

ADJUSTING BIAS IN CONCEPT LEARNING *

Paul E. Utgoff**

Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903

ABSTRACT

We approach concept learning as a heuristic search through a space of concepts for a concept that satisfies the learning task at hand. The heuristics represent bias that the concept learning program employs when forming an inductive generalization. We present a model of bias adjustment and report our experience with an implementation of the model.

I ROLE OF BIAS

A major objective in research on inductive concept learning is creation of a program that can accept training data, apply knowledge, and form inductive hypotheses of the concept, all without human intervention. The learning program searches a space of hypotheses for those that are consistent with the observed examples, and which classify the unobserved instances as indicated by heuristics. The heuristics, which we call *bias*, determine the inductive generalizations that the program will form, given some set of training examples. The initial bias may be appropriate for one learning task, yet inappropriate for another. In many concept learning programs to date, e.g. (Michalski 83), (Mitchell 77), the search for appropriate bias is done by hand. We examine an approach to mechanizing that search.

Vere (Vere 80) and Lenat (Lenat 83) have programs that adjust their bias. Vere uses the set difference operator to construct a new description $C = A - B$ when no other consistent description is available. Lenat's EURISKO learns heuristics that lead the program to discover interesting concepts. EURISKO is an advance from Lenat's AM where the search for appropriate bias was done by hand.

II AN APPROACH TO ADJUSTING BIAS

We represent bias as a restricted search space of concepts that we call the *concept description language*. We use Mitchell's Candidate Elimination Algorithm (Mitchell 77) to maintain a version space of all concept descriptions in the concept description language that are consistent with the training instances. When the trainer supplies a positive instance, all concept descriptions that exclude the instance are refuted and removed from the version space. Similarly, when the trainer supplies a negative instance, all concept descriptions that include the instance are refuted and removed from the version space. As shown in Figure 1, as the version space shrinks, the complementary set of refuted hypotheses grows. If the version space becomes empty, then all descriptions in the concept description language have been refuted. A refuted concept description language indicates a refuted bias, and is sufficient justification for adjusting bias. We

* This work was supported by National Science Foundation Grant GMCS80-08889, Rutgers University Laboratory for Computer Science Research, and Siemens Research and Technology Laboratories.

** Current address is Siemens Research and Technology Laboratories, 105 College Road East, Princeton, New Jersey 08540.

do not yet know stronger justifications. The role of bias adjustment is to enlarge the concept description language so that the version space becomes nonempty and search can continue. It may be that certain concepts are not describable within a particular formalism. If so, shift to a more appropriate formalism may be necessary.

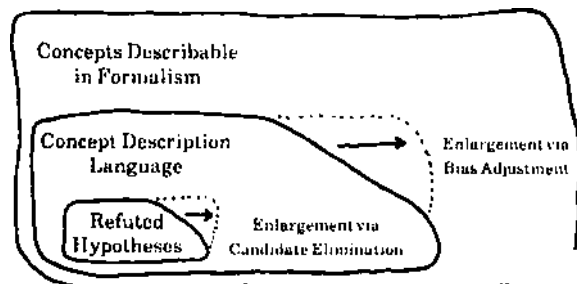


Figure 1 Model of Bias

We approach bias adjustment as a three step process:

1. Determine a membership specification for the new concept.
2. Translate the membership specification into a description that uses the formalism of the concept description language.
3. Assimilate the new concept description into the hierarchy of the concept description language.

A *membership specification* is itself a concept description, but is expressed in an unbiased language so that the heuristics for determining the specification are not compromised by what is expressible. Because the membership specification may be expressed in an alternate formalism, we require the ability to translate the membership specification into a concept description in the formalism of the concept description language. The translation step may require creation of new descriptive terms that correspond to membership specifications that are not describable in the current concept description language. A new concept must then be assimilated into the concept description language by defining its location in the hierarchy of descriptions.

III AN EXAMPLE

We have implemented a program that models bias adjustment as part of the LEX program (Mitchell 83). LEX is a concept learning program that learns heuristics that suggest whether a given operator should be applied to a given problem state. Each operator heuristic is represented as the concept "set of states to which this operator should be applied."

LEX currently learns heuristics in the domain of integral calculus. The program is initially given a set of problem-solving operators. Each operator has a domain of applicability, a rewrite rule, and a range of producible states. For each operator, the legal domain of applicability describes states to which the

Back-propagation of the set of integers k through the exponent in Op 150 caused a new description $\{x|(\exists A(x)(\text{match } x/2 \text{ 'k'}))\}$ to be constructed, the set of even integers. Because the operator could divide any real exponent by 2, the constructed set description was assimilated as a specialization of the set of real numbers. It would have been better to assimilate the set of even




Figure 3 - Solution and propagation for $\int \cos^7(x) \cdot dx$

integers as a specialization of the set of integers. Precise assimilation could have been determined with the following reasoning: An alternate definition of even integers is $\{y|\exists x \in \text{Integers } y = 2 \cdot x\}$. Because x and 2 are integers, and because we know that integers are closed under multiplication, we prove that every y is also an integer. The critical piece of knowledge for the proof is that integers are closed under multiplication. We have not implemented a mechanism for executing such proofs. Proposing and then proving subset relationships is an open problem of great interest. The backward constraint propagation method is useful for deducing which conjectures are worth proving.

2 Experiment #2

In this experiment, the set of real numbers "r" was to be back-propagated through Op22 as shown in Figure 4. The result of the back-propagation step was $r_4 \cdot (r_3/r_4)$. Because this set was describable by an existing set description, it was not necessary to add any new descriptions to the search space.




Figure 4 - Evaluating Product with Op22

Nevertheless, the back-propagation failed because the deduced set $r_4(r_3/r_4)$ did not satisfy a safety check requiring that a deduced set, in this case $r_4 \cdot (r_3/r_4)$, match the state that occurred in the original solution sequence, in this case 5-6. Although r_4 matched 5, r_3/r_4 did not match 6. The concept description language did not include knowledge that r_3/r_4 is equivalent to the set of rational numbers, and that 6 is an instance of a rational number. The inverse operator had inserted an implicit computation; division!

3. Experiment #3

A problem $\int x^2 dx$ was given to LEX, and was solved by Op2 as shown in Figure 5. Other instances were also provided that caused the bias adjustment module to be invoked. The program

needed to back-propagate the set of real numbers "r" through the exponent in Op2¹. The code created a description $\{x|KA(x)(\text{match } x+1 \text{ 'r'})\}$, and assimilated it as a specialization of r. This produced a description with infinite recursion of the form "x is a real number if x + 1 is a real number". If the program could have applied knowledge that the real numbers are closed under addition, in this case by 1, then it would have proven that adding 1 to an x will not alter whether that x is a real number.

It is important to consider not only where to assimilate a new description, but whether to assimilate a new description. To circumvent superfluous descriptions, we added an ad hoc procedure for computing closures; if the description $D_{M,W}$ is intended to be a specialization of an existing description $D_{i,i}$, but $D_{H,W}$ is more general or equal to $D_{0,m}$, then $D_{m,w}$ is superfluous




Figure 5 - Solution and propagation for $\int x^2 \cdot dx$

IV SUMMARY

We suggested that heuristics used for induction should be found by the learning program itself, not by hand. We presented a model of a concept learning program that adjusts its own bias, thereby altering the inductive generalizations it will make. We implemented an example of the model by adding a bias adjustment module to the existing LEX concept learning program. Experience with the implementation uncovered examples where the formalism was weak, or additional knowledge was needed to determine whether or where to assimilate a newly constructed description. Where these problems were not a factor, the program was able to automatically adjust its own bias.

ACKNOWLEDGEMENTS

I thank Tom Mitchell, Rich Keller, and Donna Nagel for discussions of ideas presented here. I thank Jack Mostow and Bruce Ladendorf for helpful comments on a draft.

REFERENCES

[11] Lenat, D. B., "The Role of Heuristics in Learning by Discovery," *Machine Learning*, Michalski, Carbonell, Mitchell (eds.), Tioga, Palo Alto, 1983, pp. 243-306.

[21] Michalski, R. S., "A Theory and Methodology of Inductive Learning," *Machine Learning*, Michalski, Carbonell, Mitchell (eds.), Tioga, Palo Alto, 1983, pp. 83-134.

[31] Mitchell, T. M., "Version Spaces: A Candidate Elimination Approach to Rule Learning," *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, Mass., 1977, pp. 305-310.

[4] Mitchell, T. M., Utgoff, P. E., and Banerji, R. B. "Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics," *Machine Learning*, Michalski, Carbonell, Mitchell (eds.), Palo Alto, Tioga, 1983, pp. 163-190.

[51] Utgoff, P. E. and Mitchell, T. M., "Acquisition of Appropriate Bias for Inductive Concept Learning," *Proceedings of the Second National Conference on Artificial Intelligence*, Pittsburgh, August 1982, pp. 414-417.

[61] Vere, S. A. "Multilevel Counterfactuals for Generalizations of Relational Concepts and Productions," *Artificial Intelligence*, 14:2, September 1980, pp. 138-164.