Title

Schema Acquisition from One Example: Psychological Evidence for Explanation-Based Learning

Permalink

Journal

Authors

Ahn, Woo-Kyoung
Mooney, Raymond J.
Brewer, William F.
et al.

Publication Date

Peer reviewed

Schema Acquisition from One Example:

Psychological Evidence for Explanation-Based Learning[*]

Woo-Kyoung Ahn, Department of Psychology

Raymond J. Mooney, Coordinated Science Laboratory

William F. Brewer, Department of Psychology

Gerald F. DeJong, Coordinated Science Laboratory

University of Illinois at Urbana-Champaign

Abstract

Recent explanation-based learning (EBL) models in AI allow a computer program to learn a schema by analyzing a single example. For example, GENESIS is an EBL system which learns a plan schema from a single specific instance presented in a narrative. Previous learning models in both AI and psychology have required multiple examples. This paper presents experimental evidence that people can learn a plan schema from a single narrative and that the learned schema agrees with that predicted by EBL. This evidence suggests that GENESIS, originally constructed as a machine learning system, can be interpreted as a psychological model of learning a complex schema from a single example.

Introduction

Recent explanation-based models in machine learning (DeJong & Mooney, 1986; Mitchell, Keller, & Kedar-Cabelli, 1986) allow a program to learn a concept or schema by analyzing the causal structure of a single example. Explanation-based learning (EBL) systems construct an explanation for why an instance is a member of a concept or why a particular sequence of actions achieves a goal. This explanation is then generalized, retaining only the constraints required to maintain its causal structure. (Mooney & Bennett, 1986) reviews and compares a number of similar algorithms for performing this generalization. These algorithms produce a general concept description or plan schema which can be used to improve performance on future classification, understanding, or problem solving tasks.

A major difference between EBL and other approaches to learning is the number of examples required. Similarity-based learning (Michalski, 1983; Mitchell, 1978; Quinlan, 1986) requires many examples and systems based on analogy (Carbonell, 1983; Falkenhainer, Forbus, & Gentner, 1986; Winston, 1980) require two examples, while EBL requires only a single example. Although a number of psychological experiments exist demonstrating people's ability to learn concepts or schemata from two examples using analogy (Gick & Holyoak, 1983; Spencer & Weisberg, 1986)or from many examples using similarity-based induction (Medin, Wattenmaker, & Michalski, 1986; Posner & Keele, 1968), there seems to be no experiments directed at demonstrating people's ability to learn a concept or schema from a single example. Consequently, until now, there has been no empirical evidence to support the use of EBL as a psychological model of human learning. This paper summarizes a number of recently conducted experiments which demonstrate subjects' ability to learn a new general plan schema from a single narrative describing a specific instance of the plan. The schema acquired in this way is shown to obey the variables and constraints predicted by an EBL model.

Learning from One Example in GENESIS

The idea of learning a schema by analyzing the explanation of a single narrative was first presented in (DeJong, 1981). The GENESIS system (Mooney, 1985; Mooney & DeJong, 1985) is a realization of this idea which processes short English narratives and is able to acquire new plan schemata from single specific instances.

During the understanding process, GENESIS attempts to construct explanations for characters' actions in terms of the goals their actions were meant to achieve. This process involves plan-based understanding mechanisms like those employed by previous narrative processing systems (Dyer, 1983; Schank & Riesbeck, 1981; Wilensky, 1983). When the system observes that a character has achieved an interesting goal in a novel way, it generalizes the composition of actions the character used to achieve this goal into a new schema. The generalization process (described in Mooney and Bennett (1986)) consists of an analysis of the causal model of the narrative which removes unnecessary details while maintaining the validity of the explanation. The resulting generalized set of actions is then stored as a new schema and used by the system to correctly process narratives which were previously beyond its capabilities. Currently, GENESIS has learned schemata for kidnapping-for-ransom, arson-for-insurance, murder-for-inheritance, and for a police-officer impersonating a prostitute in order to entrap solicitors. In each of these cases, it demonstrates a performance improvement by using the schema it has learned to construct explanations for narratives which it previously could not explain.

The goal of the present research is to show that GENESIS can be interpreted as a psychological model. Thus, we have carried out a series of experiments to see if people can acquire a novel plan schema from a narrative describing a single specific instance of a novel action. Specifically, we predicted that people would build a causally complete representation of the text by causally connecting instantiations of existing schemata (as in Johnson-Laird's *mental model* (1983) or van Dijk & Kintsch's *situation model* (1983)). The *explanation* is the connected portion of the model which contributes to the characters achieving important goals. This explanation is immediately generalized into a schema by changing constants to variables within the constraint that the structure of the explanation remains intact. The resulting schema is characterized by a set of *variables* which are slots which can be filled by different objects or agents in each instance, and a set of *constraints* which specify necessary properties of variables and necessary relationships between variables. The *constraints* are those properties and relations required to maintain the causal validity of the explanation.

Overview of the Experiments

GENESIS was originally constructed as a machine learning system. It was not explicitly written to model an existing set of psychological data. In order to explore the validity of GENESIS as a psychological model of human learning, four psychological experiments were conducted. The basic design was to have subjects read a passage describing a specific instance of a novel plan. Each of the first three experiments used a different task to test whether or not subjects had acquired an abstract schema from this single example. The last experiment tested whether subjects generalized the narrative as a natural part of comprehension or only produced the abstract schema when they were asked general questions about the narrative.

Three passages were constructed to present situations for which the subjects presumably did not already have a pre-established schema but which they could understand using aspects of their existing knowledge. For example, one passage involves a cooperative buying scheme used in other countries. In Korea the system is called a "Kyeah" and in India it is called a "chit fund". The experimental narrative describing a single instance of this plan follows:

> Tom, Sue, Jane, and Joe were all friends and each wanted to make a large purchase as soon as possible. Tom wanted a VCR, Sue wanted a microwave, Joe wanted a car stereo, and Jane wanted a compact disk player. However, after paying their expenses, they each only had $60 left at the end of every third month. Tom, Sue, Jane, and Joe all got together to solve the problem. They made four

slips of paper with the numbers 1,2,3, and 4 written on them. They put them in a hat and each drew out one slip. Jane got the slip with the 4 written on it, and said, "Oh darn, I have to wait to get my CD player." Joe got the slip with the 1 written on it and said, "Great, I can get my car stereo right away!" Sue got the number 2, and Tom got number 3. In February, they each contributed the $60 they had left. Joe took the whole $240 and bought a Pioneer car stereo at K's Merchandise. In May, they each contributed their money again. This time, Sue used the $240 to buy a Sharp 600 watt 1.5 cubic foot microwave at K-mart. In August, all four again contributed $60. Tom took the money and bought a Sanyo Beta VCR with wired remote at Service Merchandise. In November, Jane got the money and bought a Technics CD player at Apple Tree Stereo.

The complexity of the experimental examples prevented a complete formal analysis and computer implementation. Therefore, constraints and variables were obtained for each schema by determining a set of roles for the schema and deciding which properties and relations of these roles were important in maintaining the underlying causal structure of the narrative. Variables and constraints were determined for each schema before conducting any of the experiments. Table 1 shows the list of variables and constraints identified for the Kyeah schema.

In addition to a group given specific narratives (Example group), experiments 2 and 3 also used a control group which was given abstract descriptions of the schemata underlying each of the example narratives (Abstract group). The description of the Kyeah schema given to the Abstract group follows:

> Suppose there are a number of people (let the number be n) each of whom wants to make a large purchase but does not have enough cash on hand. They can cooperate to solve this problem by each donating an equal small amount of money to a common fund on a regular basis. (Let the amount donated by each member be m.) They meet at regular intervals to collect everyone's money. Each time money is collected, one member of the group is given all the money collected (n X m) and then with that money he or she can purchase what he or she wants. In order to be fair, the order in which people are given the money is determined randomly. The first person in the random ordering is therefore able to purchase their desired item immediately instead of having to wait until they save the needed amount of money. Although the last person does not get to buy their item early, this individual is no worse off than they would have been if they waited until they saved the money by themselves.

Since subjects in the Abstract group had been directly told the content of the schema, they were presumed to have learned the schema. Consequently, if the Example group performed as well as the Abstract group on a task requiring knowledge of the general schema, then it is reasonable to assume that the subjects in the Example group had also acquired the schema.

In addition to the Kyeah schema, the other two situations used in the experiments included a technique for making additional money by fencing copies of a stolen collectable and a confidence game called the "phony bank-examiner ploy" (the latter was taken from Wharton (1967)). None of the experiments found appreciable differences among the three schemata; consequently, the results reported for each experiment are averaged across all three examples.

Table 1: Variables and Constraints for the Kyeah Schema

Variables	Constraints
identity of participants	participants want items of similar value
number of participants (n)	participants cannot afford items
exact time of meetings	participants trust each other
interval between meetings (ι)	participants can afford m each ι
amount of donation (m)	each participant donates same amount
items bought	cost of desired items $\simeq$ n X m
stores where items bought	number of meetings = number of participants
method of determining order	order must be assigned randomly

In all of the experiments, the Example group was given only a single instance of the schema. Thus any learning that occurred in this group would be outside the domain of learning theories which required multiple instances (i.e. analogy or similarity-based induction).

In order to determine whether a learned schema agreed with that predicted by EBL, subjects' learning was always judged on how well they obeyed the *constraints* and recognized the mutability of the *variables*.

Subjects in all of the experiments were undergraduate students at the University of Illinois at Urbana-Champaign and participated in the experiments to fulfill a requirement of an introductory psychology course. After each of the experiments, subjects were asked whether they had previously heard of any of the three plans described in the passages. If so, their data were discarded.

Experiment 1: General Description Generation

In Experiment 1, in order to test whether subjects had acquired a schema from a single instance, we asked them to write a general description of the schema. We predicted that in the subjects' descriptions the schema constraints would remain but the actual objects in the story would be replaced with variables.

Method

The experiment used only one group of subjects. Each subject was given the three experimental narratives and was told for each one to "write, in abstract terms, a description of the general technique illustrated in the narrative." In order to clarify the instructions, they were given a sample narrative and its corresponding general description. We included this demonstration narrative to show subjects which level of abstraction we wanted them to generate. The demonstration narrative was about skyjacking and was selected to be unrelated to the experimental passages. The demonstration passages do not reveal which portions of the narratives are constraints and which are variables. For example, an airplane is mentioned in both the demonstration narrative and its corresponding general description; however, although a VCR is mentioned in the Kyeah narrative, it is not a part of the Kyeah schema. A correct analysis can only be determined by providing an explanation for the individual example.

After reading the instructions, all of the subjects read the first narrative and wrote a general description for it and did the same thing for the second and third story at their own pace. Data were collected from 11 subjects.

Results and Discussion

In general, subjects produced good schema descriptions. The following is the description of the Kyeah schema written by one subject:

> Suppose in a group of people, each person would like to buy something expensive, but over a period of time, each person cannot earn enough to buy what he would like. By using random selection, each person could be assigned a number, when the group had saved enough money *together* to purchase an item, the person with the first number would get his item. This would continue for the rest of the group until everyone had gotten what he wished.

To provide a more objective index of the subjects' performance, we counted the number of constraints mentioned and the number of variables identified in their general descriptions. A variable was considered to have been *identified* if either an abstract term, such as "group" or "something," was used to refer to it, or if it was simply not mentioned in the description. If a particular variable is not mentioned at all, then it is reasonable to assume that the subjects believe that its particular value was not important to the overall schema.

On the average, 84% of the constraints were explicitly mentioned and 88% of the variables were identified. These percentages indicate that subjects can acquire an abstract schema from a single instance and that the characteristics of the learned schema agrees with those predicted by EBL theory.

Experiment 2: Story Generation

In Experiment 2. we tested whether the schema that subjects acquired could deal with new instances. If both Example and Abstract groups could produce another instance equally well based on what they read, it would indicate that both groups had acquired a generative schema.

Method

Subjects in the Example group were given the three narratives and told for each one to "write another story in which characters use the general method illustrated in the story but that is otherwise as different as possible." Subjects in the Abstract group were given the three abstract schema descriptions and told for each one to "write a story in which particular individuals use the technique described in the passage in a specific case." Both groups read their first passage and wrote a story and the task was repeated for the second and the third passages. There were eight subjects in the Example group and seven in the Abstract group.

Results and Discussion

In general, both groups produced equally good narratives. The following is the new Kyeah narrative written by one subject in the example condition.

> Bill, Kim, John and Mary were all business associates. Bill wanted some land in Northern Illinois, Kim wanted a new house in Switzerland, John wanted a new Porshe 928S with all all accesories, and Mary wanted to take a trip around the world. The only problem was they each only had $25,000.00 left unspent at the end of each month. They all got together and picked random variables on Bill's business computer. Mary was farthest from her variable so she would have to wait till last to get her trip around the world. John nailed his variable and jumped enthusiastically saying, "Yea, I get to get my new Porshe 928S right now." They each talked with their banker and drew the $25 Thousand dollars out and pooled it together after the first month and the next day John drove up in this new, black, 928S with all accessories. At the end of the next month they again pooled their money and Kim got her chalet in Switzerland. Again at the end of the next month they pooled their money an Bill got his land in Northern Illinois. Finally, after the fourth month they pooled their money together and Mary left for her trip around the world.

Again, to be more objective, we counted the number of constraints obeyed and the number of variables changed in the stories generated by the Example group. However, in the Abstract group, we could only count the number of constraints obeyed since there were no constants to change in their passages (e.g. number of participants, items purchased, etc.).

Overall, the Example group obeyed 90% of the constraints while the Abstract group obeyed 87%. In addition, the Example group changed 71% of the variables. These results imply that from a single specific instance subjects can acquire a schema equivalent to that acquired directly from an abstract description of the schema.

Experiment 3: Yes/No Questions

Experiment 3 was designed to test whether subjects in both groups were equally good at detecting which portions of the narratives were variables and which were constraints. Although in Experiment 2. the Example group changed only 71% of the variables. this does not necessarily indicate that they did not identify the rest of them. It is possible that the subjects simply did not take the effort to change the values of all of the variables. For example. most of the subjects did not change the number of participants in their Kyeah narratives. Experiment 3 directly tested the ability of subjects to recognize all of the variables.

Method

We developed a yes/no question for each constraint (for example. "Can some people consistently donate less than others and have the system work?") and a yes/no question for each variable (for example. "Is there any particular number of people required for this plan?"). For each question the expected answers based on EBL were sometimes "yes" and sometimes "no". Both the Example and Abstract groups read their first passage, answered the same questions with "yes" or

"no", and justified their answers. Then they did the same thing for the second and the third passage at their own pace. There were ten subjects in the Example group and seven in the Abstract group.

Results and Discussion

The data supported our prediction that the Example group would perform as well as the Abstract group. The average percent correct for the Example group was 83% while that for the Abstract group was 81%. For constraints, the scores were 84% and 81% for the Example and Abstract groups, respectively. For variables, the scores were 80% and 81%, respectively.

Also, we examined the subjects' justifications for incorrect answers and found that most of their "errors" were not due to the subjects' failure to generalize in an explanation-based manner, but were due to the subjects' generating a schema slightly different from the one we had attempted to embody in the text. Some of the questions made certain underlying assumptions about the execution of the plan which could be relaxed to generate an even more general schema. Among those answers marked as incorrect, 87% of the Example group's justifications and 79% of the Abstract group's justifications presented arguments which were based on a causally consistent interpretation of the schema. For example, given the following question: "In the above plan, is it necessary that the number of meetings be the same as the number of people in the group?", one subject responded: "No, it's irrelevant. They could collect money every week and then at the end of the month the one person gets it all." An example of a causally *inconsistent* justification is when a subject was asked: "Is there any particular number of people required for this plan?" and responded: "Yes. Four is the only number of people that will make this plan work."

Experiment 4: Memory Test

Experiments 1-3 indicated that people could acquire a schema by generalizing the explanation of a single example; however, they did not indicate *when* generalization occurred. The subjects might have performed schema abstraction at the time they read the passage or only later when asked questions about it. For example, in Experiment 3, subjects in the Example group might have answered the questions by storing specific representations of the narratives in memory and then generalizing these representations after they were asked the questions. If this is the case, they should also be able to answer questions about *specific* facts in the narratives as well as questions about the *general* schemata.

Method

We tested subjects' memory for specific and general information one day after they read an experimental narrative. To test general information, we used the same questions from Experiment 3, including both questions on constraints and on variables. To test specific information, we developed new yes/no questions on each constraint and variable instantiated in the example. For example, to test the variable, "number of participants", we asked, "Were there five people in the group described?".

Subjects read only one of the three narratives and were asked to rate the quality and usefulness of the plan. They then left without knowing that there would be more tests related to what they had read. After one day, subjects returned expecting another experiment. Instead, each subject received questions on the narratives they had read previously and were asked to answer them with "yes" or "no" at their own pace. There were 17 subjects in total.

Results and Discussion

The percentage correct for general questions was 84% whereas that for specific questions was only 60%. These results indicate that subjects' responses to general questions were based on their general schemata not on specific representations of the narratives. This is because the hypothesis that the subjects were using a specific representation to answer the questions requires that the subject retrieve the specific information relevant for each abstract question. Also, since they were not aware that there would be a comprehension test one day later, these data provide evidence that the

general schemata were a natural product of the comprehension process.

Conclusions

In general, the experiments described above support EBL as a viable psychological model of certain types of human learning. Specifically, they demonstrate that, like GENESIS, people can learn a schema by generalizing the explanation of a single narrative.

Previous research has generally assumed that multiple examples are required for schema or concept acquisition. For example, Rumelhart and Norman (1978) claimed that there are basically two ways in which schemata can be formed. These were *pattern generation* and *schema induction* both of which require exposure to multiple examples. Brewer and Nakamura (1984), postulating schemata as containing abstract *generic* knowledge, also assumed that schema acquisition required multiple examples.

Gick and Holyoak (1983) even specifically showed that subjects could *not* learn a particular schema (the "convergence schema") from a single example and that acquiring the schema required two analogous examples. However, the convergence schema is a very abstract concept which a standard EBL mechanism could not acquire from one example. Explanation-based generalization as performed by GENESIS involves retaining the basic structure of the plan used in a specific example and removing actions, properties, and relations which are clearly irrelevant to achieving the goal. Given a story like "The General" from Gick and Holyoak (1983), such a process would acquire a schema for capturing an enemy fortress by attacking it simultaneously from all sides; however, it could not acquire an extremely abstract concept like the convergence schema. Consequently, the results presented here do not contradict the specific results or conclusions presented in Gick and Holyoak (1983). However, our results do show that there is a large class of schemata that *can* be acquired from one example.

Although explanation-based learning can be used to learn a schema from only one example, it is unsuitable for learning certain classes of concepts. EBL is only applicable when one has sufficient knowledge of the domain and the schema to be learned is solely determined by causal constraints. In this case, people use their existing knowledge of the domain to guide the schema acquisition process and distinguish relevant from irrelevant features after given only one example. This improves the efficiency of learning and results in schemata which are free from spurious correlations. However, many schemata, such as that for a wedding ceremony or a birthday party, are determined by un-explainable social conventions as well as by necessary causal relationships among their constituent actions. Learning such schemata efficiently will require a mechanism which successfully integrates the different approaches underlying current learning mechanisms.

Nevertheless, EBL is one of the first attempts to incorporate significant amounts of causal and explanatory domain knowledge into a learning mechanism. Murphy and Medin (1985) argue persuasively for the importance of "theories" in concept formation, where important features of a "theory" include "An explanatory principle common to category members" and a "Network formed by causal and explanatory links." Investigating explanation-based learning from a single example represents an important step in understanding how "theories" can be successfully employed in concept acquisition.

References

Brewer, W. F. & Nakamura, G. V., (1984). The nature and functions of Schemas. In R. S. Weyer & T. K. Srull (Eds.), *Handbook of Social Cognition: Volume 1*, Lawrence Erlbaum and Associates, Hillsdale, NJ, 119-160.

Carbonell, J. G. (1983). Learning by analogy: Formulating and generalizing plans from past experience. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine Learning : An Artificial Intelligence Approach*. Tioga, Palo Alto, CA, 137-162.

DeJong, G. F. (1981). Generalizations based on explanations. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, B.C., Canada, 67-70.

DeJong, G. F. & Mooney, R. J. (1986). Explanation-based learning: An alternative view. *Machine Learning 1*, 2.

Dyer, M. J. (1983). *In-Depth Understanding*, MIT Press, Cambridge, MA.

Falkenhainer, B., Forbus, K., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 272-276.

Gick, M. L. & Holyoak, K. L. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1-38.

Johnson-Laird, P. N. (1983). *Mental Models*, Harvard University Press, Cambridge, MA.

Medin, D. L., Watenmaker, W. D., & Michalski, R. S. (1986). *Constraints and Preferences in Inductive Learning: An Experimental Study of Human and Machine Performance*. (Technical Report ISG 86-1 UIUCDCS-F-86-952). Department of Computer Science, University of Illinois, Urbana, IL.

Michalski, R. S. (1983). A theory and methodology of inductive learning. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*. Tioga Publishing Company, Palo Alto, CA, 83-134.

Mitchell, T. M. (1978). *Version Spaces: An Approach to Concept Learning*. (Technical Report STAN-CS-78-711). Stanford University, Palo Alto, CA.

Mitchell, T. M., Keller, R., & Kedar-Cabelli, S. (1986). Explanation-based generalization: A unifying view. *Machine Learning 1*, 1, January, 47-80.

Mooney, R. J. (1985). *Generalizing Explanations of Narratives into Schemata*, Masters Thesis, Department of Computer Science, University of Illinois, Urbana, IL, May.

Mooney, R. J. & Bennett, S. (1986). A domain independent explanation-based generalizer. *Proceedings of the National Conference on Artificial Intelligence*. Philadelphia, PA, 551-555.

Mooney, R. J. & DeJong, G. F. (1985). Learning schemata for natural language processing. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*. Los Angeles, CA, 681-687.

Murphy, G. L. & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 3, July, 289-316.

Posner, M. J. & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology, 77*, 3, July, 353-363.

Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning 1*, 1, 81-106.

Rumelhart, D. E. & Norman, D. A. (1978). Accretion, tuning and restructuring: Three modes of learning. In J. W. Cotton & R. L. Klatzky (Eds.), *Semantic Factors in Cognition*. Lawrence Erlbaum and Associates, Hillsdale, NJ.

Schank, R. C. & Riesbeck, C. (1981). *Inside Computer Understanding*, Lawrence Earlbaum and Associates, Hillsdale, NJ.

Spencer, R. M. & Weisberg, R. W. (1986). Context-dependent effects on analogical transger. *Memory & Cognition, 14*, 5, 442-449.

van Dijk, T. A. & Kintsch, W. (1983). *Strategies of Discourse Comprehension*, Academic Press, New York.

Wharton, D. (1967). Five common frauds, and how to avoid them. *Reader's Digest*, December.

Wilensky, R. W. (1983). *Planning and Understanding: A Computational Approach to Human Reading*. Addison-Wesley, Reading, M.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the Association for Computing Machinery, 23*, 12, 689-703.