# Machine Learning Paradigms for Utility-based Data Mining*

Naoki Abe
IBM T. J. Watson Research Center
P.O.Box 218 Yorktown Heights, NY 10598
nabe@us.ibm.com

## ABSTRACT
In this talk, I will describe a number of machine learning paradigms that are relevant to utility-based data mining, and review some key techniques and results in each.

## Categories and Subject Descriptors
I.2.6 [**Artificial Intelligence**]: Learning - Induction; H.2.8 [**Database Management**]: Applications - Data Mining

## General Terms
Algorithms

## Keywords
Machine Learning, Cost-sensitive Learning, Reinforcement Learning, Active Learning, Data Mining

There are a number of ways to introduce *utility* in machine learning, depending on the application scenario. One natural way to introduce utility is in terms of the *cost* assigned to misclassification errors, and this is the so-called *cost sensitive learning* [4]. Another way in which utility can be introduced is by considering the *cost* of data acquisition. This aspect has been rigorously formulated as *Economic Machine Learning* by Provost (c.f. [7].) One paradigm of machine learning that pays special attention to the cost of data acquisition, in addition to the predictive quality of the obtained hypotheses, is *active/query learning* [2]. The standard active learning paradigm assumes, in effect, that acquiring each example is equally costly, but it readily admits generalizations to accout for general cost structure. Another machine learning paradigm, which we might collectively refer to as *active on-line learning* addresses the issue of optimizing the combination, and trade-off, of losses incurred during data acquisition, and those associated with the predictive quality of the final hypothesis. Some examples of learning paradigms that fall within this general class include the classic bandit problem [3] and its generalizations and *associative reinforcement learning* [5, 1]. Theories have been developed on these learning paradigms, which provide learning strategies that come with theoretical guarantee on the total losses, inclusive of the two types of losses. Finally, a comprehensive paradigm of machine learning, which includes all of the ones mentioned so far as special cases, is reinforcement learning. Indeed, some authors have embedded instances of utility-based data mining problems within the MDP framework (e.g. [6]). While the MDP formulation is the most general, it does not necessarily follow that it will be the most effective in practice. When the problem at hand falls into one of the special cases discussed, the theory and methodology in that special case may be the most effective. I hope to draw some examples of real world applications, for which some of these special cases have indeed proved to be satisfactory.



## 1. REFERENCES

[1] N. Abe, A. W. Biermann, and P. M. Long. Reinforcement learning with immediate rewards and linear hypotheses. *Algorithmica*, 37(4):263–293, 2003.

[2] D. Angluin. Queries and concept learning. *Machine Learning*, 2:319–342, 1988.

[3] D. A. Berry and B. Fristedt. *Bandit Problems.* Chapman and Hall, New York, 1985.

[4] C. Elkan. The foundations of cost-sensitive learning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 973–978, Aug. 2001.

[5] L. Kaelbling. Associative reinforcement learning: Functions in k-dnf. *Machine Learning*, 15(3):279–298, 1994.

[6] A. Kapoor and R. Greiner. Reinforcement learning for active model selection. In *Proc. ACM SIGKDD Workshop on Utility-based Data Mining*, 2005.

[7] F. Provost. Economic machine learning. In *Proc. ACM SIGKDD Workshop on Utility-based Data Mining*, 2005.