

An Inhibitory Mechanism for Goal-Directed Analogical Mapping

Barbara A. Spellman

Keith J. Holyoak

Department of Psychology
University of California, Los Angeles, CA 90024-1563

Abstract*

Theories of analogical thinking have differed in the roles they ascribe to processing goals as a source of constraint on analogical mappings. We report an experiment that examines the impact of processing goals on subjects' mappings in (a) a task involving generation of plot extensions for soap opera scripts, and (b) an explicit mapping task based on characters in the scripts. The scripts were written so that the mappings for central characters were four-ways ambiguous. Manipulations of subjects' processing goals influenced their preferred mappings, both in the plot-extension and mapping tasks. In the latter task, goal-irrelevant information contributed to the resolution of mappings that were ambiguous on the basis of goal-relevant information alone. The qualitative pattern of results was successfully simulated using ACME, a constraint-satisfaction model of mapping, in which processing goals are assumed to control an inhibitory process of selective attention. Processing goals attenuate the activation level of goal-irrelevant information, reducing or even eliminating its impact on mapping decisions.

Introduction

A crucial requirement for purposeful thinking is ensuring that inferences are *relevant* to the goals of the reasoner (see Holyoak & Spellman, 1993, for a review of recent work on thinking). We suggest that maintaining relevance of inferences depends in part on goal-directed *inhibition* of irrelevant information that might otherwise misdirect the reasoning process. The impact of processing goals and the possible role of an inhibitory mechanism for selective attention can be investigated in the context of analogical mapping. Most computational models of analogical inference have postulated (although not always implemented) some mechanism by which the reasoner's external purpose in using the analogy can influence analogical transfer (e.g., Falkenhainer, Forbus & Gentner, 1989; Forbus & Oblinger, 1991; Holyoak

& Thagard, 1989; Kedar-Cabelli, 1985; Winston, 1980).

Although theorists have generally acknowledged that analogical inference is somehow influenced by goals and context, views differ as to whether such pragmatic constraints enter into the mapping process. The clearest contrast is offered by two general theories of analogical mapping: Gentner's structure-mapping theory (1983, 1989) and Holyoak and Thagard's multiconstraint theory (Holyoak, 1985; Holyoak & Thagard, 1989). The structure-mapping theory and its computational implementation, SME (Structure-Mapping Engine) operating with structure-mapping rules (Falkenhainer et al., 1989), postulates that the mapping process is purely syntactic and is not directly influenced by pragmatic constraints. "[P]lans and goals and other aspects of current context influence the analogy process *before* and *after* the analogy engine but not during its operation" (Gentner, 1989, p. 215; italics in original).

In contrast, the multiconstraint theory assumes that goals and context affect every stage of analogical inference, including the central mapping stage, yielding the prediction that "Differing goals can lead to different mappings for what is putatively the 'same' analogy..." (Holyoak, 1985, p. 76). Holyoak and Thagard (1989) proposed a computational model called ACME (Analogical Constraint Mapping Engine), in which pragmatic and other constraints interact during the mapping process to generate correspondences by parallel constraint satisfaction. Here we will report the results of an experiment testing ACME's prediction that different goals of the analogist can lead to different mappings for the same analogy, and that goals actively interact with other types of constraints on mapping. To model our results we develop an extension of ACME in which an inhibitory mechanism for selective attention allows goals to influence mapping.

We wished to show that mapping (in addition, presumably, to pre-mapping and post-mapping processes) is a locus of goal-directed processing. Pre-mapping selection on the basis of the goal could operate by blocking goal-irrelevant information in long-term memory from the working memory used for mapping. By analogy to Broadbent's (1958) model of selective attention, we will refer to the hypothesis that pragmatics entirely screens out goal-irrelevant information prior to

* This research was supported by a UCLA Academic Senate Research Support Grant to K. Holyoak and NSF Grant DIR-9024251 to the UCLA Cognitive Science Research Program. We thank Eric Melz for his technical assistance.

mapping as the "filter" hypothesis, and by analogy to Treisman's (1964) alternative to Broadbent's model we will refer to the hypothesis that pragmatics deemphasizes goal-irrelevant information without necessarily screening it out entirely as the "filter-attenuation" hypothesis. Note that if selective attention is viewed as an inhibitory process, then complete filtering is the logical extreme of a filter-attenuation mechanism in which the deemphasized information is fully inhibited. The ACME model, which treats pragmatic importance as a continuum, is compatible with the filter-attenuation hypothesis. To distinguish these alternatives, we sought evidence that processing goals affect mapping without completely screening out goal-irrelevant aspects of the source prior to mapping. To do so we attempted to show that goal-irrelevant information actually does influence subjects' mappings, serving to disambiguate mappings that cannot be decided on the basis of the goal-relevant information alone. Paradoxically, establishing that goals influence the mapping process itself, rather than solely acting as a pre-mapping filter on inputs, requires simultaneously demonstrating a secondary influence of goal-irrelevant aspects of the analogs.

Finally, to rule out an explanation solely in terms of post-mapping processes, such as transfer and adaptation of a problem solution (Novick & Holyoak, 1991), our main dependent measure was a mapping task that did not require additional post-mapping processes.

Analogies Between Soap-Opera Plots

Our experiment involved analogies in which certain key objects were four-ways ambiguous on structural grounds alone, with two alternative mappings being supported by each of two sets of relations. The processing goal was then manipulated so as to emphasize one of the two sets of relations. If the processing goal were used to filter out all information that was not emphasized by the pragmatic manipulation, then two mappings would remain viable, with no basis for selecting between them. However, if the processing goal emphasizes the goal-relevant information but does not completely filter out goal-irrelevant information, as the filter-attenuation hypothesis allows, then the goal-relevant and goal-irrelevant information could jointly determine a single optimal mapping for each ambiguous element.

Method

Ninety-six subjects enrolled in introductory psychology at UCLA participated in the experiment, which was based on a scenario involving plagiarism of soap-opera plots. Subjects were told to pretend that they were successful

writers of a new soap opera, and that they were in court trying to prove that writers from another soap opera had stolen their ideas. First they were presented with the source -- the plot of their own soap opera, called Soap Opera University (SO-U). Characters in SO-U included an ex-alcoholic professor named Peter, Peter's research assistant Susan, who was now running Peter's lab, Peter's brother Ben, and some other minor characters. The three major relations between these characters were: Professional (Peter was Susan's boss); Romantic (Peter was in love with Susan); and Inheritance (Peter cheated Ben out of money the latter should have received from the will of a mutual relative).

The plot of the other soap opera, Soap Opera City (SO City), involved two fairly distinct sets of characters. The "lawyer set" included Leslie, an ex-addict entertainment lawyer, and Mark, a young lawyer working at her firm who had often filled in for her. The "doctor set" included David, a prominent physician who had suffered a nervous breakdown, and Felice, an intern who was now treating most of David's patients. Leslie and David were half-siblings and Mark and Felice were cousins. Both pairs had aging relatives ready to leave them money in a will; in one version of the story Leslie and Felice (the women) cheat David and Mark (the men), respectively, out of their shares of the inheritance, and in the other version the men cheat the women out of their shares. In SO City, the three analogous major relations were as follows: Professional (Leslie was Mark's boss and David was Felice's boss); Romantic (Mark was in love with Leslie and Felice was in love with David); and Inheritance (either Leslie cheats David and Felice cheats Mark or vice versa). From this description the object mappings are ambiguous; for example, if the women are the cheaters then Peter seems to map equally to Leslie and Felice.

The processing-goal manipulation was accomplished using a "plot-extension" task. At the end of the description of SO-U, subjects were told what happened in the Thursday and Friday episodes. In one of the episodes, the "Professional continuation", Peter steals the credit for Susan's successes; in the "Romantic continuation", Peter goes to Susan's apartment and tries to persuade her to let him spend the night. After reading the description of the characters on the plagiarized show, SO City, subjects were told that a judge had seen a copy of the tape of the next SO City episode, that it involved three characters, and that it was "just like" the SO-U Thursday (or Friday) show. To demonstrate that SO City was really stealing the ideas of SO-U, subjects were asked to describe what they think happened on the tape and which three characters were involved. They were told that predicting these events successfully would

Table 1. Optimal mappings for the main source characters based on goal-relevant relation (Professional or Romantic) and gender of cheater in the Inheritance relation.

Professional Plot - Extension		Role	Source Characters	Role	Romantic Plot - Extension	
Gender of Cheater Male	Female				Male	Female
DAVID	LESLIE	Boss	PETER	Pursuer	MARK	FELICE
FELICE	MARK	Underling	SUSAN	Pursued	LESLIE	DAVID

be the strongest possible evidence that the other writers were stealing their ideas. It was assumed that subjects would write plot extensions that were analogous to the appropriate subplot continuation; that is, that the pragmatic manipulation should define the choice of continuation used in the plot-extension task. The three characters would include a Peter-analog, a Susan-analog and an analog of the minor character involved in the relevant continuation. By providing all subjects with extensions of the source based on both Professional and Romantic relations, and by counterbalancing which of these extensions became the pragmatic focus, the design made it possible to control for the sheer number of possible inferences that could be produced using mappings based on each of the two sets of primary relations. This is an important control because the SME model postulates that number of possible inferences is a factor that influences mapping preferences. The present experiment held number of possible inferences constant as processing goals were manipulated.

Subjects were then given a mapping task. They were told that the judge wanted them to explicitly tell which characters of SO City were like which characters of SO-U because that would provide even more evidence of the plagiarism. The six characters from SO-U were listed and subjects were encouraged to match only one character from SO City to each character of SO-U. If goals operate in accord with the filter-attenuation hypothesis, then subjects' choices of mappings may reflect both the pragmatic manipulation and the information conveyed by the incidental Inheritance relation.

Table 1 illustrates the predicted mappings for the two central characters in the source as a function of our two main manipulations (i.e., which subplot was relevant in the plot-extension task and who cheated whom out of their inheritance). Each of these source characters -- Peter and Susan -- may map to any of four characters from the target, depending on the pragmatic manipulation and the incidental Inheritance relation. For both the plot-extension and the mapping tasks, subjects' choice of mappings for these characters made it possible to diagnose whether they were sensitive to the

pragmatic focus (either Professional or Romantic) and/or to the incidental Inheritance relation. For example, consider the possible mappings for Peter summarized in Table 1. Suppose the plot extension made Professional relations the pragmatic focus. If subjects prove sensitive to this manipulation, they should tend to map Peter to either David or Leslie, rather than Mark or Felice. Suppose also that the incidental Inheritance relation established that the cheater was female. If subjects were sensitive to the Inheritance relation, they should map Peter to either Leslie or Felice, rather than David or Mark. Finally, if subjects in this example were sensitive to both the pragmatic focus *and* to the incidental relation, they should map Peter to Leslie -- the one choice that is supported both by the Professional and the Inheritance relations. The design thus makes it possible to determine whether the impact of the pragmatic manipulation can be attributed to a mechanism that entirely screens out goal-irrelevant relations prior to mapping (the filter hypothesis), or whether the processing goal serves to emphasize the relevant relation without entirely suppressing incidental relations (filter-attenuation hypothesis)

Because the plot-extension task directly depended on a specific type of relation (either Professional or Romantic), we would expect pragmatic focus to have a strong impact on the choice of characters used in that task. To the extent that a strong pragmatic focus implies attenuation of goal-irrelevant relations, incidental Inheritance should have a reduced effect on character choice in the plot-extension task. In contrast, the mapping task does not directly demand that the goal-relevant relation be of exclusive concern; hence the impact of the pragmatic manipulation would be expected to be less extreme, and the influence of the goal-irrelevant relations would be less attenuated. The mapping task is therefore more likely than the plot-extension task to provide evidence that goal-relevant and goal-irrelevant relations can jointly guide the choice of mappings for ambiguous characters.

The complete design had the form of a 2x2x2x2x6 factorial. The factors of theoretical interest were the subplot involved in the plot-extension task (Professional or Romantic) and the




Figure 1. Percentages of subjects in the plot-extension task (left) and the mapping task (right) who mapped major characters (Peter and Susan) in accord with the goal-relevant relation (either Professional or Romantic) and with the incidental Inheritance relation.

gender of the cheater in the Inheritance relation (male or female). The other factors (order of subplot continuations in source; order of character sets in target; order of names on the mapping task) represented counterbalancing of various presentation details. Because the design generated a total of 96 possible conditions (including counterbalancing), each subject in the experiment received a four-page booklet containing a unique combination of materials.

Results and Discussion

For both the plot-extension task and the mapping task, the primary focus of the analyses was on answering two questions. First, were subjects' mappings primarily guided by the goal-relevant relation (either Professional or Romantic)? Second, were subjects' mappings also influenced by the incidental Inheritance relation?

Plot-extension task. In order to determine whether subjects' choices of characters were guided by the goal-relevant relation (either Professional or Romantic) and/or by the incidental Inheritance relation, we scored the frequencies with which various possible combinations of characters in the target soap opera were used as analogs of the Peter (boss/pursuer) and Susan (underling/pursued) characters in the source. As summarized in Table 1, there were four possible choice combinations (the two characters in each column) that would provide a structurally consistent pair of analogs to Peter and Susan. Of these four combinations, two were consistent with the goal-relevant relation (Professional or Romantic) and two were consistent with the incidental Inheritance relation. Figure 1 (left panel) presents the percentage of subjects who selected each of the four possible consistent pairs, or some other inconsistent or unrelated character set. The great majority of subjects wrote story extensions using one of the two pairs compatible with the goal-relevant relation. The preference for characters consistent with the goal-relevant rather

than the alternative goal-irrelevant relation (88% versus 2%) was highly significant, $X^2(1) = 78.2$, $p < .001$. The preference for mappings consistent with the incidental Inheritance relation was smaller (53% versus 36%) and not significant, $X^2(1) = 2.98$, $p < .10$. Overall, then, character selection in the plot-extension task was consistent with either the filter hypothesis, or with the filter-attenuation hypothesis assuming high inhibition of all goal-irrelevant relations.

Mapping task. A similar analysis of character selection was performed using data from the mapping task, which did not so directly demand that subjects focus on the goal-relevant relation. Overall, subjects preferred to map to characters based on the Professional subplot rather than the Romantic subplot, $X^2(1) = 45.8$, $p < .001$. Figure 1 (right panel) presents the percentage of subjects producing various possible combinations of mappings for Peter and Susan, collapsing over the Professional and Romantic conditions. Subjects tended to map characters on the basis of the goal-relevant relation more frequently than on the goal-irrelevant relation (49% versus 31%), $X^2(1) = 3.75$, $p < .06$, although the preference for mappings supported by the goal-relevant relation was less pronounced in the mapping task than in the plot-extension task. In contrast to the plot-extension data, data from the mapping task revealed a significant preference for mappings consistent with the incidental Inheritance relation (55% versus 25%), $X^2(1) = 10.92$, $p < .001$. Thus subjects' preferences in the mapping task provided evidence that both the goal-relevant relation and an incidental relation were used in the mapping process, supporting the filter-attenuation hypothesis.

Attentional Inhibition in ACME

The resolutions of ambiguous mappings in the two tasks revealed a trade-off between the potency of the processing goal and of the incidental

Inheritance relation. The influence of the processing goal was much larger in the plot-extension task than in the explicit mapping task. Conversely, the influence of Inheritance was somewhat larger in the mapping task than in the plot-extension task. To test whether the filter-attenuation hypothesis can account for the major qualitative aspects of the the results from both tasks, we performed a series of simulations using ACME. To do so we extended ACME to use inhibition as a mechanism for goal-directed mapping. This extension avoids some computational problems that arise when excitation is used as the sole basis for favoring goal-relevant over irrelevant information (Hummel, Burns & Holyoak, in press). In addition, empirical evidence implicates inhibition as a mechanism for selective attention in other cognitive tasks (Tipper, 1992).

Our simulations were based on predicate-calculus representations of the two soap operas, with 99 propositions representing the source (SO-U) and 136 propositions representing the target (SO City). Within the SO-U representations, we identified 32 propositions in the initial set-up that were causally related to the episode that constituted the Professional extension, and a different 17 propositions that were causally related to the episode that constituted the Romantic extension. To model the manipulation of processing goals, the propositions in the goal-relevant subset of the target were marked as "important". These propositions, as well as the objects and predicates used in them, were then spared from inhibitory activation that was directed at all other mapping units, which involved unimportant (i.e., goal-irrelevant) elements of the target. This inhibition was implemented by a negative weight from the clamped "pragmatic unit" (see Holyoak & Thagard, 1989, for a detailed discussion of ACME and its parameters). Thus when the Professional propositions were marked as important, the Romantic and Inheritance information was inhibited; whereas when the Romantic propositions were marked as important, the Professional and Inheritance information was inhibited.

To model the difference between the plot-extension and mapping tasks, we assumed that attentional inhibition was higher in the former task. This assumption seems justified by the fact that the requirement to generate an analogous episode based on the SO-U characters directly demanded attention to the goal-relevant subset of the source, whereas the explicit mapping task did not call attention to any particular subset of the source. The sole difference between simulations of the plot-extension and mapping tasks was in the parameter value for attentional inhibition, which was set at either a high (-.04) or low (-.005) value,

respectively. Other major parameter settings were held constant as follows: excitation, .005; inhibition (structural), -.16, similarity of identical predicates, .005; decay, .005; starting activation for all units, .001. The Grossberg updating rule, with maximum activation of 1 and minimum activation of -.3, was used to settle the network. The mapping network for the problem contained 4037 units interconnected by 146718 links. Simulations were performed using a version of ACME written in *LISP running on a 16000-processor CM2 Connection Machine.

In an initial set of four runs we varied which primary relation was important (Professional or Romantic) and whether attentional inhibition was high or low. In these runs the "cheater" in the Inheritance relation was female. (The symmetry of the representations ensures that identical results would be obtained if the gender assignment was reversed.) ACME is implemented as a deterministic model, and the structure of the ambiguous mapping has a "Necker cube" quality, such that a single consistent mapping will emerge given even a small initial pressure favoring one mapping over the others (cf. Spellman & Holyoak, 1992). Thus all four basic versions of the problem settled into the optimal consistent mapping for Peter and Susan (i.e., the mapping consistent with both the goal-relevant relation and Inheritance), with activations of .80 or above, after 331-503 cycles of updating.

To assess the relative potency of the various factors, we tested the robustness of each when it was pitted against a bias toward a contrary mapping. The greater the impact of a factor, the more it should resist the opposing pressure of the bias (which can be viewed as an approximation to the expected impact of adding variability to the mapping mechanism). To assess the relative impact of the Inheritance factor at different levels of attentional inhibition, a bias was introduced by adding a link with a small positive weight (.0001) from the pragmatic unit to the non-optimal mapping unit for Peter (i.e., the bias favored a match for Peter that differed from the optimal one only in terms of the Inheritance relation). When attentional inhibition was high, the bias prevailed, driving the mappings for both Peter and Susan into those contrary to the Inheritance relation. But when inhibition was low, the Inheritance relation overcame the bias, producing the optimal correspondences for both characters. This pattern was observed both when the Professional and when the Romantic propositions were treated as important. These simulations thus reproduced a qualitative aspect of subjects' mappings: a stronger impact of the Inheritance factor in the mapping task (low inhibition) than in the plot-extension task (high inhibition).

Another set of runs was performed to assess the robustness of the preference for mappings based on the goal-relevant rather than goal-irrelevant relation (i.e., Professional versus Romantic or vice versa). In these runs the bias favored mapping Peter to the person who would be consistent with the goal-irrelevant relation (as well as the Inheritance factor). The bias was increased gradually across each set of runs to determine at what point the impact of the goal-relevant relation would be overcome by the competing bias. When the Professional propositions were treated as important and attentional inhibition was high (simulating the plot-extension task), the mapping for Peter favored by the bias did not prevail until the bias weight reached -.005; moreover, no amount of bias for the Peter mapping was able to reverse the preferred mapping for Susan. In contrast, when attentional inhibition was low (simulating the mapping task), the mapping for Peter was reversed when the bias reached -.0025, and that for Susan was reversed as well when the bias reached -.045. Thus the goal-relevant Professional relation was substantially more resistant to an opposing bias when attention inhibition was high rather than low.

Comparable runs in which the Romantic propositions were treated as important produced a weaker pattern. When attentional inhibition was high, the mappings for both Peter and Susan were reversed when the bias reached -.0025. When attentional inhibition was low, the same reversal occurred at a slightly lower level of bias, -.002. Although the effect was thus much weaker when the Romantic rather than the Professional propositions were important, in both cases the goal-relevant relation proved more potent when attentional inhibition was high than when it was low, just as subjects showed a stronger preference for mappings based on the goal-relevant relation in the plot-extension than in the mapping task. Overall, the lesser potency of the Romantic than the Professional relation in the simulations (resulting from the smaller number of propositions based on the former relation) had its counterpart in subjects' strong overall preference for mappings based on the Professional relation in the explicit mapping task.

Conclusion

The results of the experiment revealed that processing goals can interact with incidental information to disambiguate mappings, supporting the claim that processing goals modulate the mapping process. Different goals can indeed yield different mappings for the same source and target analogs. Although it is no doubt possible to develop an account of the present results (or any

other pattern of mappings) in terms of some combination of pre-mapping and post-mapping processes, the obtained pattern can be explained parsimoniously in terms of an inhibitory mechanism for control of selective attention, which influences a process of mapping by constraint satisfaction.

References

- Broadbent, D. E. (1958). *Perception and communication*. London: Pergamon Press.
- Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41, 1-63.
- Forbus, K. D., & Oblinger, D. (1990). Making SME greedy and pragmatic. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. (pp. 61-68). Hillsdale, NJ: Erlbaum.
- Gentner, D. (1983). Structure-mapping: A theoretical framework. *Cognitive Science*, 7, 155-170.
- Gentner, D. (1989). The mechanisms of analogical reasoning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). London: Cambridge University Press.
- Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation* (Vol. 19). New York: Academic Press.
- Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology*, 44, 265-315.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295-355.
- Hummel, J. E., Burns, B., & Holyoak, K. J. (in press). Analogical mapping by dynamic binding: Preliminary investigations. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory, Vol. 2: Analogical connections*. Norwood, NJ: Ablex.
- Kedar-Cabelli, S. (1985). Purpose-directed analogy. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
- Novick, L. R., & Holyoak, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 398-415.
- Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler then who is George Bush?: Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, 62, 913-933.
- Tipper, S. P. (1992). Selection for action: The role of inhibitory mechanisms. *Current Directions in Psychological Science*, 1, 105-109.
- Treisman, A. (1964). Monitoring and storage of irrelevant messages in selective attention. *Journal of Verbal Learning and Verbal Behavior*, 3, 449-459.
- Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM*, 23, 689-703.