

Is there a Place for Semantic Similarity in the Analogical Mapping Process?

Michael Ramscar and Helen Pain

Department of Artificial Intelligence
University of Edinburgh
{michael,helen}@aisb.ed.ac.uk

Richard Cooper

Department of Psychology
Birkbeck College, University of London
R.Cooper@psyc.bbk.ac.uk

Abstract

Ramscar & Pain (1996) argued that the analogical process cannot be easily distinguished from the categorisation process at a cognitive level. In light of the absence of any distinction between analogy and categorisation, we have argued that analogy is supervenient upon an important part of the classification process, and that as such 'analogical' models are capable of illuminating some categorisation tasks, for instance, the way in which structural systematicity can determine not only analogical judgements, but also category decisions. Our scepticism regarding the cognitive distinction between these two processes has implications for both analogy and categorisation research: in this paper we consider two leading analogical theories, Gentner's Structure Mapping Theory and Holyoak's Multi-Constraint Theory, and argue that results from our use of analogical modeling techniques in categorisation tasks offer some important insights into exactly which elements should be included in a theory of analogical mapping.

Introduction

Ramscar & Pain (1996) argue for a position that rejects any simple distinction between analogy and categorisation at the cognitive process level (see Ramscar, Pain & Lee, 1997 for more evidence). The argument runs as follows: definitions of analogy (e.g. Holyoak & Thagard, 1995; Clement & Gentner, 1991) rely on making a distinction between 'straight' categorical judgements on the one hand, and analogical, 'extra-categorical' judgements on the other; with no convincing account of just what constitutes a 'straight' categorical judgment, this amounts to little more than hand waving, attempting to characterise one ill-defined process by contrasting it with another ill-defined process. In the absence of any compelling reason to believe in distinct cognitive processes of analogy and categorisation, there is much utility in viewing human analogical and categorical behaviour as manifestations of the same process: research into analogy has yielded a number of plausible models of the analogical process (Forbus, Gentner & Law 1995; Holyoak & Thagard, 1995), in sharp contrast to categorisation research, where process models, to the extent that they feature at all, tend to be more conjectural in nature (e.g. Medin & Ortony, 1989). A companion paper to this (Ramscar, Pain & Lee, 1997) presents further evidence of the benefits and insights that this approach can bring to investigations into the nature of categorisation decisions.

In this discussion, we wish to examine the perspective our particular view of the categorisation / analogy divide, and the results of our experiments using analogical modeling to explore categorisation, can bring to theoretical approaches to analogy (and by implication, existing models of analogy). The removal of neatly bounded 'concept domains' with which various aspects of the analogical process can interact produces a changed circumstance that seems likely to have repercussions for theories based on just such assumptions: in particular, we wish to examine the implications of our approach to the kinds of similarity constraints — structural, pragmatic and semantic — which are variously included in, or excluded from, theories and models of the analogical process.

A number of 'competing' theories of analogy exist: Holyoak & Thagard (1995); Gentner (1983; Forbus, Gentner & Law, 1995); Keane (Keane, Ledgeway & Duff, 1994); Hofstadter (1995). In this paper we examine in detail the effects of our blurred distinction perspective on the vexed question of 'semantics' in the first two of these, Gentner's 'Structure Mapping' theory and Holyoak & Thagard's 'Multi-Constraint' theory. These constitute the most explicit theories, whose supporting evidence and accompanying process models have been widely disseminated and accepted.

Mapping Without The Distinction

Gentner's Structure Mapping Theory

Gentner (Gentner, 1983; Clement & Gentner, 1991) proposed the Structure Mapping Theory as an attempt to explain how it is that two domains can be considered analogous, and in particular how it is that correspondences between analogues from two domains can be mapped.

Structure mapping proposes that the mapping and inference between two domains can be achieved by assigning correspondences between objects and attributes and then mapping predicates with identical names. In order to do this, Gentner assumes a predicate like representation (figure 1), distinguishing between *objects*, *object-attributes* and *relations*. Object-attributes are those predicates that have one argument and describe object properties, e.g. YELLOW(sun). Relations are divided into a hierarchy of orders, with those predicates with two or more arguments which are used to describe relations between objects, for example ATTRACTS(sun, planet) forming the lowest order, and those predicates describing different levels of relationships between relations e.g. CAUSE(ATTRACTS(sun,

planet), REVOLVES_AROUND(planet, sun)) forming the higher orders.

Solar System

YELLOW(sun)
MASSIVE (sun)
HOT(sun)
ATTRACTS(sun, planet)
ATTRACTS(planet, sun)
MORE_MASSIVE_THAN(sun, planet)
REVOLVES_AROUND(planet, sun)
HOTTER_THAN(sun, planet)

Hydrogen Atom

ATTRACTS(nucleus, electron)
ATTRACTS(electron, nucleus)
MORE_MASSIVE_THAN(nucleus, electron)
REVOLVES_AROUND(electron, nucleus)

Figure 1: Predicate representations of the solar system and a hydrogen atom

The theory itself comprises two parts: *mapping rules*, and the *systematicity principle*. Mapping rules state that attributes of objects are not mapped; but relations between objects are preserved. The systematicity principle requires that complex higher order relations (e.g. CAUSE above) are mapped preferentially, followed by relations that constitute the higher order arguments. This is intended to capture the notion that analogy conveys a system of connected knowledge, rather than an assortment of independent facts:

“structure mapping stems in part from the observation that useful analogies, such as those used in science or education, involve rich, interconstraining systems of mappings between two domains, rather than a set of independent correspondences”
(Clement & Gentner, 1991, pp. 91–92)

Gentner’s structure mapping theory has been implemented as a computer simulation model SME (Falkenhainer, Forbus & Gentner, 1989).

One criticism that has been made of structure mapping theory is that the representations it uses are arbitrary. The primitives used in the examples are chosen selectively. The example would fail if REVOLVES_AROUND in the Solar System domain was replaced with ORBITS. Furthermore, the inclusion of a larger range of domain relations would result in the matches appearing far less sound than they do. Moreover, Holyoak (1985) claims that Gentner can give no account as to why HOTTER_THAN in figure 1 should not be mapped onto the Hydrogen Atom domain. Holyoak argues that expanding the representation in order to utilise the systematicity principle will be of little help, since on another expansion the relation HOTTER_THAN will form part of a higher order system (involving the warming of planets, conditions for the existence of life, etc.). The structure mapping theory seems to give no explanation as to why this relation is not mapped.

To some researchers this has appeared to be a major prob-

lem. A system which uses no goal or domain information, and which relies purely on structural inferences, might appear to have trouble when, given a number of networks of higher-order relations between domains (i.e. several causal networks), giving an account as to why it selected the one that was relevant to making the desired analogy.

Holyoak and Thagard’s Multi-Constraint Theory

Holyoak’s theory (Holyoak, 1985; Holyoak & Thagard, 1989; Holyoak & Thagard, 1995) is concerned primarily with problem solving. It attempts to capture the intuition held by some researchers that goals play some part in the analogical process, particularly in mapping. (Gentner’s theory precludes any such influence.) The scope and ambition of this theory has changed considerably over time: what is presented here is a sketch of its general nature.

Holyoak *et al* advocate an explanation of analogy in terms of a goal-driven processing system — mappings are controlled by the system’s goals. The analogical mapping problem is seen as one of explaining how the large number of possible mappings between domains can be evaluated and a subset of these used for the transfer of information between domains.

They suggest that this subset emerges from an attempt to balance the different influences upon the mapping process. More specifically they regard the process as an attempt to simultaneously satisfy several constraints. The first group of these (logical compatibility, role identity, uniqueness and relational consistency) are structural, and therefore compatible with the overall thrust of Gentner’s theory. However, Holyoak *et al* also consider constraints of *pragmatic centrality* and *semantic similarity* to be integral to the mapping process. These non-structural constraints rely on information other than that found in the basic domain representations.

Structural Constraints Logical compatibility ensures that mappings are only considered if they are between entities of the same ‘type’. Thus, in the Solar System / Hydrogen Atom analogy the predicate MORE_MASSIVE_THAN cannot be matched with the object electron. Similarly, a mapping between the predicates HOT and REVOLVES_AROUND will not be considered as they take different numbers of arguments. This primarily syntactic constraint is intended to ensure that mappings between different levels of description are not attempted. For example single-argument predicates, such as HOT, tend to be purely descriptive, specifying a particular attribute of an object. Multi-argument predicates describe relationships between objects, and so can be considered to represent a higher level of description. The argument is that mappings between different levels of description are not productive and this constraint serves to eliminate any potential mappings of this kind.

A further hurdle potential mappings must overcome in this model is the role identity constraint. This assumes that the base and target domains can be divided at a higher level of description than that at which the mapping takes place. In the use of analogy in problem solving, upon which the authors focus predominantly, this means the domains may be redescribed in terms of a start state, the problem goals, and

the operators that can be used to try and achieve these goals. Role identity then limits mapping to relations and objects that appear in the same part of the domain definition. This provides a weak pragmatic influence in that elements can only be considered for mapping if they play a similar role in both domains.

Holyoak *et al* assume that each element in the base domain will ultimately map onto only one element in the target domain, and *vice versa*. Thus there is a competition between members of the set of potential mappings between one base element and a number of possible target elements. For example, if HOTTER_THAN in the base maps onto HOTTER_THAN in the target then it cannot map onto LESS_MASSIVE_THAN in the target. Accordingly any factor which serves to increase the level of support for the former mapping will consequently act to decrease support for the latter.

A final structural constraint is relational consistency, which acts to ensure that mappings between the base and target domains are consistent. If mappings between structural elements receive support, mappings between the structures themselves, and any other elements, are also supported.

Pragmatic Centrality The importance of an element (object or relation), whether in the base or target domain, is another consideration in the mapping process. An element's importance is defined in terms of how useful the element is in satisfying the current goal (or subgoal) of the 'analogiser'. Thus any mappings involving 'useful' elements receive more support than mappings involving less useful elements. When our example analogy is used to explain the relative motion of sub-atomic particles, mappings involving YELLOW and HOTTER_THAN are going to be less favourably considered than those involving REVOLVES_AROUND, since the former are not utilised in satisfying any explanatory goals.

Semantic Similarity Holyoak *et al* suggest that the most useful mappings are likely to come from elements which are semantically similar. In the Solar System / Hydrogen Atom analogy, predicates with identical names can be regarded as more similar than those with different names. In more complex examples the method of determining relative similarity is more difficult. Holyoak *et al* make no claim as to any particular model of semantics. The semantic similarity constraint is regarded more as a heuristic than a firm rule, and can be applied in differing strengths at various stages of the mapping process.

Holyoak & Thagard (1989) emphasise that the logical compatibility and role identity constraints are restrictions on the building of the mapping network, and these restrictions are regarded as less important than the three principal constraints of isomorphism (uniqueness and relational consistency), semantic similarity and pragmatic centrality.

The theory has been implemented as a computer simulation ACME, which is a constraint satisfaction network incorporation these considerations.

Comparing The Two Theories

Structure Holyoak & Thagard (1995) argue that the similarities and differences between their theory and structure mapping theory can best be illustrated by comparing Gentner's theory with the three main constraints posited by multi-constraint theory. They claim that multi-constraint theory captures Gentner's insight regarding the importance of systematic structure (in ACME, interconnected systems will have more mutually supporting links than an isolated relation), but in a more flexible manner. As implemented in SME, Gentner's theory rigidly enforces one-to-one mappings and structural consistency — potential mappings which violate these constraints are not made. In contrast, ACME, whilst preferring one-to-one mappings (by using inhibitory links to discourage many-to-one mappings) nevertheless will allow violations.

Pragmatic Constraints Gentner's theory (and SME) does not incorporate or recognise the influence of pragmatic — goal driven — constraints on the mapping process. According to structure mapping theory, the operation of goals is external to the actual mapping process, constraining the evaluation of mapping outcomes, rather than actual mappings (although I-SME (Forbus, Ferguson & Gentner, 1994) does incorporate pragmatic influences in mapping). However, whilst the evidence that goals can influence *what is mapped* in analogy is clear (Spellman & Holyoak, 1992), it is less clear that goals directly influence or constrain the *mapping process*.

We noted two criticism of structure mapping earlier: the neatness of the representations used in SME (this applies equally to ACME), and that if an analog offers up two competing possible modes of transfer with a similar level of systematicity, then the systematicity principle cannot act as a constraint in the selection of one or the other. None of this need militate against structure mapping in principle: an analogue which allows two equal mappings may be a poor choice of an analogue; the systematicity principle's yielding of two equally valid mappings may be a psychologically valid resolution of the initially poor choice of analogy. To return to the earlier objection to Gentner's structure mapping theory (Holyoak, 1985), it might well turn out that an expansion of the representations of the Solar System and the Hydrogen Atom will simply lead to a situation in which the two examples are no longer seen to be analogous. Given that the success of any analogy is contingent upon the way in which the putative analogues are represented (Gentner, 1989), it is not a failing of a theory that it cannot provide accurate mappings in situations where candidate analogues are presented in such a way as to obscure any analogous similarity between them.

A major criticism of all analogical theories is that they do not consider the considerable psychological evidence that the choice of representation is crucial to analogy (Hofstadter, 1995). If categorisation research tends to ignore processes and overconcentrate on representation, the opposite is true of analogy. Yet, in fact some of the best evidence of the influence of goals (Spellman & Holyoak, 1992) seems to support the view that goals influence representation rather than mapping.

These “results suggest one way in which analogies can be used systematically to influence people’s inferences — [the representation of] the source [analogue] can itself be massaged to encourage a desired mapping” (Holyoak & Thagard, 1995).

Semantic Constraints With regard to the semantic similarity constraint, the main focus of this discussion, the respective positions are as follows: SME only matches predicates with identical names, thus if planets (see figure 1) were represented as SMALLER_THAN the sun, and electrons as LESS_MASS_THAN a nucleus, then structure mapping would not allow, and SME would not make, a mapping between the two relations. On the other hand, whilst ACME again prefers to map identical relations, weights on the network can be adjusted to capture the semantic similarity between SMALLER_THAN and LESS_MASS_THAN. Holyoak & Thagard argue that this shows a significant weakness in Gentner’s theory:

“with its emphasis on structure to the exclusion of all other constraints, SME does not simply discourage mappings between non-identical but semantically similar items; it does not even permit them.” (Holyoak & Thagard, 1995, p. 258)

Holyoak and Thagard’s criticism of the lack of semantic considerations in Gentner’s theory carries a lot of intuitive weight. It does seem a perverse, restrictive analogical theory that rejects mappings between SMALLER_THAN and LESS_MASS_THAN in the course of an analogical mapping. However, it is not necessarily so, and from the perspective of a blurred distinction between analogy and categorisation, it might actually be that rejecting such a mapping is necessary *per se*, rather than necessarily perverse.

One reason for so-arguing stems from the results of our investigations into the effects of systematic structure upon categorisation judgements. Ramscar & Pain (1996) addressed the question of whether analogy can be distinguished from categorisation by contrasting categorisational and analogical processes by presenting subjects with Gentner *et al*’s analogy materials and asking them to categorise them. Given that Gentner *et al* define the analogical mechanism in terms of structure mapping, and given our hypothesis that this process was not distinct from a basic categorisation process, we expected structure mapping to determine categorisation. Gentner *et al* seem to implicitly assume that match items with only structural similarities (i.e. analogues) belong to different categories. We predicted that they would be categorised together. We found that 79.5% of the groupings formed by our subjects had only shared systematic structure (traditionally defined as analogy) as a common feature amongst members of the categories formed. In contrast, only 5% of groupings produced had common object descriptions as the common similarity across categories (i.e. the attribute matches often thought to be determinate of categorisation). To the 79.5% of structural congruity groupings could be added a further 8% of classifications where structural *additions* to otherwise structurally congruent representations caused them to be classed singularly. Thus we concluded that that mechanisms normally considered to be analogical could also in fact support categorisation

tasks, and in this instance, no discernible difference could be found between analogical and categorical behaviour.

This suggests that the process of classifying two terms together (mapping SMALLER_THAN to LESS_MASS_THAN) is no different to the process of determining the analogy between SOLAR_SYSTEM and HYDROGEN_ATOM. Thus the process of mapping SMALLER_THAN to LESS_MASS_THAN seems to be less a sub-process of the process of determining the analogy between SOLAR_SYSTEM and HYDROGEN_ATOM and more like the same process functioning in parallel.

Prototype Schemas? One possible objection to the above would be that our characterisation of categorisation is incomplete. It might be argued that analogical judgements may not be easily distinguished from classification judgements, but *categorisation* judgements can be. This claim relies on the idea that categorisation doesn’t involve that same mapping process because it makes use of generalized ‘prototypical’ schemas (Holyoak & Thagard advocate such a view of categories in Holland *et al*, 1984). If SMALLER_THAN and LESS_MASS_THAN share the same prototypical schema, then there will be no need to compute the similarity between SMALLER_THAN and LESS_MASS_THAN, since such similarity can be confirmed merely by reference to the prototype¹; and it is this confirmatory reference-to-prototype that is modeled by the semantic links in Holyoak and Thagard’s theory.

There are a number of objections to such an account. Practically, there is the problem of providing a convincing account of what a prototype category is (Medin & Ortony, 1989; Ramscar & Pain, 1996). Propositional schematic models of representation (some variety of which is assumed in the majority of cognitive theories) have the power to store a combination of both schemas and the exemplars from which such schemas are constructed. Moreover, what is actually stored is the subject of much debate: are just schemas stored? Fodor & Lepore (1996) give many good reasons for doubting such a ‘pure’ prototype schema theory. If both are stored, then how do exemplars contribute to schemas? Some have advocated a rejection of schemas altogether (e.g. Nosofsky, 1988), arguing that only exemplars are stored, and that new objects are classified by comparison with stored exemplars, and the calculation of some kind of fit. A further experiment (Ramscar, Pain & Lee, 1997) was designed to see whether analogical theory could shed any light on the nature of subjects’ stored representations.

In their analogical recall experiments, Gentner *et al* (Gentner, Ratterman & Forbus, 1993) showed that analogical access relied primarily upon surface attribute (object) matches:

¹Curiously, a footnote to Gentner, Ratterman & Forbus (1993) states “SME’s constraint of matching identical predicates assumes canonical *conceptual* representations, not lexical strings. Two concepts that are similar but not identical (such as “bestow” and “bequeath”) are assumed to be decomposed into a canonical representation language so that their similarity is expressed as a partial identity (... “give”), which suggests that this idea has widespread appeal; a version of it has been advocated by the authors in the past (Ramscar, Lee & Pain, 1996).

stories with shared attributes were recalled from memory far more readily than objects with shared structure, even though subjects adjudged some of these stories to be less similar, and inferences generated from them to be less sound with regards to a base than the stories which possessed only shared structure. We decided that this phenomenon might be useful in exploring representation. In the course of the classification experiment, subjects were asked to give their classes “a simple descriptive name”, which they could then associate with their classification. By examining what attributes they could recall that were associated with that name, we used Gentner *et al*'s findings about structure versus attribute to shed some light on the mental representation associated with the name. If subjects stored some kind of abstracted prototype, then our hypothesis was that attributes associated with the most prototypical stories would be most readily retrieved from memory, with other attributes recalled insofar as they were shared with the prototypical story. On the other hand, if subjects stored only exemplars, then the lack of context provided by a simple class name should make it equally likely that attributes associated with any given (or all) exemplar might be recalled.




Figure 2: Two views of the mapping process

The experiment yielded little evidence to support the hypothesis that subjects had abstracted and stored schemas from the groups they had classified, despite the fact that it appeared to be a shared structural schema that was the basis of subjects' original classification decisions (Ramscar & Pain, 1996). Instead, when Gentner, Ratterman & Forbus (1993)'s analysis was applied to subjects' behaviour, it appeared that being presented with a class name in no particular context caused the subjects to randomly recall one of the exemplars associated

with that name, and then use that exemplar as the stimulus for recalling other class members; overall, only 10% of recall tasks favour the former analysis, whereas some 71% support the latter. Thus as well as finding no evidence to support stored prototypes, Ramscar, Pain & Lee (1997) provide evidence that only exemplars (instances) of categories are stored (which strengthens the argument for not strongly distinguishing analogy and categorisation). It would seem that schemas can't supply the necessary theoretical justification for semantic links.

Discussion

If we consider the schematic representation of the mapping process in figure 2, Holyoak & Thagard model the mapping of F onto G in analogy $X \rightarrow Y$ by means of a semantic link, whereas the evidence points towards a similar parallel process for $F \rightarrow G$ and $X \rightarrow Y$. We feel that this is an important distinction: whereas Holyoak & Thagard model the semantic similarity between SMALLER_THAN and LESS_MASS_THAN as a *sub-process* of analogy, it seems more likely that this similarity is computed by the *same process*, in parallel². This can be seen most clearly via a fine-grained analysis of the predicates as in figure 3.

```

SMALLER_THAN
RELATION(smaller_than)
NO_OF_ARGUMENTS(smaller_than, 2)
COMPARITIVE(smaller_than)
NONREFLEXIVE(smaller_than)
ASYMETRIC(smaller_than)
TRANSITIVE(smaller_than)
CONCERNS(smaller_than, size)

LESS_MASS_THAN
RELATION(less_mass_than)
NO_OF_ARGUMENTS(less_mass_than, 2)
COMPARITIVE(less_mass_than)
NONREFLEXIVE(less_mass_than)
ASYMETRIC(less_mass_than)
TRANSITIVE(less_mass_than)
CONCERNS(less_mass_than, weight)

```

Figure 3: Predicate representations of SMALLER_THAN and LESS_MASS_THAN

Since from our perspective, this mapping process is what we mean when we talk about the 'analogical' process, it follows that there is no room for a semantic similarity constraint in a cognitive theory of analogy (or at least, it follows that before one can posit such a constraint, one will have to successfully distinguish 'analogy' from 'categorisation'). None of which means that we rule out the use of semantic links in mapping networks — rather, we suggest that they should be

²We are not saying that *all* similarities are calculated this way; what is true for relational predicates may not be true for others. It seems unlikely that structure mapping will feature in red vs. crimson.

seen as implementational details within a model, rather than embodiments of psychological theory.

One possible benefit to be gained from removing the semantic and pragmatic constraints³ from Holyoak and Thagard's model is that it can allow a straightforward comparison between the structural elements in both theories to be made. Isomorphic mapping in multi-constraint theory differs from Gentner's structure mapping in a number of details; our next step is to evaluate the significance of these details by specifying both theories in a common, executable, specification language. Similarities and differences between the theories are currently blurred by the ambiguities of the language (English) in which they are stated. Whilst computational implementations of both theories exists, those implementations do not fully clarify either theory because they make no distinction between theoretical claims and implementational details. This is part of a broader issue: cognitive science is in need of tools and techniques in which to precisely state theoretical proposals so that their assumptions and implications are clear and comparison is possible (c.f. Cooper, Fox, Farrington & Shallice, 1996). We are currently using the COGENT modelling environment (Cooper & Fox, 1997) to develop executable specifications of both structure-mapping theory and multi-constraint theory. It is our strong belief that this work will demonstrate that, modulo semantic similarity, the theories differ mainly in implementation details.

The COGENT models will also provide a framework in which to explore more thoroughly issues of representation. We noted above the lack of focus on representation in analogical models (a point made frequently by Hofstadter, 1995). We have demonstrated that analogical process models can yield interesting results when applied to categorisation tasks (where perhaps too much focus is given to representation, c.f. Fodor & Lepore, 1996). Is it too much to hope that a similar cross-fertilisation might ultimately flesh out analogy's decidedly anorexic picture of representation?

Acknowledgements

Many thanks to John Lee, Dedre Gentner and two anonymous reviewers for comments; thanks also to Richard White for interesting us in analogy in the first place.

References

- Clement, C.A. & Gentner, D. (1991) Systematicity as a selection constraint in analogical mapping, *Cognitive Science*, 15, 89–132.
- Cooper, R. & Fox, J. (1997) COGENT: An interactive graphical environment for cognitive modelling. Unpublished: Manuscript available.
- Cooper, R., Fox, J., Farrington, J. & Shallice, T. (1996) A systematic methodology for cognitive modelling. *Artificial Intelligence*, 85, 3–44.
- Falkenhainer, B., Forbus, K.D. & Gentner, D. (1989) The structure mapping engine: an algorithm and examples. *Artificial Intelligence*, 41, 1–63.
- Fodor, J. & Lepore, E. (1996) The red herring and the pet fish: why concepts still can't be prototypes. *Cognition*, 58, 253–270.
- Forbus K.D., Ferguson, R.W. & Gentner, D. (1994) Incremental structure mapping. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*, pp. 313–318.
- Forbus K.D., Gentner, D. & Law, K. (1995) MAC/FAC: A model of similarity based retrieval *Cognitive Science*, 19, 141–205.
- Gentner, D. (1983) Structure-mapping: a theoretical framework for analogy. *Cognitive Science*, 7, 155–170.
- Gentner, D. Ratterman, M.J. & Forbus, K. (1993) The roles of similarity in transfer. *Cognitive Psychology*, 25, 524–575.
- Gentner, D. (1989) The mechanisms of analogical reasoning. In S. Vosniadou and A. Ortony (eds.) *Similarity and Analogical Reasoning*. Cambridge University Press, Cambridge.
- Hofstadter, D. (1995) *Fluid Concepts and Creative Analogies*. Basic Books, Harper Collins Publishers.
- Holyoak, K.J. (1985) The pragmatics of analogical transfer. *Psychology of Learning and Motivation*, 19, 59–87.
- Holyoak, K.J. & Thagard, P. (1989) Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295–355.
- Holyoak, K.J. & Thagard, P. (1995) *Mental Leaps*. MIT Press, Cambridge, MA.
- Keane, M., Ledgeway T. & Duff, S. (1994) Constraints on analogical mapping: a comparison of three models. *Cognitive Science*, 18, 387–438.
- Medin, D. & Ortony, A. (1989) What is psychological essentialism? In S. Vosniadou & A. Ortony (ed.) *Similarity and Analogical Reasoning*. Cambridge University Press, Cambridge.
- Nofosky, R.M. (1988) Similarity, frequency and category representations. *Journal of Experimental Psychology: Learning Memory and Cognition*, 14, 54–65.
- Ramscar, M.J.A. & Pain, H. (1996) Can a real distinction be drawn between cognitive theories of analogy and categorisation? In *Proceedings of the 18th Annual Conference of the Cognitive Science Society*, pp. 346–351. San Diego, CA.
- Ramscar, M.J.A., Lee, J.R. & Pain, H.G. (1996) A cognitively based approach to computer integration for design systems. *Design Studies*, 17, 465–483.
- Ramscar, M.J.A., Pain, H.G. & Lee, J.R. (1997) Some observed effects of systematic structure on categorisation decisions. Technical report, Department of Artificial Intelligence, University of Edinburgh.
- Spellman, B. & Holyoak K. (1992) If Saddam is Hitler, then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, 62, 913–933.

³The question of whether pragmatic constraints affect mapping or representation remains open — our comments regarding the relative similarity between SME and ACME when it comes to structural mappings might apply equally to pragmatic factors in ACME and I-SME.