Entropy **2007**, *9*, 27-29

entropy
ISSN 1099-4300
www.mdpi.org/entropy/

Short note

An Adaption of the Jaynes Decision Algorithm

Vlad Tarko

Str. Piscului 5, bl. 81, sc. A, ap. 1, sector 4, Bucharest, Romania
E-mail: vladtarko@rdslink.ro

Received: 7 September 2006 / Accepted: 23 January 2007 / Published: 31 January 2007

Abstract: There are two types of decisions: given the estimated state of affairs, one decides to change oneself in a certain way (that is best suited for the given conditions); given what one is, one decides to change the state of affairs in a certain way (that is best suited for what one wants for oneself). Jaynes' approach to decision theory accounts only for the first type of decisions, the case when one is just an observer of the external world and the decision doesn't change the world. However, many decisions involve the wish to transform the external environment. To account for this we need to add an additional step in Jaynes' proposed algorithm.

Keywords: decision theory, loss function, MaxEnt, Bayes

Jaynes [1] proposes the following algorithm for "finding the optimal decision of any kind":

(1) Enumerate the possible states of nature θ_j, discrete or continuous, as the case might be.
(2) Assign prior probabilities $(\theta_j \mid X)$ which maximize the entropy subject to whatever prior information X you have.
(3) Digest any additional evidence E by application of Bayes' theorem, thus obtaining the posterior probabilities $(\theta_j \mid EX)$.
(4) Enumerate the possible decisions D_i.
(5) Specify the loss function $L(D_i, \theta_j)$ that tells what you want to accomplish.
(6) Make that decision D_i which minimizes the expected loss:

$$\langle L \rangle_i = \sum_j L(D_i, \theta_j) \cdot (\theta_j \mid EX)$$

However, in spite of the claim, this algorithm is not general – it applies only to the case when the decision does not change the state of nature, θ_j. But, the aim of a decision might be exactly to change the state of nature – we usually want something because we are not satisfied with what we have.

The loss function has to define the value of the loss due to decision D_i for the case when the state of nature happens to be θ_k *after the decision has been enforced and its consequences have unfolded*. Jaynes used a loss function which defined the value of the loss due to decision D_i for the case when the state of nature happened to be θ_j *when the decision was made*.

Therefore, in general, we have to add an additional step, between (4) and (5), one which takes into consideration the possible consequences of decisions, i.e. that describes the transition from θ_j to θ_k.

Let P_j be the probability that the state of nature is θ_j before any decision is made:

$$P_j = (\theta_j | EX)$$

and $Q_{kj}(D_i)$ the probability that the state of nature is θ_k after a decision has been enforced and its consequences have unfolded, given that the initial state was θ_j. $Q_{kj}(D_i)$ is a function of D_i and of the initial state θ_j – this function gives the probability that the final state is θ_k if one takes the decision D_i and the initial state happened to be θ_j:

$$Q_{kj} = (\theta_k | D_i \theta_j)$$

So, in the final step one has to take the decision that minimizes the following expected loss:

$$\langle L \rangle_i = \sum_{j,k} L(D_i, \theta_k) \cdot (\theta_k \mid D_i \theta_j) \cdot (\theta_j \mid EX)$$

(there is no summation over i)

The terms P_j are static terms which describe our probabilistic knowledge about the state of nature before the decision. The terms Q_{kj} are dynamic terms which describe our probabilistic knowledge about the consequences of each decision.

In case decision D_i does not affect the state of nature we have:

$$(\theta_k \mid D_i \theta_j) = \delta_{kj}$$

This is the special case Jaynes described.

References and Notes

1. E.T. Jaynes, Probability Theory with Applications in Science and Engineering, chapter 13, "Introduction to Decision Theory".