Noisy Information Value in Utility-based Decision Making

Clayton T. Morrison
USC Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292, USA
1-310-448-8430
clayton@isi.edu

Paul R. Cohen
USC Information Sciences Institute
4676 Admiralty Way, Suite 1001
Marina del Rey, CA 90292, USA
1-310-448-9342
cohen@isi.edu

ABSTRACT

How much is information worth? In the context of decisions, the value of information is the expected increase in utility of the decision as a result of having the information. When the information might be noisy, the model is slightly more complicated. We develop a model of the value of noisy information in the context of a plausible intelligence information gathering and decision making scenario.

Categories and Subject Descriptors

H.4.2 [**Information Systems Applications**]: Types of Systems – *decision support*; G.3 [**Probability and Statistics**]; I.2.3 [**Artificial Intelligence**]: Deduction and Theorem Proving – *uncertainty and probabilistic reasoning*.

General Terms

Economics, Management, Theory.

Keywords

Value of information, Uncertainty, Noise, Economic Utility, Decision Theory, Intelligence Analysis.

1. INTRODUCTION

In this paper we describe a method for evaluating the value of information with respect to making a decision. The method answers the question, if paying more for information ensures better information, how much should I pay? Or, what is the value of corrupted or noisy information? Our goal is to develop a rational theory of intelligence information gathering as it relates to decision-making, particularly in scenarios where information quality depends on the amount of resources spent. We develop this model within statistical decision theory, which combines probability, statistics and utility theory to provide a coherent framework for evaluating and choosing actions under conditions of uncertainty [20, 13, 15, 11]. Here we consider a class of decision-theoretic models made up of three components: a probabilistic model of the states of the world and their causal relations, decisions that link actions to consequences, and a system of values assigned to those consequences. We review the standard framework for assessing the value of information, develop the model of the value of noisy information, and provide an example of its use in an intelligence information gathering and decision scenario.

UBDM'05, August 21, 2005, Chicago, Illinois, USA.

2. OLYMPIC SECURITY

Consider the following scenario. You are the head of security for the 2004 Olympic Games in Athens, Greece. You have been tasked with protecting the games from potential terrorist attacks. You have the authority to raise a terrorist threat alert that will mobilize special forces and implement extreme counter-terrorism measures. If you raise the alert just prior to an actual attack, the attack will be thwarted and your security firm will be awarded a handsome sum. However, raising the alert is costly, especially if it is a false alarm – it will disrupt the games and your security firm will be held accountable for the resulting loss of potential revenue. On the other hand, not raising the alert prior to an attack will be devastating. The city will be unprepared, lives will be lost and the security firm will be held accountable for a very large sum of money. If no attack is imminent and you do not raise the alert, you are assumed to be doing your job and will be paid as per your contract.

At your disposal you have a team of agents to collect information; each piece of information must be purchased. Suppose one piece of information is the location of terrorist group members, and another is about how prepared they are (in the sense of having the right materials and manpower) to carry out an attack. Based on your experience, you have developed a model of how knowledge of the location or level of preparedness will affect your belief in the likelihood of an attack. You have been provided a fixed budget and it is up to you to manage your information-gathering resources and make the decision about the terror alert. You must decide which information to pay for (if any) while minimizing overall expenditures and still make the right decision about the alert level.

3. A DECISION MODEL

The value of information is always relative to some *target decision*. In the Olympic Security scenario, the target decision is a choice among the actions of raising or not raising the terror alert. The target decision is associated with a *target hypothesis* T, a state of the world that has a direct bearing on the outcome of the target decision. The target hypothesis for our security firm is whether an attack is about to take place. In a probabilistic model, T is a random variable with possible states t, and the available actions a are represented by a decision variable A. In general, the states and actions of T and A are the domain of the variables T and A. A random variable has a probability distribution over its possible states, while a decision variable is assumed to be deterministically controlled by the decision maker.

The outcomes of actions taken in the context of world states may be assigned values or *utilities*, which represent the relative desirability of outcomes. In the Olympic scenario, the outcome of deciding whether to raise the terror alert is represented in terms of money being gained or lost, depending on whether an attack is or is not about to happen. More generally, a utility function $U(a,t)$ maps action a and a target hypothesis state t to a utility value.

Given a target hypothesis, a set of actions and a mapping of target states and actions to utilities, we can frame the target decision problem. Under a widely accepted characterization of rational decision making [20, 15], the optimal decision is to take the action that maximizes the expected utility given beliefs about the target state of the world. Given a distribution over the possible values t of target state T, the expected utility of taking action a is

$$EU(T) = \sum_{t \in T} P(t)U(a,t) \text{ [1]} \quad (1)$$

Some actions may yield a higher expected utility than others. The utility of taking the *optimal* action out of the possible action choices in A is the utility of taking the action that *maximizes* expected utility, expressed as follows:

$$MEU(T) = \max_{a \in A} \sum_{t \in T} P(t)U(a,t) \quad (2)$$

4. UTILITY-BASED VALUE OF INFORMATION

In this section we present the standard approach to assessing the value of information with available utilities. In Section 5 we develop the value of noisy information out of this basic framework.

In many situations, we cannot simply observe the state of the target hypothesis T and make our decision. Instead, we must rely on other states, which are observable, and (we hope) tell us something about the state of the target. The Olympic Security scenario describes two potential sources of information that have some relation to the target hypothesis about whether there is an imminent terrorist attack: the proximity of terrorists to Athens and the capabilities or readiness of the terrorists to commit a terrorist act. These information sources are also about states of the world (e.g., how far away the terrorists are from the Games) and may themselves be represented by random variables (e.g., a distribution over whether the terrorists are within the city limits, in the country, or outside of the country). We use the generic term *indicator variable* for a variable that has some relationship with a target hypothesis. We denote an indicator random variable by I, which represents a distribution over possible states i. We also assume we have a complete joint probability distribution representing the relationship between I and T. This means that we have the information required to derive prior probabilities over I and T, as well as their conditional relationships, $P(I|T)$ and $P(T|I)$. In practice, joint distributions are efficiently represented as Bayesian belief networks, and algorithms exist for effectively deriving and estimating probability distributions from them [15, 11].

Usually the state of our target hypothesis is not directly observable. Instead, we may need to rely on one or more indicators, and determining their state may come at a cost, C_I. In this case we are faced with an *information gathering* decision, which is to be made in the service of our target decision. Out of the set of available indicators, which should we spend resources on? Making this decision requires assessing the value of the information the indicators may provide about the target hypothesis.

The value of any information source is defined as the difference between the utilities of two decision strategies, one in which we choose the optimal action after finding out the state an indicator variable is in, the other choosing the optimal action without that information [10, 13, 15, 8, 11, 12]. The expected utility of taking the optimal action given the outcome i of indicator variable I is

$$MEU(T \mid i) = \max_{a \in A} \sum_{t \in T} P(t \mid i)U(a,t) \quad (3)$$

Since the outcome of I is not known ahead of time, we can calculate the expected utility of having evidence I by marginalizing over the possible values of I:

$$MEU(T \mid I) = \sum_{i \in I} P(i) MEU(T \mid i) \quad (4)$$

$$VOI(T \mid I) = MEU(T \mid I) - MEU(T). \quad (5)$$

Taking into account the cost C_I of acquiring the information about the state of I, the net value or *expected profit* of purchasing the information is

$$netVOI(T \mid I) = VOI(T \mid I) - C_I. \quad (6)$$

If the net value is greater than zero, then the information is worth paying for.

4.1 Myopic Value of Information

Equations 3 through 6 allow us to calculate the value of information about a particular indicator given our current state of knowledge. However, once we consult one source of information, our state of knowledge may change, affecting what we may learn from other information sources, and this in turn affects their value. In general, when considering sequences of information gathering decisions, every permutation of the available information sources must be considered [15, 8, 11]. A *myopic* approximation of information value assesses each information source independently of the others. The myopic assessment is made as if the information source were the only one available, and under the assumption that immediately after gathering the information a final decision is made that incurs some utility [15]. While not perfect, this method has been found to perform well in medical diagnostic systems [7, 9, 14]. Heckerman, Horvitz and Middleton [8] have proposed an *approximate nonmyopic* method for computing information value given certain constraints. In this paper we will consider only myopic value of information.

[1] A note about notation: All of the probability terms in this paper are assumed to be in the context of all currently available evidence. That is, $P(t)$ could be expressed as $P(t|E)$ where E is the set of all other known variables, some of which are known to be in specific states. For clarity, we omit E from our equations. Similarly, all utility calculations are assumed to be in the context of a particular decision variable A and we will only explicitly note A in the context of maximization functions.

5. THE VALUE OF NOISY INFORMATION

We now add another wrinkle to the story, one that has been alluded to [10, 15, 16], but to our knowledge has not received extended treatment, except in very different terms in the economics literature [1, 6]. Suppose the amount you pay affects the quality of information you receive from an information source and the more you pay the more accurate the information is. Now our information gathering decision is to determine which level of payment is optimal for this potentially noisy information. We start by considering paying for reports at a particular cost level and then present the more general formulation of the choice of cost level.

We use R_C to represent a distribution over possible reports about the state of an indicator I at a particular cost level C. We emphasize that what we are paying for at this cost level is not a particular report, but a *distribution* over reports, as the report we receive depends on both the state of I as well as the probabilistic, and therefore noisy, relationship $P(R_C|I)$ we are paying for.

Assessing the expected utility of paying for possible reports R_C about the state of I is no different than the standard Value of Information calculation presented as Equations 3 through 6 in Section 4. Only now we're considering the state of R_C as an indicator of state I, which in turn is an indicator of our target hypothesis T. That is, we could re-write Equations 3 through 6 by replacing R_C for I. Nonetheless, it is useful to highlight the relationship between R_C and I because, again, it is *this* relationship we are paying for.

The following recasts the *VOI* calculation in terms that make the relation between R and I explicit:

$$MEU(T \mid R_C) = \sum_{r \in R_C} P(r) \max_{a \in A} \left(\sum_{i \in I} \sum_{t \in T} P(t \mid i) P(i \mid r) U(a,t) \right) \quad (7)$$

All we have done is add a term that conditions the probability of i on r, and marginalized the effects of r on expected utility by multiplying by and summing over $P(r)$. In other words, Eq. 7 represents the expected maximum utility given that our only source of information is r.

Equations 8 and 9 express the expected benefit and profit of paying for noisy reports at cost C.

$$VONI(T \mid R_C) = MEU(T \mid R_C) - MEU(T) \quad (8)$$

$$netVONI(T \mid R_C) = MEU(T \mid R_C) - MEU(T) - C \quad (9)$$

With $netVONI(T|R_C)$ we can now determine the report distribution R_C at cost level C that yields the highest expected utility:

$$maxVONI(T \mid R_C) = \max_{R_C \in \Re} \left[netVONI(T \mid R_C) \right], \quad (10)$$

where $\Re$ is a set of sources, R, of information about I, each with its own distribution $Pr(I|R)$, distinguished only by how much R costs.

5.1 A Simple Model of Noisy Reports

There are many possible representations of the relationship between I and R. This is a general topic for intelligence analysis and modeling research and will depend on the domain being represented. To demonstrate the value of noisy information in the Olympic Security scenario, we provide a simple linear noise model to generate R_C.

We define a "*noise* level" as a real-valued number between 0.0 and 1.0 (inclusive), where 0.0 means perfect information (no noise) and 1.0 means complete noise. Suppose the reports we are paying for are about an indicator with possible states {*close*, *near*, *far*}. For each possible state of I that report r could say I is in, given that the indicator is actually in state i, we determine the probability $P(R_C=r|i)$ as follows:

$$P(R_C = r \mid i) = \left(\left(\frac{1}{m} - d \right) \times noise \right) + d , \quad (11)$$

where m is the number of possible states of the indicator and $d = 1$ when r reports the same as the target state value i of the indicator, otherwise $d = 0$. We repeat this for each state of I to arrive at $P(R_C|I)$, the conditional probability distribution over reports given the states of I at a particular noise level. Finally, we provide a cost function that maps costs to noise levels, so that given a particular payment C, we can generate $P(R_C|I)$. Under the assumption that information becomes exponentially more expensive with accuracy, we chose the cost function depicted in Figure 1.

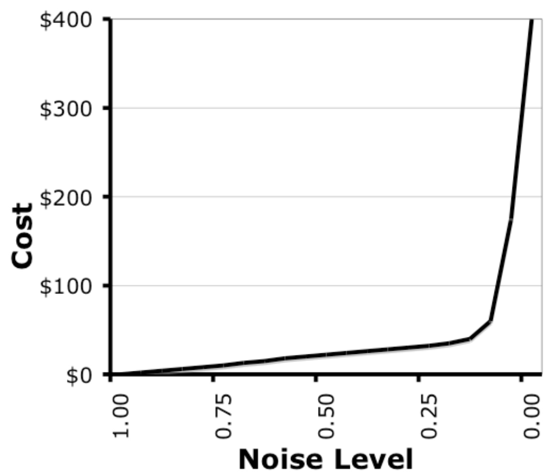

Figure 1. An exponential cost function mapping noise levels to costs of information in dollars.

5.2 Back to the Olympics

Figure 2 shows a *decision graph* representing the Olympic Security scenario. A decision graph is a useful formalism for representing relationships between variables in decision problems. In the graph, decision variables are represented by squares, utility functions by diamonds, and random variables by circles. A directed arrow indicates that the state of a parent node participates in determining the state of a child node (where the child is the node being "pointed to"). The labels in the nodes represent random variables, and we have included text near the nodes indicating which part of the Olympic Security decision problem the variable corresponds to. To complete the specification of the model, we need the prior and conditional probability relationship between random variables as well as a utility function. The tables on either side of the graph in the figure provide this information.

In this scenario, we consider purchasing a distribution over reports about the capabilities of the terrorists. With all of the information represented in Figure 2, we can calculate the value of noisy information of a report about terrorist capabilities at a given level of noise (Eq. 8). Figure 3 does this for noise levels ranging from 1.0 (complete noise) down to 0.0 (no noise). The state of knowledge about proximity affects the value of information, and proximity can be in one of four states (*unobserved, close, near* or *far*). Because of this, Figure 3 plots four different curves representing the value of information across noise levels *given* that proximity is in one of its four possible states. Figure 4 factors in the cost of information for the *net* value of information (Eq 9). The *max* value of noisy information, Equation 10, provides us with a strategy by selecting the cost level at which the expected utility peaks on each curve. We should select the level of payment for a noisy report according to the noise level that yields the greatest expected utility.

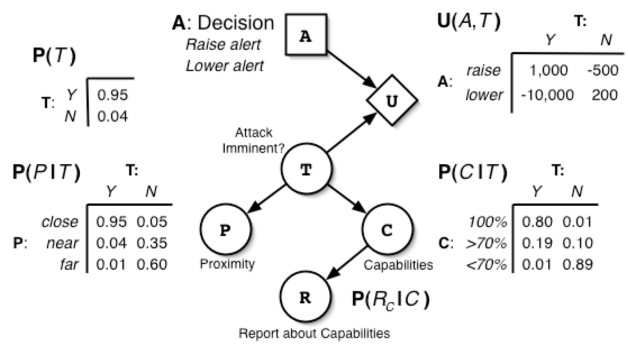

T:	
Y	0.95
N	0.04

U(A,T) A:	T: Y	N
raise	1,000	-500
lower	-10,000	200

P(P|T) P:	T: Y	N
close	0.95	0.05
near	0.04	0.35
far	0.01	0.60

P(C|T) C:	T: Y	N
100%	0.80	0.01
>70%	0.19	0.10
<70%	0.01	0.89

Figure 2. Decision graph and probability and utility tables characterizing the Olympic Security scenario.

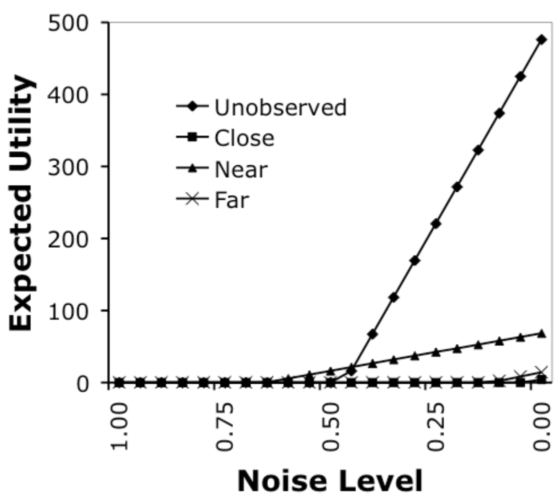

Figure 3. The VONI at varying noise levels. Each curve represents VONI given a state of proximity.

As Figure 4 shows, whether to pay for reports, and if so, how much, depends on our belief about the proximity of terrorists to the Games. When proximity is *unobserved*, paying for a report has the greatest benefit. In particular, the benefit is maximized in the peak of the curve, when noise level is 0.10, costing $60. When proximity is *near*, the benefit of paying for a report peaks at noise level 0.15, with a cost of $40. When proximity is either *close* or *far*, a report is simply not worth paying for at any level. When proximity is known to be *far*, it is likely that an attack is not about to take place; when proximity is *close*, then an attack is almost certain. Under these conditions, paying for more information is simply not worth it. However, as Figure 4 shows, when proximity is *unknown* or *near*, then knowing about the state of the terrorists capabilities is useful in determining whether an attack is about to occur, and paying the price of perfect information is not as cost effective as paying less for somewhat degraded information.

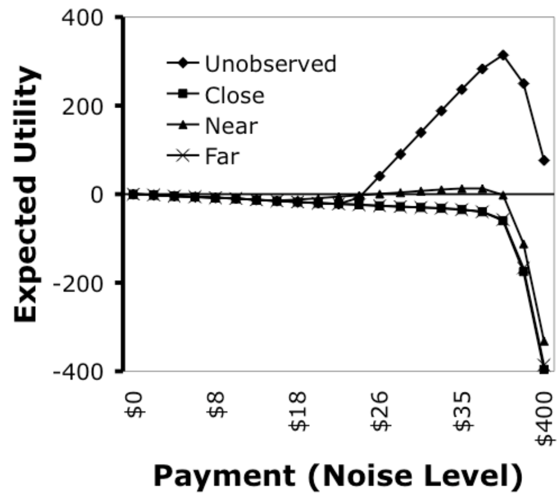

Figure 4. The netVONI at varying noise levels (as represented by their cost) given proximity.

6. CONCLUSION

We have presented a decision model based on the value of information and demonstrated its use in a simple intelligence analysis decision scenario. We demonstrated that with VONI we can determine the optimal amount to pay for information where the amount of effort or cost invested affects information quality. The value of noisy information is an incremental extension to the standard value of information framework, making it possible to assess the value of reports about an indicator that informs a target decision.

As we noted in Section 5.1, the representation of the relationship between and indicator variable I and possible reports R about the indicator is a general topic for intelligence analysis and modeling research. Our simple linear noise model is just one example. Robust models of these relationships will depend on specific scenarios, the expertise of trained analysts, and possibly learned from collected data.

Although we have not presented this framework in terms of machine learning, there are connections between the VONI framework and recent work in cost-sensitive [11, 19] and active [6, 10, 17, 18] learning. In particular, VONI makes explicit the role of data acquisition costs and the impact that acquiring costly information has on decision-making. Recently, [18] explicitly argues for the importance of making active learning decisioncentric, demonstrating in an active learning scenario that simply improving the accuracy of the target classification on which a decision is based does not necessarily lead to overall improvement in decision-making. An important next step in the VONI model is to explore how a model of the noisy relationship between reports and indicators can be learned.

7. ACKNOWLEDGEMENTS

We thank Jafar Adibi and Aram Galstyan for stimulating discussions about portions of the model presented. Work on this project was funded by the Air Force Research Laboratory, account number 53-4540-0588.

8. REFERENCES

[1] Athey, S. C. and Levin, J. D. 2001. The value of information in monotonic decision problems. MIT Department of Economics Working Papers.

[2] Ben-Bassat, M. 1978. Myopic policies in sequential classification. *IEEE Transactions of Computing* 27: 170-174.

[3] Cohn, D. A., Atlas, L. and Ladner, R. E. (1994). Improving Generalization with Active Learning.

[4] Cohn, D. A., Ghahramani, Z. and Jordan, M. I. (1996). Active Learning with Statistical Models. *Journal of Artificial Intelligence Research*, 4: 129-145.

[5] Elkan, C. (2001). The Foundations of Cost-Sensitive Learning. Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01).

[6] Gauthier, L. and Morellec, E. 1997, June. Noisy information and investment decisions: A note. MIT Department of Economics Working Papers.

[7] Gorry, G. and Barnett, G. 1968. Experience with a model of sequential diagnosis. *Computer and Biomedical Research* 1: 490-507.

[8] Heckerman, D, Horvitz, E. and Middleton, B. 1993. An approximate nonmyopic computation for value of information. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3): 292-298.

[9] Heckerman, D., Horvitz, E. and Nathwani. 1992. Toward normative expert systems: Part i. the pathfinder project. *Methods Inform. Med.* 31: 90-105.

[10] Howard, R. A. 1966. Information value theory. *IEEE Transactions on System Science and Cybernetics* 2(1): 22-26.

[11] Jensen, F. 2001. *Bayesian networks and decision graphs*. New York: Springer.

[12] Jensen, F. and Liang, J. 1994. A system for value of information in Bayesian networks. In Proceedings of the 1994 Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, 178-183.

[13] Lindley, D. V. 1971. *Making Decisions*. New York: John Wiley & Sons.

[14] Mussi, S. 2004. Putting value of information theory into practice: a methodology for building sequential decision support systems. *Expert Systems* 21(2): 92-103.

[15] Pearl, J. 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Los Altos, CA: Morgan-Kaufman.

[16] Russell, S. J. and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. New Jersey: Prentice Hall.

[17] Saar-Tsechansky, M. and Provost, F. 2004. Active Sampling for Class Probability Estimation and Ranking. *Machine Learning* 54(2), 153-178.

[18] Saar-Tsechansky, M. and Provost, F. Active Learning for Decision-Making. Working paper (CeDER-04-06, Stern School of Business, New York University) . Available at http://www.mccombs.utexas.edu/faculty/Maytal.Saar-Tsechansky/home/GOAL.pdf

[19] Turney, P. (2000). Types of Cost in Inductive Concept Learning. Proceedings of the Cost-Sensitive Learning Workshop at the 17th ICML-2000 Conference, Stanford, CA. pp.15-21.

[20] Von Neumann, J. and Morgenstein, O. 1953. *Theory of Games and Economic Behavior*. 3rd Ed. New York: John Wiley & Sons.,