

Partial Match and Search Control via Internal Analogy*

Angela Kennedy Hickman
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
ach@cs.cmu.edu

Marsha C. Lovett
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213
lovett@psy.cmu.edu

Abstract

In a previous study (Hickman & Larkin, 1990), we introduced a within-trial analogy mechanism called *internal analogy* that transfers both success and failure experiences between corresponding parts of the search tree for a single problem. In this paper, we describe powerful extensions to the learning procedure and their consequences on problem solving behavior. First, we explain how our similarity metric can be naturally augmented to provide a more flexible partial match. To overcome the need for a static measure, however, we propose a mechanism that *learns* the appropriate level of partial match through feedback from previous analogical reasoning. Second, we show how this partial match mechanism controls the problem solver's search. Protocol data from a subject working in a geometry theorem-proving domain provide support for the psychological fidelity of the extended internal analogy model.

Introduction

Analogical reasoning is a psychologically-motivated computational mechanism for efficient problem solving. In essence, it reduces search by modifying and transferring components of previous solution instantiations to guide the current problem solving. Unlike macro-operators (Anderson, 1983), chunking (Laird & Newell, 1986), and EBL (Mitchell, Keller, & Kedar-Cabelli, 1986), analogy employs partial match and adaptation of past knowledge instead of exact match and direct replay. This allows a previous experience to provide some search guidance even if it cannot provide an exact solution. Although the process is not completely understood, evidence suggests that people reason analogically when learning a new domain (Faries & Reiser, 1988) or to reduce the amount of effort required to solve a problem in a known domain.

In a previous study (Hickman & Larkin, 1990), we introduced a psychologically motivated technique called *internal analogy* that transfers both success and failure experiences from a previously solved subgoal to a current subgoal that is contained in the same problem. Most previous

studies of analogy in problem solving have investigated transfer between problems from similar domains (within-domain analogy) (Carbonell & Veloso, 1988). Psychological models of analogy have concentrated on the transfer of knowledge between similar problems in different domains (cross-domain analogy) (Gentner & Toupin, 1986). In general, both of these types of analogy require an acceptable partial match between entire problem instances and then transfer entire problem solutions. On the other hand, EUREKA (Jones, 1989) uses analogical reasoning to choose only the next problem solving step. Unlike all of these systems, internal analogy takes advantage of whatever regularity exists between the two problems in as efficient a manner as possible. Since the process may be applied to multiple subgoals while solving a single problem but does not require the entire problem definition to acceptably match a source problem, it should be better able to exploit symmetries and other regularities that exist in the search space. Being able to accomplish transfer at varying levels of granularity is an important mechanism for optimally reducing search. As a result, internal analogy appears more widely applicable, providing greater search reduction.

The current implementation of internal analogy provides only for exact match and direct replay. In this paper, however, we first describe how the *information content metric*, our similarity metric, can be modified to implement a more flexible partial match. In order to overcome the need for a static measure, though, we describe a mechanism that *learns* the appropriate level of match through feedback from previous analogical reasoning. Using this new match criterion, internal analogy transfers the solution in an efficient compiled form and also automatically identifies what information is missing and requires adaptation. The adaptation process is carried out by recursively applying internal analogy or other general problem solving methods to find the missing segments of the solution. Secondly, we show how this new partial match mechanism will provide search control for the problem solver by using it to model protocol data of a subject solving geometry problems. This analysis provides even more support for the psychological fidelity of the internal analogy model.

The next section presents the model of the extended internal analogy process, and the following section provides evidence from the protocol that supports the psychological fidelity of the extended model. The final section summarizes the results and discusses plans for future work.

*This research was supported in part by DARPA contract number F33615-87-C-1499 Amendment 20, and AIP-ONR contract number N00014-86-K-0678-N123. Angela Kennedy Hickman was partially supported by a Zonta International Amelia Earhart fellowship. Marsha Lovett was supported by a NSF Graduate Research Fellowship. We thank our collaborators, John Anderson, Jaime Carbonell, Jill Larkin, and Peter Shell for their insights, and Peter for his invaluable help in the implementation of the model.

Internal Analogy Model

Before describing the extensions to the model, we briefly review the internal analogy process. When a new subgoal is opened, the system examines previously solved goals of the same type, the candidate source goals. It then calculates something similar to the weakest preconditions for a candidate source's solution from the instantiated productions that compose that solution.¹ If these preconditions are all satisfied in the current target problem (information-content(candidate) \leq information-content(target)), then the candidate is chosen as the actual analogical source. Otherwise, it is rejected. If the candidate is chosen as the actual source, its solution is appropriately mapped by instantiating the productions that solved it in the current working memory and replaying them to solve the current subgoal. This process operationally defines the relevant structure that is important to be mapped. Since this newly solved subgoal is contained in the current problem, it automatically becomes available as a future candidate source. The system learns from its failures using the same mechanisms. The interested reader can find more details on the algorithm in (Hickman, Shell, & Carbonell, 1991).

Empirical studies and a theoretical analysis demonstrate that this version of internal analogy produces significant savings in search during problem solving (Hickman, Shell, & Carbonell, 1991). However, it is extremely conservative; it errs by undergeneralizing, not overgeneralizing. This conservative strategy may inhibit the learner from using source solutions that actually could provide some, if not complete, search guidance. To reach its full effectiveness, the internal analogy process must be extended to perform a more flexible partial match and adaptation. In the following discussion, we show how natural extensions to the mechanisms already in place can produce this behavior.

The two similarity metrics the system currently uses are the comparisons of 1) the quantity types and 2) the information contents of the source and target goals. One way to relax the current exact match is by allowing goal types that are "close" in the type hierarchy to match. Additionally, we can expand our definition of compatible information contents. For instance, if in the initial state of the candidate source, a variable had a known value which is not known in the current state, the candidate source may still be determined to be compatible with the target goal. In that case, internal analogy can apply the source solution until the value of the unknown variable is needed. At that point, it can recursively use the problem solver's operators (which include internal analogy) to find the value of the unknown variable. This affords internal analogy a way of transferring compiled solutions at the correct level of granularity to achieve maximal efficiency. The ability to adapt the solution where needed and to fail it completely if necessary are capabilities which chunking, EBL, and

¹This can be viewed as using a form of guided lazy EBL for retrieval. However, the presence of the target goal provides guidance as to 1) what to learn and 2) what amount of generalization to perform. Since our original study in 1990, Veloso has implemented a closely related mechanism for PRODIGY, a general purpose planner (Veloso & Carbonell, 1991).

macro-operators do not possess. All of these methods require complete justification before applying the compiled knowledge. Thus, by using partial match and adaptation during the reuse of experience, internal analogy will provide search control in many cases that the other methods cannot.

Next, we refine the notion of partial match even further. First, we associate a deviance measure with each pair of types in the type hierarchy. Second, we cache two deviance weights for each operator precondition. The purpose of these weights is to reflect how closely the weakest preconditions contained in the source information content are required to match the target. The first one of these weights denotes how deviant it is for an information content precondition to be totally absent from the current working memory. The second weight denotes how deviant it is for this precondition to be filled by the same type working memory element instead of by the exact same working memory element. To compute the deviance between the candidate source and target subgoals, we combine the quantity type deviance measure with the information content deviation weights, using the least deviant instantiation of the target information content.

With the partial match mechanisms in place, we can extend the functionality of the system to *learn* the correct circumstances under which solutions can transfer. Alternately, one could view this as learning the correct notion of partial match for the domain. This learning is accomplished by the system as it tunes the deviance measures described above according to feedback from the previous analogical reasoning. The feedback contains information about whether the analogy failed or succeeded and how much effort was required to adapt the retrieved solution. In summary, not only do these deviance measures provide the system with a method for choosing the best source, they provide a way of learning the definition of "best" for any given domain.

The consequence of this new flexible partial match mechanism is that it provides even greater search control for the problem solver than the mechanism did previously. Comparing the deviance measures for all the candidate sources, internal analogy can perform a least-deviant-first search over its memory to find the best candidate source to choose as the actual source. The lower the total deviance of the source, the less adaptation that will be required. As such, there is a tradeoff between the amount of search the system does to find the best source and the amount of adaptation that source requires. In addition, the acceptability of the best candidate source is determined by a threshold that rejects any source whose total deviance is too high. This threshold is based on how much search the problem solver has already undertaken. If the best source is determined to be acceptable, the problem solver reuses the solution stored with that source. As the deviance measures become "tuned" through learning, the search for the best source gets even smarter. Thus, internal analogy finds increasingly better analogical sources by a smarter search through memory instead of by a constant reorganization of it.

In addition to providing search control for backward

chaining as described above, internal analogy utilizes the new partial match metric in a different way to provide search control for forward chaining. In forward chaining, there is no target goal driving the search. However, there is a decision point at each operator application. Therefore, instead of fixing the target goal and finding the source goal via least-deviant-first search, the forward chaining problem solver picks a source goal, most likely via recency, and uses the partial match mechanism to find the least deviant target goal. In this way, the problem solver uses a known strategy to try to easily elucidate more knowledge about a problem. As can be seen in the next section, this new partial match mechanism provides complex but useful search control guidance to the problem solver, which is not easily explained using any existing theories.

Evidence from the Protocol

The original computational model was compared with the protocols of four subjects solving DC-circuit and fluid statics problems, and traces of its behavior matched the subjects' use of internal analogy. In the current paper, we explain the behavior of one subject, WB, solving a problem in the domain of geometry theorem-proving using the extended internal analogy model. This domain and the resulting problem space is more complex than the DC-circuit domain, and thereby provides an even better test of the extended model's psychological fidelity and its general applicability.

To begin the experiment, the subject, WB, studied a geometry review sheet, which remained available to him throughout the study. Next, four problems were presented to WB that had the potential to elicit both internal and inter-problem analogy. (In the analysis below, we describe his protocol for the third problem.) After the subject finished each problem, the experimenter removed the review sheet from view and presented the next problem.

WB's search for a solution to the problem in Figure 1 can be segmented into three parts: the opening, the intermediate segment, and the final path. These segments are distinguished by the types of goals WB verbalizes, the strategies he employs, and the degree of success he attains.

During the opening, WB restates the problem and weighs the potential of various solution strategies. He then embarks on a forward chaining sequence that includes all the conservative applications of internal analogy. Even though WB's opening strategy does not seem to be directed at fulfilling any specific goal, it does provide a relatively easy way of elucidating a lot of new knowledge about the problem, some of which is needed for his solution.

The intermediate segment of the protocol is generally characterized by unsuccessful backward chaining and exhaustive search. WB attempts to backward chain on the current goal three times during this segment, yet he never gets beyond the first level of goal decomposition. Instead, WB spends the main portion of this segment doing exhaustive search through the entire set of operators. He accomplishes this twice by reading through the list of theorems on the review sheet and sequentially checking the applicability of each one. During this segment, however,

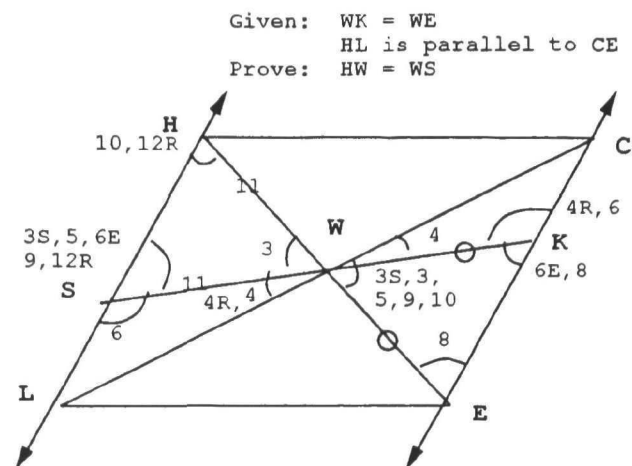


Figure 1: Problem diagram. The numbers in this diagram represent the result of WB's problem solving at each step; numbered arcs denote congruent angles, and numbered lines denote congruent line segments.

WB only gains two new pieces of knowledge.

WB's final path to solving this problem includes backward chaining with multiple applications of internal analogy. In this segment, WB solves the problem directly with no extraneous search.

In the following two subsections, we present a more detailed analysis of WB's protocol. It provides evidence for the psychological fidelity of the internal analogy model's extensions - the new partial match metric and the associated least-deviant-first search.

Partial Match

During the forward chaining opening segment of the protocol, WB first uses internal analogy as his fourth search step (See Figure 2). His third step, then, is important because it serves as a source for the subsequent internal analogy. (Indeed, it is the only available source at this point in WB's problem solving.) In the third step of the protocol, WB tries to use the Alternate Interior Angles Corollary (AIAC, Figure 3). However, there are two bugs in his application of AIAC. First he states that $\angle HSW$ and $\angle KWE$ are alternate interior angles with respect to the parallel lines HL and CE and the transversal SK (Line 3S, Figure 2; and the corresponding step, Figure 1). This is an incorrect definition of AIAC, which WB reuses throughout the protocol; we call it AIAC-av. Second, WB does not correctly write down his verbalization of this step on the solution sheet. Instead, he transcribes that $\angle SWH = \angle KWE$, thus, introducing the second bug² (Line 3, Figure 2). We call this incorrect definition of alternate interior angles AIAC-vv. (According to the AIAC rule in Figure 3, these bugs correspond to incorrect definitions of the precondition $alt-int(\langle a \rangle, \langle b \rangle, \langle line1 \rangle, \langle line2 \rangle, \langle t \rangle)$.)

Next, WB immediately applies internal analogy saying, "through using the same theorem [AIAC], I can say that $\angle CWK$, angle $\angle CWK$, is congruent to, not sure if this is going

²This second bug fortuitously yields a correct piece of knowledge since $\angle SWH = \angle KWE$ by the vertical angles theorem.

Step	Operator	Instantiation	Added knowledge
0.	Read problem	Current problem	WK=WE, HL//CE:
1.	CTCP	tri-SHW, tri-WKE, ?, ?	Suspend
2.	AIAT	HL//CE	Already given
3S.	AIAC-av	HL//CE, SK, <HSW, <KWE	<HSW = <KWE:
3.	AIAC-vv	HL//CE, SK, <SWE, <KWE	<SWH = <KWE:
4.	IA	HL//CE, SK, <CWK, <SWL	<CWK = <SWL
4R.	AIAC-av	HL//CE, SK, <CKW, <SWL	<CKW = <SWL
5.	IA	HL//CE, SK, <HSW, <KWE	<HSW = <KWE
6.	IA-aa	HL//CE, SK, <CKW, <WSL	<CKW = <WSL
6E.	Told	Replace knowledge from 5.	<HSW = <KWE

Figure 2: The Opening Segment of the Protocol. CTCP = Congruent Triangles Congruent Parts, AIAT = Alternate Interior Angles Thm., AIAC = Alternate Interior Angles Cor., IA = Internal Analogy.

```

IF    angle(<a>)
        angle(<b>)
        goal(congruent(<a>,<b>))
        parallel(<line1>,<line2>)
        transversal(<t>)
        alt-int(<a>,<b>,<line1>,<line2>,<t>)
THEN congruent(<a>,<b>)

```

Figure 3: The AIAC Rule.

to help me but it never hurts, to angle SWL.”* (Line 4, Figure 2) The target in Line 4 is a partial match to the source in Line 3, and its information content deviance with respect to the source is very low. To understand why, consider the AIAC rule shown in Figure 3. Since this internal analogy application transfers only a simple operator, the information content of the source includes instantiations for all the variables in the left-hand side of the rule. WB’s target instantiation contains all the same working memory elements, except for the pair of angles. This produces a deviation value whose only non-zero terms correspond to these new angle instantiations. Due to the physical constraints of the geometry domain, changing any other variable in this rule (e.g. the parallel lines or the transversal) would necessarily imply a change of angles as well, and hence, a greater deviation value. Therefore, the optimal choice for the target instantiation is one that only changes the instantiations of the angles. This is exactly WB’s choice in his internal analogy of Step 4.

As shown above, the partial match metric models WB’s choice of a target for this internal analogy because it uses a deviance metric that is very sensitive to the similarity of the source and target instantiations. This sensitivity eases the problem solver’s task not only by providing search control (as discussed in the next section), but by easing the burden of instantiation. By using a least-deviant-first search strategy, the problem solver is guaranteed to choose an analogy that requires making the fewest changes to the target instantiation. WB’s protocol supports this efficiency argument, for when he first applies the AIAC rule, he verbalizes the entire instantiation: “Okay, and the what I want to say in here is use line... Which do I want to use as the transversal? ... Okay, first of all, I’m going to say that HL and CE are bis- are cut by a transversal line SK.”

Upon applying the internal analogy in step 4, however, he merely states the changes to the source instantiation, the new pair of angles. (See *’d protocol excerpt above.) This same behavior was observed in the forward chainer from the previous study.

Search Control

Now we describe how the new partial match metric explains WB’s complex but useful search strategy. WB begins his problem solving by stating the givens and the goals of the problem and then considers whether to use a generalized schema that consists of proving Congruent Triangles and then proving Congruent Parts (CTCP; Line 1, Figure 2). In weighing the merits of this strategy, he says, “Ok. There’s a long way of doing and proving congruent triangles and then showing there are corresponding parts of congruent triangles which might be easier because I’ve just done that [referring to an earlier problem in the study]. But then there’s also this alternate interior angles theorem which actually probably gets me the same way just a little bit quicker.” With the statements, “long way,” and “quicker,” WB weighs the complexity, which relates to the size of the information content, and the optimality of this solution (Carbonell & Veloso, 1988). In opposition to this, he says, “might be easier because I’ve just done that.” With this statement, the subject is estimating the efficiency of finding a solution by referring to the last problem he completed, during which he successfully used the CTCP schema. Thus, WB is considering using inter-problem analogy, which can also be explained using our notion of partial match. At this point in the problem solving, none of the preconditions of this schema are known; the total deviance measure exceeds WB’s acceptability criterion; and the analogy is abandoned.

After reaching a dead end in Line 2, WB produces the buggy internal analogy of Lines 3 and 4 described in the previous subsection. At this point, the experimenter points out that <CWK and <SWL are not alternate interior angles with respect to lines HL, CE, and SK. The subject agrees and produces the behavior in line 4R. Notice that he now has returned to his original buggy definition of alternate interior angles first exhibited in Line 3S. WB’s second application of internal analogy now occurs as he reapplies in Line 5 the immediately preceding AIAC-av operator to obtain the congruence of <HSW and <KWE. For the same reasons explained in the preceding section, this is the least deviant target with respect to the source in Line 4R; WB’s behavior supports the predictions of the model once more. WB’s third application of internal analogy occurs in Line 6 when he says, “Through the same reasoning <CKW is congruent to <WSL.” In this least-deviant application, WB introduces yet another bug into his original buggy definition of alternate interior angles (AIAC-av). This bug yields the correct definition, and the subject proves <CKW = <WSL. Finally, the experimenter corrects the knowledge WB derived in Line 5 and establishes the correct knowledge <HSW = <KWE in Line 6E. At this point, there are no more analogous applications of AIAC left that involve the transversal SK. Since it is still early in the problem solving, changing the

Step	Operator	Instantiation	Added knowledge
9.	IA	HL//CE, SK, <HSW, <EWK	<HSW = <EWK
10.	IA	HL//CE, HE, <SHW, <EWK	<SHW = <EWK
11.	ITC	<HSW, <SHW, HW, WS	HW=WS
12R.	TP	<HSW, <EWK, <SHW	<HSW=<SHW

Figure 4: The Final Path Segment of the Protocol. IA = Internal Analogy, ITC = Isosceles Triangle Cor., TP = Transitive Property.

transversal raises the deviation level too much, and it is not an acceptable target. As a result, the subject has exhausted his available applications of internal analogy and must pursue another strategy.

During this opening segment, WB performs all possible applications of internal analogy, with the constraint that he allows only small deviations between the source and target. This strategy allows him to easily elucidate knowledge about the current problem, some of which is relevant to his solution. The forward chainer in the previous DC-circuit experiment employed this same strategy. However, the search space in that experiment was smaller due to the smaller number of operators (yielding smaller breadth) and the shorter solution (yielding smaller depth). As a result, this strategy was sufficient for that subject to solve the problem. WB, on the other hand, is still missing some crucial knowledge at this point in the protocol. Also, when WB uses a buggy source in this segment, he produces a buggy target solution (Lines 4 and 5, Figure 2). This too replicates a result from our previous study where Subject One used internal analogy to transfer a buggy equation to a new situation.

At the beginning of the intermediate segment, WB reconsiders using the CTCP schema from Line 1 in Figure 2. Like the forward chainer in the previous study, WB also verbalizes the information content as he says, "I do have two angles in each of these triangles." At this point, he realizes that the target's current information content is less deviant to that of its source than it was when he considered using CTCP in the opening segment. As a result, he spends most of the intermediate segment trying to use this strategy and eventually fails. The bulk of the segment is characterized by large amounts of exhaustive search through the entire set of operators.

At the beginning of the final path segment (Figure 4), WB posts the goal <SHW = <HSW and backward chains using the transitive property. As a result, he produces the analogical behavior in Line 9 of Figure 4, reestablishing the erroneous knowledge that was originally created in Line 5 of Figure 2 and later deleted by the experimenter. In this step, he chooses AIAC-av as a source from the opening segment and instantiates it with HL//CE, SK, <HWS, <EWK. This instantiation is chosen because it involves the same transversal as the source, SK, and, thus, has a low deviance level. Now, however, its deviance is even lower because one of the goal angles driving the backward chaining, <HSW, is an exact match.

Next, internal analogy explains WB's behavior in Line 10 of Figure 4 for several reasons. A different goal angle, <SHW, is an exact match this time, adding zero to the deviance measure for this instantiation. Unlike the internal

analogies of the opening segment, the parallel lines and the transversal of this internal analogy are now only a type match. Originally, this mismatch caused the total deviation to be too high, making this instantiation unacceptable. (i.e. WB never took this step in the opening segment, even though he employed internal analogy many times.) However, after the large amount of exhaustive search in the intermediate segment, WB's threshold allows for more deviation and accepts the instantiation in Line 10. Line 11 completes the subject's buggy solution, which resulted from his buggy AIAC rule. And finally, in Line 12 the experimenter asks WB to fill in the transitivity step he left implicit in his solution.

Conclusion and Future Work

In this paper, we have first shown how a powerful adaptive partial match can be integrated into the internal analogy model by extending the basic mechanisms already provided. This extended model also differentiates to a finer granularity the general notion of exact match as previously defined by macro-operators and chunking. Second, the extensions have proven useful for explaining the search control of a complex protocol in a complicated domain. They were sufficiently general to model both the exploratory forward chaining of the subject's opening segment and the efficient goal-directed backward chaining of his final path segment. In the future, we plan to implement and demonstrate the computational efficiency afforded by the new extensions.

References

- Anderson, J. R. (1983). Acquisition of proof skills in geometry. In *Machine Learning, An Artificial Intelligence Approach*, pages 191–220. Morgan Kaufmann, San Mateo, CA.
- Carbonell, J. G. & Veloso, M. M. (1988). Integrating derivational analogy into a general problem solving architecture. In *Proceedings of the First Workshop on Case-Based Reasoning*, pages 104–124. Clearwater Beach, FL. Morgan Kaufmann.
- Faries, J. M. & Reiser, B. J. (1988). Access and use of previous solutions in a problem solving situation. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, Montreal, Canada. Lawrence Erlbaum Associates.
- Gentner, D. & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10:277–300.
- Hickman, A. K. & Larkin, J. H. (1990). Internal analogy: A model of transfer within problems. In *The 12th Annual Conference of The Cognitive Science Society*, Cambridge, MA. Lawrence Erlbaum.
- Hickman, A. K., Shell, P., & Carbonell, J. G. (1991). Internal analogy: Reducing search during problem solving. In Copetas, C., editor, *The Computer Science Research Review 1991*. The School of Computer Science, Carnegie Mellon University.
- Jones, R. (1989). *A Model of Retrieval in Problem Solving*. PhD thesis, Department of Information and Computer Science, University of California, Irvine.
- Laird, J. E., R. P. S. & Newell, A. (1986). Chunking in soar: The anatomy of a general learning mechanism. *Machine Learning*, 1(1):11–46.
- Mitchell, T. M., Keller, R. M., & Kedar-Cabelli, S. T. (1986). Explanation-based learning: A unifying view. *Machine Learning*, 1(1):47–80.
- Veloso, M. M. & Carbonell, J. G. (1991). Variable-precision case retrieval in analogical problem solving. In *Proceedings of the Third Workshop on Case-Based Reasoning*, Washington, D. C. Morgan Kaufmann.