Unweaving the Analogical Rainbow with Lightweight Lexical Ontologies

1. Introduction

Analogical reasoning is a decidedly knowledge-hungry faculty, whether one is interpreting or generating new analogies. Analogy is, after all, one of the foremost cognitive tools we possess for shedding light on a poorly understood domain by importing the structure of one more clearly understood. The approach to analogical reasoning most familiar to AI researchers is undoubtedly the structure-mapping approach, first suggested by Patrick Winston and Dedre Gentner and given an algorithmic realization in the Structure-Mapping Engine (or SME) of Falkenhainer, Forbus and Gentner (1989). SME, and other models that operate on similar principles (such as ACME by Paul Thagard and Keith Holyoak, IAM by Mark Keane, LISA by Keith Holyoak and John Hummel, or Sapper by Tony Veale; see Veale, 1996 for a review), presupposes that analogy operates by systematically projecting the causal propositional structure of one domain onto another, so in effect, structure-mapping can be viewed more mathematically as a cognitive variant of the well-understood, but NP-complete, problem of finding the largest isomorphic sub-graph of two representations (see Veale and Keane, 1997 for a proof).

As the name suggests, structure-mapping is vexingly dependent on the availability of explicitly structured domain descriptions to operate effectively. For those researchers like myself whose first exposure to analogy was via structure-mapping, the 1990's was a heady time in which competing models of analogical mapping were pitted against each other on specially-crafted domain descriptions of Aesopian fables and Shakespearean plots. Indeed, so spirited were this competition that Paul Thagard has aptly referred to the whole enterprise as the "Analogy Wars". However, the need for rich domain descriptions, mostly in first-order logical form, meant that analogical research in this period relied for the most part on hand-coded representations.

For the past three years, my group and I have been attempting to implement robust and scalable models of analogy, both interpretative and generative, that rely instead on large-scale representations from third-party sources. This has lead us to consider a number of possible knowledge-sources, from Cyc (Lenat and Guha, 1990) to WordNet (Miller, 1995). Indeed, the quest for large-scale structured resources that were independent of their analogical uses lead me in 1999 to Cycorp inc. of Austin, Texas, where I spent a year applying structure-matching ideas to the propositions and axioms stored in the Cyc knowledge-base. Unfortunately, the results were not encouraging; while Cyc contains many nested predications of a causal nature, there appears to be far too much structural variation between the descriptions of different domains (usually ontologized by different engineers) to make Cyc a viable source of representations for structure-mapping.

Cyc is a heavy-duty ontology, with a rich upper-model ontology and extensive cross-linking between concepts. But given this, it still seems inadequate for structure-mapping purposes. So if we have to forego the structure-rich approach to analogy that is structure-mapping, we may as well employ resources that are not themselves structure-rich. This has lead my group to look instead to freely available, if flawed, light-weight ontologies like Princeton WordNet (PWN; see Miller, 1995).

We have found PWN to be a sufficiently rich basis for modelling lexical analogies, such as those found on Scholastic Aptitude Tests. For example, *doubloon is to coin as what is to ship?* (answer: Galleon, since a Doubloon is a *Spanish* coin and a Galleon is a *Spanish* ship). PWN and similar resources can be used to understand and generate analogies like these by unlocking the implicit references contained in the textual glosses that annotate each conceptual entry. Certain gloss terms will be shared in common between lexical analogues (e.g., *Spanish* in both Galleon and Coin) while others will be domain-shifted (e.g., *coin* to *ship* in the above analogy, or *spacecraft* to *airplane* in an analogy between Astronaut and Pilot), and recognizing which terms serve which function is the essence of lexical analogy. The best lexical analogies involve a coordination of overt similarity with constrained difference, the latter often occurring within the same semantic field.

For example, the following table summarizes the analogical mappings that can be generated with PWN in the taxonomic domain of deities from different cultures.

Difference → *Commonality* ↓	*Greek*	*Roman*	*Hindu*	*Norse*	*Celtic*
supreme	Zeus	Jove	Varuna	Odin	N/A
wisdom	Athena	Minerva	Ganesh	N/A	Brigit
beauty, love	Aphrodite	Venus	Kama	Freyja	Arianrhod
sea	Poseidon	Neptune	N/A	N/A	Ler
fertility	Dionysus	Ops	N/A	Freyr	Brigit
queen	Hera	Juno	Aditi	Hela	Ana
war	Ares	Mars	Skanda	Tyr	Morrigan
hearth	Hestia	Vesta	Agni	N/A	Brigit
moon	Artemis	Diana	Aditi	N/A	N/A
sun	Apollo	Apollo *	Rahu	N/A	Lug

Analogy is a powerful retrieval tool that allows users of PWN to locate concepts not just via synonymy or taxonomy, but through complex allusions to other concepts (e.g., "Muslim bible" can be used to retrieve Qu'ran). As such, our work centres on the re-invention of the humble thesaurus as a creative resource in itself, an analogical thesaurus capable of understanding a user's allusions and even generating creative allusions of its own. This expressive power also fuels our current uses of lexical analogy as basis of puzzles and riddles in both computer games and scholastic tests.

References

1. Winston, P. H. Learning and Reasoning by Analogy. Communications of the Association for Computing Machinery, 23(12). 1980.
2. Gentner, D. *Structure-mapping: A theoretical framework for analogy*. Cognitive Science, 7, pp 155-170. 1983.
3. Falkenhainer, B., Forbus, K. and Gentner, D. Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, pages 1-63. 1989.
4. Veale, T. Metaphor and Meaning. Unpublished PhD thesis, available online at: *www.compapp.dcu.ie/~tonyv*. 1996.
5. Veale, T. and Keane, M. T. The Competence of Sub-Optimal Structure Mapping on Hard Analogies. *The proceedings of IJCAI'97, the International Joint Conference on Artificial Intelligence*, Nagoya, Japan. Morgan Kaufmann. 1997.
6. Lenat, D. and Guha, R. V. Building Large Knowledge-based systems. Addison Wesley. 1990.

7. Miller, G. A. WordNet: A Lexical Database for English. Communications of the ACM, Vol. 38 No. 11. 1995.