Title
Evidence for the Processing of Re-representations during the mapping of Externally Represented analogies

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 22(22)

Author
Ramscar, Michael

Publication Date
2000

Peer reviewed

Evidence for the processing of re-representations during the mapping of externally represented analogies

Michael Ramscar
michael@cogsci.ed.ac.uk
Division of Informatics, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, Scotland

Abstract

High level descriptions of the analogical reasoning process in cognitive science have now converged to present a relatively unified account (Hummel and Holyoak, 1997). However, the broad, consensual account of analogy is still far from complete: whilst it is possible to give a good explanation of the mapping of larger, structured representations in analogy, accounts of the mappings of individual sub-elements in these representations are still under-specified. Here, we review some possible approaches to this problem, describe an experiment that provides some empirical support for the 're-representation' approach to sub-mapping, and then identify some shortcomings in the 're-representation' approach as it is currently conceived.

Introduction

Cognitive science has made great strides towards answering the important question of 'How do humans reason by analogy?' If we take a familiar example, the analogy between the solar system and Rutherford's model of the atom, then it is possible to explain – in broad terms – exactly how it is that two seemingly disparate objects can both remind us of one another in the first place, and then how it is that we can make meaningful correspondences between them.

Studies have shown that reminding (or retrieval) is driven by a computationally inexpensive process that initially matches surface (or semantic) elements in representations (witness the frequency – and mundanity – of most similarity based remindings, such as a lamp-shade reminding a party joker of a hat; see Gentner, Ratterman and Forbus, 1993).

Analogical mappings, on the other hand, are determined by a relatively more computationally expensive process. Global, systematic structural similarities between items to be matched need to be detected in order to make the kind of 'deeper', inference supporting correspondences that characterise analogy (Gentner, 1983; Goswami, 1992; Holyoak and Thagard, 1995; Hummel and Holyak, 1997).

Whilst theories and models of analogy are very compelling at one level of abstraction, there are certain assumptions made by *all* analogical theories that beg interesting questions if one seeks a more detailed explanation. As one increases the resolution of the question 'How do humans reason by analogy?' it appears that there are important gaps in current theories and process models. Here, we wish to consider just one aspect of one of these gaps: the problem we focus on is that of matching the 'semantics' of elements during the analogical mapping process. This problem can be summarised as follows: suppose that in your representation of the atom, you describe the motion of an electron in relation to the nucleus in terms of it "revolving around" the nucleus (perhaps this is how you ordinarily think about this motion). On the other hand, suppose that in your representation of the solar system you conceive the motion of the planets in terms of their "orbiting" the sun.

At one level of abstraction, it may be sufficient to say that similarities in the meanings – or usage – of these words determine these mappings. However, in a more detailed account – and model – of analogy we might wish to do more than appeal to humanistic intuitions about similarities of meaning. We might wish to account for the way in which these sub-elements of our representations of the atom and the solar-system are mapped onto one another with the same level of detail with which we account for the mappings between the representations themselves.

If we are to fully explain high-level mapping in analogy, we must also account for the way lexically distinct but 'semantically' similar items in representations are reconciled with each other in a way that allows high-level mappings to be made. Here, we review some possible approaches to this problem, and present some evidence that offers some support to a popular proposal in the literature: the *re-representation* hypothesis.

Semantic reconciliation and the re-representation hypothesis

Perhaps the most straightforward way to explain the mapping between "revolving around" and "orbiting"[1] would be in conceptual terms. If "revolving around" and "orbiting" could be shown to decompose into some canonical conceptual representation (say "circumnavigating"), then the link between them could be explained by reference to that concept, and the process by which it is made. This proposal is put forward by Gentner, Ratterman and Forbus (1993):

> "[the...] 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

[1] We shall refer to this as the problem of semantic reconciliation in analogy.

representation language so that their similarity is expressed as a partial identity (... "give")" Gentner, Ratterman and Forbus (1993, p 553)

The main drawback to this proposal is the lack of any specification of what a canonical conceptual representation (or a canonical representation language) is. Research into the mental representation of concepts suggests that human conceptual representations are anything but canonical; the proposals for generalised theories of representation that exist in the concepts literature fall well short of providing the kind of 'neat' account of concepts that canonical conceptual representation assumes (see Komatsu, 1992; Ramscar and Hahn, 1998 for reviews).

This problem has not gone unrecognised. In conjunction with other factors, such as evidence of the important role that structural commonalities (the 'what' of analogy) play in 'ordinary' conceptual tasks (e.g. Ahn, 1998), and the sheer difficulty of distinguishing analogy from 'ordinary' conceptual tasks (Ramscar and Pain, 1996), a widespread view has emerged that suggests that analogy *itself* may play an important role in semantic reconciliation (Forbus, Gentner, Markman and Ferguson, 1997, Hummel and Holayoak, 1997).

The basic idea behind this is outlined by Forbus, Gentner, Markman and Ferguson (1997) who propose that semantic terms might be decomposed into sub-predicate re-representations, with mapping between these being determined using the same process as similarity based transfer:

"re-representation allows relational identicality to arise out of... analogical alignment, rather than as a strict constraint on the input descriptions" Forbus, Gentner,Markman and Ferguson (1997, p 246).

A similar re-representation proposal is advanced by Hummel and Holyoak (1997):

"With the notion of chunked predicates and objects, LISA hints at a kind of recursive representation for meaning that may ultimately ground itself in basic perceptual primitives. In its current implementation, LISA can represent and map hierarchical propositions of arbitrary. Analogously, it is possible to imagine structures for roles and objects that are, themselves, deeply embedded recursive structures. The depth to which a role or object would need to be decomposed for the puroses of mapping would depend on the task at hand. For example, mapping 'John lifted the hammer' onto 'Bill raised the book' may require little or no decomposition of the predicates 'lift' and 'raise', which will have substantial overlap in their *semantic features*.[2] On the other hand, mapping 'John lifted the hammer' onto 'Bill pushed the cart' where the predicates have less feature overlap, may be more likely to depend on decomposition of 'lift' into 'caused to rise' and 'push' into 'cause to move laterally', thereby making explicit the parallelism of their internal structures. Recursively 'rise' and 'move laterally' might be decomposed into structures relating simpler predicates with basic perceptual primitives representing motion and locations in space residing at the very bottom." Hummel and Holyoak (1997, p.457).

Whilst re-representation is a popular idea in the analogy literature, its current status is largely hypothetical: re-representation proposals are usually couched in terms that relate to computational models, and as yet no evidence has been offered to support the psychological validity of the proposal.

The following experiment was designed to formulate a concrete re-representation proposal, and explore it empirically. The problem of semantic reconciliation revolves around supplying an account of what happens when two 'semantically similar' terms – "revolving around" and "orbiting" – are encountered during the mapping process. In ordinary usage, the representations of human category information involved in these processes are implicit; people know what – "revolving around "and "orbiting" mean, and they reconcile (or map between) the two terms accordingly. But the exact nature of *what* they know, and *how* such knowledge is represented appears to be inaccessible at the level of detail required to specify and model the underlying cognitive processes involved in the semantic reconciliation of the two terms.

Participants were asked to make inferences with the aid of two candidate bases (see figure 1). In both the target and each of the two candidate bases, the term that was crucial to determining the representation of higher order structure in the scenarios was a novel, artificial term. By supplying 'definitions' for that term, we hoped to be able to control the representations participants used for semantically reconciling particular terms during their analogising. By doing this, we hoped to test the prediction that participants would use the same *process* to match semantic items in their representations as they would in ultimately determining their analogies - i.e. that in these externally represented analogies, at least, participants would use and process re-representations to facilitate semantic reconciliation.

SCENARIOS

TARGET - The Guralaga

can be found in Australia
lives in Rainforests
only eats gau-gau berries
has a cronomus lucundus
the *cronomus lucundus* enables the Guralaga to eat gau gau berries.

BASE 1 - The Mongret

can be found in Australia
lives in Rainforests
only eats gau-gau berries
has a probus razoris
the probus razoris enables the Mongret to eat the gau gau berries.
Thanks to the way they eat, Mongrets live to a ripe old age and rarely suffer from cancer

BASE 2 - The Crany Dog

can be found in Papua new Guinea
lives in the grassy backlands
eats vegetation
has a remulum grandoso
because of the *remulum grandoso* the Crany Dog can eat vegetation.
Crany Dogs are particularly prone to cancer, which originates in their digestive system.

Figure 1: A base and two targets. The surface similarities between the target and base 1 are highlighted. The target and base 2 share few surface similarities

[2] The emphasis is ours

	Surface Match Base (SMB)	**Structurally Similar Base (SSB)**
Analogy Level	Shares **surface** features with target ***Structural*** overlap with target determined by dictionary mapping	Doesn't share surface features with target ***Strucutral*** overlap with target determined by dictionary mapping
Dictionary Level	Shares **surface** features with target Doesn't share structure with target	Doesn't share surface features with target Shares **structure** with target
Inference	In ***type A*** sets, the ***B*** inference is only supported by **surface** matches between target and SMB the dictionary entries In ***type B*** sets, the ***A*** inference is only supported by **surface** matches between target and SMB the dictionary entries	In ***type A*** sets, the ***A*** inference is only supported by **structural** matches between target and SSB the dictionary entries In ***type B*** sets, the ***B*** inference is only supported by **structural** matches between target and SSB the dictionary entries

Figure 2: The relationships between the base, targets, dictionary entries and inferences in the main stimulus groups.

Experiment

Participants

The participants were 170 volunteers, a mixture of postgraduate and undergraduate students from the Department of Artificial Intelligence, Centre for Cognitive Science, Department of Psychology and the Faculty of Music at the University of Edinburgh.

Materials, Design and Hypotheses

The materials comprised 5 groups of specially constructed scenarios (figure 1) with corresponding sets of novel dictionary entries (figure 3) and candidate inferences for each group (figure 4).

To control for biases towards particular inferences, each scenario group was further sub-divided into two versions of the scenario sets, and two versions of the dictionary entry sets, so that each scenario / dictionary sub-set supported one of the two different candidate inferences.

To classify the different structural / featural relation amongst the scenarios, we used Gentner, Ratterman and Forbus's (1993) taxonomy of similarity relationships:

- *Literal similarity* matches include both common relational structure and common object descriptions;
- *Surface matches*: based upon common object descriptions, plus some first order relations;
- *Structural similarity*, matches based upon common system of relations.

The relations between the various scenarios in a given scenario group can be summarised as follows (see also figure 2): In a group in which re-representation supported inference A, the target and one candidate base scenario (the SSB, or structurally supported base) shared only structural matches; mappings between the SSB's dictionary entry and the base dictionary entry also shared only structural matches.

There was a structural correspondence between the target structure supported by the target dictionary entry and the SSB's dictionary entry which in turn structurally supported the transfer of candidate inference A in the base.

Mappings between the target and the other candidate base scenario (the SMB, or surface match supported base) were supported by shared surface features, and mappings between the SMB's dictionary entry and the target dictionary entry also shared common object descriptions. However, there was a structural correspondence between the base structure supported by the base dictionary entry and the SMB's dictionary entry which supported candidate inference B. This allowed participants to use shared surface features to determine semantic reconciliation, but still use structural correspondences (c.f. Gentner, 1983) to determine their inferences (in this case, making a 'literally similar' match at the analogy level; see figure 2).

In a group where re-representation supported inference B, this pattern of correspondences was reversed.

"DICTIONARY ENTRIES"

BASE DICTIONARY ENTRY

Cronomus lucundus

are unique to certain types of bird
are important to berry eaters
is a long spleen-like organ
keeping berries in the cronomus lucundus allows the berries to slowly ferment, allowing the goodness inside the bitter skins to be released

SMB DICTIONARY ENTRY

Probus razoris

are unique to certain types of bird
are important to berry eaters
is a long plier-like bill
crushing berries in the probus razoris allows the goodness inside their bitter skins to be released without the skins having to be swallowed

SSB DICTIONARY ENTRY

Remulum grandoso

is unique to certain types of dog
are important to dogs which eat a wide range of vegetation
is a short intestine-like organ
keeping vegetation in the remulum grandoso allows it to slowly ferment, allowing the goodness inside the outer skins to be released

Figure 3: Dictionary entries for the Target and two Bases in figure 1. Surface similarities between the Target and SMB are in bold italic print; the structural match between the Target and the SSB is in normal italic

INFERENCES

A. Guralaga live to a ripe old age and rarely suffer from cancer .

B. Guralaga are particularly prone to cancer.

Figure 4: The target inferences for the stimulus group shown of the following pages. In a type A set, structural commonalities would support the A inference; surface similarities would support the B inference. In a type B set, structural commonalities would support the A inference; surface similarities would support the A inference.

To try and simplify the above: in each group of stimuli, the target and candidate base scenario, and their corresponding dictionary entries, shared surface features, and a higher order structural correspondence that corresponded with one candidate inference, whilst the target and the other candidate base scenario, and their corresponding dictionary entries, shared structural correspondences, and a higher order structural correspondence that corresponded with the alternative candidate inference. Each stimulus set was divided into two subsets: in one, structural features in the bases and their novel term dictionary entries supported one set of inferences (Type A), whilst in the second sub-set, the same kind of matches supported the contrasting inference (Type B), so that biases towards a given inference could be eliminated (see figure 2).

Experimental Hypothesis

In keeping with the analysis presented above, we expected that participants would use analogy to reconcile semantic terms in order to perform analogical mappings between the scenarios and generate support for one candidate inference. We predicted that in order to be able to carry out the top level analogy, participants would carry out another analogy in parallel - mapping structures only in the dictionary entries - reconciling semantic terms in a way that supported the top-level 'analogical' structure mapping over the top-level surface mapping, and favour the inference that corresponded to the structurally similar scenario over the scenario that shared only surface features.

Additional Controls and Control Hypotheses

In addition to the basic stimuli, 3 sets of control stimuli were also created:

1 In the main control, the dictionary entries were eliminated, and participants were given only the target and the two candidate bases. In this control, in the absence of any structural support from the dictionary entries for the SSB inference, we expected participants to use the surface commonalities between the target and the SMB to determine their inference choice (i.e the prediction was that when subjects were asked to make an inference in a situation where neither of the base inferences benefitted from any structural bias, partcipants would prefer the inference which was additionally supported at the object level over the inference that received no such support; consistent with the findings of previous studies, such as Gentner, Ratterman and Forbus, 1993, we expected weak similarity to provide more support than no similarity).

2 In the second control, participants were given materials in which the novel terms were removed, and the structural information in the dictionary entries was added to the bases and target - in effect creating 'normal' analogy materials (see figure 5). In this control, where no re-representation was required, we expected the structural commonalities between the target and the SSB to determine the choice of inference, overriding the surface commonalities between the target and the SMB (this would be consistent with previous findings such as Gentner, 1983).

Chateau Bogusse:

is a vineyard.
is in the southern French district of Pretence.
has sandy soils, with a lot of surface pebbles
has a warm microclimate which enables grapes to be produced.
the particular microclimate results in ripe grapes.
the ripeness causes the sugar level in the grapes to rise.
this makes the walls of the grapes weaken and collapse.

Domaine Fraudulent:

grows plums.
has clay soils in which wildflowers grow
is in the western Departement of Maidoop.
its warm microclimate causes melons to grow
the particular microclimate yields extremely ripe plums.
the extreme ripeness causes the plums to become very sweet
this super-sweetness makes the plums soft and squashy
Because of their squashiness Domaine Fraudulent's plums are held in low esteem, and sell poorly.

Mas de la Fiction:

grows grapes.
is in the southern Departement of Whaupper.
has sandy soils, with a lot of surface pebbles
its fine microclimate causes grapes to grow.
the particular microclimate results in ripe grapes.
the ripeness causes some of the moisture in the grapes to evaporate
this evaporation leads to extremely concentrated flavours
Because of their concentrated flavours Mas de la Fiction's grapes are prized and sell for high prices.

Inferences

A. Chateau Bogusse's grapes are highly prized and sell for high prices.

B. Chateau Bogusse's grapes are held in low esteem, and sell poorly.

Figure 5: A control set in which the structural information in the dictionary entries has been included in the bases and target to create an 'ordinary' analogical problem. Surface similarities are illustrated in bold; structural similarities are italicised.

3 In the final control set the dictionary entries were altered so that surface and structural commonalities all supported the same mapping (the LSB, or literally similar base). In this final control, both structural and surface commonalities between the target and the LSB, and their dictionary entries were aligned in support of one inference. Since structure was predicted to be the key factor in deciding inferences (in line with the findings of previous studies), we did not expect the results from this control to differ significantly from the main experimental task.

In all of the controls, the inference supported by the various similarities was again randomised to control for any inherent biases towards particular inferences.

Procedure

Participants were presented with 2 x 6-page questionnaires, each of which contained one scenario set, with its dictionary and candidate inferences, a diversionary task and a scenario set and pair of candidate inferences without a dictionary (the main control). The order in which the sets were presented ('with-dictionary' versus 'without-dictionary' control), was randomised, as was the presentation order of the targets within the sets. A second, smaller group of participants were given the other two controls in similar fashion.

Participants were asked to infer one candidate inference, and give a confidence rating (1=not at all confident; 5=very confident). They were told that the dictionary entries might be useful to them, but told explicitly that the use of them was left to participants' discretion.

Results

Consistent with the initial hypothesis, in the main control condition where no dictionary entries were provided, the inference that received common surface-feature support was favoured by 67% of participants, with only 33% preferring the inference that was not supported by any commonalities, χ^2 (1, N=140) = 17.1, $p<.001$.

However, in the main experimental condition, where definitions – which offered the possibility of structural mappings – were provided, participants reversed their preferred inference for a given target / candidate bases set. Again consistent with the initial hypothesis, in this condition, if participants had preferred the A inference in the first control, when provided with scenario sets where structural commonalities in the dictionary supported the B inference, then participants now chose the B inference. Overall the inferences which received structural support were favoured by 71.2% of participants, with only 28.8% preferring the inference that was supported by surface commonalities alone, χ^2 (1, N=125) = 20.748, $p<.001$.

Also consistent with the initial hypothesis, in the control condition with no novel terms, where structure in the dictionary entries was included in the base and targets, inferences which received common surface-feature support were favoured by only 27.0% of participants, with 73.0% preferring the inference that was supported by structural commonalities, χ^2 (1, N=26) = 3.869, $p<.05$.

There was no deviation from this pattern in the final control condition, where the dictionary entries were altered so that surface and structural commonalities all supported the same base - target (the LSB) mapping, the inference supported by the LSB was favoured by 75.0% of participants, χ^2 (1, N=28) = 5.17, $p<.05$

Analysis of participants' confidence scores in the main control show significantly greater confidence for inferences based on surface commonalities when no structure was present, t=8.72, $p<0.001$. However, this trend was reversed in the other controls and the main experiment - given the choice, participants seem to prefer structurally supported inferences. In the second control condition (analogies) inferences based upon structural commonalities received a significantly higher confidence rating than those based on surface features, t=3.982, $p<0.001$. Similarly, in the main experimental condition, when definitions were provided, inferences based upon structural support received a significantly higher confidence rating than those based only on surface commonalities, t=2.9, $p<.005$. This trend was repeated in the third control, though mean differences were not significant, t=1.02, p =0.33.

Discussion

This experiment seems to show, consistent with the re-representation hypothesis, that participants can use the same process that they used to make analogical inferences to reconcile the semantic discrepancies they encounter in the representations of base and target analogs.

Participants made inferences with the aid of two targets. By controlling the structure of the information representing the 'semantics' of the term that was in turn crucial to the determination of the representation of higher order structure in the base and each of the targets, we were able to control the representations participants used for semantically reconciling particular terms during their analogising. The re-representation processing prediction – that participants would use the same mapping process to match semantic items in their representations as they would in ultimately determining their analogies – appears to be supported by the results of this experiment.

General Discussion

Two very reasonable objections might be made to the results of this experiment:

1. Firstly, the 'dictionary entries' in the main task were artificial: there is a wealth of evidence that definitions are an inadequate basis for conceptual semantics (see Komatsu, 1992). Since the 'dictionary entries' are no more than definitions, it seems reasonable to question whether the use of definitions in exploring conceptual reconciliation affects the validity of our results.

2. A second obvious objection to the findings of the experiment is that participants were presented with the tasks on paper, and had unlimited time in which to solve the inferencing problems, and reconcile and map any 'semantics' in the various base and target specifications. It might be said in objection that since structure mapping is a computationally expensive process – especially in comparison to mapping surface features – this experiment has little relevance to the on-line demand characteristics of analogical processing 'in the wild'. Since participants in this experiment had unlimited time, and external representations of the problems, their behaviour is no predictor of the kind of processes used in making

analogical mappings in memory, where working memory limits will impose restrictions on processing.

Though we acknowledge our sympathy for these objections, neither of them should militate against our interpretation of these results: that the processing of re-representations is possible with externally represented problems. Obviously the second objection – which relates to internal representations – cannot apply to this interpretation. In respect of the first, we note that even though participants used what amounted to definitions in reconciling semantics in the main inferencing task, it is the *processing* that they used to map re-representations in semantic reconciliation (and not the particulars of the representations themselves) that is of interest here. To the extent that participants' processing matched our predictions (and what empirical findings there are in respect of natural representations in similar tasks, e.g. Ahn, 1998), it seems reasonable to assume that this processing was ecologically valid, even if the representations it worked with were not.

These objections do, however, highlight aspects of the re-representation hypothesis that are still seriously under-specified. In particular, the re-representation hypothesis lacks detail concerning the representations it supposes, and the processing demands it appears to make.

In the experiment above, we concentrated on the semantic reconciliation of *one* set of similar-but-not-identical terms, and followed this reconciliation process down through *one* level of recursion, where we saw – consistent with the re-representation hypothesis – that the same process was used to resolve semantic ambiguities as was used to determine analogical mappings.

However, it is unlikely that realistic representations of real-world analogies will contain such a small number of similar yet non-identical predicate matches to reconcile. These representations will contain many more such predicates, and the re-representations of these predicates – whose predicates will need to be matched during semantic reconciliation of the original predicates – may contain many more non-identical but semantically similar predicates, potentially as a factorial of the original number of predicates re-represented in semantic reconciliation.

Logically, at least, this seems to point to both a combinatorial explosion – in terms of the number of predicates to be reconciled, and hence individual semantic reconciliation sub-processes to be run – and a potential infinite regress: if an *identical* mapping process is to be run recursively, and if re-representation doesn't ultimately uncover *identical* predicate-decomposition representations at some level, then mapping may not terminate.

One solution to this problem might be the basic perceptual primitives posited by Hummel and Holyoak (1997; see above). We see two problems with this account: firstly, quite what 'perceptual primitives' are is unclear: at present, they offer no more explanatory clarity than 'concepts' when it comes to explaining semantic reconciliation; and secondly, and more worryingly, this proposal – like all re-representation hypotheses – assumes an almost unlimited capacity for structural mapping in memory. However, recent research (Halford, Wilson and Phillips, 1998) indicates that in reality this is far from the case: human capacity for representing and processing structured information appears to be seriously constrained.[3] In the light of these considerations, we are cautious in inferring too much from the findings reported here. We have shown that re-representation is possible in externally represented tasks. Whether these results *can* be replicated in ecological analogy tasks in memory – and the extent to which re-representation is a viable psychological account of semantic reconciliation – remain open questions in need of further empirical investigation.

Acknowledgements

Thanks to Lera Boroditsky, Ken Forbus, Dedre Gentner, Usha Goswami, Mark Keane, John Lee and Helen Pain for insightful discussions about this work. I'm grateful to Andrew Wishart, Dan Yarlett and Alex Heneveld for comments on an earlier draft of this paper.

References

Ahn W-K (1998) Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition*, 69, 135-178.

Forbus K, Gentner, D, & Law, K (1995) MAC/FAC: A model of similarity based retrieval *Cognitive Science* **19**:141-205

Forbus, K.D., Gentner, D., Markman, A.B., & Ferguson, R.W. (1998). Analogy just looks like high level perception. Why a domain-general approach to analogical mapping is right. *JETA*I 10, 231-257.

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

Gentner, D, Ratterman, MJ & Forbus, K (1993) The roles of similarity in transfer. *Cognitive Psychology* **25**: 524-575

Goswami, U. (1992) *Analogical reasoning in children.* LEA, Mahwah, NJ

Halford, G.S., Wilson, W.H. & Phillips, S. (1998) Processing capacity defined by relational complexity: implications for comparative, developmental and cognitive psychology. *Behavioral and Brain Sciences*, 21, 803-864

Holyoak, KJ & Thagard, P (1995) *Mental Leaps*. MIT Press, Cambridge, Ma.

Hummel, J.E. & Holyoak, K.J. (1997) Distributed Representations of Structure: A Theory of Analogical Access and Mapping. *Psychological Review*, 104, 427-466.

Komatsu, L. K. (1992): Recent views of conceptual structure. Psychological Bulletin, 112, 500-526.

Ramscar, M.J.A. and Hahn, U. (1998) What family resemblances are not. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society*, LEA, Mahwah, NJ 865-870

Ramscar, M.J.A. and 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*, LEA, Mahwah, NJ 346-351

[3] So far we have only been concerned with re-representation in relation to the mapping process: these problems will multiply massively in relation to the semantic reconciliations that must be made in order to facilitate retrieval.