Towards a Computational Model of Analogical Arguments

Amy Bohan & Mark T. Keane*

ABSTRACT. In everyday disputes, and especially political disputes, people often use analogical arguments to support their views. For example, many of the arguments about the War in Iraq were regularly bolstered by analogical comparisons to WWII and Vietnam. Though logic and philosophy has always viewed analogical argumentation as suspect, analogical thinking has recently been studied and modelled extensively in cognitive science (see e.g., [1], [2], [3], [4], [5]). In this paper, we assess whether computational models of analogy can be safely applied to analogical argumentation. For various reasons, we conclude that they cannot because they are too powerful.

1 INTRODUCTION

In everyday disputes, and especially political disputes, people often use analogical arguments to support their views. In Logic and Philosophy argumentation by analogy has often languished in darkness of sophistry rather than in the shining light of formally correct reasoning. However, recently, Cognitive Science has somewhat restored the reputation of analogical thinking by demonstrating its centrality in human thinking [see e.g., 3, 6].

Work like Gentner's Structure-Mapping Theory (SMT)[6] characterises analogy as an isomorphic mapping between two domains of knowledge based on matching connected sets of relational predicates {e.g., *cause [(kick (John, ball), enter(ball, goal)]*} rather than attribute predicates [e.g., *round(ball-1)*] between a base/source domain and a target domain. For example, in the 'atom is like a solar system' analogy, the connected set of relations about mass differences causing revolution are mapped from the *base* solar system domain to the *target* atom domain. This preference for connected sets of matches is called the *systematicity principle* and has been confirmed by many psychological studies. SMT's distinction between relations and attributes also allows us to define a taxonomy of comparisons from deep analogies (many connected relations and few attributes, like the solar system and atom) to literally similar comparisons (sharing relations and attributes, like our solar system and the K9 solar system) to superficial comparisons (matching attributes only, e.g. the earth and a football are both round). As we shall see in section 3, there are common computational mechanisms that capture many of these comparison types [see 1, 5].

2 ANALOGIES AND ARGUMENTATION

Given the generality of SMT's formulation, one would expect it to be able to also characterise analogical arguments. Indeed, it is possible to describe a schema for analogical arguments using the theory's precepts (see Figure 1). For example, let's assume we are arguing about the factual proposition that "Prison causes recidivism" and someone says that "Prisons are universities for criminals". Unpacked, the essence of this argument is a mapping between relations like *attend(student, university) -> attend(prisoner, prison), sleep-in(students, dormitories) ->sleep-in(prisoners, cell-blocks)*. These basic mappings support predictions that are transferred from the analogical domain to the factual domain; for example, inferences like *learn-in(student, university) -> learn-in(prisoner, prison), prepare-for(university, student, later-life)->prepare-for(prison, prisoner, later-life)*. These analogical transfers then have established new facts in the target domain that support the original proposition. The learning environment of the prison supports further crime in later life.

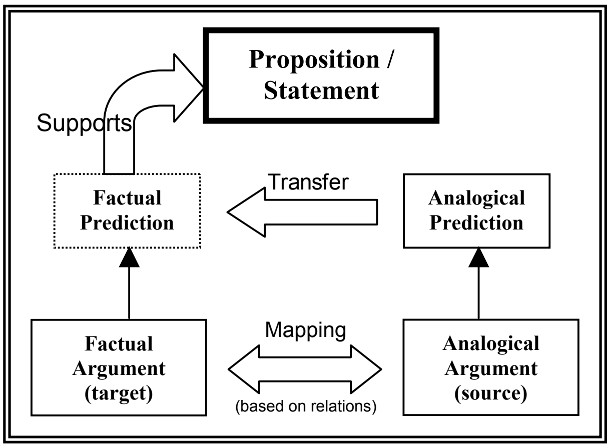

Figure 1. Schema of Analogical Arguments

On the face of it, following SMT, one would expect arguments to be better if the two domains have many mappings in common. However, it is probably more important that a suitable analogical transfer can be made, for the mappings found, to create a new factual prediction that clearly supports the original proposition (ideally, causal support). This schema presents a high-level theoretical analysis of a sort that has been implemented in several ways by various computational models, to which we now turn.

3 MODELS OF ANALOGY

Current computational models of analogy agree at the computational level about the things that need to be computed with drawing an analogy (e.g., matching relations, systematicity, etc). However, at the algorithmic level, they take on many different flavours achieving this computation in many different ways [5]. Here, we briefly review some of these different flavours.

3.1 SME

The Structure Mapping Engine [SME, 1], is based on Gentner's Structure Mapping Theory [6]. SME finds the largest inter-domain

* Department of Computer Science, University College Dublin, Belfield, Dublin 4, Ireland. E-mail: <amy.bohan><mark.keane>@ucd.ie

CMNA IV, 4th *Workshop on Computational Models of Natural Argument*
Edited by Floriana Grasso, Chris Reed and Giuseppe Carenini
ECAI 2004

mapping, using only the structure of the source and target domains. It basically finds relational mappings and any object mapping supported by these relations and then merges these root mappings into larger consistent groups. The group with the highest degree of connectiveness (i.e., systematicity) is the preferred interpretation for the analogy. Computationally, SME is essentially finding the least-common sub-graph of the two domains of connected predicates [see 1, 12]. SME performs a bottom-up search through the space of possible inter-domain mappings retaining the largest mapping found. Unmatched source items are transferred to the target domain forming the analogical inferences. This "carry over" strategy has since been formalised into the CWSG (Copy With Substitution and Generation Algorithm) algorithm [8].

3.2 ACME

The Analogical Constraint Mapping Engine [ACME, 9] models analogy according to SMT using a parallel constraint satisfaction network. It forms a localist network where each unit represents a possible mapping with activation/inhibition links between these units signifying, for example, if one mapping supports another or if one mapping violates the isomorphism of the mapping between the two domains. After the network of nodes is built activation is passed between the nodes until it settles into a stable state, then the mapping nodes with the highest activations above a given threshold are read off, reflecting the best interpretation for the analogy. Transfers are achieved by a separate CWSG mechanism. ACME works very efficiently once the network is built to find the optimal mapping, but the building of the network can be very computationally expensive, as its forms nodes for every predicate of the same parity in both domains [see 5, 11]. See LISA for a newer and better neural network approach to analogy by this group [10].

3.3 IAM

The Incremental Analogy Matching [IAM, 5] performs serial constraint satisfaction after selecting a subset of the base-domain's predicates from which to build the analogy incrementally. It performs a top-down heuristic search to find the best interpretation for the analogy making it more efficient than SME and ACME. IAM correctly predicts human performance very closely in analogy problems and that the order of information in a domain description influences solution speed and quality. More recent versions of SME and the LISA model have also adopted some form of incrementally.

3.4 Sapper

Finally, Sapper [11, 12] uses a localist semantic network with spreading activation to model analogy in a wholly different way. Sapper views semantic memory as a localist-graph in which nodes represent distinct concepts and arcs between those nodes represent semantic relations between concepts. Mappings between concept nodes are represented by putative bridging links between these concepts, along with activation may/may not flow. Sapper builds these bridges in its knowledge base in advance of the analogical episode, when two concepts are analogically related in a statement, activation is initiated from these concept nodes and flows to related concepts (and across bridges) within a fixed link horizon. Sapper has been shown to be, perhaps, the most computationally efficient analogical mapper [see 11].

3.5 Conclusion on Models

All of these programs can instantiate the schema for analogical argument described in Figure 1 and as such are candidates for a computational treatment of argumentation by analogy. The choice of one model over another could be driven by efficient considerations but there is a bigger issue to be resolved before any choice can be made. A crucial issue is whether the argument schema is actually accurate as a characterisation of analogical arguments or, more subtly, whether analogical arguments of the form outlined in the schema are more convincing than their factual equivalents. Fortunately, we have carried out some recent empirical work on this issue and can thus shed some light on these questions.

4 ARE PEOPLE ANALOGICALLY CONVINCED?

We have carried out several experiments designed to see whether analogical arguments based on the proposed schema are arguments that people find more convincing than their factual equivalents. In these experiments people were either given analogical arguments or factual arguments for various topical propositions (numbering 10 in total). The propositions dealt with topics like alcohol abuse, military service, the Iraq war and so on. For each proposition we created two supporting factual arguments and two corresponding analogical arguments. For example, for the proposition 'The War on Iraq was justified' the two factual arguments were that "Saddam was a dictator" and "Saddam committed genocide in his country". The corresponding analogical arguments were that Saddam was like Hitler, in that "Hitler was a dictator" and "Hitler committed genocide in his country".

Participants in the study were asked to rate their belief in the proposition (independent of the arguments made) and also rate the goodness/convincingness of the factual/analogical arguments. These belief-bias and goodness-ratings tasks were counter-balanced. When analysed, they showed that people's agreement/disagreement with the proposition did not predict how convincing they found the arguments to be. In other words, these undergraduate participants could consistently separate their assessment of the argument from their *a priori* belief in the proposition.

Surprisingly, we found that people found the factual arguments to be reliably more convincing than the analogical arguments, suggesting that generations of sophist politicians have been wasting good comparisons on their voting audiences. Furthermore, this result was replicated in several successive studies showing that this is a robust finding. However, several manipulations showed that part of the source of this effect lie in the amount of cognitive processing people carried out on the analogy. We found that when people were encouraged to spend more time considering the analogy (e.g., to overtly draw out the comparisons, by being asked to explicitly indicate the matching objects in both domains) they started to find the analogical arguments slightly more convincing.

Table 1 shows the results of one experiment in which there were four different groups of arguments: analogy-alone, analogy+mapping, analogy+mapping+factual and analogy+mapping+irrelevant. In the analogy-alone condition, people rated both factual and analogical arguments, where the analogical arguments were simply stated (as described above about Hitler and Saddam). In the analogy+mapping, the analogical argument was stated and people were asked to affirm the mappings established between the domain objects (e.g., that Saddam corresponded to Hitler etc). In the analogy+mapping+factual condition, the factual equivalent of the analogical argument was given in conjunction with the analogy and mappings. Finally, in the analogy+mapping+irrelevant, people were given the analogical arguments and mappings but also some additional factual material (e.g., other facts about Hitler). This latter condition was designed to control for any effects that might be due to just being given additional information (irrespective of its content) for the final two conditions. As can be seen from Table 1, the convincingness ratings for the factual arguments remain high and constant around 4 and are not reliably different. In contrast, the convincingness of the analogical arguments gets progressively higher when people are given more information and are asked to consider the mappings in more detail. Note, that in the control condition, the analogy+mapping+irrelevant condition, the ratings drop back when people are given irrelevant additional information instead of the factual arguments.

Table 1. Results of Convincingness Experiment

Condition	Analogy	Factual
analogy-alone	3.225	4.1875
analogy+mapping	3.985294	4.371429
analogy+mapping+factual	**4.0875**	**3.9375**
analogy+mapping+irrelevant	3.4125	4.05
Average:	**3.677574**	**4.136607**

These results present some interesting problems for current models of analogy. They suggest that it is not enough to lead people to the analogical water, for them to drink from the font of conviction. In the final section we consider some of these implications.

5 IMPLICATIONS FOR MODELS OF ARGUMENTATION

The result of the studies of analogical arguments that we have outlined here, present several problems for current computational models. First, it is clear that people can understand analogical arguments though, as we have seen, they do not find them especially more convincing than a straight factual argument. Second, when they are given an analogical argument they appear to be more influenced by the amount of consideration they give to the analogy, than the content of the argument *per se*. The more time they put in to thinking about the analogy the more they find it convincing. The fundamental problem that these findings present for current analogical models is that they do not easily model any of these effects. All of the models would have no difficulty performing the analogical mappings on which the analogical arguments are based. Furthermore, they would not find any of these analogies any better or more convincing by virtue of the way they processed them. It is clear that any computational consideration of analogical arguments will require quite different sorts of models than those that currently exist; though it could be said that some properties of Sapper [11, 12] suggest that it might be more amenable to handling these effects.

Finally, another very different conclusion could be drawn from these findings; namely, that the schema derived from SMT to characterise analogical arguments is inappropriate to the analogical arguments actually used by people. It may well be, that the schema is too complex for people to find convincing, it may require too much processing to be considered good as an analogical argument. We are currently pursuing this possibility by looking at much simpler analogical arguments; e.g., Astronauts will die in the attempt to reach Mars because many explorers died in the attempt to reach the New World. Such arguments would, of course, require much less computational baggage than most current models offer.

REFERENCES

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

[2] Holyoak, K.J & Thagard, P. (1989): “Analogical Mapping by Constraint Satisfaction”, Cognitive Science, 13, 295-355.

[3] Keane, M. T (1988). Analogical Problem Solving. New Yord: Simon & Schuster.

[4] Keane, M.T. (1997). What makes an analogy difficult?: The effects of order and causal structure in analogical mapping. *Journal of Experimental Psychology: Language, Memory & Cognition, 23,* 946-967.

[5] Keane, M. T, Ledgeway, T, & Duff, S. (1994): “Constraints on Analogical Mapping: A Comparison of Three Models”. Cognitive Science, 18, 387-438.

[6] Gentner, D (1983): “Structure-Mapping: A Theoretical Framework for Analogy”. Cognitive Science 7, 155-170.

[7] Kuhn, D (1991): “The Skills of Argument”, Cambridge, Cambridge University Press.

[8] Holyoak K. J. Novick L. Melz E. (1994): “Component Processes in Analogical Transfer: Mapping, Pattern completion and Adaptation”, in Analogy, Metaphor and Reminding, Eds. Barnden and Holyoak, Ablex, Norwood, NJ.

[9] Holyoak, K.J. Thagard, P. (1989): “Analogical Mapping by Constraint Satisfaction”, Cognitive Science, 13, 295-355,

[10] Hummel, J.E. & Holyoak, K. J. (1997). “Distributed Representations of Structure: A Theory of Analogical Access and Mapping”. Psychological Review, 104, 3, 427-466.

[11] Veale, T. & Keane, M.T. (1994). Belief modelling, intentionality, and perlocution in metaphor comprehension. In A. Ram & K. Eiselt (Eds.), *Sixteenth Annual Conference of the Cognitive Science Society.* Hillsdale, NJ: Erlbaum.

[12] Veale, T. & Keane, M.T. (1997). The competence of sub-optimal theories of structure mapping on hard analogies. *Fifteenth International Joint Conference on Artificial Intelligence.* Los Altos: Morgan Kaufmann.