Review

Relational knowledge: the foundation of higher cognition

Graeme S. Halford[1], William H. Wilson[2] and Steven Phillips[3]

[1] School of Psychology, Griffith University, Nathan, 4111, Australia
[2] School of Computer Science and Engineering, The University of New South Wales, Sydney 2052, Australia
[3] Mathematical Neuroinformatics Group, Human Technology Research Institute, AIST 1-1-1 Umezono, Tsukuba Central 2, Tsukuba, Ibaraki 305-8568, Japan

Accumulating evidence on the nature, function and acquisition of relational knowledge indicates a crucial role of such knowledge in higher cognitive processes. In this review, we specify the essential properties of relational knowledge, together with the role it plays in reasoning, categorisation, planning, quantification and language. Furthermore, we discuss the processes involved in its acquisition and how these processes have been implemented in contemporary neural network models. We present evidence demonstrating that relational knowledge integrates heuristic and analytic cognition, is important for symbolic processes and the creation of novelty, activates specific regions of the prefrontal cortex, and is the most recently evolved and slowest-developing cognitive process. Arguably, relational knowledge represents the core of higher cognition.

The role of relational knowledge in higher cognition

Relations are a core concept in mathematics [1], computer science [2] and cognition [3–7]. In the latter domain, there is converging evidence from both empirical research and computational modelling on the nature, function and acquisition of relational knowledge. As we will argue, this evidence indicates that relational knowledge plays a core role in higher cognitive processes.

In what follows, we first outline the core properties of relational knowledge and then discuss the role of working memory as the workspace where relational representations are constructed. We subsequently discuss the foundational role of relational knowledge in reasoning, language, categorisation and planning. We briefly present evidence on the neural basis of relational processing before discussing acquisition processes, including formal models. Taken together, the evidence we present makes a compelling case for a foundational role of relational knowledge in higher cognition. This thesis is in line with current thinking in the field [6,8], as well as with earlier proposals [4,9–11].

Core properties of relational knowledge

Relational representations can be conceptualised as a binding between a relation symbol and a set of ordered tuples of elements [3,5]. For example, the relation-symbol *larger* is bound to the set of ordered pairs: {(elephant, mouse), (pig, cat)...} (Figure 1a). The symbol represents the 'intension' of a relation and specifies which relation is intended; for example, elements are ordered by size in Figure 1a and by fear in Figure 1b. The ordered tuples represent the 'extension' of a relation, can include knowledge learned by experience (either case by case or incorporating regularities), and can provide statistical knowledge of the world. For example, representations of the *larger* relation include instances of horses being larger than dogs, and exceptions in which horses are not larger than dogs.

Corresponding author: Halford, G.S. (g.halford@griffith.edu.au).

Glossary

Categorical syllogism: Categorical syllogisms consist of two premises and a conclusion, each with a single quantifier, which can be 'All', 'Some', 'Some not' or 'None' (e.g. All *X* are *Y*, No *Y* are *Z* therefore No *X* are *Z*).
Class inclusion: if *A* and *A′* are non-empty subclasses that are included in *B*, then there are more instances of *B* than of *A* or *A′*.
Coordinate system: A system that uses a set of numbers or other ranked entities to uniquely determine the position of points or objects, where position is identified by an ordered tuple of numbers (one for each dimension of the coordinate system).
Error backpropagation: A neural network learning algorithm that allows multilayer networks to learn by computing the error (i.e. the difference between actual and target outputs), 'backpropagating' this error to compute notional errors for nodes in earlier layers, and using these (notional) errors to adjust connection strengths between nodes.
Inner (dot) product: The inner (dot) product $\mathbf{u} \bullet \mathbf{v}$ of vector $\mathbf{u} = (u_1, u_2, \ldots, u_n)$ with vector $\mathbf{v} = (_1, _2, \ldots, _n)$ in Euclidean space is the scalar $\sum_{i=1}^{n} u_{i\,i}$.
Mental model: A constructed representation incorporating instances that make a set of premises true, and which permit valid conclusions to be drawn [3].
Modus ponens: Given that *P* implies *Q*, and that *P* is true, we can infer that *Q* is true that is $P \Rightarrow Q$, *P* therefore *Q*.
Predicate: A logical formula consisting of an *n*-ary predicate symbol *P* followed by *n* terms $t_1, t_2, \ldots, t_n$, and written $P(t_1, t_2, \ldots, t_n)$. Terms can contain variables, and $P(t_1, t_2, \ldots, t_n)$ can evaluate to either true or false when all variables are instantiated. A predicate is essentially a truth-valued *n*-ary function.
Predicate calculus: A form of logic where formulas contain variables that can be quantified, such as the formula $\exists x$ *Man*(*x*), where *x* is an existentially quantified variable, and *Man* is a predicate.
Premise: Propositions that appear in the antecedent position of an argument (distinct from conclusion which appears in the consequent position).
Proposition: An expression that has a truth value, for example 'dogs purr' has a truth value (false) but 'dogs' has no truth value, and is not a proposition.
Relation: An *n*-ary relation is a subset of the Cartesian product of *n* sets. The ordered *n*-tuples that comprise the extension of a relation can be generated from the Cartesian product.
Schema: A structured cognitive representation for example a restaurant schema represents the relation between ordering, receiving and paying for food; the ordering schema in Figure 2 represents relations between the positions top, middle and bottom.
Structure: A set of elements together with the relations defined between them.
Tensor product: The tensor product of two vectors $\mathbf{u} = (u_1, u_2, \ldots, u_m)$ and $\mathbf{v} = (_1, _2, \ldots, _n)$ is a 2-subscript object similar to a matrix, $\mathbf{B} = (b_{ij})$ where $b_{ij} = u_{i\,j}$. Higher rank tensor products are defined analogously, for example the product $\mathbf{T} = (t_{ijk})$ of $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w} = (w_1, w_2, \ldots, w_p)$ would have $t_{ijk} = u_{i\,j} w_k$. Tensor products have several convenient mathematical properties.
Transitive inference: Given that the relation ρ holds between *A* and *B*, and between *B* and *C*, then if ρ is a transitive relation, we can infer that ρ holds between *A* and *C*, that is $A \rho B$ and $B \rho C$ implies $A \rho C$.

  doi:10.1016/j.tics.2010.08.005 Trends in Cognitive Sciences, November 2010, Vol. 14, No. 11

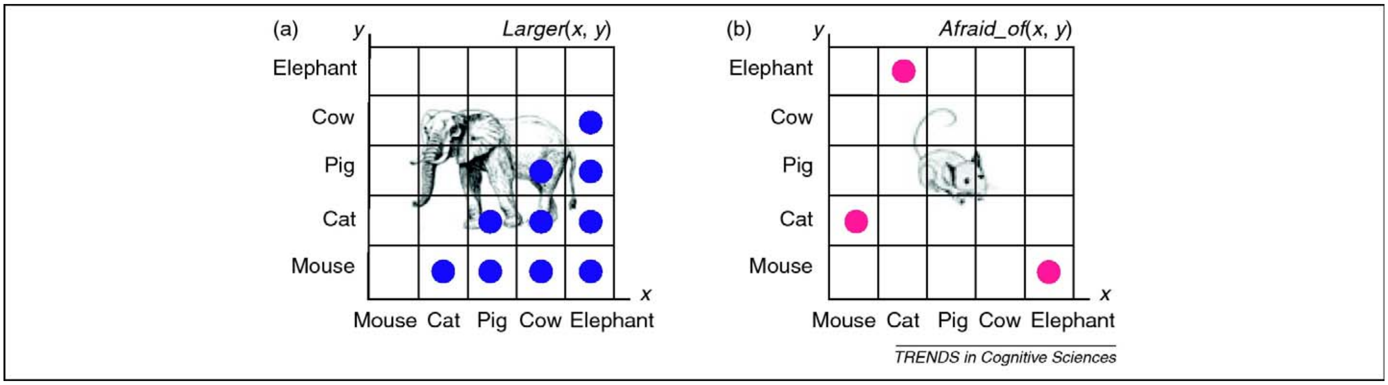

Figure 1. Relational representations can be conceptualised as a binding between a relation symbol and a set of ordered tuples of elements **(a)** Dots indicate the ordered pairs belonging to the relation *larger* (*larger* = {(elephant, mouse), (pig, cat)...}). **(b)** Dots indicate the ordered pairs belonging to the relation *afraid_of* (*afraid_of* = {(mouse, cat), (cat, elephant), (elephant, mouse)}).

These representations provide a knowledge base for estimating conditional probabilities (e.g. of a horse being larger than a dog). Therefore relational knowledge can serve as the basis for heuristics consistent with Bayesian cognitive theory [12].

Relational representations have several core properties, which we outline below. By core properties, we intend those properties that are essential to relational knowledge and which distinguish it from other forms of cognition such as association, or automatic and modular processes, to which we will refer collectively as nonanalytic processes [13].

Structure-consistent mappings, a crucial property of relations, are correspondences between representations that preserve structure in the sense illustrated below. One example would be the way relations between dots on a map correspond to relations between the places they represent, even though the dots have no resemblance to the places. Structural correspondence is defined by consistent mapping of elements and relations [11]. This is illustrated in Figure 2: a spatial schema comprising the elements 'top', 'middle', and 'bottom', with the relation *above* defined between them, used to represent order or position. The premises 'Tom is taller than Peter' and 'Bob is taller than Tom' are mapped onto this schema. In Map (a), the elements 'top', 'middle' and 'bottom' in the schema are mapped to the elements 'Bob', 'Tom' and 'Peter' respectively in the premises. The relations are mapped to correspond with the element mappings, so the relation above is consistently mapped to the relation *taller*. Although mappings can also be influenced by other factors, including element similarity and relational similarity, structure-consistent mappings are crucial both to cognitive representations and to analogies [10]. They enable analytic cognition that has some degree of independence from similarity of content, and they promote selection of relations that are common to several relational instances, which is a major step towards abstraction and representation of variables [14,15]. In fact, structure-consistent mapping between relational representations is a core property of higher cognition, and has been identified as the process that best distinguishes human cognition from that of other animals [16], with chimpanzees possibly being transitional in that they can only perform simple analogies such as 'small square is to large square as small circle is to large circle' [17].

Compositionality is another core property of relations. Representations of complex entities are compositional

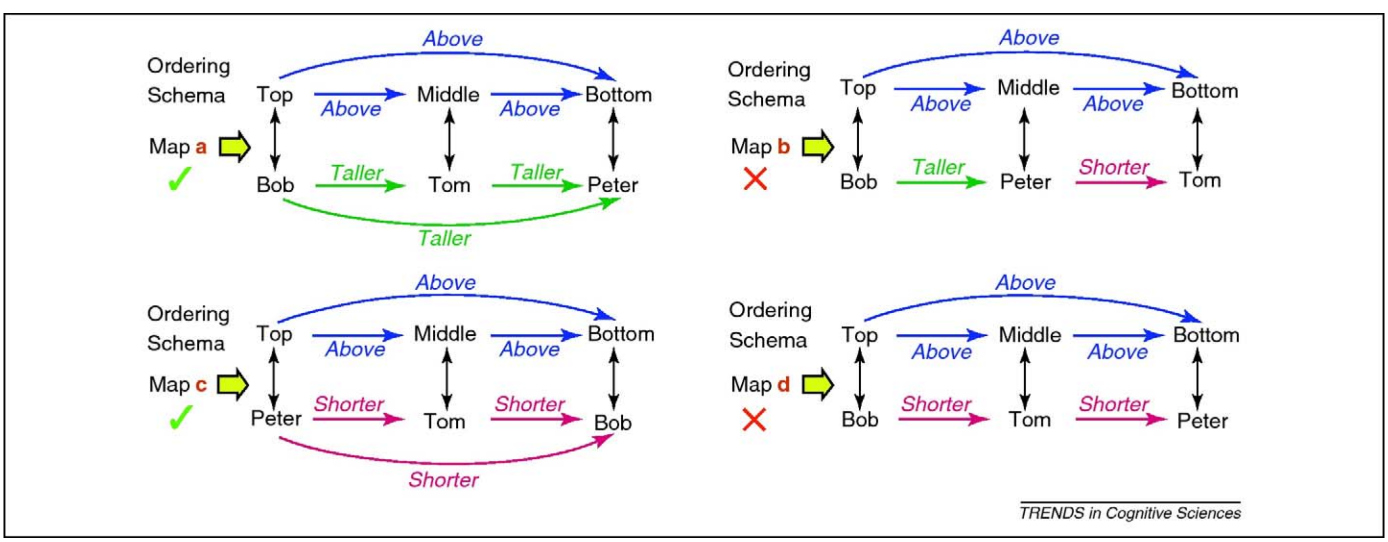

Figure 2. An example of explicit transitive inference. The premises 'Tom is taller than Peter' and 'Bob is taller than Tom' are presented once only, and a mental model is formed in working memory by mapping onto an ordering schema, such as top-middle-bottom. Four mappings are shown. Map **(a)** is based on abstract structural correspondence: the *taller* relation consistently corresponds to the *above* relation. Map **(b)** is *not* structure-consistent because *above* corresponds to *taller* in one case and *shorter* in the other case, and a valid transitive inference cannot be made. Map **(c)** is structure-consistent because *shorter* consistently corresponds to *above*. Map **(d)** is structure-consistent but the truth of the premises is not preserved (taller has been replaced by shorter without converting the order of the elements).

when the constituent entities from which they are constructed retain their identity in the compound representation and are accessible. Given *larger* (*horse*,*dog*) we can determine the answer to “what is larger than a dog?” (*horse*) and the answer to “what is the specified relation between horse and dog?” (*larger*) and so on.

Systematicity means that certain cognitive capacities are intrinsically connected in that, for example, the capacity to understand ‘John loves Mary’ implies the capacity to understand ‘Mary loves John’ [18]. Systematicity enables generation of novel instances. With a representational format *loves*(-, -), whereby agent (lover) and patient (loved) slots are bound together by the predicate ‘loves’, we can generate an unlimited set of propositions by binding elements to the slots: *loves*(*Mary*, *John*), *loves*(*Wendy*, *Bob*) and so on. A more formal explanation is given by [19].

These core properties of relations (as well as the implications and applications of relational knowledge presented in Box 1) help explain their foundational role in higher cognitive functions which we address in detail in a subsequent section.

The role of working memory in constructing relational representations

Working memory is recognised as the workspace where relational representations are constructed [5,15,20,21]. It plays a role in the determination of structural correspondence [22,23] because the mapping between elements is temporary and because validity of a mapping can be established by activating the representations without external input. Notice that, in Figure 2, a structurally consistent mapping can be determined from the correspondence between the premises and the ordering schema.

A contemporary working-memory theory [21], based on dynamic binding to coordinate systems (i.e. ordered sets of elements), grounds relational knowledge theory in working memory. The ordering schema in Figure 2 would be an example of a coordinate system [5,15,20,21,23].This reformulated working-memory theory provides a general-purpose mechanism that can build, maintain, manipulate and update structural representations. The representations are based on activation of long-term memory, and there is a region of direct access, a focus of attention, and a procedural working memory that together constitute an analytic processing subsystem [21]. New structural representations can be transferred to long-term memory. The amount of information required to establish consistency, and therefore the working memory load, can be quantified by the relational complexity metric [5,8] summarised in Box 2. The mapping process is modelled in the Discovery Of Relations by Analogy (DORA) [20] and Structured Tensor Analogical Reasoning (STAR) [5] models, considered later.

Working memory accounts for approximately 50% of variance in fluid intelligence [24,25] and it shares substantial variance in reasoning that is not accounted for by processing and storage demands, or by processing speed [26]. This indicates that the shared variance at least partly reflects ability to form structured representations.

In other words, working memory is the workspace where relational representations are constructed and it is influenced by knowledge stored in semantic memory. Therefore, it plays an important role in the interaction of analytic and nonanalytic processes in higher cognition.

The foundation of reasoning, language, categorisation and planning

Reasoning and structure-consistent mappings

Relational knowledge integrates nonanalytic and analytic cognition, sometimes called Type 1 and Type 2 respectively

> **Box 1. Applications and implications of relational knowledge**
>
> - An instance of a **relation** is a proposition that is true, for example *larger*(elephant, mouse) is an instance of the *larger* relation and is a proposition.
> - **Symbols** inhabit symbol systems, that is symbol structures in a relational environment [4,79].
> - An **operation** is a kind of function, and a function is a special kind of relation. Thus arithmetic addition is a function of two addends, and is a ternary relation comprising a set of ordered 3-tuples; $+=\{(x, y, z)|x, y, z \in \mathbf{N} \,\&\, x+y=z\}$. A function is a relation in which every element in the domain (input) is mapped to only one element in the co-domain (output). Thus addition is a function, but *parent_of*(−,−) is not.
> - **Higher-order** relational representations, with lower levels embedded within higher levels and the same operation being performed at each level, form a basis for recursion [80].
> - **Relations differ from associations**. Relations are distinguished by symbols for example *contains*(cup, drink) and *stands_upon*(cup, saucer), whereas associations *cup* → *drink*, and *cup* → *saucer* are not [5]. An associative link *per se* cannot be an argument to another association, so associations cannot be recursive [80]. However, relational frames can be superimposed on associative learning [81].
> - **Rules** are widely used to describe higher cognitive processes, but rules can be expressed as relations that have the advantage of an established mathematical definition. In practice, rules are often compressed or compact representations of relations.
> - **Structure** is defined as a set of relations between elements.
> - **Novel** instances can be created, as noted under systematicity. Mapping new phenomena into existing representations, as in the transitive inference in Figure 2, also creates novelty. Furthermore, substituting a different relation symbol activates a different set of relational mappings, for example changing from + to × switches from mappings 3,5 → 8 to 3,5 → 15 and so on [5]. Change of representations in working memory can change strategies ([21] § 2.2.2).

> **Box 2. The relational complexity metric**
>
> The complexity of relations is assessed by the number of entities that are related, which corresponds to the number of slots in the representation of a relation [5,8]. Thus the binary relation larger-than has two slots, one for a larger entity and one for a smaller. In the transitive inference in Figure 2, the mental model is a ternary relation, *monotonically_taller*(*Bob, Tom, Peter*), relating three entities, and it imposes a higher processing load than the premises *taller*(*Bob, Tom*) and *taller*(*Tom, Peter*) that are binary relations [5].
>
> Human processing capacity is limited to one quaternary relation in parallel [84]. Concepts too complex to be processed in parallel are handled by segmentation (decomposition into smaller segments that can be processed serially) and conceptual chunking (recoding representations into lower complexity, but at the cost of making some relations inaccessible). These principles lead to objective assessment criteria embodied in the Method for Analysis of Relational Complexity (MARC) [8].Working-memory limitations causing complex relations to be segmented makes serial processing inevitable in complex relational cognition [84]. The relational complexity metric has wide applicability [35,43,47,54,85–88] (see also main text).

[13,27]. Implicit transitive inferences [28] are nonanalytic to the extent that they reflect associative learning. Explicit transitive inference, based on construction of a mental model in working memory as in Figure 2a, is analytic because it reflects the logical consequences of the relevant relations. However, transitive inferences based on relational knowledge can also be influenced by beliefs. Given premises 'the tractor is faster than the train' and 'the car is slower than the train' participants tend to make the invalid inference 'the car is faster than the tractor' [29] because of acquired knowledge of the relative speeds of cars and tractors.

Although it is assumed that analytic and nonanalytic processes are supported by one system [13,27], there is a need for more information about the basis and properties of this integrated system. Relational knowledge processes incorporate real-world knowledge and also the ability to construct mental models that reflect the logical consequences of premises. Therefore, relational processes have properties of both analytic and nonanalytic cognition, and provide a natural way of dealing with their interdependency and interaction. We will consider some of the cognitive processes that depend on relational processing.

Analogy is a basic reasoning process that has a role in a wide range of higher cognitive processes [6,30] and exemplifies well the core properties of relational processing, including structure-consistent mapping [31,32]. In analogy, the mapping can be influenced by element or relational similarity, by semantic knowledge, by motivational factors and by structural correspondence [33]. For example, consider a picture analogies task [34] that comprises picture A, in which a boy restrains a dog that chases a cat, and picture B, in which a tree restrains the dog that chases the boy. In a relational match, 'boy' in A is mapped to 'tree' in B because both restrain the dog. In an element (featural) match, 'boy' in A is mapped to 'boy' in B. Similarity also depends on structural correspondence, so 'woman feeding squirrel' is not similar to 'squirrel feeding woman' because the elements, although identical, are not in corresponding slots. Mappings based on abstract structural correspondence are also important because they permit inferences that go beyond experience, and have some independence of semantic knowledge.

Mental models, which account for many forms of reasoning [35–38], are iconic in that they have a structure that corresponds to the structure of the subject matter that is represented [3]. Consequently, structure-consistent mappings play an important role. Mental models also have other properties of relational knowledge including compositionality and systematicity, they are constructed in working memory, and they reflect content and context [3,39]. A mental model for transitive inference, referred to by William James [40] as the fundamental principle of inference, can be formed by mapping premises into existing schemas, as shown in Figure 2. The mapping in Figure 2a forms a mental model 'Bob_taller than Tom_taller than Peter', from which a valid transitive inference, 'Bob is taller than Peter', can be made. However, premises such as 'Bob *was taller* than Tom, Tom *was taller* than Peter' are not transitive because of possible change over time (e.g. if Peter grew faster than Bob). Consequently, participants tend not to form an integrated mental model, thereby reflecting the semantic content of the premises [39]. This reflects the interaction of content and structure that is inherent in relational knowledge.

Mental models of conditional reasoning are also relational and reflect semantic and syntactic information. The initial mental model of a basic conditional $P \rightarrow Q$ (P implies Q) comprises only the conjunction that is seen as making the conditional true, that is P Q. Fleshing out into a full model entails adding *not-P Q* and *not-P not-Q*, and depends on information retrieved from semantic memory [41]. However with a promise, the initial model contains the conjunctions P Q and *not P not Q*, reflecting the semantic properties of a promise, which implies a reward if, and only if, the required task is performed [42]. This also reflects the content–structure interaction that characterises relational knowledge.

Mental models of categorical syllogisms can be expressed as relations between categories implied by the premises [43] and their difficulty can be predicted by the relational complexity metric, summarised in Box 2 [5,8]. Mental models are typically constructed by the reasoner, but reasoning can be performed by mapping into pre-existing logical schemas, such as *modus ponens* [28,44,45]. Therefore structure-consistent mappings are also relevant to reasoning by logical schemas.

Mental models embody the core properties of relational knowledge, including structure-consistent mapping, compositionality, systematicity, and construction of mental models in working memory. Relational knowledge, as embodied in mental models, also incorporates nonanalytic knowledge and provides an account of its interaction with analytic knowledge.

Language

Relational knowledge is increasingly recognised as important in cognitive linguistics. Structural alignment facilitates learning of both word meanings and grammar [14]. The acquisition of verbs includes frames or slots [46], which are components of relational representations, and the relational complexity metric accounts for difficulty of sentence comprehension [47] as well as specialised language metrics do [48]. The syntactic structure of language, although conventionally expressed as rules can, in principle, be expressed as relations. Our proposal that the core properties of relational knowledge incorporate syntactic and semantic information for use in decision making and reasoning is consistent with parallel phonological, syntactic and semantic architectures, with interfaces, for language [49]. Furthermore, recursion has been proposed as an essential foundation for language [50] and relational knowledge is recursive (see Box 1). This might be the most important implication of relational knowledge for language.

Categorisation

Relational categories, such as parenthood, which is defined by relation to an offspring, have high frequency and importance [51]. Theory-based categories [52] can be conceptualised as relational categories: for example, the category of orbiting objects can be based on relations such as gravitational attraction and difference in mass causing planets to orbit the sun. Theory-based categorisation is sensitive to alignment of explanatory structures [53]. Equivalence classes of cognitive processes can be defined by mappings between mathematical structures. Thus, although transitivity is superficially dissimilar to class inclusion [54] analysis of mathematical structures shows underlying commonalities [55].

Planning

Planning depends on creating sequences of actions or operations that transform the current state into a goal state as well as on representing relations [5]. Strategy acquisition can be based on relational representations. For example, strategies for ordering objects can be acquired using an ordered set of three or more elements as a template [56]. Cognitive control can be based on representation of motivationally relevant relations [57], such as those that predict rewards or punishments. An example would be the 1 – AX and 2 – BY task in which context cues 1 or 2, change the contingencies AX or BY that lead to reward [58].

Processes with the core properties of relational knowledge, including structure-consistent mappings, play a significant role in a wide range of higher cognitive functions, including reasoning, which depend on analogy, mapping onto mental models or logical schemas, as well as language, categorisation and planning.

The neuroscience of relations

Anterior regions of the prefrontal cortex that are known to be involved in higher cognition are activated by processing relations [59–61]. The frontopolar cortex (lateral Brodmann area 10) is involved in relational integration [62–65], maintaining dynamic binding between representations in working memory [66] and when processing highly abstract information [67]. The prefrontal cortex is also involved in dynamic mappings between representations in analogy [68,69]. The growing evidence that relational processing involves these most recently evolved and slowest developing regions of the brain [70] is consistent with relational processing being the core of higher cognition.

The acquisition of relational knowledge

The acquisition of relational knowledge depends on processes that are relatively unique and are crucial to understanding the nature of relational knowledge itself. Relational knowledge can be acquired from experience with examples and is partly autonomous and self-supervised [20]. Implicit learning, including artificial grammar acquisition [71], can be triggered by relatively automatic recognition of regularities, without explicit rules. Implicit learning is evolutionarily and developmentally early, is robust to neurological damage and makes low information-processing demands [72]. Further development of relational representations can occur through theory revision and redescription [73] and by structural alignment, recoding, binding to symbols and analogical mapping [15]. Relational representations can be established autonomously in working memory, based on structure-consistent mappings [15,20]. The following subsections discuss three neural net models that implement explicit relational knowledge acquisition processes (Figure 3).

The Semantic Cognition model

The Semantic Cognition model [74] demonstrates that some relational properties can be captured without formal representations of structure. It consists of five layers of interconnected units, as shown in Figure 3a. There are two layers of input units: item units represent plants and animals, and relation units represent relational inputs or 'contexts' (ISA, is, can, has). Each item unit is connected to every unit in a layer of Representation units. Every Representation and Relation unit is connected to all units in a layer of Hidden units, which in turn are connected to all units in a layer of Attribute (output) units that represent attributes (e.g. living thing).

A backpropagation learning algorithm adjusts the connection weights between layers so that activation of an item (e.g. *robin*) and a relation (e.g. *ISA*) activates appropriate attribute(s) (e.g. *bird*) in the output layer. This results in patterns of activations in Representation and Hidden layers that reflect similarities within and between categories (e.g. *robin* and *salmon* are more similar than *robin* and *daisy*). Similarity is simulated even when not directly perceptible (e.g. *robin* and *salmon* are both l*iving*, *grow* etc.). This mechanism exemplifies the way representations that code structural features can form in internal layers of feed-forward networks as a result of learning to compute input–output links [75]. It is consistent with the idea [76] that categorisation can develop in infants through observation of events in which entities take part, sometimes independent of perceptible similarity. The model potentially explains children's ability to generalise propositions (e.g. *dog has a spleen*) on the basis of category membership rather than perceptual similarity [77,78]. Other relational properties include implicit representation of slots because predicates can bind different items and attributes (e.g. the predicate *has* can bind *robin* to *wings*, *fish* to *gills etc.*). The model can also interpret new information, so if taught *sparrow is a bird*, it generalises to *sparrow, can fly*.

The importance of these achievements is partly in showing how relational knowledge can be based on semantic similarity of activations in Representation and Hidden layers. These layers do not include formal representation of structure, so they would not enable solution of a problem such as: 'If there are both boys and girls in city X, are there more children or more boys?' This is because there is no representation of the inclusion relation between the superordinate set of children and the (non-empty) subordinate sets of boys and girls [78]; there is no mapping between representations on the basis of structural correspondence, and there is no active construction of relational representations in working memory. Representations in this type of model have restricted accessibility, due to lack of symmetry between input and output layers. Given inputs such as *Robin* ISA the output is *bird* but there is no way to generate *Robin* in response to queries such as ____ *ISA bird*. Symbolic connectionist models, including DORA and STAR, complement semantic knowledge models.

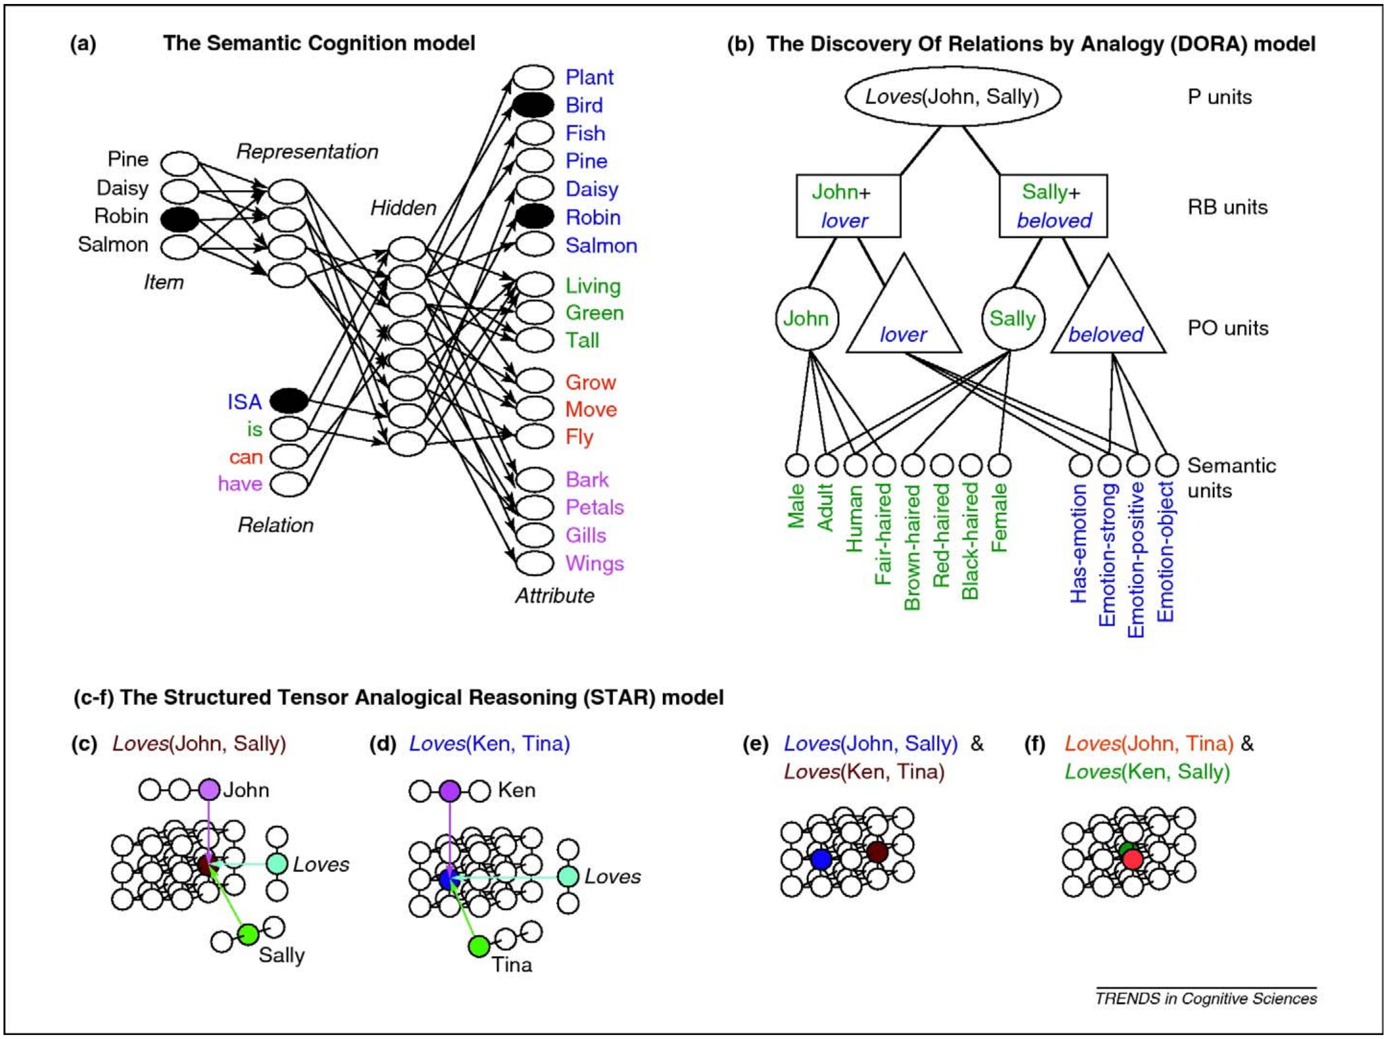

Figure 3. **(a)** The Semantic Cognition model (adapted, with permission, from [74]). The net comprises five layers of units: *item* (input) units represent plants (pine, daisy) and animals (robin, salmon). Each item unit is connected to every unit in a layer of *representation* units. *Relation* units represent relational inputs or 'contexts' (ISA, is, can, has). Every representation and relation unit is connected to all units in a layer of *hidden* units, which in turn are connected to all units in a layer of *attribute* (output) units that represent attributes (e.g. living thing). The activations shown (filled circles) represent 'robin ISA bird'. **(b)** Outline of the architecture of the DORA model (adapted from [20]). Semantic units that represent features of objects (e.g. male, adult…has-emotion) are linked to P (e.g. *lover*) and O (e.g. *John*) units that are linked to RB units that code role-filler bindings. The relation instances so formed are represented by P units. **(c–d)** The STAR model [5] representation of binary relations. Throughput is multidirectional so input of any two components (using the dot product operator) produces the third component. **(e–f)** Two relations superimposed on the same representation. The input units, representing the relation symbol (*loves*) and arguments (*John and Sally*), are shown only in figures **(c–d)**, whereas figures **(e–f)** show only binding units, to avoid clutter.

The DORA model

Relations are represented in DORA [20] by four layers of units, as shown in Figure 3b. At the top are Proposition (P) units that are linked to Predicate and Object (PO) units via Role-Binding units (RB). The PO units are connected to Semantic units that represent features of objects (e.g. male, adult…has-emotion). Relations are represented by binding roles to fillers, so *loves*(*John*, *Sally*) would be represented by binding *John* to the *lover* role and *Sally* to the *loved* role; these bindings are coded by the RB units. Role-filler bindings are dynamic, and are coded by firing in close temporal proximity: units representing *lover* fire in close temporal proximity with units representing *John*, but out of temporal proximity to units representing *loved* or *Sally*. Retrieval from long-term memory occurs due to activation that originates with the P unit, passes through RB and PO units to Semantic units, which excite units in long-term memory. Structure-consistent mapping occurs by concurrent activation of units in two analogues, taking account of structure and semantic similarity. The model accounts for several important properties of relational knowledge, including acquisition from examples, and integrates semantic and syntactic information.

The STAR model

In the STAR model [5] (Figure 3c-f), the relation symbol and arguments are each represented by activations on sets of neural units that correspond to vectors: *loves(John, Sally)* is represented by activations on three vectors of input units, representing *loves, John, Sally* (Figure 3c). The relation symbol is bound to elements by an array (tensor) formed by the outer products of the vectors representing the inputs. Activations of the internal units in the tensor are formed dynamically in working memory as the direct result of activations in the input units. The relation symbol and roles are represented by positions in the representation, so the symbol *loves* and roles *lover* and *loved* correspond to different axes of the array. This gives a natural correspondence to relational representations in predicate calculus expressions such as *loves*(*John*, *Sally*), in natural language ("John loves

Sally") and to representation of structure in the working-memory model of Oberauer [21]. Relational instances can be accumulated by superimposition, as shown in Figures 3e and 3f, enabling recognition of their commonalities. Accessibility is implemented by the dot product operator ([5], §4.2.1) and throughput is multidirectional, so input of any two components produces the third component. Relational instances comprising up to four related elements could be learned initially, whereas in DORA single-place are learned initially.

Complementary but distinct processes

These models represent, in different ways, the combination of semantic and syntactic cognition that characterises relational knowledge [13]. The overlap between nonanalytic and analytic cognition can make these processes difficult to distinguish, and has resulted in protracted debate as to the nature, or even the existence, of distinct cognitive systems [13]. However, relational knowledge has unique properties, including structure-consistent mapping and the ability to construct new representations in working memory. These confer the ability to break free from previous experience and are the basis of much human inference [13,38]. These unique properties are best observed in tasks that cannot be solved by semantic cognitive processes, including some instances of categorical syllogisms, explicit transitive inference and conditional reasoning [28]. These unique properties are captured in DORA and STAR to a greater extent than in the Semantic Cognition model. However, the Semantic Cognition model simulates acquisition and coding processes that could form the basis of structure-based cognition, so the two types of models are complementary. At the same time, they are also distinct, so evidence of one type of knowledge might not imply existence of the other.

Concluding remarks

The properties of relational knowledge correspond in many ways to those of higher cognition. Relational knowledge also provides an integrative framework for a broad range of fields, including inference, categorisation, quantification, planning, language, working memory and knowledge acquisition. Relational knowledge integrates semantic and syntactic information used in decision making and reasoning, indicating that some theoretical rivalries could be replaced by research that identifies conditions conducive to each process, and defines their respective strengths and limitations and the nature of their interaction (Box 3).

The core feature that is unique to relational processes entails construction of representations in working memory based on structure-consistent mappings. This property, together with research on the nature and limitations of working memory, might offer explanations for the correlation with fluid intelligence, for the serial nature of higher cognitive processes, for semantic tasks evolving earlier and being mastered at an earlier age, and for the flexibility and adaptability of higher cognition. Relational properties provide a complexity metric that has been shown to have wide generality, and recent computational models and empirical research provide an account of acquisition processes. Relational knowledge can be operationalised by a wide range of empirical paradigms; it can be compared directly with association, as well as with semantic networks and the compositional syntax and semantics of classical cognitive theories. Analogical reasoning, which incorporates relational processing, is linked to an impressive array of cognitive functions. The extensive literature that we have reviewed supports the contention that relational knowledge provides the foundation for higher cognitive processes.

Box 3. Questions for future research

- What is the precise nature of the link between dynamic binding in working memory and acquisition of relational knowledge?
- What specific working-memory process determines whether two representations are in structural correspondence?
- Given that some relational processing can be accomplished without construction of representations in working memory, as in the Semantic Cognition model, what limits (if any) apply to semantic models of relational cognition?
- What are the possible correspondences between relational knowledge and executive functions?
- What cortical subregions are involved in dynamic mappings between representations on the basis of structural correspondence?
- The identification of symbolic processes with relational knowledge raises new questions regarding the 'subsymbolic to symbolic' transition. Does dynamic binding in working memory have a crucial role in the transition to symbolic processes, in analogy, in the development of strategies, or in the transition from Theory of Mind based on image schemas to Theory of Mind based on elicited responses [89]?
- Should relational representations be based on role-filler bindings [20] or on bindings of symbols to elements, as in STAR [5]?
- Are relations learned one component at a time (as in DORA) or can relational knowledge acquisition begin with links between up to four entities (as in STAR)?

Acknowledgements

We would like to acknowledge valuable contributions by Brentyn Ramm, Mark Chappell, Glenda Andrews, Trevor Hine, Arthur Poropat and three anonymous reviewers.

References

1 Mac Lane, S. and Birkhoff, G.D. (1967) *Algebra*, Macmillan
2 Codd, E.F. (1990) *The relational model for database management: Version 2*, Addison-Wesley
3 Goodwin, G.P. and Johnson-Laird, P.N. (2005) Reasoning about relations. *Psychol. Rev.* 112, 468–493
4 Halford, G.S. and Wilson, W.H. (1980) A category theory approach to cognitive development. *Cogn. Psychol.* 12, 356–411
5 Halford, G.S. *et al.* (1998) Processing capacity defined by relational complexity: Implications for comparative, developmental, and cognitive psychology. *Behav. Brain Sci.* 21, 803–831
6 Kokinov, B. et al. (2009) *New frontiers in analogy research: Proceedings of the Second International Conference on Analogy*, New Bulgarian University Press
7 Suppes, P. and Zinnes, J.L. (1963) Basic measurement theory. In *Handbook of mathematical psychology* (Luce, R.D., ed.), pp. 1–76, Wiley
8 Halford, G.S. *et al.* (2007) Separating Cognitive Capacity from Knowledge: A new hypothesis. *Trends Cogn. Sci.* 11, 236–242
9 Gentner, D. (1983) Structure-mapping: A theoretical framework for analogy. *Cogn. Sci.* 7, 155–170
10 Holland, J.H. *et al.* (1989) *Induction: Processes of inference, learning and discovery*, MIT Press
11 Holyoak, K.J. and Thagard, P. (1989) Analogical mapping by constraint satisfaction. *Cogn. Sci.* 13, 295–355
12 Oaksford, M. and Chater, N. (2007) *Bayesian rationality: The probabilistic approach to human reasoning*, Oxford University Press

13 Evans, J.S.B.T. (2008) Dual-Processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol.* 59, 255–278
14 Gentner, D. and Namy, L.L. (2006) Analogical processes in language. *Curr. Dir. Psychol. Sci.* 15, 297–301
15 Halford, G.S. and Busby, J. (2007) Acquisition of structured knowledge without instruction: The relational schema induction paradigm. *J. Exp. Psychol. Learn. Mem. Cogn.* 33, 586–603
16 Penn, D.C. *et al.* (2008) Darwin's mistake: Explaining the discontinuity between human and nonhuman minds. *Behav. Brain Sci.* 31, 109–130
17 Premack, D. (2010) Why humans are unique: Three theories. *Perspect. Psychol. Sci.* 5, 22–32
18 Fodor, J.A. and Pylyshyn, Z.W. (1988) Connectionism and cognitive architecture: A critical analysis. *Cognition* 28, 3–71
19 Phillips, S. and Wilson, W.H. (2010) Categorial Compositionality: A Category Theory Explanation for the Systematicity of Human Cognition. *PLoS Comput. Biol.* 6 (7), e1000858 DOI: 10.1371/journal.pcbi.1000858
20 Doumas, L.A. *et al.* (2008) A theory of the discovery and predication of relational concepts. *Psychol. Rev.* 115, 1–43
21 Oberauer, K. (2009) Design for a working memory. In *The psychology of learning and motivation: Advances in research and theory* (Ross, B.H., ed.), pp. 45–100, Elsevier Academic Press
22 Kubose, T.T. *et al.* (2002) The role of textual coherence in incremental analogical mapping. *J. Mem. Lang.* 47, 407–435
23 Waltz, J.A. *et al.* (2000) The role of working memory in analogical mapping. *Mem. Cognit.* 28, 1205–1212
24 Kane, M.J. *et al.* (2004) The Generality of Working Memory Capacity: A Latent-Variable Approach to Verbal and Visuospatial Memory Span and Reasoning. *J. Exp. Psychol. Gen.* 133, 189–217
25 Kane, M.J. et al. (2005) Working Memory Capacity and Fluid Intelligence Are Strongly Related Constructs: Comment on Ackerman, Beier, and Boyle (2005). Psychol. Bull. 131, 66–71
26 Oberauer, K. *et al.* (2008) Which working memory functions predict intelligence? *Intelligence* 36, 641–652
27 De Neys, W. (2006) Two systems but one reasoner. *Psychol. Sci.* 17, 428–433
28 Goel, V. (2007) Anatomy of deductive reasoning. *Trends Cogn. Sci.* 11, 435–441
29 Andrews, G. (2010) Belief-based and analytic processing in transitive inference depends on premise integration difficulty. *Mem. Cognit.* 38, 926–938
30 Hofstadter, D.R. (2001) Analogy as the core of cognition. In *The analogical mind: Perspectives from cognitive science* (Gentner, D. *et al.*, eds), pp. 499–538, MIT Press
31 Gentner, D. and Kurtz, K.J. (2006) Relations, Objects, and the Composition of Analogies. *Cogn. Sci.* 30, 609–642
32 Holyoak, K.J. (2005) Analogy. In *The Cambridge Handbook of Thinking and Reasoning* (Holyoak, K.J. and Morrison, R.G., eds), pp. 117–142, Cambridge University Press
33 Blanchette, I. and Dunbar, K. (2000) How analogies are generated: The roles of structural and superficial similarity. *Mem. Cognit.* 28, 108–124
34 Tohill, J.M. and Holyoak, K.J. (2000) The impact of anxiety on analogical reasoning. *Think. Reas.* 6, 27–40
35 English, L.D. and Halford, G.S. (1995) *Mathematics education: Models and processes*, Erlbaum
36 Evans, J.S.B.T. (2005) Deductive Reasoning. In *Cambridge Handbook of Thinking and Reasoning* (Holyoak, K. and Morrison, R.G., eds), pp. 169–184, Cambridge University Press
37 Halford, G.S. (1993) *Children's understanding: The development of mental models*, Erlbaum
38 Johnson-Laird, P.N. (2005) Mental Models and Thought. In *Cambridge Handbook of Thinking and Reasoning* (Holyoak, K. and Morrison, R.G., eds), pp. 185–208, Cambridge University Press
39 Goodwin, G.P. and Johnson-Laird, P.N. (2008) Transitive and pseudo-transitive inferences. *Cognition* 108, 320–352
40 James, W. (1890) *The Principles of Psychology* (Vol. 2), Holt
41 Markovits, H. and Barrouillet, P. (2002) The development of conditional reasoning: A mental model account. *Dev. Rev.* 22, 5–36
42 Gauffroy, C. and Barrouillet, P. (2009) Heuristic and analytic processes in mental models for conditionals: An integrative developmental theory. *Dev. Rev.* 29, 249–282
43 Zielinski, T.A. *et al.* (2010) Complexity of categorical syllogisms: An integration of two metrics. *Eur. J. Cogn. Psychol.* 22, 391–421
44 Rips, L. (1994) *The psychology of proof: Deductive reasoning in human thinking*, MIT Press
45 Stenning, K. and van Lambalgen, M. (2008) *Human reasoning and cognitive science*, MIT Press
46 Naigles, L.R. *et al.* (2009) Flexibility in early verb use: Evidence from a Multiple-N diary study. *Monogr. Soc. Res. Child Dev.* 74, 1–112
47 Andrews, G. *et al.* (2006) Relational processing and working memory in the comprehension of complex relative clause sentences. *Mem. Cognit.* 34, 1325–1340
48 Gibson, E. (2000) The dependency locality theory: A distance-based theory of linguistic complexity. In *Image, language, brain* (Miyashita, Y. *et al.*, eds), pp. 95–126, MIT Press
49 Jackendoff, R. (2007) *Language, consciousness, culture*, MIT Press
50 Fitch, W.T. *et al.* (2005) The evolution of the language faculty: Clarifications and implications. *Cognition* 97, 179–210
51 Gentner, D. and Kurtz, K.J. (2005) Relational categories. In *Categorization inside and outside the laboratory* (Woo-Kyoung, A. *et al.*, eds), pp. 151–175, American Psychological Association
52 Murphy, G.L. and Medin, D.L. (1985) The role of theories in conceptual coherence. *Psychol. Rev.* 92, 289–316
53 Jameson, J. and Gentner, D. (2009) Explanation and structure-mapping in theory-based categorization. In *New Frontiers in Analogy Research: Proceedings of the Second International Conference on Analogy* (Kokinov, B. *et al.*, eds), pp. 222–228, New Bulgarian University Press
54 Andrews, G. and Halford, G.S. (2002) A cognitive complexity metric applied to cognitive development. *Cogn. Psychol.* 45, 153–219
55 Phillips, S. *et al.* (2009) What do transitive inference and class inclusion have in common? Categorical (co)products and cognitive development. *PLoS Comput. Biol.* 5, e1000599 DOI: 10.1371/journal.pcbi.1000599
56 Halford, G.S. *et al.* (1995) Modelling the development of reasoning strategies: The roles of analogy, knowledge, and capacity. In *Developing cognitive competence: New approaches to cognitive modelling* (Simon, T. and Halford, G.S., eds), pp. 77–156, Erlbaum
57 Speed, A. (2010) Abstract relational categories, graded persistence, and prefrontal cortical representation. *Cogn. Neurosci.* 1, 126–137
58 Reynolds, J.R. and O'Reilly, R.C. (2009) Developing PFC representations using reinforcement learning. *Cognition* 113, 281–292
59 Golde, M. *et al.* (2010) Differential role of anterior prefrontal and premotor cortex in the processing of relational information. *NeuroImage* 49, 2890–2900
60 Kroger, J. *et al.* (2002) Recruitment of anterior dorsolateral prefrontal cortex in human reasoning: A parametric study of relational complexity. *Cereb. Cortex* 12, 477–485
61 Smith, R. *et al.* (2007) Localizing the rostrolateral cortex at the individual level. *NeuroImage* 36, 1387–1396
62 Cho, S. *et al.* (2010) Common and dissociable prefrontal loci associated with component mechanisms of analogical reasoning. *Cereb. Cortex* 20, 524–533
63 Fangmeier, T. *et al.* (2006) fMRI Evidence for a Three-Stage Model of Deductive Reasoning. *J. Cogn. Neurosci.* 18, 320–334
64 Jung, R.E. and Haier, R.J. (2007) The Parieto-Frontal Integration Theory (P-FIT) of intelligence: Converging neuroimaging evidence. *Behav. Brain Sci.* 30, 135–154
65 Ramnani, N. and Owen, A.M. (2004) Anterior prefrontal cortex: Insights into function from anatomy and neuroimaging. *Nat. Rev. Neurosci.* 5, 184–194
66 Prabhakaran, V. *et al.* (2000) Integration of diverse information in working memory within the frontal lobe. *Nat. Neurosci.* 3, 85–90
67 Christoff, K. *et al.* (2009) Prefrontal organization of cognitive control according to levels of abstraction. *Brain Res.* 1286, 94–105
68 Bunge, S.A. *et al.* (2005) Analogical Reasoning and Prefrontal Cortex: Evidence for Separable Retrieval and Integration Mechanisms. *Cereb. Cortex* 15, 239–249
69 Green, A.E. *et al.* (2006) Frontopolar cortex mediates abstract integration in analogy. *Brain Res.* 1096, 125–137
70 Stone, V.E. (2006) The evolution of ontogeny and human cognitive uniqueness: Selection for extended brain development in the hominid line. In *Evolutionary Cognitive Neuroscience* (Platek, S.M. *et al.*, eds), MIT Press
71 Shanks, D.R. (2010) Learning: From association to cognition. *Annu. Rev. Psychol.* 61, 273–301

72 Litman, L. and Reber, A.S. (2005) Implicit cognition and thought. In *Cambridge Handbook of Thinking and Reasoning* (Holyoak, K. and Morrison, R.G., eds), pp. 431–455, Cambridge University Press
73 Dixon, J.A. and Kelly, E. (2007) Theory Revision and Redescription: Complementary Processes in Knowledge Acquisition. *Curr. Dir. Psychol. Sci.* 16, 111–115
74 Rogers, T.T. and McClelland, J.L. (2004) *Semantic cognition: A parallel distributed processing approach*, MIT Press
75 Thomas, M.S.C. and McClelland, J.L. (2008) Connectionist models of cognition. In *Cambridge Handbook of Computational Psychology* (Sun, R., ed.), pp. 23–58, Cambridge University Press
76 Mandler, J.M. (2000) Perceptual and conceptual processes in infancy. *J. Cogn. Dev.* 1, 3–36
77 Gelman, S.A. and Kalish, C.W. (2006) Conceptual Development, In *Handbook of Child Psychology: Cognitive, language and perceptual development* (6th edn) (Kuhn, D. and Siegler, R., eds), John Wiley & Sons Inc
78 Halford, G.S. and Andrews, G. (2006) Reasoning and problem solving, In *Handbook of Child Psychology: Cognitive, language and perceptual development* (6th edn) (Kuhn, D. and Siegler, R., eds), pp. 557–608, John Wiley & Sons
79 Newell, A. (1980) Physical symbol systems. *Cogn. Sci.* 4, 135–183
80 Phillips, S. et al. (1995) The processing of associations versus the processing of relations and symbols: A systematic comparison. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (Moore, J.D. and Lehman, J.F., eds), pp. 688–691, Lawrence Erlbaum
81 Gross, A.C. and Fox, E.J. (2009) Relational frame theory: An overview of the controversy. *Anal. Verbal Behav.* 25, 87–98
84 Halford, G.S. *et al.* (2005) How many variables can humans process? *Psychol. Sci.* 16, 70–76
85 Andrews, G. and Halford, G.S. (1998) Children's ability to make transitive inferences: The importance of premise integration and structural complexity. *Cogn. Dev.* 13, 479–513
86 Andrews, G. *et al.* (2003) Theory of mind and relational complexity. *Child Dev.* 74, 1476–1499
87 Birney, D.P. and Halford, G.S. (2002) Cognitive complexity of suppositional reasoning: An application of the relational complexity metric to the knight-knave task. *Think. Reas.* 8, 109–134
88 Halford, G.S. *et al.* (2002) Young children's performance on the balance scale: The influence of relational complexity. *J. Exp. Child Psychol.* 81, 417–445
89 Baillargeon, R. *et al.* (2010) False-belief understanding in infants. *Trends Cogn. Sci.* 14, 110–118