

## Tradeoffs in the Utility of Learned Knowledge

SMADAR KEDAR  
KATHLEEN MCKUSICK  
ARTIFICIAL INTELLIGENCE RESEARCH BRANCH  
MS 269-2  
NASA AMES RESEARCH CENTER  
MOFFETT FIELD, CA 94035

(NASA-TM-108111) TRADEOFFS IN THE  
UTILITY OF LEARNED KNOWLEDGE  
(NASA) 3 p

N93-16750

Unclass

G3/61 0135329

**NASA** Ames Research Center  
Artificial Intelligence Research Branch

Technical Report FIA-92-22

June, 1992

---

# Tradeoffs in the Utility of Learned Knowledge

---

Smadar T. Kedar and Kathleen B. McKusick

Sterling Software, Inc.

NASA Ames Research Center

AI Research Branch, Mail Stop 269-2

Moffett Field, CA 94035

kedar@ptolemy.arc.nasa.gov, mckusick@ptolemy.arc.nasa.gov

## Abstract

Planning systems which make use of domain theories can produce more accurate plans and achieve more goals as the quality of their domain knowledge improves. MTR, a multi-strategy learning system, was designed to learn from system failures and improve domain knowledge used in planning. However, augmented domain knowledge can decrease planning efficiency. We describe how improved knowledge that becomes expensive to use can be approximated to yield calculated tradeoffs in accuracy and efficiency.

## 1 INTRODUCTION

Successful planning and control systems in realistic domains depend on the ability to improve with experience. One characteristic of such systems is the ability to recover gracefully from failures, and avoid similar failures in the future. The long term objective of our machine learning research (Kedar *et al.*, 1991) is to improve planning and control systems by autonomously and systematically detecting failures, and refining domain knowledge to correct them.

Adding knowledge to a system via machine learning methods is not without consequent cost to the system making use of this knowledge. Recent research in machine learning has begun to address this cost in addition to considering system performance improvement which results from the added knowledge. The notion of a *utility problem* was first presented in (Minton, 1988), to refer to the degradation of system performance by machine learning (specifically Explanation-Based Learning). Holder (1988) generalized this idea to other learning paradigms and performance metrics.

Most approaches to utility analysis focus on a single performance system, a single learning paradigm, and a single measure of utility (e.g. *efficiency* in Minton, 1988; Tambe 1990; or *accuracy* in Holder, 1991). The utility of learned knowledge in more complex integrated systems needs to be measured along several di-

mensions at once. In this paper, we present a case study of a multi-strategy machine learning system, *mutual theory refinement*, which refines knowledge for an integrated reactive system, the Entropy Reduction Engine (Drummond, *et al.*, 1991). We describe a method for trading off two conflicting utility metrics, system accuracy and system efficiency, in order to achieve particular global performance objectives.

## 2 LEARNING IMPROVES PLAN ACCURACY

Our case study is cast within the Entropy Reduction Engine (ERE), a system which integrates planning and scheduling with reaction. ERE uses *operators* to model actions, and *domain constraints* to model physical laws (e.g., "the agent cannot be in two locations at once"). The operators and constraints are only approximate models, and therefore may not always correctly predict the results of actions. Prediction failures drive the learning system, *mutual theory refinement* (MTR) (Kedar *et al.*, 1991), to refine these two world models.

MTR distinguishes itself from other analytic theory refinement methods (e.g. Hammond, 1986; Chien, 1989) in the ability to use an approximate model, rather than a fully correct and complete one, to refine other approximate models. MTR is also unique in its ability to switch from analytic to inductive refinement when the approximate models are insufficient. While reducing prediction failures, the ultimate aim of MTR is to improve the overall performance of the associated system (e.g. ERE). We have demonstrated experimentally that MTR increases the accuracy of the associated ERE system, but does so while degrading its efficiency (Kedar & McKusick, 1992). That is, overall performance involves an accuracy/efficiency tradeoff.

## 3 APPROXIMATION IMPROVES PLANNING EFFICIENCY

Learning in an integrated system needs to promote some global performance objectives, e.g. a certain level




Figure 1: Tradeoffs in Efficiency and Accuracy While Approximating Operator Preconditions.

of system goal achievement given an efficiency constraint. Unfortunately, an augmented domain theory may be too inefficient to use given such a constraint. Our objective here is to show that by approximating the refined theory in an informed manner, we can improve system efficiency while maintaining an acceptable level of accuracy. Through experimentation, we can anticipate how effective a particular approximation is likely to be with respect to the global accuracy and efficiency objectives.

We illustrate this process using data from our case study. We use two methods of approximating our theory: first, to improve efficiency in operator match cost once missing preconditions have been learned, the system approximates certain preconditions by truifying or nullifying them (as in Keller, 1987). Second, to improve efficiency in planning search once multiple outcomes have been learned, the system approximates the operator model by pruning some of the outcomes.

Figure 1 shows accuracy and efficiency results, averaged for a set of 100 test problems, for all the approximate theories generated using the first approximation method. The horizontal axis plots efficiency, as measured in match cost. The vertical axis plots accuracy in terms of percent goal achievement. Each point on the scatter plot represents the average tradeoff yielded by a particular approximated theory. Boundary points, also known as *pareto-optimal* points (Ellman, 1988), are circled. Each point represents a version of the refined knowledge that cannot be improved in one dimension without degradation in the other dimension. A system can attain global objectives if a pareto-optimal point exists which meets or exceeds these objectives.

For example, consider global objectives where desired accuracy on a set of problems is at least 60% goal achievement, with match cost below 700 function calls. We find the pareto-optimal point which best satisfies the global objectives at 67% goal achievement. By explicitly measuring and plotting the tradeoffs for particular approximations, the system is able to identify one yielding a tradeoff that is likely to achieve the performance objectives on new tasks.

## 4 CONCLUDING REMARKS

The goal of approximating refined knowledge is to achieve improvement in one utility dimension without unacceptably degrading another. In different situations different approximations of the same knowledge may be appropriate to satisfy particular performance objectives. We are currently implementing an ERE/MTR performance system monitor that will enable the performance system to *dynamically* approximate the knowledge, sensitive to various performance measures and performance system components. Such an approach could lead to a more flexible system which achieves goals efficiently without having to limit or destructively modify its store of learned knowledge.

## References

- Chien, S. (1989). Using and refining simplifications: Explanation-based learning of plans in intractable domains. *Proc. of the Eleventh IJCAI* (pp. 590-595). Detroit: Morgan Kaufmann.
- Drummond, M., Bresina, J., and Kedar, S. (1991). The entropy reduction engine: integrating planning, scheduling, and control. *AAAI Spring Symp. on Integrated Intelligent Architectures* (pp. 48-53). Palo Alto, CA.
- Ellman, T. (1988). Approximate theory formation: An explanation-based approach. *Proc. of the Seventh Natl. Conf. on AI* (pp. 570-574). St. Paul, MN: AAAI Press.
- Hammond, K. (1986). Learning to anticipate and avoid planning problems through the explanation of failures. *Proc. of the Fifth AAAI Conf.* (pp. 556-560). Philadelphia: Morgan Kaufmann.
- Holder, L. B. (1988). *Maintaining the utility of learned knowledge using model-based adaptive control*. Master's thesis, Dept. of Comp. Sci., University of Illinois, Urbana.
- Holder, L. B. (1991). Selection of learning methods using an adaptive model of knowledge utility. *Proc. of the First Internatl. Workshop on Multistrategy Learning*. (pp. 247-254). Princeton, NJ.
- Kedar, S., Bresina, J., and Dent, L. (1991). The blind leading the blind: Mutual refinement of approximate theories. *Eighth Internatl. Workshop in Machine Learning* (pp. 308-312). Evanston, IL: Morgan Kaufmann.
- Kedar, S. and McKusick, K. (1992). There is No Free Lunch: Tradeoffs in the Utility of Learned Knowledge. (Tech Rep FIA-92-04): NASA Ames Research Center.
- Keller, R. (1987). Defining operationality of explanation-based learning. *Proc. of the Sixth Natl. Conf. on AI* (pp. 482-487). Seattle, WA: AAAI Press.
- Minton, S. (1988). Quantitative results concerning the utility of explanation-based learning. *Proc. of the Seventh Natl. Conf. on AI* (pp. 564-569). St. Paul, MN: Morgan Kaufmann.
- Tambe, M. and Rosenbloom, P. (1990). A framework for investigating production system formulation with polynomially bounded match. *AAAI Spring Symp. on Case-Based Reasoning* (pp. 693-700). Palo Alto, CA.