Proceedings of the Annual Meeting of the Cognitive Science Society

Analogical Transfer by Constraint Satisfaction

Proceedings of the Annual Meeting of the Cognitive Science Society, 13(0)

Melz, Eric R.
Holyoak, Keith J.

1991

Analogical Transfer by Constraint Satisfaction

Eric R. Melz **Keith J. Holyoak**
Department of Psychology
University of California, Los Angeles, CA 90024

Abstract*

The robustness of analogical transfer based on the ACME modeling of mapping by constraint satisfaction (Holyoak & Thagard, 1989) was investigated in a series of computational experiments using Hinton's (1986) "family tree" problem. Propositions were deleted randomly from the full representations of either both analogs (descriptions of an English and an Italian family) or just the target, and after mapping a "copy with substitutions" procedure was used to generate transfer propositions intended to restore the full representational structures. If as many as 50% of the propositions in the target analog were deleted, the system was able to recreate all of the missing information without error; significant recovery was obtained even if as many as 80% of the target propositions were deleted. Robustness was only slightly reduced when the two analogs lacked any similar predicates, so that mapping depended solely on structural constraints. Transfer was much more impaired when deletions were made from both analogs, rather than just the target. The results indicate that for isomorphic representations, analogical transfer by constraint satisfaction can exceed the regenerative capacity of general learning algorithms, such as back-propagation.

Introduction

Theoretical accounts of analogical transfer distinguish various component processes, of which *mapping* is considered central. Mapping entails establishing systematic correspondences between the objects and relations of a source analog and a target analog. But in addition to mapping, analogical transfer depends on post-mapping processes in which new inferences are drawn using the correspondences established by the mapping, in conjunction with the structure of the analogs. Such procedures can effectively fill in gaps of missing information when knowledge from a well-understood domain is transferred to an analogous but less well-understood domain. By analogy to perceptual mechanisms for "filling in" missing information, we will refer to the generation of inferences based on analogical mapping as *analogical pattern completion*. The simplest basic mechanism for analogical pattern completion is *copying with substitution* (CWS), which in some form has been included in all computational models of post-mapping analogical transfer (e.g., Carbonell, 1983; Falkenhainer, Forbus & Gentner, 1989; Winston, 1980). The basic idea of CWS is simple: for a known proposition about the source, construct an inference about the target by substituting the corresponding predicate and argument(s) in the target for those in the source proposition.

Although several models of analogical transfer that include the CWS principle have been proposed, no systematic tests of robustness have been reported for models of either mapping or post-mapping transfer processes. Mapping models have typically been applied to representations that are nearly isomorphic; it is unclear how well the systems could map less orderly representations. Naturalistic use of analogy typically involves situations in which at least one of the analogs -- the novel target -- is imperfectly understood. A model of human analogical transfer must be sufficently robust as to be able to identify systematic correspondences between analogs despite gaps in the initial representations, and then proceed to generate plausible inferences to fill those gaps. In this paper we examine the inferential power of a CWS mechanism as an extension of ACME, an model of analogical mapping by constraint satisfaction (Holyoak & Thagard, 1989).

Learning Family Trees by Back-Propagation

Hinton (1986) describes a back-propagation network which he trained on propositional representations of family trees. The primary purpose of Hinton's study was to determine if the hidden units of the network could develop intuitively meaningful representations of abstract features of a corpus of propositions. The basic family trees that the network learned, also used in the present study, are depicted in Figure 1. As is visually apparent, these English and Italian families have an isomorphic structure (e.g., Christopher enters into the same pattern of kinship relations as does Roberto). Using 12 common relational terms (father, mother, husband, wife, son, daughter, brother, sister, uncle, aunt, nephew and niece), each family can be described by a set of 56 propositions about relationships among the 12 individuals.

In Hinton's project, propositions of the form (person1 relation person2) (e.g., (Emilio has_father Roberto)) were translated into a localist connectionist representation and presented to a multi-layer network. Using back-propagation, the hidden units were able to abstract useful features, such as nationality and generation, from the training set. When tested on its ability to complete the 4% of the propositions that had not been used in training, the network was correct on 100% of these in one run and 75% in a second run.

---

* This research was supported by Contract MDA 903-89-K-0179 from the Army Research Institute, NSF Biological Facilities Award BBS 87-1420, and a grant from the Keck Foundation to the UCLA Cognitive Science Research Program.

Figure 1: Two isomorphic family trees. The symbol "=" means "married to". (From Hinton, 1986)

Analogical Mapping and Pattern Completion in ACME

Mapping

Holyoak and Thagard (1989) frame the problem of analogical mapping in terms of parallel satisfaction of multiple constraints that jointly determine the optimal correspondences between elements of the source and target analogs. Their ACME (Analogical Constraint Mapping Engine) model receives as input a source analog and a target analog represented in predicate-calculus notation. Each analog consists of a set of propositions, where each proposition consists of an n-place predicate, a list of constituent arguments of the predicate, and a proposition label. For example, the fact that Emilio's father is Roberto would be represented by the proposition

(has_father (Emilio Roberto) I1).

The full representations of each family would consist of 56 propositions, each involving a 2-place relation.

ACME constructs a network of units representing potential elemental correspondences between the two analogs. The constraint that good analogies should tend to be isomorphic is enforced by installing excitatory links between consistent mapping hypotheses and inhibitory links between inconsistent hypotheses. Units representing mappings of semantically similar predicates are linked to a "semantic unit", the activation of which is clamped to a maximum activation value of 1. In the experiments reported in this paper, only mappings of identical predicates have connections to the semantic unit. Units representing correspondences presumed on the basis of prior knowledge are linked to a special "pragmatic unit". The network is then allowed to settle using a relaxation algorithm (see Holyoak & Thagard, 1989). The updating algorithm is implemented in *LISP on a CM2 Connection Machine.

Pattern Completion

Once a set of mappings for objects, predicates, and propositions have been obtained by relaxing the system, we invoke a simple CWS procedure to generate candidate inferences based on the mappings and the structure of the analogs. The intuition behind this procedure is that in analogical transfer, there is a pressure for propositions in the source to have corresponding propositions in the target, and vice versa. If some proposition exists in one analog, but has no corresponding proposition in the the other, yet all of the constituent elements of the existing proposition map to elements in the other analog, we may reasonably conjecture a new proposition. Formally, the procedure we use to generate propositions in the aftermath of mapping is as follows:

- *If a proposition P consisting of relation r and objects* $o_1, o_2, \ldots, o_n$ *(notated* $P{:}r(o_1, o_2, \ldots, o_n)$*) exists but does not map to any corresponding proposition, and*
- *if P's relation and objects have the mappings*
  $r \text{->} r'$
  $o_i \text{->} o_i'$ *(for* $i=1 \ldots n$*)*
- *then create the new proposition* $P'{:}r'(o_1', o_2', \ldots, o_n')$.

The criteria for generateinga proposition are that (a) the best mapping of P must have an activation below some threshold value, and (b) the relation and objects of P must have activations above the threshold. For all the experiments reported in this paper, this threshold was chosen to be .2. Note that the above procedure for inference generation is inherently symmetrical: the new proposition can be added to either the source or the target analog.

Transfer Tests

The family-tree problem has several virtues as the basis for computational tests of an analogical transfer model. First, the two complete family structures are in fact isomorphic, so analogical mapping should be possible. Second, the full representations are a well-specified set of propositions, so we can quantify the degree to which analogs have been corrupted by eliminating propositions from the input representations. By deleting propositions from the inputs, we can systematically reduce the degree to which the inputs are actually isomorphic, and hence examine the robustness of the mapping and pattern completion mechanisms.

Mapping the Intact Family Trees

The first requirement was to demonstrate that ACME could in fact map the two analogs if the complete representations (i.e., 56 propositions for each family) were provided as inputs. This is a non-trivial computational problem simply because the mapping network formed is very large (3424 mapping units interconnected by 381224 symmetrical links), due to the fact that all propositions involve 2-place relations (so that any proposition in one analog could potentially map to any proposition in the other). The major parameter values used were .005 for decay, excitation, and similarity of identical predicates, and -.16 for inhibition. The network settled into a stable asymptotic state after 196 cycles of activation updating, producing a complete and correct set of correspondences between elements of the two structures. All the correct mappings had asymptotic activations close to the maximum possible value of 1.

Reconstruction of Damaged Analogs With Identical Predicates

Having established that ACME can map the intact analogs, we next performed a series of computational experiments in which we randomly deleted propositions from the family-tree representations provided as inputs to ACME, mapped the damaged representations, and then applied the CWS pattern-completion mechanism to attempt to reconstruct the complete analogs. The results of a series of experiments on analog reconstruction are presented in Figures 2 and 4.

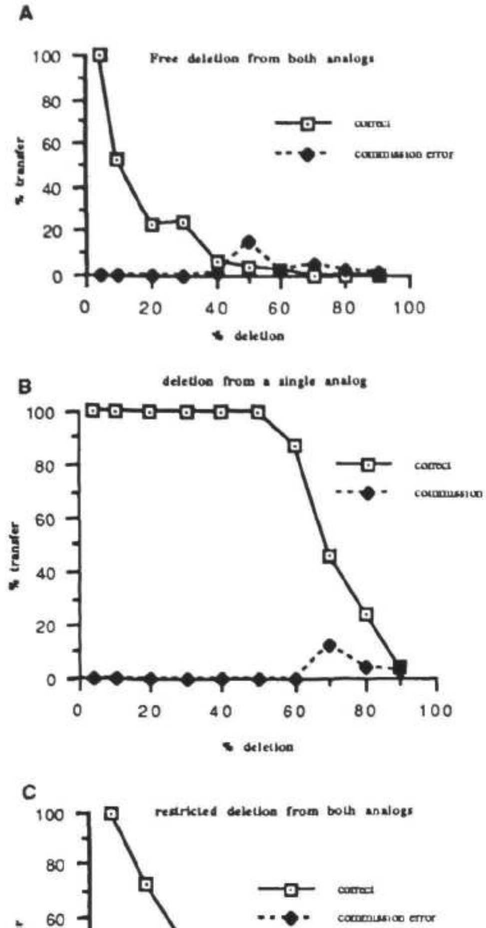

Figure 2: Robustness of transfer (identical predicates)

In our first experiment, we randomly deleted propositions from both analogs, and observed the proportion of deleted propositions that were correctly or incorrectly created by the CWS mechanism. The results of this experiment are shown in Figure 2A. The abscissa represents the proportion of propositions that were deleted from the entire set of propositions in the two analogs. The ordinate represents the proportion of the deleted propositions that were either created correctly (labeled "correct"), as well as the frequency with which incorrect propositions (labeled "commission error") were generated, also expressed as a percentage on the deletions in the inputs. Each data point on this and subsequent graphs represents the average of the results of two runs. ACME was able to reconstruct 100% of the missing propositions when 4% had been deleted, and 53% of those missing when 10% had been deleted, without making any commission errors. Correct restorations diminished to 23% when the deletion rate was increased to 20-30%, and at higher levels of deletion correct inferences were essentially eliminated. Commission errors were very infrequent even at the highest levels of damage to the analogs.

The above experiment involved symmetrical damage to the two analogs, with bidirectional transfer between the two analogs. In contrast, naturalistic analogical transfer typically involves asymmetric transfer from a well-understood source to a poorly-understood target. To more closely approximate the naturalistic asymmetry of analogical transfer, we ran a second experiment in which we restricted proposition deletions to only a single analog (in this case, it happened to be the Italian family). As the results in Figure 2B clearly indicate, transfer was far better than in the previous experiment. Full recovery of deleted propositions was possible at deletion rates of up to 50%, and even at a deletion rate of 80% ACME was able to recover 24% of the missing propositions.

The difference in robustness between the two deletion procedures is extreme indeed. For example, deleting 60% of the propositions in one analog produces the same quantity of missing information as does deleting 30% of the propositions across both analogs. Yet the former procedure yields almost perfect recovery (Figure 2B), whereas the latter procedure allows recovery of less than 1/4 of the missing information (Figure 2A). We next explored potential structural explanations for this difference in robustness as a function of whether deletions were made from one or two analogs.

One possible explanation is that when random deletions occur across both analogs, it is possible that corresponding propositions (e.g., E1 and I1) may both be deleted, in which case the CWS mechanism is guaranteed to fail (because there will be no proposition from which to generate an analogous inference). (The proportion of deleted propositions for which we can expect such an event to occur is $1 - (1 - p)^d$, where p is the proportion of the total propositions that are deleted, and d is the number of propositions that are deleted.) In contrast, if one analog is left intact, it is guaranteed that one member of each proposition pair is available (namely, the proposition in the source). To test the effect of this structural advantage for the latter procedure, we introduced a third deletion scheme that allowed propositions to be deleted from both analogs, with the restriction that at most one proposition from each pair of corresponding propositions could be deleted. The results, shown in Figure 2C, indicate that although this procedure produces somewhat more robust transfer than does free deletion from both analogs (Figure 2A), it remains much worse than when deletions are performed from only a single analog (Figure 2B). These results indicate that some other structural factor must account for the greater robustness of transfer when deletions are restricted to a single analog.

Another structural factor that varies when deletions are made from one versus two analogs involved differences in the potential for generating incorrect proposition mappings, as illustrated in Figure 3. In the case where deletion is restricted to a single structure (Figure 3a), for any proposition such as I5 in the target structure that is deleted, the strongest mapping of its corresponding source proposition, such as *E5=I10*, will be inhibited by the correct mapping unit, here *E10=I10*. Since *E10=I10* represents the mapping of corresponding propositions, it will have a high activation, and drive the mapping *E5=I10* well below the threshold required for generating transfer candidates. Hence, proposition E5 in the source will be left unmapped, making it a candidate for generation of an inference about the target by CWS.

In contrast, when proposition deletions occur in both analogs (Figure 3b), incorrect proposition mappings do not necessarily experience the devastating inhibition described above. For example, consider the case in which proposition E10 is deleted from the source and I5 is deleted from the target. If E10 and I5 share a common relation or object, they may produce a reasonably strong mapping. Since neither *E10=I10* nor *E5=I5* exists, the mapping *E5=I10* faces no serious competition, and thus is able to produce an activation level above the transfer threshold, preventing subsequent transfer from E5. The greater potential for erroneous proposition mappings to emerge when deletions are made from two analogs rather than one appears to be the main reason for the reduced robustness of transfer for the former case.

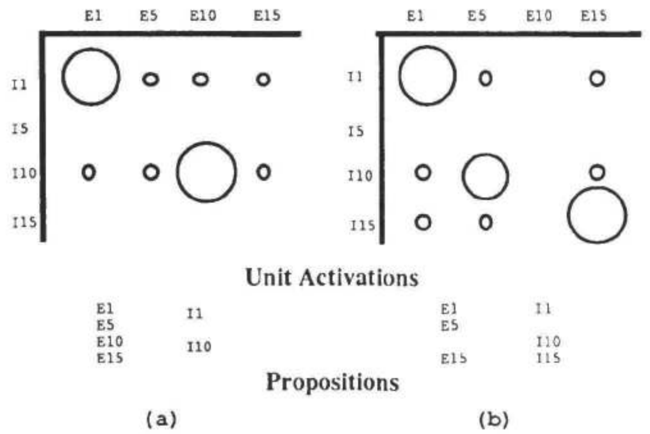

Figure 3: The units created for (a) single-analog deletion and (b) two-analog deletion are represented by circles in the Unit Activation tables. Activation levels of the proposition mapping units *column=row* are indicated by the size of the circles. Mapping units on the diagonal, which represent correct mappings, strongly inhibit the units in their row and column. Hence in (a), the unit E5=I10 is strongly inhibited by E10=I10, so that E5 is left unmapped, and hence a candidate for generating an inference. In contrast, in (b) the unit E5=I10 does not encounter any such inhibition, and hence its activation is above the threshold required for inference candidacy.

Reconstruction of Damaged Analogs Without Identical Predicates

The previous three experiments have shown that missing information can be reconstructed under a variety of deletion conditions. Such transfer is critically dependent on the quality of mappings ACME produces. As previously discussed, the mapping process is influenced by both structural and semantic pressures. Since in this problem we are mapping two isomorphic structures, the semantic pressure generated by the identical predicates in the two analogs (e.g., *has_father=has_father*) might be regarded as superfluous, or as unduly "setting up" the mapping.

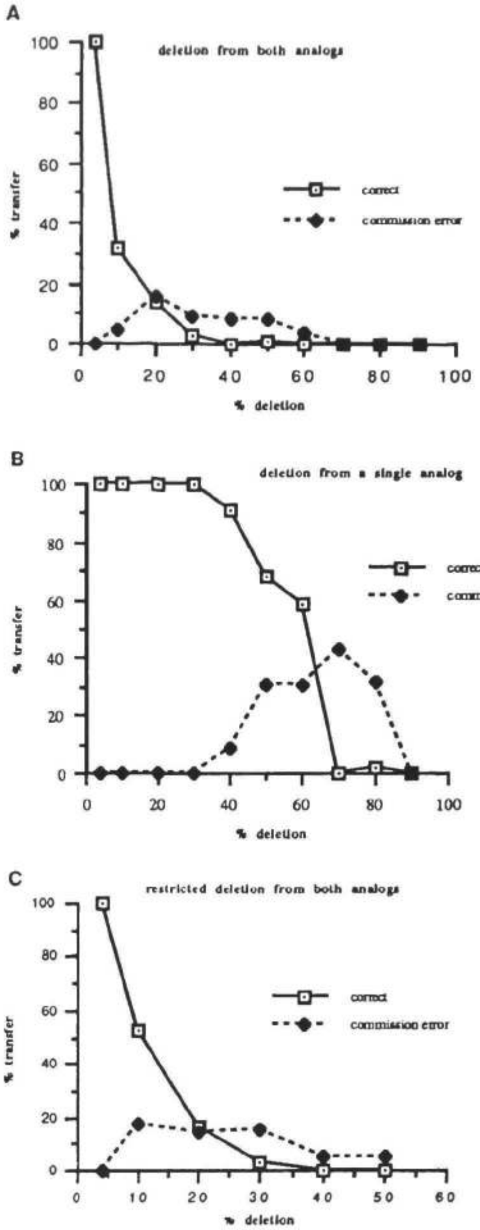

Figure 4: Robustness of transfer (no identical predicates).

In order to eliminate the semantic pressure, we translated all the relation names in the Italian family tree into their Italian equivalents. For example, the relation "has_father" became "ha_padre". Since ACME only creates a strong link between the semantic unit and a predicate mapping units if the mapped predicates have the same name, this modification prevents any useful semantic links from being created. In addition, since the English family tree can be consistently mapped to either the Italian family tree shown in Figure 1 or a vertical mirror image of the Italian family tree[1] (e.g., Christopher maps to Francesca, etc.), we weakly "presumed" one of the appropriate mappings (by giving it a weight of .005 on its link from the pragmatic unit) in order to provide a minimal impetus towards the "correct" set of mappings.

---

[1] We thank Ed Stabler for making this observation.

We next reran our three basic deletion experiments, eliminating semantic pressure. The results, displayed in Figure 4A-C, represent a qualitative replication of the previous patterns of transfer, with some degradation in the amount of transfer. This degradation is expected for two reasons. First, the overall activations of the proposition, relation and object mapping units are weaker when semantic pressures are removed. Second, clusters of mappings which are locally consistent but globally inconsistent can emerge under this scheme, especially when a significant amount of relational information has been removed from the representations.

Conclusion

We found that transfer was far more robust when deletions were restricted to a single analog than if deletions were made in both. This result leads to the prediction that transferring knowledge from a well-understood source to a poorly-understood target will be easier than transferring knowledge between two moderately-understood analogs. As far as we know, this prediction has not yet been directly tested for human analogical transfer.

It would be interesting to test the robustness of other models of analogical transfer using the same basic materials used here. It does not appear likely, for example, that comparable results could be obtained using Falkenhainer et al.'s (1989) SME program with their structure-mapping rules. ACME is able to perform well even when the analogs lack any similar or identical predicates, whereas SME with structure-mapping rules requires that mapped relations be identical.

It is of interest to compare the performance of ACME with Hinton's (1986) learning by back-propagation as applied to the family-tree problem. Our system appears better able to recover implicit missing information for the family tree problem; however, obvious and significant differences exist between the two systems. On one hand, ACME requires that the propositions about the two families be explicitly separated into target and source analogs, whereas Hinton's system received all propositions intermixed, and in fact *learned* that the distinction between English and Italian people was an important regularity. On the other hand, Hinton's generalization task involved giving the system the first argument and relation and asking it to generate the second argument, whereas ACME was asked to generate entire new propositions without any explicit partial cues. In addition to these differences in the transfer tasks performed by the two systems, the general aims of each system are quite different. Hinton's system is primarily intended to abstract general features from a body of propositions. ACME, on the other hand, has no such generalization capability, but rather conjectures the existence of unstated information based solely on structural correspondences between two sets of propositions.

The present version of ACME has a number of significant limitations as a model of the post-mapping processes involved in analogical transfer. The success of the CWS procedure depends upon the model's tacit assumption that the source and target situations, despite any apparent "gaps" in their representations, are in fact isomorphic. If the source includes propositions that actually lack parallels in the target situation, commission errors (i.e., erroneous inferences about the target) are likely to result. Successful transfer between non-isomorphic analogs will depend upon additional pragmatic information that either prevents erroneous inferences from being generated, or weeds them out after they are generated. In addition, post-mapping transfer processes must also deal with cases in which some useful source propositions are based on predicates or objects that are initially left unmapped (in contrast to CWS, which can apply to a proposition only if all its elements have been mapped).

Finally, a full model of analogical transfer between non-isomorphic situations will require processes of *adaptation* (Novick & Holyoak, 1991) that modify the inferences generated by analogical pattern completion in order to take account of unique requirements of the target. For example, an analogical solution to a target problem may have to be adapted to incorporate extra constraints required in the target but not the source, or to accommodate variations in which problem elements are specified as unknown by the goal description (Carbonell, 1983). We view the present extension of ACME as only a modest initial step toward accounting for the full range of transfer processes that depend upon, but potentially go far beyond, the information provided by analogical mapping.

Acknowledgements

We thank Ed Stabler for suggesting the idea for this project. Reprint requests may be directed to K. J. Holyoak, Dept. of Psychology, UCLA, Los Angeles 90024-1563.

References

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41*, 1-63.

Carbonell, J. G. (1983). Learning by analogy: Formulating and generalizing plans from past experience. In R. Michalski, J. G. Carbonell & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach.* Palo Alto, CA: Tioga Press.

Hinton, G. E. (1986). Learning distributed representations of concepts. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Novick, L. R., & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17*, 398-415.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM, 23*, 689-703.