The Structure of Conceptual Competence

Chris Thornton
COGS/Informatics
University of Sussex
Brighton
BN1 9QH
UK
c.thornton@sussex.ac.uk

July 15, 2010

Abstract

This should now refer to rather than reproduce 'enumerative analysis of conceptual structure', ie the first half.

Conceptual abilities which seem to be significant in cognition include categorization, concept learning, prototype formation, abstraction, analogy use, schema acquisition, relation discovery, reasoning and abstract thought. Investigation of these abilities draws inspiration from a variety of sources and exploit a range of paradigms and methodologies. The study of concept learning, for example, is informed by ideas drawn from philosophy (e.g., the 'classical' model), by empirical and experimental work, by computer modeling and by formal analysis. Work on a particular competence normally makes use of a particular blend of approaches and is carried out in a dedicated subfield of research, tending to give the impression that there are fundamental discontinuities separating one competence from another. However, as the present paper shows, a significant subset of conceptual competences are very closely related. Classical concept learning, analogy exploitation, schema acquisition and reasoning can be shown to form a hierarchical sequence with well-defined subsumption relations holding between more complex competences (such as analogy use) and simpler ones (such as concept learning). On the basis of this, individual conceptual competences may be treated as special cases of generic, conceptual functionality.

Keywords: concept, concept learning, relation discovery, analogy, schema acquisition, theoretical cognitive science, cognitive informatics

1 Introduction

The necessity to discriminate forms of conceptual competence pervades many areas of cognitive science. In the study of conceptual development, there may be a need to define boundaries between stages, e.g., (Piaget and Inhelder, 1958; Karmiloff-Smith, 1992). In comparative studies of primates, there may be a need to discriminate human from non-human competences, e.g., (Penn *et al.* 2008), or the competences of one non-human species from another, e.g., (Thompson and Oden, 1998). In theoretical cognitive science, there may be a need to discriminate competences required for the completion of a certain reasoning or learning task (Clark and Thornton, 1997). In studies of creativity, there may be a need to define the competences required for creative activity, e.g., (Hofstadter, 1995; Boden, 2004). In studies of language evolution, there may be a need to delinate conceptual competences required for particular types of language use, e.g., (Fodor, 1975; Donald, 1991; Bickerton, 1996).

Investigators meet the need to differentiate conceptual competences in a range of ways. Use may be made of common-sense terms of the language. For example, reference may be made to the 'learning of features' (Smith and Medin, 1981) or 'use of metaphor' (Lakoff and Johnson, 1980), with dictionary definitions brought to bear. Reference may be made to competences relating to informally-defined processing models, e.g., acquisition of 'proxytype representations' (Prinz, 2002), or acquisition of 'symbolic' representations (Gardenfors, 2000). Competences may be characterized in terms of formally defined or computationally implemented processing models. For example, a competence might be characterized in terms of the aquisition of 'decision tree' representations as implemented in C4.5 (Quinlan, 1993), 'relation representations' in DORA (???) or in terms of acquisition of 'distributed concept' representations (Hinton, 1986).

In some cases, a conceptual competence may be characterized as the ability to aquire or make use of a particular entity within an ontology of concepts. Here again, the degree of formality may vary. Informal distinctions may be made between 'abstract' and 'non-abstract' concepts (e.g., Gardenfors, 2000), 'relational' and 'non-relational' concepts (e.g., Prinz, 2002) and 'object' and 'superordinate' concepts (e.g., Rosch, 1977). Reference may be made to intuitively characterized conceptual constituents such as 'attributes', 'dimensions' and 'properties' (e.g., Rosch and Lloyd, 1978).

In other cases, a formally-defined ontology may be invoked such as the 'conceptual structures' of (Sowa, 1984) or the compositional objects of 'formal concept analysis' (Ganter *et al.* 2005). Common in recent years has been the use of predicate logic as a surrogate concept ontology, e.g., (Gentner, 1983; Gentner and Kurtz, 2006). In this approach, concept models are equated with the structures which can be built in a particular form of predicate logic. Competences are then differentiated in terms of the ability to build or use particular forms of structure, e.g., (Holyoak and Thagard, 1995; Gentner and Kurtz, 2005).

A wide diversity of strategies are used to discriminate conceptual competence, then. But this raises a number of problems. Comparison of one competence against another may be infeasible if they are characterized in different ways or in terms of different ontologies. Equivalence or subsumption relationships between competences may be hard to discern. Where a competence is informally characterized — as in 'learning of features' (Smith and Medin, 1981) or 'reasoning over higher-order relations' (Penn *et al.* 2008) — there is often the potential for different interpretations to be applied (??? bbs commentators). Even where such a competence is identified with the operation of a computer simulation (e.g., Hummel and Holyoak, 2005) ambiguities may remain. Although the modeling of a competence in terms of a computer system provides both a working illustration and evidence of internal consistency, this falls short of formalization. The size and complexity of such systems generally allows for modeled competences to be interpreted in several ways. Even the guarantee of internal consistency must be taken with a grain of salt: non-trivial computer implementations are rarely bug-free.

Recourse to use of concept ontologies is also problematic. There is controversey over what such an ontology should contain (Fodor, 1998; Gardenfors, 2000; Prinz, 2002). Formally defined ontologies (Sowa, 1984; Ganter and Wille, 1999) have the benefit of precision but may be seen as obscuring critical distinctions. For example, formal concept analysis (Ganter *et al.* 2005) does not clarify the distinction between low-order and higher-order concepts. Use of predicate logic as a surrogate ontology has several advantages but also incurs drawbacks. Competences are deemed to involve use/acquisition of particular concepts, which are taken to be (interpretable as) logical structures. Objects are typically represented as variables and features/categories as 1-place predicates. Relations are then treated as multi-place predicates. This allows aquisition of higher-order concepts to be modeled in terms of construction of higher-order predicates. But there is then no clear way to formalize a 'superordinate' comptence, such as 'reasoning over higher-order relations' (Penn *et al.* 2008). Predicate logic offers no natural way to discriminate between categories of relations and higher-order relations. The natural way to represent a category of relations (i.e., as a higher-order predicate) is also the natural way to represent a higher-order relation.

There is no completely satisfactory approach, then, for differentiating conceptual competence. Rather, a diversity of strategies are applied. These vary in their formal precision and may be difficult to relate one to another. In some cases they entail contradictory or antagonistic models, as in the 'classical' v. 'probabilistic' contrast of (Smith and Medin, 1981). It is problematic to provide a firmly grounded classification of a particular competence and generally impossible to characterize precisely the relationship that holds between one competence and another.

What is wanted is a breakdown of the possible forms of conceptual competence that shows how those forms fit together. This should not be context-sensitive to any particular body of empirical data or philosophical tradition. It should map-out the space of possible functionalities in a way which does not presuppose any particular ontology of concepts or any equivalence with logic formalisms. It should serve to explicate subsumption relationships and be sufficiently detailed to enable classification and analytic comparison of models. Provision of a taxonomic scheme which meets these desiderata is the main goal of the present paper. The subsidiary goal is use of the scheme for purposes of highlighting the hierarchical relationship that exists between a core subset of competences.

The argument is set out in three stages. In Section 2, an enumerative/combinatorial analysis is used to divide up the space of possible conceptual structures according to their functional and structural characteristics. Examples are provided to illustrate how particular conceptual structures populate the space. In Section 3, a metric is defined for conceptual 'productivity', i.e., the representational coverage of a structure viewed as a proportion of its implementational cost. This is then deployed to characterize productivity-oriented exploration of conceptual-structure space as an abstract supercompetence. In Section 4, a set of conceptual competences are located within the space and equivalence and subsumption relationships are highlighted where relevant. The final section (Section 5) provides a summary and some concluding comments.

2 Enumerative analysis of conceptual structure

It is widely accepted that concepts are normally nested structures: any one concept can typically be seen as constituted of several others (Smith and Medin, 1981; Prinz, 2002). An ability to form new concepts can then be viewed as the ability to construct new concepts from existing constituents, i.e., other concepts. But care needs to be taken in distinguishing constructions which combine constituents as alternatives, and constructions which bind them into a relation. Special terms will be used to flag these two cases. Constructions which combine constituents as alternatives will be termed *categorical*. Constructions which combine constituents into a relation will be termed *compositional*.

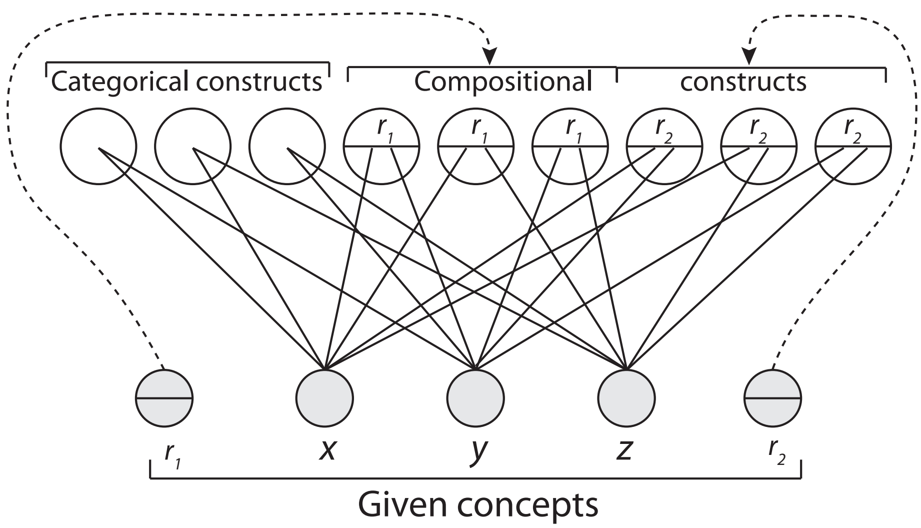

Figure 1: Categorical and compositional concept-combination.

Taking this distinction into account, it is then possible to enumerate the constructs which can be formed from any basis of given concepts. For each subset of givens, a categorical construct can be formed that combines the givens in the subset. If there are some relational concepts among the givens, there are also those constructs which can be formed by combining one of the relations[1] with one of the possible subsets. Figure 1 illustrates the combinatorial possibilities. Five concepts are assumed to be given, comprising two relational and three non-relational concepts. These are shown as circles at the bottom of the figure. The non-relational concepts (x, y and z) are shown in the middle of the bottom row. The two relational concepts (r_1 and r_2) are on the outside.

Circles in the top row then represent constructs which can then be formed. The three constructs on the left (top row) are categorical. These are cases in which the constituents are treated as alternatives. In this case, connecting arcs from the constituents are brought together at a point. Next to these we have the constructs in which constituents are bound into a relation: first, the three constructs obtained by application of relationship r_1; then, the three formed by applying r_2. The arcs here connect to a horizontal bar labeled with the relation applied. The figure enumerates all constructions which combine exactly two (non-relational) constituents. Given that there are two relations and three ways to form pairs out of the three non-relational constituents, there are nine possible constructions in all.

The graphical approach taken here is followed throughout the paper. Circles are used to represent concepts. Connecting arcs show whether the concepts are categorically or compositionally formed. In the case of categorical construction, arcs are brought together at a point, indicating alternation. In the case of compositional construction, the arcs abut a horizontal line labelled with the relation applied. The distinct characteristic of a compositional construct is thus the inclusion of a labeled bar. This differs from the approach taken in graph-based logic frameworks such as (Sowa, 1984) which use circles to represent relations and boxes to represent concepts. It also deviates from the custom in some analogy research, which is to use circles to represent categorical constructions and triangles to represent compositional constructions. (e.g., Doumas *et al.* 2008, p. 7). The approach taken here better serves the goal of representing the results of recursive construction, as will be seen.

2.1 Illustration

The distinction between categorical and compositional construction provides a way of enumerating the possible constructs which can be formed directly from a set of given concepts. It also models the difference between concept construction which is specifically generalizing in nature (e.g., abstraction and category formation), and construction which is more related to schema-formation or modeling. Generalization involves the creation of an identify for a set of alternatives: this is precisely the effect achieved in categorical construction. Shema formation involves creation of a structure in which certain parts are assembled into a certain relation. This is the effect achieved in compositional construction.

[1] The terms ‘relation’ and ‘relational concept’ are used interchangeably.

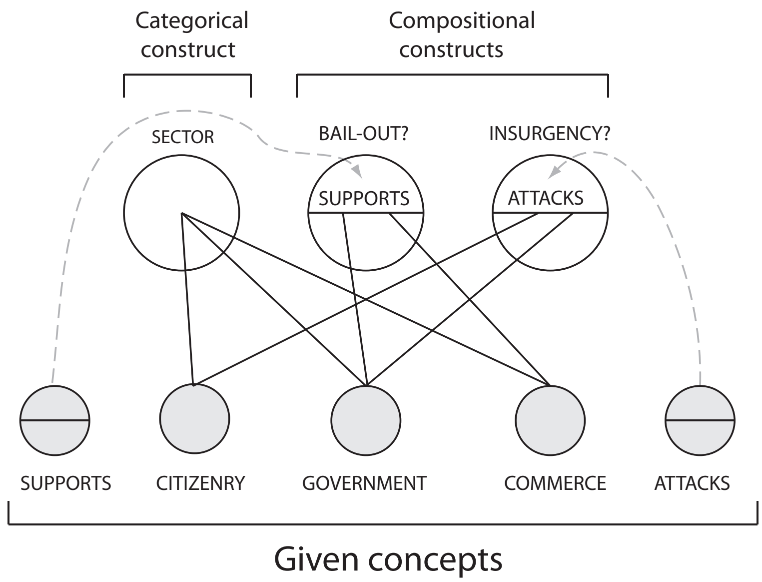

Figure 2: Examples of categorical and compositional construction.

As an illustration, consider Figure 2. This depicts a scenario in which the given concepts are taken from a domain of social institutions and relations. The given non-relational concepts are CITIZENRY, GOVERNMENT and COMMERCE. These are shown in the bottom row. The given relational concepts are SUPPORTS and ATTACKS. These appear in the bottom row on the outside. Possible constructions correspond to the circles in the top row, as before.

A categorical construction on all three, non-relational concepts can be viewed as forming the generalization SECTOR, i.e., the abstract notion of a sector of society. A compositional construct which applies SUPPORTS to GOVERNMENT and COMMERCE might be a way of forming BAIL-OUT, i.e., the idea of government covering financial losses in the commercial sector. Applying the relation ATTACKS instead might be a way of constructing the concept CLAMP-DOWN. Applying the relation ATTACKS to CITIZENRY and GOVERNMENT might be a way of constructing the concept of INSURGENCY, while application of the relation SUPPORTS to the same constituents might be a way of constructing TAXATION. These examples illustrate the ways in which categorical and compositional construction can mediate forms of generalization and schema formation. However, it continues to be simply the combinatorial possibilities which are depicted. There is no sense in which the constructions in this example are *informed* by the semantic properties of the labels.

2.2 Hierarchical construction

So far we have considered constructs which can be generated directly from given concepts. With a reasonable number of these, there is the the potential for a large number of constructions.[2] However, these remain the 'tip of the iceberg'. The result of any conceptual construction is necessarily a new concept, with the potential to be treated as a constituent in further construction. Beyond those structures which can be built directly, then, we also should take account of the possibilities for recursive (i.e,. hierarchical) construction.

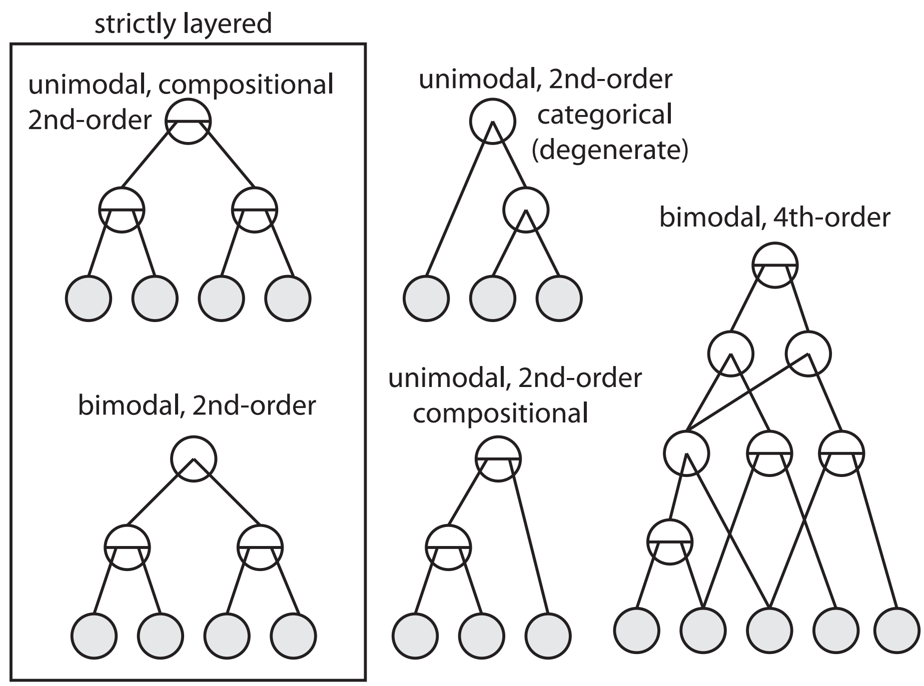

Figure 3: Unimodal and bimodal variants of hierarchical structure.

Figure 3 exemplifies the possibilities and introduces some of the dimensions along which hierarchical construction can vary. The key way in which hierarchical structures can vary is in their *depth*, i.e., the number of constructive steps embodied. In characterizing this dimension, use will made of the notion of conceptual *order*. Concepts which are given will be termed '0th order'. Any concept which is constructed directly in terms of 0th-order concepts, will be termed '1st-order'; any concept constructed from 1st-order concepts will be termed '2nd-order' and so on. A complication is the possibility of constructions which combine elements of mixed order. However, normal practice will be fol-

[2] A lower-bound on the number is $n2^m$, where m is the number of non-relational primitives and n is the number of (any-arity) relations.

lowed: the order of a construct will be deemed to be one greater than the order of its highest-order constituent.

Another way in which hierarchical conceptual structures can vary is in *modality*. We can have unimodal structures utilizing a single type of construction, or bimodal ones utilizing both types. Here, again, there is a minor complication. This relates to the class of categorically unimodal hierarchies. These are implicitly degenerate in the sense that they logically reduce to a non-hierarchical form. Recall that categorical construction treats constituents as entities which have no relationship with each other. On that basis a hierarchical construction constituted solely of categorical constructs is equivalent to a non-hierarchical construction which subsumes the same primitives. To illustrate, a categorical construction on concepts b and c where b is a categorical construction on $\{x, y, z\}$ and c is a categorical construction on $\{y, z\}$ is, by definition, equivalent to a single categorical construction on $\{x, y, z\}$. Hierarchical elaboration is thus redundant unless there is involvement of compositional construction.

A third way in which conceptual structures can vary is in their treatment of *derived* relational concepts. Where we have a categorical construction applied to compositional constructs on given relations, the resulting concept is itself inherently relational. Potentially, this new concept can be pressed into service in the formation of further compositional constructs. Such constructs may then be said to be *bootstrapped*: they are constructed using relational constructs which are themselves the result of previous construction. This is a technical use of the term 'bootstrapped', however. Kilverstein and Clark (2008) use the term more generally, as does Gentner, who comments 'Our results suggest that analogical encoding is a pervasive and important *bootstrapping* process.' (Gentner and Kurtz, 2005, p. 256).

A fourth way in which hierarchical structures can vary is in the strictness of hierarchical layering. Figure 3 illustrates this by making a distinction between the strictly layered constructions contained in the box on the left and the non-strictly layered concepts shown elsewhere. In a strictly-layered construct, any concept of order k is constructed solely from constituents of order $k - 1$. In non-strictly-layered constructs, constituents can be of any order less than that of the construct itself.

Taking these distinctions into account, it is possible to define the set of possible hierarchical constructions in a formal way. A convenient medium for this is a system of rewrite rules. The simplest case of the bimodal, strictly-layered, non-bootstrapped construction can be specified using the following rules:

$$C_i \Rightarrow \hat{C}_i | \bar{C}_i$$

$$\hat{C}_i \Rightarrow C^*_{i-1}$$

$$\bar{C}_i \Rightarrow \bar{C}_0, C^*_{i-1}$$

Here $\hat{C}_0$ is assumed to be any non-relational given concept and $\overline{C}_0$ is any relational, given concept. The first rule states that a concept of arbitrary order

i may take the form of either a categorical ($\hat{C}$) or compositional ($\overline{C}$) construct. In the former case, the construct consists of any number of constructs whose order is one less than than i. In the latter, it is defined as being any combination of a 0th-order relational concept and some collection of constructs whose order is one less than i. Note that the vertical bar is used here to indicate disjunction and the star superscript is used to denote two or more instances.

The more general case, which allows for both non-strict layering and bootstrapping can be specified as follows.

$$C_i \Rightarrow \hat{C}_{j<i} | \bar{C}_{j<i}$$

$$\hat{C}_i \Rightarrow C^*_{j<i}$$

$$\vec{C}_i \Rightarrow \bar{C}^*_{j<i}$$

$$\bar{C}_i \Rightarrow \bar{C}_0 | \vec{C}_{j<i}, C^*_{j<i}$$

Here, subscripts are defined so as to allow for constructs that are constructed in a non-layered way. There is also an additional rule (the third rule) which deals with the construction of bootstrapped relations.

2.3 Illustrations

As an illustration of the way in which hierarchical construction might lead to

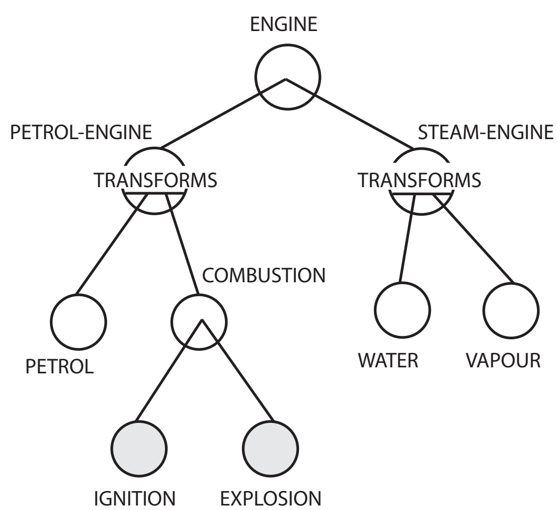

Figure 4: Hierarchical construction for an ENGINE concept.

the generation of meaningful concepts, consider Figure 4. Here the concept ENGINE is modeled as a 3rd-order, bimodal, strictly layered construct. The only relation used is TRANSFORMS. This is given rather than derived, meaning that ENGINE (in this case) is not bootstrapped. The given concepts, IGNITION and EXPLOSION, are combined categorically to form COMBUSTION. The given concepts WATER and VAPOUR are combined with the TRANSFORMS relation to produce STEAM-ENGINE. Applying the same relation to given concepts PETROL and derived concept COMBUSTION produces PETROL-ENGINE. PETROL-ENGINE and STEAM-ENGINE are then combined categorically to produce ENGINE.

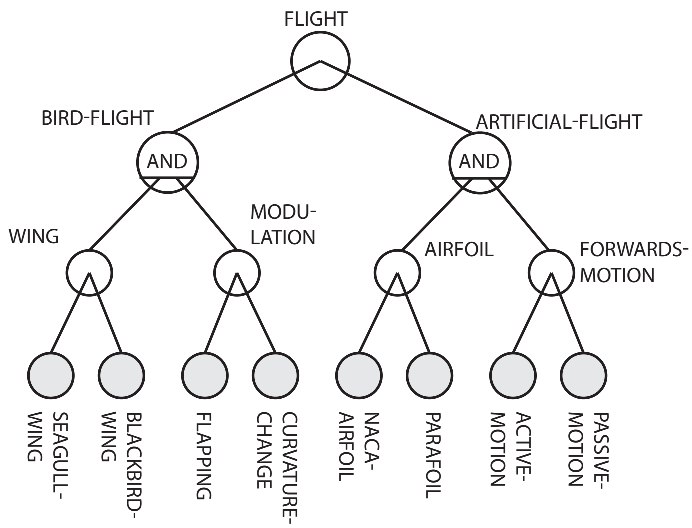

Figure 5: Bimodal, 3rd-order construction of FLIGHT.

Figure 5 provides another illustration. Here, the FLIGHT concept is formed using the AND relation and constructions on (given) concepts representing wing-shape and motion. FLIGHT in this case is portrayed as a bimodal construction of a 3rd-order, categorical entity. The immediate constituents of FLIGHT are the 2nd-order concepts BIRD-FLIGHT and ARTIFICIAL-FLIGHT. These are themselves constituted of categorical constructs on the given primitives.

2.4 Higher-order abstraction

Under the enumerative analysis of conceptual structure, any categorical construct formed solely from categorical constructs is degenerate — it is equivalent to one, all-encompassing categorical construct. This rules out the possibility of 2nd- and higher-order conceptual structure consisting solely of categorical constructs. Compositional constructs are therefore required to provide the 'step up' for higher orders of development. Figure 6 visualizes what this means in practice. All circles here (of all sizes) represent concepts as usual. The figure shows a hierarchical structure of five levels, but with constructive arcs ommitted

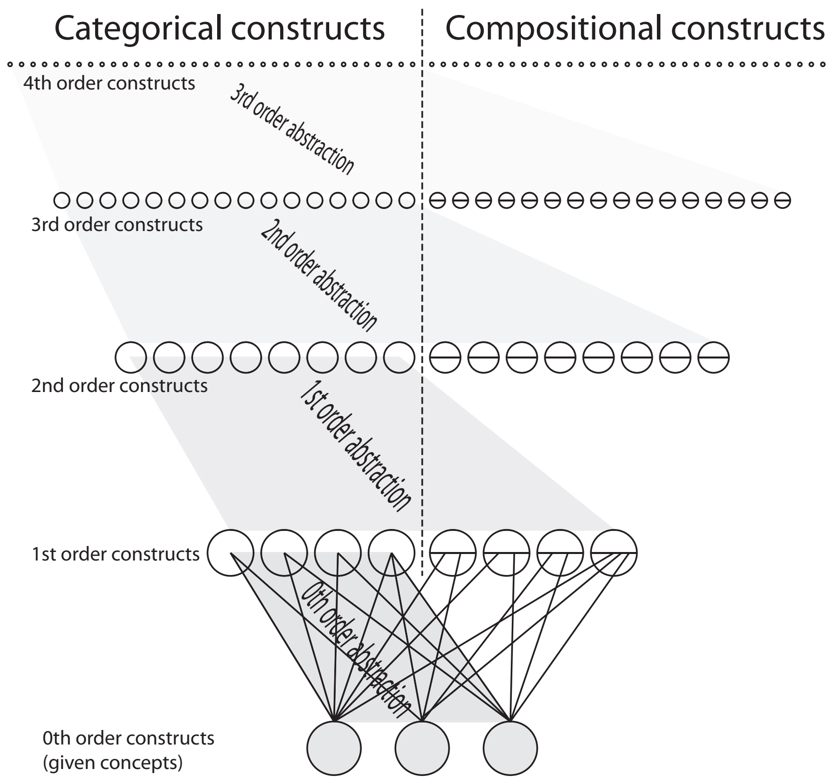

Figure 6: Compositional constructs viewed as the 'step up' for higher-order abstraction.

above first-order. The point to note is that it is the *compositional* constructs at one level of the hierarchy which provide the basis for constructions at the level above. The trapezoidal, shaded region above each layer of concepts represent the constituencies involved in these critical constructions.

As has been noted, categorical construction gives an identity to a set of alternatives and therefore models generalization and abstraction. But under hierarchical development, the abstractive effect becomes cumulative in nature. Categorical construction applied to 1st-order constructs provides an abstractive effect which is specifically 1st-order (it applies to 1st-order objects). Categorical construction applied to 2nd-order constructs provides an abstractive effect which is specifically 2nd-order, and so on. Hierarchical development thus serves to superimpose one level of abstraction over another.

Higher-order constructions necessarily provide higher-order abstractions. (Subscripts exhibit an offset, however. An nth-order construct provides an $n-1$th-order abstraction). A higher order of abstraction inevitably yields a more complicated mapping between the top and the bottom of the hierarchy. Where a concept is formed as a categorical construct of given concepts, the concept can be seen as constituted of the subsumed entities. But even a single level of abstraction obstructs intuitive interpretation. In terms of Figure 4's illustration of the ENGINE construct, we can readily see COMBUSTION as constituted of IGNITION and EXPLOSION. But the abstractive effects produced by the 2nd-order compositional constructs mean that it is more difficult to envisage the way in which ENGINE is 'constituted of' VAPOUR, or PETROL-ENGINE is 'constituted of' IGNITION. Similar remarks apply to the relationship between the FLIGHT construct of Figure 5, and the given concepts for that example.

Boostrapping makes the situation even worse. Boostrapping produces generalizations of relation-applications which then figure in subsequent, compositional construction. The process is abstractive in itself. But is also paves the way for subsequent compositional construction which makes use of abstracted relations. In this sense, bootstraping is *doubly abstractive*.

Even with boostrap-free structure, the grounding of top-level concepts in bottom-level entities may be counter-intuitive. Where bootstrapping *is* involved, there may be no way to form an intuitive interpretation of the way in which a higher-order concept is constituted in the primitive entities. Only when the relevant categorical and compositional constructions are taken into account (along with the way in which they interact to produce abstractive structure) does the top-to-bottom connection become comprehensible.

For example, consider the concept CAUSE. This has long been cited as a case where it is extremely difficult to see how there could be any grounding in perceptual or sensory data (as empiricist epistemologies might suggest). However, invoking the notion of higher-order abstraction, sensorily-grounded constructions for this concept can be envisaged. Figure 7 sketches out a possible construction. In this scheme, CAUSE is portrayed as a categorical construct of at least 6th-order. One of its constituents incorporates a relation that is bootstrapped from 4th-order constructs. (Ellipsis is used in various places to indicate missing but inferrable structure.) The SM1, SM2 constructs are taken to be primitive sense data, or associations of sense and motor data. These figure in compositional constructions relating to basic, perceptual entities and motor activities (e.g., SMALL-OBJECT and ARM-MOTION). Compositional constructions are then formed for a variety of behavioral contingencies (e.g., relation CONNECT applied to ARM-MOTION and SMALL-OBJECT). Categorical constructions on these produce relational abstractions such as LAUNCH, which then provide the basis for bootstrapping at a higher level. The top-level construct for CAUSE is then constructed categorically from applications of bootstrapped relations.

In this scheme, abstractive effects of different types are brought into play at various levels. The general effect is a a cumulative distancing of top from bottom. Any attempt to fomulate an intuitive interpretation of the way in which the top-level construct is 'constituted of' the givens, is doomed to failure. There is no informal sense in which CAUSE can be considered to be 'constituted of' SM1, SM2 etc. To appreciate the form of the constituency, it is necessary to take account of the extent and character of the embodied constructive operations.

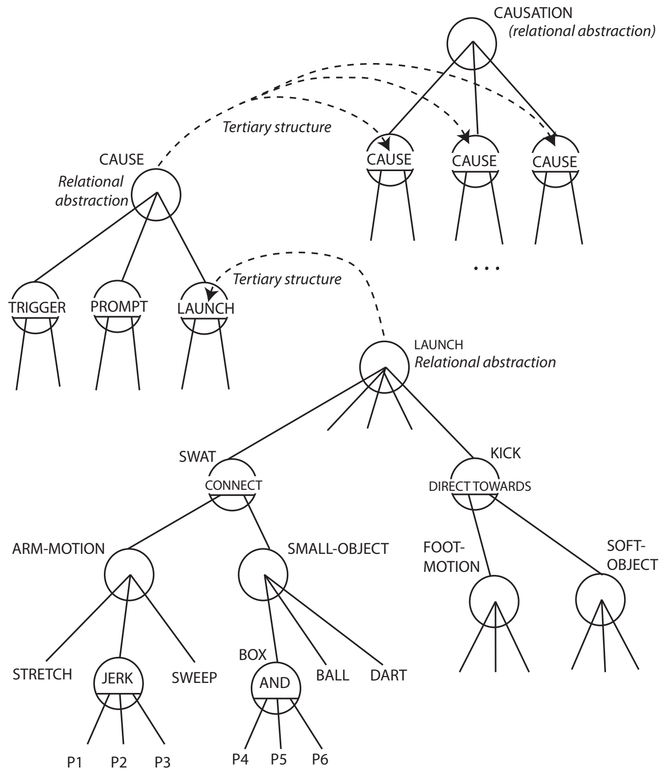

Figure 7: Sketch of a structure for CAUSE.

3 Instantiation, activation and productivity

The initial task of the paper has now been completed. Combining the principle of concept-nesting with the categorical/compositional distinction, enumerative analysis has been used to break-up the space of conceptual structures into clearly defined parts. Within the scheme, concepts may be categorically or compositionally constructed. Conceptual structures can be unimodal or bimodal and may or may not feature bootstrapping. Particular constructions also have a well-defined order. The requirement that the taxonomy should be context-free has also been met. Derived solely on the basis of first principles, the scheme does not relate to, or depend on any particular empirical approach.

A point that emerges from the analysis is that 'order' is a strictly relative attribution. In terms of the taxonomy, we can never specify the order of a concept in an absolute way. We can only say that it has a particular order relative to a particular constructive context. This point is also stressed by Palmer, who notes 'the order of a relation is not God-given or dictated by logic alone but something that results from representational choices' (Palmer, 1989,

p. 340).)

Another point emerging from the analysis relates to abstraction. While it makes sense to equate abstraction in its most basic form with categorical construction, we have seen how compositional and categorical construction may interact to produce a cumulative process. On this basis, it becomes possible to discriminate orders of abstraction. Account then needs to be taken of how higher-order abstraction may be associated with constitutive mappings which are opaque and counter-intuitive.

The most significant point to come out of the analysis, however, relates to instantiation. Normally, we take the *instances* of a concept to be things 'in the world', e.g., objects, properties or features of some domain of real phenomena. But in the enumerative analysis, this idea becomes problematic. Concepts may be categorical in nature, in which case the instances would seem to correspond to the constituents of the construct. But, in general, these are themselves conceptual structures that may embody further categorical constructs.

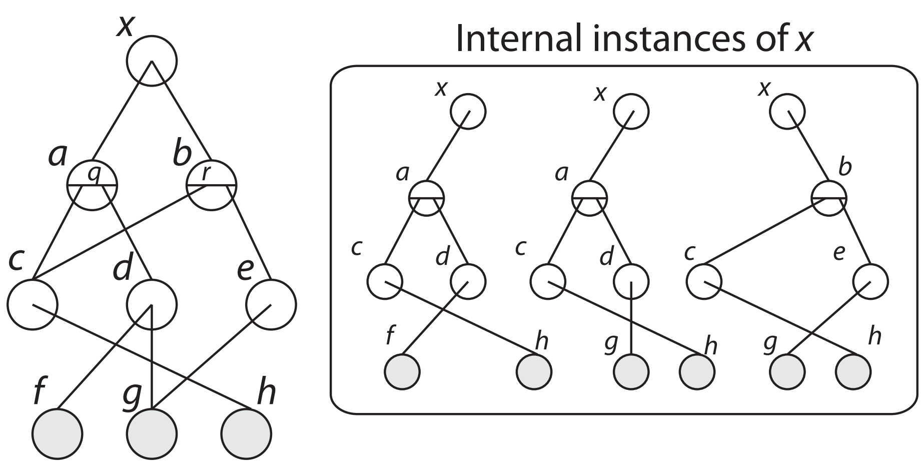

Figure 8: Instances as ANDed subtrees.

To clarify the situation, a distinction is introduced between external and internal instances, with the former being the aforementioned 'in the world' entities and the latter being the ways in which a particular concept can manifest itself *within* conceptual structure. The form that a concept takes (in the enumerative analysis) is, in general, a hierarchy of categorical and compositional constructions. Categorical constructions are disjunctive — they represent alternation — while compositional constructions are essentially conjunctive. To a first approximation, then, a concept forms an AND-OR tree; the ways in which it can manifest itself internally correspond to the ANDed subtrees that the AND-OR tree encapsulates.

An illustration appears in Figure 8. On the left, we see the categorical/compositional structure (AND-OR tree) for concept x. To the right we then have the three encapsulated AND trees which form its internal instances. Building on the idea of an internal instance, we can define order-specific *projections*, the set of *concepts* of a particular order contained with a particular instance. For example, $\{g, h\}$ is the set of 1st-order concepts involved in the left and middle instances of x, while $\{g, i\}$ is the set of 1st-order concepts involved in the rightmost instance. The complete projection of concept x on 1st-order concepts is thus $\{\{g, h\}, \{g, i\}\}$.

In what follows, sets of internal instances will be denoted by placing a double dot over the concept label. Thus $\ddot{x}$ identifies the set of all internal instances (ANDed subtrees) of concept x. Projections will then be specified by appending an integer superscript to the double-dotted identifier. In general, it is convenient to specify relative rather than absolute projections. The projection of x on 1st-order concepts is thus described as x's 2nd-order projection (because there are two intervening levels of constitution) and denoted by appending an appropriate superscript to the double-dotted identifier. The 2nd-order projection of x in Figure 8 is written

$$\ddot{x}^2 = \{\{c, d\}, \{c, e\}\}$$

while the 1st-order projection is

$$\ddot{x}^1 = \{\{a\}, \{b\}\}$$

An important case is the *base projection.* This is the projection on 0th-order (i.e., given) concepts. It is denoted using a dash superscript:

$$\ddot{x}^- = \{\{f, h\}, \{g, h\}\}$$

Determination of projection provides a way to connect internal with external instances. For this exercise, D will represent the mapping between 'in the world' phenomena and internal concepts. Specifically, D will be a set of sets, such that each subset contains the given concepts corresponding to some external phenomenon. The external instances of some concept x can be identified with the intersection that D makesx with x's base projection. Members of this intersection are those patterns of primitive instantiation that correspond to external phenomena.

This way of handling the relationship between internal and external instances also enables measurement of explanatory *productivity*, i.e., the degree to which a particular conceptual structure covers phenomena of a domain. Here, the notation $|c|$ will be used to denote the total number of concepts encapsulated in concept c. In the case of concept x in Figure 8 we have

$$|x| = 8 + 2 = 10 \tag{1}$$

since x encapsulates 8 non-relational concepts and two non-relational concepts. (The two relational concepts q and r are not shown explicitly in the diagram.)

The explanatory productivity $v(c)$ of concept c may then be defined as

$$v(c) = \frac{|\ddot{c}^{-} \cap D|}{|c|} \quad (2)$$

where D is the set of sets representing domain phenomena. This defines productivity to be the number of external phenomena contained in the base projection, viewed as a proportion of the total number of concepts utilized. In simple terms, it defines representational value as a proportion of representational cost.

In the case of Figure 8, concept x embodies 10 concepts in all with a base projection of size 2. Letting D be

$$D = \{\{f, g\}, \{g, h\}\}$$

and recalling that

$$\ddot{x}^{-} = \{\{f, h\}, \{g, h\}\}$$

we then have a situation where x covers one of the domain phenomena, at a cost of 10 concepts in all. This yields an explanatory productivity of

$$v(x) = \frac{1}{10}$$

This highlights the fact that a relatively small number of domain phenomena are covered at a cost of a relatively large number of concepts.

4 Mapping the space of constructive competence

Up to this point, the paper has been primarily concerned with development of a taxonomic framework. Enumerative analysis has been used to sudivide the space of conceptual structure into well defined parts. Necessary revisions to the notion of concept instantiation were then formulated and a measure of explanatory productivity was introduced. In this final section, the paper aims to make use of this framework for purposes of showing how a range of conceptual competences fit together.

In this exercise, attention will focus exclusively on *constructive* competences (i.e., competences relating to the construction of conceptual structure) rather than on applicative ones (i.e., competences relating to the use of conceptual structure). (Analysis of applicative competences will will form the topic of a future paper.) For each competence considered, the aim will be (a) to locate the structure(s) it builds within the enumerative taxonomy and (b) to consider the degree to which the processing it entails is is modeled by productivity maximization.

5 Generalization and category formation

At the start, it is worth reviewing the situation with regard to the competence of generalization. As already noted, this involves creation of an identify for a set of alternatives and is thus precisely the effect achieved by categorical construction. Within the taxonomy, then, generalization (and category-formation) equates to the act of categorical construction. The situation relating to abstraction is more complex, however. As noted, the process of abstraction may be treated as identical to the process of generalization, on which basis we would equate it with categorical construction. However, as has been seen (in Section 2), the taxonomy provides the means of discriminating different orders and degrees of abstractive effect. It makes more sense, then, to say that it is only abstraction in its most elementary form which equates to categorical construction.

5.1 Classical concept learning

Fundamental among constructive, conceptual competences is the functionality of 'concept learning': the acquisition of new concepts on the basis of experience. The long-standing model is the Aristotelian or 'classical' notion of concepts being formed on the basis of definition of necessary and sufficient features. In Smith and Medin's view (Smith and Medin, 1981, p. 23), the heart of the classical view is that 'The features that represent a concept are (1) singly necessary and (2) jointly sufficient to define that concept.' The distinguishing characteristic of a classical concept is that it defines a conjunction of necessary and sufficient features.

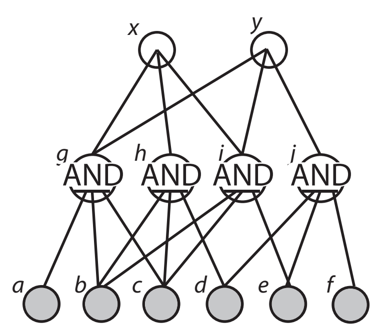

Figure 9: Classical/prototype definitions viewed as 2nd-order constructs.

Within the enumerative taxonomy, definition of a conjunction requires use of a compositional construct using the AND relation. On the basis of treating features as domain-instantiated primitives (as described in Section 3), a classical definition must then constitute either a 1st-order, AND-based, compositional construct or a categorical construct formed from concepts of that type. The latter possibility is illustrated by concept x in Figure 9. As a categorical construct of three AND-based, compositional constituents, x generalizes the relevant sets of primitives. The generalization of g, h and i, x thus implicitly defines the subset c and d. Concept x has the form of a classical concept definition and the act of constructing a classical definition can be equated with the act of constructing a 2nd-order categorical generalization of 1st-order, AND-using, compositional constructs. This shows where classical definitions are located within the enumerative taxonomy. But to what extent is the competence of learning such definitions modeled by productivity maximization?

Treating this competence as the identification of that set of features which are held in common by a set of instances, and treating instances as combinations of features, a categorical combination of AND-based constructs *corresponding* to those combinations is necessarily the least complex conceptual structure representing the domain. The concept must then have the highest achievable explanatory productivity. On that basis, it will be prioritized by any productivity maximizing process. The competence of learning classical concept definitions is thus modeled by productivity maximization.

For example, assume that the domain relating to the structure of Figure 9 is

$$D = \{\{a, b, c\}, \{b, c, d\}, \{b, c, e\}\}$$

Note that

$$\ddot{x}^{-} = \{\{a, b, c\}, \{b, c, d\}, \{b, c, e\}\}.$$

It then follows that

$$v(x) = \frac{3}{8} = 0.375$$

Given inclusion of the five primitive and one compositional construct for each exhibited combination, a concept representing the entire domain can use embody no less than eight concepts. On this basis, the productivity of x is the maximum that can be achieved. Construction of x will then be the result of productivity maximization.

5.2 Theoretical evaluation of 'non-classical' concepts

One objection made against the classical model is that there are concepts which appear to be *impossible* to define in terms of any set of necessary and sufficient features. In Wittgenstein's widely noted example, there can be no classical definition for the GAME concept, since there is no set of necessary features which all games have in common (Wittgenstein, 1958). It is worth noting, however, that in terms of the enumerative taxonomy, such assertions are inherently ambiguous. Features can only correspond to sub-concepts appearing in one or more projections. The assertion that a concept has no set of necessary and sufficient features must thus be positing an empty *intersection* of projection sets. But the validity of this necessarily depends on what order of projection is assumed. If the assertion is made without specifying order, there is the possibility that it might be true at one order and false at another. Indeed it is possible that there might be a set of orders for which the asssertion is true and a set for which it is false.

An unambiguous reformulation of the claim might be that, where a concept is represented as a 2nd-order categorical construct over compositional elements, base projection sets always have an empty intersection. (This is probably what Wittgenstein intended.) A stronger and more general claim might be that, regardless of the order of the concept's representation, projection sets at every available order have empty intersections. Whether any such claims can be valided in practice is not clear. But for statements of this type to have an unambiguous meaning, assumptions relating to projective order must be spelled out.

6 Similarity-based learning

While the process of identifying the set of features common to the instances of a concept provides a rudimentary concept acquisition method, a more powerful approach for this type of concept any many others entails some use of the similarity-based strategy. This has been envisaged to take various forms, e.g., the version space method of (Mitchell, 1977), focussing (Wielemaker and Bundy, 1985) decision-tree induction (Quinlan, 1983; Quinlan, 1993) and some neural network models, such as perceptron learning (Minsky and Papert, 1988). Its essential characteristic is maximization of the similarity of any (constructed) concept's instances. Regarding construction of the type of 2nd-order categorical construct here treated as constituting a classical definition, it can be shown that productivity-maximization will also tend to promote this strategy.

As an illustration, consider Figure 9. In this example we have

$$\begin{aligned} \ddot{x}^- &= \{\{a,b,c\},\{b,c,d\},\{b,c,e\}\} \\ \ddot{y}^- &= \{\{b,c,d\},\{b,c,e\},\{d,e,f\}\} \end{aligned}$$

On this basis, costs and projection sizes are

$$\begin{aligned} |x| &= 8 \\ |y| &= 9 \\ |\ddot{x}^-| &= 3 \\ |\ddot{y}^-| &= 3 \end{aligned}$$

It then follows that the productivity of x must be greater than the productivity of y:

$$\frac{|\ddot{x}^-|}{|x|} > \frac{|\ddot{y}^-|}{|y|} \Rightarrow v(x) > v(y)$$

The greater similarity between the primitive constituents of x increases this concept's productivity relative to y. Featural similarity among a set of instances will always tend to reduce concept usage in any over-arching categorical constructs. Productivity is then increased. Productivity-maximization in construction of this type of concept therefore conforms to the similarity-based strategy.

6.1 Analogical mapping

Another significant constructive competence is analogical mapping. This is the formation of structure which models the analogical relationship between two subsidiary structures. Figure 10 is the schematic used by Gentner (1983)

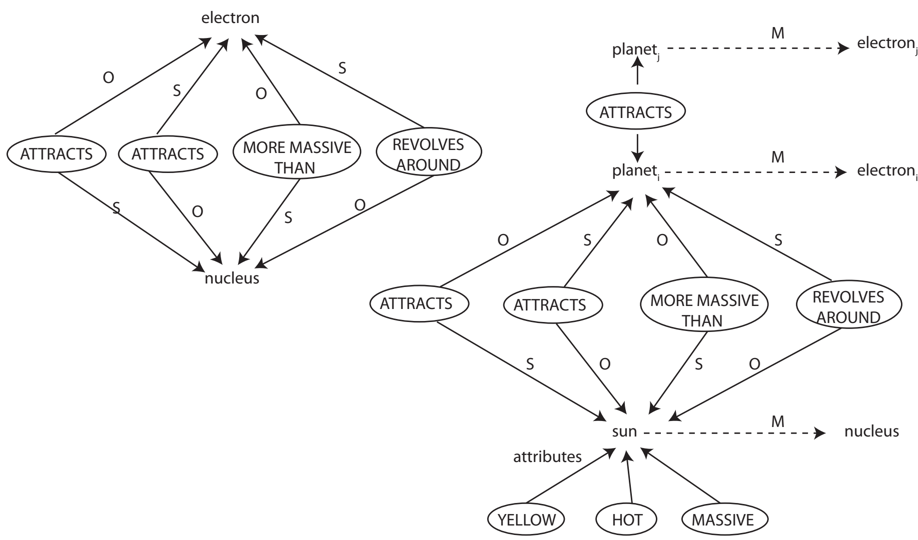

Figure 10: Gentner's model of the Rutherford analogy, after (Gentner, 1983, p. 160).

to illustrate how the mapping process works in the case of the atom/solar-system analogy. In Gentner's analysis, the solar-system structure is treated as the 'base' for the analogy, and the atom structure is treated as 'target'. The schematic shows how the structure of relations affecting planetary motion (i.e., the base) can be mapped onto those affecting electrons orbiting a nucleus (i.e., the target). The strength of the analogy between base and target depends on the correspondence between the first-order relations (ATTRACTS, REVOLVES-AROUND etc.).

Construction of the mapping serves to make connections between primitive objects in base and target (e.g., sun -> nucleus). It also has the potential to transfer relational knowledge from base to target. In this case, knowledge that planets attract each other is transferred (through the mapping) to the target, enabling the inference that electrons attract each other.[3]

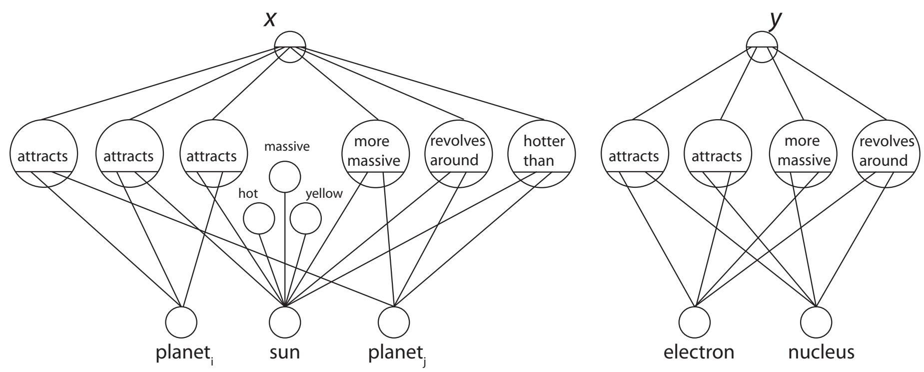

Figure 11: Atom/solar system analogy.

Figure 11 shows how base and target would be represented using the conventions of the enumerative framework. Properties are represented as categorical constructs and relations as compositional constructs. Base and target structures are treated as distinct concepts. The solar-system structure is treated as constituting concept x, and the atom structure as constituting concept y. Subject/object roles are also represented differently. In the Gentner's schematic, subject/object roles are explicitly labeled. Here, they are represented implicitly using left-to-right ordering.

In Gentner's model, formation of the mapping between base and target is accomplished using the 'structure-mapping engine' (SME) (Falkenhainer *et al.* 1989). In terms of the enumerative framework, this serves to form a categorical construct which *generalizes* the compositional constructs representing base and target. Mapping x onto y serves to equate constituents of x with corresponding constituents of y. This process of generalizing x and y equates to the formation of a categorical construct, as illustrated in Figure 12. The critical property of the over-arching categorical construct is that it is at least 2nd-order with respect to the highest-order compositional constructs (i.e., relations) involved in the mapping. In this particular example, the construct is 3rd-order. In general it might be of 3rd-order or above.

6.2 Abstractness and systematicity

The analogical mappings of structure-mapping theory equate to higher-order categorical constructs. The competence of forming such mappings involves building higher-order constructs with the relevant properties. In characterizing

[3]Holyoak calls this process 'copy with substitution and generation' or CWSG, noting that 'All major computational models of analogical inference use some variant of CWSG' (Holyoak, 2005, p. 128).

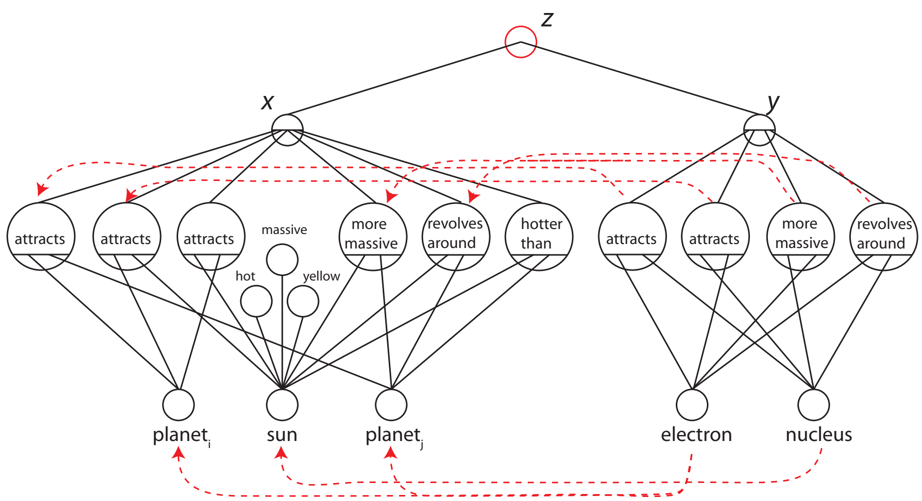

Figure 12: Analogy and transfer mediated by recursive generalization.

analogy-mapping, structure-mapping theory posits the *systematicity principle* (Gentner, 1983). This states that the strength of analogy depends on the degree to which systems of relations in the base can be mapped onto systems of relations in the target. Also posited is the *abstractness principle*. This is the idea that strength also depends on the level of abstraction at which correspondences are formed. The competence of analogical mapping is then characterized as a process which exploits both these principles. To what extent is this functionality modeled by productivity maximization?

Consider first the systematicity principle. A process applying this will be disposed to map one structure onto another just in case they exhibit relational commonality. The degree of preference will correspond to the degree of commonality. In the enumerative framework, relations are compositional constructions and mapping is mediated by formation of an over-arching categorical construct. Where there is a greater degree of relational commonality between base and target, there must be lesser use of relational concepts in the over-arching construct. Other things being equal, we then expect a higher conceptual productivity in any over-arching construct which encapsulates structures exhibiting greater commonality of relations. Maximization of productivity then models maximization of systematicity.

A similar argument can be used with regard to abstractness. Where a mapping makes correspondences between relations of higher order, these must subsume a greater degree of matching sub-structure. Where they are made between relations of lower order, they must subsume a lesser degree of matching sub-structure. A greater degree of matching sub-structure implies more commonality between base and target and a lower representational cost for the over-arching categorical construct. Productivity will always be higher where there is a greater degree of abstractness in the mapping made. Application of the abstractness principle within analogical mapping is thus also modeled by productivity maximization.

6.3 Schema acquisition and relation discovery

A competence often seen as related to analogy exploitation is schema-acquisition. In the terms of (Gentner and Colhous, Forthcoming) this is the task of finding 'common relational structure'. In the enumerative framework, it equates to the formation of higher-order generalizations of constituents which are *themselves* analogically constructed. In the context of Figure 12, this would mean construction of a concept which uses concept z (the generalization of x and y) as a constituent. For example, envisage a categorical construct which has z as a constituent. As a representation of the structural commonalities ('common relational structure') between z and the other constituents, this would form a relational schema. The capacity of progressive conceptualization to generate such schemas thus models schema-acquisition.

A construct such as z can also be viewed as generalizing a set of relations and relation applications. From this point of view, formation of z appears to be an act of relation discovery. But bringing to bear the concept of bootstrapping (see ???), there is the potential for discriminating different degrees of the process. Where there is greater utilization of bootstrapped relations in the innovation of an analogy-mediating categorical construct, we have a more substantive element of *discovery*. It becomes meaningful to treat schema-acquisition and relation-discovery as continuous, then, with the degree of relation discovery depending on the degree of bootstrapping. A system like DORA, described as 'the first detailed, computationally instantiated account of how comparison can serve to bootstrap the discovery and predication of structured relational concepts' (Doumas *et al.* 2008, p. 30), might then be interpreted as preferentially biased towards exploration of bootstrapped constructs.

6.3.1 Archimedes

An advantage of the analysis is that it treats functional and structural aspects of conceptualization in the same way. In principle, all aspects are functional. But since all functional competences are grounded within the structural operations of categorical and compositional construction, it then becomes possible to portray a progressive conceptualization process in purely structural terms. There is then the potential to 'draw' out the process which leads to the emergence of a particular conceptual result, e.g., the discovery of a particular analogy.

As an illustration, consider Figure 13. This depicts Archimedes' discovery of immersive volume measurment, i.e., the discovery that there is a way of treating immersion as analogous to volume measurement. In this schematic (an alternative interpretation is presented by Koestler, 1964, p. 106-7), construction of the analogy is mediated by formation of a categorical construct linking two subordinate structures. One of these conceptualizes the behavior of bath water;

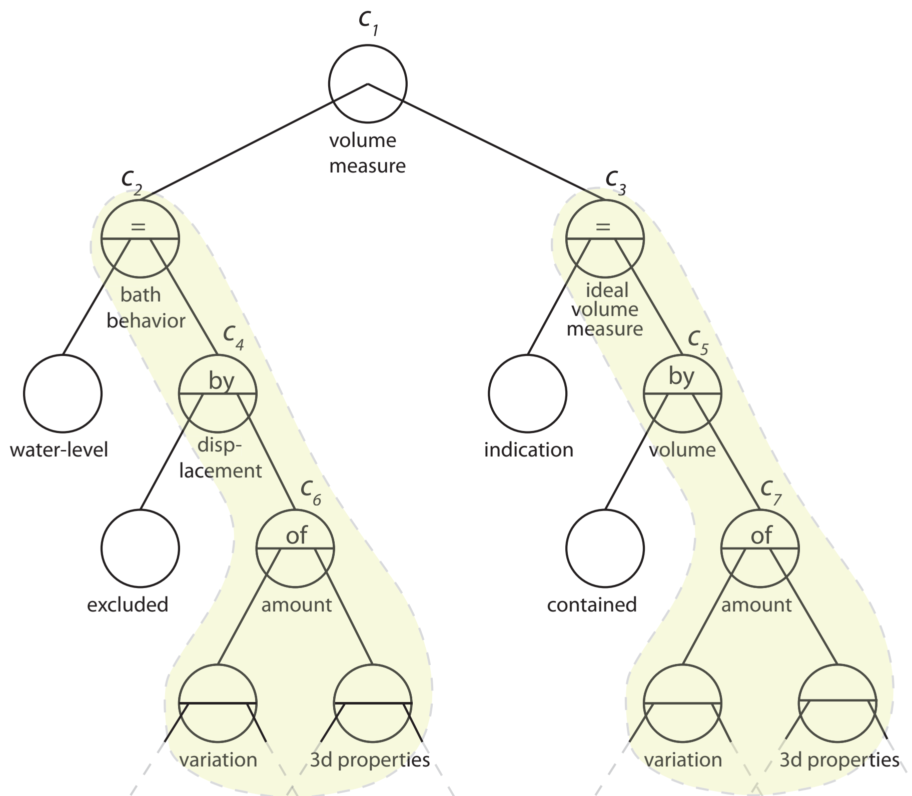

Figure 13: Archimedes's volume/immersion measurement analogy.

the other conceptualizes the behavior of an ideal measure of volume. Formation of the over-arching generalization then places displacement-measurement and volume-measuring into a single category, implicitly incorporating the fact that the former can be used to achieve the latter. The diagram captures both the mapping aspect of the operation (the mapping of c_2 onto $c3$) and the transfer aspect — the implicit substitution of 'water-level' for 'indication'. Correspondences between structures of relations ('systematicity') are indicated using shading.

6.4 Prototype-related competences

A problem with classical definitions is that they do not naturally account for typicality effects in human classification (Rosch, 1973; Rosch, 1975). Humans are generally willing to rate the typicality of a particular entity with regard to a particular class. In the familiar example, ROBIN is typically rated as more typical of the BIRD concept than PENGUIN. This seems contrary to the idea of a concept being represented in terms of necessary and sufficient features, a strategy which seems to ordain a rigid dichotomy between instances and non-instances.

For purposes of remedying this deficiency in the classical model, various forms of prototype theory have been put forward (Mervis and Rosch, 1981). Rather than posit that concepts are representated in terms of necessary features, these suggest representation is in terms of some kind of metric framework, such that particular instances can be rated according to their typicality for particular concepts. Too many competences have been described for any kind of case-by-case consideration. However, it can be shown that a basic 'typicality' competence does emerging naturally in a generalization of classical-concept use.

Maximum activation of a categorical concept implementing a classical definition (such as x in Figure 13) models the competence of accessing or using a classical definition. Where the domain exhibits an instance of the concept, precisely one of the compositional constructs is activated. This must fully activate the over-arching categorical construct. Maximum activation of this concept then implements application of the classical definition. Where the domain exhibits some close approximation of an instance, we expect a lesser degree of activation, with this corresponding to the degree of match between exhibited phenomena and one of the represented instances. Degree of activation thus implements a kind of typicality assessment, in the way proscribed by prototype theories. On this basis, graded activation can implement the rating of typicality.

This model does not assume, however, that typicaly rating process is based on measurment of the number of common features which an instances exhibits. Nor does it entail introduction of a capacity limit in the detection mechanism, as in the 'complexity model' of (Smith and Medin, 1981). Rather it allows for the fact that typicality rating is based on (implicit) assessment of the degree of structure which a phenomenon has in common with a represented instance. This is the factor which influences degree of activation of the over-arching categorical construct.

6.4.1 Pasteur

An illustration which more specifically highlights the competence of schema-acquisition is Figure 14. This represents Pasteur's discovery of vaccination. The background in this case is the story of Pasteur's experiments with chicken cholera. By chance, Pasteur had injected a number of chickens with an old and non-virulent sample of the bacillus of chicken cholera. When all the chickens survived, Pasteur repeated the experiment using more chickens and a more virulent strain of the bacillus. Unexpectedly, all the chickens in the original sample survived. Pasteur then reasoned that the survival of the initial group of chickens was analogous to the survival of those smallpox sufferers who had been exposed to the related disease 'cow-pox'. He reasoned that the analogy between the two cases was indicative of a general relational effect characteristic of immunization.

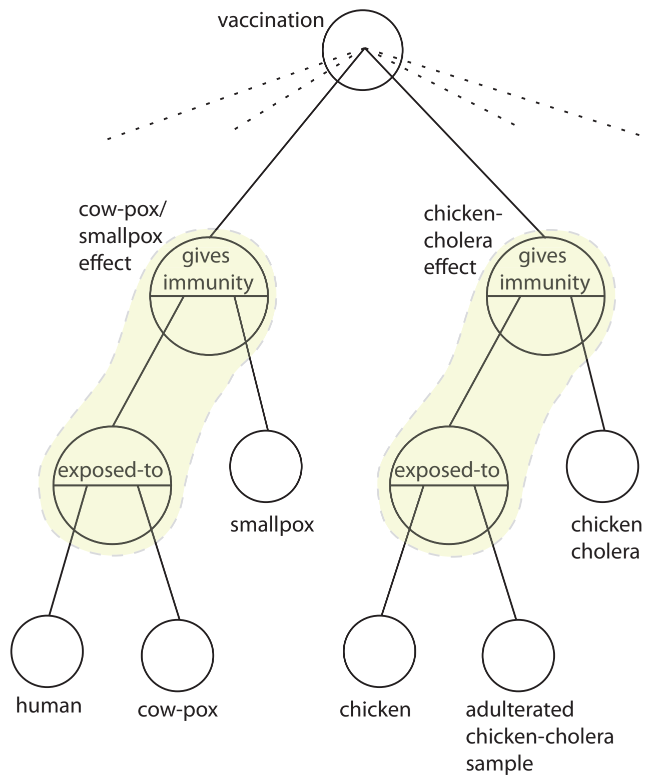

Figure 14: Pasteur's vaccination analogy.

In this discovery,[4] Pasteur formed an analogical mapping between the situation affecting smallpox/cow-pox immunity and the situation affecting cholera immunity in his chickens. The result was then a transfer of knowledge about the smallpox domain to the cholera domain. It can alternatively be seen as a case of schema acquisition: Pasteur developed a new relational schema relating certain types of exposure and certain types of immunity.

6.5 Transformational construction

For completeness, it is also noted that progressive construction has the potential to model competences involving 'transformation'. Where we have some existing hierarchy, we can create combinations of elements of different order (i.e., marking arbitrary selections from the hierarchy) and use them as the givens for a further phase of progressive conceptualization. One progressive conceptualiza-

[4]It seems appropriate to use the word 'discovery' in this context. As Koestler notes, it 'has been said that discovery consists in seeing an analogy which nobody had seen before.' (Koestler, 1964, p. 104).

tion process can thus be applied in different ways to the resutls of another. This functionality does not appear to correspond to any widely recognized competence. However, it clearly exists as a distinct possibility.

6.6 Reasoning

Conceptualization can also model elementary reasoning. For example, consider

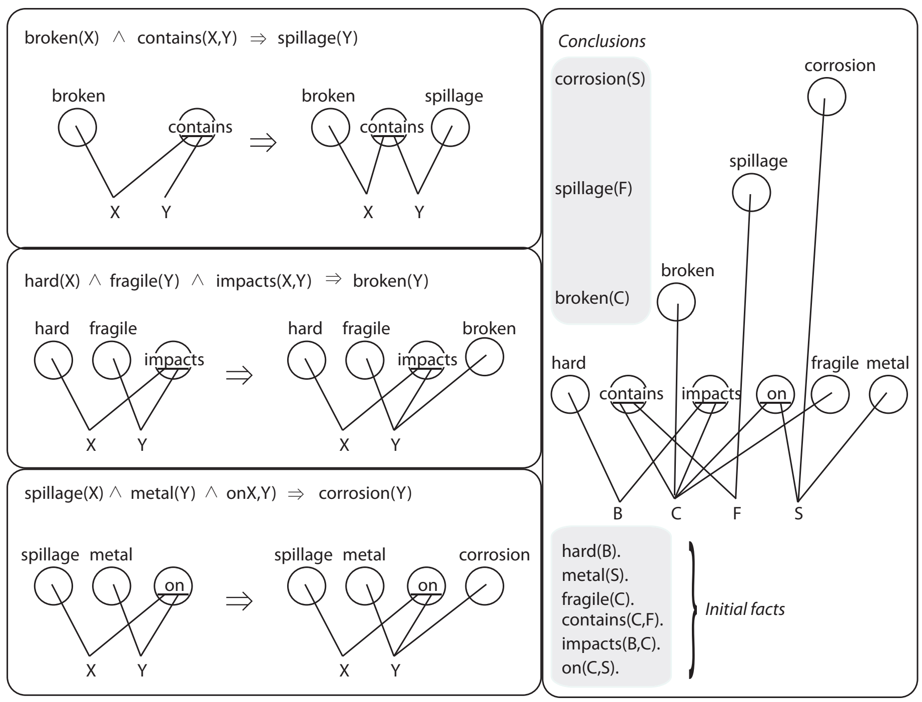

Figure 15: Reasoning viewed as concept construction.

the following rules:

$$\text{hard}(X) \wedge \text{fragile}(Y) \wedge \text{impacts}(X,Y) \rightarrow \text{broken}(Y)$$
$$\text{broken}(X) \wedge \text{contains}(X,Y) \rightarrow \text{spillage}(Y)$$
$$\text{spillage}(X) \wedge \text{metal}(Y) \wedge \text{on}(X,Y) \rightarrow \text{corrosion}(Y)$$

Given the facts $\text{hard}(B) \wedge \text{fragile}(C) \wedge \text{impacts}(B,C)$ we can then infer $\text{broken}(C)$, which in combination with the fact $\text{contains}(C,F)$ leads to the inference $\text{spillage}(F)$, which in combination with the facts $\text{metal}(S) \wedge \text{on}(C,S)$ yields $\text{corrosion}(S)$. But derivation of a particular conclusion (in a reasoning process) can always be viewed as the construction of an inferential structure. We can then view the reasoning in this example as a concept-construction sequence.

This is illustrated in Figure 15. The rules in this example refer to objects, surfaces and relations involved in 'spillage' events. Each panel on the left shows how a particular rule from the rulebase can be modeled in terms of an act of concept construction. The larger panel on the right then shows the complete reasoning process (using these rules/acts) in which the production of conclusions corresponds to acts of construction. Application of the construction rules

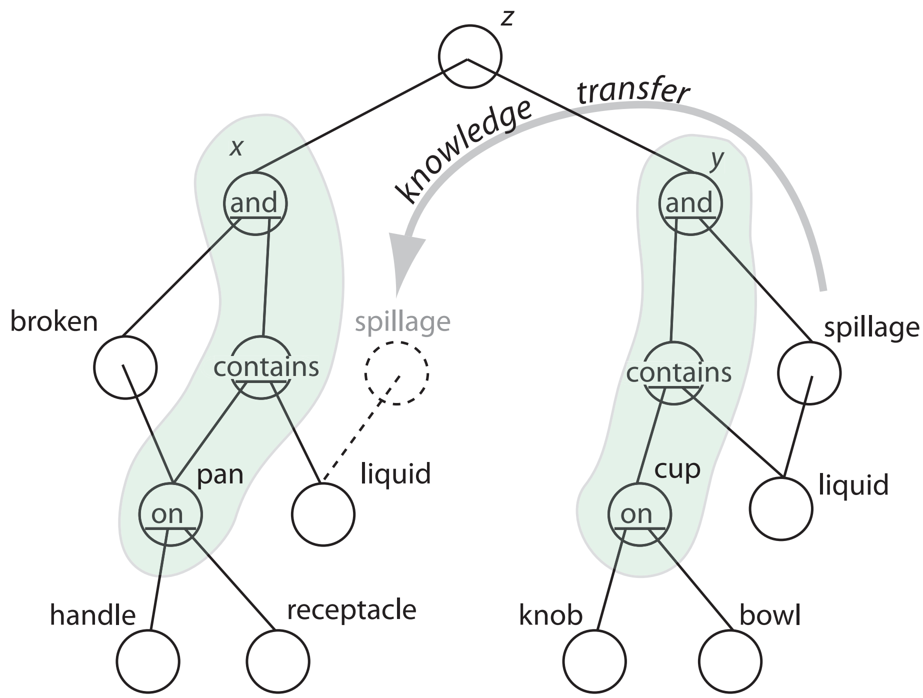

Figure 16: Inference viewed as knowledge transfer.

(illustrated in the left panels) involves the addition of a new concept to an existing structure. Thus inference can be modelled as a kind of knowledge transfer. Figure 16 illustrates how this might work in the case of the inference from the top-left panel. (The shaded areas here indicate corresponding relational structure.) Knowledge-transfer serves to add a new concept to an existing structure. This implements what we see as an inferential step in the reasoning. Knowledge transfer is able to model inference and progressive conceptualization is therefore able to model basic processes of reasoning.

7 Conclusions

Analysis of conceptual competences reveals their underlying, taxonomic structure. Competences that may have been viewed as forming contrastive or even antagonistic models, such as the classical and prototype models of conceptualization, are shown to be in a hierarchical relationship. The analysis also brings into focus some of the relationships suspected to exist, such as the relationship between analogy and similarity (Ramscar and Pain, 1996). As Gentner and Kurz note (2006, p. 613), 'there is considerable support for the claim that "similarity is like analogy", cf. (Gentner and Markman, 1995; Medin *et al.* 1993). The analysis puts this claim on a firmer footing.

The analysis also reveals that exploitation of analogy (through, e.g., structural alignment and knowledge transfer) can be understood as a recursive form of similarity-based learning. Again, this chimes with general intuition, e.g., that 'the process of drawing analogies has important similarities to the processes underlying the computation of categories and concepts, and that rather than being distinct processes a common mechanism may underlie both categorization and analogy (Goswami, 2002, 448).

Another advantage of the analysis is the connection it makes between cases of analogy exploitation and cases where a conceptual mapping appears to be of value without there being any knowledge transfer. This is salient in the case of linguistic metaphors. As Fauconnier notes, in constructing a metaphor such as "digging one's own grave", 'we are not exploiting analogical transfer from the domain of "dying and graves" to the domain of "action and failure," because in fact the two domains in question are *not* structurally analogical in the relevant respects. Yet the power of the metaphor is as great as, or greater than, in cases of simple transfer' (Fauconnier, 2002, p. 263).

For Fauconnier, the difficulty of accounting for benefits in terms of knowledge-transfer motivates innovation of a specifically combinational theory, cf. the similarly motivated approaches of (Ward *et al.* 1999) and (Lakoff and Johnson, 1980). However, there is a reasonable prospect that metaphor (and related scenarios) can be treated in the same way as analogy-use (cf. Gentner *et al.* 2002; Holyoak, 2005; Gentner, 1982; Gentner *et al.* 1988). The analysis supports this theory directly. Insofar as benefits are equated with conceptual productivity, processes which promote exploitation of analogy fall into the same category as processes promoting use of metaphor. Innovation of an over-arching categorical construction can enhance productivity without there being any implied transfer of knowledge. This is illustrated in the vaccination analogy (Figure 14) where the benefits obtained are mediated primarily by schema acquisition.

Treating conceptual competences as mediated by a productivity preference also enhances generality in other ways. The innovation of concepts of higher productivity increases representational power for a lesser 'expenditure' of represented concepts. A preference for conceptual productivity can thus be viewed as a parsimony heuristic: the productivity preference is akin to the goal of minimizing storage. Rather than envisaging an analogy-related competence, say, to be mediated by special-purpose principles, we can then view it is oriented towards cognitive enconomy. Gentner and Kurtz comment, that 'It is as though we had an implicit aesthetic built into our comparison process that likes connected systems better than lists of separate matches.' (Gentner and Kurtz, 2005, p. 252). On the basis of the analysis, this aesthetic is really just a preference for minimized representational cost.

Another benefit of the analysis is its ability to ground and formalize judgements of conceptual *capacity*. There is the possibility of using the hierarchical structure as a way of discriminating different degrees of competence. For example, we might identify an absolute limit on the number of constructs which can be formed by a certain type of agent. There is also the potential for limiting capacities in a specifically functional way. We might place a limit on the maximum number of constituents which can be combined in a construct; or on the maximum number of hierarchical levels which can be assembled; or on the degree to which bimodality (combinatorial use of categorical and compositional constructs) can be exploited.

As noted in the introduction, the differentiation of conceptual competence in cognitive science has been progressed using a variety of strategies and methodologies.

Including use of intuition, computer models, philosophical tradition. Invoking ontolies, formal, semi-formal and formal. Invoking logic formalisms.

In contrast, the present paper has adopted the methodology of *reverse engineering*. Given the distinction between compositional and categorical construction, enumeration of combinatorial possibilities then exposes a rich taxonomy of functionality in which the recognized conceptual competences (e.g., classical concept learning, prototype formation, relation learning, schema acquisition, higher-order reasoning) can all be given a formal classification. The taxonomy and its classifications are shown in Figure 17. (Note that single-starred labels in the diagram refer

Illucidation of the hierarchical structure of conceptual competence reveals that competences which have often been regarded as contrary (i.e., embodying competing theories) are actually in well-defined subsumption relationships. The functionality of forming classical concept represntations is contained within the functionality of similarity-based learning. This is contained within the functionality of prototype formation. The competence of prototype formation is contained within the competence of analogical mapping/knowledge transfer. This is contained within the competence of schema acquisition. The situation with regard to relation learning is more a matter of degree. The competence of analogy exploitation can be seen to be on a continuum with relation discovery, with the degree of relation discovery being governed by the degree to which embodied relational constructs are bootstrapped.

The generic form of bimodal, progressive conceptualization can also be viewed as modeling basic, conceptual reasoning. Prototype formation can then be viewed as being contained within either analogy exploitation or within reasoning. At a higher level of the hierarchy, the competence of transformational construction is identified. This is the functionality in which features of a derived hierarchical structure are re-selected in arbitrary ways for use as primitive data within a further process of progressive construction.

The analysis also idenfities a hierarchical relationship applying to processing principles such as Gentner's absrtractness and systematicity (cf. Hofstadter's 'temperature'). Respect for the principles of abstractness and systematicity are

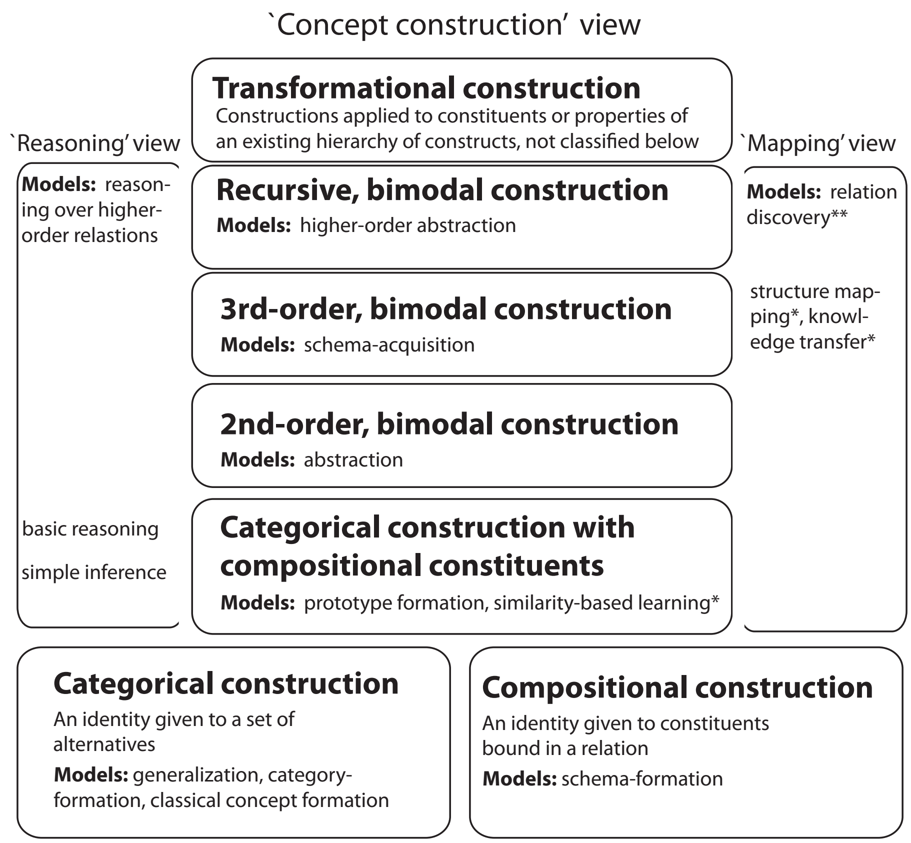

Figure 17: Analytic decomposition of conceptual competence.

shown to emerge naturally when a process of prgressive conceptualization is pursued so as to maximise economy, i.e., minimise storager and representation costs. ...

Regarding the fundamental nature of *concepts*, the analysis is also able to offer some predictive evidence. Derivation of the analysis does not entail adopting any particular ontology of concepts. But its results embody precise predictions about the structural properties that concepts must necessarily have. A key entailment is that there is no *context-free* way of distinguishing between concepts and conceptual structures. A concept is always a conceptual structure at some other level of analysis (and vice versa). (This is broadly in line with Smith and Medin (1981, p. 17), who note that there 'are no a priori means for distinguishing between a concept and a feature, for any feature can itself be treated as a concept.')

A second entailment is that there is no meaningful way to discriminate structure from function in conceptual entities. Functional properties are necessarily entailed in structural properties and vice versa. On this basis, it is wrong to think of a conceptual system as being comprised of some static, symbolic structures which are then processed by a seperate functional module. Rather the structural and functional properties always go hand in hand. Conceptualization is then a progressive process of function-structuring. Its effect is to generate a multi-leveled processing/representation structure which serves to consolidate different levels of conceptual description.

On this basis, the analysis is very closely related (and may even be equivalent) to Prinz's version of concept empiricism (Prinz, 2002), which views concepts structured perceptual detection mechanisms. In Prinz's model, there is an equivalent merging of functional and structural properties and a similar expectation of multi-level abstraction. Prinz does not portray combinational processes in details. As he notes (Prinz, 2002, 35) 'Rather than explain how concepts combine compositionally, I show that compositional combination is compatible with emergent features and context sensitivity.' However, his general view that 'concepts are mechanisms of detection, which allow us to track things, and which enable us to simulate them when they are not impinging upon our senses' would appear to be broadly compatible with the present analysis.

The assertion of function/structure hybridization also chimes with (Gardenfors, 2000), in which functional and structural properties of concepts are taken to be combined. Gardenfors particular emphasizes three levels of representation, being the ... level, the symbolic level and the conceptual level and is particularly committed to the geometric or spatial representations being used at the subconceptual level. There are differences in detail here but the general thrust of Gardenfors analysis is compatible with present proposals. Indeed, with regard to Figure 15, were one to introduce one division between 2nd and 3rd-order concepts and another at the bootstrapping level, the conceptual system depicted would then break down into three sub-system broadly analogous to Gardenfors' connectionist, symbolic and conceptual levels.

One final advantage of the present proposal deserves mention. Being derived purely through enumeration of combinatorial possibilities, the analysis entails no particular assumption to be made regarding the character or constituency of the entities which are 'given'. It makes no difference (to the conclusions drawn) whether we take them to be primitive concepts, sense data, sensory-motor associations or some other entity. On this basis, we have the promise of reasonable compatibility with enactive and embedded/embodied approaches. These tend to reject the 'input/output' image of cognition (Clark, 2008) in which sense data are accommodated (by some conceptual processing module) to static conceptual structures, prior to being re-processed into motor commands. On the enactive view, exploitation of advantageous ecolical feedback loops necessitates integrated processing of sensory-motor associations, on which basis any kind of conceptual phenomena would seem to require the possibility of grounding in sensory-motor associations. The analysis is fully open to this possibility. Indeed, in the portrayal of CAUSE (Figure 15), the incorporation of sensory-motor associations as primitive constituents would seem to be essential.

The analysis is thus not specifically empiricist or rationlist. Rather it is compatible with both approaches.

References

Bickerton, D. (1996). *Language and Human Behavior*. London: UCL Press.

Boden, M. (2004). *The Creative Mind: Myths and Mechanisms* (2nd edition). London: Weidenfeld and Nicolson.

Clark, A. (2008). *Supersizing the Mind: Embodiment, Action and Cognitive Extension*. Oxford: Oxford University Press.

Clark, A. and Thornton, C. (1997). Trading spaces: computation, representation and the limits of uninformed learning. *Behaviour and Brain Sciences*, *20* (pp. 57-90). Cambridge University Press.

Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, Mass.: Harvard University Press.

Doumas, L., Hummel, J. and Sandhofer, C. (2008). A theory of the discovery and predication of relational concepts. *Psychological Review*, *115* (pp. 1-43).

Falkenhainer, B., Forbus, K. and Gentner, D. (1989). The structure-mapping engine: algorithm and examples. *Artificial Intelligence*, *41* (pp. 1-63).

Fauconnier, G. (2002). Conceptual blending and analogy. In D. Gentner, K.J. Holyoak and B.N. Kokinov (Eds.), *The Analogical Mind: Perspectives from Cognitive Science* (pp. 255-285). London: The MIT Press.

Fodor, J. (1975). *The Language of Thought*. New York: Crowell.

Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Ganter, B., Stumme, G. and Wille, R. (Eds.) (2005). *Formal Concept Analysis*. New York: Springer-Verlag.

Ganter, B. and Wille, R. (1999). *Formal Concept Analysis: Mathematical Foundations*. New York: Springer-Verlag.

Gardenfors, P. (2000). *Conceptual Spaces: The Geometry of Thought*. Cambridge, Mass.: The MIT Press.

Gentner, D. (1982). Why nouns are learned before verbs: linguistic relativity versus natural partitioning. In S.A. Kuczaj (Ed.), *Language development: Vol. 2. Language, Thought and Culture* (pp. 301-334). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

Gentner, D., Bowdle, B., Wolff, P. and Boronat, C. (2002). Metaphor is like analogy. In D. Gentner, K.J. Holyoak and B.N. Kokinov (Eds.), *The Analogical Mind: Perspectives from Cognitive Science* (pp. 199-253). London: The MIT Press.

Gentner, D. and Colhous, J. (Forthcoming). Analogical processes in human thinking and learning. In A. von Mu.ller E. Po.ppel (series eds), A. von Mu.ller, E.P.. Ppel, B. Glatzeder, V. Goel, A. Von and .M.. Ller (Eds.), *On Thinking: Vol. 2. Towards a Theory of Thinking.* Berlin, Heidelberg: Springer-Verlag.

Gentner, D., Falkenhainer, B. and Skorstad, J. (1988). In D.H. Helman (Ed.), *Analogical Reasoning: Perspectives of Artificial Intelligence, Cognitive Science, and Philosophy.* Kluwer Academic Pub.

Gentner, D. and Kurtz, K. (2005). Relational categories. In W. K., Ahn, R.L. Goldstone, B.C. Love, A.B. Markman and P.W. Wolff (Eds.), *Categorization Inside and Outside the Lab* (pp. 151-175). Washington, DC: APA.

Gentner, D. and Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, *30* (pp. 609-642).

Gentner, D. and Markman, A. (1995). Similarity is like analogy: structural alignment in comparison. *Similarity in Language, Thought and Perception.* C. Cacciari (pp. 111-147). Brussels, Belgium: Brepols.

Gentner, D. and Sagi, E. (2006). Does "different" imply a difference? a comparison of two tasks. In R. Sun and N. Miyake (Eds.), *Proceedings of the Twenty-eighth Annual Meeting of the Cognitive Science Society* (pp. 261-266).

Goswami, U. (2002). Analogical reasoning in children. In D. Gentner, K.J. Holyoak and B.N. Kokinov (Eds.), *The Analogical Mind: Perspectives from Cognitive Science* (pp. 437-470). London: The MIT Press.

Hinton, G. (1986). Learning distributed representations of concepts. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society.* Amherst, Mass.: Lawrence Erlbaum Associates.

Hofstadter, D. (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought.* Basic Books.

Holyoak, K. (2005). Analogy. In K.J. Holyoak and R.G. Morrison (Eds.), *The Cambridge Handbook of Thinking and Reasoning* (pp. 117-142). Cambridge, UK: Cambridge University Press.

Holyoak, K. and Thagard, P. (1995). *Mental Leaps: Analogy in Creative Thought.* Cambridge, MA: MIT Press.

Hummel, J. and Holyoak, K. (2005). Relational reasoning in a neurally-plausible cognitive architecture: an overview of the LISA project. *Current Directions in Cognitive Science*, *14* (pp. 153-157).

Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on Cognitive Science*. Cambridge,Ma.: MIT Press/Bradford books.

Kilverstein, J. and Clark, A. (2008). Bootstrapping the mind: response to open peer commentary on (hurley 2008). *Behavioral and Brain Sciences*, *31* (pp. 1-58).

Koestler, A. (1964). *The Act of Creation*. London: Hutchinson.

Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. Chicago: The University of Chicago Press.

Medin, D., Goldstone, R. and Gentner, D. (1993). *Respects for Similarity Psychological Review*, *100* (pp. 254-278).

Mervis, C. and Rosch, E. (1981). Categorization of natural objects. *Annual Review of Psychology*, *32* (pp. 89-115).

Minsky, M. and Papert, S. (1988). *Perceptrons: An Introduction to Computational Geometry* (expanded edn). Cambridge, Mass.: MIT Press.

Mitchell, T. (1977). Version spaces: a candidate elimination approach to rule learning. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence* (pp. 305-310).

Palmer, S. (1989). Levels of description in information-processing theories of analogy. In S. Vosniadou and A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 332-345). New York: Cambridge University Press.

Penn, D., Holyoak, K. and Povinelli, D. (2008). Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, *31* (pp. 109-178).

Piaget, J. and Inhelder, B. (1958). *The Growth of Intelligence in the Child*. London: Routledge & Kegan Paul.

Prinz, J. (2002). *Furnishing the Mind: Concepts and their Perceptual Basis*. Cambridge, MA: MIT Press.

Quinlan, J. (1983). Inferno: a cautious approach to uncertain inference. *The Computer Journal*, *26*, No. 3.

Quinlan, J. (1993). *C4.5: Programs for Machine Learning*. San Mateo, California: Morgan Kaufmann.

Ramscar, M. and Pain, H. (1996). Can a real distinction be made between cognitive theories of analogy and categorization?. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society*.

Rosch, E. (1973). On the internal structure of perceptual and semantic categories. In T. Moore (Ed.), *Cognitive Development and the Acquisition of Language*. Academic Press.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, No. 3 (pp. 192-233).

Rosch, E. (1977). Classifications of real-world objects: origins and represenations in cognition. In P. Johnson-Laird and P. Wason (Eds.), *Thinking: Readings in Cognitive Science* (pp. 212-222). Cambridge: Cambridge University Press.

Rosch, E. and Lloyd, B. (Eds.) (1978). *Cognition and Categorization.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Smith, E. and Medin, D. (1981). *Categories and Concepts.* Cambridge, Mass.: Harvard University Press.

Sowa, J. (1984). *Conceptual Structures: Information in Mind and Machine.* Reading, Mass.: Addison-Wesley.

Thompson, R. and Oden, D. (1998). Why monkeys and pigeons, unlike certain apes, cannot reason analogically. In K.J. Holyoak, D. Gentner and B.N. Kokinov (Eds.), *Advances in Analogy Research: Integration of Theory and Data from the Cognitive, Computational, and Neural Sciences* (pp. 38-48). Sofia: Bulgarian New University Press.

Ward, T., Smith, S. and Finke, R. (1999). Creative cognition. In R.J. Sternberg (Ed.), *Handbook of Creativity* (pp. 189-212). Cambridge University Press.

Wielemaker, J. and Bundy, A. (1985). *Altering the Description Space for Focussing.* DAI Research Paper 262, University of Edinburgh.

Wittgenstein, L. (1958). *Philosophical Investigations.* Oxford: Basil Blackwell.