A Survey of Delayed Action Models for Open-Ended Gain Problems

Roby Joehanes
Kansas State University
324 Nichols Hall
Manhattan, KS 66502–2302
robbyjo@cis.ksu.edu

Abstract

We survey a class of decision problems characterized by open-ended gain for recommender systems and how the delay on actions can affect the gain. We believe that these problems differ from the much surveyed delayed action cost problems in terms of the interactivity and the time frame. We conjecture that the estimated cost of delayed action (ECDA) model in cost problems can be used in this scenario. We survey some existing methods on making the model less dependent on the formulation of utility and probability function. We also expose an open problem on determining the right time for a certain action to be executed.

Background

Timely action is often critical in facing real world challenges. Time-critical contexts are situations where the expected value of an outcome diminishes over time.

When a doctor advises a patient to start exercising to prolong his life expectancy, the doctor uses statistical data to convince him on how many extra years he would gain. Unfortunately, this data often assumes that the patient exercises regularly right away, which may not be the case for procrastinators. It would probably be more convincing if the doctor can present a tailored graph specific to patient explaining how much of his life expectancy be reduced upon the delay. This way, the patient can decide which delay time versus gain tradeoff he desire. This type of graph may be useful to convince patients to quit smoking or stop abusing alcohol.

Similar problems have been studied [Horvitz] on giving decisions in emergency situations. The case used is triaging patients. Here, they assess from trauma experts the time-dependent probability of a patient's survival as a function of delay between the initiation the injury and the receipt of attention at an emergency center, as shown in figure 1. Note that the decrease in survival is not linear with delayed treatment.

We believe that open-ended gain problems differ from the problems above in three ways:

1. They are gain problems, not cost problems.

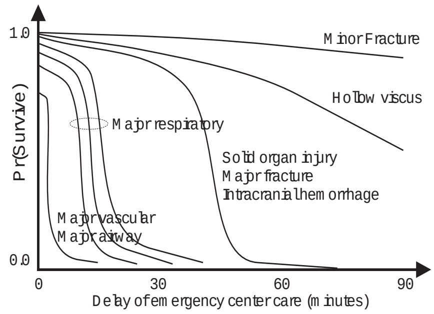

Figure 1: An example of time-dependent probabilities of survival. The value of the probability function decreases over time. Adapted from [Horvitz].

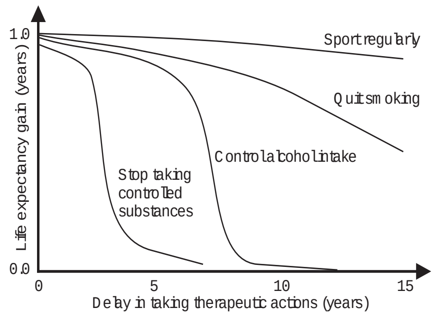

Figure 2: A hypothetical graph on life expectancy gain in delayed therapeutic actions.

2. The time frame may be unbounded in such problems while not only the deadlines but also the time frame addressed in emergency situations are finite.

3. Open-ended frameworks for decision-making such as health maintenance entail interactive characteristics.

The authors in [Horvitz] propose a model to estimate cost of delayed action (ECDA), which has a basic form as follows:

$$ECDA = \max_{A} \sum_{j} u(A_i, H_j, t_0) p(H_j|E, \xi) -$$

$$\max_{A} \sum_{j} u(A_i, H_j, t) p(H_j | E, \xi) \quad (1)$$

Interestingly, although the model is called a "cost" model, the graph plotted in figure 1 can be considered a gain graph over a period of time if we reckon the remaining survival probability as a gain. Intuitively, the ECDA formulation is the utility difference if we take an action now instead of at time t. If we consider the gain to be diminishing, i.e. gain by doing the action now is greater than doing it later, then this formulation also fits for our problem.

Although the time frame of the emergency situations is limited, the formulation of the model is not a limiting restriction. Therefore, this model can be used to address open-ended situations. In order to stop plotting the decision graph for the recommender system, we can impose a limit on the time frame ($t_0 \leq t \leq t_k$, for a reasonable k) or we can stop if the gain is small enough. Alternatively, we can apply a discounting factor to enforce the diminishing outcomes.

Open Questions

The ECDA model depends highly upon the formulation of the utility and the probability functions. However, this is typically unavoidable as these are often domain dependent. For some problems where these formulations are not well-established, such as the health care problem we present earlier, we can apply several approaches:

1. We can learn the function from the available recorded cases.
2. Use some form of discounted version of the utility function if we know that the gain is diminishing.
3. If the function is sparsely defined, we can smooth out the function.

Which of three is best in certain domains is an application dependent issue.

Although the ECDA model is good at estimating the cost of delayed action, it does not address the interval during which the action should be carried out. We can claim that the appropriate time t to carry a certain action A is the one that has the lowest ECDA (or the highest, if it is a gain problem). However, doing so requires computationally expensive scheduling and dynamic planning or replaning step. To make matters worse, we must proliferate the state parameter H_j to consider all possibilities. This can be overcome by restricting the time frame and time granularity, limiting the set of actions and abstracting the states. How to do these the right way is the topic of our intended research.

Future Work

We intend to test our conjectures and conduct empirical studies on how ECDA model could fit the gain model on open-ended framework.

We also try to investigate which abstraction is best applied for gain problems and modify the model formula if necessary.

Acknowledgements

We would like to thank Dr. William Hsu for all of his patience and guidance.

References

Resnick, Paul and Varian, Hal. *Recommender Systems*, introduction to special section of Communications of the ACM, March 1997, vol. 40(3).

Horvitz, Eric and Seiver, Adam. *Time-Critical Action: Representations and Application*, Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, August 1997.

Kaelbling, L. P.; Littman, M. L.; and Moore, A. W. *Reinforcement Learning: A Survey*, Journal of Artificial Intelligence Research 4, 1996.

Norvig, Stuart and Russell, Peter. *Artificial Intelligence A Modern Approach*, Prentice-Hall, Inc., 1995.