‘Soft’ Approaches to Analogical Alignment and Sub-Graph Isomorphism

Tony Veale, School of Computer Applications, Dublin City University, Dublin, Ireland.
Email: Tonyv@compapp.dcu.ie

Abstract: Structure-mapping is a process of isomorphic sub-graph alignment that is exploited in metaphor, analogy and exemplar-based reasoning. Once a mapping between two knowledge structures from different domains is created in this way, knowledge from one domain (the well-understood *source*) can be transferred to the other (the less-understood *target*). However, in many cases both domains might be semantically analogous yet structurally non-isomorphic, requiring an element of structural *slippage* or *warping* to accommodate an isomorphic mapping between both. This paper considers a principled approach to the problem.

1. Introduction

Subgraph isomorphism is an NP-hard problem central to many areas of Artificial Intelligence— particularly those considered to be *knowledge-intensive*—such as analogy-driven problem-solving and theorem proving, case-based reasoning and memory-based translation [1]. Each of these domains requires a *match-and-retrieve* mechanism capable of searching a large memory of past cases for an exemplar structure (called the *source*) that most resembles the current problem (the *target*), and of generating an isomorphic mapping between both that guides the adaptation of the source to suit the target. Since large case-bases can contain heterogeneous knowledge specified at differing levels of detail and redundancy, realistic systems must seek the best near-isomorphism between the largest subgraphs of the source and target, rather than a complete isomorphism.

Of the existing approaches to this problem the most notable make use of the *‘soft’* computing paradigm. For example, the ACME model of [2] is an analogical mapping engine that employs a Hopfield-style connectionist network to generate both isomorphisms and homomorphisms between subgraphs of the source and target representations (essentially semantic networks). However, even when a complete isomorphism between source and target is possible, ACME is not guaranteed to find it; neither is it always guaranteed to return a mapping that is entirely coherent and systematic. Overall then, the flexibility of the soft paradigm remains underconstrained in ACME.

Another interesting approach is the Copycat model of [3], which employs a stochastic, temperature-driven view of the mapping process that resembles the directed randomness of simulating annealing. Copycat is built upon a slipnet of concepts (semantic labels) that indicates how, under certain stochastic conditions and temperatures, one particular concept or relation can be coerced into another. This allows Copycat to fluidly change its perspective on a given structure, and to subsequently warp this structure to achieve a satisfying mapping. However, while this stochasticity yields cognitively-plausible results, it is also responsible for Copycat’s inherent non-determinism: when more than one mapping is possible, one cannot predict which will be chosen. Similarly, Copycat is not constrained to return the same mapping in identical situations.

In this paper we describe some ‘soft’ extensions to an existing model of analogical structure mapping—called *Sapper* (see [4])—which allow Sapper to generate systematic mappings between non-isomorphic graph structures (i.e., semantic networks). Sapper is a *strong* structure-matcher in the mold of the Structure-Mapping-Engine (SME) of [5], inasmuch as it will never generate a mapping that is internally inconsistent (unlike ACME, for instance). The challenge of accommodating structural slippage in such strong models is to allow warping of the source or target domains in a principled fashion, such that the end mapping, while not an isomorphism, is nevertheless semantically and logically sound.

2. Sapper: Memory-Situated Mapping

The Sapper model of Veale *et al.* ([4,6,7,8]) views semantic memory as a localist graph in which nodes represent distinct concepts, and arcs between those nodes represent semantic/conceptual relations between concepts. Memory management under Sapper is pro-active toward structure mapping, that is, it employs rules of structural similarity to determine whether any two given nodes may at some future time be placed in systematic correspondence in a metaphoric context. If so, Sapper notes this by laying down a *bridge* between these nodes, which can be exploited in some future structure-mapping session. The rules Sapper employs to lay down bridges, termed *Triangulation* and *Squaring*, are defined below:

Triangulation: *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 conceptual bridge linkage B_{ik}.*

```
Spread Activation from nodes (T)arget and (S)ource in memory to a horizon H
   When a wave of activation from T meets a wave from S at a bridge T':S'
          linking the tenor domain concept T' to the vehicle domain concept S' Then:
       Find a path of semantic relations R that links both T' to T and S' to S
       If R is found, then the bridge 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 tenor concept t between T' and T as linked by R Do:
         |        Put t in alignment with the equivalent concept s between S' and S
         |        π ← π ∪ {<t : s>}
         | Φ ← Φ ∪ {π}
   Once the set Φ of all pmaps within the horizon H have been found, Do
       Evaluate the richness of each pmap π ∈ Φ
       Sort the collection Φ of pmaps in descending order of richness.
       Pick the first (richest) interpretation Γ ∈ Φ as a seed for overall interpretation.
       Visit every other pmap π ∈ (Φ - Γ) in descending order of richness
          If it is coherent to merge π with Γ (i.e., without violating 1-to-1ness) then
               Γ ← Γ ∪ π
          Otherwise discard π
       When Φ is exhausted, Γ will contain the overall Sapper interpretation of T:S
```

Figure 1: The Sapper Algorithm, as based on the exploitation of cross-domain bridges in semantic memory.

Squaring: *If B_{jk} is a bridge, and if there already exists the linkages L_{ij} and L_{lk} of the semantic type L, forming three sides of a square between the concept nodes C_i, C_j, C_k and C_l, then complete the square and augment memory with a new bridge linkage 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 (for Source), it uses the algorithm of Figure 1 above. The algorithm comprises two main phases: the first of these seeks out the set Φ of all well-formed and balanced semantic pathways (of length $\leq$ 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 Φ of pmaps into a coherent global whole (i.e., an SME *gmap*); it does this using a *seeding* algorithm which starts with the structurally richest pmap Γ as its seed, and then attempts to fold every other pmap into this seed, if it is coherent to do so, in descending order of the richness of those pmaps. This algorithm is computationally equivalent to the greedy merge algorithm of [9].

3. Principlesof StructuralSlippage

Underlying this soft variant of Sapper is a Copycat-like *slipnet* in which different semantic relations (i.e., graph labels) are probabilistically connected (e.g., $P_{slip}(Part \rightarrow Contain) = 0.9$). Operating in conjunction with this slipnet is a collection of semantically motivated structure-warping rules, essentially soft variants of the standard Sapper rules, which allow non-isomorphic structures to be mapped. Soft-Sapper thus allows analogies to be created between domains that have been defined at different levels of detail and redundancy: for instance, in the SportsCar domain one might state that the Engine contains Pistons which control the Wheels, or alternately, that the Pistons control the Crankshaft which in turn controls the Wheels. In mapping this source structure then to that of either Jaguar or Puma say (an analogy used by Ford for two of their sports cars), it may be necessary to either contract or stretch the target structure to accommodate the possible occurrence of the node Crankshaft (which might or might not map to Leg-Muscle, say).

Given two pmaps of equal depth (i.e., each composed of paths of a given length), a probabilistic *rigidity* measure of how much slippage each involves can be ascertained, as a product of the necessary slippage probabilities entailed by each. Thus, a pmap that maps *X—part→Y—contains→Z* to *A—part→B—contains→C* has a rigidity measure of 1.0, while one that maps the same path to *A—contains→B—part→C* has a rigidity measure of $0.9 \times 0.6 = 0.54$. These measures can in turn be incorporated into a quality metric that prefers rigid pmaps over their looser variants that have *slipped*.

If $S \rightarrow \ldots \rightarrow S_1 — R_1 \rightarrow S_2 — R_2 \rightarrow S_3$ *is a path under investigation in the source domain,*
and $P_{slip}(R_1, R_2) > \varepsilon$ *(a minimal rigidity threshold)*
Then
$S \rightarrow \ldots \rightarrow S_1 — R \rightarrow S_3$ *is another path in S that should also be pursued in the source*
Where
$R = R_1$ *if* R_1 *is a causal relation, otherwise* $R = R_2$

Figure 2: The Core Slippage Principle employed in Sapper.

4. Mechanicsof StructuralSlippage

Given the existence of a relational slipnet to handle label slippage, the complementary problem of structural warping can be handled with the single, compositional rule of Figure 2. The action of this rule is simple yet effective: two successive semantic relations R_1 and R_2, linking two concepts S_1 and S_3 via an intermediary S_2, can be *snipped* to produce a path that links S_1 and S_3 directly; if R_1 is a causal relation (such as *Cause, Enable, Support*, etc.) then it is favoured as the relation that directly connects S_1 and S_3; otherwise R_2 is chosen. If applied at every stage of a given pathway's development, this rule is capable of removing a significant number of linkages, as many are as needed to make the pathway structurally isomorphic with a mirror pathway in the target domain. For instance, *Part→Cause* reduces to *Cause*, as does *Cause→Part*, while *Part→Substance* and *Part→Contains* both reduce to *Part*. As illustrated in Figure 3, the concepts *Engine* and *CrankShaft* are temporarily removed from the source picture to accommodate a mapping between *Muscle* and *Piston*.

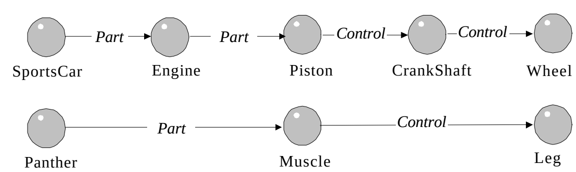

Figure 3: Path simplification in the SportsCar domain yields a path isomorphism with the Panther domain.

Note that this convenient deletion of *Engine* is indeed both temporary and non-destructive, inasmuch as it effects just this single pathway. Other pathways that ultimately find a mirror partner in the target domain may instead provide a mapping for *Engine* (for instance, a pathway between *Fuel-cap* and *Exhaust-pipe* will necessarily pass through *Engine*, mapping it to either *Brain* or *Heart*). Figure 4 presents the complete mapping for the *SportsCar* as *Jaguar* metaphor when the slippage rule is active. Comparing those mappings which are underlined in Figure 3 (depicting those produced without the benefit of the structural and label slippage) to those which are not yields a clear argument of the use of such slippage mechanisms in the interpretation of real metaphors and analogies.

If we view **Sports_car** as a **Jaguar**
Then **car_bodywork** is a **jaguar_head**
and **petrol_cap** is a **jaguar_mouth**
and **petrol_tank** is a **stomach**
and **combustion** is **digestion**
and **petrol** is **blood**
and **tyre** is a **paw**
and **car_wheel** is a **jaguar_leg**
and **crankshaft** is a **heart_muscle**
and **engine** is a **heart**
and **fuel_circulation** is **blood_circulation**
and **engine_rev** is a **heart_beat**
and **main_fuel_line** is an **aorta**
and **glass_plate** is a **cornea**
and **headlight** is a **jaguar_eye**
and **car_seat** is a **pelvis**
and **chasis** is a **jaguar_body**
and **exhaust_pipe** is a **rectum**
and **catalytic_converter** is a **brain**
and **combustion_chamber** is a **bowel**
and <u>**vehicle** is a **creature**</u>
and <u>**car_paint** is a **cat_fur**</u>

Figure 4: Sapper interpretation of the analogy SportsCar as Jaguar, which requires slippage rules. Only those mappings which are underlined are generated without the use of slippage rules.

The slippage-modified Sapper algorithm (which we dub *'Slapper'*) simply combines the compositional warping rule of Figure 2 with the basic Sapper algorithm of Figure 1. At each stage of the development of a semantic pathway in the (T)arget domain (as constructed via spreading activation), Slapper applies the rule of Figure 2 to generate one or more additional *'warped'* paths, which are also extended by spreading activation until a bridge-point is reached. Though many of the pathways that reach a cross-domain bridge-point will be *virtual pathways*, in the sense that they do not correspond to an actual sub-structure of semantic memory, such pathways may nonetheless be seen as semantically-valid *summaries* of other, longer sub-structures of memory. In effect then, Slapper dynamically summarizes one domain while attempting to structurally align it with another, less structurally/semantically-enriched domain.

5. Conclusions

Knowledge is a variable commodity, differing in resolution and scope from one cognitive agent to another. Take for example the domains of car mechanics and surgery: one naturally expects a professional mechanic to possess more expertise that a surgeon with regards to car maintenance, and vice versa — we pray that our surgeon knows more about the workings of the human body than does a mechanic. Though it follows that experts will know more about their chosen domains than non-experts, this does not preclude a non-expert making analogies that involve the pet domains of others. For instance, a car mechanic may train a novice by comparing the inner-workings of a car to the human body (e.g., the engine is the heart; oil is blood; the sump is the bowel; etc.), while a surgeon may educate a freshman medical student using the reverse analogy. A good deal of metaphors and analogies will therefore juxtapose domains that are defined at differing levels of causal complexity: just as we expect our tutor mechanic to overlook specific domain vagaries in his surgical analogy (e.g., how to order a hard-to-obtain spare for a *1972 VW Beetle*), we expect our surgical lecturer to overlook equivalent details in the medical domain (e.g., the biochemical distinction between harmful and not-so-harmful variants of cholesterol). In accommodating two such ill-fitting domains, one will have to be simplified if it is to be isomorphically reconciled with the other.

In this simplification lies the problem of principled structural slippage for analogical alignment. The *'soft'* mechanism outlined in this paper works in conjunction with the path-based Sapper algorithm to introduce semantically-motivated warping into the structure of (S)ource knowledge structure to ensure an isomorphic fit with the intended (T)arget. The approach is semantically-valid inasmuch as impossible or nonsensical semantic structures are never created by the warping rule; rather, structures which essentially summarize the meaning of the Target domain are instead generated *on-the-fly*.

Ultimately, such techniques will be required whenever structural comparisons are sought between knowledge domains that have been defined at varying levels of detail, redundancy and resolution. Sapper's path-based approach to structure-mapping, in contrast to the tree-based approaches of SME and ACME ([2,5,9]), allow for these levels of detail to be dynamically equalized. We expect that with further research, such path-based techniques will find greater application in future knowledge-based system design.

In closing, we note that Prolog implementations of Sapper and its slippage variant Slapper, with the example knowledge domains used to test each, are available from the metaphor-related metaphor site: *http://www.compapp.dcu.ie/~tonyv/metaphor.html*.

References

[1] Veale, T. and A. Way. (1997). Gaijin: A Template-Driven Bootstrapping Approach to Example-Based Machine Translation, in *the Proceedings of NeMNLP'97, New Methods in Natural Language Processing,* Sofia, Bulgaria, September 1997.

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

[3] Hofstadter,D., and The Fluid Analogies Research Group (1995). *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought.* New York: Basic Books.

[4] Veale, T. and M. T. Keane. (1997). The Competence of Sub-Optimal Structure Mapping on 'Hard' Analogies, *submitted to IJCAI'97, the International Joint Conference on Artificial Intelligence.*

[5] Falkenhainer, B, K. D. Forbus, & D. Gentner. (1989). Structure-Mapping Engine. *Artificial Intelligence*, 41, pp 1-63.

[6] Veale, T. and M. T. Keane. (1994). Belief Modelling, Intentionality and Perlocution in Metaphor Comprehension, in *the Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*, Atlanta, Georgia. Hillsdale, NJ: Lawrence Erlbaum.

[7] Veale, T., Smyth, B., O'Donoghue,D. and M. T. Keane. (1996). Representational Myopia in Cognitive Mapping, *in the Proceedings. of the AAAI workshop on Source of the Power in Cognitive Theories*, Portland, 1996.

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

[9] Oblinger, D. and K. D. Forbus. (1990). Making SME Pragmatic and Greedy, in *the Proceedings of the Twelfth Annual Meeting of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.