ENTROPY AND SEARCH THEORY*

E. T. Jaynes

Arthur Holly Compton Laboratory of Physics

Washington University

St. Louis, Missouri 63130

*Presented at the First Maximum Entropy Workshop, University of Wyoming, June 1981.

ENTROPY AND SEARCH THEORY

1. INTRODUCTION

A recent article by Pierce (1978) has brought search theory to the attention of workers in related fields which also use statistical theory. In recounting history, he noted that early workers tried to relate detection probability p_D and search effort to the posterior entropy H_{ND} conditional on nondetection [Eq. (4) below] or to the "expected posterior entropy" $H_E = p_D H_D + (1 - p_D) H_{ND}$, discovered quickly that no general relation exists; and concluded that information theory has no useful connection with search theory.

As Pierce stated: "These negative findings had a clearly inhibiting effect on research, and relatively little effort has been devoted to the connections between information and search for the past fifteen years. Nonetheless, the intuitive appeal of information theory remains strong." He then presented some numerical analyses showing that in some cases maximum posterior entropy did, after all, correspond closely with maximum detection probability, although exceptions were also found. After analyzing the available evidence, Pierce concluded that the relation between search theory and information theory remains complex, but that the situation is promising enough to justify further study.

To an information theorist that intuitive appeal is so strong that one is convinced from the start; there _must_ exist a close relation between information and optimal search policy; and not just a numerical coincidence holding in some cases. There must be an exact, analytically demonstrable, and very general relation pertaining not only to search theory, but to optimal planning for any objective. For any optimal strategy is only a procedure for exploiting our prior information in order to achieve whatever goal is set, as quickly (or as efficiently as measured by any cost assignment) as possible.

Indeed, Shannon's original creation of information theory arose from a special case of this: optimal encoding of a message so as to transmit it most efficiently by the cost assignment of channel capacity. As we have pointed out (Jaynes, 1978), all of presently known Statistical Mechanics is included in the solution that Shannon proposed for this problem. In any such problem, the attainable efficiency must be related to--because it is determined by--the amount of prior information available. If past efforts to find this relation have failed, it can be only from a technical failure to ask the right questions.

We show here that such a relation does indeed exist, but it involves different entropies than H_{ND}. We develop it here by analyzing the simple search model studied by Pierce (single stationary target, no false alarms, independent detection probability for successive looks), after which we speculate on generalizations. One of our entropy connections was given by Barker (1977); the other is possibly new. However, our purpose here is "introductory tutorial" rather than reporting new research.

2. THE SIMPLE MODEL

There is a hidden "target" in region R. Each time we look at R we have, independently, the probability q of detecting it, so the probability that it will have been detected in k looks is $[1 - (1-q)^k]$. We generalize by replacing the discrete number of looks k by a continuous "search effort" variable z, and define a "search parameter" s by $(1-q) = \exp(-s^{-1})$.

Then the probability that a search effort z will result in detection is

$$p(D|z) = 1 - \exp(-z/s) \quad . \tag{1}$$

Now consider that, instead of a single region R, the target is known to be in one and only one of n different "cells" with various search parameters $\{s_1 \ldots s_n\}$. With prior probability P_i that it is in cell i $(1 \leq i \leq n)$, the predictive prior probability that the search allocations $\{z_1 \ldots z_n\}$ will result in detection, is

$$P_D = \sum_{i=1}^{n} P_i[1 - \exp(-z_i/s_i)] \quad . \tag{2}$$

If, after this search, the target has not been located, the posterior probability that it is in cell i will be

$$p_i = \frac{P_i \exp(-z_i/s_i)}{\sum P_j \exp(-z_j-s_j)} \quad . \tag{3}$$

In this notation, we use $p_i = p_i(z_1 \ldots z_n)$ as the "running variable" cell probabilities that evolve continuously throughout the search, and $P_i = p_i(0)$ for their fixed initial values. The aforementioned entropy is then

$$H_{ND} = -\sum p_i \log p_i \quad . \tag{4}$$

3. RELATIVE ENTROPY

The cell parameter s_i is a measure of the search effort required to achieve a given detection probability in cell i. If the cells consist of various areas to be searched, then one expects that a cell of twice the area will require twice the search effort; thus we may think of s_i quite generally as the "size" of cell i. Of course, in different problems this size may be measured in various terms: not only area, but equally well man-days, gasoline consumption, number of microscope slides, film footage, computer time, etc. In whatever units cell size s is measured, search effort z will be measured in the same units.

Now the original definition of our cells arose presumably in some natural way out of circumstances of the problem; but in principle their definition is arbitrary. They may be combined or subdivided in various ways and the cell sizes are additive. In particular, we can find integers N_i such that

$$\frac{s_i}{\sum s_i} = \frac{N_i}{\sum N_i} , \qquad 1 \leq i \leq n \tag{5}$$

to any desired accuracy (by choosing N_i sufficiently large). But then a cell with size, probability, and search allocation $\{s_i, p_i, z_i\}$ may be subdivided into N_i equal cells, each of size, probability, and search allocation

$$r_k = s_i/N_i \ ; \qquad w_k = p_i/N_i \ ; \qquad y_k = z_i/N_i \ ; \qquad 1 \leq k \leq N_i \tag{6}$$

and the detection probability $p_i[1 - \exp(-z_i/s_i)]$ may be written equally well as

$$\sum_{k=1}^{N_i} w_k[1 - \exp(-y_k/r_k)] \quad . \tag{7}$$

But by construction all the new cells (k) are the same size, from whatever old cell (i) they were derived. Therefore we have refined the problem to one where we have

$$N = \sum_{i=1}^{n} N_i \tag{8}$$

equal cells. At this point, we could generalize further by relaxing the requirement of equal w_k, y_k in (6).

Clearly, in view of the symmetry, the correct entropy which measures our information about the refined cells must be

$$H = -\sum_{k=1}^{N} w_k \log w_k = -\sum_{i=1}^{n} N_i(p_i/N_i) \log(p_i/N_i) \tag{9}$$

which has the upper bound $H_{max} = \log N$. It is customary to subtract off this irrelevant additive constant by defining the new entropy $I \equiv H - \log N$, or

$$I(z) = \sum_{i=1}^{n} p_i \log(m_i/p_i) \tag{10}$$

where

$$m_i \equiv \frac{N_i}{N} = \frac{s_i}{S} \quad , \qquad S \equiv \sum s_i \tag{11}$$

are the cell sizes, normalized to $\sum m_i = 1$. We may, equally well, define an entropy in which the roles of the distributions $\{p_i\}$, $\{m_i\}$ are interchanged:

$$J(z) \equiv \sum_{i=1}^{n} m_i \log(m_i/p_i) \tag{12}$$

These satisfy the Gibbs inequalities $I \leq 0$, $J \geq 0$, with equality in each case if and only if $\{p_i = m_i, \; 1 \leq i \leq n\}$.

The quantity I may be called the entropy of the distribution $\{p_i\}$ relative to the basic "measure" m_i (so called to suggest still further generalizations not needed here), while (-J) is the entropy of $\{m_i\}$ relative to $\{p_i\}$. These quantities go by various other names--"cross entropy," "directed divergence," "minimum discrimination information statistic," "essergy," etc.--but we think those terms should be discouraged because they imply that I is a different kind of object than H_{ND}. In fact, I is simply the entropy over the symmetric refined cells, and is every bit as much as a "true entropy" as H_{ND}. Indeed, it is not H_{ND}, but I and J, that have a simple and general relation to search theory, as follows:

Consider a search that starts from initial values I(0), J(0) which are measures of our prior information about the target location. At any subsequent stage $\{z_1 \ldots z_n\}$ of the search effort--whether optimal or not--the present values are I(z), J(z). The change [I(z) - I(0)] is the measure of the amount of prior information utilized up to that point, while [J(0) - J(z)] is the measure of the saving in search effort thereby achieved. The optimal policy is then the one that trades off initial information for reduced search effort, as quickly as possible.

The connection of $I(z)$ with information was indicated in the derivation of (10). To demonstrate the connection of $J(z)$ with search effort, note that from (2), the denominator of (3) is just $(1 - p_D)$. Therefore, at any stage where we have allocated the search effort $\{z_i\}$, $J(z)$ is from (12)

$$J(z) = \sum_{i=1}^{n} m_i \log\left[\frac{m_i}{p_i}(1 - p_D)\exp(z_i/s_i)\right]$$

$$= J(0) + \log(1 - p_D) + (z/S) \qquad (13)$$

where $z = \sum z_i$ is the total search effort used. But (13) states only that at this stage the detection probability is

$$p_D = 1 - \exp\left(-\frac{z + z^*}{S}\right) \quad . \qquad (14)$$

where

$$z^* \equiv S[J(0) - J(z)] \quad . \qquad (15)$$

Since $J(z) \geq 0$, if we start from prior ignorance, $J(0) = 0$, then clearly the best we can do is to conduct the search so as to keep $J(z) = 0$; then the detection probability will be

$$p_D = 1 - \exp(-z/S) \quad , \qquad (16)$$

i.e., just the original detection function (1), in which we have lumped all cells together into one large cell of size $S = \sum s_i$. Thus, z^* is precisely the saving in search effort, for a given detection probability, that has been achieved up to that point by exploiting the prior information.

All details of an optimal policy are easily set forth if we may assume the following property.

4. DYNAMICAL CONSISTENCY

In a real-life situation the problem of deciding on a search allocation will be almost hopelessly complicated, or even indeterminate, unless the following property holds. Consider two different problems:

(A) You are allotted a total search effort $\sum z_i = C$. Decide on the optimal allocation $\{\hat{z}_1 \ldots \hat{z}_n\}$ to maximize the probability of detection.

(B) The authorities have divided the search effort C into two portions, $C_1 + C_2 = C$. You are allotted first the amount C_1, and must decide on the optimal allocation $\{\hat{z}_i^{(1)}\}$ with $\sum z_i^{(1)} = C_1$ on the assumption that no further search effort is available. This fails to locate the target; but you then learn that you may apply for permission to use the additional search effort C_2. If this is granted, you must then decide on the optimal allocation $\{z_i^{(2)}\}$ with $\sum z_i^{(2)} = C_2$ for the second try.

The problem has <u>dynamical consistency</u> if the optimal total search allocation is the same in problems (A) and (B); i.e., if

$$\hat{z}_i^{(1)} + \hat{z}_i^{(2)} = \hat{z}_i \quad , \qquad 1 \leq i \leq n \tag{17}$$

for all C_1 in $(0 \leq C_1 \leq C)$. This is a highly desirable property for psychological, practical, and mathematical reasons.

Psychologically, it is a comfort to the decision maker; for then he can face his problems one at a time, making at each step the decision that is optimal for the search effort being then committed--secure in the knowledge that whatever the final outcome, no critic full of hindsight can later accuse him of blundering (this remains true even if he has inherited the job from a blundering predecessor). Put differently, we are supposing that "global" optimization can be found by a sequence of "local" optimizations.

Practically, it is a useful property; for even if one knows in advance exactly how much total search effort can be used, it may be necessary to search the cells one at a time. Then one must in any event decide on the optimal order of searching, which amounts to a sequence of problems of type (B). Without dynamical consistency, the optimal action for today would in general depend on imaginary contingencies that might or might not arise tomorrow, and a "global" optimum would be very hard to find.

Mathematically, dynamical consistency reduces the problem for any amount of search effort to successive allocations of infinitesimal amounts δz, for which the optimal allocation is obvious. Given any previous search allocation $\{z_1 \ldots z_n\}$, whether optimal or not, which has reduced the cell probabilities (3) to $\{p_1 \ldots p_n\}$, if the new increment δz is used in cell j, the probability that it will result in detection is

$$p_j[1 - \exp(-\delta z/s_j)] = (p_j/s_j)\delta z \quad . \tag{18}$$

But if detection does not result, then according to (3) the probability of the i'th cell is changed by

$$\delta p_i = (p_i - \delta_{ij})(p_j/s_j)\delta z \tag{19}$$

and from (10), (12) the entropies I(z), J(z) will receive the increments

$$\delta I = [I + \log(p_j/m_j)](p_j/m_j)\delta z \quad , \tag{20}$$

$$\delta J = S^{-1}[1 - (p_j/m_j)]\delta z \quad . \tag{21}$$

Since $\sum p_i = \sum m_i = 1$, we have always $(p_j/m_j)_{max} \geq 1$, and from (10), $[I + \log(p_j/m_j)_{max}] \geq 0$. Thus the allocation of δz which maximizes the detection probability (18) leads to $\delta I \geq 0$, $\delta J \leq 0$; from (20), (21) it therefore also maximizes the posterior entropy I and minimizes J. By all three criteria, the optimal policy allocates each new increment δz to whatever cell has at that time the greatest value of (p_j/m_j); as noted, this is the optimal present policy even if the previous allocation $\{z_i\}$ was not optimal.

This optimal search policy always takes us toward the condition of "complete ignorance" $I = J = 0$; and thus (as noted by Pierce, in agreement with an earlier conjecture of Richardson) it "uses up" the prior information, as rapidly as possible. The limiting state $I = J = 0$ is actually reached, at a finite total search effort, if the search continues long enough without detection [Eq. (37) below]. Such a search therefore has a fundamental division into an "early phase" in which $I < 0$, $J > 0$ and the prior information is being used to determine policy; and a "final phase" in which it is all used up: $I = J = 0$, and the optimal policy is independent of the prior information.

We now examine in some detail the course of the optimal search policy for this model. In this we necessarily repeat a few facts well known in the literature of search theory; our object is to point out their interpretation in the light of the entropies $I(z)$, $J(z)$. The search process now appears very much like an irreversible process in thermodynamics, in which an initially nonequilibrium state relaxes into the equilibrium state of maximum entropy. But now it is only our <u>state of knowledge</u> that relaxes to the "equilibrium" condition of maximum uncertainty, $I = J = 0$.

5. AN EXAMPLE OF OPTIMAL SEARCH

On the assumption of dynamical consistency, the entire course of the optimal search effort is clear; since according to (19) the probability of the searched cell is always lowered, that of the others raised, the optimal strategy is the one that equalizes the numbers

$$a_i \equiv p_i/m_i \tag{22}$$

starting from the top, as quickly as possible. We follow the aforementioned notation of writing $a_i = a_i(z_1 \ldots z_n)$ for the "running variables" that evolve during the search, and $A_i = a_i(0)$ for their fixed initial values. Number the cells according to those initial values, so that

$$A_1 \geq A_2 \geq \cdots \geq A_n \quad . \tag{23}$$

Then the optimal search proceeds as follows:

<u>Stage 1</u>. All the initial effort should go into cell 1 until its probability is reduced to the point where $a_1 = a_2$. The search effort required to do this is, from (3),

$$z_1^{(1)} = s_1 \log(A_1/A_2) \tag{24}$$

and the prior probability of detection at or before this point is

$$p_D^{(1)} = P_1[1 - \exp(-z_1^{(1)}/s_1)] = m_1(A_1 - A_2) \ . \tag{25}$$

Thus from (13) the entropy J has changed by

$$J^{(1)} - J(0) = \log[1 - m_1(A_1 - A_2)] + m_1 \log(A_1/A_2) \ . \tag{26}$$

That this must be negative if $A_1 > A_2$ is evident from (21); to prove it directly from (26) one must take into account also the inequalities $\{A_2 \geq A_k, 3 \leq k \leq n\}$.

At any stage in the search, the entropy I(z) may be written in the form

$$I(z) = \log(1 - p_D) + (1 - p_D)^{-1} K(z) \tag{27}$$

where K(z) is an analytically simpler expression. Therefore we indicate the entropy changes by giving the values of K at each stage of the search. Initially, K(0) = I(0), but after the search effort (24) we find

$$K^{(1)} = I(0) - P_1 \log(m_1/P_1) + (m_1/m_2)P_2 \log(m_2/P_2) \ . \tag{28}$$

That is, the right-hand side of (28) is the expression (10) for I(0) in which the first term $[P_1 \log(m_1/P_1)] = -m_1A_1 \log A_1$ has been replaced by $-m_1A_2 \log A_2$. In effect, this lumps the first two cells together into a single cell of measure $(m_1 + m_2)$.

Stage 2. According to (18), (20), (21) cells 1 and 2 are now equally search-worthy, a further small search effort yielding equal detection probability and equal entropy increase in either. The next efforts should therefore by allocated to both, in the ratio which maintains the equality $a_1 = a_2$; that is, in the ratio $m_1:m_2$, which amounts to equal allocation to the $(N_1 + N_2)$ refined cells derived from cells 1, 2. The second stage continues until $a_1 = a_2 = a_3$, at which point we have used the additional search effort

$$(s_1 + s_2) \log(A_2/A_3) \tag{29}$$

and the total amounts expended in cells 1 and 2 up to this point are

$$z_1^{(2)} = z_1^{(1)} + s_1 \log(A_2/A_3)$$

$$= s_1 \log(A_1/A_3) \tag{30}$$

$$z_2^{(2)} = s_2 \log(A_2/A_3) \quad . \tag{31}$$

The prior probability of detection at or before this point is

$$p_D^{(2)} = m_1(A_1 - A_3) + m_2(A_2 - A_3) \tag{32}$$

and the entropy $I^{(2)}$ is given by (27) with

$$K^{(2)} = I(0) - P_1 \log(m_1/P_1) - P_2 \log(m_2/P_2) + \frac{m_1 + m_2}{m_3} P_3 \log(m_3/P_3) \tag{33}$$

that is, by I(0) with the first two terms replaced, in effect lumping the first three cells into a single cell of measure $(m_1 + m_2 + m_3)$.

<u>Stage 3</u>. At this point, cells 1, 2, and 3 are equally search-worthy, and so the next effort is allocated to them in the ratios $m_1:m_2:m_3$ until $a_1 = a_2 = a_3 = a_4$, at which point $K^{(3)}$ is given by I(0) with the first three terms replaced--and so on.

This initial "equalization phase" continues until for the first time $a_1 = a_2 = \ldots = a_n$, at which point we have used up the total search effort

$$z' = \sum z_i = \sum_{i=1}^{n-1} s_i \log(A_i/A_n) = -S[J(0) + \log A_n] \tag{34}$$

but have not searched at all in cell n: $z_n = 0$. The prior probability of detection has reached

$$p_D^{(n-1)} = 1 - A_n \tag{35}$$

and $K^{(n-1)}$ is I(0) with the first (n-1) terms replaced: i.e.,

$$K^{(n-1)} = - A_n \log A_n \quad . \tag{36}$$

From (27), (35), the entropy I(z) is now reduced to

$$I^{(n-1)} = 0 \tag{37}$$

and from (13), (34), (35) we have also $J^{(n-1)} = 0$. The posterior probabilities (3) have completed the relaxation into their "equilibrium" values $\{p_i = m_i, 1 \le i \le n\}$; i.e., the refined cells now have equal probabilities $\{w_k = N^{-1}, 1 \le k \le N\}$.

Final Phase. All cells are now equally search-worthy, so if detection is not yet achieved, any further search effort z" is allocated to all cells in the proportions $z_i'' = m_i z''$ which maintain that condition; i.e., it is allocated equally to the refined cells. The posterior probabilities remain at their equilibrium values, the entropies I, J remain zero, and the detection probability with any further amount of search effort (i.e., for total effort $z = z' + z'' \ge z'$) is

$$p_D(\infty) = 1 - \exp\left[- \frac{z + z^{**}}{S}\right] \tag{38}$$

where, comparing with (14), (15),

$$z^{**} \equiv S\,J(0) \qquad (39)$$

is the maximum possible saving in search effort that can be "bought" with the prior information.

6. CONCLUSION — THE MORAL

We have shown how entropy maximization is related to optimal strategy in one simple case. This can, of course, be generalized in many different ways--in fact, the situation is open-ended because there is no end to the variety of new problems that could arise. So it is impossible to give a "most general" case once and for all. But before one can extend the theory to some particular new case, it is necessary to understand the moral of what we have just learned.

Why did it require nearly thirty years after Shannon's work to find this (maximum entropy)-(optimal search) connection, in spite of the fact that many workers suspected its existence and tried to find it? The answer was given about 130 years ago by George Boole, who remarked: "I think it one of the peculiar difficulties of probability theory, that its difficulties sometimes are not seen". What has not always been seen here is that the simple, unqualified term "entropy" is meaningless; entropy is always defined with respect to some basic "measure" and the result of maximizing it depends not only on the constraints, but also on the measure.

The difficulty in applying maximum entropy to problems outside thermodynamics is not in deciding what constraints should be applied, but in deciding what is the underlying measure--or, as I prefer to call it, what is the "hypothesis space" on which our entropies are defined? More informally, what is the field on which our game is to be played?

This problem does not loom large in thermodynamics--in fact most writers seem hardly aware of it--but this is only because it was solved over 100 years ago by Liouville. Classical phase volume is invariant under canonical transformations (of which the equations of motion are a special case), and so equal weighting to equal phase volumes was the field on which Gibbs' game was played.

This leads to many correct predictions (equations of state, susceptibilities, high-temperature specific heats of solids and monatomic gases), but at low temperatures Nature persisted in giving lower specific heats--and therefore states of lower entropy--than Gibbs predicted. In Nature, therefore, there must be further constraints operative, beyond those imposed by Gibbs. This was the first clue pointing to quantum theory.

The resolution, found by Einstein, Debye, von Neumann, and Brillouin, was quite simple. It seems that not all classically allowed energies are used by Nature, and equal weighting to orthogonal quantum states of a system--which goes asymptotically into Liouville's weighting--is the new field on which we play the game of quantum statistics. According to all present knowledge, maximum entropy on this hypothesis space leads unerringly to correct predictions. Still, I keep trying to find a case where it fails, because then we would have a clue pointing to the new theory that will someday replace our present quantum theory, and so history would be repeated.

In applications outside thermodynamics we are still at a phase corresponding to--if one can imagine it--statistical mechanics before the discovery of Liouville's theorem. The originally tried entropy H_{ND} of Eq. (4) was defined with respect to uniform weighting of all search cells regardless of their size. Such a weighting simply ignores the cogent information about cell sizes. Our proceeding to the refined cells of equal size restored the symmetry of our hypothesis space--and corresponded to the discovery of Liouville's theorm. As soon as we play our game on the field defined by (9), the connection of entropy with optimal search appears immediately.

Moral: In any new problem, one must face anew, what is the underlying symmetrical hypothesis space on which our entropies are defined? The strategy is:

(1) Think hard about the appropriate hypothesis space. Look for some symmetry/invariance property.

(2) Try out your best choice. If the desired kind of useful results appear, then well and good--there is no evidence pointing to a different hypothesis space and you are done--at least for the time being.

(3) If you get unsatisfactory results, then if you are convinced that all relevant constraints have been taken into account, this is evidence that Nature is using a different hypothesis space than yours. Go to step (1).

In spectrum analysis, the Burg solution implied independent uniform weighting to all possible values of $\{y_0 \ldots y_N\}$. Its success thus far indicates that we are now at step (2). However, the future may bring some surprise here. Any persistent failures would point to a new hypothesis space--and therefore to the possibility of still better predictions.

In image reconstruction, the present solutions seem to be based on uniform independent weighting to all values of luminance for each pixel. I have a suspicion--perhaps shared by John Skilling, although he expresses himself in very different terms--that a deeper hypothesis space that to some degree "anticipates" correlations of adjacent pixels, may be still better. Of course, we would have to accumulate a great deal of further experience before we could be sure that we were at step (3).

We hope that entropy considerations will be brought to bear on other problems of optimal strategy, and perhaps with enough experience we shall learn how to define our hypothesis space for such problems, just as confidently as physicists now do in Statistical Mechanics.

REFERENCES

Wm. H. Barker (1977), "Information Theory and the Optimal Detection Search", Operations Research, v.25, pp. 304-314.

E. T. Jaynes (1978), "Where Do We Stand on Maximum Entropy?", in The Maximum Entropy Formalism, R. Levine and M. Tribus, Editors; M.I.T. Press, Cambridge MA, pp. 15-118.

John G. Pierce, "A New Look at the Relation Between Information Theory and Search Theory", in The Maximum Entropy Formalism, R. Levine and M. Tribus, Editors; M.I.T. Press, Cambridge MA, pp. 339-402.