The Competence of Sub-Optimal Theories of Structure Mapping on Hard Analogies

Tony Veale,
School of Computer Applications,
Dublin City University,
Dublin, Ireland.

Mark Keane,
Dept of Computer Science,
University of Dublin, Trinity College,
Dublin, Ireland.

Abstract

Structure-mapping is a provably NP-Hard problem which is argued to lie at the core of the human metaphoric and analogical reasoning faculties. This NP-Hardness has meant that early attempts at optimal solutions to the problem have had to be augmented with sub-optimal heuristics to ensure tractable performance. This paper considers various grounds for qualifying the competence of such heuristic approaches, and offers an evaluation of the sub-optimal performance of three different models of analogy, SME, ACME and Sapper.

1; Introduction

Metaphor interpretation and Analogical reasoning are two, closely related, cognitive faculties which rely upon structure mapping to generate coherent and systematic correspondences between two domains of discourse. But since structure-mapping is clearly a *graph-isomorphism process* which must consider a combinatorial number of such correspondences to generate an optimal mapping, it is both intuitively and provably an NP-hard problem.

A variety of computational approaches to the problem have been described in the AI literature, such as the *Structure Mapping Engine* (SME) of [Falkenhainer *et ah* 1989], the *Analogical Constraint Mapping Engine* (ACME) of [Holyoak and Thagard 1989] and the *Sapper* model of [Veale *et al* 1996a,b]. The first of these, SME, provided an optimal (and thus potentially exponential) solution to the problem, followed by a heuristic, sub-optimal greedy-merge approach (see [Oblinger and Forbus 1990]) and later, an incremental approach (see [Forbus *et al.* 1994]). Falkenhainer *et al.* [1989] provide a complexity analysis of SME that identifies several factors leading to factorial explosion, but argued that analogies producing such difficulties were unlikely to occur. At its heart the original SME is a forest-matching mechanism, which extends known results regarding the $0(N^2)$ complexity of determining sub-tree isomorphism (e.g., see [Garey & Johnson 1979], [Akutsu, 1992]) to forests of inter-tangled tree representations. Layered on top of this forest matcher is a factorial merge process which combines the results of the polynomial sub-tree matching phase (called partial maps, or *pmaps)* into larger, global mappings *(gmaps).* This merge process is clearly $0(2^N)$ where N is the number of pmaps (isomorphic sub-tree matches) involved. SME's designers state that flat representations (i.e., non-nested) that cause N to be large will cause SME to be overly factorial, but that analogies leading to this situation would be rare or incoherent *'jumbles of unconnected expressions'* (p28, 1989).

However, this is shown not to be the case. Veale *et al.* [1996b] demonstrate that many concepts—most notably those that underlie nouns (such as Composer and War) but also story-based or narrative-structured concepts—are essentially *object-centred,* and are most naturally represented as a multitude of shallow trees. These trees *are* highly-connected in a coherent manner by means of shared arguments (common leaves). A mapping between two such domains is illustrated in Figure 1 below:

(control Composer Orchestra)		(control General Battle_Plan)
(control Composer Baton)		(isa Theatre_of_War Theatre)
(part Orchestra Percussion)		(control General Sabre)
(control Theatre_Manager Concert_Theatre)		*(perform Soldier Warfare)*
(part Orchestra Musician)		(part Army Artillery)
(attr Composer Artistic)		(control General Army)
(control Musician Musical_Instrument)		(part Army Soldier)
(part Percussion Drum)		(part Artillery Cannon)
(attr Drum 1812-Overture)		*(attr General Ruthless)*
(control Composer Musical-Score		(control Soldier Musket)
(isa Concert_Theatre Theatre)		(attr Cannon Russian_Campaign)

Figure 1: *Partial domain descriptions relating to the concepts Composer and* General. *Noisy predications which do not contribute to the metaphor "Composers are Generals" are shown in italics.*

This type of commonly occurring metaphor exacerbates SME's original $0(2^N)$ complexity. For example, the metaphor *Surgeons are Butchers* requires less that 15 seconds of processing time in Sapper, yet generates enough pmaps to keep an optimal SME busy for many billions of years. Veale *et al.* report the results of Figure 2 for an empirical test involving over 100 object-centred representations (drawn from the domain of professions):

Avg. # pmaps	386 per metaphor	12,657	18
Avg. Time per Metaphor	N/A - worst case $0(2^{386})$ seconds	N/A	12.5 seconds

Figure 2: Evaluation of Sapper, SME and ACME. The unavailability of times for SME and ACME reflects the inability of these models to run in a matter of days.

Sapper out-performs SME in these domains because it is designed to seek out structure *laterally* from shallow trees that are connected via common elements, while SME seeks structure vertically from the hierarchical nesting of deep trees. Sapper can also encode the kind of verb-centred story analogues upon which SME demonstrates its strongest competence (such as the *Karla the Hawk* and *Fortress: Tumor* stories). This paper builds upon these results to show that while heuristic, sub-optimal *Greedy-SME* and *Incremental-SME* avoid factorial time performance, they are still very sensitive to tree organization, producing poor mappings when dealing with object-centred representations. We also apply our intuitions to the Sapper and ACME models, demonstrating that while the latter is from the outset a sub-optimal model, it also exhibits diminished competence on hard problems. This discussion will allow us to outline in greater detail exactly what we mean by a *hard* analogical problem.

2. Cognitive Theories of Structure Mapping

ACME approaches the structure-mapping problem from a different perspective than either SME or Sapper, pursuing what might be called a *natural computation approach* to analogy and metaphor. ACME models structure-mapping as a problem of parallel constraint satisfaction, in which the demands of 1 -to-1 coherence and structural systematicity are coded as soft constraints, or pressures, on the system. Ultimately, it is a sub-optimal approach which offers no guarantees of mapping quality.

ACME employs a Hopfield-style connectionist network to encode mapping constraints (see [Hopfield and Tank,

1985]). Every structure-mapping hypothesis—either between a source and target predicate or between a source and target entity—is coded as a distinct neuron. Likewise, structural entailments among these hypotheses are coded as bi-directional excitatory links between the corresponding nodes, while inhibitory links are used to connect mutually exclusive hypotheses.

Such an arrangement is the connectionist equivalent of a 2-CNF SAT formula, raising the question of ACME's logical soundness. Indeed, it happens that the use of bi-directional linkages in ACME—which makes all implications *mutual* implications—means that an ACME representation is logically unsound. Because argument mappings can dictate predicate mappings, ACME is sound only when the source and target structures are trees (hence argument mappings *do* imply predicate mappings), but as noted, structure mapping is polynomially bounded anyway in such situations.

Overall, the complexity prognosis of ACME is not good: as a feedback-based neural network, there is no guaranteed polynomial bound on its time performance. Yet, because the network size is polynomially-bounded (i.e., $0(n^2)$ nodes and $0(n^4)$ linkages, where n is the number of distinct symbols in the source domain), the theoretical results of [Brack and Goodman 1990] apply, who prove that a Hopfield-style network of polynomial size can only optimally solve NP-hard problems if NP = p. So, since an ACME network realistically embodies a polynomial algorithm, why should it be allowed to consume an exponential amount of time doing so?

2.1. Sapper: A Memory-Situated Model

The Sapper model of [Veale *et al* 1996a,b] views semantic memory as a localist graph in which nodes represent distinct concepts, and arcs between those nodes represent semantic / conceptual relations between those concepts. Memory management under Sapper is pro-active toward structure mapping, that is, it employs rules of structural similarity—called Triangulation and Squaring—to determine if any two nodes may at some future time be placed in systematic correspondence in a metaphoric context. If so, Sapper notes this fact by laying down a *bridge relation* between these nodes, to be exploited in some future structure-mapping session.

Triangulation Rule: *If memory already contains two linkages L_{ij} and L_{kj} of semantic type L forming two sides of a triangle between the concept nodes C_k, C_i and C_j, then complete the triangle and augment memory with a new bridge linkage B_{ik}.*

Spread Activation from node T in long-term memory to a horizon H
Spread Activation from node S in long-term memory to a horizon H

When a wave of activation from T meets a wave from S at a bridge T' :S'
   linking the target domain concept T' to the source domain concept S' Then:

Find a chain of semantic relations R that links both $\mathbf{T'}$ to T and S' to S
If R is found, then the bridge $\mathbf{T':S'}$ is balanced relative to T:S, so Do:

Generate a partial interpretation (pmap) π of the metaphor T:S as follows
   For every target concept t between $\mathbf{T'}$ and T as linked by R Do
      Put t in alignment with the equivalent concept s between S' and S
      Thus, $\pi \leftarrow \pi \cup \{<t: s>\}$
Let $\Phi \leftarrow \Phi \cup \{\pi\}$

Once the set Φ of all pmaps within the horizon H have been found, Do
   Evaluate the richness of each pmap $\pi \in \Phi$
   Sort the collection Φ of pmaps in descending order of richness.
   Pick the first (richest) interpretation $\Gamma \in \Phi$ as a seed for overall interpretation.

   Visit every other pmap $\pi \in (\Phi - \Gamma)$ in descending order of richness
      If it is coherent to merge π with T (i.e., without violating 1-to-1ness) then
         $\Gamma \leftarrow \Gamma \cup \pi$
      Otherwise discard π

   When Φ is exhausted, Γ will contain the overall Sapper interpretation of T:S

Figure 3: The Sapper Algorithm, as based on the exploitation of cross-domain bridge-points in semantic memory.

Squaring Rule: *If Bjk is a bridge, and if there already exist the linkages L_{ij} and L_{ik} of the semantic type L, forming three sides of a square between the concept nodes C_i, C_j, C_k and Cl, then complete the square and augment memory with a new bridge B_{il}.*

At some future time, if Sapper wishes to determine a structural mapping between a target domain rooted in the concept node T (for Target) and one rooted in the node S (Source), it applies the algorithm of Figure 3.

The Sapper algorithm comprises two main phases: the first of these seeks out the set Φ of all well-formed and balanced semantic pathways (of length < 2H) that originate at the root node of the target (T), and terminate at the root node of the source (S), crossing a single conceptual bridge (i.e., the domain cross-over point) at its mid-point. Each such pathway corresponds to a partial interpretation (a pmap in SME parlance) of the metaphor/analogy. The second phase coalesces this collection of pmaps Φ into a coherent global whole; it does this using a *seeding algorithm* (see [Keane and Brayshaw, 1988]) which starts with the structurally richest pmap Γ as its seed, and then attempts to fold each other pmap into this seed, if it is coherent to do so, in descending order of the richness of those pmaps. This seeding phase is directly equivalent to the greedy merge phase of Greedy-SME (see [Oblinger & Forbus 1990).

3. Proof: Structure-Mapping is NP-Hard

In this section we place our arguments on a solid footing by proving the NP-Hardness of the structure mapping problem. Though the known NP-complete problem LCS (Largest Common Sub-Graph) is perhaps a more immediate match, we instead employ here 3DM (3-Dimensional Matching) as a proof basis, a problem which seeks to obtain a non-overlapping matching of points in a 3-D space. A consideration of 3DM will shed light on the worst case scenario as encountered by the greedy heuristics employed by greedy-SME. Garey and Johnson ([1979]) define 3DM as follows:

Unique 3-Dimensional Matching (3DM): Given a set M of points in 3-D space, i.e., $M \subseteq X \times Y \times Z$, where X, Y and Z are disjoint sets of integers and $|X| = |Y| = |Z| = q$, find the largest set $M' \subseteq M$ such that no two elements of M' agree in any coordinate.

Proof: To reformulate 3DM as a problem of structure-mapping, it is necessary to represent each 3-D point $<X_i, Y_i, Z_i> \in M$ as a pair of predicates, one in each of the source S and target T domains, such that these predicates are only allowed to map onto each other. Furthermore, any isomorphic mapping must not contain two different predicate matches that arise from two points sharing one or more coordinates. We can ensure this using the following polynomial transformation:

$$\forall <X_i, Y_i, Z_i> \subseteq M \; Do$$
$$add \; P_{XYZ}(X_i, Z_i) \; to \; S$$
$$and \; add \; P_{XYZ}(Y_i, \Omega^* X_i + Y_i) \; to \; T$$
$$where \; \Omega = max(max(X \cup Y \cup Z), |min(X \cup Y \cup Z)|)$$

Now, because the predicate P is uniquely tagged with the subscript XYZ which ties it to a particular 3-D point, these two predicate structures can map only to each other. When so mapped during the analogy process, such a mapping results in the creation of the following structure *(a pmap* in SME parlance):

$$map_i = \{<X_i, Y_i>, <Z_i, \Omega^* X_i + Y_i>\}$$

In this manner a root mapping will be created for each point in M. Note also that $\Omega^* X_i + Y_i$ is unique for each pairing of X_i and Y_i, thus X_i, Y_i and Z_i are *tied together* and cannot be cross-mapped with any other point coordinate. Suppose we have two such pmaps, mapi = $\{<X_i, Y_i>, <Z_i, \Omega^* X_i + Y_i>\}$ and map_k = $\{<X_k, Y_k>, <Z_k, \Omega^* X_k + Y_k>\}$, arising out of the two points $<X_i, Y_i, Z_i>$ and $<X_k, Y_k, Z_i>$ which share a Z-coordinate Z_i. These maps cannot therefore be merged to create a larger mapping as such a merge results in Z_i being mapped to both $\Omega^* X_i + Y_i$ and $\Omega^* X_k + Y_k$, a clear violation of mapping isomorphism.

Once a maximal gmap is found for the analogy, each pair $<X_i, Y_i>$ and $<Z_i, \Omega^* X_i + Y_i>$ of this gmap can then be decomposed and reassembled (in polynomial time) to recreate a point $<X_i, Y_i, Z_i>$ that is added to M'. Since the gmap is maximal, so is M'. Because it solves 3DM, structure-mapping is thus NP-Hard. $\square$

4. Problem Reorganization for Tractability

A large body of problem instances may nevertheless be tractably amenable to an optimal Sapper variant. If an optimal Sapper solution can be obtained for a large enough body of problem examples, these solutions can be used as ceilings to measure the competence of sub-optimal heuristics like greedy merging / seeding.

The domain descriptions in the Sapper profession corpus contain on average over 120 predications each. Test metaphors in the profession corpus thus generate too many partial mappings to make optimal evaluation tractable. Yet, some problem re-organization can be applied to reduce the number of pmaps to frequently make an Optimal-Sapper interpretation feasible, without losing the combinatorial scope of the interpretation. This reorganization process, whereby redundant areas of the combinatorial search space are pruned, is the equivalent

of *arc-consistency testing* in satisfaction problems to a priori remove contradictory variable assignments (see [Mohr and Henderson, 1986]).

For each metaphor (whose pmap set is denoted Φ) a conflict graph is constructed in $O(|\Phi|^2)$ time, by determining for each pmap the set of other pmaps with which it cannot be combined. This set is similar to the set of *NoGoods* calculated by the SME algorithm, though it used differently to achieve more extensive reductions in performance time. The conflict set CF_i for a particular pmap $\pi_i \in \Phi$ is thus defined as:

$$CF_i = \{\pi_k \mid k \neq i \land \neg systematic(\pi_i, \pi_k)\}$$

Compatibility between pmaps can thus be defined as:

$$compatible(\pi_i, \pi_k) \; iff \; CF_i \subseteq CF_k \land \pi_i \notin CF_k$$

In contrast to SME, Optimal-Sapper uses this information to recognize any compatibility-based redundancies, and redistribute them accordingly *before* entering the punishing factorial merge-stage, as follows:

$$\forall \pi_i \pi_j, \; i \neq j, \; if \; compatible(\pi_i, \pi_j) \; then$$
$$\forall \pi_k \in CF_j - CF_i \; do$$
$$\pi_k \leftarrow \pi_k \cup \pi_i$$
$$\Phi \leftarrow \Phi - \pi_i$$

Given that the combinatorial merge stage of an Optimal-Sapper algorithm is $O(2^{|\Phi|})$, each such pmap factored out a priori lowers the eventual cost another exponential notch. On our corpus of profession metaphors, we have found that problem reduction of this form reduces the number of pmaps for each metaphor by an average of 60%, pruning the search space of the most intractable instance, *Generals are Surgeons,* from $O(2^{39})$ to one more manageable by Optimal-Sapper, $O(2^{17})$.

5. Experiment: Sapper Vs.& Greedy-SME

We can now quantify the competence of sub-optimal heuristics such as seeding and greedy-search as a percentage of optimal performance. But first, we consider the nature of the interpretations that structure-mapping algorithms will generate for these test metaphors. The mapping of Figure 1 is the Sapper interpretation of the metaphor *Composers are Generals,* while the mapping of Figure 4 is that returned by greedy-SME for the same metaphor.

Since an official implementation of greedy-SME is not yet publicly available ([Forbus, 1996]), we therefore simulate greedy-SME by feeding the pmaps generated by the available optimal-SME through the Sapper seeding stage. This a computationally equivalent process.

If Composer is like General
Then Drum is like Cannon
and Powerful is like Loud
and Loud is like Powerful
and Conductor_Baton is like Sword
and Tchaikovsky is like Napoleon
and Libretto is like Plan
and Narrow is like Dangerous
and 19th_Century is like French
and Music_ecital is like Cavalry_Charge
and Long is like Sharp
and Orchestra is like Army
and Listener is like Soldier
and W_A_Mozart is like George_Patten
and Percussion is like Artillery
and Theatre is like Influential
and Russian is like 19th_century
and Music_Composition is like Bomb_Raid
and Musical is like Healthy
and Music_Note is like Enemy_]Soldier
and Sudden is like Dead
and Piano is like Snub_Fighter
and Fictional is like On_Target
and Character is like Smart_Bomb
and 18th_Century is like Arrogant
and Symphony is like Military_Propaganda
and Violin is like Musket
and Musical_Score is like Enemy_Army
and Operatic_Act is like Medal
and Opera is like Militaiy_Uniform
and Inspiration is like Corpse

Figure 4: Simulated Greedy-SME interpretation of "Composers are Generals".

A selection of the mappings in Figure 4 above are displayed in an italics face to convey their 'ghost' status: 'ghosts' are essentially noisy mappings that might work in another metaphoric context but which are not systematic here. But why does greedy-SME generate so many ghosts while Sapper produces none, when both employ equivalent merge processes? To see why, consider that SME and Sapper agree on three tacit assumptions for seeding: first, that a goodness ordering can be placed upon the set of pmaps; secondly, that the pmap chosen as seed for the merge is rich enough to justify its own inclusion in the global mapping; and thirdly, that this seed is rich enough to nudge the overall merge process toward a good to optimal global mapping. However, Greedy-SME—unlike Sapper—does not generate sufficiently rich (and thus differentiable) pmaps in object-centred domains to make these assumptions work. As these domains are best represented as a broad

forest of many shallow trees rather than a tight forest of few, deep trees (see [Veale et al 1996]), the pmaps generated by SME for object-centred metaphors are equally shallow and numerous. In fect, these impoverished SME pmaps resemble the geometric pmaps generated in section 3 when reposing 3DM as structure-mapping. One clearly would not expect a greedy approach to work in this geometric context as no one pmap would have enough structure to successfully guide the merge process to a good solution.

The competence of Sapper and greedy-SME has been determined over the test corpus of 100+ profession metaphors, where the optimal-Sapper of section 4 is used as a savant: a mapping of a sub-optimal interpretation is considered valid if it is also contained in the optimal Sapper interpretation. The sub-optimal competence of Sapper and greedy-SME is thus calculated as 100 * (No. of valid mappings) / (Total No. of mappings). If this validity criterion seems overly strict and all-or-nothing, it needs to be for tractability reasons. If one were to evaluate a noisy interpretation on the basis of its largest systematic subset, the partition of the interpretation into signal and noise would in itself be an intractable problem of combinatorial dimensions. Comparative results are displayed in Figure 5 below:

Competence	95.2%	18.7%	100%	80.5%
% of Times Optimal	77%	0%	100%	45%

Figure 5: Comparative trials of Sapper and Sub-Optimal greedy-SME with a random control.

Greedy-SME performs disappointingly on these trials, significantly trailing even the random control trial, in which a random merging of coherent Sapper pmaps is generated as an interpretation for each metaphor. These results speak for the importance of structurally rich pmaps, for when these are rich enough even a random coalescence of pmaps will generate a good interpretation. But if the set of pmaps is structurally impoverished, as with SME in object-centred domains, not even a best-first sorting will compensate. These random trials indicate that a system's true competence is to be found in the processes which generate pmaps, more so than in those which combine them.

6. Where the Hard Analogies Are

What do the results of section 5 say about the identifiable qualities of hard analogies/metaphors? Clearly, when employing an optimal mapping

algorithm, the number of distinct roots in the forest-of-trees representation of each domain is a direct indicator of the exponential requirements of the algorithm. What can be said of the hardness of analogies as perceived by sub-optimal approaches such as Sapper and SME?

In complexity terms no problem instance is—strictly speaking—*hard* to a sub-optimal structure matcher, as the number of pmaps is largely irrelevant in a $O(N^2)$ greedy merging / seeding process. However, if one measures hardness in terms of the likelihood of generating a quality (i.e., ghost-free and accurate) interpretation, the best indicator of hardness is the average structural richness of each pmap (i.e., the average number of mappings in each pmap). The lower this average richness, the more probable it is that any two pmaps can be coherently merged, and thus the more likely that the final interpretation will be noisy and ghost-ridden. In contrast, the higher this average, the more probable it is that a final interpretation will be near optimal, and less likely to contain ghosts (as each pmap merge operation will have a greater chance of failure).

If we have side-lined ACME'S sub-optimal approach to structure-mapping in this paper, it is due to the belief that ACME represents an excessive approach to the problem. Recall that ACME can be characterized as a 2-SAT problem, where network nodes mirror SAT variables, and network linkages mirror SAT clauses, From the ratio of ACME nodes to linkages for any given metaphor/analogy, we can determine the equivalent SAT ratio of clauses to variables as $0(N^2)$, thus making an ACME problem hugely over-constrained (see [Mitchell *et al.* 1992]). Given the large networks which ACME can construct for a hard problem (> 12,000 nodes), existing relaxation techniques based on constraint prioritization do not seem practical (see [Bakker *et al.* 1993]). ACME thus reduces to a difficult subclass of *maximal 2-SAT,* with the size of that subset of clauses it must leave unsatisfied growing exponentially with the extent of network over-constraint, which itself grows quadratically with metaphor size. In this case, sub-optimality certainly thus not imply tractability.

In closing, we note that the profession corpus upon which our experiments are based is available from the following URL, in Sapper, ACME and SME formats:
http://www.compapp.dcu.ie/~tonyv/metaphor,html

References

[Akutsu, 1992] T. Akutsu. An (RNC) Algorithm for Finding a Largest Common Subtree of Two Trees, *IEICE Transactions on Information and Systems,* 75-D, pages 95-101, 1992.

[Bakker *et al.* 1993] R. Bakker, F. Dikker, F. Templeman, and P. Wognum. Diagnosing and Solving Over-determined constraint satisfaction problems, in the Proceedings of IJCAI'93, the thirteenth International Joint Conference on Artificial Intelligence. Morgan Kauffinan, 1993.

[Bruck and Goodman, 1990]. J. Brack and J. W. Goodman. On the Power of Neural Networks for Solving Hard Problems, *Journal of Complexity* 6, pages 129-135, 1990.

[Falkenhainer *et al.* 1989] B. Falkenhainer, K. D. Foibus, and D. Gentner. Structure-Mapping Engine. *Artificial Intelligence,* 4 1, pages 1-63, 1989.

[Forbus *et al.* 1994]. K. D. Foibus, R. Ferguson and D. Gentner. Incremental Structure-Mapping, in *the Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society,* Atlanta, Georgia. Hillsdale, NJ: Lawrence Erlbaum, 1994.

[Forbus, 1996]. K. D. Foibus, personal communication, August 1996.

[Garey and Johnson, 1979] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* Freeman, NY, 1979.

[Holyoak and Thagard, 1989] K. J. Holyoak and P. Thagard. Analogical Mapping by Constraint Satisfaction, *Cognitive Science* 13, pp 295-355, 1989.

[Hopfield and Tank, 1985] J. J. Hopfield and D. W. Tank. "Neural" Computation of Decisions in Optimization Problems. *Biological Cybernetics* 52, pp 141-152, 1985.

[Keane and Brayshaw, 1988] Keane, M. T. and M. Brayshaw. The Incremental Analogical Machine: A computational model of analogy. *In D. Sleeman (Ed.), European Working Session on Learning.* Pitman, 1988.

[Mitchell *et al.* 1992] D. Mitchell, B. Selman and H. J. Levesque. Hard and Easy distributions of SAT problems, *in the proceedings of AAAI'92, the 1992 conference of the American Association for AI,* 1992.

[Mohr and Henderson, 1986] R. Mohr, and T. Henderson. Arc and Path Consistency Revisited. *Artificial Intelligence* 25, pages 65-74, 1986.

[Oblinger and Foibus, 1990] D. Oblinger, and K. D. Foibus. Making SME Pragmatic and Greedy, in *the Proc. of the Twelfth Annual Meeting of the Cognitive Science Society.* Lawrence Erlbaum, 1990.

[Veale *et al.* 1996a] T. Veale, B. Smyth, D. O'Donoghue and M. T. Keane. Representational Myopia in Cognitive Mapping, *in the Proc. of the AAAI workshop on Source of the Power in Cognitive Theories,* Portland, 1996.

[Veale *et al.* 1996b] T. Veale, D. O'Donoghue and M. T. Keane. Computability as a Limiting Cognitive Constraint: Complexity Concerns in Metaphor Comprehension, *Cognitive Linguistics: Cultural, Psychological and Typological Issues (forthcoming).*