The use of a high-dimensional, "environmental" context space to model retrieval in analogy and similarity-based transfer

Michael Ramscar and Daniel Yarlett
{michael,dany}@cogsci.ed.ac.uk
Division of Informatics
University of Edinburgh
Edinburgh, Scotland EH8 9LW

Abstract

Current models of the retrieval of analogies from a long-term memory store assume mental representations that are generally either underspecified or implausible. In this paper we conduct two experiments which demonstrate that an 'environmental' approach to retrieval can produce appropriate retrieval patterns on cognitively plausible styles of representation, utilising information that can be easily learned from a linguistic environment.

Introduction: Similarity-Based Transfer

Analogy (and similarity-based transfer) is a central cognitive process that represents a versatile problem-solving and reasoning strategy, allowing agents to bring previous experience to bear on novel problems. Its operation embodies two distinct processes: (i) reminding, or *retrieval*, of appropriate analogs from a long-term memory store; after which (ii) candidate analogs are *mapped* onto the representation of the current problem (the target) to determine deeper relational matches, and to allow inferences to be made (Gentner, Ratterman & Forbus, 1993; Forbus, Gentner & Law, 1995; Holyoak & Thagard, 1995).

The latter mapping process has been shown to rely largely on structural commonalities (Gentner, 1983; Holyoak & Thagard, 1995; Hummel & Holyoak, 1997), and computational models of the mapping processes that determine structural commonalities have been subject to much critical scrutiny (Falkenhainer, Forbus & Gentner, 1989; Holyoak & Thagard, 1989; Keane, Ledgeway & Duff, 1994, Hummel & Holyoak, 1997). In contrast *retrieval* has been subject to less investigation. Here, we subject the relatively more neglected issue of modelling analog retrieval to a more focussed theoretical examination.

Four Constraints on Retrieval

Empirical studies by Gentner, Ratterman and Forbus (1993) established four primary constraints on the patterns that an appropriate theory of retrieval should produce given a specific context or probe:

1. *Primacy of the mundane*: The majority of retrievals evoked should be literally similar to the context, sharing both surface and structural characteristics (e.g. a bicycle should call to mind memories of other bicycles).
2. *Surface superiority*: Retrievals based on surface similarity alone (without structural similarity) should also be frequent (e.g. a fairy story about a frog might call to mind other stories about frogs, although the structure of the stories might differ greatly).
3. *Rare insights*: Memories that are structurally similar to the target context should be retrieved only occasionally (e.g. the orbits of the solar system reminding one of electrons orbiting an atom).
4. *Scalability*: The model must plausibly extend to realistically sized memory pools because people typically have vast numbers of memories, and are able to access them in a matter of seconds.

Gentner, Ratterman and Forbus' (1993) investigation demonstrated that retrieval is sensitive to surface (or 'semantic', Hummel and Holyoak, 1997) similarities between a target representation and a base analogy that needs to be to be retrieved. (As opposed to the shared relational structure that determines an analogical match.) The retrieval process, being relatively computationally cheap, acts as an efficient prefilter to the more expensive process of structural alignment (albeit at the expense of potentially passing over useful analogies that share structural commonalities with the target domain).

Meeting the Constraints

MAC/FAC (Forbus, Gentner & Law, 1994) and LISA (Hummel & Holyoak, 1997) are the two foremost models of similarity-based transfer. Below we review the approach taken by both models with regards to retrieval, and examine the theoretical basis for each.

MAC/FAC: Content Vectors

MAC/FAC models retrieval by generating a *content vector* for each representation that is stored in its memory-pool. A content vector summarises the surface features of a representation by recording the frequency with which each *lexically distinct* predicate occurs in it. Thus, the following proposition:

```
(CAUSE (STRIKES-WITH JOHN CUE CUE-BALL)
       (AND (POTS CUE-BALL) (POTS BLACK)))
```

would be assigned the following content vector:

```
((CAUSE . 1) (STRIKES-WITH . 1)
     (AND . 1) (POTS . 2))
```

A measure of the degree that two representations share the same surface features can then be derived by calculating the dot-product of their content vectors (if a particular predicate does not appear in a representation then it is implicit, adopting a sparse-encoding approach, that it has a frequency of zero). It is important to note that only predicates that are *identical* from one another can contribute to the magnitude of a dot-product between two content vectors: there is no potential for multiplying the frequencies of distinct predicates in the dot-product calculation.

Forbus, Gentner and Law (1994) argue that the dot-product between two content vectors provides an empirically adequate measure of the *retrievability* of one representation, given another as a context, because it satisfies the four constraints on retrieval performance.

A Critique of Content Vectors

In order to model the way that *lexically distinct* items in stimuli prime one another for retrieval, the content vector theory makes a commitment to a theory of mental representation we shall call canonical representation (CR) theory. This presupposes a translation procedure that allows tokens that are lexically distinct but share similar semantic "meanings" to be re-encoded using identical tokens. This translation procedure accounts for cross-lexeme priming effects by identically encoding distinct lexemes that should prime for one another, thus ensuring that they can contribute to the dot-product score between the two representations in which they feature. CR theory assumes that during the process of comprehension (representation building):

> "Two concepts that are similar but not identical (such as 'bestow' and 'bequeath') are decomposed into a canonical representation language so that their similarity is expressed as a partial identity (here, roughly, give')." Forbus, Gentner and Law, 1994, pp. 153

'Canonical Form'?

According to CR theory, complex semantic elements can be recursively decomposed -- or *re-represented* -- until a canonical measure of their semantic significance is reached. Hence CR theory assumes that the mental encoding of semantically complex concepts can ultimately be analysed in terms of a stock of canonical forms. Clearly the correctness or otherwise of this assumption is an empirical matter. However, it does seem worth noting that 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 & Hahn, 1998 for reviews). Lacking as it does an account of what a canonical conceptual form is, in its current form CR theory is under-specified, and thus fails to operationalise the notion of semantic similarity in a sufficiently tight manner. This prevents specific predictions being made from the theory (e.g. how strongly do 'cat' and 'dog' prime for one another based on an analysis of the overlap in their shared semantic features?).

LISA: Semantic Features

The other leading model of analogy in the literature, LISA (Hummel & Holyoak, 1997) also relies upon the notion of semantic units (or links) – and re-representations into 'semantic primitives' – in its structured representations to model retrieval. These semantic elements are largely constrained by the representation strategy adopted in LISA (e.g. `\verb+likes1+` or `\verb+likes2+`). Hummel and Holyoak's claim is that these allow appropriate patterns of retrieval to be produced by their model. However, they offer no empirical support for the selection of their particular set of primitive semantic features. At present, the semantic information in LISA's representations is *hand-coded*, and ultimately reliant upon humanistic intuitions about similarities of meaning.

Summary of Current Approaches

Both MAC/FAC and LISA present models of retrieval that are theoretically under-specified. Both accounts rely on the problematic (i.e. currently undefined) notion of re-representation, either into 'canonical conceptual representations' (MAC/FAC) or 'semantic primitives' (LISA). Ultimately, this means that both models rely on hand-coded information to drive their retrievals. Neither LISA nor MAC/FAC actually *models* the representation of lexical information. They rely instead on imported information (primarily intuition) to underpin their behaviour, thus neither can be said – at present – to offer any real explanation of the role of lexico-conceptual knowledge in retrieval.

None of this means, of course, that the shortcomings that we describe in each of the two theories could not ultimately be addressed. We do, however, feel that in the light of these shortcomings there is room for an investigation of whether another approach to the representation of lexico-conceptual knowledge might be used to ground an alternative theory of retrieval.

Co-occurrence Models of Semantics

One approach to lexico-conceptual knowledge that seems promising in this respect is the high dimensional modelling of context spaces. This is a data-intensive technique that analyses a set of corpora, and from this derives a summary of the variety of different contexts that different words can be *used* in. There is a growing body of evidence that the frequency with which different lexemes co-occur with one another (that is, are used together within a particular context, such as a paragraph or moving-window) can provide useful information about the semantic properties of those lexemes.

In co-occurrence analyses, a contextual distribution is calculated for each lexeme encountered in a corpus analysis by counting the frequency with which it *co-occurs* with every other lexeme in the corpora being analysed. The contextual distribution of a lexeme can then be summarised by a vector showing the frequency with which it is associated with the other lexemes in a common linguistic environment. One can think of this information as defining a model containing a network of links between the lexemes in a language, each with varying strengths,

and representing the varying contextual co-occurrences of lexemes in that language. Two such co-occurrence models are the *Latent Semantic Analysis* (LSA) model (Landauer and Dumais, 1997; Landauer, Foltz & Laham, 1998), and the *Hyperspace Analog to Language* (HAL) model (Burgess & Lund, 1997).

There is good evidence that co-occurrence analysis extracts information from corpora that can be used to model certain linguistic behaviour. For example, Landauer and Dumais (1997) report that the LSA model can pass a multiple-choice TOEFL synonym test. Lund, Burgess and Atchley (1995) present evidence that co-occurrence data can act as a good predictor of various priming effects. Burgess and Lund (1997) demonstrate that the HAL model can produce clustering in its high-dimensional space of lexemes from differing grammatical categories.

Whilst the exact parameters of LSA and HAL are different, they both adopt the general approach outlined above to generate co-occurrence vectors. We feel that there are a number of attractive benefits to be gained from modelling the semantic information used in analogical and similarity based retrieval in this way:

1. The proposed semantic metric is clearly specified. By proposing that the semantic information used in retrieval is learned from observing the varying contextual co-occurrences of lexemes in a language, we avoid having to postulate entities – such as 'semantic primitives' whose theoretical and psychological nature is massively under-specified.
2. The semantic information used could be easily learned from the environment,[1] thus avoiding the problems inherent in positing entities whose learnability is somewhat controversial, and whose innateness might otherwise have to be treated as axiomatic (as canonical concepts seem to be; see Laurence & Margiolis, 1999; Fodor, 1981).
3. An environmental context model contains representationally cheap, summarised information, the usage of which makes only limited processing demands. Thus it allows one to avoid the theoretical problems inherent in theories of re-representation which explain cheap surface matches in terms of semantic decomposition and expensive structural alignment (c.f. Holyoak & Hummel, 1997; Forbus *et al*, 1997).
4. Environmental context models are relatively objective: they do not require that a particular set of 'semantic features' are defined before textual analysis begins. Instead the co-occurrence technique takes the lexemes themselves as features, and uses frequency relations between them to define their associativity. This is an advantage given the difficulty we have already highlighted of empirically grounding claims as to the identity of semantic features. Furthermore, the use of dimensional reduction techniques on the vectors associated with each lexeme (Landauer & Dumais, 1997) offers evidence that, in fact, there may not be a *unique* set of semantic features used in the encoding of semantic relations, but rather that multiple encodings can provide sufficient information to meet empirical constraints
5. Because co-occurrence techniques do not rely on a predefined set of semantic features (such as gender, plurality, animacy and so on), this eliminates subjectivity from the decisions that are made during the process of hand-coding representations during the modelling process.

The success of co-occurrence techniques in accounting for priming effects (c.f. Lund, Burgess and Atchley, 1995), has shown them to be useful models of lexical retrieval. Here, we seek to establish whether these models can be used to account for the retrieval of structured composite representations, and not just individual lexemes, from a memory-pool.

The 'Karla the Hawk' Stories

The experiments detailed below use the 'Karla the Hawk' materials as originally used by Gentner, Ratterman and Forbus (1993). The Karla materials consist of twenty sets of stories written in natural language. Each set consists of a base story, and four systematic variations of that story. Two factors are crossed over the four variant stories, as shown in Table 1.

	+ST	-ST
+SF	Literal Similarity	Surface similarity
-SF	Analogy	1st Order Relations

Table 1: The Karla materials

The four story categories systematically vary the commonalities that are shared with the base-story from which they are derived. Each variant can either share or not share surface (±SF) and structural (±ST) commonalities with the corresponding base-story. This 2 x 2 materials design allows for the controlled examination of the sensitivity of various putative measures of retrieval. Gentner, Ratterman and Forbus (1993) found that the prime determinant of retrievability was shared surface commonalities, whilst shared structural commonalities had a nonsignificant effect. This is the pattern of results that we will look for in our experiments. The empirical results reported in Gentner, Ratterman and Forbus (1993) are summarised in Table 2.

	LS	SS	AN	FOR
Retrieval Scores	1.92	1.64	0.44	0.27
Inferential Soundness	4.41	2.70	4.16	2.58

Table 2: The results of the experiments conducted by Gentner, Ratterman and Forbus (1993).

Below, we report two experiments that compare the performance of the content vector (CV) theory of retrieval, as implemented in MAC/FAC, against the measure provided by the LSA model.

[1] Indeed, despite some of the stronger claims made for co-occurrence models of language (c.f. Landauer & Dumais, 1997) we feel that they are best characterised as being essentially models of the associativity of lexemes in a common linguistic environment, such that we prefer to call them "environmental context models". It is also worth noting that co-occurrence techniques are also compatible with a neural implementation. Lowe (1997) demonstrates that a co-occurrence model can easily be implemented as a self-organising Kohonen map, and this offers some support for the idea that some form of co-occurrence counting could occur in the brain.

Experiment 1: Stripped Natural Language.

Experiment 1 was designed to determine whether there is sufficient informational content in a reduced representation of the Karla the Hawk stories to produce retrieval patterns conformable to the empirical data.

It is clear from experimental studies that in addition to the accretion of structural information during comprehension, there is a concomitant loss of superficial verbatim information as propositional representations are built up (Sachs, 1967; Gernsbacher, 1985). Since we wanted to simulate retrieval of what subjects in Gentner *et al*'s studies actually stored (and there is good evidence that people do not store texts verbatim), we decided to initially test retrieval on versions of Gentner *et al*'s stimuli that had all of the *closed-class*[2] lexemes removed from them.

Applying this principle resulted in a set of words for each story which constituted the words which are, in some sense, maximally informative about the context that the representation defines. For example, some words (generally the closed-class words) may occur in almost any (and every) possible context (e.g. 'the' can co-occur plausibly with an extremely diverse set of lexemes). Thus encountering such a word in a probe representation has little informational utility with respect to retrieval because it fails to narrow the set of candidate retrievals at all. Such lexemes are unlikely to influence the kind of retrieval studied by Gentner, Ratterman and Forbus (1993).

The original Karla the Hawk base-story after it had been pruned of all closed-class lexemes is given below, as an example of the characteristic 'bag of words' that remained once the natural language representations had been stripped:

```
Karla old hawk lived top tall oak tree
afternoon saw hunter ground bow crude arrows
feathers hunter aim shot hawk missed Karla
knew hunter wanted feathers glided down hunter
offered give hunter grateful pledged shoot
hawk shot deer
```

Method

The base story for each story-set of the reduced representations was compared with each of its four variants in turn, using the LSA and CV (MAC/FAC content vector) models. This was done in order to reproduce the experimental format embodied in Gentner's original retrieval experiments. The LSA model was set to compare items in *document-to-document* mode, using the 300 most significant factors extracted by the model from a corpus that approximates the general reading a first year college student is exposed to (which seemed appropriate given the participants in Gentner *et al*'s studies). Because of the 2 x 2 design of the experiment, a repeated-measure ANOVA analysis is the appropriate test to determine which of the factors, ±SF or ±ST, the two metrics are sensitive to.

Results

The results of the inter-story comparisons conducted with the LSA and CV models of retrieval are recorded in Table 3. As noted above, each variant story either exhibits ±SF and ±ST, depending on whether it shares or does not share object-attributes and higher-order relations (structure) with the base story it is derived from. The ANOVA analysis revealed that the CV metric was sensitive to both ±SF ($F(1,19) = 11.965$, $p<0.01$) and ±ST ($F(1,19) = 10.027$, $p<0.01$), with no significant interaction effect ($F(1,19) = 3.717$, $p>0.05$). For the LSA metric there was a main effect of ±SF ($F(1,19) = 68.985$, $p<0.01$); no effect of ±ST ($F(1,19) = 2.611$, $p>0.05$), and no significant interaction between the factors ($F(1,19) = 2.428$, $p>0.05$).

	LS	SS	AN	FOR
CV Metric	0.116	0.084	0.057	0.053
LSA Metric	0.442	0.412	0.151	0.152

Table 3: Experiment 1 -- The category means for the CV and LSA scores derived from comparing each base-story with its four variants on the stripped ('bag-of-words') representations. All twenty story-sets had closed-class lexemes removed from them, and were used in the comparison.

Discussion

The clustering in the mean LSA scores for each category of variant (LS-SS and AN-FOR) mirrors the subject data in Gentner, Ratterman and Forbus's (1993) study closely. The same pattern is not observable in the CV metric. Furthermore, the only significant factor in Gentner's original retrieval experiments was ±SF and only the LSA scores conform to this pattern. The CV metric was also sensitive to the ±ST factor, which indicates that it is sensitive to a factor which has been shown to have little significant impact on retrieval performance. It appears that there is sufficient information remaining in the reduced representation to allow different contexts for retrieval to be discriminated from one another in a way that simulates the empirical findings discussed.[3] Moreover, it seems clear from these results that LSA models the original empirical data more accurately than CV.

Experiment 2: Faithful Dgroups

Experiment 2 investigated the performance of the CV and LSA measures on a style of representation that explicitly encodes the structural features implicit in the original stories. This structural information is required to be able to complete the mapping phase of similarity-based transfer, and so these experiments were conducted to determine whether a single style of representation would be sufficient to underpin both the retrieval and mapping processes of similarity-based transfer. The style of representation that we chose shares the substantial core of its form with that used in SME and MAC/FAC, but we developed a series of constraints for translating text into these structured representations whilst avoiding any commitment to the CR theory (we call these representations Faithful Dgroups, 'Dgroup' being the usual term used to describe individual – "chunked" – structured representations in the SME literature.).

[2] Closed-class words belong to the set of words which are closed under the grammatical rules of a language.

[3] It should be noted here that the LSA retrieval scores remain more or less unchanged from pilot testing on the full NL versions. The CV scores, however, are significantly reduced from the original NL materials. This seems to indicate that the LSA model is more robust across representations.

Producing The Faithful Dgroups

Humans are capable of extracting more meaning from language than the basic information that is encoded in the surface structure of texts and dialogues might suggest. To take the following as an example:

```
John hit Mary; Mary cried. The Headmaster
        expelled John.
```

In interpreting this passage, a reader has to infer firstly that John's hitting Mary caused her to cry, and secondly that the relationship between John's hitting Mary, and her crying, caused the Headmaster to expel John. We might express this information in terms of the following nested propositional structure:

```
cause( cause( hit(john,mary), cry(mary) ),
        expel(headmaster,john))
```

None of this causal information appears explicitly in the original utterance, so it is clear that it must in some way be inferred from a prior source. (The need for inference here is uncontroversial: all theories of comprehension agree that language comprehension requires a great deal of active involvement on the part of the comprehender when it comes to inferring information that is not explicitly encoded in language (e.g. McKoon & Ratcliff, 1992); where they disagree is on what, and how much, inference actually happens.)

Whilst we haven't attempted to make a commitment to a particular theory of comprehension in specifying the procedure for translating texts into Faithful Dgroups, what we have tried to do is to provide the beginnings of a method that requires a minimal amount of inference, and is broadly compatible with the bulk of the available data in this area (again, see McKoon & Ratcliff, 1992).

The basic outline of a procedure for forming the Faithful Dgroups from natural language samples is described below.

Algorithm for Construction of Faithful Dgroups

Seeking to maximally preserve closed-class lexical information:

1. Identify the objects that are referred to in the text, and list them using `(sme:defEntity ...)` commands.
2. Identify all the lexeme structures used to express attributes of the objects in the text, and express these as unary expressions.
3. Identify the lexeme structures used to express relations between the identified objects, and express these in the Dgroup form as expressions with two or more arguments, taking only objects as arguments.
4. Now deal with higher-order information (i.e. temporal and causal information that is frequently *implicit* in NL representations). Express this information as expressions taking other expressions as arguments. Note that because this information is often implicit in the NL forms of the stories, a standard (or canonical) lexical identity for each expression must be adopted (this has the effect of minimising the influence of inferred structures on retrieval, which is in accordance with Gentner's empirical findings). The set of inferred relations should be the minimum set required to articulate the narrative structure of the story.[4]

Thus we sought to minimise unwarranted inferences, and the addition of features not warranted by their inclusion in the original materials. In contrast to the original Dgroups, the Faithful Dgroups incorporate much of the lexical information that is present in the original natural language representations.

Method

Faithful Dgroups representing nine of the original story-sets were created.[5] The faithful Dgroup representing the base story for each story-set was then compared with each of its four variants in turn, again using both the CV and LSA models. The LSA model was again set to compare items in *document-to-document* mode, using the 300 most significant factors extracted by the model from the "first year college student, general reading'' corpus.

Results

The result of the CV and LSA comparisons on the Faithful Dgroups are presented in Table 4 below.

For the CV method there was no significant effect of ±SF ($F(1,8) = 3.647$, $p>0.05$), no significant effect of ±ST ($F(1,8) = 3.383$, $p>0.05$), and no interaction effect ($F(1,8) < 1$). For the LSA method there was an effect of ±SF ($F(1,8) = 66.091$, $p<0.01$); no significant effect of ±ST ($F(1,8) = 2.190$, $p>0.05$); and no significant interaction between the factors ($F(1,8) = 1.094$, $p>0.05$).

	LS	SS	AN	FOR
CV Metric	0.751	0.718	0.735	0.688
LSA Metric	0.670	0.633	0.466	0.456

Table 4: Experiment 2 -- The category means for the CV and LSA scores derived from comparing each base-story with its four variants in the Faithful Dgroups. Nine of the NL story-sets were encoded in this format

Discussion

As expected, on representations make no commitment to CR theory – using instead the lexico-semantic information derived from the external representations to drive retrievals – these results demonstrate that the CV method is insensitive to the surface-features of the stories, and thus fails to produce empirically adequate retrieval patterns. This is because the CV method only permits priming between lexically identical items. The LSA method, however, performs much better: its retrievals are only sensitive to the ±SF factor, which is what is required to model the empirical evidence.

It is particularly noteworthy that the LSA method assigned high retrieval scores to the LS and SS categories in this experiment, when their representations need not share any *identical* lexemes with their corresponding base representation. It follows that the LSA model is not simply relying on identical lexemes in distinct

[4] Thus, as with other models of similarity-based transfer, some hand coding of representations does occur (though the freedom to make unprincipled coding decisions is greatly reduced in comparison with other models). This procedure was designed to minimise the influence of such hand coding, although our ultimate goal is the automation of this process.

[5] For comparison purposes, we encoded the same set of stories that Forbus, Gentner and Law (1994) coded for MAC/FAC.

representations to facilitate retrievals, but is modelling instead a more complex kind of relationship between the ways that individual lexemes are used in differing linguistic contexts.

Conclusion

The performance of the LSA measure on both styles of representation offers concrete evidence that it can act as a good predictor of retrieval. That it can do so even when operating on a style of representation that remains faithful to the natural language source of information, and relies on only a psychologically plausible range of inferences for its structure (i.e. a structured, propositional representation that handles lexeme-encoding realistically) is encouraging. As is the fact that we were able to model the empirical data *without* hand tailoring a model of semantics, instead using an objectively, and independently, derived model of lexico-semantic information.

We alluded above to a potential problem in employing the idea of re-representation in retrieval: that studies have shown retrieval to act as a cheap pre-filter for the more computationally expensive – and conceptually rich – process of analogical mapping. Yet the use of re-representation in this process will result in multiple structural mappings being carried out at the conceptual decomposition stage (as many as there are lexically distinct but "semantically" similar items in representations to be mapped). It doesn't take much reflection to realise that will lead to a situation where more structural mapping is required in reconciling semantic differences than in mapping an analogy itself.

At some point mappings between richly represented structure will have to stop, if only because cognitive processing capacity is limited. Our contention is that re-representation – in retrieval at least – is expensive and unnecessary. Structure mappings can be retrieved – and conceptualised – using a far cheaper source of information. Not only does the use of high-dimensional, "environmental" context space to model retrieval in analogy and similarity-based transfer appear to be a plausible approach, it also seems to satisfy Gentner, Ratterman and Forbus' *scalability* constraint better than other models as well.

Given the role structure appears to play in concepts, any *conceptual* solution to matching semantics may suffer from to re-representation problem as well. It may be that *all* conceptualisation – analogical and literal – is about retrieving and mapping the right information in context. Gentner, Ratterman and Forbus (1993) showed that an inexpensive source of information was all that was needed to contextualise retrieval: our results indicate that a of high-dimensional, "environmental" model can provide that context in analogy and similarity-based transfer. Our suspicion is that it might also serve to contextualise broader conceptual processing as well.

Acknowledgements

We are grateful to Andrew Wishart for his help in coding the Faithful Dgroups, and also for insightful comments on an initial draft of this paper. We would like to thank Lera Boroditsky, Ken Forbus and Dedre Gentner for helpful discussion of these issues. This work was supported in part by EPSRC Grant GR/M59846

References

Burgess, C., and Lund, K. (1997). Modelling Parsing Constraints with High-Dimensional Context Space. *Language and Cognitive Processes*, 12, 177-210.

Falkenhainer, B., Forbus, K.D., and Gentner, D. (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1-63.

Fodor, JA (1981) The present status of the innateness controversy. In J Fodor (ed) *RePresentations*, MIT Press

Forbus, K., Gentner, D., and Law, K. (1994). 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, 155-70.

Gentner, D., Ratterman, M., and Forbus, K. (1993). The Roles of Similarity in Transfer: Separating Retrievability from Inferential Soundness. *Cognitive Psychology*, 25, 524-575.

Gernsbacher, M. A. (1985). Surface Information Loss in Comprehension. Cognitive Psychology, 17:324-363.

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

Hummel, J.E., and Holyoak, K.J. (1997). Distributed Representations of Structure: A Theory of Analogical Access and Mapping. *Psychological Review*, 104, 427-66.

Keane, M., Ledgeway, T., and Duff, S. (1994). Constraints on Analogical Mapping: A Comparison of Three Models. *Cognitive Science*, 18, 387-438.

Komatsu, L K (1992) Recent views of conceptual structure. *Psychological Bulletin*, **112**(3), 500-526

Landauer, T.K., and Dumais, S.T. (1997). A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge. *Psychological Review*, 104, 211-40.

Landauer, T.K., Foltz, P.W., and Laham, D. (1998). An Introduction to Latent Semantic Analysis. *Discourse Processes*, 25, 259-84.

Laurence, S. & Margiolis, E. (1999) Concepts and cognitive science. In Laurence & Margiolis (eds.) *Concepts*. MIT Press, Cambridge, Ma.

Lowe, W. (1997). Semantic Representation and Priming in a Self-Organizing Lexicon. *Proceedings of the 4th Neural Computation and Psychology Workshop* pp. 227-39. Springer-Verlag.

Lund, K., Burgess, C., and Atchley, R.A. (1995). Semantic and Associative Priming in High-Dimensional Semantic Space. *In Proc. 17th Annual Conference of the Cognitive Science Society* LEA pp. 660-65. LEA.

McKoon, G., and Ratcliff, R. (1992). Minimal Inference: A Framework for Discourse Processes. *Psychological Review*, 99:440-66.

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

Sachs, R.S. (1967) Recognition Memory for Syntactic and Semantic Aspects of Connected Discourse. *Perception and Psychophysics*, 2, 422-437.