

The Role of Experiences and Examples in Learning Systems

Edwina L. Rissland+
Oliver G. Selfridge
Elliot M. Soloway*

Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003

Abstract

In this paper, we discuss the role of experiences and examples in learning systems. We discuss these issues in the context of three systems in particular: Rissland and Soloway's Constrained Example Generation (CEG) System, Selfridge's COUNT, and Soloway's BASEBALL.

1. Introduction

Examples and experiences, by which we mean concrete instances, situations or problems, are critical to any system, man or machine, that learns. Examples provide the basis from which generalizations, concepts and conjectures are made. They also provide the criticisms needed to refute and refine.

For instance, in Winston's learning program [Winston 1975], examples of the concept to be learned, e.g., an arch, and non-examples, e.g., "near misses", are the critical input from which his program builds a structural description of a concept. In Lenat's concept discovery program, AM, [Lenat 1977], examples help direct the discovery process, by providing evidence of the reasonableness and interestingness of new concepts.

+Supported in part by the National Science Foundation under grant IST-80-17343.

*Supported in part by the Army Research Institute for the Behavioral and Social Sciences under ARI grant MDA903-80-C-0508. Opinions expressed in this report are those of the authors, and do not necessarily reflect the views of the U.S. Government.

Examples are critical in human learning and discovery, whether it be in children [Hawkins 1980] or sophisticated adults. Lakatos [1976] gives a detailed exposition of the historically important Euler's formula; there, examples like various "monsters" (i.e., counter-examples) play a central role in concept refinement.

Thus, examples are grist for the learning process in a critical way. In this paper, we consider the role of examples in learning systems, in particular issues such as the richness of the base of examples upon which the system runs. Questions about the generation of examples are discussed in [Rissland 1980, Rissland and Soloway 1980, 1981].

2. Three Learning Systems

We now restrict our discussion to three learning systems to illustrate some general issues in learning. Briefly, the three systems work as follows:

1. BASEBALL possesses both high level schemas which describe the intentions of people in action-oriented competitive games and low level schemas which provide a common sense understanding of spatio-temporal events. From observations of activity in a baseball game, the system first interprets that activity, generalizes from these hypothesized rules, and finally accepts or rejects these rules based on their predictive utility. Rules for concepts like "out", "hit", "single" are learned.
2. CEG generates an example to meet posted desiderata by modifying known examples from its knowledge base of examples, its "Examples-space". In the version of the system being used to study learning, the system possesses several operators that can modify a given feature; its task is to explore both the space of examples and the space of modification operators not only to arrive at a solution -- a base example plus a sequence of modifications -- but also to gain experience in using the operators in order to allow later learning about the operators themselves.
3. COUNT possesses a repertoire of primitive number and string manipulation routines, such as "increment by 1", "move the pointer right by 1", and control routines "repeat", and "do N times", from which it is to build procedures, i.e., strings of primitives, and ultimately a "count" procedure to count the number of symbols in a string. The system learns by

solving problems posed by its user who acts as its teacher.

3. Examples and Learning

Each of the above three learning systems is provided experiences and examples upon which it bases its learning. The provider of these examples, in effect, acts as its teacher.

BASEBALL: the ensemble of observed games
CEG: the initial Examples-space, posed problems
COUNT: posed problems

Thus, the systems gain experience that ranges from a set of example games, to examples and problems, to just problems.

In each of the systems, there is a classic trade-off between the richness and size of the initial knowledge (not only of examples) and the amount and care of search that must be made for solutions and conclusions. The initial knowledge is used to control the size of the search space. The amount of search in the system varies:

BASEBALL: small
CEG: small-medium
COUNT: medium-large

Within CEG itself, for instance, the richer the initial Examples-space, the less care was needed to explore "adequately" the space of operator sequences. COUNT works with very little embedded knowledge and expends a large effort in search. BASEBALL generates a small number of interpretations and generalizations.

All of the systems make use of evaluation and judgement mechanisms:

BASEBALL: uses a teacher-specified threshold to accept/reject hypotheses based on their predictive utility
CEG: possessed by an explicit JUDGE module
COUNT: performed by the system for the ultimate counting task, and by the teacher for all others

The order of problems can be an important feature of the learning:

BASEBALL: importance of order of observations depends on the threshold setting for accepting hypotheses as true
CEG: order of problems is important when solutions are saved
COUNT: order of problems is very important

In all cases, when there is an ordering of problems, it is the responsibility of the teacher. In COUNT, the order of problems is critical; if COUNT is over-faced with too hard a problem, it will exceed its search tolerance without finding a solution. The art in teaching COUNT is selection of a sequence of problems that challenge it enough to learn things it couldn't do before but not to ask it to make too large a leap in one problem. BASEBALL is sensitive to ordering if the hypothesis acceptance threshold is set low; early acceptance of a mistaken hypothesis can cause difficulties in subsequent interpretation and generalization.

4. General Issues in Learning

Thus, we see that learning systems in addition to requiring various submodules [Smith et al 1977] can be described along several dimensions. Some of the dimensions of this description are:

1. Presence of Teacher: From strongly taught systems such as COUNT, CEG and Winston's to minimally taught systems such as BASEBALL and Lenat's. In the minimally taught systems, the teacher is often implicitly embodied as built-in evaluation functions, focus of attention thresholds, and heuristics.
2. Richness of Experience: From systems that need a rich experience such as COUNT and CEG to those like BASEBALL and Winston's that don't. For the latter systems, experience is often implicitly given to the system in the form of schemas and descriptors for the domain. The former systems can potentially deal with more diversity in their discoveries, although they might not know what to do with it, while the latter have already been focused to interpret their experiences within a given framework or model. For these latter systems, to handle more diversity, say a Roman arch, the system needs to know when to let the model worlds bifurcate.
3. Style of Learning: The styles can range from focussed to exploratory. When one knows something about the domain or general area in which the learning takes place one can be more focussed and directed;

BASEBALL knows a lot about competitive games in general and Winston's program, about the blocks world. They both have access to symbolic descriptions and frameworks like "action" entities and "must/must not" links. At the other end of the spectrum, COUNT is like a tyro just beginning to explore its world; it needs to gather lots of experience with its primitive capabilities. CEG is somewhere intermediate on the tyro-expert learner spectrum; it is able to harness its knowledge somewhat symbolically (by knowing links between examples and relations between procedures and features), but must still do a large amount of exploration.

4. Grainsize of Knowledge: There is a spectrum of knowledge grainsize ranging from atomic primitives in COUNT, to mid-size entities and relations in CEG, to larger chunks in Winston, to large frameworks in BASEBALL.

5. Conclusions

From our own and others' experience with learning systems, it is clear that examples and experiences play a critical role in learning. While the importance of examples in learning is often overlooked, the number, variety and order of examples cannot be since they so clearly influence the style and content of learning.

In addition, while it might be fine for a system to do high-level processing when it knows something, it might be more appropriate to rely on low-level processing (e.g., trial-and-error, success-failure correlations) when it is just beginning. Perhaps, such a low-level style is the only way for the inexperienced learner and perhaps, it is a way for him to discover larger clusters of knowledge.

References

- Hawkins, D., "The View from Below". For the Learning of Mathematics. Volume 1, No. 2, FLM Publishing Association, Quebec, Canada, November 1980.
- Lakatos, I., Proofs and Refutations. Cambridge University Press, London, 1976.
- Lenat, D. B., Automatic Theory Formation in Mathematics, Proceedings Fifth IJCAI, 1977.
- Rissland, E. L., Example Generation. Proceedings Third National Conference of the Canadian Society for Computational Studies of Intelligence, Victoria, B.C., May 1980.
- _____, Understanding Understanding Mathematics. Cognitive Science, Vol. 2, No. 4, 1978.
- Rissland, E. L., and E. M. Soloway, Overview of an Example Generation System. Proceedings First National Conference on Artificial Intelligence, Stanford, August 1980.
- _____, Constrained Example Generation: A Testbed for Studying Issues in Learning. Submitted to Seventh IJCAI, 1981.
- Selfridge, O.G., Learning to Count: How A Computer Might Do It. Bolt Beranek and Newman, Inc, 1979.
- Soloway, E. M., Learning = Interpretation + Generalization: A Case Study in Knowledge-Directed Learning. COINS Technical Report 78-13, University of Massachusetts, 1978.
- Smith, R.G., T. M. Mitchell, R. A. Chestek and B. G. Buchanan, "A Model for Learning Systems". Proceedings Fifth IJCAI, 1977.
- Winston, P. H., "Learning Structural Descriptions from Examples" in The Psychology of Computer Vision, Winston (ed), McGraw Hill, 1975.