which it differs from the models discussed in section 4 is instructive. The fundamental difference is that McClelland's model does not incorporate the

omnidirectional access property (see sects. 2.2.6 and 4.2.6). Weight and distance on left and right must always be inputs, and only one output, balance/left-down/right-down can be calculated. An implementation of a quaternary relation that met the specifications given in section 2.2.6 can take *any* subset of $N - 1$ variables as input and generate the N th variable as output. For example, given weight and distance on the left, distance on the right, and the fact that the beam is balanced, the implementation can decide what weight must be on the right. This is realistic because such tests are used in assessment (see, e.g., Surber & Gzesh 1984) and because we would be unwilling to attribute understanding to a child who could compute only one type of output (such as whether the beam would balance) but could say nothing about (for example) which weights and distances were required to produce a given state of balance or imbalance. Thus, although McClelland's three-layered net efficiently computes a specific function of four variables, it does not meet the criteria for relational knowledge.

6.3. Capacity development redefined

The question of whether processing capacity changes with age can be reformulated by proposing that relational complexity of representations would increase because representations become differentiated into more vectors, with appropriate reconnection, as was noted in section 4.2.4. This would not necessarily change the total amount of information that can be processed, but it would increase the complexity of the relations that could be represented.

6.3.1. Predictions in advance. It has not been common practice for information processing theorists to publish predictions of developmental performance prior to obtaining data. However, if the conceptual basis for capacity limitations advanced in this paper is more objective than previous proposals, it should be possible to do this. Halford (1993, Ch. 9) made a number of such predictions. One was that 2- and 3-year-olds should be able to make balance scale judgments based on weight or distance, but not on both. The reason is that comparison of weights (or distances) on the two sides of the balance entails a binary relation, and norms indicate that this is possible at age 2 years (see sect. 6.2.3). By contrast, Siegler (1981) found no evidence of weight or distance rules before approximately age 5 years, and Case (1985; 1992) predicts that children will not be able to judge which side will go down until age 3.5–4 years. Weight is likely to be easier initially, because children have more experience with the downward force of weights. With appropriate experience, however, children should be able to make either weight or distance comparisons. This prediction was confirmed by Halford and Dalton (1995).

A further prediction was that taking account of weight and distance requires integration of binary relations, which is equivalent to at least a ternary relation, depending on the basis of the integration (see Halford 1993, pp. 413–422 for details of the prediction). Hence the ability to consider both weight and distance in a single judgment should develop at a median age of 5 years, and should be predicted by performance on other measures of ternary relations. This was tested by Harper (1996) using three tasks that require ternary relations processing but are from domains other than balance (cardinality, class inclusion, transitivity). Both predictions were confirmed.

6.3.2. Capacity and cognitive development. Halford (1993) has suggested that the observations that gave rise to cognitive developmental stage theory (Piaget 1950) might be attributable to progressive differentiation of representations with age. In very broad terms, Piaget's secondary circular reactions, with minimal representation, correspond to single vector representations, pre-conceptual reasoning corresponds to unary relations (a binding of two vectors), intuitive reasoning to binary relations (a binding of three vectors), concrete operational reasoning to ternary relations (a binding of four vectors), and formal operational reasoning to quaternary relations (a binding of five vectors).

Tasks that have been considered to belong to a particular stage tend to have a common level of relational complexity. For example, transitivity and class-inclusion, which are considered to be concrete operational, entail ternary relations, as was noted in section 6.2.4. Each increase in complexity of relations that can be processed in parallel would allow processing of a new level of tasks. However, the fact that the theory can explain some phenomena that have been attributed to stages does not mean that all aspects of stage theory are automatically entailed. Because there is considerable potential for misinterpretation on this point, we will amplify implications for cognitive development theory.

It is a common assumption that theories of cognitive development that assign a role to capacity *ipso facto* have no role for experience. However, nothing in this theory implies that attaining a given level of processing capacity automatically furnishes the mind with all concepts at that level. Defining a role for capacity in no way diminishes the importance of learning, induction, categorization, and other acquisition processes. We have proposed that development depends on the interaction of processing capacity and acquisition processes, so that what is acquired depends on both experience and capacity (Halford 1971; 1980; 1995; Halford & Fullerton 1970; Halford & Wilson 1980). Acquisition of transitive inference, for example, depends on experience with relations, but children who can process ternary relations will develop strategies that are more powerful and comprehensive than those of children who are restricted to processing one binary relation at a time. This effect has been simulated in the model of Halford et al. (1995).

Another common misconception is that capacity theories emphasize what children cannot do and imply insurmountable barriers to performance. On the contrary, good complexity analyses can actually lead to previously unrecognized capabilities, as our work on the balance scale discussed in section 6.3.1 illustrates. Also, the discussion of chunking and segmentation in section 3.4 shows that capacity limitations do not constitute barriers to performance. Ultimately, capacity theories are about processes rather than barriers. Saying that a particular group of participants process relations of a given complexity in parallel implies that the task must be chunked or segmented to keep within this capacity. It therefore leads to predictions about the kinds of strategies that must be used. Thus saying that 3-year-olds process binary relations in parallel implies that their strategies will differ from those of 6-year-olds, who process ternary relations in parallel. It does not imply that they can never process transitive inferences, or other ternary relations tasks, any more than evidence that adults process quaternary relations in parallel implies that we can never understand force, which ultimately depends on more than four dimensions (as discussed in sect. 3.4.1). Capacity limitations imply only inability to perform when chunking and segmentation are inhibited, as with center-embedded sentences discussed in section 6.1.4. Capacity theory points the way to questions that need to be investigated; for example, what chunking and serial processing abilities do children of a given age and background have in a particular domain, and how does this influence their performance? We can investigate new questions and reexamine old questions with this orientation.

The influence of relational complexity on cognitive development does not imply that development is discontinuous, and it is important that relational complexity theory not be confused with traditional stage theory in this respect. Processing capacity is an enabling factor, but development is experience-driven and continuous, so the acquisition of concepts at a given level will occur gradually after capacity becomes sufficient. There is an unlimited number of concepts belonging to a given level of complexity, and acquisition of each will be a function of experience in the relevant domain, with what is learned being influenced by capacity. Furthermore, the acquisition of less complex concepts does not cease once capacity increases to a higher level of dimensionality because there is an unlimited number of concepts at all levels, so when a child becomes capable of processing, say, ternary relations,

he or she does not cease acquiring concepts based on unary or binary relations.

The ages at which each level of relational complexity is typically attained should be seen as medians, with the proportion of children who attain a given level increasing gradually, in accordance with a biological growth function. The specific ages are determined empirically and are essentially normative. Thus, if it could be shown that, say, 3-year-olds could represent ternary relations, this would revise the age norms, but it would not in itself invalidate the theory.

There should be correspondence between the attainment of different concepts of the same level of complexity, provided that domain knowledge is adequate. For example, transitivity, class inclusion, and other concepts requiring ternary relations, about which there is plenty of opportunity to learn, should have the same acquisition function, a prediction confirmed by Andrews (1996). Capacity to process relations of a given level of complexity should predict ability to acquire concepts at that level. Thus, the training asymptotes for concepts at a given level of complexity are the best data for testing the theory (this methodology has been used, for example, by Halford 1980 and Halford & Leitch 1989).

The theory would be invalidated if relational complexity did not predict processing demand. This can be tested in many ways, however, including some that are not developmental. Research discussed in section 6.1 illustrates how relational complexity can be manipulated precisely, while other factors are controlled, resulting in clearcut effects on processing demand as indicated by objective indicators such as concurrent probe reaction time. There is potential for using the same methodology with brain imaging techniques.

The cognitive developmental aspect of relational complexity grew out of levels of representational structure defined by Halford and Wilson (1980) and is consistent with neo-Piagetian theories (Case 1985; 1992; Chapman 1987; Pascual-Leone 1970). Consensus exists here that the growth of processing capacity is an enabling factor that has an explanatory role in cognitive development, with the important caveat that it is not the only factor, as noted above. The relational complexity metric might refine theories of processing capacity by providing a clearer mathematical definition, which should facilitate objective task analyses. It also opens up possibilities for computational modeling of growth in capacity, by differentiating neural nets into more dimensions, with consequent increases in the complexity of interactions.

Although no other cognitive developmental theory has used the metric proposed here, there is a broad parallel to the major stages defined by Case (1985; 1992), Fischer (1980), and Piaget (1950); the M-space levels defined by Pascual-Leone (1970); and the number of representational schemes defined by Chapman (1987), but there are also many differences. We will focus on the theory of Case (1985; 1992), who proposes that cognitive development progresses through four major stages, the sensorimotor, relational, dimensional, and vectorial. There are four substages – operational consolidation (preliminary), operational coordination, bifocal coordination, and elaborated coordination – which recur in each major stage. The executive processing load (equivalent to demand) is defined as $OP + S$. OP is specific to the major stage, so there are $OP_{\text{sensorimotor}}$, $OP_{\text{relational}}$, $OP_{\text{dimensional}}$, and $OP_{\text{vectorial}}$ but no quantitative value is specified for OP . One role of the present theory is to fill that gap.

The substages are quantified by the demands they make on short-term memory:

The number of goals children can maintain (and hence the complexity of problem they can solve) is determined by the size of their short-term memory for the particular class of operations in question. . . . This short-term storage space (STSS) can hold 1, 2, 3, and 4 items at the preliminary, first, second, and third substage of each period, respectively. (Case 1992, p. 32)

Notice that processing capacity, defined in terms of number of goals, ultimately depends on short-term *storage* capacity, rather than being defined in terms of information that is being processed,

as in the relational complexity metric. More important, this progression described by Case (1985; 1992) occurs recursively in each of the major stages, so S is maximal when the highest substage is reached, and minimal when transition is made to the next major stage. This means that a task imposing an executive processing load of 4 when performed at, say, the relational stage imposes a load of 1 when performed at the dimensional stage. The metric applies to progression within each major stage but does not transcend major stages, so it is not possible to compare, for example, the operational coordination substage of the relational stage and the bifocal coordination substage of the vectorial stage using a common metric (the values would be 2 and 3, respectively, but they are on different scales). Task demands are therefore assessed according to the major stages to which they are considered to belong.

However, the relational complexity metric proposed here is not stage dependent in this way. It applies throughout the age range and is also applicable to nonhuman primates (to be discussed in sect. 6.4). If participants of any age perform the same task, encoded in the same way and using the same strategy, task demand is the same. Processing demand may well change as expertise is acquired, because rules can be discovered that simplify decision making. For example, demand might be reduced in the TOH once participants realize that the difficult decisions need only be made on the first and every fourth step thereafter, but this can be taken into account in the analysis of processing demands. Demand can vary with coding (e.g., if dimensions are chunked) or as a function of strategy (e.g., if the task is segmented into more steps, requiring less information to be processed in parallel), but these factors are not stage dependent. A well-validated model of task performance allows processing demand to be analyzed objectively without assignment to substages.

Case (1985) also proposes that total processing space is constant across ages, and is flexibly allocated to operating space plus short-term storage space, but the expected trade-off between processing and storage does not occur (Halford et al. 1994). Perhaps the most important difference between the present theory and that of both Case (1985) and Pascual-Leone (1970) is that in the present theory processing capacity is not only a matter of space availability but is linked to the way vectors representing dimensions of a task are connected together. As Figures 1 and 2 illustrate, processing capacity ultimately depends on connections between related entities.

This formulation also differs from Piagetian and neo-Piagetian theories in that it does not postulate substages. This is not an oversight but is a natural consequence of the way the theory is formulated. As was noted above, development is continuous and depends on the growth of processing capacity as well as on acquisition processes. Substages may be used for descriptive convenience, but they are no more necessary to account for development than they are to account for, say, acquisition of expertise in adulthood. This should not be misinterpreted to mean that performance does not change over short periods of time. The way in which a task such as the TOH is performed might well change radically, even over a few trials, because of changes in the encoding of subproblems (e.g., chunking into pyramids, as noted in sect. 6.1.3). Acquisition mechanisms, including learning, induction, and categorization, can operate over a short or a long time frame and are modified by processing capacity only insofar as it operates as an enabling factor.

6.4. Relational complexity and other primates

Although there is little doubt that nonhuman mammals have representations (Gallistel 1990), only the primates appear to be able to process explicit relational representations that meet the criteria outlined in section 2.2. We will briefly review this evidence.

Premack (1983) reports that tasks requiring symbolic representations differentiate chimpanzees and monkeys from lower

mammals, whereas tasks based on perceptible similarity, or on inferences about spatial location, do not. We propose that relational representations subsume symbolic processes, because of the properties of relational knowledge given in section 2.2. One of Premack's procedures requires chimpanzees (*Pan troglodytes*) to identify the relation, given a pair of arguments, or to identify an argument, given the relation and the other argument (omnidirectional access as defined in sect. 2.2.6). In the former case, for example, chimpanzees would be shown two objects and asked to produce a symbol indicating whether the objects were the same or different. They also seemed to know that a knife corresponds to the relation between an intact and a cut object, a key to the relation between a closed and an open lock, and so on. These tasks seem to require recognition of a symbol for a binary relation; a knife is a symbol for the relation between an intact and a cut object (the knife functions as a relation symbol; see sect. 2.2.4 and sect. 4.2.3).

Another task used by Premack has been analyzed in terms of relational complexity by Holyoak and Thagard (1995). Chimpanzees were required to choose a pair of objects that had the same relation as a sample pair. We will call this the "relational-match-to-sample task." In one variant, the sample comprised two objects that were the same (XX). The participant had to choose another pair of objects that were the same (YY) in preference to a pair that were different (CD). The task is a form of analogical reasoning, as Premack (1983) points out, and requires a representation of the relation between elements in the pair. The sample is the base, and the comparison pair is the target. Following Holyoak and Thagard (1995), we can code the sample as O-same(X, X), where "O-same" means "same object." The correct comparison (target) object is represented as O-same(Y, Y). In the alternative task, the base would be coded as O-different(X, Y) and the correct target as O-different(C, D). Only chimpanzees, and only those that had been language trained, could perform this task. It seems reasonable to conclude that chimpanzees can process binary relations, albeit only after extensive experience with symbols. However, there appears to be no evidence that analogies based on binary relations can be made by lower animals.

The one-object match-to-sample task entails presenting a sample, X, and rewarding animals for choosing the comparison object that is the same as the sample, X, in preference to the one that is different, Y. Chimpanzees and some monkeys can generalize the performance beyond examples on which they have been trained. This requires responding to an attribute or a category label, which can be coded as a unary relation (as shown in sect. 2.3). For example, if the sample was an apple, this can be represented as apple(object₁). If the comparison objects were an apple and an orange, they would be represented as apple(object₂), orange(object₃). The one-object match-to-sample task is an analogy based on a unary relation: it is a mapping from apple(object₁) to apple(object₂). Some species of monkeys thus appear to process unary relations.

6.5. Role of frontal lobes

In a review of the literature, Robin and Holyoak (1995) propose that the prefrontal cortex functions as an overall system for constructing and maintaining relational representations that guide thought and action. They argue that many of the functions impaired in those with frontal lobe lesions, including planning and control, entail two- and three-dimensional representations, with dynamic binding. According to this hypothesis, the species and age differentiations discussed earlier would be attributed to the fact that the frontal lobes evolve later and are slower to myelinate. This formulation, together with associated empirical work, allows relational complexity of tasks to be manipulated while other factors are controlled. These techniques have reached high levels of refinement (see, e.g., Andrews 1996) and are now ready to be used with brain imaging techniques. It would be predicted that tasks requiring more complex relations to be processed in parallel

would produce more activation in the prefrontal cortex, with other factors controlled.

7. Conclusions

The empirical database in cognitive psychology and the current neural net models of relational knowledge indicate that processing capacity is limited not by amount of information or number of items per se but by the number of independent dimensions that can be related in parallel. Relational complexity, defined as the number of independent sources of variation that are related, constitutes a major factor underlying the difficulty of higher cognitive processes. It is related to processing load, to differences between higher animal species, and to age in children. The potential exists to explain processing loads by modeling neural net representations of relations.

This theory provides a way of blending serial and parallel processes, with serial processing being necessitated by limitations in the complexity of structures that can be processed in parallel. Empirical data and contemporary neural net models of relational knowledge indicate that the most complex structure that can be processed in parallel, and without cross talk, is equivalent to one quaternary relation. For more complex representations, either the representation must be chunked into fewer components (with the result that some of the relational structure becomes temporarily inaccessible) or the task must be segmented into smaller components that are processed serially, or both. Thus, the need for serial processing strategies can be seen as a consequence of processing capacity limitations.

The theory implies that the traditional approach of defining limitations in terms of items is inappropriate for processing capacity, although some common ground is found with Miller's (1956) suggestion that the limit is defined by the number of independent components, rather than the amount of information. The concept of a chunk is retained but extended to include conceptual chunks, which represent compressed relational instances. The definition of capacity in terms of relational complexity, and the exploration of possible neural net implementations, has provided a new way of looking at many issues, one that integrates a wide-ranging database. Regarding cognitive development, this means that the issue is no longer simply whether capacity changes with age but whether representations become differentiated into higher dimensionalities, so that more complex relations can be processed.

APPENDIX A. RELATIONS DERIVED BY DECOMPOSING HIGHER RANK RELATIONS

We will consider the example monotonically-larger(a, b, c) discussed in section 2.2.9. Let us abbreviate monotonically-larger to ML. It is interesting to look at the derived relations (projection relations) ML₁, ML₂, and ML₃. ML₁(b, c) means that b > c and there is some item x such that x > b. Similarly, ML₂(a, c) means that a > c and there is some item x such that a > x > c. Thus, each ML_i is a subrelation of the relation > (greater-than). The ternary relation induces a number of binary relations, and in this case ML(a, b, c) can be reconstructed from the induced relations in the sense that ML(a, b, c) ≡ ML₁(b, c) and ML₂(a, c) and ML₃(a, b). By contrast, the ternary relation R(x, y, z) ≡ x > yz, where x, y, and z are positive rational numbers, is not reconstructable; each of the induced relations is the "trivial relation"; for example, for each pair of numbers y and z, there exists an x such that x > yz (for example, yz + 1). Thus, R₁(y, z) is true for all y and z, and the same is

true for R_2 and R_3 . R is not effectively decomposable, basically because of the presence of the binary operation yz .

APPENDIX B. RETRIEVAL FROM TENSOR PRODUCT REPRESENTATION OF SYMBOL-ARGUMENT-ARGUMENT BINDINGS

The *outer product* set of three vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$, is a multidimensional object like a generalized matrix with three subscripts i, j, k . The (i, j, k) -th entry in the outer product is defined to be the product of the corresponding entries in each vector: $T_{ijk} = \mathbf{u}_i \mathbf{v}_j \mathbf{w}_k$. This definition generalizes in a natural way to as many vectors as required.

Regarding *generalized inner products*, for a tensor T of a given rank, a number of retrieval operations can be defined. For example, with a rank four tensor $T = T_{pqrs}$ [storing relational instances of the form $r(a_1, a_2, a_3, \text{say})$], one may wish to:

- (0) Check k the validity of a particular relational instance $r(a, b, c)$;
- (1) find the Xs such that $r(X, a_2, a_3)$ is stored in the tensor;
- (2) find the Xs such that $r(a_1, X, a_3)$ is stored in the tensor;
- (3) find the Xs such that $r(a_1, a_2, X)$ is stored in the tensor;
- (4) find the Xs such that $X(a_1, a_2, a_3)$ is stored in the tensor.

Other retrievals are possible, such as finding the pairs X, Y such that $r(X, Y, a_3)$ is stored in the tensor, but we shall restrict ourselves to modes 0–4 here. Each of the five operations required is a kind of generalized inner product. For convenience, let us define notation that distinguishes between the five kinds of generalized inner product.

- (0) $v_r \otimes v_{a_1} \otimes v_{a_2} \otimes v_{a_3} \cdot T$ checks validity
- (1) $v_r \otimes - \otimes v_{a_2} \otimes v_{a_3} \cdot T$ retrieves a_1
- (2) $v_r \otimes v_{a_1} \otimes - \otimes v_{a_3} \cdot T$ retrieves a_2
- (3) $v_r \otimes v_{a_1} \otimes v_{a_2} \otimes - \cdot T$ retrieves a_3
- (4) $- \otimes v_{a_1} \otimes v_{a_2} \otimes v_{a_3} \cdot T$ retrieves r

Now we give the computation required for each kind of retrieval. Let r_p signify the p th component of the vector v_r representing r , and let $(a_1)_q, (a_2)_r,$ and $(a_3)_s$ signify, respectively, the q th, r th, and s th components of the vectors $v_{a_1}, v_{a_2},$ and v_{a_3} representing $a_1, a_2,$ and a_3 .

(0) Let $D = \sum_{pqrs} r_p(a_1)_q(a_2)_r(a_3)_s T_{pqrs}$. If $D = 1$, then $r(a_1, a_2, a_3)$ is stored in T . Otherwise $D = 0$, and $r(a_1, a_2, a_3)$ is not stored in T . For 1–4, let \mathbf{v} be the vector representing the missing concept X , e.g., a_1 in operation 1.

- (1) $v_q = \sum_{prs} r_p(a_2)_r(a_3)_s T_{pqrs}$
- (2) $v_r = \sum_{pqs} r_p(a_1)_q(a_3)_s T_{pqrs}$
- (3) $v_s = \sum_{pqr} r_p(a_1)_q(a_2)_r T_{pqrs}$
- (4) $v_p = \sum_{qrs} (a_1)_q(a_2)_r(a_3)_s T_{pqrs}$

Similar operations can be specified for tensors of lower and higher ranks. For rank two, the operations reduce to matrix pre- and/or post-multiplication by vector(s).

It might be useful to spell out the details of retrievals for a set of binary operations represented in a rank three tensor product space, as another example. We assume a rank three tensor $T_{ijk} \in V_p \otimes V_1 \otimes V_2$ and consider the three possible “directions” of access: (1) We know the two arguments, represented by vectors $\mathbf{u} \in V_1$ and $\mathbf{v} \in V_2$, and we want to know the relation symbol(s) $\mathbf{p} \in V_p$ such that $\mathbf{p} \otimes \mathbf{u} \otimes \mathbf{v}$, that is, $\mathbf{p}(\mathbf{u}, \mathbf{v})$, is a fact “known to the tensor”; (2) we know \mathbf{p} and \mathbf{u} , and we want to find the \mathbf{v} for which $\mathbf{p}(\mathbf{u}, \mathbf{v})$ is known; (3) we know \mathbf{p} and \mathbf{v} , and we want to find out the \mathbf{u} for which $\mathbf{p}(\mathbf{u}, \mathbf{v})$ is known. In fact, $\sum_j \sum_k T_{ijk} u_j v_k$ is the answer to (1), in the sense that this expression is a vector (subscripted by i) that is the sum of all the vectors representing symbols \mathbf{p} such that $\mathbf{p}(\mathbf{u}, \mathbf{v})$ is known. Similarly, $\sum_i \sum_k T_{ijk} p_i v_k$ is the answer to (2), and $\sum_i \sum_j T_{ijk} p_i u_j$ is the answer to (3). Because the value of any argument can be computed given the values of all the other arguments, variations in any dimension as a function of the others can be computed.

APPENDIX C. PARSING OF CENTRE-EMBEDDED SENTENCES

The decision process for parsing the sentence *The boy the girl the man saw met slept* is shown in Figure 5. The rules used are spec-




Figure 5. Decision process for parsing center-embedded sentence.

ified below.

Combinatorial rules:

- 1. Subjects of distinct verbs must be distinct.
- 2. Objects of distinct verbs must be distinct.

Combinatorial/grammatical rule:

- 3. A single noun phrase cannot be both the subject and the object of the same verb.

Case/grammatical rules:

- 4. $NP \cdot TV \rightarrow \text{SNP}_{\text{OBJ}}; \text{SUBJ}(TV) = NP$.
- 5. $NP_1 \cdot \text{SNP}_{\text{OBJ}} \rightarrow NP_2; \text{OBJ}[\text{VERB}(\text{SNP}_{\text{OBJ}})] = NP_1$.
- 6. $TV \cdot NP \rightarrow VP; \text{OBJ}(TV) = NP; NP \cdot VP \rightarrow S$.
- 7. $NP \cdot IV \rightarrow S; \text{SUBJ}(IV) = NP$.

In the case/grammatical rules, TV signifies transitive verb, and IV signifies intransitive verb. The rules are context-free rules written back to front, augmented with case assignments. Thus rule 5 could be read: “If you find an NP_1 (noun phrase) followed by an SNP_{OBJ} (sentence with omitted object NP), then you have found a (higher level) noun phrase NP_2 : Set the object slot of the verb in the SNP_{OBJ} to be the noun phrase NP_1 .”

The tree at the top in Figure 5 shows the *combinatorial* choices for analysis of the sentence *The boy the girl the man saw met slept*. Subject to rules 1–3, there are 18 combinatorial possibilities, labeled P1–P18. Levels are shown at right; the 3-way split at level A shows the three-way choice of subject for the verb *slept*.

The split at level B for the subject of *met* is 2-way, because rule 1 eliminates one of the possibilities; for example, in the leftmost subtree at level B, the SUBJ cannot be *man* because *man* has already been used as the subject of *slept*. At level C, similarly, there is only one choice in each case for the subject of *saw*.

At level D, insofar as the subject of *met* has already been chosen, because of rule 3, there are only two remaining choices for the object of *met*. At level E, when choosing the object of *saw*, the NP both must not have been chosen as the object of *met* (rule 2) and must not have been chosen as the subject of *saw* (rule 3). The intersection of these possibilities is sometimes just one NP, sometimes two. For example, with P16, $\text{OBJ}(\text{saw})$ cannot be *girl* because $\text{OBJ}(\text{met})$ is *girl*, and it cannot be *man* because $\text{SUBJ}(\text{saw}) = \text{man}$. However, with P17, P18, $\text{SUBJ}(\text{saw}) = \text{OBJ}(\text{met}) = \text{man}$, so both *girl* and *boy* are possible for $\text{OBJ}(\text{saw})$.

Now, the case/grammatical rules.

Apply rule 4 to *the man saw* $\rightarrow \text{SUBJ}(\text{saw}) = \text{man}$ [possibilities satisfying this: P10, P11, P12, P16, P17, P18].

Apply rule 5 to *the girl the man saw* $\rightarrow \text{OBJ}(\text{saw}) = \text{girl}$ [possibilities: P10, P12, P17].

Apply rule 5 to *the boy the girl the man saw met* $\rightarrow \text{OBJ}(\text{met}) = \text{boy}$ [possibility: P10].

Apply rule 7 to *the boy the girl the man saw met* $\rightarrow \text{SUBJ}(\text{slept}) = \text{boy}$ [possibility: P10].

Rule 6 is not used in this example sentence, but would be used in *The boy the girl met hit Tom*.

ACKNOWLEDGMENTS

This work was supported by grants from the Australian Research Council. The authors give special thanks to Keith Holyoak, whose continuing interest in the problems and penetrating observations contributed greatly to the development of this work. We also thank many other colleagues who provided valuable feedback concerning ideas discussed in this target article. Special mention is due to John Hummel, Pat Carpenter, Marcel Just, Robert Siegler, John Flavell, Robbie Case, Ken Kotovsky, James McClelland, Alan Baddeley, John Bain, Julie McCredde, Glenda Andrews, Tracey Zielinski, Roger Wales, Doug Saddy, Campbell Dickson, Robert Dennys and Kerry Chalmers, Wayne Wicklegren, Andrew Biemiller, John Tsotsos, Lokendra Shastri, Daniel Berch, and one anonymous reviewer.

NOTES

1. More precisely, “ $x:\text{dog}(x)$ is bigger than $y:\text{cat}(y)$ ” has no truth value because of the variables present. However, the corresponding family of propositions referring to particular dogs and cats will be true in some cases (presumably most cases) and false in others, as when x is a Chihuahua and y is a decent-sized cat.

2. Note that the structural correspondence principle does not imply that the person in the lover role must always be first in the expression, but it does mean that entities in a given role must always be in the same position relative to other roles. Thus, if “John loves Mary” is represented by $\text{loves}(\text{John}, \text{Mary})$ then “Peter loves Angela” must be represented by $\text{loves}(\text{Peter}, \text{Angela})$, and so on.

3. A dimension, A, is independent of another dimension, B, if, when classifying entities according to their values on dimensions A and B, knowing the value of an entity according to dimension A does not always determine the value of the entity according to dimension B.

4. This is true even though their significance could not be determined.

5. Actually the problems of instance identification are even greater than this. Not only must each relational instance be bound to a unique context vector (i.e., the context in which the relational instance was memorized), but each vector must be orthogonal (dissimilar) to avoid the problem of cross-talk (role vectors bound to fillers from different relational instances). The orthogonality requirement introduces a dilemma: if we choose random dissimilar vectors, then in general relational instances are not distinguishable on the basis of their contents. On the other hand, if we generate identification vectors on the basis of contents, these will no longer be orthogonal, because many relational instances share the components. These problems can be overcome by defining each instance by its components and linking them together – but note that this effectively entails adopting symbol–argument–argument bindings.

6. In the special case, when each role is represented by a unique local basis vector (e.g., $[1\ 0\ 0]$), the tensor role–filler method also maps each relation dimension onto a separate tensor subspace. However, role–filler methods may entail the psychologically unrealistic assumption that each relational instance can be uniquely identified by a context independent vector.

7. This property can be linked to compositionality (Fodor & Pylyshyn 1988) but there is no space to develop that link in this target article.

8. The class $[f(n)]$ is the class of functions that are asymptotically of the same complexity as $f(n)$, that is, functions $g(n)$ such that there is a constant C for large enough n , $g(n) \leq C \cdot f(n)$, and there is a constant D such that, for large enough n , $f(n) \leq D \cdot g(n)$.

9. Order notation $\Theta(\cdot)$ identifies the term in n with the largest power, so that $100n^2$, $n^2 + 10n$, and $0.001n^2$ are considered to be of the same order (i.e., $\Theta(n^2)$ – quadratic).

10. Frye et al. (1995) refer to this as a “setting condition.”

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

ACT-R: A higher-level account of processing capacity

John R. Anderson, Christian Lebiere, Marsha Lovett, and Lynne Reder

Psychology Department, Carnegie Mellon University, Pittsburgh, PA 15213.
 ja@cmu.edu cl@cmu.edu lovett@cmu.edu reder@cmu.edu
 act.psy.cmu.edu

Abstract: We present an account of processing capacity in the ACT-R theory. At the symbolic level, the number of chunks in the current goal provides a measure of relational complexity. At the subsymbolic level, limits on spreading activation, measured by the attentional parameter W , provide a theory of processing capacity, which has been applied to performance, learning, and individual differences data.

In their target article, Halford, Wilson & Phillips propose that cognitive limitations on information processing capacity should be defined in terms of relational complexity. They argue that limits on activation, as introduced by Anderson et al. (1996), do not provide a general metric for processing complexity. In this commentary, we argue otherwise, reviewing the ACT-R theory of processing capacity and examining how it relates to the relational complexity theory of Halford et al. [See also Anderson: “Is Human Cognition Adaptive?” *BBS* 14(3) 1991.]

A central concept in the ACT-R production system (Anderson 1993; Anderson & Lebiere 1998) is the current goal, which represents the focus of attention. At each cycle, a production must first match the state of the current goal before performing memory retrievals and modifying the goal state. When a goal is successfully achieved, that goal chunk¹ becomes a declarative memory fact. Chunks are composed of a number of labeled slots, each of which holds a value which can be another chunk. Each chunk is an instance of a particular chunk type, which defines the name and number of slots. The mapping to relational knowledge is therefore fairly straightforward. Chunk types correspond to relations, with slots as arguments. Chunks correspond to relational instances, with slot values as fillers. The dimensionality of a relation equals the number of slots in the corresponding chunk type. Operations on relations, from basic omni-directional access to more complex ones such as analogy, are implemented in the manipulation of chunks by productions. The mechanisms to reduce the dimensionality of relations, chunking and segmentation, can also be used to reduce the size of goals. A new chunk (e.g., cat) can be defined as the combination of several slot values (e.g., c, a, t), then used as a single slot value in other chunks. Segmentation consists in performing a complex goal by pushing several smaller subgoals on the goal stack. The quaternary limit on relational dimensionality is generally compatible with the goal size in published ACT-R models.

We just sketched the correspondence between ACT-R and the relational account at the symbolic level. Some properties of relations, such as strength and asymmetry of access, result from subsymbolic activation computations in ACT-R, which control the retrieval of declarative chunks by productions. It is those activation computations that provide ACT-R’s account of processing complexity. The activation of a chunk, which controls its availability, is the sum of a base-level activation, reflecting its past frequency of use, and an associative activation, reflecting its relevance to the

current goal. Associative activation spreads from the current goal to declarative chunks. The associative activation of a chunk is the sum for each activation source of its source activation times the strength of association between the source and the chunk. The activation sources are defined as the slot values of the current goal, and a fixed amount of source activation W (1 by default) is divided evenly among the sources. Therefore, as the goal becomes larger, W will be divided among more sources and the resulting activation will be spread among more chunks, diluting the effect of the focus. As established first by Anderson et al. (1996) and more generally in Anderson and Lebiere (1998), this dilution of activation will result in poorer performance, that is, longer latencies and more frequent errors. In addition to affecting performance, large goal sizes also hinder learning. Lebiere (in preparation) establishes that a set of related chunks can be reliably learned only if the source activation for each chunk component is higher than the activation noise level. Therefore, given a particular noise level, simple facts (e.g., counting) might be learned but more complex facts (e.g., addition) might not because their components have lower source activation. This suggests that the gradual increase in processing capacity reported in the developmental data described by Halford et al. could be accounted for by a continuous increase in W . Finally, Lovett et al. (1997; in press) relate W to individual differences. They fit a range of subject performance on working memory tasks using a single ACT-R model, with high-performance subjects modeled by larger W and low-performance subjects modeled by smaller W . All those results point to W as the basic measure of processing capacity in ACT-R.

ACT-R is not unrelated to the neural network models presented by Halford et al. Lebiere and Anderson (1993) presented ACT-RN, a neural network implementation of ACT-R which uses essentially a positional encoding of symbol-argument-argument bindings, with clean-up memories as in convolution models. But we are not committed to any specific connectionist representation since, as Halford et al. report, they have the same basic properties. Just as theoretical computer science proved the equivalence of various computational paradigms in order to establish proofs of complexity valid for all, ACT-R aims to provide a higher-level definition of processing capacity independent of any lower-level neural representation.

NOTE

1. This meaning of chunk is somewhat different from Miller's (1956).

Processing demands associated with relational complexity: Testing predictions with dual-task methodologies

Daniel B. Berch^a and Elizabeth J. Foley^b

^aChild Development and Behavior Branch, National Institute of Child Health and Human Development, Bethesda, MD 20892-7510; ^bDepartment of Psychology, University of Cincinnati, Cincinnati, OH 45221-0376. db254g@nih.gov

Abstract: We discuss how modified dual-task approaches may be used to verify the degree to which cognitive tasks are capacity demanding. We also delineate some of the complexities associated with the use of the “double easy-to-hard” paradigm for testing claim of Halford, Wilson & Phillips that hierarchical reasoning imposes processing demands equivalent to those of transitive reasoning.

Halford, Wilson & Phillips note that their theory would be invalidated if relational complexity did not predict processing demands (sect. 6.3.2). However, carrying out tests of such predictions is by no means straightforward. Halford et al. provide examples of various ways in which this may be done, suggesting the use of concurrent probe reaction time (sect. 6.3.2). This variant of the standard dual-task approach has certainly enjoyed extensive utility for

indexing processing resources in experimental studies with adults, but it has seen only modest applicability to investigations of cognitive development. Numerous criticisms have been leveled at this approach, however, because dual-task deficits can frequently be attributed to various types of concurrence costs.

Halford (1993) has also noted some of the same difficulties and has espoused the use of the easy-to-hard paradigm developed by Hunt and Lansman (1982; see also Lansman & Hunt 1982). Essentially, this paradigm constitutes a modified dual-task approach where in addition to performing a secondary task and a relatively easy version of the primary task in isolation as well as concurrently, subjects perform a harder version of the primary task in isolation only. The latter manipulation eliminates interpretations based on concurrence costs.

The notable advantages of the easy-to-hard paradigm notwithstanding, it has thus far been used in only a handful of studies for examining capacity limitations in children. This may be partly owing to how labor-intensive it is to carry out this technique with children. We have found that computerized tasks are superior in this regard, in part because it is easier to maintain young children's interest throughout the administration of several different tasks (Foley & Berch 1997). Furthermore, distinct efforts must be made to adhere closely to the recommendations for enhancing dual-task methodology, such as varying the modality of the stimuli in the primary and secondary tasks and reducing the likelihood of interference at input.

However, even the successful use of this paradigm has limited applicability in testing more sophisticated facts of the Halford et al. model, in that it is designed primarily to determine whether any single given task is capacity limited. For example, Halford et al. claim that class inclusion and transitivity are essentially equivalent with respect to the capacity demands they impose because they both putatively require the processing of ternary relations (sect. 6.3.2). As it turns out, Halford (1993) has previously suggested that one could use what he refers to as the “double easy-to-hard” approach for this purpose. One must administer seven tasks: two hard primary tasks presented alone, two easy primary tasks presented alone, two easy tasks presented concurrently with a secondary task, and the secondary task presented alone. There are numerous complexities associated with this approach, such as the need to make the level of difficulty of the easy and hard versions of both tasks as equivalent as possible.

Attempting to use the double easy-to-hard approach raises a number of theoretical and methodological issues that have not yet been considered, with important implications for successful implementation of this paradigm. First, in addition to meeting the dual-task assumptions of the basic easy-to-hard paradigm, the double easy-to-hard approach requires that an additional assumption be satisfied prior to performing the critical tests that permit one to verify whether two hard primary tasks are equivalent in the processing demands they impose. Specifically, one must determine whether the dual-task measures critical for the double easy-to-hard correlations reflect capacity demands that draw on one general resource pool or on multiple pools. The former assumption would be preferred, but if it proved to be incorrect, then application of the double easy-to-hard approach could yield inaccurate conclusions.

We developed a method for testing this assumption in attempting to compare the processing demands of class inclusion and transitive inference tasks, and obtained evidence that these measures tap a general resource pool (Foley 1997). We had hypothesized that these tasks may not be equivalent because of differences in dimensionality. Recall that Halford et al. define dimensionality as the number of factors that vary independently (sects. 2.3 and 3.2). Although this definition is sufficient for the transitive inference task, it does not seem applicable to the class inclusion problem. This is because the concept of class inclusion is based on understanding that there is an inherent hierarchical relationship between the superordinate class and both the subclass and the complementary class, meaning that certain components of the

problem are dependent on others. For example, if the size of either subset is changed, the size of the superordinate class changes. The added complexity associated with processing this dependent relationship may yield increased capacity demands as compared with the resources needed for solving a transitive inference problem. Although the results did not conclusively support our position, there was some evidence to indicate that class inclusion and transitivity may not require equivalent processing demands.

In sum, we argue that modified dual-task approaches are going to be of crucial importance in testing claims of equivalent processing demands for cognitive tasks presumed to be at the same level of relational complexity. Such tests are critical for validating the theoretical model of Halford et al.

Synchronization of neural activity and information processing

Roman Borisyyuk,^{a,b} Galina Borisyyuk,^b and Yakov Kazanovich^b

^aSchool of Computing, University of Plymouth, Plymouth, Devon PL4 8AA, United Kingdom; ^bInstitute of Mathematical Problems in Biology of the Russian Academy of Sciences, Pushchino, Moscow Region 142292, Russia
borisyyuk@soc.plym.ac.uk; kasanov@impb.serpuhov.su

Abstract: Synchronization of neural activity in oscillatory neural networks is a general principle of information processing in the brain at both preattentive and attentional levels. This is confirmed by a model of attention based on an oscillatory neural network with a central element and models of feature binding and working memory based on multi-frequency oscillations.

Two types of neural network models of relational knowledge are considered in Halford et al.'s target article: (1) models based on product operations to bind entities; and (2) oscillatory neural networks that use synchronization for feature binding. We would like to demonstrate that the synchronization of neural activity can be considered as a general principle of information processing in the brain at both the preattentive and the attentional levels.

Oscillatory neural network models, synchronization, and attention modeling. Kryukov (1991) has proposed an oscillatory neural network with a central executive as a model of attention. In Kryukov's model the central executive is an oscillator, the central oscillator (CO), which is coupled with other oscillators, the peripheral oscillators (PO) by feed-forward and feedback connections. It is presumed that the septo-hippocampal region plays the role of the CO, while the POs are represented by cortical columns sensitive to particular features. This concept is in line with Damasio's (1989) hypothesis that the hippocampus is the top vertex of the convergent zone pyramid and Miller's (1991) theory about the representation of information in the brain based on cortico-hippocampal interplay.

Attention is realized in the network in the form of synchronization between the CO and some POs. We have found (Kazanovich & Borisyyuk 1994) that selective attention can be associated with the regime of partial synchronization when the attentional focus is supposed to be formed by those POs that work synchronously with the CO. One of the results obtained is the formulation of conditions in which decreasing the interaction of the CO with the oscillators representing one of two stimuli that form the attentional focus can lead not to focusing attention on the other stimulus but to destruction of the attention focus. For some parameter values it has been found that the CO is capable of synchronizing alternately with one or the other of two groups of POs. This can be interpreted as a spontaneous switching of attention which is observed in some psychological experiments.

Estimations of complexity of neural dynamics. Halford et al.'s idea of measuring the complexity of a cognitive process by the number of interacting variables that must be represented in par-

allel seems very promising. In the last few years a number of investigations have shown that some modes of brain activity can be characterized as low-dimensional chaos and that the complexity of the dynamics of neural activity can be measured by the dimensions of an attractor (Babloyantz 1989).

We have applied the correlation dimension to studying the properties of the attentional model with a central oscillator. In that context the model has been used as a generator of EEG activity. The main result of our studies is that the correlation dimension of the simulated EEG can be used as a measure of the complexity of neural activity in the system. This measure is equal to the number of groups of cortical oscillators working at different frequencies (Borisyyuk et al. 1994), an important feature of the attentional system because it shows the diversity of stimuli that are out of the focus of attention.

Multi-frequency oscillations, feature binding, and working memory. It is known that frequency encoding of stimuli is hampered by insufficient informational capacity. For example, in the target article the synchronous oscillation model suggests that only about 5–10 entities can be processed in parallel due to limits in frequency resolution. Indeed, the range of permissible frequencies is not large; hence one frequency can be separated from another with only limited accuracy. This implies restrictions in frequency and phase encoding. The application of double-frequency oscillations makes it possible to extend frequency encoding because the second frequency plays the role of the second variable. Thus, two coordinates can be used for encoding instead of one.

Borisyyuk et al. (1995) suggest that feature binding is organized hierarchically: simple features of one type are bound in the primary cortex while combinations of different simple features are derived in higher cortical areas. The main idea of this approach is to use multi-frequency (in particular, two-frequency) envelope oscillations in such a way that simple features are bound at a high frequency while low frequency synchronization is used to bind compound features of a complex stimulus.

Halford et al.'s statement that models of associative memory always need a large number of training examples to learn appropriate representations seems mistaken. A counterexample is an oscillatory model of the hippocampus that is able to memorize several sequences of "lessons" and to recall the sequence related to initial associations (Borisyyuk & Hoppansteadt 1998). We think this can be used as a working memory model suitable for relation learning.

On the psychological reality of parallel relational architectures: Whose knowledge system is it anyway?

Margaret Chalmers and Brendan McGonigle

Department of Psychology, Edinburgh University, Edinburgh EH8 9QT, Scotland, U.K. m.mcgonigle@ed.ac.uk or ejua48@ed.ac.uk
www.psy.ed.ac.uk/nis

Abstract: We argue that Halford et al.'s characterisation of relational complexity offers an unadaptive principle in terms of cognitive economy, that its relation with the empirical evidence is highly selective, and that the task behaviours used in support of a multivector processing space are better described by linear serial processes which do not require n-dimensional mappings for their emergence.

Halford et al. make a commonplace claim that relations provide knowledge which is "symbolic, content-independent, flexible and modifiable." The novelty in their thesis is in laying such conceptions directly at the feet of capacity models of working memory and in offering a gradable construct by which the processing in the central executive can be both characterised and measured. This is a timely and well justified exercise in our view, but one which will

surely fail (as others have done before) if the construct it offers is wrong. We believe that it is wrong.

Taking its lead from a formal rather than an empirically founded conception of what relational complexity is, Halford et al. commit what we have argued (Chalmers & McGonigle 1998) to be the (same) fundamental mistake as Piaget, and in so doing provide a model for complexity which is top-heavy in computational cost and consequentially restricted in adaptive value.

Reminiscent of the Piagetian concept of logical mother-structures, the complexity construct described in this target article is based on a co-present system of relations used as a template against which cognitive level can be gauged. Implying that long-term architectures (of varying levels of dimensionality) are available to the subject during each step of problem solution places a huge burden of simultaneity on working memory. Yet the fact is that even adult humans show normatively what appears to be a restricted moment to moment access to relationally based understanding – a partiality which calls not for hand-wringing but for insight. As many have argued (e.g., Clark 1969), premise encoding during the 3–5 term series problem produces psychologically privileged representations, whose selectivity trades off in the long term against the computational cost of “omni-directional access” (sect. 2.2.6; see also Clark & Thornton 1997). Unlike the characterisation of complexity which Halford et al. appear to derive from possible relational mappings, psychological representations (e.g., the linear spatial devices) are often also interpreted as the product of unidirectional serial architectures (see e.g., Potts 1972). Such representations do not preclude the recovery of other connected or implied relations (McGonigle & Chalmers 1986), but only when considered as secondary and derivative, that is, not part of the same working memory content. Similarly, more normative examples of sentence production (than the tortuous examples of parallel parentheses Halford et al. cite (sect. 2, para. 1) would illustrate these twin themes within natural language: directional partiality in word order and the derivative nature of the more complex syntactic forms (as in the passive).

Normative matters aside, we would argue that seriality is a useful design feature, conferring inductive power and processing economy, yet leading to the self-same examples of cognitive achievement discussed by Halford et al. In the case of transitivity, binary level success ($A \vee C$), can lead to success at a triadic level ($A \vee B \vee C$) by the serial extension of a binary solution (McGonigle & Chalmers 1977; 1992). Considered a ternary relation by Halford et al. (sect. 3.4.3, para. 7), the number of items (arguments) in a transitivity test should not affect its level of solution as the attribute (e.g., a size relation) is unidimensional; but when modelled using a rule stack based production systems approach in which binary decisions were serially prioritised (Harris & McGonigle 1992), it was formally demonstrated that relational transitivity is guaranteed by the binary choice structure in a n -term series task, whereas triadic tests requires greater depth of serial search for their resolution. Thus, we would argue that complexity is related to the number of items in the decision space, not to transitivity per se; representing transitivity as a vector in its own right thus imposes a dimensional load with no proven psychological reality.

For Halford et al., for whom young children and monkeys are in any case unlikely to be capable of representing transitivity, such results could be viewed as examples of an a priori working memory restriction (e.g., context), however, evidence from seriation development in older children is highly instructive: Success at seriation emerges normatively at around the age when Halford et al. claim young children make “genuine” transitive inferences using ternary mappings. It could therefore be solved at this age as a series of transitive relations taking several arguments bound along a single dimension such as size. However, on Halford et al.’s reduction argument (sect. 3.4.3, para. 7), if no information is to be “lost” and the place of every item in the set is made fully determinate, then a task analysis of 8 item seriation would suggest an oct-ternary relational solution – well beyond the working memory bounds of an adult human, let alone a seven-year-old child.

This example exposes the unnecessary burden imposed by the parallel analysis offered by Halford et al. A task like seriation, requiring the recruitment of a linear ordering architecture, could be performed on Halford et al.’s model by parsing into “appropriate chunks” (sect. 3.5, para. 1), by using, for example, a set of ternary evaluations carried out serially. But if this is an option, why, given the added information processing burdens imposed by their model, should the agent not “go serial” at the binary level and thereby reduce resource demands without cost to the actual organisational achievement? In classification tasks, for example, extended search through a hierarchically nested set of relations leaves semantic differentiation fully preserved (“mother of” first, then “loves,” “feeds,” etc.) – but at much less cost to working memory.

In short, we reject the idea that cognitive competences in evolution and development can be reduced to a look-up table based on restrictions on working memory capacity. We would argue instead that, rather than loading up our working memory capacity to the point of extinction, cognitive complexity is designed to diminish information processing demands through selective serial processing.

Discontinuity and variability in relational complexity: Cognitive and brain development

Donna Coch and Kurt W. Fischer

*Department of Human Development and Psychology, Harvard University
Graduate School of Education, Cambridge, MA 02138.*

cochdo@hugse1.harvard.edu kurt_fischer@harvard.edu

Abstract: Relational complexity theory has important virtues, but the present model omits key aspects and evidence. In contrast, skill theory specifies (1) a detailed series of developmental changes in relational complexity from birth to age 30, (2) processes of interaction of content and structure that produce variability in complexity, (3) the role of cortical development, and (4) empirical criteria for complexity levels, including developmental discontinuities. Many findings support these specifications.

The relational complexity theory outlined by Halford, Wilson & Phillips resonates well with similar concepts in dynamic skill theory (Fischer 1980), which goes beyond the “broad parallel” that the authors acknowledge (sect. 6.3.2, para. 10). Indeed, relational complexity defines the major cognitive developmental changes that skill theory specifies in detail. Unary, binary, ternary, and quaternary relations are reminiscent of single set, mapping, system, and system of systems in skill theory, which describe a developmental sequence of increasing complexity and connectivity among skills. However, skill theory elaborates how these levels move beyond the four types, specifying an extended series of relations that are evident in development and variability in action and thought and which relate to cortical growth. These levels have been assessed in multiple domains across wide age ranges, with specific empirical criteria for each type of relational complexity.

The skill levels are cyclically recurring ones (substages) within a superstructure of four tiers (reflex, sensorimotor, representational, and abstract), which involve applying the four types of relations to themselves. In this way, each cycle through the four relations produces movement to a higher-order version of relation. Both the levels within a tier and the tiers themselves instantiate the four relations, but at a different order of complexity. Each level and tier is marked by a cluster of discontinuities in growth functions for optimal performance (Fischer & Bidell 1997).

Positing such a hierarchy is more than a “descriptive convenience” (sect. 6.3.2, para. 15), because it predicts and explains changes in both development and variability as evidenced in learning and problem solving. It explains, for example, qualitative differences between binary relations in young children and adolescents (across tiers). A 4-year-old girl functioning at her upper limit can understand interactions in which two characters (dolls or peo-

ple) interact to form reciprocal parental roles – a binary relation (representational mapping) that maps activities in the mother role to those in the father role (Watson 1984). However, not until 15 to 16 years can a person understand the binary relation (abstract mapping) involved in the motivational concept of a social lie, in which action is motivated by honesty in relation to kindness – lying in order to protect a friend's feelings (Lamborn et al. 1994). Such findings demonstrate Halford et al.'s generalization that “what is acquired depends on both experience and capacity” (sect. 6.3.2, para. 3), but more precise concepts and methods can predict such developmental changes.

Skill theory also explains how people of any age use skills from a wide range of levels, varying from their upper limit on down. For example, development at lower levels continues when people are capable of higher levels. In much learning and problem solving, people build new lower-level skills to deal with novel contexts and issues. Global higher-level shells are used as bridges to construct new lower-level skills, forming foundations for novel actions and concepts (Granott et al. 1997). A person's movement to a lower level, elaboration of skills at that level, and subsequent gradual construction of higher-level skills (relations) can be directly analyzed, producing a portrait of the construction (microdevelopment) of a new skill.

With its empirical criteria for emergence of types of relations (levels and tiers), skill theory has made possible the first model relating cognitive development to specific changes in cortical functioning (Fischer & Bidell 1997; Fischer & Rose 1996). The model outlines how neurocognitive networks are reorganized over developmental time. The basic mechanism for building skills of higher relational complexity is (a) learning two simpler skills or networks, (b) mastery to allow parallel sustaining of the two, and (c) coordinating them in a new relation through multiple steps specified by skill/network combination rules.

As Halford et al. note (sect. 6.5), ample evidence indicates that frontal cortex is well-suited to holding information on-line, which is required to support coordination and increase in complexity. The development of networks seems to involve prefrontal coordination of components to form new networks, as described by Thatcher (1994). The development of these networks is evidenced by changes in cortical activity and connectivity measured by EEG power and coherence, which show growth cycles moving systematically through cortical regions. In research to date, the correspondence between cortical growth cycles and behavioral discontinuities for behavioral levels and tiers is remarkable.

One of the cortical discontinuities involves abrupt shifts in oscillation frequency in cortical coherence, giving the appearance of cusp catastrophes. These oscillation shifts suggest major change in neural network functioning and correspond closely to ages of emergence of skill levels. Such networks consist of the same basic components (skills or neurons) related in dynamic new ways, illustrating simultaneous continuity and discontinuity in development.

Halford et al. neglect these kinds of evidence in arguing that relational complexity “does not imply that development is discontinuous” (sect. 6.3.2., para. 5). Guttman sequencing of tasks and cross-domain correlations do not provide sufficient evidence for the reorganizations in Halford et al.'s version of relational complexity theory. Discontinuities and variations in development of behavioral complexity and cortical activity provide valuable tools for research on relational complexity.

What is more explanatory, processing capacity or processing speed?

Nelson Cowan

Department of Psychology, University of Missouri, Columbia, MO 65211.

psycowan@showme.missouri.edu www.missouri.edu/~psycowan/

Abstract: Halford et al. have sharpened the concept of processing capacity as applied to various complex tasks. This commentary examines the apparent contradiction between capacity theories and theories in which it is processing speed, rather than capacity, that presumably limits cognitive performance. It explains how capacity and speed often are interrelated and suggests how one might examine whether capacity or speed is the more elementary in processing.

Halford, Wilson & Phillips have sketched out a powerful and sensible mechanism that predicts performance constraints, as well as individual and developmental differences, on a wide array of cognitive tasks, all based on “processing capacity defined by relational complexity.” At the level of analysis that the authors chose in order to examine the phenomena, I am enthusiastic about the theory.

In the broader view, it seems likely that Halford et al. will soon have to confront other levels of analysis. One level lower would be a microanalysis of performance on the tasks. The question would be just *how* the capacity limit causes failures. An analogy might help here. Suppose one learns that a restaurant has the capacity to seat 50 people at once. A natural assumption is that the limit is in the size of the restaurant. However, that might not be the case. Instead, perhaps the kitchen is unable to serve more than 50 people without intolerable delays, and only enough tables were installed to match the kitchen's capability rather than exceeding it. In that case, the supposed capacity limit is actually a speed limit when viewed at a finer level of analysis. To pursue the analogy further, the answer might change again when viewed in even more detail. Perhaps a limit in the size of the kitchen is the fundamental reason that customers cannot be served more quickly, so that one capacity limit (in kitchen size) causes a speed limit (in serving food) that in turn causes another capacity limit (in the available seating).

This example was selected for an important reason. A leading contender for human processing limits is the speed of processing. Numerous investigators have championed changes in the speed of processing as a primary cause of processing limits (e.g., Fry & Hale 1996; Kail & Salthouse 1994; Salthouse 1996). The notion would be that the limit in capacity can ultimately be accounted for by a limit in speed. For example, a certain processing speed seems necessary to reactivate items in working memory before they become deactivated to the point of being inaccessible (cf. Baddeley 1986; Kail & Park 1994; Schweickert & Boruff 1986).

The speed versus capacity issue looks at first glance like a classic chicken-and-egg problem (i.e., which came first?) but I suspect that the issue of causal priority can eventually be resolved. One potentially relevant type of evidence is a fine-grained analysis of the timing of spoken recall. Results of such analyses suggest that capacity limits may sometimes cause speed limits rather than the reverse. Cowan (1992) tested auditory word span in 4-year-old children, measuring the duration of each word and each inter-word pause in the spoken responses in the span task (only for the trials in which the response was totally correct). The response word durations did not differ for children with lower versus higher spans, however, there were differences in the durations of the pauses between words in the responses (which changed little across serial positions). Longer lists resulted in longer inter-word pauses, suggesting that before speaking each word a child had to select the item to be pronounced next from other items in the list. Moreover, for lists of a fixed length, children with a lower span had longer inter-word pauses in their responses. When children were compared on lists at their own span length, this difference between children disappeared; speaking rates were comparable, and more capable children just spoke for a longer total duration.

Cowan et al. (1994) obtained similar results with a develop-

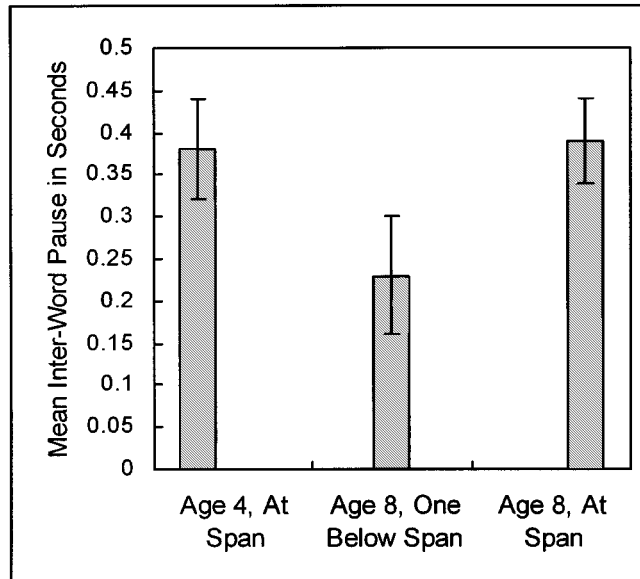


Figure 1 (Cowan). Inter-word pause times in responses to spoken lists in a word span task. Data are from Cowan et al. (1994). For the 4-year-olds (N = 16), the mean span was 3.00; for the 8-year-olds with recorded below-span data (N = 15), the mean list length for the displayed pause data was 3.51. The third bar is based on all 8-year-olds in the sample (N = 23), whose mean span was 4.02. Notice that the younger and older children's pause durations within the recall of span-length lists were nearly the same, but that the older children produced shorter pauses within lists of a more comparable length (1 below span). Error bars reflect standard error.

mental sample of 4- and 8-year-old children. For lists of comparable lengths, 4-year-old children paused between words longer than 8-year-old children. When the children in the two age groups were examined on lists equal to their own span lengths, this difference in pause durations disappeared. The pattern of results is shown in Figure 1. The importance of these findings is that the memory load limited the speed of responding, in contrast to the a priori expectation that speed of responding would remain fixed across list lengths and span would be limited by how many words could be spoken in a fixed period before memory decayed (e.g., Schweickert & Boruff 1986). Capacity actually determined speed.

Further research suggests that some speeds may be elementary whereas other speeds may be derived from capacity limits. Cowan et al. (1998) examined several measures of each of several constructs: (a) forward digit span, (b) interword pauses at fixed list lengths in the span tasks to estimate the speed of retrieval from working memory, and (c) rapid speaking tasks to estimate the possible speed of covert rehearsal (Baddeley 1986). Span turned out to be correlated with both retrieval and rehearsal speeds (at about .4 in each case). However, there was no correlation at all between retrieval and rehearsal speeds. In one account of these results, retrieval speed would be influenced by storage capacity limitations. It would be related to limitations in central executive processes in the model of Baddeley (1986) and limitations in the focus of attention in the framework of Cowan (1995). In contrast, rehearsal speed would be unrelated to these factors and related instead to phonological loop processes (Baddeley 1986) or activation outside of the attentional focus (Cowan 1995).

This body of research is not definitive but does illustrate that one man's cause becomes another man's effect (with either capacity or speed playing the role of either cause or effect), and that more work is needed to map out the true causal paths between such variables.

ACKNOWLEDGMENT

This work was supported by National Institutes of Health grant R01 HD-21338.

Complexity: From formal analysis to final action

Douglas Frye^a and Philip David Zelazo^b

^aGraduate School of Education, University of Pennsylvania, Philadelphia, PA 19104; ^bDepartment of Psychology, University of Toronto, Toronto, ON, Canada M5S 3G3. doug@psych.nyu.edu zelazo@psych.utoronto.ca

Abstract: Relational complexity provides a metric for measuring task demands, and in this respect has much in common with the cognitive complexity and control theory. However, relational complexity does not account for the relative difficulty of different relational types, and appears to underestimate the importance of changes in children's ability to act on the basis of their understanding.

Complexity: From formal analysis to final action. The recent, recurrent pattern in developmental psychology has been one of rigid debates about whether any given development occurs early or late, and whether different developments can better be explained by domain-general or domain-specific accounts. Developmental psychology is commonly recognized as the study of change, but without a method for ordering the changes, the phenomena become as disorganized as those in the physical sciences would be without a periodic table.

Halford's target article represents a reply to this problem through the consideration of complexity (see also "Cognitive complexity and control theory," Zelazo & Frey 1997). Complexity provides a means of ordering achievements within and across domains. The article takes this approach a long way. It gives a well formed definition of complexity in terms of the number of relations or dimensions that can be processed in parallel, and shows how the definition makes unique predictions for the difficulty of given problems. It provides a standard procedure – reduction analysis or whether a set of relations can be decomposed and recombined without a loss of information – for calculating complexity. It also identifies two processes, chunking and segmentation, that can alter the complexity of a problem. These processes complicate the picture, but there is overwhelming evidence that they occur and the theory is right to include them.

The progress offered by this approach is illustrated in one area in which it might be least expected. A staggering amount of research in developmental psychology has lately been devoted to children's theory of mind – that is, children's understanding of their own and others' mental states (for a review, see Astington 1993). It is commonly assumed that theory of mind is a domain-specific acquisition and hence, cannot be a function of domain-general cognitive changes, such as changes in relational complexity.

Nonetheless, both the target article and Cognitive Complexity and Control (CCC) theory (Frye et al. 1995) have proposed that children's theory of mind is in part an outcome of changes in relational complexity. The accounts agree that the 3- to 5-year change in mental state understanding involves a specific cognitive change that the target article identifies as the inception of ternary relations and CCC theory as embedded rules. Both argue that the understanding of someone's false belief about a situation, for example, requires that the child form a three-term relation. In this new relation, the child's own view of the situation and the other's conflicting view are prefaced by a conditional that makes it possible to select between them, while also recognizing that the two perspectives apply to the same thing.

This agreed-upon complexity proposal has several important advantages over domain-specific approaches. One is that it explains why young children initially make realist mistakes. When 3-year-olds try to make sense of false belief, they are only able to

consider two of the terms in the three-relation (see sect. 3.5). Hence, 3-year-olds default to the more familiar judgment (Frye et al. 1996) and make the realist error of concluding that everyone will see the situation as it really is.

Another advantage of complexity, one that may be unique, is that it manifests what the target article labels “omni-directional access” (sect. 2.2.6). When children master a set of relations, there are a number of different, equivalent problems they will be able to solve. Consequently, it becomes possible to predict that seemingly unrelated developments will have a common source. For example, it is well established that children’s understanding of false belief is connected to their ability to act deceptively (Wimmer & Perner 1983) although this link could be conceptual, given belief is involved in both. However, theory of mind has now been shown to be related to executive function (for a review, see Zelazo et al. 1997) – that is, the development of children’s abilities to control their own actions. Mental state understanding and control of one’s own actions do not share an obvious conceptual link, yet their similarity can be explained in a complexity account (Frye, in press; Zelazo et al. 1997).

Despite the similarities, there are some points of disagreement between the complexity proposals of the target article and CCC theory. Both now indicate that a conditional (basically, an *or* that selects between perspectives – the child’s and the other person’s in false belief) forms the ternary relation that supports the acquisition of theory of mind. However, the relational complexity account fails to recognize the hierarchical dependency that obtains among dimensions in this instance (although it could in principle; see sect. 4.2.5). For example, one cannot tell by looking at the ternary relation, Seen-Object (<condition>, <object-colour>, <percept>), whether one argument controls the others. In contrast, on the CCC account, the task is difficult precisely because it requires a hierarchical rule structure.

Moreover, the relational complexity account does not appear to consider the relation’s type when predicting developmental difficulty. All that seems important is the complexity of the relation. CCC theory distinguishes among relational types in addition to structural complexity. Theory of mind involves choosing between perspectives, ignoring a familiar one to judge from the other, which will be simpler than combining perspectives, even though both require ternary relations. Being able to select flexibly between dimensions occurs earlier than the feat of combining them (Frye et al. 1995, Experiment 3). The distinction among relational types makes it possible to explain why theory of mind appears at about 4 years, rather than at the end of preschool when conservation and matrix classification – the classic examples of combining dimensions – occur.

The target article also appears to underestimate other control or performance factors, with the consequence that it neglects the variety of variables that affects the likelihood children will actually use their knowledge to guide action in any particular situation. Success on problems such as the A-not-B task is not simply a matter of unary relations. There is ample evidence that children who can master the relational complexity nonetheless commit the A-not-B error under certain circumstances. For example, Zelazo et al. (1998) found that 24-month-olds perseverated on a multistep, multilocation version of the A-not-B task. Presumably these children understood the relational complexity of the problem: adding steps and locations does not involve adding a new dimension. Yet, they apparently had trouble maintaining attention to the relevant information and inhibiting a tendency to repeat a previously rewarded response. The point is that actual action cannot be explained by logical relations and domain-specific knowledge alone. The relational complexity approach adds a great deal to our understanding of understanding, but it under-emphasizes the importance of the control processes that extend from understanding to action.

Deep thinking in children: The case for knowledge change in analogical development

Dedre Gentner^a and Mary Jo Rattermann^b

^aDepartment of Psychology, Northwestern University, Evanston, IL 60208;

^bDepartment of Psychology, Franklin and Marshall College, Lancaster, PA 17604. gentner@nwu.edu m_rattermann@acad.fandm.edu

Abstract: Halford, Wilson & Phillips argue that cognitive development is driven by increases in processing capacity. We suggest that changes in relational knowledge are as important or more so. We present evidence that 3-year-olds’ analogical abilities are sharply improved by teaching them relational labels; over a 30-minute experimental session children gained approximately 2 years in effective performance. These results mandate caution interpreting age-related change as indicating maturational change, and call for a deeper consideration of the role of epistemological change in cognitive development.

In their fascinating and provocative target article, Halford et al. present developmental and comparative evidence for the thesis that cognitive development is driven by increases in processing capacity. We challenge this account as it applies to the development of analogy. In Halford’s relational complexity account, children’s ability to carry out analogical matches increases maturationally with increasing processing capacity, from unary relations (object matches) through binary relations, ternary relations, and finally quaternary relations. We also have theorized a shift from objects to relations in the development of analogy (Gentner 1988; Gentner & Rattermann 1991; Gentner & Toupin 1986). However, we postulate a change in the depth of the relational structures that can be matched (Gentner 1993; Gentner & Markman 1997), rather than in the dimensionality of the relations. The sequence we propose is, first, object similarity matches (e.g., red ball/apple); then, relations between objects (e.g., ball rolling on table/toy car rolling on floor); and then higher-order relations between relations (e.g., ball’s rolling causes glass to fall/car’s rolling causes vase to fall). We further suggest that change of knowledge, not change of processing capacity, is the main driver of this evolution (Brown 1989). This knowledge-based claim is supported by the observation that the relational shift occurs at different ages – ranging from infancy to adulthood – in different domains, depending on level of knowledge.

A key prediction in Halford et al.’s account is that children under 4–5 years cannot process ternary relations nor integrate two binary relations. Halford et al. consider one challenge to this claim, namely, Goswami’s (1995) studies suggesting that 3-year-olds can map ternary relations in analogies. Their reanalysis shows that in Goswami’s task the correct relational response was often the best object-similarity match (based on closest absolute size), compromising her evidence for relational processing. However, our studies using similar methods pose a stronger challenge (Gentner & Rattermann 1991). We placed object similarity in direct competition with relational similarity (using Gentner and Toupin’s cross-mapping technique), so that the relational match was never supported by object similarity. We showed 3-, 4- and 5-year-olds two triads of objects, each arranged in monotonically decreasing size (e.g., 4 3 2 → 3 2 1) in a fanlike pattern. The child watched the experimenter place a sticker under an object in her set and then searched in his own set for a sticker hidden under the corresponding object. Because of the cross-mapping, matches based on object similarity (e.g., 3 → 3) competed against matches based on relational similarity (e.g., 3 → 2). The child received feedback on the correct answer, which was always based on relational similarity.

The results show a strong developmental change, from 43% correct in 3-year-olds to 82% correct in 5-year-olds. So far this pattern is consistent with either a maturational or a knowledge-based account. However, further results suggest that the difference between the 3-year-olds and the 5-year-olds lay in their knowledge of the relevant relations: specifically, the higher-order relation of




Figure 1 (Gentner). Proportion correct relational responses for different ages on a cross-mapped analogy task with relational labels (solid line) and without them (dashed lines).

monotonic change that governs the two lower-order relative size relations.

According to the knowledge-based account, increasing children’s relational knowledge can lead to improved analogical performance. To test this, we taught a group of 3-year-olds to use the relational labels “Daddy/Mommy/Baby,” which convey the relation of monotonic change. When the labels were used on each trial, 3-year-olds were 84% correct, performing comparably to the 5-year-olds in the original task (see Fig. 1). (Similar results were obtained with the relational labels “big/little/tiny.”) The beneficial effects of learning relational labels were not dependent on direct modeling. After experience with the “Daddy/Mommy/Baby” labels, 3-year-olds maintained much of their gain (57%) when given new stimuli on which no labels were used. Indeed, 3-year-olds who returned four to eight weeks after the initial session produced significantly more correct relational responses (62%) than a control group without label experience (28%; Rattermann & Gentner, in preparation).

Why would these labels be so effective that the benefits persist four to eight weeks later? Our explanation is that the “Daddy/Mommy/Baby” labels act to invite a higher-order relational structure. With this structure, the relational mapping is sufficiently deep and internally constraining that it can compete successfully with the local object match. (Markman & Gentner [1993] found a similar relation between relational depth and object similarity in adults.) These striking changes in performance are a testament to the importance of learned relational structure in children’s analogical mapping, and by extension in the development of relational reasoning in general. These findings challenge Halford’s account of development, first, because they show that 3-year-olds can carry out mappings based on what in his account are ternary relations; and, second, because they establish that striking changes of performance – equivalent to that of a 2-year-age gain – can occur even over the course of one experimental session, based on the acquisition of relational knowledge.

Halford has suggested that 3-year-olds succeed in our cross-mapping task not by representing the relationship of monotonic change across the three objects (a ternary relation that should be beyond their capacity), but rather by chunking the three objects into two sets – *the chosen object* and *the other two objects* – thus permitting a simple binary relation (Halford et al. 1995). For example, the child would process a Mommy → Mommy match by encoding two chunks: [Mommy] vs. [Daddy and Baby]. But this encoding seems implausible. First, it would require the children to repeatedly rechunk the stimulus set, as the chosen object

changes on each trial. Second, in middle-object examples like the above, a split chunk would be required. Third, it is hard to see how such a binary encoding would result from the use of the triadic relationship “Daddy/Mommy/Baby.”

Halford and his colleagues acknowledge that knowledge accretion plays a role in development, but in their theory the major impetus is maturational change in processing capacity. Clearly, neurological change influences children’s developing abilities. But as the above results demonstrate, one needs to be cautious about interpreting age-related change as resulting from capacity change. We believe that our evidence calls for deeper consideration of the role of epistemological change in cognitive development.

Is relational complexity a useful metric for cognitive development?

Usha Goswami

Institute of Child Health, University College London, London WC1N 1EH, England. u.goswami@ich.ucl.ac.uk www.ich.bpmf.ac.uk

Abstract: This commentary focusses on the evidence used by Halford et al. to support their postulated links between relational complexity and age differences in children’s understanding of concepts. None of their developmental claims is consistent with recent cognitive-developmental research. Relational complexity must be an important variable in cognition, but it does not provide a satisfactory metric for explaining cognitive development.

Halford et al. have presented an elegant theory of how relational complexity can be modelled. Their view has impressive support from neural net applications. Their “core argument,” however, concerns the applicability of their theory to human cognition (sect. 1). This core argument must stand or fall on the basis of empirical evidence. My comments will focus on the postulated links between relational complexity and age differences in children’s understanding of concepts. The key issue is whether relational complexity is a useful metric for *explaining* cognitive development. I will show that, although Halford et al. point out that a capacity approach does not imply insurmountable barriers to performance, and although they make a number of caveats (which limit the testability of their theory), a survey of current research in developmental psychology does not fit neatly into their theoretical framework.

Halford et al.’s first claim is that in infancy representation is limited to objects in the immediate spatio-temporal frame. “There is no evidence that semantically interpretable relations are represented, however, or that inferences go beyond the perceptible properties of objects” (sect. 6.2.1). This is not the case. Recent infancy research has shown that some forms of relational reasoning are present in the first year. For example, Baillargeon has shown that 4- to 5-month-old infants are surprised by “impossible” physical events in situations where surprise must be based on inferences about cause-effect relations that occur outside the immediate spatio-temporal frame (e.g., Baillargeon 1994; Baillargeon & Graber 1987; Baillargeon et al. 1990). Infants can even impose an “intentional stance” onto simple perceptual events, making assumptions of agency on the basis of causal analyses of impoverished physical situations such as computer displays of moving circles (e.g., Gergely et al. 1995; see Goswami, 1998, for a review). Baillargeon’s conclusion that the infants in her experiments are engaging in a knowledge-based, conceptual analysis of the physical world is increasingly accepted in developmental psychology.

Halford et al.’s second claim is that the disappearance of well-documented phenomena in infancy such as the A-not-B search error can be explained by their theory. They suggest that the A-not-B error should be understood in terms of “an inability to treat hiding place as a variable” (sect. 6.2.2). They argue that perseverative searching at location A occurs because infants treat hiding place as a constant. Yet similar perseverative errors are made when

infants attempt to crawl around a barrier to reach their mothers (McKenzie & Bigelow 1986; Rieser et al. 1982). In a few trials, the same infants learn to correct these errors. This would not be predicted by Halford et al.'s theory, but it is easily explained within a cognitive-developmental framework that accounts for perseverative errors in terms of a difficulty in inhibiting prepotent motor tendencies. Recent analyses suggest that the A-not-B error reflects a deficit in gaining *control* of one's motor behaviour, allowing behaviour to *reflect* knowledge of the fact that hiding place is a variable (e.g., Diamond 1988).

Halford et al.'s third claim concerns children's supposed difficulty with ternary relations (sect. 6.2.4). Most space is devoted to this difficulty, which is said to explain children's failure in Piagetian tasks such as class inclusion and transitivity as well as their difficulties with appearance–reality tasks and theory of mind tasks. I have discussed various problems with this explanation of performance in Piagetian tasks elsewhere (Goswami 1996). The extension of relational complexity theory to appearance-reality and theory of mind tasks is novel (and ambitious). The example discussed by Halford et al. concerns children's apparent inability to understand that objects placed behind coloured filters only *appear* to change colour. Halford et al. argue that this difficulty is explained by children's inability to represent ternary relations. Although children can represent the relation between a person's knowledge and the properties of an object/situation, they cannot represent the fact that this is conditional on a third variable such as the presence or absence of a coloured filter.

Recent developmental research suggests that the problem does not lie with conditional understanding, but with the format of the traditional appearance-reality task. For example, Cutting (1996) tested children's understanding of the appearance-reality distinction indirectly, by using tasks in which 3-year-old children were shown (for example) two crayons, one green and one yellow. The yellow crayon was placed under a blue filter and the green crayon under a clear filter, so that both crayons appeared to be identical shades of green. The children were asked "Which crayon can I use to draw green grass?" Over 80% of the children answered correctly. Furthermore, following experience with natural examples of the appearance-reality distinction in their nursery schools (e.g., making sunglasses with coloured tissue paper), 3-year-olds improved in the traditional appearance-reality tasks as well. Thus an inability to represent ternary relations at 3 years cannot explain failure in appearance reality tasks.

Finally, Halford et al. provide a critique of some of my own data showing that ternary relations can be processed earlier than a median age of 5 years if familiar relations are used (sect. 6.2.4.4; see also Goswami et al. in press; Goswami 1996, for other evidence not discussed by Halford and Phillips). They argue that the impressive levels of performance found by Goswami (1995) can be explained in terms of binary relational mappings if certain assumptions about chunking strategies are made. Of course, the counterargument is that it is more parsimonious to assume that these levels of performance reflect successful ternary relational mappings, but such a debate has no resolution. This highlights the key problem with Halford et al.'s analyses of cognitive developmental research, which is that there is no *independent* test of whether a child is solving a given task on the basis of relational mappings at the postulated level of complexity. Although failure in a given task is assumed to prove an inability to use ternary relational mappings, there is no independent positive test of whether children *are* using ternary relational mappings.

This flaw places Halford et al.'s concluding arguments concerning cognitive development on very shaky ground. They note that "the fundamental problem here is that cognitive development research must take account of actual cognitive processes to be theoretically meaningful" (sect. 6.2.4.5). The present commentary takes the opposite view. Theories about potential cognitive processes are only meaningful if they take account of actual results in cognitive development research.

Why is capacity limited? Missing dynamics and developmental controversies

Richard A. Heath and Brett K. Hayes

Department of Psychology, University of Newcastle, Callaghan, NSW 2308, Australia. {heath,hayes}@psychology.newcastle.edu.au
psychology.newcastle.edu.au/~heath
psychology.newcastle.edu.au/~hayes

Abstract: The discovery of a quaternary complexity limitation to mature cognitive operations raises important theoretical issues. We propose that cognitive limitations arise naturally in a complex dynamic information processing system, and consider problems such as the distinction between parallel and serial processes and the representativeness of the empirical data base used by Halford et al. to support their relational complexity scheme.

Halford et al. provide a convincing argument that constraints on the simultaneous processing of relations generate developmental limitations in human cognition. Their preferred mathematical representation is tensor calculus, an extension of associative memory models (Humphreys et al. 1989). However, their target article is somewhat dismissive of alternative models based on synchronised oscillations and it is unclear whether either model can be evaluated experimentally, since no model fitting is attempted.

Cognitive dynamics, learning, and the parallel-serial processing distinction. Halford et al.'s model is descriptive rather than explanatory because there is no mechanism for explaining cognitive limitations. Furthermore, no explicit learning mechanisms are presented, there being a general disregard of the underlying cognitive dynamics. Complexity-theoretic research, based on cellular automata, suggests that optimal performance for an evolving cognitive system occurs when the number of parallel operations is small, perhaps four (Heath, in press). This result, arising naturally from dynamic considerations, is consistent with the representational limitations expressed in the target article.

An advantage of dynamic developmental models is the relatively stable stagelike behaviour emerging from this interaction. Such models, previously used in motor development (Thelen & Ulrich 1991), also explain developmental change in cognitive tasks involving memory, reasoning, and problem solving (Howe & Rabinowitz 1994; van der Mass & Molenaar 1992; van Geert 1994), the key focus of Halford et al.'s model.

Cognitive dynamics resolve the dilemma that cognitive developmental change within an individual appears abrupt, whereas "the acquisitions of concepts . . . will occur gradually after capacity becomes sufficient" (sect. 6.3.2). In van Geert's (1994) dynamic model, continuous parameter changes generate qualitative changes in cognitive behaviour as bifurcations, which allow children to progress from binary to ternary relationships, for example. Since the relationship between the number and magnitude of the transitions in an evolving fractal system is governed by a power law, only a small number of transitions of major cognitive significance will occur within a lifetime (van Geert 1994, p. 137). This alternative explanation of the resource limitations in cognitive development is appealing. Using van Geert's idea that cognitive change can be represented by connected growers with varying (carrying) capacity, Halford et al.'s representational scheme might be enhanced by using nonlinear dynamical concepts. For example van Leeuwen et al. (1997) showed how a coupled nonlinear system accounts for synchronised and hierarchically self-organised cognitive activity.

In much of Halford et al.'s discussion, learning processes play a minor role. The tensor idea can be extended to support an adaptive process in both space and time (Heath 1991) using an update rule similar to that employed in associative memory models. Performance constraints result from the interplay between the stability and plasticity of the memory update process, using parameters estimated from serial position data (Heath & Fullham 1988).

The definitions of capacity in various sections (sects. 2.1, 3.3)

depend critically on whether performance decrements are observed when the number of simultaneous tasks increases. Although indirect evidence of such limitations is obtained using a probe reaction time task (sect. 6.1.1), considerable ingenuity is required to dissociate parallel and serial information processing mechanisms experimentally (Townsend 1990).

The developmental data base. Halford et al. avoid any strong predictions regarding the ages at which each relational complexity level is attained. Nevertheless, some of their developmental claims ignore evidence indicating the attainment of higher levels of relational complexity at a much earlier point in development than that proposed by their model. Binary relations involving the relative position of two or more objects are encoded by infants as young as 4½ months of age (Needham & Baillargeon 1993) and are used to infer physical support relationships between these objects (Baillargeon et al. 1995). By 6 months, infants discriminate between causal and noncausal physical events, even when the spatio-temporal properties of these events are equivalent (Leslie & Keeble 1987). Even more problematic for the theory, by 18 months, infants can predict the intentions and motor behaviour of an adult who is observed solving a block construction task (Meltzoff 1996). This demonstrates the early development of an understanding that people can have at least two representations of an object, the current observable state and a future altered state. In Halford et al.'s terms this involves a ternary relation and should not be evident until 3–4 years. Taken together, such findings refute the claim that in infant cognition “there is no evidence that semantically interpretable relations are represented, however, or that inferences go beyond the perceptible properties of objects” (sect. 6.2.1).

More generally, it is disappointing that Halford et al. provide so little detail concerning the mechanisms producing age changes in relational representation complexity. They suggest (sects. 2.6.3, 4.2.4) that with age there is an increase in vectorial differentiation of representations. No specifics are given, however, about the processes triggering such differentiation during development. This represents a limitation of the model, because competing dynamic approaches (e.g., van der Maas & Molenaar 1992; van Geert 1994) make explicit the processes by which stable stagelike behaviour patterns are disrupted by organismic and/or environmental events and restabilise into strategies that enable the child to deal with more complex processing tasks.

Another curious omission is any reference to the considerable body of evidence for an exponential decline across the lifespan in the processing time in both simple visuo-spatial tasks (Kail 1991) and in more complex problems involving memory and analogical reasoning (Kail & Park 1992). We are skeptical that such age changes in processing speed can explain all of the cognitive developmental phenomena described by Halford et al., but this robust body of data deserves discussion in any relational complexity model.

Conclusion. Although Halford et al. have provided those interested in capacity limitations with a solid representational framework for their investigations, there is insufficient consideration of alternative viewpoints, a lack of appreciation of the role played by dynamical processes, and an underrepresentation of some critical empirical phenomena. Despite these reservations, the target article contains an interesting compendium of ideas to focus and advance our knowledge of cognitive developmental processes.

Is multi-tasking complex?

W. Bentley MacLeod

Department of Economics and The Law School, University of Southern California, Los Angeles, CA 90089-0253. wmacleod@usc.edu

Abstract: In a simple economic decision problem with multi-tasking the dimensionality of the problem is neither a necessary nor a sufficient measure of complexity. Rather, dimension is good measure of complexity when

there is an aggregate resource constraint that creates an interaction between the different activities, resulting in a problem with high algorithmic complexity.

Introduction. Halford et al. argue that the complexity of a task or problem depends on the dimensionality of the problem. As an economist, I shall approach this from the point of view of a manager deciding how to allocate resources for several tasks. A simple example shows that while dimensionality is both suggestive and useful, its relevance depends on the computational complexity of the problem, which in this example depends on whether there is an aggregate resource constraint. (See MacLeod, 1996, for a discussion of this model and further references to the economics literature addressing the implications of cognitive limitations for decision making.)

Consider a manager who must determine the extent of resources $y_i \geq 0$ to be allocated to activity $i \in \{1, \dots, N\}$. For simplicity, suppose that the total benefit from the manager's decision is given by:

$$B(y : \theta, \beta) = \sum_{i=1}^N \theta_i y_i - \beta_i y_i^2 / 2 - \delta(y_i) \tag{1}$$

where $\theta_i \in \{\bar{\theta}, \underline{\theta}\}$, $\beta_i \in \{\bar{\beta}, \underline{\beta}\}$ are random variables that can take the high or low value with equal probability ($\bar{\theta} > \underline{\theta} > 0$ and $\bar{\beta} > \underline{\beta} > 0$), while $\delta(y_i) = \begin{cases} 0, & \text{if } y_i = 0 \\ k, & \text{if not.} \end{cases}$ denotes a fixed cost if

positive resources are allocated to this activity. Let $\Omega = \{\bar{\theta}, \underline{\theta}\}^N \times \{\bar{\beta}, \underline{\beta}\}^N$ denote the space of possible parameter values. The parameters θ_i represents the benefit from allocating resources to this activity, while β_i is a measure of the costs. For example, the benefit might be saving a building that is on fire, while the costs are given by the danger involved. Both the benefit and the cost are likely to change from period to period, requiring the manager to determine a plan for each new situation. The manager's optimal response is to make a resource allocation decision as a function of the activities and is defined by:

$$R(\theta, \beta) = \arg \max_{y \in \mathbb{R}_+^N} B(y; \theta, \beta). \tag{2}$$

$$\arg \max_{x \in X} f(x) = \{x \in X \mid f(x) \geq f(x') \text{ for all } x' \in X\}.$$

One can put this model into the framework of the Halford et al.'s target article by defining the relationship:

$$r(y : u, \beta) = \begin{cases} 1, & \text{if } y = R(u, \beta), \\ 0, & \text{if not.} \end{cases} \tag{3}$$

Representation. For this simple model the state space defined by the parameters $\{\theta, \beta\} \in \Omega$ has size 2^{2N} , and hence grows very quickly with the dimension. However Halford et al. point out in the target article that dimension rather than size is the crucial issue. If one drops the size criteria, then exactly what do we mean by dimension? For example, one can embed Ω into the real line by letting

$$x = \sum_{i=0}^N \frac{(\theta_i - \underline{\theta})}{(\bar{\theta} - \underline{\theta})} 2^{2i} + \frac{(\beta_i - \underline{\beta})}{(\bar{\beta} - \underline{\beta})} 2^{2i+1} \tag{4}$$

and then define the optimal response by $\bar{R}(x) = R(\theta, \beta)$. In that case the response \bar{R} depends only upon a one dimensional real number x . However, the optimal response function \bar{R} is now highly nonlinear, and is itself very complex.

The functions in the target article are maps to responses from the space of words/sentences, hence from a finite collection. What is not clear is whether ternary relations are more complex than bi-

nary ones because of dimensionality or because ternary relations use longer strings as inputs? Halford et al. discuss this problem in section 3.4, and state that “it is implausible that appropriate encoding strategies are immediately available for novel cognitive tasks” (sect. 3.4.3). In the target article, the input strings are characters separated by commas, with the number of dimensions equal to the number of commas plus 1. What makes dimension important here does not seem to be dimension per se, but some aspect of the computational problem faced by the decision maker.

Interaction. In the simple resource allocation model, the optimal rule is defined by:

$$y_i^* = \begin{cases} u_i / \beta_i, & \text{if } u_i^2 / 2\beta_i \geq k \\ 0, & \text{if not.} \end{cases} \quad (5)$$

This is an illustration of the concept of segmentation discussed in the target article. In this case, each activity can be considered independently of the other. In particular, if C represents the cost of computing the appropriate response for this activity, then the total cost of computing the optimal allocation is $N \cdot C$. Thus, even though the size of the state space is increasing exponentially in N , the computational costs are linear in N .

Now suppose that the total amount of resources is finite, so that

$$\sum y_i \leq Y. \quad (6)$$

If we let λ be the Lagrange multiplier associated with this constraint, then those activities receiving the a positive level of resources have output given by:

$$y_i(\lambda) = (\theta_i - \lambda) / \beta_i. \quad (7)$$

However, due to the fixed cost of allocating a positive level of resources to an activity, for each λ one must determine which activities receive resources. The problem is made more difficult by the fact that due to variations in both θ and β , there is no unique way to order the productivity of the tasks as λ varies. There are a number of algorithms that one can use to solve this problem, however, all of them are significantly more complex than the case without the resource constraint because the resources allocated to a task depend on the parameters for the other tasks. One algorithm would iterate on λ , starting at zero, and search for the maximum value of B . For each λ one first computes the payoff:

$$B_i(\lambda) = \theta_i y_i(\lambda) - \beta_i y_i^2(\lambda) / 2 - \delta(y_i(\lambda)), \quad (8)$$

and sort the B_i 's from highest to lowest, which takes $O(n \log n)$ using the best technique. One would then select the m tasks such that $\sum_{j=1}^m y_{i(j)}(\lambda) \leq Y$, and $B_{i(j)}(\lambda) \geq 0$, where $i(j)$ is the index as a function of rank j . The maximum value for λ is θ , and given that the problem is nonconvex, the solution is found by search for the best payoff using this algorithm as λ varies on $[0, \theta]$. Thus in general the time complexity of this problem is at least a higher order polynomial function of n , the dimension.

This simple example illustrates the point that although these two problems have exactly the same representation and objectives, the addition of the aggregate resource constraint (6) makes the problem significantly more difficult. The problem of algorithmic complexity is discussed in section 5 of the target article. (In that section Halford et al. state incorrectly that NP-complete problems are of exponential complexity, whereas this is in fact an open question. It is suspected that they are of exponential complexity with the current upper bound given by $O(2^{p(n)})$, where $p(n)$ is a polynomial function of the size, n , of the problem.) Halford et al. point out that the use of parallel processing can solve the time complexity, but only at the expense of space complexity.

The simple economic example illustrates that thinking in terms of an overall aggregate budget constraint may be a useful way to think about the problem. If one supposes that each computation requires resources, any increases in real time speed via parallel

processing would require a corresponding increase in current resources. Hence parallel processing does reduce the resource constraint, but is simply a way to reallocate resources to be used in the future. In the simple economic example, parallel processing would correspond to computing the value of $B(\lambda)$ for many different values of λ simultaneously. What makes the economic model complex is not the number of ingredients but the way the resource constraint requires trading off decisions along several dimensions.

Finally, I had great difficulty in following the model of section 4.2, which seemed to hold the core of the argument, and one that I have attempted to explore with this stylized economic model. Essentially, the complexity of a problem instance depends not so much upon the dimension per se, but upon the need to relate several different elements, with the complexity rising with the number of interrelationships that need to be explored. I am sympathetic to the intuition presented in the target article, but would have been able to follow the argument more easily with a more detailed presentation of one of the mathematical models, along with some theorems on the algorithmic complexity of the proposed mechanisms.

Is it processing capacity that is being defined?

David Navon

The Laboratory for Perception and Attention, Department of Psychology, University of Haifa, Haifa 31905, Israel. dnavon@psy.haifa.ac.il

Abstract: Halford et al. are not redefining capacity in the sense of limit on resources but in the sense of limits on what resources can do. Furthermore, the necessity of using resource theory as a theoretical frame is questionable.

What is it that determines task demands on working memory? I share Halford et al.'s frustration with the traditional analysis in terms of the information metric. The alternative suggested by Halford et al. is not unreasonable, but I would like to comment on the level at which that was done.

It is one thing to point out that the relevant factor is the number of independent sources of variation that are related, or to tie that conclusion to the empirical observation that human adults cannot process without segmentation a relation that has more than four arguments. It is another thing to claim that processing capacity is defined by relational complexity. There appear to be two problems with that argument.

Capacity of output limit? Although the word “capacity” has been used by psychologists in various senses, it has acquired a more circumscribed sense in resource theory – as a scientific term denoting the limit on resources. Halford et al. accept this definition (sect. 2.1), but they do not seem to have assimilated it into their analysis.

The term “resources” is used in cognitive psychology to refer to provisions for, or internal inputs to, processing (e.g., Gopher 1986; Kahneman 1973; Navon 1984; Navon & Gopher 1979; 1980; Navon & Miller 1998; Norman & Bobrow 1975; Wickens 1984). Theorists typically regard resources as pools of multipurpose units, analogous to gallons of fuel, workers in a factory, or storage location. Such pools have been defined in resource theory as having several properties: aggregate nature, exclusive usage of units, distributability, effectiveness, and scarcity (see Navon 1985). In other words, each resource unit can potentially serve different users, can actually be allotted at any time to any given user independently of the allotment of any other unit, can serve only one user at a time, contributes to the output in a cumulative manner, and comes from a limited pool.

Within this context the term “capacity” has a definite meaning – the limit on resources (viz., the total number of units of internal

input that may be expended on mental work) rather than the limit on work (viz., output that those units of internal input can produce). By analogy, the limit on the resources of a shoe factory is better expressed in terms of, say, labor hours than in terms of shoes produced. Likewise, the capacity of a library is better expressed in terms of meters of shelf space than in terms of number of books stored. Actually measuring capacity would be a little hard to do in psychology unless we had a better idea of what the mental resources were. We might settle for characterizing work, namely *processing complexity*, in the most general way we can, and then specifying the limit in terms of work – that is, on *output limit* rather than a *processing capacity* limit.

This is often a formidable problem, and not only in psychology. How, for example, would one succinctly specify maximal work or maximal production in a shoe factory, if it produced all kinds of shoes? For the same reason, it would clearly be unsatisfactory to characterize the output limit of working memory in terms of, say, the maximum number of digits making up a mental arithmetic task, or the maximum number of items that can be recalled in a digit span task. This is probably why people were excited by the notion of chunks and hoped to gain much from the measure of information. Four decades after George Miller published his classic article (Miller 1956), we are somewhat disenchanted with chunks and bits. But the lesson could be a little more general: it may not be a good idea to develop too much hope with respect to *any* such measure of processing complexity.

This may also be true of measures of relational complexity. How many independent sources of variation are there in memory scan? What is the relational complexity of a digit span task, and how does it grow when the number of digits is increased from 3 to 7? What is the relational complexity of a dual-task situation, and how can it be related to the complexities of the component tasks? How can the confusability between the tasks (which certainly affects performance) be captured by a measure of relational complexity? I wish Halford et al. lots of luck in trying to answer these questions, but I suspect it would not be easier to find a common denominator for all the shoes produced in that mental factory called working memory than it would be to learn more about its laborers and about the extent to which their production depends on their number.

Is processing resource limited? And here is the second reason I would hesitate to use the term “processing capacity”: At present we do not know enough about those laborers. Empirical effects of load or task interference are too often described using a theory-laden jargon, even when the data do not afford evidence that selectively supports the theory. Typically the description is in terms of resources and capacity. If the concept of mental resources is not taken for granted – and there is some reason for not taking it for granted (cf. Navon 1984; 1985) – then the use of the terms “mental resources” and “processing capacity” – at least in conjunction – must imply that there is empirical evidence supporting some version of resource theory rather than any of its competitors (for one, models of outcome conflict, Navon & Miller 1987; 1998).

Is there a good reason for Halford et al. to do that? They present neural net models that account for limits on processing concepts with high relational complexity. The relevant question is then whether these models assume that processing is resource limited.

In general, it is doubtful that neural nets satisfy the defining properties of resources: (1) Neural net units evolve to permanently subserve the specific function(s) they do subserve. (2) The contribution of each unit is special in a way (e.g., its specific connections to other units, its vector of gains on emanating connections, possibly its transfer function). (3) The units operate in a concerted manner, thus contributing synergistically rather than cumulatively (more the way different body organs contribute to accomplishing a certain biological function than the way oxygen molecules do). (4) Although the overall number of units must be limited, the fact that they are function-dedicated prevents local scarcity, in the sense of a unit being temporarily unable to serve one function because at the time it serves another one. (5) It seems

more likely that two functions that are both served by some units will interfere with each other due to crosstalk, the extent of which depends on the overlap in units.

Even if the use of the notion of resources in the context of neural nets were appropriate, however, this would fall short of convincing me that processing is resource limited. It is not clear that we have to take neural nets for granted any more than we do resource theory. Whether human information processing is best captured by neural net modelling is still an open question. One could imagine an information processing system in which processing provisions were never scarce but outcome conflict was prevalent. What would the term “capacity” denote then?

Chaotic dimensionality of hand movements define processing capacity by relational complexity

Danko Nikolic

Department of Psychology, University of Oklahoma, Norman, OK 73019.

dnikolic@ou.edu

www.ou.edu/class/psy2503/danko/danko.html

Abstract: Measurements of the dimensionality of chaotic attractors obtained on behavioral data represent the task complexity and also could be hypothesized to reflect the number of synchronized neural groups involved in the generation of the data. The changes in dimensionality for different experimental conditions suggest that limited processing capacity, task complexity, demand, and synchrony in neural firing might be closely related.

Results on chaotic dimensionality of hand movements support processing capacity definition by relational complexity.

The chaotic dimensionality of a signal is determined by a method that includes two steps: (1) reconstructing the chaotic attractor or phase space, and (2) determining the dimensionality of the reconstructed attractor. Such an analysis provides us with the “dimension” of a data set. The dimension provides information about the minimal number of first-order differential equations (or active degrees of freedom) that are necessary to generate the signal, and therefore represents a measure of the complexity of the process that generated that data. In addition, it will be assumed that the number of dimensions is a rough estimate of the number of synchronous neural groups involved in the process. In other words, each differential equation that generates the motor movements corresponds to one synchronous neural group. The larger the complexity of the motor movements, the larger the number of synchronous neural groups involved. Any change in dimensionality can therefore be interpreted as a change in the number of synchronous neural groups used to carry out the task. With these two interpretations of the attractor’s dimension in mind (complexity and number of synchronized neural groups), the results obtained on repetitive hand movements provide support for several fundamental assumptions underlying Halford, Wilson & Phillips’s theory.

First, it has been found that repetitive movements (i.e., swinging a stick) actually result from a low-dimensional chaotic attractor (Mitra et al. 1997) whose dimensionality does not seem to be larger than 5, a range similar to that proposed by Halford et al. and estimated by others (Hummel & Holyoak 1997; Lebedev 1980; Shastri & Ajjanagadde 1993a). This finding suggests that the small variations in hand movements are not noise but come from a deterministic system. Therefore, any change in the dimensionality might reflect a change in the complexity of the processes underlying the hand movements.

When participants were asked to produce a simple repetitive swinging-like movement with either a heavy or light stick, the swinging of the heavy stick was judged to be a much easier task

(Riley 1997). Mitra et al. (1997) found that more difficult tasks (the lighter stick) have larger dimensionality ($d_L = 4$) than the easier task ($d_L = 3$). This finding suggests that a more difficult task involves more complex processing and probably a larger number of synchronized neural groups. This finding seems to be in line with Halford et al.'s assumption that the processing effort (demand) and processing complexity are related (where complexity is defined as the number of variables interacting in parallel that were attributed to the synchronous neural groups).

Another finding supports Halford et al.'s proposal that one way to overcome the large dimensionality/complexity of the task is through a form of learning called "conceptual chunking." Such learning results in a decrease in the complexity of a concept's representation, and consequently in involvement of a smaller number of synchronized groups in processing. An experiment by Mitra et al. (1998) somewhat supports this assumption. They had participants practice a relatively difficult swinging task that showed performance improvement over time. The improvement in performance was accompanied by a decrease in the complexity of the movements. Although there was not an obvious decrease in the number of low dimensional active degrees of freedom (the method did not allow for measurement of dimensionality more precise than an integer), there was a substantial decrease in the high dimensional component (or "noise"). Although the high dimensional chaos was abundant at the beginning of training, it nearly vanished at the end. This is an empirical example of how the complexity of behavior decreases with practice/learning. Furthermore, the more speculative interpretation suggests that the number of synchronized neural groups involved in the task might have decreased as well.

These results concur with Halford et al.'s hypothesis that the limited processing capacity (maximal number of variables), the task complexity (requisite number of variables), and the difficulty of the task (effort), are closely related. Finally, if the obtained number of active degrees of freedom is interpreted as the number of synchronous neural groups involved, then their hypothesis about the mechanism underlying limited processing capacity would also be supported (see Nikolic, in press, for similar conclusions). Hence, the decrease in processing complexity, possibly caused by the decrease in the number of synchronous neural groups necessary to carry out the task, might be the mechanism by which the brain copes with the complexity of computational tasks. This would mean that through learning and practice the neural network restructures in such a way that the processing of the task requires fewer synchronous neural groups. As a consequence, part of the limited processing capacity is free for processing of additional information (Nikolic, in press). The same mechanism might underlie both simple skill acquisition and higher cognitive processes.

To appraise developmental difficulty or mental demand, relational complexity is not enough

Juan Pascual-Leone

Department of Psychology, York University, Toronto, Ontario M3J 1P3, Canada. juanpl@yorku.ca

Abstract: Two assertions of Halford et al. are critiqued: their claim of priority in relational complexity analysis and the sufficiency for cognitive development of their relational-complexity analysis of tasks. Critical discussion of concrete task analyses (i.e., the relational complexity of proportionality problems, of balance scale problems, and the Tower of Hanoi) serves, by way of counterexamples, to highlight problems in their method.

Halford et al. make three general claims in this important target article. The first concerns the adequacy of their tensor-product

method for computational modelling of cognitive developmental processes. Elsewhere (Pascual-Leone 1994) I have discussed the psychological limitations of this method, although I recognize its heuristic interest. The second claim concerns the sufficiency of the Halford et al.'s developmental method of task analysis, that is, evaluation of a task's mental demand via objective (ideal observer's) estimations of its relational complexity. The third claim is that authors are the first to propose the method of relational complexity analysis in developmental psychology. Not so: Nassefat (1963) and, independently, de Ribaupierre & Pascual-Leone (1979); Pascual-Leone (1970); Pascual-Leone & Goodman (1979); Pascual-Leone & Smith (1969) were the first developmental researchers to quantify complexity by means of a form of dimensional analysis: the highest number of schemes – task relevant dimensions of variation – that subjects must consider simultaneously to solve the task. In their effort to formalize the method, Halford et al. have made a number of oversimplifications, which hinder their ability to assess task complexity in mental processing – what I call *mental demand*. This commentary illustrates some of these problems.

Here is the main problem: Halford et al. believe (sect. 2.3) that their "objective" relational complexity estimate equals mental demand, so they assert that a human adult can process, within one mental act, at most a four-argument relation. Strict tests of this claim require paradigms in which, within the same task and using the same executive processes, one has multiple homogeneous classes (scales) of items, with each class varying the relational complexity of items by one unit (1, 2, 3, 4, . . . , 7, 8), relative to the previous class. Then, with items randomly ordered within the test, normal subjects of various ages would be examined. With this design, Halford's model makes a clear prediction: at 11 years of age (when processing a relation with 4 arguments becomes possible) the performance level of subjects will reach an asymptote – that is, older subjects should not perform any better. This is the paradigm of M-capacity measurement that I initiated (Pascual-Leone 1970; Pascual-Leone & Baillargeon 1994). Using this paradigm, we have found repeatedly that the asymptote predicted by Halford's model does not materialize (e.g., Pascual-Leone & Goodman 1979). Our results are consistent with developmental findings about fluid intelligence and Piagetian formal operations (de Ribaupierre & Pascual-Leone 1979), which do not reach asymptote until late adolescence.

Halford's error reflects a basic flaw in this task analytical methods: confusion between an *observer's* perspective (i.e., objective analysis from without) and an *organismic* perspective looking at processes from within (Pascual-Leone 1978; 1994). Halford et al.'s own examples illustrate this error. According to them (sects. 2.3 and 3.6), a proportionality relation such as $4/12 = 32/96$, or $a/b = c/d$, has the relational complexity of 4: P(a, b, c, d). This cannot be so, however, if by complexity they mean "effective complexity" as defined by this new science. Gell-Mann (1994, p. 96), one of its architects, defines the "effective complexity of an entity, relative to a complex adaptive system that is observing it and constructing a schema, as the length of a concise description of the entity's regularities identified in the schema." A scheme or schema for proportionality *must* contain all essential, internal and external functional relations (all the regularities) that proportionality prescribes for terms in the argument set (a, b, c, d). Thus the two ratios (a/b and c/d) whose equality makes the proportion valid are part of the complexity count. After reducing redundancy, the relational complexity in proportionality should be equal to 4 (number of arguments) *plus* 2 (one for the proportionality operator P and another for at least one of the two internal ratio-functors f). In principle, the relational complexity of proportionality is equivalent to 6 units, as shown by its mental-structure formula: $P(f_1^L(a,b), \{f_2\}_L(c,d))$. In this formula the brackets around the functor f_2 should be interpreted psychologically as meaning that f_2 is part of the same structural chunk L as f_1 , and so does not increase the attentional demand. Surprisingly, Halford et al. (sect. 2.3) assign to proportionality the structure P(a,b,c,d), with a complexity count of 4; they also

give the same structural complexity to the much simpler propositional structure (sect. 3.2) “played(John, cricket, oval, Sunday)” – John played cricket at the Oval on Sunday – which, unlike proportionality, requires no internal functors.

This confusion of two very different relational structures shows a fault in Halford et al.’s task-analyzing method. But this is just one counterexample; there are several others. The Piagetian balance-scale task (sect. 6.2.5.1) is another one: the authors adopt a simple structure like P(a,b,c,d), that is, Balance (W_1, D_1, W_2, D_2), when the internal functional relations that tie together weights and distances within each arm demand the complex structure just discussed: Balance($f_1^L(W_1, D_1), f_2^R(W_2, D_2)$), with a mental demand of 6. Pascual-Leone (1994) has discussed Halford’s problems with this task in detail. De Ribaupierre and Pascual-Leone (1979) give detailed organismic task analyses of the balance showing other factors. They also show, theoretically and empirically, that most others Piagetian formal-operational tasks have a mental demand greater than 4.

Halford et al.’s analyses of the Tower of Hanoi also have methodological problems (sect. 6.1.3; Table 1), for example the 4-disc game (see Table 1). In Move 1, goal sequence 4C 3B 1C 2C 1B (notice that in Table 1, goal 1C has been omitted!) has for the authors a dimensionality of 10 (in Table 1 they write 8 because goal 1C is missing) – well beyond the maximum complexity of 4 prescribed by their model. They justify this frequently found anomaly – that is, tasks being solvable even when their analysis predicts excessive dimensionality – by claiming that chunking and segmentation reduce (de Ribaupierre & Pascual-Leone 1979) complexity. But since the theory has no explicit rules constraining their occurrence, chunking and segmentation turn into theoretical loopholes for explaining away empirical anomalies. Pascual-Leone and Goodman (1979) show how rules for constraining the occurrence of chunking (L learning) in task analysis can be formulated.

Chunks, bindings, STAR, and Holographic Reduced Representations

Tony A. Plate

School of Mathematical and Computing Sciences, Victoria University, Wellington 6001, New Zealand. tony.plate@vuw.ac.nz
www.mcs.vuw.ac.nz/~tap

Abstract: Much of Halford et al.’s discussion of vector models for representing relations concerns the perceived inadequacies of alternative methods with respect to chunking, binding, systematicity, and resource requirements. Vector-based models for storing relations are in their infancy, however, and the relative merits of different schemes are not so clearly in favor of their STAR scheme as Halford et al. portray.

Halford, Wilson & Phillips make three strong claims about the processing and representation of relational knowledge in the human brain and in neural network models. First, information processing capacity is more restricted by the complexity of relations that can be processed in parallel than by capacity for short-term storage. Second, in vector-based models, relations are better represented as predicate-argument bindings than role-filler bindings. Third, these bindings are best computed as tensor products. Halford et al.’s goal of producing a vector-based neural-network model of relational processing that incorporates capacity limits is thoroughly worthwhile and the model they propose has many interesting properties. Apart from a slight lack of clarity about the relationship between tensor-product representations for relations and such diverse phenomena as parsing center-embedded sentences and solving word puzzles, my main disagreements are with the second and third claims, which lie at increasingly deep levels of implementation detail.

In Halford et al.’s story there is an unwavering path leading from the first through the second to the third claim, constrained by the

inadequacies of alternative methods with respect to chunking, binding, systematicity, and resource requirements. However, vector-based models for storing relations are in their infancy, and the relative merits of different schemes are not so clearly in favor of their STAR scheme as Halford et al. portray. I will formulate my discussion as a defense of holographic reduced representations (HRRs) (Plate 1994; 1995) against Halford et al.’s charges, although in many cases the defenses also apply to related methods such as Smolensky’s (1990) role-filler tensor product representations and Kanerva’s (1996) binary spattercodes.

Chunks. Chunking is one of the most important issues in any model that aspires to process hierarchical relational structure. Such a model must be able to form modest-sized chunks of information that can be presented to its processing machinery. In symbolic architectures “cons” cells or records can be viewed as chunks, and pointers serve to record relationships among chunks. In neural networks, chunking is less straightforward because there are no addresses to point to. Neural network chunking techniques are potentially more powerful, however, because a chunk can give some indication of its contents,¹ as Hinton (1990) suggested in his proposal for “reduced representations.”

Chunking is intimately bound up with HRRs. Contrary to Halford et al.’s claim in section 4.2, HRRs easily handle conceptual chunking and satisfy the three general principles of chunking listed in section 3.4.1. Every vector in an HRR model is already a chunk and no further compression (as sought by Halford et al. in sect. 4.2.4) is necessary. For example, the HRR $L1 = \text{love} + \text{loveagt} * \text{john} + \text{loveobj} * \text{mary}$ represents the relational instance loves (John, Mary) and is also a chunk that can be used as a filler in other relational instances such as knows(Sam, loves(John, Mary)) represented by $L2 = \text{know} + \text{knowagt} * \text{Sam} + \text{knowobj} * L1$. An important property of the HRR scheme is that there is only one unitary representation for a relational instance (as opposed to the nonunitary representation, which is the collection of representations for the components), which serves both as a chunk and records relational structure. In contrast, in the STAR model there are two possible unitary representations for a relational instance: the tensor product version and the chunked version.² This is not a fatal flaw, but it does seem to be a case of multiplying entities without necessity.

The claim in section 4.2.4 that HRRs rely on randomly generated component vectors is not exactly correct: the component vectors merely have to have elements whose statistics are those of a Gaussian distribution, an important difference that allows similar objects to be represented by similar vectors. Chunked representations are easily made similar to their components by adding the components directly into the chunk, as in $L1' = \text{love} + \text{john} + \text{mary} + \text{loveagt} * \text{john} + \text{loveobj} * \text{mary}$.

Systematicity. Halford et al. point out that HRRs do not provide a complete solution to systematicity (sect. 4.2.9). This is neither surprising nor a significant drawback – contextualized representations³ were only intended to support an estimation of the degree of systematicity present. In general, detecting whether systematicity is possible is equivalent to solving the graph-isomorphism problem, for which no polynomial-time algorithm is known (Garey & Johnson 1979). Thus, it would be very surprising if any polynomial time algorithm operating on fixed width vectors (such the dot-product of HRRs, or some tensor operation in STAR) were able to detect systematicity in general. It is a bonus that the degree of systematicity between small structures represented in contextualized HRRs can be estimated using a vector dot-product. As regards the general case, even uncontextualized HRRs do represent the structural information required to make a decision about whether systematicity is present. What is missing is a computational process that can act on that information. Given the computational complexity of detecting systematicity, this process must involve sequential computations in addition to parallel vector operations.

Role-filler and predicate-argument bindings. Halford et al. are right in claiming that role-filler binding schemes are susceptible to confusion if vectors representing relations are added in memory. However, this is a minor problem. There is no compelling

reason why relation vectors must be added. This allows a very simple solution: keep relation vectors separate in an associative content-addressable memory. There is no shortage of associative-memory models which can store a few or many items, for example, the memories described in Baum et al. (1988) and in Kanerva (1988).

There is considerable room for discussion of the computational and psychological implications of different types of bindings. One consequence of predicate-argument bindings is that there is no explicit representation of roles, which makes it difficult to represent similarity between roles of different relations, as we might want to do for the relations (X buys Y from Z) and (U sells V to W). Another consequence is that roles must be known at storage and recall time, which may not always occur. Furthermore, access based on one missing argument requires a different computation depending on which argument is missing, which seems cumbersome. In addition, recall based on more than one missing argument is difficult. Finally, dealing with relations having optional or symmetric roles also seems troublesome. The two schemes result in quite different patterns of similarity among relations: the similarity of two relations with role-filler bindings is the average of the similarities of their arguments, but with predicate-argument bindings it is the product of the similarities of their arguments. The two schemes are not mutually exclusive – it would be possible to use both together and thus combine their respective advantages.

A role-filler binding scheme could lead to resource bounds similar to the ones that Halford et al. identify. If looking up a relation in content-addressable memory is what is involved in making a single decision, then the number of independent pieces of information that could influence a decision would be the same as the size of chunks, that is, the number of role-filler bindings that can be reliably superimposed in the same vector. This explanation leads to a soft bound on processing capacity rather than the hard bound implied by predicate-argument binding. This capacity would also depend on the similarity of the particular pieces of information involved – one can reliably superimpose more dissimilar items than similar ones.

The issues regarding the relative merits of role-filler versus predicate-arguments bindings are largely orthogonal to those regarding the vector operations for implementing the bindings: either type of binding can be implemented with different vector operations (Halford et al. discuss only tensor products and convolution but there are other possibilities; see Plate [1994] and Kanerva [1996]).

Resource requirements. Halford et al. cite the requirement of a clean-up memory as a major computational cost unique to HRRs. However, the space complexity of the clean-up memory is linear in the number of entities and relations stored, which is hardly unreasonable. Suitable fast auto-associative memories are easily implemented in neural architectures. The simplest scheme uses a “grandmother cell” to store each item (Baum et al. 1988): one item is encoding in the incoming and outgoing weights of a cell and retrieval involves a winner-take-all competition over the grandmother cells. In any case, a realistic implementation of the STAR model will almost certainly need to be equipped with a clean-up memory to clean up the results of processing nonorthogonal vectors for items and chunks. Realistic implementations demand nonorthogonal vectors for two reasons. The first is based on simple capacity considerations: denser, nonorthogonal representations must be used if thousands of different possible items (say of the order of number of nouns in the English language) are to be represented. It is unreasonable to suppose that a rank-4 tensor might be formed over vectors 10,000 elements wide, for it would require 10^{16} bits of storage. The second reason is that representing similarity among items (e.g., that John is more similar to Bob than Flying-Saucer) requires nonorthogonal vectors and is essential for reasoning and generalization.

Halford et al.’s “worst case” situation for a clean-up memory where the majority of instances in a relation are stored (sect. 5.2.3) will not occur when there are even a moderate number of items.

Consider populating the space of rank 4 tensors where there are 1,000 possibilities for the predicate name and each of the three arguments. Even after storing one instance per second for 100 years, only 3.15% of the relational space would be filled (and this is a gross overestimate – Simon [1974] quotes an estimate of 5 to 10 seconds to store a chunk in long-term memory). In fact, for storing sparse relations a HRR scheme with a clean-up memory is far more economical than a STAR scheme because the HRR scheme only requires resources proportional to the actual number of relational instances stored, whereas the STAR scheme requires resources proportional to the total number of possible relational instances.

NOTES:

1. A pointer usually gives no indication of what it points to; it must be dereferenced to find out anything about the contents. A simple example of the idea of having a pointer say something about its contents can be seen in some LISP architectures where a pointer indicates the type of the object, integer, cons-cell, and so on, that it points to.

2. Halford et al.’s two proposals for implementing chunking appear to be representationally inadequate. The first in which items are chunked by concatenating vectors (illustrated in Fig. 1E), can result in confused argument bindings: the result of superimposing the tensors for the two unary relations $p(a/b)$ and $p(d/e)$ (i.e., $p \times c(a,b) + p \times c(d,e)$, where $c(x,y)$, a function which concatenates the vectors x and y end-to-end) is indistinguishable from the superposition of the vectors for the relations $p(a/e)$ and $p(d/b)$ ($p \times c(a,e) + p \times c(d,b)$). The second proposal, in which items are chunked by convolving vectors, loses the order of components within the chunk: $p(a/b)$ has the same tensor product as $p(b/a)$. This is because convolution is commutative: $a*b = b*a$. Both these objections to Halford et al.’s proposals are minor because there are other operations which have the required properties, for example, the noncommutative variants of convolution discussed in Plate (1995).

3. Halford et al. describe contextualization incorrectly: what contextual features are added depends on the situation, so in one situation features for smiling might be added to the representation for Jane, but in another features for frowning might be added. This does not change the general representation of Jane, only the representation in a particular situation. One of the motivations for the particular implementation of contextualization was that over the long term, if Jane was consistently frowning, the frowning feature might become part of her representation. However, this would not prevent the smiling feature being added in some particular situation.

Can we measure working memory without contamination from knowledge held in long-term memory?

John Sweller

School of Education Studies, University of New South Wales, Sydney 2052, Australia. j.sweller@unsw.edu.au

Abstract: The metric devised by Halford, Wilson & Phillips may have considerable potential in distinguishing between the working memory demands of different tasks but may be less effective in distinguishing working memory capacity between individuals. Despite the strengths of the metric, determining whether an effect is caused by relational complexity or by differential levels of expertise is currently problematic.

Working memory measures can be used to distinguish between both (a), tasks that impose different working memory demands and (b), individuals or classes of individuals who differ in working memory capacity. I will discuss these two aspects of Halford, Wilson & Phillips’s important target article.

Measuring task demands. Halford et al.’s metric provides a promising a priori measure of cognitive load imposed by various tasks. My research group has used a similar, though less detailed and sophisticated, structure (e.g., Tindall-Ford et al. 1997), but we do run into a common problem that is of general theoretical and practical interest: unless one knows for a particular individual

what relational dimensions have been chunked (what interacting elements have been formed into schemas), one does not know what relational dimensions one is dealing with.

Consider someone solving distance/rate/time problems of the sort discussed in section 3.4.1. Halford et al. indicate that this relation is three dimensional. In fact, there are more than three dimensions for novices who must define distance, rate and time before dealing with the problem and, as Halford et al. point out, fewer than three dimensions for experts. What constitutes a dimension will depend on the level of expertise of the person who is solving the problem, making it difficult to determine problem dimensionality.

How does expertise develop on these problems? We know from the work on chess (e.g., Chase & Simon 1973) that chess experts do not generate moves based on the rules of chess but rather learn to recognise tens of thousands of configurations (by chunking the pieces that make up each configuration) and learning the best moves associated with each configuration. Similarly, expertise in learning to use equations such as $V = s/t$ develops by people learning to recognise that for a particular combination of givens and goal (e.g., V and t known, find s), a particular procedure needs to be followed and that procedure changes when the combination of givens and goal changes (e.g., Sweller et al. 1983). The givens, goal and solution, are chunked into a single entity but separated (segmented?) from other entities with a different combination of givens and goal.

First, I am not sure that any relations become inaccessible by this procedure (sect. 3.4.1), but more important, I am not sure how to measure relational complexity. I think that on this analysis expertise consists of the development of large numbers of unary, or perhaps binary, relations. If so, that would fit with an increased working memory capacity available to experts. It would nevertheless leave us with a major empirical problem of having to determine the precise level of expertise before we could assign a level of relational complexity to a problem.

An indication of the difficulty of relating logical problem solving structures to working memory demands comes from Kotovsky et al. (1985) who showed that Tower of Hanoi isomorphs (all of which have an identical logical structure) can vary in difficulty by a factor of 16 depending on the cover story. The differences are due to automated knowledge held in long-term memory. Isomorphs that tap into previous experience are very easy whereas those that do not can be inordinately difficult. All have the same relational complexity.

Similarly, whereas the days of the week problem described in section 2 of Halford et al. is difficult, the isomorphic problem: *Suppose adding 5 to a particular number and then subtracting 2, gives you 10. What is 1 more than the original number?*, is substantially easier ($x + 5 - 2 = 10$, where x is the original number). We have highly automated schemas for dealing with mathematical entities. We have not needed to learn such procedures for days of the week.

It is reasonable to assume that relational complexity affects task difficulty in the manner outlined by Halford et al., but considerable work may be needed to determine what constitutes a relation for a given combination of learners and materials and to distinguish between relational complexity and other effects. Of course, the difficulty of dealing with the effects on working memory of schemas or chunks held in long-term memory is hardly unique to relational complexity theory. Halford et al.'s work has increased the sophistication of the metric available to us to measure differences in task demands; as such, it is potentially very valuable.

Measuring individual differences. Halford et al. suggest that young children's working memory capacity increases with age and that relational complexity theory can be used to measure the changes (sects. 6.2 and 6.3). I have doubts. Because expertise is an individual-difference factor, it is a much more serious contaminant when considering other individual difference factors such as working memory size than when distinguishing between the

memory demands of tasks. I doubt there is any current reliable procedure that can disentangle the effects of schemas and chunks held in long-term memory from an individual's working-memory size.

How do we know that older children have not acquired more automated knowledge than younger children for any experimental materials used and that all differences are due solely to this factor? If, as is surely the case, children of different ages differ in the extent to which they are familiar with names such as John, Tom, and Peter, and with the concept of ordering three people according to size, should they not differ in their ability to deal with the logic of size relations? Age differences may be due entirely to differential familiarity with the entities and procedures used rather than to differential ability to deal with relational complexity. In short, how can our experimental design isolate logical structures and their consequences for working memory from knowledge held in long-term memory?

Relational complexity, the central executive, and prefrontal cortex

James A. Waltz, Barbara J. Knowlton, and Keith J. Holyoak

Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095-1563. waltz@psych.ucla.edu
knowlton@lifesci.ucla.edu holyoak@lifesci.ucla.edu

Abstract: Halford et al.'s analysis of relational complexity provides a possible framework for characterizing the symbolic functions of the prefrontal cortex. Studies of prefrontal patients have revealed that their performance is selectively impaired on tasks that require integration of two binary relations (i.e., tasks that Halford et al.'s analysis would identify as three-dimensional). Analyses of relational complexity show promise of helping to understand the neural substrate of thinking.

Halford et al. have provided an extremely useful analysis of the cognitive and computational requirements for the representation and manipulation of explicit relational knowledge. Although numerous studies have established a role for working memory in reasoning, the precise actions of components of working memory in higher-level cognition have not been specified. One of the benefits of developing a theoretical taxonomy of relational complexity is that it could clarify the functions of the central executive component of working memory in humans and other primates. Working memory (Baddeley 1992) involves both the active maintenance of information and its processing, with the executive being responsible for control of processing. A wide range of evidence implicates the dorsolateral prefrontal cortex as a critical part of the neural substrate of working memory (e.g., Cohen et al. 1997; D'Esposito et al. 1996; Shallice & Burgess 1991). In humans, lesions to dorsolateral prefrontal cortex impair performance on a wide variety of tasks that have been identified with executive processing and fluid intelligence, including memory monitoring, management of dual tasks, rule application, and planning sequences of moves in problem solving. Yet, although much has been learned about what tasks require executive processing and what brain structures subservise working memory, it has been far from clear how executive processing should be characterized. What exactly is the kind of "work" performed by working memory?

One possibility, as Halford et al. note, is that executive processing involves the formation and active manipulation of relational knowledge. Based on a review of the human and other primate literature on frontal functions, Robin and Holyoak (1995) argued that the major functions of the prefrontal cortex can be understood as aspects of an overall system for reasoning with and learning about explicit relational concepts. In particular, the prefrontal cortex may be required for dynamic binding of fillers to relational roles. A theory such as the one outlined by Halford et al., which

defines cognitive complexity in terms of the number and types of mental elements that must be simultaneously manipulated in cognitive processing, provides a framework for characterizing the work done by those components of working memory that involve prefrontal cortex. Executive control will be required in complex planning and problem solving because these intrinsically involve high levels of relational complexity. As Halford et al. note in their analysis of the Tower of Hanoi problem, complexity measures based on levels of embedding in a goal hierarchy can be subsumed as special cases of variation in relational complexity. Contextual elements that need to be considered in combination with other elements will increase relational complexity, as will conditional contingencies based on multiple elements, especially when these are difficult to chunk due to their temporal or spatial separation. Abstract concepts will also have relatively high dimensionality, and hence be more dependent on prefrontal functions. Prefrontal cortex will be critical in dynamically binding fillers to roles. Within any mammalian species, tasks requiring the highest complexity level attainable by the species will be maximally dependent on frontal functions. Once complex relations are chunked and stored in long-term memory, and appropriate task segmentations and specific role-filler bindings have been learned, frontal involvement will be sharply reduced.

Halford et al. point out the potential for testing these claims using neuroimaging techniques coupled with tasks in which relational complexity is systematically varied. We have applied similar logic in recent neuropsychological studies comparing patients with widespread prefrontal damage to patients with nonfrontal lesions (in anterior temporal cortex) and to normal controls. One reasoning task we have used (Waltz et al., in press) is transitive inference, which Halford et al. discuss in detail. The problems we presented to subjects involved between two and four propositions, each stating a “taller than” relation between two individuals; the subjects’ task was to arrange cards corresponding to the individuals in descending order of their heights. In the less complex version, the pairs introduced the names in order of height (e.g., Bob taller than Fred; Fred taller than Sam). The correct ordering could therefore be achieved using a chaining strategy that proceeds one link at a time: to build a link, only one relation – that between the name currently at the end of the chain and its successor – need be considered. In terms of Halford et al.’s analysis, this order-restricted version of transitive inference can be performed by considering one binary relation at a time, and hence is two dimensional. In the more complex version of the task, the pairs were introduced in a scrambled order (e.g., Fred taller than Sam; Bob taller than Fred) so that the item currently at the end of the chain was not in the subsequent pair, making the chaining strategy inapplicable. The subject would therefore need to consider two binary relations simultaneously to determine the overall ordering, making the task three-dimensional according to Halford et al.’s analysis (i.e., the complexity of unrestricted transitive inference).

Our results provided striking support for the hypothesis that prefrontal cortex is necessary for complex relational reasoning. On the simpler two-dimensional version of transitive inference, all groups achieved accuracy levels above 80% and did not differ statistically from one another. The prefrontal patients were able to solve transitive inference problems when as many as five names had to be ordered, as long as the names were introduced in order of height. But on the more complex three-dimensional version of the task, the performance of the prefrontal patients was catastrophically impaired: their accuracy was reduced to chance level, whereas temporal patients and normal controls continued to achieve over 80% correct. Importantly, the simple and complex problems were equated in terms of number of propositions and number of names to be ordered. In keeping with Halford et al.’s analysis and Robin and Holyoak’s (1995) hypothesis, the relevant metric for predicting sensitivity to prefrontal damage appears to be number of independent dimensions of variation that

have to be considered simultaneously, rather than a simple count of the number of problem elements.

We obtained very similar results when we manipulated complexity in variants of the Raven’s matrices task, with problems similar to geometric analogies (Waltz et al., in press). As in the case of transitive inference, the performance of prefrontal patients was selectively and catastrophically impaired on a version of the matrix task requiring integration of two binary relations (i.e., a three-dimensional version). These deficits on the more complex forms of the transitive-inference and matrix tasks cannot be attributed to some overall difficulty factor, as the frontal patients outperformed the temporal patients on a number of tests that required semantic or episodic memory. Such double dissociations indicate that frontal impairment is selectively found on tasks requiring dynamic binding of at least two binary relations in working memory.

Our findings support the hypothesis that prefrontal cortex is necessary for managing high levels of relational complexity, which for adult humans corresponds to what Halford et al. consider to be three-dimensional (or higher) representations. Although the neuropsychological evidence lends support to their computational-level analysis of relational complexity, it does not decide between alternative algorithmic accounts of how relations might be represented in a neural code, such as tensor products or neural synchrony (but see Holyoak & Hummel, in press, and Hummel & Holyoak, 1997, for computational arguments in favor of the latter approach). It should, however, be possible to use current techniques of cognitive neuroscience to address issues such as testing alternative metrics of complexity, teasing apart domain-specific and domain-general aspects of relational representations, and distinguishing the neural substrates of relational versus nonrelational knowledge. It will be fruitful to identify additional tasks that allow relational complexity to be systematically varied, and to use both neuropsychological and neuroimaging techniques to determine their neural basis. Whether Halford et al.’s metric of complexity is ultimately supported by behavioral and neural evidence, it is this kind of theoretical work that is needed to guide empirical investigations of the brain basis for thinking.

And if the developmental data doesn’t quite fit . . .

Barlow Wright

Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, England. barlow.wright@psy.ox.ac.uk

Abstract: Halford et al. seek to provide a framework that unifies distinct developmental phenomena. However, in pursuit of this goal, they sidestep crucial aspects of some well-known developmental benchmarks (most notably Transitive Inference and Object Concept), and they do not acknowledge “repeated” or “direct” experience as possibly being more fundamental than relational complexity, instead, ascribe all experience a secondary role.

Problems with crucial developmental phenomena. Halford et al. claim (sect. 6.3.2) that if we were to find the chronological frame of the emergence of some specific ability to be different from that given in the theory, this would not invalidate the theory, which would simply need to be realigned to the new age. This claim is not necessarily correct. Realignment would be difficult to justify if both ages were shown to clearly index distinct phenomena, as there would be no concrete reason to choose one age or another. This scenario is a reality in respect to Transitive Inference and Object Concept, respectively.

Transitive inference reoccurs throughout the target article (e.g., sects. 2.2.8, 3.6, 5., 6.1.1, 6.2.4.1, and 6.3.2). The floor in the age at which transitive inference first emerges appears to be 4 years,

with 3-year-olds incapable of solving such tasks (Halford 1984). When we ensure that children need to reason deductively in order to correctly solve for transitive questions (Wright 1996), tasks become more challenging for the over-fours.

Wright (1998a) found that the age at which this “logical” transitive inference ability emerges is nearer 7 and 4. Markovits et al. (1995) replicated the Pears and Bryant experiment (cited in the target article), controlling for height cues, and again determined that children under around 7 years of age do not possess true transitivity. These experiments not only give differing results to the standard story but also contribute actively to an explanation of experiments that have led to lower age estimates for the emergence of the transitive inference ability. This is in terms of the utility of nonlogical absolute cues to correct solutions together with a general memory effect (Wright 1998b). It would therefore seem that the age from which children are able to solve transitive inference tasks, through the use of genuinely deductive strategies is nearer 7 than 4, with the Pears and Bryant results still indexing the beginning of a transitive ability but not that supposed to be of interest.

Halford et al. guard indirectly against possible problems that could stem from the finding that the age for transitive reasoning differs from the one they advance by giving a secondary role to “experience.” However, they do not allow for the finding that there are two distinct types of transitive inference, with the wrong one currently being widely viewed as the only one (Halford & Kelly 1984). A theory that does not recognize these as distinct and does not itself give an account of both abilities is weakened.

A similar conclusion can be reached in respect to the object concept, which is also cited in support of the theory (sects. 6.2.1 and 6.2.2). To again present a single account, the authors needed to argue against the validity of the Baillargeon paradigm. However, although there are potential problems with it, they dismiss this kind of experiment too readily. Baillargeon (1986) provides an experiment that does seem to call for the conclusion that 6-month-olds do understand “meaning” over-and-above being able to discriminate between two types of event (later extended to younger infants).

Halford et al. also do not appear to have given enough credence to the possibility, raised by many studies, that the difference between the Piagetian task and the Baillargeon task may simply be that in the former infants must coordinate their manual motor apparatus between the two hiding places; whereas in the latter they need only co-ordinate their visual motor apparatus. Thus, the problem here might be related more to differential motor-planning or motor-control than to differential conceptual understanding.

The dismissal of the Baillargeon paradigm as irrelevant to object permanence is therefore premature; with this fact bringing about a situation analogous to that discussed for transitivity – there being two not one distinct Object Concept phenomena and the theory being unable either to discriminate between them or invalidate one of them.

What role does experience play? Halford et al. discuss implications for developmental psychology. Cognitive development is seen as children gaining a new mode of processing at certain points in time, becoming newly capable of handling an additional dimension in parallel with those already available. Experience allows them to capitalise on this new mode, such that they can now acquire new concepts that demand this increased capacity.

However, Halford et al. do not offer a credible account of how and why new modes of processing become apparent in the first place. Perhaps experience rather than relational complexity should be assigned the primary role in the theory here, and viewed as acting recursively upon a cognitive system that, in a way right from the start, has always been capable of supporting complex modes of processing up to some limit.

For example, repeated experience with identical binary problems would give rise to more and more efficient binary processes until ternary processing is born, with quaternary processing being

born in much the same way. We might further theorise that this occurs as a result of processing sequential binary steps closer and closer together in time (increased automaticity) until they are essentially processed in parallel. At this point the new mode of processing is born. On this view cognitive development ultimately arises out of repeated experience.

Experience almost certainly plays another more “direct” role, allowing new concepts to be formed without the need for primitives of fewer dimensions. Thus, to use the example of force (sect. 3.4.1), this concept can be (and probably is) created without the need to first understand “acceleration” or “velocity” (sect. 6.3.1 embraces this assertion).

In summary, Halford et al. find a need to link the theory directly to currently accepted developmental phenomena; but if multiple forms of those phenomena are a fact, then these cannot be relied on for unambiguous support unless the theory can discriminate each instantiation. Halford et al. also cite relational complexity as the primary cause of cognitive development, when this, itself, would seem likely to arise out of repeated experience. Direct experience almost certainly also allows the construction of concepts that, theoretically, require higher level dimensionality, without first acquiring related concepts of lower dimensionality. Hence, although the theory allows the integration of a great many independent cognitive domains of discourse, in its current form, it does not offer an adequate account of cognitive growth.

Authors’ Response

Relational complexity metric is effective when assessments are based on actual cognitive processes

Graeme Halford,^a William H. Wilson,^b and Steven Phillips^c

^aDepartment of Psychology; University of Queensland, Brisbane, Queensland, 4072 Australia; ^bSchool of Computer Science and Engineering, University of New South Wales, Sydney, New South Wales, 2052 Australia;

^cCognitive Science Section, Electrotechnical Laboratory, Tsukuba, 305, Japan. gsh@psy.uq.edu.au

www.psy.uq.edu.au/people/department/gsh billw@cse.unsw.edu.au; www.cse.unsw.edu.au/~billw stevep@etl.go.jp

www.etl.go.jp/etl/ninchi/stevep@etl.go.jp/welcome.html

Abstract: The core issue of our target article concerns how relational complexity should be assessed. We propose that assessments must be based on actual cognitive processes used in performing each step of a task. Complexity comparisons are important for the orderly interpretation of research findings. The links between relational complexity theory and several other formulations, as well as its implications for neural functioning, connectionist models, the roles of knowledge, and individual and developmental differences, are considered.

R1. Introduction

The target article proposed that relational complexity, defined as the number of arguments of a relation, provides the best measure of complexity in higher cognitive processes. The relational complexity metric per se does not seem to have been challenged in most of the commentaries, but they do raise many other issues, which we will consider in

turn. First, to avoid misconceptions, we would like to briefly summarise our position, placing particular emphasis on what we are *not* attempting to explain.

It is important to emphasise that performance on cognitive tasks depends on a number of factors, including familiarity with the task, domain knowledge, availability of appropriate strategies, storage and retrieval of relevant information, and so on. We fully accept the importance of these factors, and much of our work in other contexts has been devoted to them. (See, for example, work on learning, Halford 1995, induction, Halford et al., in press, and strategy development (Halford et al. 1995). However in the target article we are concerned with observed effects of *complexity* on higher cognition. There are many cognitive phenomena that require a complexity metric for their explanation. We propose that these phenomena can be accounted for by the complexity of relations that have to be processed in a single decision.

We also wish to re-emphasise that the capacity limitations we defined in the target article apply in cases where a cognitive process operates on a representation of relations in the task. This representation should have the properties of relational knowledge defined in section 2.2. Associative processes and modular processes that are specialised for processing a limited range of inputs do not have these properties (Anderson 1992; Halford 1996b; 1996c; Leslie & Keeble 1987; Phillips et al. 1995) and therefore are not expected to show the same complexity effects.

R2. How complexity should be analysed

Sweller confirms that he has found in empirical work that the number of interacting variables processed by a performer is a good measure of task demand. Relational complexity subsumes number of interacting variables because, as noted in sections 2 and 2.3, each argument of a relation effectively functions as a variable, and an n -ary relation is a set of points in an n -dimensional space. Sweller has also found that analysis of complexity depends on what dimensions are chunked, and this depends on expertise.

We agree, and in section 2.1 we make the more general point that complexity depends on the cognitive processes that are being employed. The processes employed by a particular person in performing a task undoubtedly depend on expertise, age, problem presentation, goals, and so forth. All of these factors affect the way a task is chunked and/or segmented. We said in section 6 that analysis of relational complexity depends on having a process model of the way the task is performed. The process model will need to be verified independently to make predictions based on complexity testable. Some of the commentaries indicate that the implications of these requirements have not been fully recognised.

Formal similarity of tasks does not guarantee that processes are similar, so **Sweller's** statement that isomorphic tasks have the same relational complexity is incorrect. This can be illustrated by contrasting two isomorphic tasks, one based on days and the other based on numbers: one, *Suppose 5 days after the day before yesterday is Friday. What day of the week is tomorrow?* and two, $x + 5 - 2 - 1 = 6$. If we number the days of the week consecutively, so that Sunday is 1 and Friday is 6, these problems are isomorphs,

but they differ markedly in difficulty. At first sight this might seem to imply that relational complexity cannot account for the difficulty, but in fact the tasks differ substantially in the relations that are processed. The arithmetic task is easily segmented: we can perform $5 - 2 = 3$, which is a ternary relation (binary operation); similarly for $3 - 1 = 2$. Now we have $x + 2 = 6$. Subtracting 2 from both sides involves simply cancelling the 2 and performing $6 - 2 = 4$, another ternary relation, so $x = 4$. Mapping this into the word problem, tomorrow is Wednesday. The arithmetic problem can be solved by a series of steps that require, at most, ternary relations. Because we have already learned the steps and know how to apply to them to equations of this kind, the planning process imposes a negligible load. By contrast, success in the word problem depends on devising a strategy that segments the problem into a series of manageable steps. This is difficult to do, however, in part because planning a suitable strategy depends on first representing the structure of the task as a whole. It is the planning process that imposes high processing loads in this case, not performance of the individual steps. For example, to translate the days problem into the isomorphic arithmetic problem requires recognition of the correspondence between the two structures. The load imposed by planning some other strategies might not be as high, but it is still substantial.

The crucial point therefore is that relational complexity analysis has a good chance of accounting for task difficulty when it is applied to the processes used. **Sweller** is right that determining these processes can involve a lot of work, but this is not a burden imposed specifically by our model. It is inherent in any genuine attempt to analyse cognitive complexity. On the positive side, it is a task for which cognitive psychology and cognitive science are now well equipped. There is a rich array of techniques for theoretical modeling and empirical analysis. Perhaps most important of all, clear and coherent accounts of cognitive processes are a major benefit in themselves, and it might not be unreasonable to suggest that they should be the ultimate goal of our research. Because complexity is a major factor effecting performance, albeit by no means the only factor, having a precise way of defining complexity can be of considerable benefit in our efforts to understand cognitive processes.

Chalmers & McGonigle attribute many ideas to us that form no part of our theory. One example is that, in common with both **Wright** and **Goswami**, they suggest that our conception of relational complexity is not empirically founded, which is simply untrue. In fact we have examined an extensive data base in arriving at our conclusions (see, for example, the references cited by Halford 1982; 1993), and this data base includes the transitivity of choice paradigm on which Chalmers & McGonigle's position is based. Chalmers & McGonigle also attribute to us the assumption that people adopt representations that maximise the amount of information processed in parallel. They apparently missed our statement in section 6.1.4:

We assume that participants normally segment sentences into constituents which are processed serially as far as possible. In all our modeling, in this and other contexts, we have found it a fruitful assumption that participants tend to minimise processing demand, implying that they never process more information in parallel than necessary.

Perhaps because these commentators missed this point, they mistakenly suggest that we regard multiply centre-embedded sentences as the norm. On the contrary, we stated that the interest in these sentences is that they constrain participants to process more of the sentence in parallel. This enables processing limitations to be observed. This logic has also been used by Henderson (1994) and Just and Carpenter (1992), whom we cited in section 6.1.4.

R3. Why tasks impose loads

Although it is obvious that people will tend to minimise processing loads, high loads are imposed by the nature of some tasks. The reason is that where sources of information interact, they must be considered jointly. This principle is routinely applied when we analyse effective task complexity. It can be illustrated using analysis of variance: interacting factors cannot be interpreted singly, because the effect of any factor is modified by the others. All factors that enter into a particular interaction must be interpreted jointly. This principle provides an objective criterion for complexity and has been the basis for our analyses. Premise integration in transitive inference (discussed in sect. 6.1.1) illustrates the point. Premises such as “John is smarter than Stan, Tom is smarter than John” interact in that the interpretation of one premise is influenced by the other. Therefore, each premise contains some ambiguity when processed alone, but the ambiguity is removed when they are considered jointly. Notice first that “John is smarter than Stan” assigns John to first or second position, whereas “Tom is smarter than John” assigns John to either second or third position. However, when considered jointly, the two premises assign John uniquely to second position. Tom and Stan are then unambiguously assigned to first and third position, respectively. Because neither premise can be fully interpreted in isolation, there is a limit to segmentation. Relational complexity theory predicts that the need to consider premises jointly will impose a processing load, which has been verified empirically (Maybery et al. 1986).

The most immediate implication of this is that tasks cannot always be segmented into steps that are performed serially. Ability to segment depends on a number of factors, including task structure, conditions of performance, and expertise. The analysis above would be consistent with the expertise of most people in an industrial society aged between roughly 5 years and late adulthood. It would apply where the premises have to be integrated mentally after a single presentation, as occurs in many experiments with human participants. However, there are other conditions that make segmentation easier. For example, we could present premises one at a time and let children order blocks. Thus, given $a > b$, the child can place ab , then when given $b > c$ they can concatenate c yielding abc , and so on. By this procedure, they never have to consider more than one relation in any decision (Halford 1984). Consequently, the processing load is less, and performance of children is much better (Andrews & Halford, in press). We routinely use tasks that have this property as control procedures in our experimental studies of transitivity. Thus segmentation can be influenced by task conditions, but is also constrained by structure that is inherent in the task.

Chalmers & McGonigle contend that premise integration is unnecessary in transitive inference tasks and that

processing is serial. To support their claim they rely on a paradigm derived from the work of Bryant and Trabasso (1971). Participants are systematically trained on pairwise comparisons (e.g., A versus B, B versus C, etc.) then tested on untrained pairs (e.g., A versus C). Transitivity of choice has been demonstrated in species from pigeons to humans. Having been trained, usually over many trials, on $A+B-$ (i.e., to select A rather than B) and $B+C-$, and so on, they show a transitive bias when presented with non-adjacent pairs, choosing B in preference to D, and so forth. However, the transitive bias is reduced when triads such as BCD are presented. This would not be expected if participants constructed an ordered array (A,B,C,D,E), because they would have stored the ordinal position of all elements. Harris & McGonigle (1994) claim that performance of primates is consistent with production rules that select items near one end of the series and avoid items at the other end. Children perform this paradigm in essentially the same way (Chalmers & McGonigle 1984). Wynne (1995) has shown that pigeon data are well explained by associative learning models that do not entail processing relations at all. These models do not entail representation of the relations in the task (e.g., they do not entail representation of the relation between A and B, B and C, etc.), there is nothing equivalent to a transitivity principle, no ordered array is constructed, and performance depends solely on relative strengths of item preferences, which are acquired by learning. Therefore this research shows that at least some version of the *transitivity of choice* task can be performed without a cognitive representation of relations between task elements; that is, without processing structure. In fact it can be performed by much lower level cognitive processes.

The processes used in this paradigm reflect task demands. The paradigm shows that if the task can be performed using serial processes that impose low demands on capacity, then most participants will perform it that way. **Chalmers & McGonigle** then ask why these low demand serial processes are not used all the time. The answer is that they suffice for only a very restricted range of tasks. The transitivity of choice task cannot be adopted as a paradigm for all human cognition. Many cognitive tasks require more information to be processed in parallel and more elaborate cognitive representations. Even transitive inference requires this when premises have to be integrated mentally after a single presentation, as we saw above. Furthermore, there is strong evidence that ordered arrays are constructed under some conditions, even by young children. Andrews & Halford (in press) presented young children with premises coded as pairs of coloured blocks, with A above B, B above C, . . . , D above E. They were asked to infer the relative position of B and D, then to place C. If children had constructed an ordered array, it would be easy to insert C between B and D. The proportion of correct placements of C, given that BD had been placed correctly, increased from 42% at age 4 years to 95% at age 6 years. This is added to several other lines of evidence suggesting that integration of premises to form an ordered array not only occurs, but is a cognitive achievement that is related to age.

Chalmers & McGonigle's contention that mental integration of information to form reasonably complex relational representations is an unnecessary burden simply reflects their reliance on tasks that make low demands. The transitive choice paradigm can be accounted for either by associative learning models or models that postulate pref-

erence for certain stimuli. It does not require relational processing that would conform to the principles of relational knowledge that we defined in section 2.2. Furthermore, McGonigle and Jones (1978) found that discriminating the middle item was more difficult than transitivity of choice. “Middle” is a ternary relation, though it can be chunked to binary in some circumstances, as we will consider below. Either way it is more complex than the associative processes entailed in transitivity of choice, and the finding that it is harder is clearly consistent with our position, but Chalmers & McGonigle do not acknowledge this. Criteria that distinguish between associative and relational knowledge have been defined by Halford et al. (1998). We think it is unlikely that the very restricted processes advocated by Chalmers & McGonigle would meet these criteria.

Chalmers & McGonigle are incorrect when they suggest that eight-item seriation requires an oct-ternary relation. To make this error they must have missed a major component of our argument (e.g., in sect. 2.1), which is that processing load is determined by the complexity of relations processed in a given step, not simply by the total amount of information contained in the task. We regard it as fundamental that processing load depends on the amount of information processed in any given step, not on the total amount of information in the task. Our reduction technique for defining relational complexity of a concept in section 3.4.3 (which they cite to support their claim) means that forming an eight-item series requires ternary (not oct-ternary) relations to be processed in a single decision. The point can be made more simply by considering seriation of five items. Suppose premises $A > B$, $A > C$, $A > D$, $A > E$, $B > C$, $B > D$, $B > E$, $C > E$, $D > E$ are presented in random order (noting that premises do not necessarily have to be restricted to adjacent items, nor do they have to be presented in any particular order). Given premise $A > C$, the string AC can be formed. If $A > B$ is presented, possible strings are ABC and ACB. We need premise $B > C$ to decide that the correct order is ABC; that is, $A > B$ and $B > C$ jointly determine that the order is ABC. The task can be performed by a series of steps, the most complex of which entails dealing with ordered three-tuples. That is, it only requires processing ternary relations.

It might be asked which is the most appropriate or most valid test of transitive inference? Our contention is that any of the tests is capable of providing information about capabilities of humans and other animals, but the less demanding tasks do not demonstrate complete mastery of the concept. Given a low demand task, such as transitivity of choice, participants utilise low-level cognitive strategies. This simply says that performers are rational in the sense of Anderson (1990) because they use the least demanding strategy that suffices for the task at hand. However, we would be unwilling to concede that participants who could perform no other strategy had mastered transitivity. This contention might be unwelcome to those who want to attribute transitivity to very young children or nonhuman animals, who show no evidence of using the more demanding strategies. It would not be appropriate to dismiss performances based on tendencies to select specific stimuli, as in the associative learning model of Wynne (1995) or the production rules of Harris and McGonigle (1994). We are glad to acknowledge that the careful and ingenious research conducted in that paradigm has yielded valuable knowledge of the cognitive processes utilised there by a wide range of

participants. There is no reason, however, to accept those models as comprehensive accounts of human reasoning. Even relatively simple reasoning tasks can be demonstrated easily to require more complex processes, as our example with transitivity above illustrates. It also seems reasonable to suggest that ability to mentally integrate premises, forming mental representations of the relations in a task, is a landmark cognitive achievement. In the case of transitive inference, it is this that we contend requires processing ternary relations.

The foregoing discussion should demonstrate the error in **Wright's** commentary when he contends that the relational complexity metric cannot handle more than one version of a task. The theory does *not* imply that there is one correct version of a task. There are many ways of assessing even simple cognitive functions such as transitive inference, and probably at least as great a variety of cognitive processes that can be employed. Complexity is assessed on the basis of the processes employed. This is not a problem because techniques for determining strategies have appeared in abundance over the last few decades. Let us note in passing that one reason why transitive inference is useful for complexity analyses is that we have well validated process models. Indeed, transitive inference research is arguably a great, though unrecognised, success story in cognitive psychology. Furthermore, complexity of processes employed is constrained by the fact that sources of information that interact must be interpreted jointly.

This point was evidently missed by **Pascual-Leone** when he contended, incorrectly, that “since the theory has no explicit rules constraining their occurrence, chunking and segmentation turn into theoretical loopholes for explaining away empirical anomalies.” In fact general principles of chunking were given in section 3.4.1, the reduction technique outlined in section 3.4.3 provides an objective way to determine the effective relational complexity of tasks, and principles for complexity analysis were given in section 6. However, the principle that sources of information cannot be processed serially when they interact is the core of our method. It is first necessary, as we have noted all along, to have a clear model of how a task is performed. Given this, we have consistently found it useful to analyse the number of interacting variables in a given decision. Using this technique we have predicted complexity effects before they were observed in the balance scale, hypothesis testing, and concept of mind (Halford 1993). We also have been able to analyse tasks as diverse as classification, Tower of Hanoi, knights and knaves (Rips 1989), and Raven's matrices. We also have found very good correspondences across domains.

Pascual-Leone's position does not appear to have any comparable way of removing subjectivity from complexity analyses. We consider that explicit computational models are a more objective way to determine the nature and complexity of processes used in a task. For example, the demonstration that a particular task can be performed by a typical three-layered net is a good indication that it is basically associative and does not require representation of explicit relations. Thus the demonstration by Quinn and Johnson (1997) that prototype formation can be achieved by a three-layered net suggests that it does not require relational processing.

Pascual-Leone argues that complexity estimates for proportion $a/b=c/d$ must include the ratios a/b and c/d and

also suggests that we do not allow for these ratios. We contend that both arguments can be refuted. Our position is that whether the ratios must be included depends on whether they are processed by the performer. Despite his suggestion to the contrary, we do take this into account. Processing of proportion can be simplified by first computing $a/b = x$ and $c/d = y$, then comparing x and y . Thus, if asked whether $4/8 = 12/24$, one could compute $4/8 = .5$ and $12/24 = .5$ and conclude the expression is true. This segments proportion into two ternary relations performed in succession. Notice, however, that to plan this procedure, or to understand why it is valid, the structure of proportion must be represented. A proportion is truly a quaternary relation and is defined by links between four variables. It is defined in four-dimensional space, not six-dimensional space, as contended by Pascual-Leone. Note that Pascual-Leone's position implies that from age 15 years humans can process seven dimensions in parallel. Pascual-Leone does not seem to offer any evidence that this is so, nor does he consider its theoretical consequences. What kind of cognitive processes could be performed by a system that was processing seven dimensions in parallel? This seems to be quite an unexplored territory.

Pascual-Leone misrepresents our position in at least two other ways. He incorrectly puts us in opposition to the view that effective complexity is defined relative to an adaptive system and can be represented by a concise description of the regularities in the task. Our position is quite closely related to this view and also to contributions such as that of Leeuwenberg (1969), who was early in his recognition of the potential for defining psychological complexity in this way. Leeuwenberg's metric was explicitly applied in some of our earlier work (Halford & MacDonald 1977). Because we define complexity by the space in which a cognitive process is operating, our position seems quite consistent with the one that Pascual-Leone wants to put in opposition to it. His claim that in section 3.2 we assigned "*played(John, cricket, oval, Sunday)*" to an effective complexity of four is also incorrect. That relational instance is simply a list, and its relational complexity as defined by the reduction technique in section 3.4.3 is indeterminate.

R4. Links to other formulations

We agree with **Anderson et al.** that there are points of contact between relational complexity theory and ACT-R. In particular, we agree that the name for a production could be used to represent the symbol (name) of a relation, and the slots of a production can correspond to the arguments of a relation. However, a means of binding the relation symbol and arguments would still be required, and it is not clear how that would be done in a production. We also agree that ACT-RN uses separate memories in a way that is analogous to our use of different sets of units in a tensor product to represent arguments of a relation. There is partial correspondence between the roles for activation in the two models. In ACT-R, activation is the main cause of capacity limitations. In relational complexity theory, the demand for activation increases with the rank of a tensor product (see sect. 5.2.1.2). However, it is not the only factor that limits complexity in our model, because number of arguments corresponds to the rank of a tensor, which constrains the number of connections between units.

Although the points of contact offer interesting potential,

considerable work is still needed to achieve a genuine integration, and neither relational complexity nor any comparable metric is incorporated into ACT-R as it stands. One of the clearest differences is that in ACT-R complexity is assessed by the number of symbols (Anderson et al. 1996), whereas in our model it is based on the dimensionality of the decision space. The difference can be simply illustrated by contrasting the following sets of problems in both of which participants must solve for x :

$$\begin{aligned} \text{Set 1: } & x + 4 = 7; x + 3 = 7; x + 4 = 8; x + 3 = 8 \\ \text{Set 2: } & x + 1 = 7; x + 1 = 8; x + 1 = 5; x + 1 = 6 \end{aligned}$$

A count of the number of symbols, following Anderson et al. (1996), yields 5 for both sets. However, the relational complexity of the first set is 3 (because there are three dimensions of variation) but for the second set relational complexity is 2 (two dimensions of variation). It seems intuitively likely that set 2 would also be easier. As we noted above, there are no simple ways to arrive at a valid assessment of the dimensionality of the cognitive processes performed in a task, and counting the number of symbols is not adequate.

Incorporation of the relational complexity metric into a model with the power and generality of ACT-R would be a very exciting development. However, considerable work is required to accomplish this. One of many benefits would be that ability to estimate the amount of information that humans processed in parallel should provide a useful constraint on ACT-R, just as it has on models of analogy (Halford et al. 1995; Hummel & Holyoak 1997).

MacLeod addresses the problem of the optimum allocation of resources to a set of activities (or tasks), given that each activity has costs and benefits. We should note that this is not the core question addressed by the target article, which is concerned with assessment of cognitive complexity, even in single tasks. MacLeod's approach may be productive as a theory of dual task performance, perhaps by extension of the work of Navon and Gopher (1979), who also applied economics theory to the resource problem in psychology. MacLeod's model of resource allocation, like the commentaries by **Anderson et al.** and **Cowan**, provides an interesting source of hypotheses for the cause of complexity at a lower level of analysis. We also note that our position is quite unequivocal on several of the points raised by MacLeod. We have proposed that the complexity of a task depends on the arity of a relation processed in parallel. This is related to dimensionality because, as we pointed out in section 2.3, an n -ary relation is a set of points in n -dimensional space and can represent an interaction between n variables. Therefore, complexity is related to dimensionality, not to the length of the input string. Furthermore, complexity effects occur because computational cost is a function of the number of dimensions processed in parallel. We agree that a high dimensional problem can be mapped into fewer dimensions. This is the essence of chunking, discussed in section 3.4.1, though it should be recalled that we specified limits to this process.

We do not agree that complexity is a direct result of resource limitations, though resources that are available affect our ability to deal with complexity. This discrepancy seems to arise because the problem addressed by **MacLeod** – optimum allocation of resources to activities – is not the problem addressed by us. As discussed in an earlier section of this reply, complexity effects arise because when dimen-

sions of a task interact they need to be processed jointly and this increases processing load. Therefore, processing load results from constraints that are inherent in the process being performed, not from resource limitations per se.

In relation to NP-completeness, NP stands for nondeterministic polynomial time problem. Similarly, P stands for the class of problems solvable by a polynomial-time algorithm. Problems in NP have a structure such that any hypothesis can be checked in polynomial time, but the total number of hypotheses is exponential in the size of the problem.

MacLeod is, of course, correct in pointing out that no one has yet proved that P and NP are distinct. It is possible that there are as yet unknown algorithms to solve NP-complete problems in polynomial time, just as more efficient algorithms have been discovered from time to time for various tasks such as sorting. However, at present there is no known efficient (polynomial time) algorithm for such problems, so the time (or space) required to compute a solution is exponential in the size of the problem.

MacLeod indicates he would have preferred “a more detailed presentation of one of the mathematical models, along with some theorems on the algorithmic complexity of the proposed mechanisms.” We will be gratified if our formulation opens up opportunities for further theoretical development, but the aim of the target article was to find a complexity metric that would be broadly applicable to psychological tasks, taking account of the real problems of conducting psychological research. In sections 5.2.1.1 and 5.2.1.2 we offered analyses of the algorithmic complexities of the basic processes of tensor product network operations. The tasks in which we currently are most interested are those that can be completed in a few such operations, and the processes in which we are most interested are the simplest ones for performing such tasks.

R5. Frontal lobe impairment

Our suggestion in section 6.5, based on the review by Robin and Holyoak (1994), that processing explicit relations might be a major function of the frontal lobes has now received some empirical support. The commentary by **Waltz et al.** reviews several empirical studies, including their own ongoing research, indicating that patients with damage to the dorsolateral prefrontal cortex show selective impairment in processing complex relations. Patients with damage to the anterior temporal lobes showed no such impairment. The possibility that relational complexity offers a new way of defining the functions of a major region of the frontal lobes is clearly quite exciting. The techniques developed in our laboratory, discussed in the previous section, for manipulating complexity while holding other factors constant, should have considerable utility in testing this hypothesis. Relational complexity measures also can be used to assess deterioration due to other causes, including ageing.

R6. Cause or effect: Speed versus capacity

The complex issue of whether speed or capacity is the causal factor influencing performance in complex tasks has been well elucidated by **Cowan**. He presents some interesting evidence that tends to favour capacity as cause. On the other hand, there is very extensive developmental work

by **Kail** (1991) that indicates a general processing speed factor that increases with age. We have reviewed this work elsewhere (Halford 1993, pp. 119–22). We think both speed and capacity hypotheses will remain viable for the foreseeable future and we can envisage mechanisms that are appropriate to either.

Assuming that a cognitive phenomenon can be modelled by a neural net with a settling phase, it is reasonable to expect that a network that is closer to its capacity limit will take longer to settle. For example, if more items are superimposed on a fixed set of units, the items become less discriminable. Discriminability can be increased by increasing the number of units, which is one measure of capacity in a net. A net with higher capacity will therefore have a shorter decision time. Some distributed memory theories appear to model this effect fairly directly (McNicol & Stewart 1980; Murdock 1983). In these models memory items are represented as vectors that are superimposed on a set of units. Recognition time increases linearly with number of items, thereby predicting the set size effect in memory scanning (Sternberg 1975). The implication is that more items impose higher demands on the available capacity, which increases decision time.

Although **Cowan's** example is orthogonal to the relational complexity metric (a list treated as a unary relation is still unary regardless of its length), Cowan's data may be interpretable in terms of this model. If we assume that inter-word pauses in Cowan's experiments represent decision time, then pauses would increase with list length if retrieval from longer lists imposes higher demand on available capacity. Pause time also would increase with lower capacity. If we assume participants with lower span have less capacity in some sense, then the longer pauses observed by Cowan with lower span participants can be explained. Although this argument is speculative, our model may be consistent with Cowan's data. It also suggests a testable hypothesis. This is that participants with a lower relational complexity limit, as defined by the measures developed by Andrews and Halford (in preparation) should have longer inter-word pauses.

In case of the symbol-argument tensor modeling, transmission speed could affect capacity in the following way: Suppose role units R1 and R2 and binding units for computing the appropriate relations. Post synaptic activation is only possible when all presynaptic activations from the various role units occur within the same presynaptic activation window (w). Suppose the distances from R1 and R2 to the binding units are D1 and D2, respectively, and transmission speed is s . A signal sent from unit R1 at time t will arrive at the binding unit at time $t + D1/s$. A signal sent from unit R2 at time $t + d$ will arrive at the binding units at time $t + d + D2/s$, assuming a difference d between signal transmission from units R1 and R2. Therefore, the binding unit will be active when $d + (D1 - D2)/s < w$. Clearly, the faster the transmission speed the more likely the two signals will arrive within the activation window of the binding unit, permitting the representation of the binary relation, or other higher arity relations.

In the case of processing capacity (number of units), higher ranked tensor units may simply not be connected to represent higher arity relations. From Thatcher (1994), note that although total number of units remains relatively constant through the first decade, connectivity does not, possibly suggesting the establishment of progressively

higher order connections within the cortex. These two examples are meant as illustrations. Although we cannot decide the issue at present, relational complexity does not lead to a conceptual dead-end at finer levels of analysis. Rather, it suggests more detailed hypotheses to be tested by future experiments.

R7. Complexity versus interference

Navon is concerned with the longstanding issue of whether performance is a function of outcome conflict rather than resource limitations. We have pointed out elsewhere (Halford 1993, Ch. 3) that the easy-to-hard paradigm (Hunt & Lansman 1982) avoids the interpretational difficulties associated with dual task deficit methodologies. This is because the hard version of the primary task is never performed concurrently with a dual task, so the variance it shares with the predictor, which is taken as evidence of a resource limitation, cannot reflect outcome conflict. **Berch & Foley** acknowledge that the easy-to-hard paradigm (Hunt & Lansman 1982) is an effective way to test hypotheses about capacity limitations, but point out that it is methodologically complex. We agree that this is so, but their excellent work (Foley 1997; Foley & Berch 1997) has demonstrated that it is effective for this purpose. The effort of using the easy-to-hard paradigm, though significant, is worthwhile.

In general, there have been many advances in methodology since Navon (1984) made his claim that resource limitations were difficult to distinguish from outcome conflict. These include experiments (e.g., Andrews & Halford, in press; Maybery et al. 1986) where both input and output are tightly controlled while the complexity of the intervening processes are manipulated. It seems very unlikely that the effects of relational complexity, some of which are very large, that have been demonstrated in these studies could be attributed to outcome conflict. Outcome conflict only applies at best to the dual task interference paradigm. However, the target article goes well beyond this paradigm and is primarily concerned with complexity-based limitations. Even if outcome conflict were a viable explanation in some studies, the number and diversity of complexity effects is so great that it is very implausible that they could all be explained by outcome conflict.

The question of relational tasks performed concurrently, raised by **Navon**, was addressed in section 3.3. Neural net models of the type we used to implement our theory are capable in principle of dealing with similarity effects. In particular, the shortcomings that Navon sees as applying to neural network models with respect to resource conflict are not relevant to complexity-based limitations. The complexity of the memory scanning task that is raised by Navon is addressed by Halford et al. (1988) and referred to in section 3.3 of the target article. An unordered string of digits is not relational and can be represented by a set of superimposed vectors, as noted in the previous section. An ordered string of items in short term memory is a binary task, assuming that the person's strategy is to store the list as position-item pairs; $\{(position1,item1), (position2,item2), \dots\}$. There is no growth in relational complexity as the number of digits is increased, though there may be an increased load on working memory if the items are held in an active state. Criteria for this were discussed by Halford et al. (1988).

R8. Connectionist models of relations

We agree with **Plate** that there is considerable scope for further connectionist proposals for relational processing. Note that, as Plate's holographic reduced representations (HRRs) are projections of tensor product nets (in the mathematical sense of "projection"), it is not surprising that systems based on these two approaches share some properties. We also agree that role-filler binding methods, suitably augmented with content-addressable memory, are sufficient to represent and access relational information. However, we do not agree that HRRs implement what we intended by the term *conceptual chunking*. Rather, chunking entails loss of information through dimension reduction, not noise reduction through compression. Furthermore, in light of the double dissociation reported by **Waltz et al.**, there is now a neurologic reason to favour symbol-argument over role-filler binding methods: symbol-argument bindings used different ranked tensors (i.e., different unit types) for different arity relations, whereas role-filler methods use the same tensor units. The former seems more consistent with the existence of a region in the brain, the dorso-lateral prefrontal cortex, that is specialised for processing high arity relations.

Plate argues that HRRs do, in fact, implement chunking. Although HRRs satisfy our principles (1) and (3) in section 3.4.1, they do not satisfy principle (2). Plate states that "every vector in an HRR model is already a chunk and no further compression . . . is necessary." In this case, there is no sense in which one can have both a chunked and an unchunked representation of the same concept. Thus, either satisfying principle (2) "no relations can be represented between items within a chunk" implies not being able to represent relations, or more likely, because HRRs represent relations, they cannot satisfy principle (2).

The difference is that with an appropriate (learned) chunking strategy chunks identify unique equivalence classes of concepts. In our $v = s/t$ example, the chunk "60" represents a class of relational instances: $Div(60,1)$; $Div(120,2)$, and so on. A chunked version of the corresponding rank two tensor representing the relations may be the vector determined by actually calculating the velocity. In **Plate's** method, by contrast, these equivalence classes occur by accidental collisions between vectors, at the mercy of the statistics of vector generation and manipulation. Of course, it would be possible to augment HRRs with a chunking mechanism and represent both the chunked and unchunked representations as HRRs. However, Plate's HRR model, as we understand it, does not currently do this.

R9. Symbol-argument versus role-filler binding

Waltz et al. provide further experimental support for our relational complexity metric, but remarked that the data are neutral on the issue of algorithms (e.g., whether relational information is computed by tensor product or synchronous activation). We are excited by these findings, but suggest there is also some room for interpretation at the algorithmic level of analysis.

The double dissociation between the prefrontal cortex (integration of two binary relations) and the anterior temporal lobe (single binary relation) suggests different underlying neural architectures for relations of different arities.

It has been noted elsewhere (Phillips & Halford 1997) that the tensor symbol-argument binding method, but not the role-filler method (which includes **Plate's** HRRs), permits such double dissociations. For the symbol-argument method, binary and ternary relations require distinct unit types; for example, rank 2 tensor units multiply two incoming sources of activation, compared with three sources for the rank 3 units. By contrast, role-filler methods, and for that matter synchronous activation, use the same type of unit regardless of arity. With role-filler methods, it may be the case that additional units must be recruited for higher arity relations. However, this would only demonstrate a dissociation, not a double dissociation, because first and second positions are shared by both relation types.

However, because we do not know exactly what part of the integration process the prefrontal cortex is responsible for, this interpretation is at best suggestive. We agree with **Waltz et al.** in that much more work at the neurologic level is required before definitive claims can be made about neural mechanisms.

R10. Relational complexity and chaotic attractors

Three commentaries, those by **Heath & Hayes**, **Nikolic**, and **Borisyuk et al.**, have drawn attention to some possible correspondence between relational complexity and dynamic systems theory. Both formalisms are characterised by a common limitation, in that the range of dimensionalities is low, from one to four. The basis for this common limitation is well worth further exploration. At the present time the correspondence, although intriguing, is difficult to interpret. **Nikolic** finds the same limitation in motor processes that we find in higher cognitive processes. It may be, as **Heath & Hayes** suggest, that the limitation is a general property of neural processes. However, at present it is difficult to determine whether these common limitations are coincidental or represent a genuine underlying phenomenon. The potential importance of correspondence between relational complexity and dynamic systems theory is great enough, however, to warrant more extensive investigation.

An integrated treatment could be very productive for another reason, which is that the approaches have complementary advantages. Dynamic systems theory has elegant ways of predicting discontinuities, whereas relational complexity theory is more closely tied to observable phenomena. The levels of complexity defined in the target article have been identified through detailed analyses of tasks across a wide range of domains, and complexities have been manipulated precisely while holding other factors constant. If these levels of complexity are found to correspond to dimensionalities of chaotic attractors, the methodology of relational complexity theory would seem to open up a lot of opportunities for empirical testing of dynamic systems models of cognitive processes.

R11. Knowledge and higher-order relations

A number of commentators, including **Wright** and **Gentner & Ratterman**, claim that we give priority to capacity rather than knowledge as an explanatory factor. In fact, we con-

sider that they interact. See, for example, Halford (1993, p. 272), Postulate 1.0: "Cognitive development depends on the interaction of learning and induction processes with growth in the capacity to represent concepts." To the extent that these factors interact, neither can be given priority. In statistical analysis an interaction cannot be decomposed to determine the relative importance of the constituent factors. Similarly, where knowledge and processing capacity interact to determine performance of a task, it is meaningless to ask which factor is more important.

In empirical research either factor may have a larger effect in a specific study, depending on how it is designed. A popular design is to use tasks that are well within the capacity of the participants, but that demand knowledge that has only been partially mastered by the sample selected. For example, studies of analogical reasoning in young children frequently use A:B::C:D analogies, which require binary relations to be mapped, and capacity to process binary relations appears to develop at a median age of 2 years. However, relations such as melting snow are used (Goswami & Brown 1989), which children of 3 to 4 years are just beginning to understand. Such studies typically show that knowledge has a large effect, and they are inherently incapable of showing an effect of capacity. Alternatively, one could use tasks in which children have thoroughly mastered the prerequisite knowledge, either by selecting familiar materials or by using extensive training (Halford 1980). One can then manipulate the capacity requirements of the task from, for example, binary to ternary with 4- to 6-year-old children. In these circumstances, complexity produces a large effect on performance (Andrews & Halford, in press). Neither type of study shows any general priority for knowledge or capacity. The overall picture is that knowledge and capacity both have effects and, not surprisingly, investigators design their studies to reveal the factors they want to investigate.

Gentner & Ratterman make the more specific claim that analogical reasoning in children depends on higher-order relations, which in turn depend on knowledge. Our response to this is that we have fully acknowledged the importance of higher-order relations. For example, in section 6.1.3 we analysed the Tower of Hanoi in terms of the higher-order relation "prior"; $\text{Prior}(\text{shift}(2,C), \text{shift}(1,B))$. In this context more complex tasks entail deeper structures, with more levels of embedding of higher-order relations. **Gentner & Ratterman** postulate that depth of relational structures is more important than their dimensionality. However, we showed that these tasks also entail higher dimensionality. Therefore, some of the effects that **Gentner & Ratterman** attribute to depth of structure can equally be attributed to dimensionality. This issue is amenable to test, at least in principle, because depth of structure can be separated from dimensionality in certain cases. Thus, we can have a deep structure in which each level comprises a unary relation; for example, $R_3(R_2(R_1(A)))$. Depth is three because there are three levels of relations. Dimensionality depends, as usual, on the number of variables. If A is the only variable dimensionality is one. However, if any or all of R_3, R_2, R_1 is a variable, dimensionality may be as high as four. In other cases depth may be 1 and dimensionality 4, for example; $R(a,b,c,d)$. The implication is that higher-order relations are important, but dimensionality is the more general criterion of complexity, and can be applied to structures of any depth.

R12. Individual differences

Sweller suggests that the theory is less useful for assessing individual differences in processing capacity. In fact, we are developing a test for complexity of relations that can be processed in parallel, and results so far are very encouraging. Andrews (1996) tested children aged 4 to 8 years in transitivity, hierarchical classification, cardinality, comprehension of relative clause sentences, hypothesis testing, and class inclusion. Relational complexity was manipulated within each domain. All tasks loaded on a single factor and factor scores were correlated with age ($r = .80$), fluid intelligence ($r = .79$), and working memory ($r = .66$). This result was replicated (Andrews & Halford, in preparation) with a slightly different set of domains in 1997.

Sweller also contends that age differences that we attribute to capacity are indistinguishable from differences due to expertise. However, proponents of this rather common view do not appear to have considered the type of experimental design we have used. First, we have used procedures and materials that are highly familiar to the children, and in some studies we have trained them to asymptote on all components of the task (Halford 1980; Halford & Leitch 1989; Halford & Wilson 1980). The typical result of these studies is that the younger children completely master all aspects of the procedure. Complexity is manipulated, but procedure and materials are tightly controlled. The result has invariably been that complexity has had a large effect on performance, especially with younger children. It might now be argued that there might have been some undetectable residual difference in expertise. It seems common to argue that this knowledge explanation must be the right one, even though the nature of the expertise differences is unspecified. It should be obvious that such an argument makes the knowledge explanation untestable. It also fails to take account of evidence, using the easy-to-hard paradigm, that some of the tasks in question are in fact capacity limited (Foley & Berch 1997; Halford & Leitch 1989; Halford et al. 1986).

Ultimately the strongest argument against the view that “knowledge explains all cognitive development, therefore there is no role for capacity” is that complexity has very real and very powerful effects on performance of young children. If these effects are entirely attributable to lack of knowledge, why can they not be removed by adequate experience or training? In the tasks where we have identified complexity effects, training has an effect only at or above the age at which capacity becomes sufficient for task demands (Halford 1980; Halford & Leitch 1989; Halford & Wilson 1980). At some point it is clearly incumbent on proponents of the *knowledge only* view to demonstrate that knowledge does account for these effects.

R12.1. Developmental issues. Many of the issues considered in earlier sections have important developmental implications. In the following sections we will deal with those specific developmental issues not considered earlier.

R13. Comparison with other developmental theories

Pascual-Leone disagrees with relational complexity theory but also attempts to identify it with his position by

equating schemes with dimensions: “the highest number of schemes – task relevant dimensions of variation – that subjects must consider simultaneously to solve the task.” However, his position would need major transformations before it could be considered equivalent to relational complexity theory. Pascual-Leone’s contention that our model predicts a performance asymptote at age 11 years is incorrect for several reasons. It fails to take account of the fact that capacity is a soft limit, and that 11 years is a median age, some children achieving capacity to process ternary relations later than this. Most important, it fails to recognise that capacity is an enabling factor and that attainment of any cognitive function depends on developing the relevant knowledge, including procedural knowledge or strategies. Any attempt to dismiss this as hypothesis-saving would have to explain away our extensive studies of learning, especially the self-modifying production system model of strategy development (Halford et al. 1995).

Coch & Fischer, like **Pascual-Leone**, emphasise the similarity between our position and theirs. However, to the extent that Pascual-Leone’s and Fischer’s theories are similar to relational complexity theory, they must be similar to each other, although neither commentary seems to claim this. As with Anderson’s ACT-R theory, discussed earlier, we welcome these comparisons, and we see some potential for constructive integration. We were, of course, well aware of the theories by Fischer and Pascual-Leone when our model was being formulated, and though we have already acknowledged some similarities, we believe it is clear that the differences are very real. We would like to note in passing that Chapman’s (1987) formulation is possibly the closest to relational complexity theory of all the developmental theories, but even here the differences greatly outweigh the similarities. Many of the differences between relational complexity theory and Case’s (1985; 1992) position, outlined in section 6.3.2, would apply to Fischer’s skill theory. To consider some other instances: Does Fischer’s skill theory make the predictions, discussed above, that **Chalmers & McGonigle**, **Goswami**, **Gentner & Ratterman**, or **Heath & Hayes** find so controversial? Does skill theory conform to the criteria for relational knowledge in section 2.2, and can it be implemented by the neural net architectures in section 4? We see no grounds for believing that either Fischer’s skill theory or Case’s model offer general conceptual complexity metrics as they stand, but it may well be an interesting undertaking to build such a metric, possibly based on relational complexity, into those models.

R14. Capacity in cognitive development

In the last two decades it has been fashionable to dismiss theories that postulate a role for capacity in cognitive development on a number of grounds, some of which appear in the commentaries. One is that capacity limitations in young children are inconsistent with a number of observed cases of precocious development. This argument is sometimes linked to another view, which is widespread though not often stated explicitly, that capacity theories are pessimistic because they imply limits to young children’s capabilities. A third argument is simply that knowledge acquisition is “preferable” to capacity as an explanation of cognitive development. We suggest that existence of capacity limita-

tions is not inconsistent, either with observed precocities or a major role for knowledge in cognitive development.

In order to clarify these issues, it will be helpful to first shift our orientation from “what is done” to “how it is done.” It then becomes clearer that the definition of processing capacity limitations is no more pessimistic about children than it is about adults. We have argued that adults typically process a maximum of four dimensions in parallel. This does not imply that adult cognition is limited, nor does it imply that adults cannot understand complex concepts. Our position is fully consistent with the observation that adults understand a lot of concepts that entail more than four dimensions. It implies that structures of more than four dimensions are processed by segmentation and conceptual chunking. This provides insights into the way cognitive processes operate by indicating which processes can be performed in parallel and which must be serial.

The definition of capacity limitations can be equally productive in cognitive development once certain inhibiting misconceptions are cleared away. It does not say that children cannot perform certain tasks, any more than it says adults cannot understand the concept of acceleration because it is based on more than four dimensions. Furthermore, as noted in section 6.3.1, relational complexity theory actually predicts some previously unrecognised capabilities of young children. It also makes a lot of predictions about how young children must be performing tasks, and it is clear that neither **Goswami** nor **Wright** has recognised this. Goswami says our caveats limit the testability of our theory. Although it may mean that the theory is not amenable to some of the more simplistic tests, it generates a lot of highly testable predictions, a few of which are indicated in the target article or in this response.

R15. Capacity versus precocity: How versus what

Ever since the early 1970s the plausibility of capacity as an explanatory factor in cognitive development has been apparently undermined by a stream of data indicating precocious intellectual performance by infants and young children. This work has generated much excitement and is even regarded by some people as the *raison d'être* of cognitive development research. Contrary to the claims of **Wright** that we dismiss these data, we have no reason to doubt that children perform as reported in these studies. Furthermore, contrary to the suggestion by **Goswami** that we have not taken into account actual results in cognitive development research, we not only knew about those studies, but have given them very serious consideration. Contrary to the claim by **Chalmers & McGonigle** that our position is not empirically based, if any studies of precocious intellectual performance had posed a real challenge to our claims, we certainly would have abandoned or modified them.

The conflict between capacity theory and demonstrations of precocity is more apparent than real, and it is maintained by two deficiencies in the field. The first is our ignorance of the processes entailed in many of these performances, particularly when they are first discovered, and the second is a tendency to interpret children's performances by attributing complex processes, comparable to those used by adults, to them. We will illustrate.

Wynn (1992) showed that infants recognise the effect of adding objects to and subtracting objects from a display

(i.e., $1 + 1 = 2$, $2 - 1 = 1$, $1 + 1 \neq 3$). It has been shown (Wynn 1995) that this can be accounted for by an innate accumulator mechanism that we share with other species. This work provides an important basis for development of children's understanding of number, though, of course, infants' number comprehension is very different from that which occurs in middle childhood (Wynn 1995). Suppose, however, that we had been ignorant of the cognitive processes entailed in infants' number performance. It could then have been argued that because infants can perform addition/subtraction, they process ternary relations at 5 months, which is wildly inconsistent with our norms of 5 years. The absurdity of such a claim would be easily recognised because we have reasonably good insights into the basis of infants' number knowledge. But how would these data have been interpreted if we had had none of these insights? Psychologists could have used our own understanding of addition as a model of infants' performance. We suggest that this is what tends to happen when we do not know the cognitive basis for precocious performances. In those cases it often seems reasonable to make very optimistic assumptions about the processes infants and young children use. This creates an apparent disconfirmation of capacity limitation, though in fact the case for disconfirmation has not been made.

Perhaps the most famous instance of precocious performance seeming to disconfirm a developmental theory is Bryant and Trabasso's (1971) demonstration of transitive inference in 3- and 4-year olds. As noted in relation to **Chalmers & McGonigle's** commentary, some of these tasks require only primitive, nonrelational processes. However, although Chalmers & McGonigle argue enthusiastically that it is performed by primitive cognitive processes, they are unwilling to acknowledge a very important implication of this, which is that the task is not performed by, and is not a measure of, processing ternary relations. The effect is to provide a misleading and confused picture of what happens in cognitive development.

We pointed out in section 6.2.4.4 that Goswami (1995) had over-interpreted her data as indicating ternary relational processing when they only require unary or binary relations, because of the widespread failure to analyse cognitive processes entailed in tasks. The same kind of failure appears again in **Goswami's** commentary. She cites work by Cutting (1996, cited by Goswami) in which 3-year olds are shown a green crayon and a yellow crayon covered with a blue filter, so both look green. Children recognise that the green crayon is better for drawing green grass. Although we have no reason to doubt the finding, at least on the basis of her description in the commentary, it does not appear to require processing of ternary relations. The yellow crayon provides one cue for yellow and (when covered with the blue filter) one for green. The green crayon provides two cues for green. Analysing the task in terms of our model in section 6.2.4.3, it does not entail having a cognitive representation of a ternary relation, in which the link between object-attributes and percept is conditional on a third variable, the viewing condition. Rather the percept depends on a single variable, degree-of-greenness. Therefore, the claim that the observation contradicts the relational complexity metric is another case of failing to consider the processes entailed in a task. It has been a persistent error in cognitive development research. Relational complexity analysis, insofar as it can be made on the basis of the brief description,

indicates that the cutting task is structurally simpler and should be performed earlier than theory of mind tasks, precisely what appears to have been observed. Therefore the data should probably be interpreted as *support* for relational complexity theory.

In section 6.2.4.4 we showed that tasks that **Goswami** claims show 3- and 4-year olds process ternary relations (Goswami 1995) could have been performed by unary or binary relations. Goswami does not contest our demonstration, but wants ternary relations to be the default explanation and attempts to justify this by invoking parsimony. This seems to interpret parsimonious as simplistic. The range of phenomena that can be explained by the relational complexity metric, on the basis of a fairly small number of principles, may mean that it is the more parsimonious explanation. Either way, parsimony cannot be used to escape the obligation to provide evidence. This study is presented as though it were a major disconfirmation of our hypothesis, whereas in fact the data are perfectly consistent with our position. An interesting side issue is that Goswami claims her participants processed ternary relations earlier than 5 years because familiar relations were used. The title of the paper (Goswami 1995) specifically attributes the success to the analogy of Goldilocks and the three bears. We are in complete agreement with the importance of familiar, interesting materials when assessing children; in fact, we have used the same analog in our laboratory (Rees 1994). However, we have noticed a tendency to claim that child-appropriate formats overcome all difficulties, *without supporting evidence*. Goswami's claims are another instance of this because the three bears analog did *not* improve performance (Goswami 1995, p. 883). Despite this, the alternative interpretation that children's impressive performances should be attributed to use of simple relations was not considered.

Gentner & Ratterman argue that it is implausible that children could code an ordered three-tuple by chunking into binary relations. We think this underestimates the degree to which chunking is routinely employed by both children and adults. In tasks that entail complex relations, it is normal to chunk components that are not needed for the current decision. These chunks are then unpacked to make further decisions, though, of course, chunked and unchunked relations cannot be processed in the same decision. In this way complex tasks are decomposed into a series of simpler tasks that are performed successively. Furthermore, the chunking process would be facilitated by the visual presentation used by Gentner & Ratterman, because a child can scan the display and identify element(s) about which a decision must be made, chunking the rest. The chunks can then be unpacked to enable further decisions to be made. An ordered three-tuple of the sort used by Gentner & Ratterman can be chunked into binary relational instances *first versus rest* (A,B/C), *middle versus rest* (B,A/C), and *last versus rest* (A/B,C). Chunking in this way would lead to a success rate of approximately 67% correct, which seems consistent with what 3- and 4-year olds can do in mapping ternary relations without additional support. In general, this kind of chunking can be used provided the relations between chunked terms do not need to be processed in the same decision as the rest of the relation. For example, chunking A,B,C into A,B/C, we can process A versus B/C in one decision, then unpack B/C to process B versus C. We cannot do this if relations between all three entities must be processed in one decision. Therefore, as noted in sections

3.4.1 and 6, testing the theory requires careful design to ensure that relational complexity cannot be reduced by chunking and segmentation.

R16. Infancy

Three commentaries – **Goswami, Heath & Hayes**, and **Wright** – suggest that we have not taken account of the large body of literature that indicates impressive cognitive performances in infants. In fact, this literature was considered extensively in developing our position, and we have written some reviews of it (e.g., Halford 1993). However, neither the commentaries nor the papers cited in them provide sufficient evidence that our criteria for relational processing were met. The problem is that, as illustrated above, when the cognitive processes used in a task are unknown we tend to interpret them in the most adult-like way, and this interpretation has often not been supported by fuller investigation. We appreciate the importance of innate cognitive mechanisms for quantification and identification of causes, but we should not assume without evidence that they have the same conceptual structure as later concepts. We cannot review this literature thoroughly here, but we will consider some representative cases.

Heath & Hayes suggest that the ingenious work of Leslie and Keeble (1987) showing that 6-month-old infants recognise cause demonstrates that they understand binary relations. To us adults cause is, at least *prima facie*, a binary relation, *cause(a,b)*. Heath & Hayes then make the common error described above of attributing adult cognition to infants: to adults *cause* is a binary relation, infants recognise cause, therefore, infants appreciate binary relations. In fact, Leslie and Keeble themselves suggest that cause could be recognised by a modular process that is essentially perceptual (see Leslie & Keeble 1987, pp. 283–86). Are we engaging in hypothesis-saving by dismissing recognition of cause as a module? Demonstrably no. In section 2.2.1 we said that relational schemas represent the structure of aspects of the world, and in section 2.2.12 we said they provide a basis for planning and for analogical mapping, and can be modified on-line. The kind of module that Leslie and Keeble postulate can do none of these things, and there seems no reason to believe it would have any of the other properties of relational knowledge defined in section 2.2. In fact, Leslie and Keeble themselves state: “A modular process, though it may be computationally very complex, nevertheless occurs in a fixed, automatic and mechanical way without being influenced by information or reasoning abilities that lie outside the module” (Leslie & Keeble 1987, p. 285). Thus, rather than their data being an embarrassment to us, their position is quite consistent with ours.

In this case the argument is clear because a plausible mechanism has been specified. The mechanism involved in infants' understanding of vanished objects is currently more obscure, but we contend that our position is at least as consistent with the evidence as is **Goswami's**. Our theory does not entail denial of processes such as learning detours or inhibition of responses. We also accept that infancy research in the last two decades has been very successful in assessing competence uncontaminated by performance limitations. However, our position is consistent not only with what infants can do, but also with what they cannot. How, for ex-

ample, does Goswami explain the failure of 10-month-olds to discriminate number of objects on the basis of nonspatial attributes (Xu & Carey 1996)? Furthermore, if infants and young children have such sophisticated relational understanding, why do they have such difficulty with conditional discrimination?

Wright states repeatedly, and incorrectly, that we dismiss the work of Baillargeon. On the contrary, we have long recognised the importance of Baillargeon's observations of infants' object concept (see, e.g., Halford 1989), but we do not think it has been demonstrated to require the kind of explicit relational knowledge that we have defined. We have developed a theory that fits all these observations (Halford 1996b; 1996c) and generates many new predictions. We are not asking anyone to take our arguments on trust, but we strongly believe that they should not be dismissed because of poorly substantiated claims that they are inconsistent with developmental data.

R17. Why processes matter

It seems appropriate to ask whether processes matter. After all, are we not interested in what children can do? Perhaps process models are matters for cognitive science and are irrelevant to cognitive development. Has cognitive development not done well without addressing the complex and difficult question of processes, by defining achievements in terms of observable performances? This view has considerable appeal, if only because it seems to simplify the research task. However, it amounts to a kind of neo-behaviourism, which was found to be unworkable in general cognition research by 1960. By this argument our investigations would go no deeper than concluding that 5-month-olds understand arithmetic, that 3-year-olds, monkeys, and pigeons all understand transitivity (even, apparently, understand it in the same way), and so on. It also means that we would lack objective means to determine whether two tasks measure the same thing. This can cause difficulties if on finding that children cannot perform a particular test we keep simplifying it until we find one they can perform. The problem may be that we are unable to decide whether the simplified test assesses different cognitive processes than the original test. Consequently we might decide that the new test is "right" or "fair" and all the others are deemed to underestimate children's performance, oblivious to the fact that the tests measure different cognitive processes. This fallacy can have a seriously misleading effect on the field (Halford 1989; 1993).

One riposte to this might be that capacity theory, like Piaget, emphasises what children cannot do. We know of no scholarly basis for this criticism of Piaget, and we consider it not valid for capacity theory either. Surely the scientific investigation of cognitive development requires that we adopt an impartial approach, recording children's successes and failures with equal interest, as Piaget did. Then, and only then, will we be in a position to determine the factors that influence performance and how these change with age. However, this has not always been the dominant orientation in the cognitive development literature.

Analyses of cognitive processes increase the information yielded by research paradigms, which can be compared on the basis of the cognitive processes they entail. Consider, for example, the conflicting evidence about capability for

transitive inference. Rather than arguing, as does **Wright**, that one transitivity paradigm is "right" and the others "wrong," we suggest it is more productive to ask why, for example, the transitivity of choice paradigm advocated by **Chalmers & McGonigle** can be performed by a wider range of participants (pigeons to university students) than the paradigm used by Andrews and Halford (in press), which causes great difficulty for children under a median age of 5 years. We suggest the explanation is that the paradigms require different cognitive processes, which make different cognitive demands. The transitivity of choice paradigm can be performed by basically associative processes (Wynne 1995; Harris & McGonigle 1994), whereas the paradigm used by Andrews and Halford requires a representation of relations in the task, and this imposes a load that is a function of the complexity of relations performed in each decision. From this and hundreds of similar comparisons, we have carefully constructed a theory of conceptual complexity, which makes a lot of testable predictions. The essence of the theory is outlined in the target article.

There is a tendency to discount complexity effects, no matter how large they are, because of an overriding concern with precocity. However, there is no reason why these goals – to discover precocious performances and to understand the effects of complexity – need be antagonistic. It is not a matter of emphasising what children can or cannot do, but of recognising that we are dealing with two sides of the coin. A thorough understanding of complexity, and of the way both children and adults deal with it, can be a major benefit in the goal of overcoming limitations.

Frye & Zelazo have demonstrated complexity effects in the important area of concept of mind (Frye et al. 1995). Though their work was originally independent of ours, their effects are consistent with those predicted by Halford (1993) and with the analysis in section 6.2.4.3. Their findings are supported by work in our laboratory (Halford et al. 1998). Frye & Zelazo prefer to model these effects in terms of the cognitive complexity and control (CCC) theory. We have no objection to this because there may be benefits in choosing a formalism that is best suited to a particular set of phenomena. However, it should be borne in mind that a hierarchically structured control process can be expressed as an n -ary relation. This translation is analogous to the way we defined the relational complexity of another hierarchical task, the Tower of Hanoi, in section 6.1.3. The benefit of doing so is that relational complexity provides a metric that is applicable to tasks with any kind of structure. This permits complexities of tasks in different domains, and with different surface structures, hierarchical or otherwise, to be compared directly. The only other response we would make is that we do not believe that, in general, we have neglected the link between cognition and performance. We have modelled the acquisition of performance strategies, guided by the person's concept of the task as well as other factors (Halford et al. 1995).

R18. Complexity metric and interpretation of developmental data

Frye & Zelazo recognised an important benefit of the complexity metric when they commented that "Developmental psychology is commonly recognized as the study of change, but without a method for ordering the changes, the

phenomena become as disorganized as those in the physical sciences would be without a periodic table.” Many of the arguments we have made in earlier sections illustrate the point that without an objective procedure for analysing complexity, the orderly interpretation of developmental data is virtually impossible. If we cannot analyse the complexity of the processes entailed in cognitive performances, how can we decide whether the performance of infants on quantification, cause, or representation of vanished objects is the same as, or simpler than, that of young children? Furthermore, how can we decide whether the transitive inference abilities of older children and adults surpass those of 3-year-olds, or nonhuman primates? In the past this issue has sometimes been resolved simply by attributing the most sophisticated cognitive processes. Without a complexity metric developmental psychology is inevitably locked into unresolvable debates about which test(s) provide the truest or fairest indications of children’s abilities. Continued refinement of assessments is essential, of course, but given that most investigators are well aware of this, considerable benefits can be derived from objective assessments of the complexity of cognitive processes that underlie demonstrated performances.

R19. Conclusion

Complexity effects are very real in the cognition of all higher animals, including human adults and children. A metric is needed that permits the cognitive complexity of tasks to be analysed in such a way that it can be manipulated experimentally, unconfounded by other factors. Relational complexity theory provides such a metric. It has so far shown potential for dealing with species, age, and individual differences. Such a wide-ranging theory inevitably makes contact with a number of other theories. We acknowledge some points of contact, and welcome attempts at constructive integration, but we are not aware of any other theory at present that offers the kind of general conceptual complexity metric that we have proposed, applicable to tasks that can be performed in a single processing step as well as to tasks that entail more than one step. Complexity effects are based on the cognitive processes actually used in performing a task, so valid process analyses are a prerequisite to reliable and testable assessments of complexity.

One consequence, not always recognised by commentators, is that although we accept empirically valid demonstrations of infants’ and children’s performances, their implications for our position depend on analysing the processes they entail. That is, although we accept data about *what* infants and children can do, we believe the field needs to be equally concerned with *how* they do it. Lack of interest in processes seems to have been the main feature that has distinguished cognitive development from general cognition research in the last three decades. Arguably, the cognitive revolution has had much less impact on cognitive development, and this may have been to the detriment of the field. Claims of performances more precocious than our norms appear to be based on assumptions about processes used by infants and children, rather than on evidence that they are really processing relations of a given complexity. Where process analyses have been performed, it has generally been found that the data are quite consistent with our suggested norms.

References

Note: the letters “a” and “r” appearing before author’s initials refer to target article and response respectively.

- Anderson, J. R. (1990) *The adaptive character of thought, vol. 1*. Erlbaum. [rGSH]
 (1993) *Rules of the mind*. Erlbaum. [JRA]
 Anderson, J. R. & Lebiere, C. (1998) *The atomic components of thought*. Erlbaum. [JRA]
 Anderson, J. R., Reder, L. M. & Lebiere, C. (1996) Working memory: Activation limitations on retrieval. *Cognitive Psychology* 30(3):221–56. [JRA, arGSH]
 Anderson, M. (1992) *Intelligence and development: A cognitive theory*. Blackwell. [arGSH]
 Andrews, G. (1996) Assessment of relational reasoning in children aged 4 to 8 years. Paper presented at the Conference of the International Society for Study of Behavioral Development, Quebec City, August 1996. [arGSH]
 Andrews, G. & Halford, G. S. (1994) Relational complexity and sentence processing. *Proceedings of the 21st Annual Experimental Psychology Conference, University of Sydney, vol. 24*. [aGSH]
 (in press) Children’s ability to make transitive inferences: The importance of premise integration and structural complexity. *Cognitive Development*. [arGSH]
 (in preparation) Relational complexity: Evidence for domain generality and age related changes during childhood. [rGSH]
 Astington, J. W. (1993) *The child’s discovery of mind*. Harvard University Press. [DF, aGSH]
 Attneave, F. (1959) *Applications of information theory to psychology: A summary of basic concepts, methods and results*. Henry Holt & Company. [aGSH]
 Babloyantz, A. (1989) Estimation of correlation dimension from single and multichannel recordings. In: *Brain dynamics*, ed. E. Basar & T. H. Bullock. Springer-Verlag. [RB]
 Baddeley, A. D. (1986) *Working memory*. Clarendon Press. [NC, aGSH]
 (1990) *Human memory: Theory and practice*. Allyn & Bacon. [aGSH]
 (1992) Working memory. *Science* 255:556–59. [JAW]
 Baddeley, A. D. & Hitch, G. (1974) Working memory. In: *The psychology of learning and motivation: Advances in research and theory, vol. 8*, ed. G. H. Bower. Academic Press. [aGSH]
 Baillargeon, R. (1986) Representing the existence and the location of hidden objects, object permanence in 6- and 8-month-old infants. *Cognition* 23:21–41. [BW]
 (1987a) Object permanence in 31/2- and 41/2 month-old infants. *Developmental Psychology* 23:655–64. [aGSH]
 (1987b) Young infants’ reasoning about the physical and spatial properties of a hidden object. *Cognitive Development* 2:179–200. [aGSH]
 (1994) Physical reasoning in young infants: Seeking explanations for impossible events. *British Journal of Developmental Psychology* 12:9–33. [UG]
 Baillargeon, R. & Graber, M. (1987) Where is the rabbit? 5.5-month-old infants’ representation of the height of a hidden object. *Cognitive Development* 2:375–92. [UG]
 Baillargeon, R., Graber, M., De Vos, J. & Black, J. (1990) Why do young infants fail to search for hidden objects? *Cognition* 36:255–84. [UG]
 Baillargeon, R., Kotovsky, L. & Needham, A. (1995) The acquisition of physical knowledge in infancy. In: *Causal cognition: A multidisciplinary debate*, ed. D. Sperber, D. Premack & A. J. Premack. Oxford University Press. [RAH]
 Baum, E. B., Moody, J. & Wilczek, F. (1988) Internal representations for associative memory. *Biological Cybernetics* 59:217–28. [TAP]
 Bitterman, M. E. (1960) Toward a comparative psychology of learning. *American Psychologist* 15:704–12. [aGSH]
 (1975) The comparative analysis of learning. *Science* 188:699–709. [aGSH]
 Borisyuk, G., Borisyuk, R., Khibnik, A. & Roose, D. (1995) Dynamics and bifurcations of two coupled neural oscillators with different connection types. *Bulletin of Mathematical Biology* 57:809–40. [RB]
 Borisyuk, R., Casaleggio, A., Kazanovich, Y. & Morgavi, G. (1994) Some results on the analysis of dimension of time series by a network of phase oscillators. In: *Proceedings of European Conference on Artificial Neural Networks ICANN’94, vol. 1*, ed. M. Marinaro & P. Morasso. Springer-Verlag [RB]
 Borisyuk, R. & Hoppensteadt, F. (1998) Memorizing and recalling spatial-temporal patterns in an oscillator model of the hippocampus. *BioSystems* (in press). [RB]
 Breslow, L. (1981) Reevaluation of the literature on the development of transitive inferences. *Psychological Bulletin* 89:325–51. [aGSH]
 Broadbent, D. E. (1975) The magic number seven after fifteen years. In: *Studies in long-term memory*, ed. A. Kennedy & A. Wilkes. Wiley. [aGSH]
 Brown, A. L. (1989) Analogical learning and transfer: What develops? In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Orlony. Cambridge University Press. [DG]
 Bryant, P. E. & Trabasso, T. (1971) Transitive inferences and memory in young children. *Nature* 232:456–58. [arGSH]

- Carpenter, P. A. & Just, M. A. (1996) Functional and neuroscience approaches to working memory. *International Journal of Psychology* 31:325. [aGSH]
- Case, R. (1985) *Intellectual development: Birth to adulthood*. Academic Press. [aGSH]
- (1992) *The mind's staircase: Exploring the conceptual underpinnings of children's thought and knowledge*. Erlbaum. [aGSH]
- Chalmers, M. & McGonigle, B. (1984) Are children any more logical than monkeys on the five-term series problem? *Journal of Experimental Child Psychology* 37:355–77. [rGSH]
- (1997) Capturing dynamic structuralism in the laboratory. In: *Piaget, Vygotsky and beyond*, ed. L. Smith, J. Dockrell & P. Tomlinson. Routledge. [MC]
- Chapman, M. (1987) Piaget, attentional capacity, and the functional limitations of formal structure. *Advances in Child Development and Behaviour* 20:289–334. [arGSH]
- Chase, W. G. & Simon, H. A. (1973) Perception in chess. *Cognitive Psychology* 4:55–81. [JS]
- 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. [aGSH]
- Clark, A. & Thornton, C. (1997) Trading spaces: Computation, representation and the limits of uninformed learning. *Behavioral and Brain Sciences* 20:57–90. [MC]
- Clark, H. H. (1969) Linguistic processes in deductive reasoning. *Psychological Review* 76:387–404. [MC]
- Codd, E. F. (1990) *The relational model for database management: Version 2*. Addison-Wesley. [aGSH]
- Cohen, J. D., Perlstein, W. M., Braver, T. S., Nystrom, L. E., Noll, D. C., Jonides, J. & Smith, E. E. (1997) Temporal dynamics of brain activation during a working memory task. *Nature* 386:604–07. [JAW]
- Cowan, N. (1992) Verbal memory span and the timing of spoken recall. *Journal of Memory and Language* 31:668–84. [NC]
- (1995) *Attention and memory: An integrated framework*. Oxford University Press. [NC]
- Cowan, N., Keller, T., Hulme, C., Roodenrys, S., McDougall, S. & Rack, J. (1994) Verbal memory span in children: Speech timing clues to the mechanisms underlying age and word length effects. *Journal of Memory and Language* 33:234–50. [NC]
- Cowan, N., Wood, N. L., Wood, P. K., Keller, T. A., Nugent, L. D. & Keller, C. V. (1998) Two separate verbal processing rates contributing to short-term memory span. *Journal of Experimental Psychology: General* 127:141–60. [NC]
- Cummins, D. D., Kintsch, W., Reussler, K. & Weimer, R. (1988) The role of understanding in solving word problems. *Cognitive Psychology* 20(4):405–38. [aGSH]
- Cutting, A. L. (1996) Young children's understanding of representation: A problem solving approach to the appearance-reality distinction. Unpublished Ph. D. dissertation. University of Cambridge. [UG]
- Damasio, A. R. (1989) The brain binds entities and events by multiregional activation from convergent zones. *Neural Computation* 1:123–32. [RB]
- Davis, H. L. & Pratt, C. (1995) The development of children's theory of mind: The working memory explanation. *Australian Journal of Psychology* 47(1):25–31. [aGSH]
- de Ribaupierre, A. & Pascual-Leone, J. (1979) Formal operations and M power: A neo-Piagetian investigation. In: *Intellectual development beyond childhood*, ed. D. Kuhn. (Source books on New Directions in Child Development). Jossey-Bass. [JPL]
- D'Esposito, M., Betre, J. A., Alsop, D. C., Shin, R. K., Atlas, S. & Grossman, M. (1996) The neural basis of the central executive system of working memory. *Nature* 378:279–81. [JAW]
- Diamond, A. (1988) Differences between adult and infant cognition: Is the crucial variable presence or absence of language? In: *Thought without language*, ed. L. Weiskrantz. Clarendon Press. [UG]
- English, L. D. & Halford, G. S. (1995) *Mathematics education: Models and processes*. Erlbaum. [aGSH]
- Fischer, K. W. (1980) A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review* 87:477–531. [DC, aGSH]
- Fischer, K. W. & Bidell, T. R. (1997) Dynamic development of psychological structures in action and thought. In: *Handbook of child psychology, vol. 1: Theoretical models of human development*, ed. R. M. Lerner & W. Damon. Wiley. [DC]
- Fischer, K. W. & Rose, S. P. (1996) Dynamic growth cycles of brain and cognitive development. In: *Developmental neuroimaging: Mapping the development of brain and behavior*, ed. R. Thatcher, G. R. Lyon, J. Rumsey & N. Krasnegor. Academic Press. [DC]
- Fisher, D. L. (1984) Central capacity limits in consistent mapping, visual search tasks: Four channels or more? *Cognitive Psychology* 16(4):449–84. [aGSH]
- Flavell, J. H., Green, F. L. & Flavell, E. R. (1990) Developmental changes in young children's knowledge about the mind. *Cognitive Development* 5(1):1–27. [aGSH]
- Fodor, J. A. (1983) *Modularity of mind: An essay on faculty psychology*. MIT Press. [aGSH]
- Fodor, J. A. & Pylyshyn, Z. W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28:3–71. [aGSH]
- Foley, E. J. (1997) Assessing conceptual complexity in hierarchical reasoning: A dual-task approach. Unpublished doctoral dissertation, University of Cincinnati, Ohio. [DBB, rGSH]
- Foley, E. J. & Berch, D. B. (1997) Capacity limitations of a classic m-power measure: A modified dual-task approach. *Journal of Experimental Child Psychology* 66:129–43. [DBB, rGSH]
- Fry, A. F. & Hale, S. (1996) Processing speed, working memory, and fluid intelligence: Evidence for a developmental cascade. *Psychological Science* 7:237–41. [NC]
- Frye, D. (in press) Development of intention: The relation of executive function to theory of mind. In: *Developing theories of intention: Social understanding and self control*, ed. P. D. Zelazo, J. W. Astington & D. R. Olson. Erlbaum. [DF]
- Frye, D., Zelazo, P. D., Brooks, P. J. & Samuels, M. C. (1996) Inference and action in early causal reasoning. *Developmental Psychology* 32:120–31. [DF]
- Frye, D., Zelazo, P. D. & Palfai, T. (1995) Theory of mind and rule-based reasoning. *Cognitive Development* 10:483–527. [DF, arGSH]
- Gallistel, C. R. (1990) Representations in animal cognition: An introduction. *Cognition* 37:1–22. [aGSH]
- Garey, M. R. & Johnson, D. S. (1979) *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman. [aGSH, TAP]
- Gell-Mann, M. (1994) *The quark and the jaguar*. W. H. Freeman. [JP-L]
- Gentner, D. (1983) Structure-mapping: A theoretical framework for analogy. *Cognitive Science* 7:155–70. [DG, aGSH]
- (1988) Metaphor as structure mapping: The relational shift. *Child Development* 59:47–59. [DG]
- Gentner, D. & Markman, A. M. (1997) Structural alignment in analogy and similarity. *American Psychologist* 52:45–56. [DG]
- Gentner, D. & Rattermann, M. J. (1991) Language and the career of similarity. In: *Perspectives on thought and language: Interrelations in development*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [DG]
- Gentner, D. & Toupin, C. (1986) Systematicity and surface similarity in the development of analogy. *Cognitive Science* 10:277–300. [DG]
- Gergely, G., Nadasdy, Z., Csibra, G. & Biro, S. (1995) Taking the intentional stance at 12 months of age. *Cognition* 56: 165–93. [UG]
- Gick, M. L. & Holyoak, K. J. (1983) Schema induction and analogical transfer. *Cognitive Psychology* 15:1–38. [aGSH]
- Gopher, D. (1986) In defence of resources: On structures, energies, pools and the allocation of attention. In: *Energetics and human information processing*, ed. G. R. J. Hockey, A. W. K. Gaillard & M. G. H. Coles. Martinus Nijhoff. [DN]
- (1994) Analysis and measurement of mental load. In: *The state of the art, vol. 2*, ed. G. d'Ydewalle, P. Eelen & P. Bertelson. Erlbaum. [aGSH]
- Goswami, U. (1992) *Analogical reasoning in children*. Erlbaum. [aGSH]
- (1995) Transitive relational mappings in 3- and 4-year-olds: The analogy of Goldilocks and the three bears. *Child Development* 66:877–92. [DG, UG, arGSH]
- (1996) Analogical reasoning and cognitive development. *Advances in Child Development and Behaviour* 26:91–138. [UG]
- (1998) *Cognition in children*. Psychology Press. [UG]
- Goswami, U. & Brown, A. L. (1989) Melting chocolate and melting snowmen: Analogical reasoning and causal relations. *Cognition* 35:69–95. [rGSH]
- Goswami, U., Leevers, H., Pressley, S. & Wheelwright, S. (in press) Causal reasoning about pairs of relations and analogical reasoning in young children. *British Journal of Developmental Psychology*. [UG]
- Granott, N., Fischer, K. W. & Parziale, J. (1997) Bridging to the unknown: A fundamental mechanism in learning and problem-solving. Cognitive Development Laboratory Research Report, Harvard University. [DC]
- Greeno, J. G., Riley, M. S. & Gelman, R. (1984) Conceptual competence and children's counting. *Cognitive Psychology* 16:94–143. [aGSH]
- Grieve, R. & Garton, A. (1981) On the young child's comparison of sets. *Journal of Experimental Child Psychology* 32:443–58. [aGSH]
- Halford, G. S. (1971) Acquisition of conservation through learning a consistent classificatory system for quantities. *Australian Journal of Psychology* 23:151–59. [aGSH]
- (1980) A learning set approach to multiple classification: Evidence for a theory of cognitive levels. *International Journal of Behavioral Development* 3:409–22. [arGSH]
- (1982) *The development of thought*. Erlbaum. [arGSH]
- (1984) Can young children integrate premises in transitivity and serial order tasks? *Cognitive Psychology* 16:65–93. [arGSH, BW]
- (1989) Reflections on 25 years of Piagetian cognitive developmental psychology, 1963–88. *Human Development* 32:325–87. [arGSH]

- (1992) Analogical reasoning and conceptual complexity in cognitive development. *Human Development* 35:193–217. [aGSH]
- (1993) *Children's understanding: The development of mental models*. Erlbaum. [DBB, DG, arGSH]
- (1995) Learning processes in cognitive development: A reassessment with some unexpected implications. *Human Development* 38(6):295–301. [rGSH]
- (1996a) Children's understanding of the mind: An instance of a general principle? [Review of *The child's discovery of the mind*]. *Contemporary Psychology* 41:229–30. [aGSH]
- (1996b) Capacity limitations in processing relations: Implications and causes. Paper presented at IAS 3rd Brain and Mind International Symposium on Concept Formation, Thinking and their Development, Kyoto, Japan. [rGSH]
- (1996c) Relational knowledge in higher cognitive processes. Paper presented in symposium entitled Relational Knowledge in Higher Cognitive Processes at the XIVth Biennial Meeting of the International Society for the Study of Behavioral Development, Quebec City, August 12–16, 1996. [rGSH]
- Halford, G. S., Andrews, G. & Bowden, D. (1998a) Complexity as a factor in children's theory of mind. Poster session presented at the XVth Biennial Meetings of the ISSBD, Berne, Switzerland, July 1998. [rGSH]
- Halford, G. S., Bain, J. D. & Maybery, M. T. (1984) Does a concurrent memory load interfere with reasoning? *Current Psychological Research and Reviews* 3:14–23. [aGSH]
- Halford, G. S., Bain, J. D., Maybery, M. T. & Andrews, G. (1998b) Induction of relational schemas: Common processes in reasoning and learning set acquisition. *Cognitive Psychology*, 35:201–45 [rGSH]
- Halford, G. S. & Dalton, C. (1995) Performance on the balance scale by 2-year-old children. ERIC Document Reproduction Service No. ED 385 355. [aGSH]
- Halford, G. S. & Fullerton, T. (1970) A discrimination task which induces conservation of number. *Child Development* 41:205–13. [aGSH]
- Halford, G. S. & Kelly, M. E. (1984) On the basis of early transitivity judgements. *Journal of Experimental Child Psychology* 38:42–63. [aGSH, BW]
- Halford, G. S. & Leitch, E. (1989) Processing load constraints: A structure-mapping approach. In: *Psychological development: Perspectives across the life-span*, ed. M. A. Luszcz & T. Nettelbeck. North-Holland. [arGSH]
- Halford, G. S. & MacDonald, C. (1977) Children's pattern construction as a function of age and complexity. *Child Development* 48:1096–1100. [rGSH]
- Halford, G. S., Maybery, M. T. & Bain, J. D. (1986) Capacity limitations in children's reasoning: A dual task approach. *Child Development* 57:616–27. [arGSH]
- (1988) Set-size effects in primary memory: An age-related capacity limitation? *Memory and Cognition* 16(5):480–87. [arGSH]
- Halford, G. S., Maybery, M. T., O'Hare, A. W. & Grant, P. (1994) The development of memory and processing capacity. *Child Development* 65:1338–56. [RB, aGSH]
- 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: *Developing cognitive competence: New approaches to cognitive modelling*, ed. T. Simon & G. S. Halford. Erlbaum. [DG, arGSH]
- Halford, G. S. & Wilson, W. H. (1980) A category theory approach to cognitive development. *Cognitive Psychology* 12:356–411. [arGSH]
- Halford, G. S., Wilson, W. H., Gray, B. & Phillips, S. (1997) A neural net model for mapping hierarchically structured analogs. *Proceedings of the Fourth Conference of the Australasian Cognitive Science Society, September 1997, University of Newcastle*. [aGSH]
- 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: *Advances in connectionist and neural computation theory, vol. 2: Analogical connections*, ed. K. J. Holyoak & J. Barnden. Ablex. [aGSH]
- Halford, G. S., Wilson, W. H. & McDonald, M. (1995) Complexity of structure mapping in human analogical reasoning: A PDP model. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society, July 1995, Pittsburgh, Pennsylvania*. [aGSH]
- Halford, G. S., Wilson, W. H. & Phillips, S. (1996) Human analogical reasoning capacity: Toward a neural net model. *International Journal of Psychology* 31:443. [aGSH]
- Harper, M. (1996) Complexity as a factor in children's performance on the balance scale. Unpublished Honours Thesis, University of Queensland, Brisbane, Australia. [aGSH]
- Harris, M. R. & McGonigle, B. O. (1994) A model of transitive choice. *The Quarterly Journal of Experimental Psychology* 47B(3):319–48. [rGSH]
- Heath, R. A. (1991) A nonlinear model for human associative memory based on system identification. In: *Proceedings of the Second Australian Conference on Neural Networks*, ed. M. Jabri. Sydney University Electrical Engineering. [RAH]
- (in press) From simple processes does complex behaviour emerge: A methodology for cognitive research based on nonlinear system theory. In: *Perspectives on cognitive science II: Theories, experiments, and foundations*, ed. T. Dartnell & J. Wiles. Ablex Press. [RAH]
- Heath, R. A. & Fulham, R. (1988) An adaptive filter model for recognition memory. *British Journal of Mathematical and Statistical Psychology* 41:119–44. [RAH]
- Henderson, J. (1994) Connectionist syntactic parsing using temporal variable binding. *Journal of Psycholinguistic Research* 23(5):353–79. [arGSH]
- Hinton, G. E. (1990) Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence* 46:47–76. [TAP]
- Hitch, G. J. (1980) *Developing the concept of working memory*. Routledge & Kegan Paul. [aGSH]
- Hodkin, B. (1987) Performance model analysis in class inclusion: An illustration with two language conditions. *Developmental Psychology* 23:683–89. [aGSH]
- Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986) *Induction: Processes of inference, learning and discovery*. Bradford Books/MIT Press. [aGSH]
- Holyoak, K. J. & Hummel, J. E. (in press) The proper treatment of symbols in a connectionist architecture. In: *Cognitive dynamics: Conceptual change in humans and machines*, ed. E. Dietrich & A. Markman. MIT Press. [JAW]
- Holyoak, K. J. & Thagard, P. (1989) Analogical mapping by constraint satisfaction. *Cognitive Science* 13(3):295–355. [aGSH]
- (1995) *Mental leaps: Analogy in creative thought*. MIT Press. [aGSH]
- Howe, M. L. & Rabinowitz, F. M. (1994) Dynamic modeling, chaos, and cognitive development. *Journal of Experimental Child Psychology* 58:184–99. [RAH]
- Hummel, J. E. & Holyoak, K. J. (1997) Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review* 104:427–66. [arGSH, DNI, JAW]
- 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):208–33. [aGSH]
- Humphreys, M. S., Pike, R., Bain, J. D. & Tehan, G. (1989) Global matching: A comparison of the SAM, Minerva II, Matrix, and TODAM models. *Journal of Mathematical Psychology* 33:36–67. [RAH]
- Hunt, E. & Lansman, M. (1982) Individual differences in attention. In: *Advances in the psychology of human intelligence*, ed. R. J. Sternberg. Erlbaum. [DBB, rGSH]
- James, W. (1890) *Principles of psychology*. Holt, Rinehart & Winston. [aGSH]
- Johnson-Laird, P. N., Byrne, R. M. J. & Schaeken, W. (1992) Propositional reasoning by model. *Psychological Review* 99:418–39. [aGSH]
- Just, M. A. & Carpenter, P. A. (1992) A capacity theory of comprehension: Individual differences in working memory. *Psychological Review* 99(1):122–49. [arGSH]
- Just, M. A., Carpenter, P. A. & Hemphill, D. D. (1996) Constraints on processing capacity: Architectural or implementational? In: *Mind matters: A tribute to Allen Newell*, ed. D. Steier & T. Mitchell. Erlbaum. [aGSH]
- Kahneman, D. (1973) *Attention and effort*. Prentice-Hall. [aGSH, DN]
- Kail, R. (1991) Processing time declines exponentially during childhood and adolescence. *Developmental Psychology* 27(2):259–66. [rGSH, RAH]
- Kail, R. & Park, Y.-S. (1992) Global developmental change in processing time. *Merrill-Palmer Quarterly* 38:525–41. [RAH]
- (1994) Processing time, articulation time, and memory span. *Journal of Experimental Child Psychology* 57:281–91. [NC]
- Kail, R. & Salthouse, T. A. (1994) Processing speed as a mental capacity. *Acta Psychologica* 86:199–255. [NC]
- Kallio, K. D. (1982) Developmental change on a five-term transitive inference. *Journal of Experimental Child Psychology* 33:142–64. [aGSH]
- Kanerva, P. (1988) *Sparse distributed memory*. MIT Press. [TAP]
- (1996) Binary spatter-coding of ordered k-tuples. In: *Artificial neural networks – ICANN Proceedings, Berlin; Lecture notes in computer science, vol. 1112*, ed. C. von der Malsburg, W. von Seelen, J. Vorbruggen & B. Sendhoff. Springer. [TAP]
- Kazanovich, Y. B. & Borisuyk, R. M. (1994) Synchronization in a neural network with a central element. *Biological Cybernetics* 71:177–85. [RB]
- Kimball, J. (1973) Seven principles of surface structure parsing in natural language. *Cognition* 2:15–47. [aGSH]
- Kintsch, W. & Greeno, J. G. (1985) Understanding and solving word arithmetic problems. *Psychological Review* 92(1):109–29. [aGSH]
- Klahr, D. & Wallace, J. G. (1976) *Cognitive development: An information processing view*. Erlbaum. [aGSH]
- Klapp, S. T., Marshburn, E. A. & Lester, P. T. (1983) Short-term memory does not involve the “working memory” of information processing: The demise of a common assumption. *Journal of Experimental Psychology: General* 112:240–64. [aGSH]
- Kotovsky, K., Hayes, J. R. & Simon, H. A. (1985) Why are some problems hard? Evidence from the Tower of Hanoi. *Cognitive Psychology* 17:248–94. [JS]

- Kryukov, V. I. (1991) An attention model based on the principle of dominance. In: *Neurocomputers and attention. Neurobiology, synchronization and chaos*, ed. A. V. Holden & V. I. Kryukov. Manchester University Press. [RB]
- Lamborn, S. D., Fischer, K. W. & Pipp, S. L. (1994) Constructive criticism and social lies: A developmental sequence for understanding honesty and kindness in social interactions. *Developmental Psychology* 30:495–508. [DC]
- Lansman, M. & Hunt, E. (1982) Individual differences in secondary task performance. *Memory and Cognition* 10:10–24. [DBB]
- Lebedev, A. N. (1980) A mathematical model for human visual information perception and storage. In: *Neural mechanisms of goal directed behavior and learning*, ed. R. F. Thompson, L. H. Hicks & V. B. Shvyrkov. Academic Press. [DNi]
- Lebiere, C. (in preparation) *An ACT-R model of cognitive arithmetic*. Ph. D. Thesis. [JRA]
- Lebiere, C. & Anderson, J. R. (1993) A connectionist implementation of the ACT-R production system. *Proceedings of the Fifteenth Annual Meeting of the Cognitive Science Society*, 635–40. [JRA]
- Leeuwenberg, E. L. L. (1969) Quantitative specification of information in sequential patterns. *Psychological Review* 76:216–20. [rGSH]
- Leslie, A. & Keeble, S. (1987) Do six-month-old infants perceive causality? *Cognition* 25:265–88. [rGSH, RAH]
- Logan, G. D. (1979) On the use of a concurrent memory load to measure attention and automaticity. *Journal of Experimental Psychology: Human Perception and Performance* 5:189–207. [aGSH]
- Loveday, W. (1995) The effect of complexity on planning in the Tower of Hanoi problem. Unpublished Honours Thesis, University of Queensland, Australia. [aGSH]
- Lovett, M. C., Reder, L. M. & Lebiere, C. (1997) Modeling individual differences in a digit working memory task. *Proceedings of the Nineteenth Conference of the Cognitive Science Society*, 460–65. [JRA]
- (in press) Modeling working memory in a unified architecture: An ACT-R perspective. In: *Models of working memory: Mechanisms of active maintenance and executive control*, ed. A. Miyake & P. Shah. Cambridge University Press. [JRA]
- MacLeod, W. B. (1996) Decision, contract and emotion: Some economics for a complex and confusing world. *The Canadian Journal of Economics* 26:788–810. [WBM]
- Markman, A. B. & Gentner, D. (1993) Structural alignment during similarity comparisons. *Cognitive Psychology* 25:431–67. [DG]
- Markman, E. M. & Seibert, J. (1976) Classes and collections: Internal organization and resulting holistic properties. *Cognitive Psychology* 8:561–77. [aGSH]
- Markovits, H., Dumas, C. & Malfait, N. (1995) Understanding transitivity of a spatial relationship: A developmental analysis. *Journal of Experimental Child Psychology* 59:124–41. [BW]
- Maybery, M. T., Bain, J. D. & Halford, G. S. (1986) Information processing demands of transitive inference. *Journal of Experimental Psychology: Learning, Memory and Cognition* 12:600–13. [arGSH]
- McClelland, J. L. (1995) A connectionist perspective on knowledge and development. In: *Developing cognitive competence: New approaches to cognitive modelling*, ed. T. Simon & G. S. Halford. Erlbaum. [aGSH]
- McGarrigle, J., Grieve, R. & Hughes, M. (1978) Interpreting inclusion: A contribution to the study of the child's cognitive and linguistic development. *Journal of Experimental Child Psychology* 26:528–50. [aGSH]
- McGonigle, B. O. & Chalmers, M. (1977) Are monkeys logical? *Nature* 267:694–96. [MC]
- (1986) Representations and strategies during inference. In: *Reasoning and discourse processes*, ed. T. Myers, K. Brown & B. O. McGonigle. Academic Press. [MC]
- (1992) Monkeys are rational! *The Quarterly Journal of Experimental Psychology* 45B(3):189–228. [MC]
- (1996) The ontology of order. In: *Critical readings on Piaget*, ed. L. Smith. Routledge. [MC]
- (1998) Rationality as optimised cognitive self-regulation. In: *Rational models of cognition*, ed. M. Oaksford & N. Chater. Oxford University Press. [MC]
- McGonigle, B. O. & Jones, B. T. (1978) Levels of stimulus processing by the squirrel monkey: Relative and absolute judgements compared. *Perception* 7:635–59. [rGSH]
- McKenzie, B. E. & Bigelow, E. (1986) Detour behaviour in young human infants. *British Journal of Developmental Psychology* 4:139–48. [UG]
- McNicol, D. & Stewart, G. W. (1980) *Reaction time and the study of memory*. Academic Press. [rGSH]
- Meltzoff, A. N. (1996) Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Annual Progress in Child Psychiatry and Child Development* 67–98. [RAH]
- Miller, G. A. (1956) The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63:81–97. [JRA, aGSH, DN]
- Miller, R. (1991) *Cortico-hippocampal interplay*. Springer-Verlag. [RB]
- Mitra, S., Amazeen, P. G. & Turvey, M. T. (1998) Intermediate motor learning as decreasing active (dynamical) degrees of freedom. *Human Movement Science* 17:17–66. [DNi]
- Mitra, S., Riley, M. A. & Turvey, M. T. (1997) Chaos in human rhythmic movement. *Journal of Motor Behavior* 29:195–98. [DNi]
- Murdock, B. B. (1983) A distributed memory model for serial-order information. *Psychological Review* 90:316–38. [rGSH]
- Nassefat, M. (1963) *Etude quantitative sur l'évolution des opérations intellectuelles*. These No. 12. Institute des Sciences de l'Éducation. Université de Genève. Delachaux and Niestlé. [JP-L]
- Navon, D. (1984) Resources: A theoretical soup stone? *Psychological Review* 91:216–34. [rGSH, DN]
- (1985) Attention division or attention sharing? In: *Attention and performance XI*, ed. M. I. Posner & O. S. M. Marin. Erlbaum. [DN]
- Navon, D. & Gopher, D. (1979) On the economy of the human information processing system. *Psychological Review* 86:214–55. [rGSH, DN]
- (1980) Task difficulty, resources and dual task performance. In: *Attention and performance, vol. 8*, ed. R. S. Nickerson. Erlbaum. [aGSH, DN]
- Navon, D. & Miller, J. O. (1987) The role of outcome conflict in dual-task interference. *Journal of Experimental Psychology: Human Perception and Performance* 13:435–48. [DN]
- (1998) How single is the single bottleneck? (submitted). [DN]
- Needham, A. & Baillargeon, R. (1993) Intuitions about support in 4.5-month-old infants. *Cognition* 47:121–48. [RAH]
- Nikolic, D. (in press) A dual processing theory of brain and mind: Where is the limited processing capacity coming from? *InterJournal of Complex Systems*. [DNi]
- Norman, D. A. & Bobrow, D. J. (1975) On data-limited and resource-limited processes. *Cognitive Psychology* 7:44–64. [DN]
- Pascual-Leone, J. A. (1970) A mathematical model for the transition rule in Piaget's developmental stages. *Acta Psychologica* 32:301–45. [aGSH, JP-L]
- (1978) Computational models for metasubjective processes: Commentary to Z. W. Pylyshyn's "Computational models and empirical constraints." *Behavioral and Brain Sciences* 1:112–13. [JP-L]
- (1994) An essay review of Halford's *An experimentalist's understanding of children*. *Human Development* 37:370–84. [JP-L]
- Pascual-Leone, J. & Baillargeon, R. (1994) Developmental measurement of mental attention. *International Journal of Behavioral Development* 17(1):161–200. [JP-L]
- Pascual-Leone, J. & Goodman, D. (1979) Intelligence and experience: A neo-Piagetian approach. *Instructional Science* 8:301–67. [JP-L]
- Pascual-Leone, J. & Smith, J. (1969) The encoding and decoding of symbols by children: A new experimental paradigm and neo-Piagetian model. *Journal of Experimental Child Psychology* 8:328–55. [JP-L]
- Pears, R. & Bryant, P. (1990) Transitive inferences by young children about spatial position. *British Journal of Psychology* 81(4):497–510. [aGSH]
- Phillips, S. & Halford, G. S. (1997) Systematicity: Psychological evidence with connectionist implications. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society, Stanford University*, 614–19. [rGSH]
- Phillips, S., Halford, G. S. & Wilson, W. H. (1995) The processing of association versus the processing of relations and symbols: A systematic comparison. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society, Pittsburgh, PA*, 688–91. [arGSH]
- (1998) Representational redescription: From associative to relational systems. (submitted). [aGSH]
- Piaget, J. (1950) *The psychology of intelligence*, trans. M. Piercy & D. E. Berlyne. Routledge & Kegan Paul. [Original published in 1947]. [aGSH]
- Plate, T. A. (1994) *Distributed representations and nested compositional structure*. Ph. D. Thesis, Department of Computer Science, University of Toronto. Available at <http://mcs.vuw.ac.nz/~tap>. [aGSH, TAP]
- (1995) Holographic reduced representations. *IEEE Transactions on Neural Networks* 6(3):623–41. [aGSH, TAP]
- (in press) Estimating analogical similarity by vector dot-products of holographic reduced representations. *Cognitive Science*. [aGSH]
- Pollack, J. B. (1988) Recursive auto-associative memory: Devising compositional distributed representations. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, 33–39. [aGSH]
- Posner, M. I. & Boies, S. J. (1971) Components of attention. *Psychological Review* 78:391–408. [aGSH]
- Potts, G. R. (1972) Information processing strategies used in the encoding of linear orderings. *Journal of Verbal Learning and Verbal Behaviour* 11:727–40. [MC]
- Premack, D. (1983) The codes of man and beasts. *Behavioral and Brain Sciences* 6:125–67. [aGSH]
- Quinn, P. C. & Johnson, M. H. (1997) The emergence of perceptual category representations in young infants: A connectionist analysis. *Journal of Experimental Child Psychology* 66:236–63. [rGSH]

- Rattermann, M. J. & Gentner, D. (in preparation) The effect of relational language on children's performance in an analogical mapping task. [DC]
- Rees, B. E. (1994) Children's ability to map ordered sets. Unpublished Honours Thesis, University of Queensland, Brisbane. [rGSH]
- Rieser, J. J., Doxey, P. A., McCarrell, N. J. & Brooks, P. H. (1982) Wayfinding and toddlers' use of information from an aerial view of a maze. *Developmental Psychology* 18:714–20. [UG]
- Riley, M. A. (1997) Personal communication. [DNI]
- Rips, L. J. (1989) The psychology of knights and knaves. *Cognition* 31:85–116. [rGSH]
- Robin, N. & Holyoak, K. J. (1995) Relational complexity and the functions of prefrontal cortex. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [arGSH, JAW]
- Salthouse, T. A. (1996) The processing-speed theory of adult age differences in cognition. *Psychological Review* 103:403–28. [NC]
- Schneider, W. & Detweiler, M. (1987) A connectionist/control architecture for working memory. *Psychology of Learning and Motivation* 21:53–119. [aGSH]
- Schweickert, R. & Boruff, B. (1986) Short-term memory capacity: Magic number or magic spell? *Journal of Experimental Psychology: Learning, Memory and Cognition* 12:419–25. [NC, aGSH]
- Shallice, T. & Burgess, P. (1991) Higher-order cognitive impairments and frontal lobe lesions in man. In: *Frontal lobe function and dysfunction*, ed. H. S. Levin, H. M. Eisenberg & A. L. Benton. Oxford University Press. [JAW]
- Shastri, L. & Ajanagadde, V. (1993a) From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* 16(3):417–94. [RB, aGSH, DNI]
- (1993b) A step toward modeling reflexive reasoning. *Behavioral and Brain Sciences* 16(3):477–88. [aGSH]
- Shiple, E. F. (1979) The class-inclusion task: Question form and distributive comparisons. *Journal of Psycholinguistic Research* 8:301–31. [aGSH]
- Siegler, R. S. (1981) Developmental sequences within and between concepts. *Monographs of the Society for Research in Child Development* 46:1–84. [aGSH]
- Simon, H. (1974) How big is a chunk? *Science* 183:482–88. [TAP]
- Smith, L. B. (1989) From global similarities to kinds of similarities: The construction of dimensions in development. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony. Cambridge University Press. [aGSH]
- Smolensky, P. (1990) Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence* 46:159–216. [RB, aGSH, TAP]
- Sternberg, R. J. (1980) Representation and process in linear syllogistic reasoning. *Journal of Experimental Psychology: General* 109:119–59. [aGSH]
- Sternberg, S. (1975) Memory scanning: New findings and current controversies. *Quarterly Journal of Experimental Psychology* 27:1–32. [rGSH]
- Sugarman, S. (1982) Developmental change in early representational intelligence: Evidence from spatial classification strategies and related verbal expressions. *Cognitive Psychology* 14:410–49. [aGSH]
- Surber, C. F. & Gzesh, S. M. (1984) Reversible operations in the balance scale task. *Journal of Experimental Child Psychology* 38:254–74. [aGSH]
- Sweller, J. (1993) Some cognitive processes and their consequences for the organisation and presentation of information. *Australian Journal of Psychology* 45(1):1–8. [aGSH]
- Sweller, J., Mawer, R. & Ward, M. (1983) Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General* 112:634–56. [JS]
- Tesar, B. B. & Smolensky, P. (1994) Synchronous firing variable binding is a tensor product representation with temporal role vectors. *16th Annual Conference of the Cognitive Science Society, Atlanta, Georgia*, 870–75. [aGSH]
- Thatcher, R. W. (1994) Cyclic cortical reorganization: Origins of human cognitive development. In: *Human behavior and the developing brain*, ed. G. Dawson & K. W. Fischer. Guilford. [DC, rGSH]
- Thayer, E. S. & Collyer, C. E. (1978) The development of transitive inference: A review of recent approaches. *Psychological Bulletin* 85:1327–43. [aGSH]
- Thelen, E. & Ulrich, B. D. (1991) Hidden skills: A dynamic systems analysis of treadmill stepping during the first year. *Monographs of the Society for Research in Child Development* 56 (No.1, Serial No. 223). [RAH]
- Tindall-Ford, S., Chandler, P. & Sweller, J. (1997) When two sensory modes are better than one. *Journal of Experimental Psychology: Applied* 3:257–87. [JS]
- Touretsky, D. S. (1990) BoltzCONS: Dynamic symbol structures in a connectionist network. *Artificial Intelligence* 46:5–46. [aGSH]
- Townsend, J. T. (1990) Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science* 1:46–54. [RAH]
- Trabasso, T. (1975) *Representation, memory, and reasoning: How do we make transitive inferences?* University of Minnesota Press. [aGSH]
- (1977) The role of memory as a system in making transitive inferences. In: *Perspectives on the development of memory and cognition*, ed. R. V. Kail & J. W. Hagen. Erlbaum. [aGSH]
- Trabasso, T., Isen, A. M., Dolecki, P., McLanahan, A. G., Riley, C. A. & Tucker, T. (1978) How do children solve class-inclusion problems? In: *Children's thinking: What develops?*, ed. R. Siegler. Erlbaum. [aGSH]
- Tsotsos, J. K. (1990) Analyzing vision at the complexity level. *Behavioral and Brain Sciences* 13:423–69. [aGSH]
- Van der Maas, H. L. J. & Molenaar, C. M. (1992) Stages of cognitive development: An application of catastrophe theory. *Psychological Review* 99:395–417. [RAH]
- Van Geert, P. (1994) *Dynamic systems of development: Change between complexity and chaos*. Harvester Wheatsheaf. [RAH]
- Van Leeuwen, C., Steyvers, M. & Nooter, M. (1997) Stability and intermittency in large-scale coupled oscillator models for perceptual segmentation. *Journal of Mathematical Psychology* 41:319–44. [RAH]
- VanLehn, K. (1991) Rule acquisition events in the discovery of problem-solving strategies. *Cognitive Science* 15(1):1–47. [aGSH]
- VanLehn, K. & Brown, J. S. (1980) Planning nets: A representation for formalizing analogies and semantic models of procedural skills. In: *Aptitude learning and instruction. vol. 2. Cognitive process analyses of learning and problem solving*, ed. R. E. Snow, P. A. Federico & W. E. Montague. Erlbaum. [aGSH]
- Waltz, J. A., Knowlton, B. J., Holyoak, K. J., Boone, K. B., Mishkin, F. S., de Menezes Santos, M., Thomas, C. R. & Miller, B. L. (in press) A system for relational reasoning in human prefrontal cortex. *Psychological Science*. [JAW]
- Watson, M. W. (1984) Development of social role understanding. *Developmental Review* 4:192–213. [DC]
- Wellman, H. M., Cross, D. & Bartsch, K. (1986) Infant search and object permanence: A meta-analysis of the A-not-B error. *Monographs of the Society for Research in Child Development, vol. 51*. [aGSH]
- Wickens, C. D. (1984) Processing resources in attention. In: *Varieties of attention*, ed. R. Parasuraman & D. R. Davies. Academic Press. [DN]
- Wilkinson, A. (1976) Counting strategies and semantic analyses as applied to class inclusion. *Cognitive Psychology* 8:64–85. [aGSH]
- Wilson, W. H. & Halford, G. S. (1994) Robustness of tensor product networks using distributed representations. *Fifth Australian Conference of Neural Networks, Brisbane*, 258–61. [aGSH]
- Wilson, W. H., Street, D. J. & Halford, G. S. (1995) Solving proportional analogy problems using tensor product networks with random representations. *IEEE International Conference on Neural Networks Proceedings, Perth, Australia*, 2971–75. [aGSH]
- Wimmer, H. & Perner, J. (1983) Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13:103–28. [DF, aGSH]
- Wright, B. C. (1996) What underlies transitive inference performance? Paper given at the British Psychological Society Annual General Conference, September 1996. [BW]
- Wright, B. C. (1998a) Absolute cues to differential size artificially boost performance on transitive inference tasks (submitted). [BW]
- (1998b) Reconceptualizing the transitive inference ability: A framework for research (submitted). [BW]
- Wynn, K. (1992) Addition and subtraction by human infants. *Nature* 358:749–50. [rGSH]
- (1995) Origins of numerical knowledge. *Mathematical Cognition* 1(1):35–60. [rGSH]
- Wynne, C. D. L. (1995) Reinforcement accounts for transitive inference performance. *Animal Learning and Behavior* 23(2):207–17. [rGSH]
- Xu, F. & Carey, S. (1996) Infants' metaphysics: The case of numerical identity. *Cognitive Psychology* 30:111–53. [rGSH]
- Younger, B. (1993) Understanding category members as "the same sort of thing": Explicit categorization in ten month infants. *Child Development* 64(1):309–20. [aGSH]
- Yuille, A. L. & Geiger, D. (1995) Winner-take-all mechanisms. In: *The handbook of brain-theory and neural networks*, ed. M. A. Arbib. MIT Press. [aGSH]
- Zelazo, P. D., Carter, A., Reznick, J. S. & Frye, D. (1997) Early development of executive functions: A problem-solving framework. *Review of General Psychology* 1:198–226. [DF]
- Zelazo, P. D. & Frye, D. (1997) Cognitive complexity and control: A theory of the development of deliberate reasoning and intentional action. In: *Language structure, discourse, and the access to consciousness*, ed. M. Stamenov. John Benjamins. [DF]
- Zelazo, P. D., Reznick, J. S. & Spinazzola, J. (1998) Representational flexibility and response perseveration in a search task. *Developmental Psychology* 34:203–14. [DF]