

On the Proper Treatment of Noun-Noun Metaphor: A Critique of the Sapper Model

Ronald W. Ferguson

Department of Computer Science
ferguson@ils.nwu.edu

Kenneth D. Forbus

Department of Computer Science
forbus@ils.nwu.edu

Dedre Gentner

Department of Psychology
gentner@ils.nwu.edu

Northwestern University, Evanston, IL 60201, USA

This paper responds to the claim of Veale, O'Donoghue and Keane (1995) that SME (Falkenhainer, Forbus & Gentner, 1989; Forbus, Ferguson & Gentner, 1994) performs poorly on noun-noun comparisons, such as *A surgeon is like a butcher*. Veale *et al.* argue that such noun-noun comparisons involve "object-centered representations" that do not contain higher-order relations, and that Sapper's chaining of expressions that share entity arguments provides a better model than SME for such comparisons. They also literally claim that SME will require longer than the lifetime of the universe to align Sapper's noun representations. We refute both their cognitive and computational claims here, focusing on (1) the psychology of noun-noun metaphors; (2) how Sapper works; and (3) SME's actual performance.

The psychology of noun-noun metaphors

Veale *et al.*'s claim that noun-noun metaphors require object-centered representations is inconsistent with psychological findings. People readily interpret noun-noun metaphors in terms of common relational structure; for example, "Cigarettes are like time bombs." is interpreted to mean that both cause damage after a quiescent period. Gentner and Clement (1988) found that adult ratings of metaphor aptness are positively correlated with the amount of relational information in their interpretations (as independently judged) and are non-correlated or negatively correlated with the amount of object-attribute information. Like humans, SME can produce both relational and attributional matches, but generally prefers the former.

How Sapper works

Veale *et al.* claim that Sapper's mapping ability comes from laterally chaining relational links, which exploit "complex causal structure" via "sideways systematicity." Our analysis of published mappings, however, instead suggests that Sapper depends heavily on implicit category and attribute sets to determine what to map. Sapper does this because its relations are too widely-scoped to map reliably via identity alone. For example, the relation *part* is used to relate *head* to *patient*, *sword* to *cavalry-charge*, and *security* to *bank*. To distinguish between instances of this and other ambiguous relations, Sapper implicitly treats some relational expressions as attributes based on their direction (see Figure 1). Because Sapper's triangulation rule matches objects linked to but not from a common node, an expression whose second argument is a terminal node (such as *head* and *security* but not *sword*) acts like a unary attribute over its first argument, since the second must self-

match. Accordingly, *head* acts as an attribute of *patient* which aligns *patient* with other "headed" entities, while *security* aligns *bank* with other secure entities. That terminal nodes must self-match strongly constrains alignment, since such nodes constitute the majority (57%) of the objects in Sapper's representations. Further evidence is found in the 92% of Sapper object matches based at least partially on shared categories (which are almost exclusively represented by terminal nodes). In short, categories and attributes deeply influence Sapper.

This dependence on categories and attributes need not rule out alignment of long lateral relational chains (the theoretical basis of Sapper). However, in practice only 3% of Sapper's mappings align objects more than two links away from the initial base and target nodes. In fact, 59% only align entities *directly linked* to the base and target nodes. Sapper's structural mapping ability, it seems, is severely limited.

SME's performance on Sapper's representations

In the interest of computational clarity, we ran SME on Sapper's noun representations, despite our reservations about their psychological plausibility. We rewrote terminal nodes as SME attributes, their closest functional equivalent. For Sapper's professions representations, SME produced mappings in an average of 17.5 seconds on a Pentium-133¹. SME's mappings are as reasonable as Sapper's, and on average overlap 72% of Sapper's entity matches.

Other limitations of Sapper

We note that Sapper contains other cognitive implausibilities. It assumes that all potential match hypotheses are created between all comparable objects. It provides no mechanism for generating structural inferences from base to target. Finally, Sapper models all metaphors as permanent local bridges between entities, limiting context sensitivity.

Acknowledgments

This work is supported by a grant from the Office of Naval Research.

References

- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The Structure-Mapping Engine. *Artificial Intelligence*, 41, 1-63.
- Forbus, K. D., Ferguson, R. W., & Gentner, D. (1994). Incremental structure mapping. *16th Annual Conference of the Cognitive Science Society*. Atlanta, GA: Erlbaum.
- Gentner, D., & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation*, 307-358. New York, NY: Academic Press.
- Veale, T., O'Donoghue, D., & Keane, M. T. (1995). *Epistemological pitfalls in metaphor comprehension*. Paper presented at International Cognitive Linguistics Association Conference, Albuquerque, NM.




Figure 1: How Sapper terminal nodes act like attributes

¹ Veale *et al.* apparently predict the universe has only seconds left.