REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE JULY 1996	3. REPORT TYPE AND DATES COVERED TECH REPORT

4. TITLE AND SUBTITLE
MEASURING INFORMATION GAIN IN TACTICAL OPERATIONS

5. FUNDING NUMBERS

6. AUTHOR(S)
DONALD BARR
E. TODD SHERRILL

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
OPERATIONS RESEARCH CENTER
USMA
WEST POINT, NY 10996

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
A measure is proposed for analytically determining the amount of information gained by a tactical battle commander as a result of intelligence, scouting or reconnaissance reports. The measure is based on concepts from information theory, and involves modeling a commander's uncertainty in terms of probability distributions over sets of possible states his adversary may occupy. As the commander gets information, these distributions are updated, by various means, to represent his current state of uncertainty. The information gain is defined in terms of the "distance" between the initial and updated states of uncertainty.

Based on a set of plausible assumptions about how an information gain measure should behave, it is shown the measure must be of a certain form involving the decrease in entropy (as defined by Shannon) from the prior to updated distributions. Implications of this characterization are presented and illustrated with examples. Applications to experiments performed at the U.S. Military Academy are described, including

o a Janus combat simulation study of the relative reconnaissance performances of the Comanche helicopter and an unmanned aerial reconnaissance system;
o a Janus-based experiment designed to establish links between the level of information possessed by a combat commander and the degree of success he enjoys against his adversary;
o a simulation-based design study of intelligent minefields; and
o development of an information gain MOE for Janus analyses.

We believe these examples demonstrate the potential utility of the information gain measure for a wide variety of applications.

14. SUBJECT TERMS
MEASURING INFORMATION GAIN IN TACTICAL INFORMATION

15. NUMBER OF PAGES
47

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	

MEASURING INFORMATION GAIN IN TACTICAL OPERATIONS

Professor Donald R. Barr
Department of Systems Engineering

and

Maj E. Todd Sherrill
Operations Research Center

U.S. Military Academy
West Point, NY 10996

A TECHNICAL REPORT
OF THE
OPERATIONS RESEARCH CENTER
UNITED STATES MILITARY ACADEMY

Directed and Approved by
LTC David W. Hutchison
Director
Operations Research Center

July 1996

PREFACE

The purpose of this report is to provide a summary of recent results we have obtained in the area of measuring information, as part of a project of the Operations Research Center, United States Military Academy. We are interested in both the measurement of information gained through activities such as reconnaissance and scouting, and the operational implications of possessing varying amounts of information. We have developed a measure of effectiveness of information processes which we call "information gain." We believe our modest applications of this measure demonstrate its potential utility. We hope this report will place our work on a sound theoretical footing, and will serve as a guide on how it can be applied.

The theoretical properties of information gain, together with results from several experiments, lead us to hope the concept has a generic quality. If so, the behavior of the measure implies interesting fundamental properties of information in operational terms.

The information gain measure we discuss is relatively straight-forward when one models a decision maker's state of uncertainty about his adversary in terms of discrete probability distributions over a space of possible states the adversary may occupy. This model seems adequate for applications in which it makes sense to imagine a finite set of possible states and a probability distribution over this set which may be updated as information about the state occupied is received. In these circumstances the measure depends only on the probabilities, and not upon any choice of how states are labeled.

But the situation becomes less obvious when one considers continuous probability distributions and associated random variables. Indeed, in using a random variable, one implicitly invokes a coordinate system, and hence a specific labeling of points in the state space. This can bring seeming paradoxical results, so care must be exercised in use of continuous models. Nevertheless, it frequently is useful to employ random variables to notationally represent their corresponding distributions, so we are motivated to study behavior of the information gain measure in terms of random variables with continuous as well as discrete distributions.

Aside from the applications we describe, perhaps the most interesting result in this report is the characterization of the mathematical form of the information gain measure. This characterization, together with several associated corollaries, provide insights into the nature of information and attributes of information gain. We believe there may be applications of some of these ideas to many facets of managing information processes, such as optimally allocating information gathering resources, determining the marginal value of information and timing decision points, assessing the operational value of alternative information levels or processes, and training decision makers to properly use information at hand, particularly in cases of very high information levels.

The report is organized into three parts, as follows. In Part I we give a general description of information gain, and some of the underlying ideas and techniques. Part II presents a number of examples of a slightly technical nature, and describes several applications. Part III is devoted to the characterization theorem and presentation of properties of information gain from a more theoretical perspective. Even though we frequently cross-reference the parts, they are more-or-less independent. With minor referrals to clarify notation and terminology, they may be read in part, in any order.

THE OPERATIONS RESEARCH CENTER

The United States Military Academy's Operations Research Center (ORCEN) provides a small, full-time analytical capability to both the United States Army and the Academy. It typically employs about five full-time Army analysts; at any point in time, about a half dozen Systems Engineering Department military and civilian faculty, together with students of the Military Academy, are working on a part-time basis on ORCEN projects. The ORCEN is co-located with the Department of Systems Engineering in Mahan Hall, West Point, NY and is sponsored by the Assistant Secretary of the Army (Financial Management). Fully staffed and funded since Academic Year 1990-1991, the ORCEN has made significant contributions to the Army's analytical efforts.

CONTENTS

EXECUTIVE SUMMARY

A measure is proposed for analytically determining the amount of information gained by a tactical battle commander as a result of intelligence, scouting or reconnaissance reports. The measure is based on concepts from information theory, and involves modeling a commander's uncertainty in terms of probability distributions over sets of possible states his adversary may occupy. As the commander gets information, these distributions are updated, by various means, to represent his current state of uncertainty. The information gain is defined in terms of the "distance" between the initial and updated states of uncertainty.

Based on a set of plausible assumptions about how an information gain measure should behave, it is shown the measure must be of a certain form involving the decrease in *entropy* (as defined by Shannon) from the prior to updated distributions. Implications of this characterization are presented and illustrated with examples. Applications to experiments performed at the U.S. Military Academy are described, including

- a Janus combat simulation study of the relative reconnaissance performances of the Comanche helicopter and an unmanned aerial reconnaissance system;
- a Janus-based experiment designed to establish links between the level of information possessed by a combat commander and the degree of success he enjoys against his adversary;
- a simulation-based design study of intelligent minefields; and
- development of an information gain MOE for Janus analyses.

We believe these examples demonstrate the potential utility of the information gain measure for a wide variety of applications.

In developing the measure and working on its application, we found evidence to support several tentative observations about information gain. These technical and operational observations are drawn in a tactical setting:

- the information gained in finding an enemy target is independent of the enemy's force size;
- finding a terrain cell to be target-free gives relatively greater information gain if Red has a very large force than if he has a small force;
- assuming independence in target locations that are actually correlated appears to give error in information gain that is much smaller than the respective errors in the individual entropy values;
- for a mobile target located at some time t_0, but unobserved for subsequent times t, the shape of the information loss curve may generally be of the form $-\ln(t^2)$, independent of movement rate of the target;
- usually, information gain for a given target will be positive over time, as recon is conducted; however, there are cases in which there can be increases in uncertainty as additional recon occurs; and

- from the point of view of extending the definition of entropy for discrete distributions, $-\Sigma p_i \ln(p_i)$, to continuous distributions, the expression $-\int f(x)\ln(f(x))dx$ may not be appropriate for measuring uncertainty, but integrals of this form may be employed in computing information gain for continuous distributions, with an interpretation identical to that for discrete distributions.

In addition we observed an operational behavior in one of our experiments that we believe may have more general applicability:

- a tactical commander does not require information beyond a moderate level in order to accomplish his mission, but he can achieve mission success at reduced cost when he has additional information.

INTRODUCTION

The Army has recently expressed heightened interest in managing information processes related to combat operations. This has generated a need in the analytic community for information-related analysis methods. Almost everyone intuitively believes information has value in combat, but it is not obvious how the underlying relationships might be quantified. It is desirable to measure information and its implications upon combat processes so the operational commander might accomplish actions such as:

- assessing the contributions of information to the likelihood of success of various battle operating systems;
- evaluating the information implications of courses of action during wargaming; and
- allocating effort among alternative reconnaissance, scouting and intelligence systems.

Incoming data are not information to a decision maker until they inform. In our case, the decision maker is assumed to be the tactical commander in conjunction with his staff. Commonly used analytic measures of the information a commander receives are based on the volume or rate of messages, the message quality, or characteristics of the data given in the messages. The cognition of, and response to, information conveyed in a given set of data depends upon the receiving commander. A commander mentally processes acquired data into his picture of the tatical situation. This human process depends on the personality, training, and experience of the commander. Therefore, attempting to measure information gain, either by looking at parameters of the raw message traffic or attempting to model a commander's mental processes, seems to us to be a very difficult approach.

A more tractable approach might attempt to capture the amount by which a commander has been informed as a result of receiving reconnaissance and other similar data. One would like to resolve the issue, "How does a commander's state of knowledge change when he receives new data containing information?" Rather than attempting to measure the level of information a commander possesses at a given point in time, we attempt to model the amount of *uncertainty* the commander faces. We shall describe an approach involving modeling a commander's uncertainty about his enemy's disposition in terms of probability distributions. As the commander gains information about his adversary, the probability distributions are updated to reflect the new state of the commander's uncertainty. Using this approach, along with information theoretic measures related to the probability distributions, we can measure the changes in uncertainty brought about by the receipt of new data. This change in uncertainty is a measure of the amount of information gained through receipt of the data. Figure 1 depicts how receipt of information leads to decreased uncertainty about the enemy disposition.

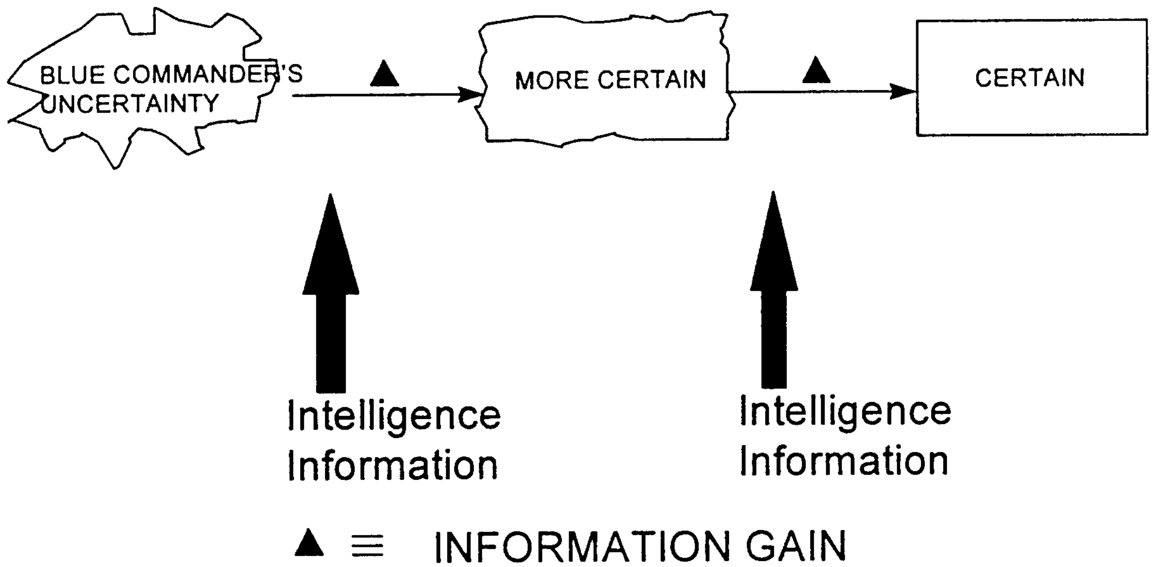

Figure 1. Decrease in uncertainty due to information gain.

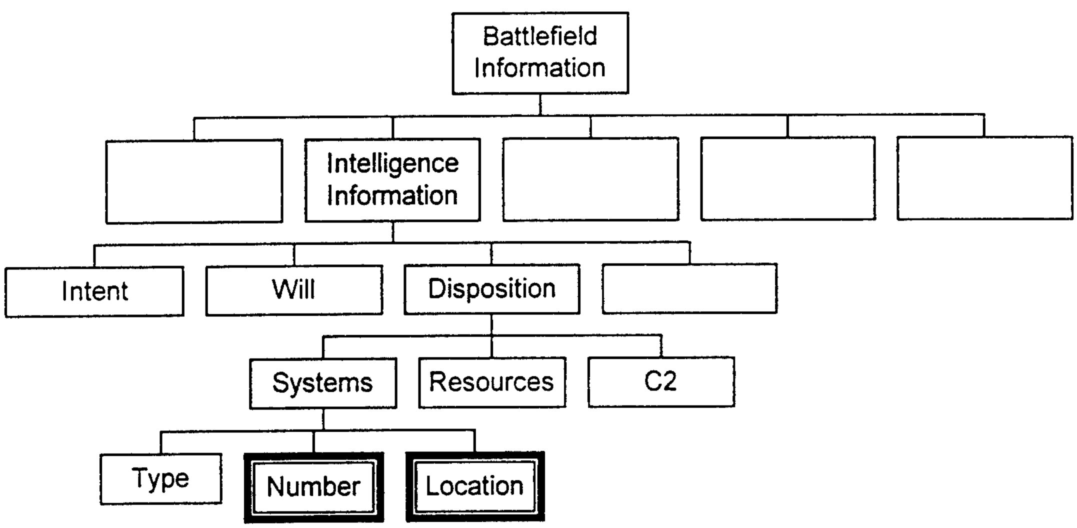

Figure 2. Subset of battlefield information considered in our examples.

We believe our approach in modeling information gain in terms of decreased uncertainty falls somewhere between approaches that model the characteristics of the physical communications system and those that attempt to model human cognition and response of the decision maker.

Our work has focused on tactical intelligence information. Figure 2 shows how our considerations concern only a subset of potential battlefield information. We deal here only with the number and locations of enemy, arguably the most important of battlefield intelligence data. Enemy size and disposition data are typically gained through reconnaissance, scouting and other defined information system activities. They generally do not depend on the intuition and experience of the commander, so there is a generic character to the approach we propose.

Earlier results are reported in [1, 2, 3]. We have employed the concept of information theory, developed by Shannon [10], to define an information gain metric that measures reduced uncertainty due to reconnaissance and other intelligence activities in an area of concern to a commander. The idea of using methods of information theory to measure effects of reconnaissance, scouting, intelligence, and other activities related to preparing for and conducting military operations seems quite natural.

This paper documents our work on the information gain measure and reports several examples of its application in experiments conducted by our students and ourselves during the past several years. One of the experiments was aimed at comparing reconnaissance platforms operating in a Bosnian scenario; a second experiment was aimed at establishing the operational value of information in simulated combat at the National Training Center; and a third application was associated with evaluating information obtained by an intelligent anti-armor minefield.

One objective of this work is to facilitate studies of relationships between information gained about the enemy's disposition and various measures of combat effectiveness. In [1] we describe a modest experiment along these lines, where information available to the tactical commander increased in a sequence of stages. A plot of a relationship between information gain and combat success is shown in Figure 3, for this experiment conducted at USMA last year. A plot illustrating comparison of the Comanche helicopter and an unmanned aerial vehicle (UAV) in gaining information about the disposition of an enemy force in a hypothetical Bosnian scenario [5] is depicted in Figure 4. These experiments and several results are described further in Part II of this report.

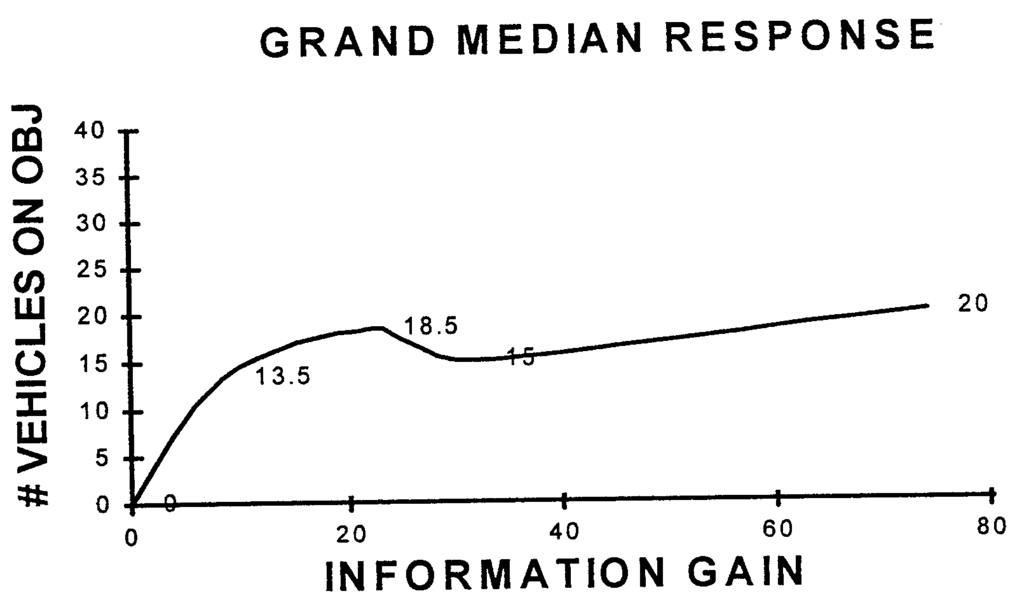

Figure 3. Plot of "mission success" versus information gain, reported in [1].

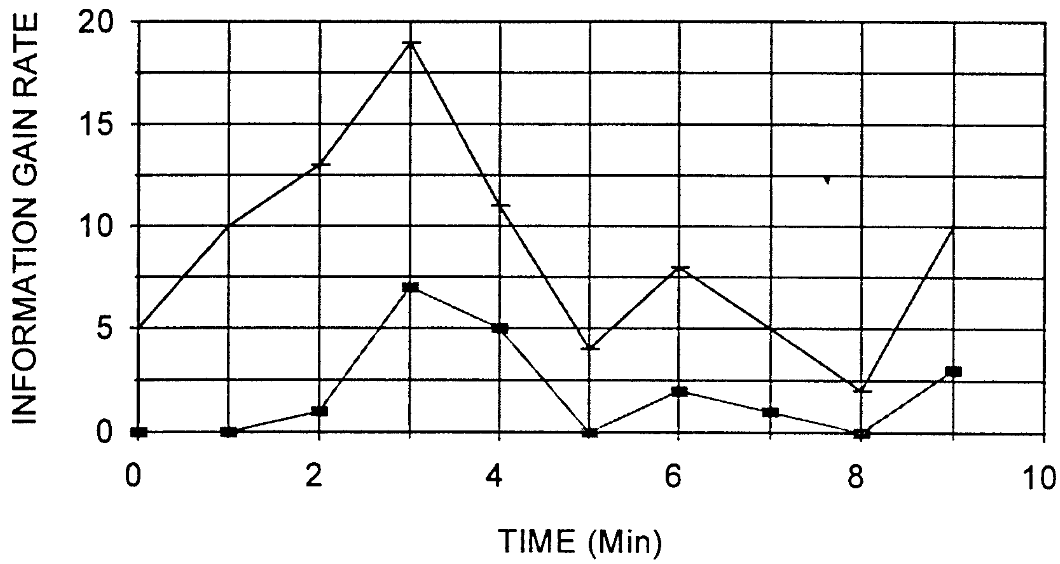

Figure 4. Sample comparison of a Comanche (lower curve) and UAV (upper curve) in a hypothetical scenario [5].

An additional objective of our work is to automate the information gain measure of effectiveness (MOE) in combat simulations. This should facilitate comparisons of various alternative reconnaissance platforms, information gathering tactics, information system organizations, and sensors. We are currently working on implementation of information gain in the Janus model [11]; a summary of these efforts is given in Part II.

This report is divided into three parts. Part I gives a general description of information gain, Part II describes several applications carried out here at the U.S. Military Academy, and Part III is devoted to presenting a characterization of the measure, developing several of its properties, and discussing several analytical issues. An operationally oriented reader may wish to concentrate on Part I and one or two of the examples in Part II. An analyst wishing to apply the measure may benefit from careful reading of most of Part II. Readers interested in basic properties and behavior of information gain may find Part III of interest. We number sections within parts; Section II-3 is the third section within Part II, for example.

PART I. OVERALL APPROACH

I-1. Information Gain

We have made a small extension of Shannon's information theoretic development of entropy to give a characterization of the information gain measure. Suppose **p** is the prior distribution, representing the commander's uncertainty at some specific time, and suppose the uncertainty he has at some later time is represented by the posterior distribution, **p***. In Section III-1 we show the information gained in resolving the uncertainty in **p** to that in **p***, measured by the information gain function, δ(**p**,**p***), must have a certain specific form, under the assumption of four plausible conditions.

Let S be a finite sample space (representing the set of possible terrain cells that might contain a vehicle, for example), and let Ω be the set of all (discrete) mass functions over S. We denote any "uniform" distribution in Ω having exactly n non-zero mass values equal to 1/n by the symbol "**n**", let **p**, **p***, and **q** be arbitrary members of Ω, and suppose X, Y, Z and I are jointly distributed random variables on S. We denote the mass values in the prior distribution **p** by $p_1, p_2, \ldots p_n$ and similarly for the posterior distribution, **p***.

Theorem (Characterization of the information gain function; see Section III-1): If the information gain function, δ(**p**,**p***), satisfies certain reasonable technical conditions and has the properties:

(a) if the outcome Z of an experiment having distribution **p** is represented as a compound experiment where an initial outcome I is observed then the remainder X of the experiment is observed conditioned on the value of I, the information gain in observing Z can be expressed as the information gain in observing I, plus the average (over values of I) of the information gain in observing X, given I; and

(b) the information gain, δ(**p**,**p***), in resolving the uncertainty in **p** to that in **p*** may be computed in terms of any intermediate stage of information which gives uncertainty represented by the distribution **q**, as follows:

$$\delta(\mathbf{p},\mathbf{p}^*) = \delta(\mathbf{p},\mathbf{q}) + \delta(\mathbf{q},\mathbf{p}^*);$$

then the information gain function must be of the form

$$\delta(\mathbf{p},\mathbf{p}^*) = -\Sigma_{i\in S}[p_i \ln(p_i) + p^*_i \ln(p^*_i)].$$

The two conditions described above can be paraphrased in operational terms, as follows:

(a') receiving a message containing the location of the enemy's 40 tanks has the same information content as two messages in which the first says the enemy has 40 tanks and the second reports their locations;

(b') the gain function measures cumulative changes in uncertainty so that the information gain from **p** (TOC shift change brief #1) to **p***(TOC shift change brief #2) is independent of intermediate levels of uncertainty **q** (snapshots of a fluid battle space that occur between shift change briefs) [1].

As the above expression for information gain reveals, the information gain function measures the difference between the randomness of the prior and posterior distributions, using Shannon's entropy definition of randomness [10]. The characterization theorem in Section III-1 thus asserts that, under plausible conditions, information gain must be given by the decrease in entropy from **p** to **p***; that is, $\delta(\mathbf{p},\mathbf{p}^*)$ = (entropy of **p**) - (entropy of **p***).

Notes On Entropy:

- If a discrete system can be in state j with probability $p(j); j=1,2,\ldots,n$, the *entropy, e,* of the system is defined to be $e = -\Sigma p(j)ln(p(j))$, where the sum is over all states j for which $p(j)>0$. In information theory, the logarithm is often taken to have base 2 (and the measure is in units of *bits*, for "*b*inary dig*its*"), but any other logarithm will differ from this by a multiplicative constant, so that need not concern us. We shall use natural logarithms in our work, so our information gain units might be called "nits" (for "*n*atural dig*its*"). Entropy, in information theory, has a connection with the thermodynamic concept of entropy [12], but this is not particularly useful in our applications. In a somewhat related application, entropy has been used in scoring the accuracy achieved by learning sensor models [9].
- If a system can be in any of n possible states, the entropy of the system can range between 0 (when the exact state of the system is known) to $ln(n)$ (when the state of the system has maximal "randomness," which occurs when the state of the system is uniformly distributed over the possible states). In the first case, where $p(1)=1$ (say) and the remaining $p(i)$'s are zero, the sum $-\Sigma p(j)\ ln(p(j))$ collapses to a single term so $e = -1\ ln(1) = 0$. If the searcher knows the enemy vehicle's location, there is no randomness and the entropy is 0.0. The second case, where $p(i) = 1/n$ for all i, gives
  $e = -\Sigma p(j)\ ln(p(j)) = -n\ (1/n)\ ln(1/n) = -ln(1/n) = ln(n)$.
- If the uncertainty represented by the prior distribution is totally resolved, so the specific Red state is determined by the information received, the posterior mass function will be degenerate at the state in question; we denote such a distribution by **1**, in keeping with the notation conventions listed above. Then $\delta(\mathbf{p}, \mathbf{1})$ represents the information gain in totally resolving the uncertainty in **p**, which can be interpreted as the randomness in the prior situation, represented by **p**. It is readily seen that $\delta(\mathbf{p}, \mathbf{1}) = -\Sigma p_i \ln(p_i)$ = (entropy of the prior distribution **p**). In our applications, this is interpreted as a measure of the degree of the Blue commander's uncertainty about a specific aspect of Red's disposition.
- The claim that $ln(n)$ is the maximal value of e is easy to prove (using mathematical induction, for example; an alternate proof, based on Jensen's inequality, is given in [7]). If the searcher knows nothing at all about the location of the enemy vehicle, he

[1] Tactical Operations Centers work around the clock. Normally there is a briefing between shift changes designed to update the new crew as well as the commander on the current state of operations as well as those events that transpired during the shift.

may assume it is equally likely to be in any one of the n possible cells, and the entropy takes on its maximum value, $\delta(\mathbf{n}, \mathbf{1}) = \ln(n)$.

- Entropy and information gain are dimensionless, and they do not depend on the labels used for outcomes in the sample space (cells). These measures depend only upon the probabilities of the possible outcomes. This is entirely reasonable in our application, because the labels of cells and the coordinate system of the battle area are inventions of the analyst; they are not inherently relevant to the amount of, or gain in, information about target location. Because entropy depends only on the probability masses in a discrete distribution, it is possible to denote such mass functions by vectors of mass values, such as **p**, as we have been doing. Such "probability vectors" have non-negative components that sum to one.
- Since zero is not in the domain of the logarithm function, we define $0 \cdot ln(0)$ to be 0. This extends the domain of $f(x) = x \cdot ln(x)$ by right continuity to include $x=0$, and simplifies notation in what follows.

I-2. Modeling Uncertainty

We model a Blue Commander's uncertainty about his Red adversary's state in terms of probability distributions relating to Red's size and disposition. For example, the distribution of the location of a given Red system from the Blue commander's perspective can often be represented as a bivariate probability distribution over the battle area. Figure 5 shows a hypothetical example of such a distribution over a 1km X 1km battle area. In this figure the battlefield is partitioned into one hundred 100m X 100m terrain cells. The height of the plotted surface above a given terrain cell represents Blue's model of the

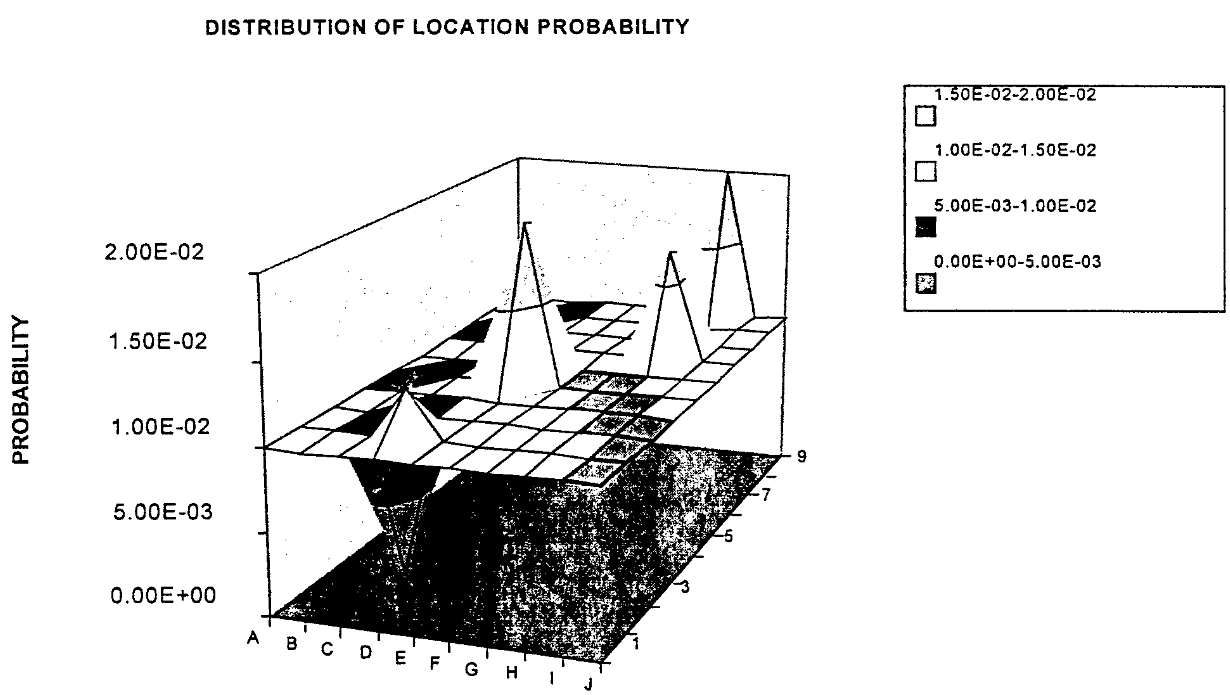

Figure 5. Prior distribution of location of a target.

relative likelihood that a given Red system, a tank, for example, is in that particular cell.

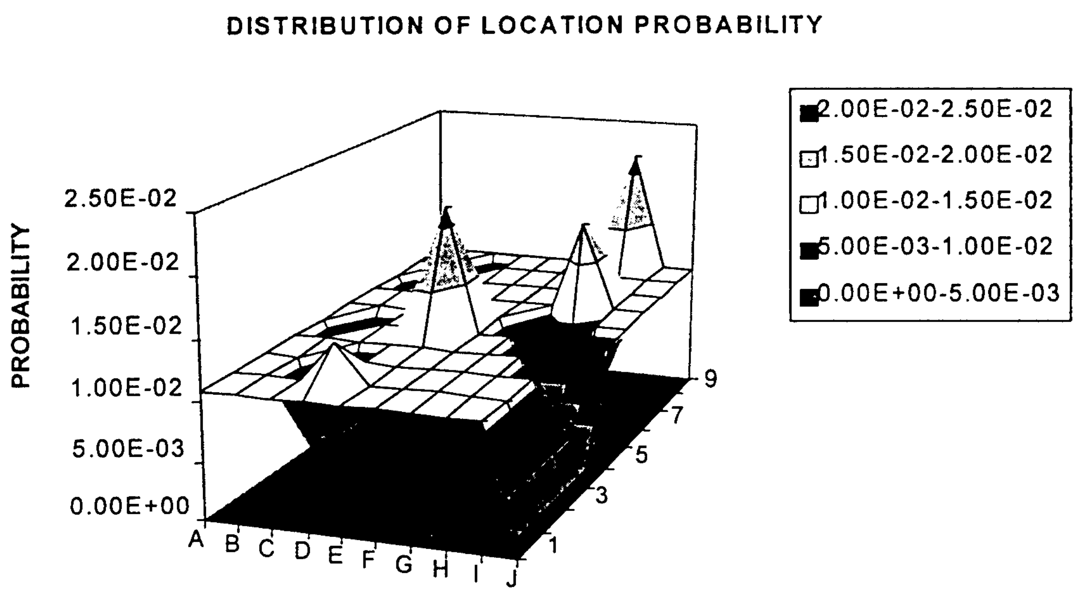

Figure 6. Posterior distribution of location, after search of the shaded cells.

As intelligence data are collected, the commander may gain information causing the values of the probabilities in his model of target presence in the cells to change. For example, if a terrain cell is scanned by a Blue sensor and a Red target is not detected, the probability the target is present in the cell should decrease from the value it had before the search. The process of changing the probability of target presence as a result of sensor activities is referred to as "updating" the probability distribution. An update of the prior distribution shown in Figure 5, to take into account a search of the shaded cells that indicated no target presence, is shown in Figure 6. The decrease in uncertainty from the prior distribution to the "updated" (more informed) posterior distribution reflects the amount of information gained from the report that a target is not present in the cells inspected. If information is received in sequence over time, a corresponding sequence of updates can be implemented, as we describe in more detail below.

I-3. Measuring Information Gain

We first consider the problem of measuring the information gained through reconnaissance in military operations. Traditional measures are usually based on detections of enemy forces by reconnaissance units or platforms [8]. Measures of effectiveness such as "percent of enemy vehicles detected," "time required to detect a tank company in hull defilade", and "average range at detection" do not give credit for reconnaissance efforts that suggest targets are *not* located in certain areas of interest to the battle commander. Finding that the enemy isn't located in a certain area can be of considerable value, and it is desirable to devise measures of information that quantify such results. As demonstrated in the example below, the information gain measure does take into account both indications of target presence and target absence in searched areas.

The method we propose can be described as follows. Compute the entropy of the prior distribution, $e_p = -\Sigma_{i\in S}\, p_i \ln(p_i)$. This may represent the uncertainty Blue has about the location of a particular Red system, for example. We compute Blue's total entropy or uncertainty about the locations of all Red systems by summing the entropy measures of the individual Red systems. (This is valid under an assumption of independence; in [2] we argue it provides a reasonable approximation of total entropy when unit locations are mildly correlated; an example is given in Section III-2.) The total entropy represents Blue's level of uncertainty about the location of Red units.

When one or more of the terrain cells are searched through reconnaissance, and the results are transmitted to the commander, he may gain information causing the values of the probabilities of target presence in the areas to change. A reconnaissance indication that a particular cell does not contain a target drives the location probability for that cell towards zero, accompanied by increases in the probabilities over all other cells. In several of our applications, it is possible to compute such changes using Bayes' formula [13]. This takes into account the capabilities of the recon system's ability to detect and its false alarm rate, both given as functions of the sensor, target, characteristics of the area searched, the search geometry, and sensor-target kinematics. Updating the target location probabilities by Bayes' formula is appropriate whether the target is located in a given search or not. Next, compute the entropy of the posterior distribution, $e_{p^*} = -\Sigma_{i\in S}\, p^*_i \ln(p^*_i)$. Typically (but not always) entropy will decrease with the changes in location probabilities associated with search effort, so the information gain, $(e_p - e_{p^*})$, usually will be positive. This measure can incorporate the combined effects of searches by many individual systems, each with a set of sensors looking in assigned or selected areas, working together as a system mix against an array of enemy targets.

Employing the information gain measure requires the analyst to model uncertainty in a way most suited to his or her particular application. As the discussion above illustrates, computing information gain is a simple summation of natural logs once the probabilities are known. The analyst's challenge is in developing and updating the probability distributions as the tactical process plays out.

The original prior distribution may be estimated or assumed, based on terrrain features (water bodies, terrain slope, etc.). Lacking that, one may begin with a uniform distribution over the sample space. Updating the distribution, however, may require modeling and analysis. The next section suggests several possible approaches to the updating process.

I-4. Approaches to Updating

Information Gain with Bayesian Updating

Suppose a battle area is considered to be composed of a large number of small cells C_1, C_2, ... C_N, and suppose reconnaissance or observation during combat can provide information implying a given cell C_j holds a given target, T, with detection probability

$P_D \in (0,1)$ (given $T \in C_j$). Similarly, suppose the false alarm rate for this recon platform on this target in this area is P_F. To simplify notation, let "$I(j)$" denote "recon information indicates T is in C_j," and let "$T(j)$" denote the event "T is in cell C_j." The current state of information, intel and recon about the location of T is represented by the current probability distribution for the location of T (which is the *prior* distribution for updating purposes). Let p_j denote the prior probability of $T(j)$: $j=1,2,\ldots,N$. We may use Bayes' formula to update the current distribution to take into account new information about whether T is in cell j.

To summarize: $P[\, I(i) \mid T(j)\,]$ depends on the scenario, recon tactics and capabilities of the sensors involved. We are assuming that, for the current search of cell C_j,

$P[\, T(j)\,] = p_j$;

$P[\, I(j) \mid T(j)\,] = P_D$; and

$P[\, I(j) \mid T(i)\,] = P_F$, $i \neq j$.

Then by Bayes' formula,

$$P[T(j)|I(j)] = \frac{P_D p_j}{P_D p_j + P_F(1-p_j)},$$

and

$$P[T(i)|I(j)] = \frac{P_F p_i}{P_D p_j + P_F(1-p_j)};\ i \neq j.$$

As a special case, relevant for Janus play of combat, suppose the false alarm probability of Blue's sensor system is zero. Then application of Bayes' formula gives

$P[T(j) \mid I(j)] = 1.0$;

$P[T(i) \mid I(j)] = 0.0$;

$$P[T(i)|\sim I(j)] = \frac{p_i}{1-P_D p_j};$$

and

$$P[T(j)|\sim I(j)] = \frac{(1-P_D)p_j}{1-P_D p_j}.$$

Here, "$\sim I(j)$" indicates the event "recon in cell j fails to detect the target."

To compute values of information gain resulting from specific recon activities, the following procedure can be used:

a. divide the region of interest into cells which might contain Red targets and which may be searched by Blue sensors;

b. determine Blue's prior probability distribution representing the marginal distribution of location of each Red target, before the search begins;

c. assume the search proceeds as a sequence of searches in designated cells, in specified time intervals;

d. when a set of cells has been searched, use Bayes' formula to update the current "prior" distribution of each target's location to obtain the posterior distributions for all targets;

e. compute the information gain, for each Red target, resulting from the search of the designated cells;
f. accumulate and store the sum of information gains for all Red targets, and the time of completion of the search of the specified cells;
g. loop through steps (d) - (f) for the duration of the search activity;
h. plot the composite (over targets) cumulative (over time) information gain as a function of time into the search. The result is the "information trace" (similar to the "battle trace" introduced in [4]), which gives an overview of the cumulative gain in information over time, as the search progressed.

Notes:

- One can plot the *rate* of information gain versus time to display how well the search activity is doing at various points in time.
- One can compute only the "end of battle," final cumulative information gain, if it is not desired to track gains (or gain rates) over time.
- It is easy to accommodate searches of sets of cells, rather than single cells, in each time period, using Bayes' formula in much the same way as shown above. We give some details in Section II-1.
- One can also measure the information gain due to receipt of information about the number of Red targets present. This can be done using the compound experiment properties described in Section III-3. We show an example of this in Section I-4 and discuss an application in Section II-3.

Application of Bayesian Updating to a Target Detection Process

Figure 5, on page 9, represents Blue's prior distribution of the location of a Red target in a hypothetical situation. The area of regard is partitioned into 100 cells, corresponding to row and column designations in the domain below the plotted surface. The "peaks" on the plotted surface represent spikes of probability over the corresponding cells, and can be interpreted as areas where the Blue commander feels Red is most likely to be. Likewise, depressed regions on the plotted surface represent cells where Blue believes Red is least likely to be deployed. Note that the relative probabilities represented by these spikes are actually discrete values over particular terrain cells. The smoothing into continuous peaks is a function of the graphical software for presentation purposes and is not represented in the calculations.

Further suppose that Blue deploys a reconnaissance sensor to search for the enemy system and the recon systems searches a path represented by the darkened terrain cells in the lower right portion of the battlefield. In this example we employ a reconnaissance system armed with a .83 probability of detection and no false alarm sensor. Given a Red system is in a particular terrain cell that the sensor searches, the probability the sensor detects the Red system is .83; given a target is not present the probability the sensor "detects" a target is zero.

Suppose the reconnaissance sensor progresses from cell to cell and does not find the Red system. Then the probability the enemy system is in the searched cells is driven towards updated values closer to zero, as shown in Figure 6, on page10. These new probabilities are calculated using Bayes' formula. In this example the Red system was not found during any of the searches of the shaded terrain cells by Blue's sensor. Even so, because there is information in sensor indications a target is *not* present in certain cells, the commander's uncertainty (and hence the entropy) decreased with the updated distribution, so there is positive information gain. The values of information gained as these cells are searched is shown in the following table.

CELL SEARCH #	INFORMATION GAIN
1	0.00557525
2	0.005598649
3	0.005622117
4	0.005645573
5	0.005669005
6	0.005692402
7	0.005715751
8	0.005739039
9	0.00576225

Information Gain With Combinatorial Updating

Suppose Red has N targets that are distinguishable with at least one of Blue's sensors, and the area of concern to Blue consists of R cells. If no cell could be occupied by more than a single target, Red could deploy his forces in any of ${}_RP_N = R!/(R-N)!$ (permutation of R areas taken N at a time) ways so this is the number of points in the sample space (set of possible Red states). From Blue's point of view, before recon begins (and lacking intel, prior knowledge, etc.) we assume each of these deployments is equally likely, so the initial entropy of Red's deployment from Blue's perspective is

$$e_0 = \ln({}_RP_N) = \sum_{j=R-N+1}^{R} \ln(j).$$

Now let us consider the information gain as recon proceeds from this starting point.

Case 1: Recon detects a target in one of the cells.

We now have $N-1$ targets in $R-1$ cells, so the entropy drops to $\ln({}_{R-1}P_{N-1})$ and the information gain is

$$\sum_{j=R-N+1}^{R} \ln(j) - \sum_{j=(R-1)-(N-1)+1}^{R-1} \ln(j) = \ln(R).$$

Note there is substantial information gain if the number of cells R is large and recon discovers a target. Note also this gain *does not depend on N*. Thus, under that stated assumption, from an information theory point of view, *the information gained in finding an enemy target is independent of the enemy's force size!*

Case 2: Recon determines a cell is target-free.

In this case, we have N targets in R-1 cells, so the entropy decreases to $ln({}_{R-1}P_N)$ and the information gain is

$$\sum_{j=R-N+1}^{R} ln(j) - \sum_{j=(R-1)-N+1}^{R-1} ln(j) = ln(\frac{R}{R-N}).$$

Note this gain is greater for N closer to R. Thus from an information theory point of view, *finding a cell to be target-free gives relatively greater information gain if Red has a very large force than if he has a small force.* Indeed, if Red's force is very small and is deployed over a large area (i.e., $N<<R$), finding that a specific cell contains no target gives little gain in information. Similarly, if Red's force is very large and is deployed over the same area, finding that a specific cell contains no target gives a greater amount of information gain.

Information Gain With Subjective Updating

In Section II-2 we discuss an application where information gain was computed using subjective estimates of an expert to update the probability distributions.

I-5. Non-Monotonicity of Entropy

Usually, entropy for a given target will decrease over time, as recon is conducted. However, there are cases in which there can be increases in entropy as additional recon occurs. This is caused by the fact that when a cell having high prior probability of containing the target is searched without success, the posterior may actually be projected toward a more uniform distribution, hence increasing entropy. This can be illustrated by a simple example involving one target placed at random in one of three cells and a sensor with detection probability *1/2*.

Suppose the prior vector is *(1/3, 1/3, 1/3),* so the initial entropy is e = 1.099. If a search of cell 1 fails to detect the target, the posterior distribution of target location becomes *(.2, .4, .4),* so entropy drops to 1.055. If a subsequent search of cell 2 is unsuccessful, the updated posterior is *(.25, .25, .50),* so entropy is further reduced to 1.040. So far, so good. The entropy has decreased steadily as recon has searched the first two cells. Now suppose a search of cell 3 fails to find the target. Then the updated posterior is *(1/3, 1/3, 1/3),* so entropy has suddenly *increased* back to 1.099. This is not a contradiction. In this case, the search of the three cells has not yielded information about the target location, so the information gain is zero. The recon system has been working diligently but has achieved absolutely nothing in terms of gaining information about the whereabouts of the target.

Numerical evaluations for examples similar to the foregoing show it is possible to have entropy oscillating between increases and decreases as recon is conducted. These situations are probably not of practical concern, since the largest effect in reducing entropy is associated with detecting and locating a target. It is to be noted again that nothing is wrong with increases in entropy as recon proceeds, under certain conditions. Indeed, a negative entropy decrease is correctly showing the amount of information lost through the recon conducted up to that point. Overall information gain may be nil or even negative if a large fraction of the available areas have been searched without success or if much time has passed since mobile targets have been located [3] (see Section III-5).

If entropy is used to measure progress on an optimal search for a target for which strong prior information is available, we expect entropy to *increase* as the search proceeds. For example, in a search for the sunken ship *SS Central America* [13], the prior distribution of the ship location was carefully developed, taking into account information from communications before the ship sank, survivor accounts of the sinking, and ocean currents and winds in the general area at the time of the sinking. This gave a prior that may be envisioned as a set of hills, where the target is thought to be most likely located under the areas where the hills are highest. The search was sequenced so as to search where the prior was highest, which should minimize search time required to locate the ship. As areas were searched, Bayes' formula was used to update the prior, and search was next directed to where the updated prior had the highest hill. As this process is followed, the search effort is directed so as to drive the updated prior toward a uniform distribution, hence in the direction of increasing entropy. It would appear that in this application, entropy could again be used to measure the progress of the search. In this case, however, the amount of *increase* in entropy in a time period might be a measure of the search progress. Of course, once the ship is located, the posterior distribution has a single spike of probability mass at the known location, and entropy drops to zero.

I-6 Combining Information Gain on the Number of Targets with Information Gain on Location

Recall that we are measuring number as well as location of enemy systems. Let us illustrate the compounding property described in Section III-3 by removing the assumption in the preceding example that it is known a target is present in one of the three cells. Now, suppose we believe there is a target present with probability 0.5, and no target present with probability 0.5. With each search of a cell, we must now update both the prior distribution of the number of targets present and the prior distribution of the location of a target, given one is present. Recall we are assuming the sensor used in the search has detection probability 0.5 and false alarm probability zero. The total entropy is computed using the relationship shown in Section III-3.

To illustrate for the first step (search of cell 1 without detecting a target), the posterior probability there is a target present in one of the cells may be calculated as follows, where we let "C1" denote the event "Cell 1 is searched and no target is found." I denote the (random) number of targets present, T denote "target position" and p denote the prior probability that a target is present:

$$P[I=1|C1]=\frac{P[C1|I=1]\cdot p}{P[C1|I=1]\cdot p+P[C1|I=0]\cdot(1-p)}$$
$$=\frac{\{P[C1|I=1,T\in cell1]\cdot(1/3)+P[C1|I=1,T\notin cell1]\cdot(2/3)\}\cdot p}{P[C1|I=1]\cdot p+P[C1|I=0]\cdot(1-p)}$$
$$=\frac{\{\frac{1}{2}\cdot\frac{1}{3}+\frac{2}{3}\}\cdot p}{\{\frac{1}{2}\cdot\frac{1}{3}+\frac{2}{3}\}\cdot p+(1-p)}=\frac{\frac{5}{6}p}{1-\frac{1}{6}p}$$

For p=1/2, this gives posterior probability a target is present equal to 5/11 (so the posterior probability that I=0 is 6/11). The location prior (1/3,1/3,1/3) is updated to (1/5, 2/5, 2/5) as shown below.

The total entropy before search is

$$e_I + e_{T|I} = \ln(2) + (1/2)\ln(3) = 1.242$$

and that after the first search is

$$e_I + e_{T|I} = -(5/11)\ln(5/11) - (6/11)\ln(6/11) - (5/11)[(1/5)\ln(1/5) + 2(2/5)\ln(2/5)] = 1.168.$$

Therefore the information gain as a result of the first search is 1.242 - 1.168 = 0.074.

Results for similar computations for searches of cells 2 and 3 in turn (each resulting in a "no target present" indication) are summarized in the following table:

Stage	p	loc. dist'n	e_I	$e_{T\|I}$	e	inf. Gain
Prior	1/2	1/3,1/3,1/3	.6931	1.0986	1.2424	-
Search 1	5/11	1/5,2/5,2/5	.6890	1.0549	1.1685	.0739
Search 2	4/10	1/4,1/4,2/4	.6730	1.0397	1.0889	.0796
Search 3	7/19	1/3,1/3,1/3	.6581	1.0986	1.0629	.0260

Note that total entropy, taking into account uncertainty in the number of targets present, is monotone decreasing as the search proceeds, in contrast with the case for $e_{T|I}$ discussed in the preceding example.

I-7. Extensions and Further Research

Further work is needed concerning the information gain on a mobile target as time since detection increases. In Section III-5 we present an example involving a crude bivariate normal model of the posterior distribution of location as a function of time since detection. It appears quite easy to incorporate more realistic models of possible target movement in the Bayesian updating of location distributions. This would give useful models of losses in information (negative information gains) over time, with applications such as optimal scheduling of resources to reestablish target location.

Several analysts have suggested the weighting of location information by "importance factors" representing some attribute of interest to the Blue commander. For example, the Blue commander may wish to weigh location information about a target that represents a threat to his own forces more heavily than that for a non-threatening target. It would be quite feasible to devise a weighted information measure, provided a credible weighting scheme could be developed.

We have discussed information gain primarily in terms of knowledge about locations and numbers of enemy targets. But the principle can be applied to any sort of "unknown" quality, provided it is feasible to list possible values of the quality, and that it makes sense to model state of uncertainty about the quality in terms of probability distributions. For example, commanders talk about gaining information about "the enemy's intent." If one could list the sample space of plausible enemy intents, place a prior distribution over this space, and update the distribution as the scenario plays out, then information gain about enemy intent could be measured.

PART II. APPLICATIONS OF INFORMATION GAIN

We describe three applications we have carried out in experiments here at the U.S. Military Academy, and discuss an on-going project aimed at implementing an information gain measure of effectiveness with the Janus combat simulation. These examples illustrate the variety of methods that can be used to "update" prior distributions. The first involves search for targets by two competing reconnaissance systems, played in the Janus simulation. It uses Bayesian updating described in Section I-4 in connection with a target detection process. The second application uses target location "probability contour maps" developed by an expert at several stages of an experiment designed to investigate how information links to combat success. The third application uses continuous probability distribution models of target location with a single stage of updating based on information conveyed by intelligent mines.

II-1. Comparing Reconnaissance Systems [3, 5]

We employed information gain as a measure of reconnaissance results in conducting a modest Janus driven experiment designed to compare two recon platforms, crudely representing the Comanche helicopter and an unmanned aerial vehicle (UAV) [5]. The experiment involved a Bosnian scenario developed by LTs Carroll, Glaser, and Mitchell, who, as cadets, worked along with one of the authors in the Operations Research Center (ORCEN) at the US Military Academy. The cadets carried out a series of experiments using the Janus combat simulation. They collected data and carried out data reduction using the ORCEN facilities as part of a capstone course in Systems Engineering. Each simulated recon battle lasted ten minutes of game time and involved a single recon platform searching for 50 identifiable targets hidden among 400 500m X 500m terrain squares, or "cells." The recon systems were able to search 261 of these cells in each trial, following the assigned routes in the scenario.

The entropy associated with each individual target was computed at times 0, 1, ..., 10 minutes, and the total entropy was calculated as the sum of the individual target entropies. The following assumptions were made:

1. As far as Blue knows, each Red target could be placed in any of 400 cells by Red. Actually, Red has placed all 50 targets in cells that will be sear ched by Blue (i.e., somewhere within the set of cells Blue will search during the recon battle).
2. For Janus runs, the false alarm probability, P_D, is zero (i.e., Janus does not play false alarms by weapon system sensors).
3. Target locations were independent, from Blue's point of view, and targets were stationary.
4. Each recon system had detection probability at least 0.05 against each Red target.
5. Blue had no initial information about target location and thus the prior distribution was taken to be uniform over the 400 possible cells involved.

Therefore, the starting entropy for each target (at time zero) was *ln(# cells)* $= ln(400) = 5.991$.

For each individual target, the following comments hold:

1. Only 261 cells were searched by Blue during the recon battle.
2. With false alarm probabilities equal to zero for each recon system, entropy drops to zero when the target is detected and located (because the posterior distribution of the target's location then becomes a vector of the form $\mathbf{1} = (0,0,...,0,1,0,...,0))$, and remains at this value (stationary target).
3. The detection probability of a given recon system against a given target was taken to be the relative frequency of detections in the ten Janus runs with that system. If a given target was never detected in the ten runs, the detection probability was set equal to 0.05.

Method of Bayesian Update Used in the Janus Experiment

The probability vector $(p_1, p_2, ..., p_{400})$ was updated at the end of each minute, using Bayes' formula. As mentioned above, if the target was detected and located during a minute period, the posterior distribution is of the form $(0,0,...1,...,0)$, so the entropy for that target drops to zero at that point in time. If cells in a set $K = \{k, k+1, ..., k+m\}$ were searched during the minute and the target was not detected, the posterior was computed as follows.

Let $T(j)$ denote "target in cell j," and $I(K)$ denote "target found in the set $K=\{k, k+1, ..., k+m\}$." Let p be the detection probability and p_j be the prior probability of the event $T(j)$, as before.

Case (a): posterior for cell j, j∉K.

$$P[T(j)|\sim I(K)] = \frac{P[\sim I(K)|T(j)]\cdot p_j}{\sum_{j\notin K} P[\sim I(K)|T(j)]\cdot p_j + \sum_{j\in K} P[\sim I(K)|T(j)]\cdot p_j}$$

$$= \frac{p_j}{\sum_{j\notin K} p_j + \sum_{j\in K}(1-p)p_j} = \frac{p_j}{D}$$

Case (b): posterior for cell j, j∈K.

$$P[T(j)|\sim I(K)] = \frac{P[\sim I(K)|T(j)]\cdot p_j}{\sum_{j\notin K} P[\sim I(K)|T(j)]\cdot p_j + \sum_{j\in K} P[\sim I(K)|T(j)]\cdot p_j}$$

$$= \frac{p_j(1-p)}{\sum_{j\notin K} p_j + \sum_{j\in K}(1-p)p_j} = \frac{p_j(1-p)}{D}$$

where D is the common value of the denominator in the last expressions in the two cases.

The computation of the posterior distribution is easily accomplished by exploiting the fact that the denominator D is the same in both cases above. We proceed as follows: for all $j \in K$, in the current prior probability vector, replace the current prior probability the target is in cell j, p_j^*, by $p_j^* \cdot (1-p)$, where p is the detection probability of the given recon system against the target in question. Then sum the elements of the resulting vector and unitize the vector by dividing each element of the vector by the sum. This vector is the current posterior at the time point in question, and it becomes the prior for the succeeding time period. Note *the same posterior results if one imagines the cells were searched one at a time, in any given order, and the posterior was computed in a sequence of corresponding "one-cell" updates.*

For each target at each time point in each run with each recon system the entropy is either zero (if the target has been detected and located) or the value $e = -\Sigma p_j^* \ln(p_j^*)$, the summation extending over the 400 elements corresponding to the 400 boxes available. A simple computer program was written to carry out these computations.

Plots of the average entropy value, over ten Janus runs, are shown in Figure 7, for the UAV and Comanche recon platforms. Plots of information gain, (averaged over ten Janus runs) for the UAV and Comanche are shown in Figure 8. The similarity in shapes of the information gain plots for the UAV and Comanche indicates both systems were performing best around minutes 2 to 4 in our scenario, with another period of increasing performance near the end of the recon battle. These observations are in accord with results expected by the experimentation team, based on details of the scenario design and experience with other simulated recon battles using Janus. Note the plot of information gain for the UAV is considerably higher than that for the Comanche, indicating the UAV performed better in this scenario. This is a counter-intuitive result, and subsequent investigation revealed this might have been caused by inaccuracies in the Janus modeling of the Comanche.

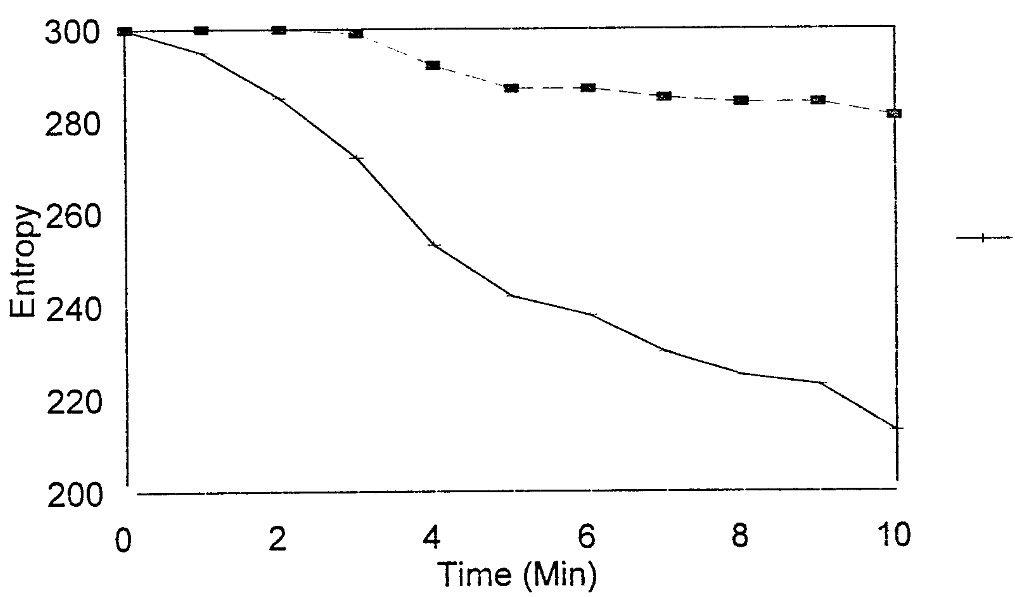

Figure 7. Average Entropy for the UAV (lower curve) and the Comanche (upper curve) in ten Janus runs.

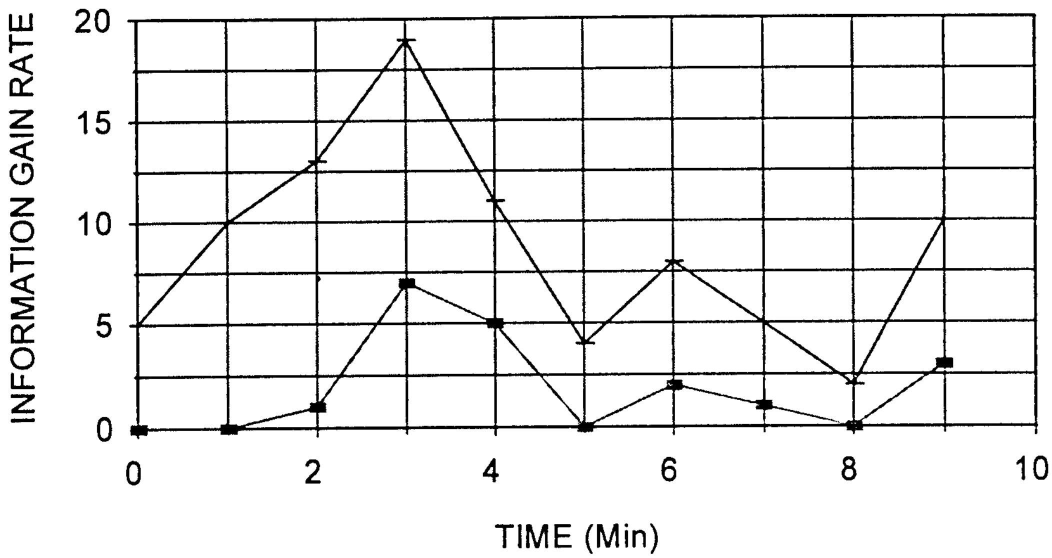

Figure 8. Average information gains for the UAV (top curve) and Comanche (bottom curve) in ten Janus runs.

II-2. Studying Relationships Between Tactical Intell And Battle Results

We investigated the effects on combat results of varying levels of information a combat commander has about his adversary [1]. We performed an experiment in which individual subjects, playing the role of task force commander, developed detailed plans for conducting operations against an enemy defender. Each commander ultimately prepared five combat plans for conducting the same operation against the same enemy force, but with increasing levels of information about the enemy's composition and disposition. We designed these information levels to correspond closely to doctrinally realistic increments of information available during the planning process.

For each phase of the experiment we gave subjects the respective information set. We required subjects to produce the following products of their battle planning exercise: 1) Task Organization, 2) Concept Sketch and Graphics, 3) Fire Support Plan, and 4) Synchronization Matrix. Once a subject turned in his plan we issued him the additional information set for developing his next plan. We then entered the subject's plans in the Janus Simulation Model. Once a subject completed all five plans, we conducted ten Janus runs with each plan. Subjects were not allowed to see results of these runs until the experiment was completed.

Measures of Effectiveness

We captured data to support computation of over a dozen measures of effectiveness (MOEs) in order to measure the effectiveness of simulated combat operations. Below we summarize results for only two MOEs: Blue Losses (BL) and Number of Combat Vehicles on the Objective (VO).

Computation of Information Gain

To compute information gain, we estimated probability distributions of Red unit locations for each phase by employing an expert. Our expert used the cumulative information set available at each phase, knowledge of Army doctrine, and knowledge of the terrain to construct the probability distributions. In this regard our expert played a role very similar to that of the intelligence officer preparing decision support products for the commander. The final product of our expert's estimate was what we called a "probability contour map" (PCM).

As its name implies, the probability contour map (PCM) partitions the total area of operations into sub-areas having given relative probabilities of containing a Red unit. Just as elevation contours on a topological map display areas of (approximate) given distinct elevation, a probability contour map uses probability contours to display areas of fixed distinct probability. In his expression of relative likelihood of Red unit locations, our expert first expressed the likelihood of containing a Red unit as a categorical variable taking values: 1) very unlikely, 2) unlikely, 3) likely, 4) very likely. We represented these categories of likelihood numerically as 0, 1, 4, and 9, respectively. This numerical scale is a subjective assignment; in [1] we discuss the lack of sensitivity of entropy decreases to changes in this representation.

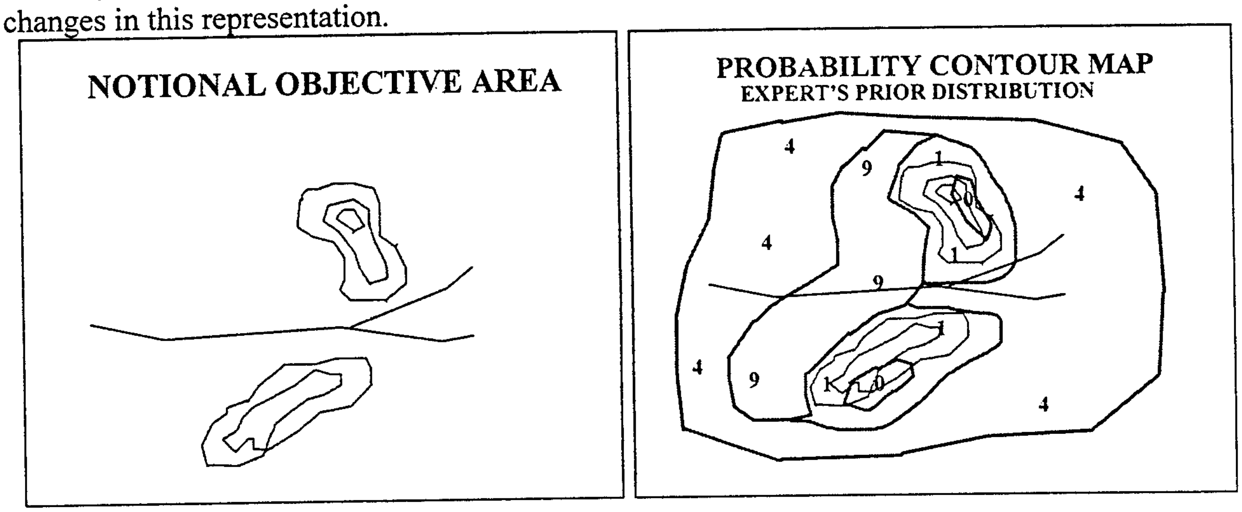

Figure 9. Hypothetical prior PCM (right) related to dessert terrain data depicted graphically at the left.

For our particular scenario, we developed PCM's for fighting systems and separate PCM's for obstacles. The process of building a PCM is analogous to developing the modified combined obstacle overlay. As an example, consider the case where nothing is known about the enemy. In this case the assignment of probabilities may depend entirely upon the terrain, as represented in maps of the area of operations. We asked our expert to consider each portion of the terrain and answer questions such as, "If the enemy had tanks, what is the likelihood they would be deployed here?" The answers to such questions determine the relative likelihood assignment of the particular area (0, 1, 4, or 9 for our application). For example, areas of terrain that are obviously unusable by a particular type of unit will receive a likelihood of zero. Areas of terrain that are obviously key terrain will receive a weighting of 9, and so forth. For each phase, our expert continued this process until the entire area of interest had been assigned a relative likelihood value. A hypothetical prior PCM is shown in Figure 9.

The area of interest in our application was a 10 km by 14 km zone bounded by the line of contact and friendly maneuver graphics. As the intelligence and reconnaissance process progressed we developed subsequent PCM's in the same way. Once the expert learns the location of one tank he may change his assessment of the probability distribution of the location of other tanks. The same holds for obstacles. Knowing the location of one obstacle helps one deduce the location of others and thus update the PCM for obstacles. Likewise, knowing the location of an obstacle helps one estimate the locations of vehicles and knowing the location of vehicles helps one estimate the location of obstacles. Thus the PCM's for fighting systems and obstacles are interdependent. A hypothetical posterior PCM representing an update of the prior shown in Figure 9, upon receipt of intelligence confirming the locations of two enemy tanks, is shown in Figure 10.

Figure 10. Posterior PCM (right) reflecting intelligence depicted at the left.

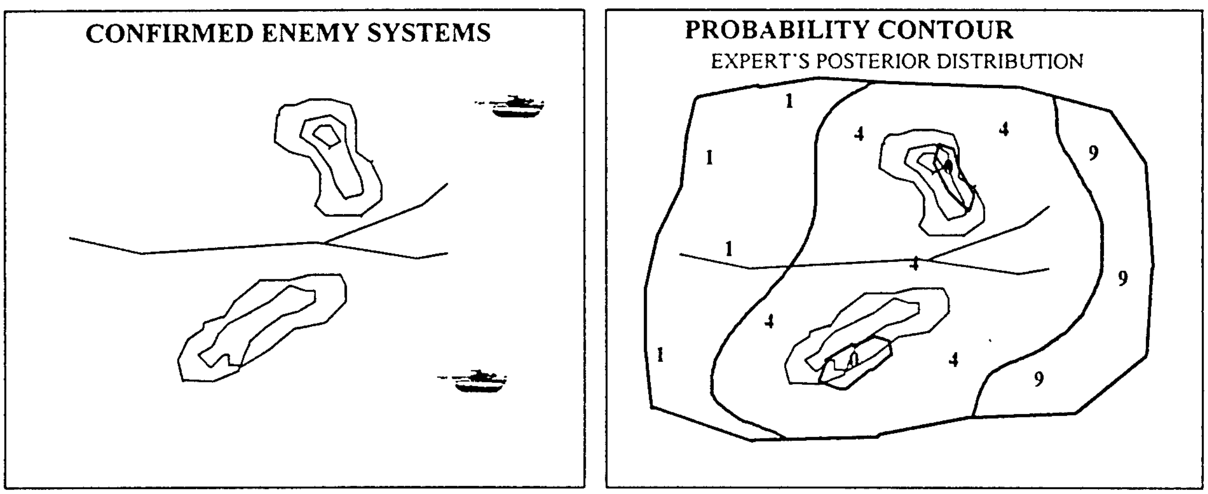

Determination of Target Density and Approximating Mass Function

Let R0, R1, R4 and R9 denote the regions over which the bivariate density described above has value proportional to 0, 1, 4, or 9, respectively, and let A(R0),...A(R9) be the areas of these regions. Estimates of these areas were obtained as follows. We estimated the areas of the portions of the four regions falling within each 1 km square, to the nearest .1 km^2, and then individually summed these estimates for each region type over the 140, 1 km squares comprising the area of interest. We felt unjustified in attempting to estimate the areas of the four region types within any 1 km square to any resolution smaller than .1 km^2. This was due to the problem of visually estimating such areas. Note, for example, the Janus display of a Red unit location uses an icon that, with location error, locates a precise unit location only to a resolution of about .1 km^2. Thus, there is some error in our area estimates, but we believe they are sufficiently accurate for entropy calculation purposes. (We report results of a modest sensitivity analysis below.) In what follows, we refer to the imaginary .1 km^2 sub-regions of each 1 km square as "cells."

We determined a bivariate density function over the 10 km x 14 km area of interest so that:

- the integral of the density over the 10 km x 14 km area of interest equals 1.0;
- the density is constant over each region R_i; and
- the ratio of the density at a point in R_i to that at a point in R_j is b_i / b_j, where b_i and b_j are elements of the set {0,1,4,9}.

It follows that the density function value (height) at any point within R_i is

$$b_i \Big/ \sum_j b_j \cdot A(R_j)\text{; for } b_i = 0,1,4,9.$$

Now consider a discrete approximation of the forgoing density function, based on the .1 km^2 cells described above. We note the probability a particular Red unit is located in a given cell within region R_i is

$$p_i = (.1)b_i / \sum_j b_j A(R_j) = .1b_i / K,$$

where we let K denote the (constant) value in the denominator. The approximating mass function is defined to have values equal to the p's at the center points of the corresponding .1 km^2 cells. This mass function therefore is defined at 1400 points within the area of concern. Note it has fixed value p_i on $10 \cdot A(R_i)$ of these points.

Let us model the position of a given Red unit as having this discrete distribution over these 1400 cell centers. Then the entropy measure of Blue's uncertainty about this location is

$$e = -\sum_{i=1}^{1400} p_i \ln(p_i) = -\sum_{j=1}^{4} \frac{.1b_j}{K} \ln\left(\frac{.1b_j}{K}\right) 10 \cdot A(R_j)$$

$$= -\sum_{j=1}^{4} \frac{b_j A(R_j)}{K} \left[\ln(.1b_j) - \ln(K)\right]$$

$$= \ln(K) - \sum_{j=1}^{4} b_j A(R_j) \ln(b_j) / K - \ln(.1)$$

$$= \ln(K) - \frac{L}{K} - \ln(.1),$$

where $L = \sum_{j=1}^{4} b_j A(R_j) \ln(b_j)$.

It can be seen that the term -ln(.1) = ln(10) = 2.3 is related to our division of each square km into 10 cells, in the formation of the discrete approximation of the density of Red unit location. The approximation of the continuous density by a discrete mass function with "resolution" 1/10 km^2 introduces the constant term ln(10) into the entropy value. This may at first seem troubling, because the level of resolution employed in the discrete approximation step is somewhat arbitrary; we could have used cells of area .01 km^2 and gotten entropy values differing by a constant value involving ln(100), for example. However, our application involves taking the *difference* of the calculated entropy at two successive phases to be the information gain between the phases. The constant ln(10) adds out (as would the constant corresponding to any fixed level of resolution in the approximation) when the decrease in entropy is computed. Therefore, for our application, the level of resolution has only the minor effect of adding some noise in estimating the values of the $A(R_i)$, as mentioned above. Note this is closely related to the behavior of information gain when continuous distributions are approximated by discrete distributions, discussed in Section III-4.

Information Gain Calculation

To complete the computation, we determined the entropy for a single Red unit of a given type, multiplied by the number of units of that type (42 for vehicles and 22 for obstructions), then added these values to obtain the total entropy for the phase in question. A summary of the total entropy at each phase and the information gain from the previous phase is shown in Table 1. The rightmost column entries are actually cumulative information gain; we used the label "cum. inf." for simplicity here.

Phase	Total Entropy	Information Gain	Cum. Inf.
maximum*	463.63		
1	299.23	164.40	0
2	290.79	8.43	8.43
3	277.03	13.76	22.19
4	267.54	9.49	31.68
5	224.98	42.56	74.24

Table 1. Total entropy and information gain by phase of the experiment.
* Based on a uniform distribution over 1400 cells.

Some Results

Because one subject appears to have performance markedly different from most of the others, we estimated the over-all (subjects) mission accomplishment profile by the median VO (vehicles on objective) response at each information level. A plot of median VO against information level (as measured by cumulative information gain) is shown in Figure 11.

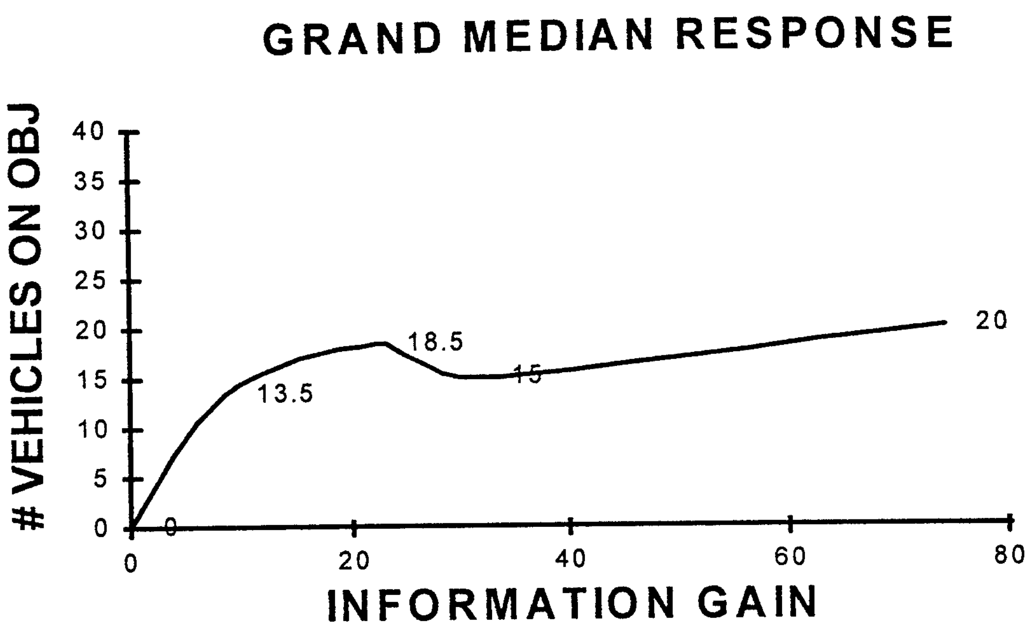

Figure 11. Plot of median VO versus information gain.

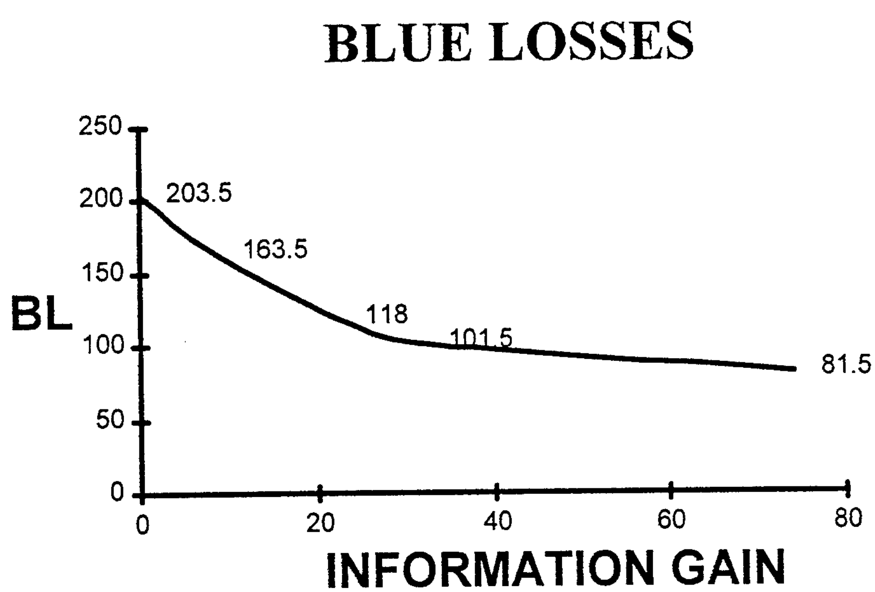

Figure 12. Plot of median blue losses (BL) versus information gain.

Figure 12 shows a plot for the MOE, "Blue Losses," which is a resource consumption MOE. The plot shown in Figure 11 suggests there is not significant improvement in the ability of the Blue commanders to achieve success in the mission objective, as information level increases beyond that available at about the third phase.

However, there is continuing decrease in blue losses over the entire span of five phases, as shown in Figure 12. This suggests that *a commander does not require information beyond a moderate level in order to accomplish his mission, but he can achieve mission success at reduced cost when he has additional information*

II-3. Intelligent Minefield Design

Possible allocations of "intelligent mines" (IMs) were investigated in connection with a systems engineering design effort by faculty and cadets at USMA involving intelligent minefield design [6]. The intelligent mine is an anti-armor system which incorporates features of the "wide area mine" (WAM) with the addition of "smart" engagement planning and communications capabilities. Both the WAM and IM sense and attempt to track targets using acoustic and seismic sensors. If a target is deemed by these mines to be within range, they attack the target. The IM has the added ability to communicate with other Ims. If two or more IMs simultaneously track the same target, then the tracking accuracy is considerably improved, resulting in higher hit probability against the target. The IM can also communicate to the Blue commander limited information about approaching targets (number of targets sensed, approximate locations and whether a target was engaged by the IM).

We developed a computer simulation (in Visual Basic) which could be used to evaluate the relative performance characteristics of alternative anti-armor minefield configurations. Each configuration was defined in terms of the numbers of conventional mines (CMs), WAMs, IMs, and artillery-deployed mines (FASCAMs) and their placement in a square minefield 1.5 km on a side. Data generated by the simulation supported computation of a variety of measures of effectiveness, such as "fraction of attacking tanks killed," "average delay time," and "average penetration distance of an attacking tank." Since one of the presumed advantages of the IM over the WAM is its ability to convey information to the Blue commander, it was deemed desirable to devise measures of effectiveness appropriate for measuring this feature.

We applied the information gain measure, based on Bayesian updating in a single, "end of battle," stage. Two types of prior distributions were devised; one representing the prior distribution of location of each attacking tank and the other representing Blue's prior notion of the number of Red tanks that might be approaching the minefield. Posterior distributions were computed, in a single stage, for the "post attack" picture. During an attempt by a platoon of Red tanks to traverse the minefield, only a fixed set of response types by the IMs was possible:

1. no IM detected a Red tank;
2. WAM or CM detonations were sensed, but no IM tracking;
3. WAM or CM detonations were sensed, some IM tracking occurred with possible engagements, but no simultaneous IM tracking; and
4. two or more IMs track and engage targets, in simultaneous pairs at one or more points in the attack.

In each of these cases, certain inferences can be drawn about the number and locations of targets based on information transmitted to Blue by the IMs. We developed simple models of the posterior distributions, using elementary conditioning to update the distribution of the number of Red tanks and assigning "equivalent areas of resolution" for the locations of the Red tanks. For example, in the first case above, no information about the attacking Red tanks was transmitted to the Blue commander, so the posterior distributions are the same as the prior distributions, and the information gain is zero.

In the second case, Blue can surmise that the number of targets in the minefield area is at least some number d*, based on the number of detonations sensed and the probabilities tanks were killed as a result. Blue also gains location information for this number of targets equivalent to locating them with a uniform distribution over a certain fraction of the minefield area (which depends on the number of detonations and the number of IMs present in the minefield). If $d^* < 4$, the remaining tanks (of the platoon of four tanks) are assumed to have location distributions uniform over a region the size of the minefield. The posterior distribution of the number of Red tanks is taken to be the conditional prior distribution, given at least d* tanks are attacking the minefield.

The third case is similar to the second. The posterior distribution of the number of attacking tanks is just the conditional distribution described for case 2. With IM tracking, the location of d* tanks is captured within a uniform distribution over a circle with radius related to the tracking radius of the IM and the number d*.

In the fourth case, again we simply conditioned the prior distribution of the number of Red tanks on the event [# tanks >= d*]. Due to the improved location accuracy with simultaneous tracking by IMs, we assume the posterior distribution of d* tanks is uniform over circles of radius related to the lethal radius of an IM, the number of IMs present, and d*.

The initial and final values of entropy were calculated using the relationship $e = E_T e_{X|T} + e_T$ discussed in Section III-3, where "E_T" denotes expectation with respect to the appropriate distribution (prior or posterior) of the number of attacking tanks, T, and $e_{X|T}$ denotes the joint entropy of location (X) of T tanks, where $T <= 4$, since there are four tanks in the attacking platoon. The conditional entropy for the prior situation was computed as the sum of the marginal location entropies of the four attacking tanks. The posterior value was computed in a similar manner. The sum of the marginal entropies of tank locations over the four attacking Red tanks given T was taken to be

$e_{X|T} = \min\{T, d^*\} \cdot$[entropy over reduced area for the given case] +

$(4 - \min\{T, d^*\}) \cdot$[entropy over area of minefield].

These entropies are computed with assumed uniform posterior distributions over the "reduced" areas described above.

II-4. Information Gain (MOE) in a Janus Postprocessor

We are currently working on a project to incorporate information gain as a measure of effectiveness available to users of Janus [11]. We hope to incorporate this measure within a post-processor being developed at USMA, called "Jets." The following discussion outlines the computational approach we are following, which is an extension of the presentation in Section II-1. Our goal is to exploit searches by sensors on Blue vehicles as a battle is simulated within Janus. We hope to obtain a non-intrusive capture of detection probabilities for each cell of the battle area, for each Blue sensor, for each time period. The cells are those employed by the Janus software; a typical battle area includes about 2,000 cells, so we are dealing with updating each component of a 2,000-element probability vector at the end of each time period.

Now, consider the case (relevant for Janus computations) in which the probability of detection is a function of Blue's sensor, s, and any cell, c, in which it looks sometime during the time increment. Denote this probability by $D_{s,c}$. Moreover, let $D_{s,c} = 0$ for any cell c not "inspected" by the given sensor, s, during the time interval. The probability of non-detection by all sensors looking in the j^{th} cell is the product of the probabilities of non-detection by each sensor looking in that cell in the given time period, assuming independence among the sensors. Let p_c denote the prior mass in cell c. Then the posterior probability vector for the target in question, given it was not located in the time increment under consideration, is found by unitizing[2] the vector whose c^{th} element is

$$\Pi_{sensors,s}(1 - D_{s,c})p_c$$

using the convention mentioned above for cells not inspected by the various sensors. This posterior updating can be carried out in one operation for all entries in the prior vector (corresponding to cells making up the battle area). Thus, if $\mathbf{p}_t$ denotes the prior vector at time t (a stochastic vector having k elements), and $\mathbf{d}_{t+1}$ denotes the non-detection probability vector for the t^{th} time interval, whose elements are composed of the values $\Pi_s(1-D_{s,c})$, then the posterior at time t+1 is $\mathbf{p}_{t+1} = \mathbf{p}_t \otimes \mathbf{d}_{t+1} / | \mathbf{p}_t \otimes \mathbf{d}_{t+1} |$ where "$\otimes$" denotes component-wise multiplication and "$| \cdot |$" denotes the sum of components in the vector involved (so this division constitutes unitization of the vector $\mathbf{p}_t \otimes \mathbf{d}_{t+1}$). As mentioned before, this holds only for targets not located in the time interval; otherwise the posterior is of the form $\mathbf{1} = (0,0,\ldots,1,0,\ldots,0)$, where the "1" is in the location corresponding to the cell in which the target was found.

Updating and Incorporating Information about the Number of Red Units

We now exploit the compound experiment result presented in Section III-3 to incorporate uncertainty the Blue commander has about the number of Red assets as well as their locations.

For a given type of Red asset, say the i^{th} type, let T_i denote the number of such assets placed in the Blue commander's area of concern by the Red commander. From Blue's point of view, T_i is unknown; Blue considers a random variable A_i representing possible values of T_i. Thus from Blue's point of view, A_i has a distribution defined by a mass function p_i which assigns probabilities to the possible numbers 0, 1, 2, ... of type i assets present in his area. Then the total entropy for type i units, e_i, is given by

$$e_i = E_{Ai}[e_{\mathbf{x}|Ai}] + e_{Ai} ,$$

where $\mathbf{x}$ is the position vector of the Red assets. Applying the conditioning argument, and assuming independence among targets, gives total entropy for type i units,

$$\begin{aligned} e_i &= E_{Ai}[T_i \cdot e_{\mathbf{x}|1}] + e_{Ai} \\ &= T_i\, e_{\mathbf{x}|1} + e_{Ai} , \end{aligned}$$

where e_{Ai} is the entropy of the (current posterior) distribution of A_i. In this equation, $e_{\mathbf{x}|1}$ could be computed as discussed above, T_i is known[3], and e_{Ai} is easily computed from the

[2] A non-zero vector with non-negative components is *unitized* by dividing each component by the sum of the components. This process transforms a vector proportional to a probability vector into a probability vector (i.e., a vector whose components form a probability mass function).

[3] The value of T_i is the total number of Red assets of type i that remain undetected by Blue. We note the value of T_i will generally decrease as time into the battle increases. As noted earlier, the location entropy of detected assets is zero, which accounts for the decrease in T_i after each detection of a type i asset.

posterior distribution of A_i, formed by updating the prior distribution at the end of each time period.

Bayesian Updating of the Distribution of the Number of Red Units

If m units of type i are detected in a time period Δt, then the prior distribution of A_i at the beginning of that period, p_i, can be updated using Bayes' formula as before. Let us illustrate the essentials of this argument for the first time period, for which the prior distribution of A_i is the mass function $p_i(x) = P[A_i = x]$; $x = 0, 1, 2, ..., T_i$. If Blue's sensors looked in s cells during the first time period and m units of type i were detected, we wish to compute the conditional probabilities $P[A_i = x \mid \text{m units detected}]$; $x = m, m+1, ..., T_i$. This posterior distribution would then become the prior going into the second time period. Note the conditional mass function evaluated at any of the values 0, 1, ..., m-1 is equal to zero, for Blue has detected m such units.

If Blue's sensors looked in s cells during the period Δt and did not detect any units of type i then a crude approach to computing the posterior of A_i could be based on combinatorial arguments. For the purposes of updating p_i, imagine (temporarily) that the Red units of type i are uniformly distributed over all the cells, say N in number, of the area of regard. Then the number of units of type i present in a sample of s cells out of N is hypergeometric distributed. Since the number of cells sampled is generally quite small relative to the total number of cells present, this distribution is approximated well by a binomial distribution with parameters s and T_i/N, where T_i is the total number of Red units of type i that remain undetected by Blue, and N is the total number of cells in the area of concern to the Blue commander [14].

Given no type i units were seen in searching s cells in the time interval Δt, the conditional probability that $A_i = k$ can be approximated using a binomial distribution and Bayes' formula. The binomial model is used to calculate the probability no type i unit was detected, given s cells were searched (no "successes" in s "trials", with a certain success probability per trial. Specifically,

$$P[A_i = k|none\,seen] \propto (1-\frac{k}{N}D)^s\, p_i(k)\ ; k = 0,1,2,\dots,$$

where D denotes the detection probability in each of the s cells searched (by all the sensors looking in the cells). The "one-trial" success probability parameter in the binomial distribution is approximately equal to kD/N because, given k out of N cells contain targets, and assuming equally likely locations, the probability a randomly selected cell contains a target is k/N, and the probability the target is detected given it is present is D. Under these assumptions, the overall probability of detecting a target in a given cell is [k/N]D.

The argument for finding the posterior given m units of the type in question were detected in the first time period is similar, and Bayes' formula gives posterior density values proportional to the binomial mass evaluated at m, multiplied by the prior value. In the more realistic case where the detection probabilities vary over cells and the units aren't uniformly distributed over the N cells, the expression above should be revised. We can model the probability a target is found in a given cell c by $T_i \cdot p(c)$, where T_i is the total number of undetected assets of the given type and p(c) is the current prior probability a given asset of this type will be in cell c. Similarly, the detection probability in cell c by any Blue sensor can be written in the form $1 - [\Pi_{sensors,\, s}(1 - D_{s,c})]$, where $D_{s,c}$ is the probability of detection of a target in cell c by sensor s, given a target of the given type is actually in the cell. It then follows the expected number of detections may be modeled by

$$T_i \cdot \Sigma_{cells,\, c}\, p(c)[1 - \Pi_{sensors,\, s,\, looking\ in\ given\ cell,\, c}(1 - D_{s,c})]$$
$$= T_i \,|\, \mathbf{p} \otimes (\mathbf{1\#} - \mathbf{D'}) \,|$$
$$= T_i\, [\, |\, \mathbf{p} \otimes \mathbf{1\#} \,| - |\, \mathbf{p} \otimes \mathbf{D'} \,|\,]$$
$$= T_i\, [\, 1 - |\, \mathbf{p} \otimes \mathbf{D'} \,|\,],$$

where **1#** is the N-component vector of 1's, **D'** is the vector of expressions of the form $\Pi_{sensors,\, s,\, looking\ in\ given\ cell,\, c}(1 - D_{s,c})$, and **p** is the vector of location probabilities (with components p(c); c = 1, 2, ..., N). Note the value of $|\, \mathbf{p} \otimes \mathbf{D'} \,|$ is available from the computation of the posterior of **p**, as discussed above.

PART III. DEVELOPMENT AND PROPERTIES

III-1. Characterization of the Information Gain Function.

We show the information gain function, $\delta(\mathbf{p},\mathbf{p}^*)$, must have the form of entropy decrease, under the assumption of four plausible conditions (which are discussed at the end of this section). This result implies there is no room, under our assumptions, for the question, "Why use this specific mathematical definition for measuring information gain?"

Notation: Let S be a finite sample space, and let Ω be the set of all (discrete) mass functions over S. We denote any "uniform" distribution in Ω having exactly n non-zero mass values equal to 1/n by the symbol "**n**", let **p**, **p***, and **q** be arbitrary members of Ω, and suppose X, Y, and I are jointly distributed random variables on S. We denote the mass values in **p** by $p_1, p_2, \ldots p_n$ and similarly for **p***.

Theorem (Characterization of the information gain function):
Let $\delta(\mathbf{p},\mathbf{p}^*)$ be a function mapping $\Omega\times\Omega$ into the set of real numbers, satisfying the following four axioms:

(1) δ is continuous;
(2) for any fixed $\mathbf{p}^*$, $0 < n < m$ implies $\delta(\mathbf{n},\mathbf{p}^*) < \delta(\mathbf{m},\mathbf{p}^*)$;
(3) for a compound experiment I followed by X (denoted "I;X"),
$$\delta(I;X,\mathbf{1}) = \delta(I,\mathbf{1}) + E_I\delta(X,\mathbf{1} \mid I)^{4};$$
(4) for any $\mathbf{q} \in\Omega$, $\delta(\mathbf{p},\mathbf{p}^*) = \delta(\mathbf{p},\mathbf{q}) + \delta(\mathbf{q},\mathbf{p}^*)$.

Then
$$\delta(\mathbf{p},\mathbf{p}^*) = k[\Sigma p^*_i \ln(p^*_i) - \Sigma p_i\ln(p_i)],$$
where the summation is over positive masses in **p*** (**p**, respectively) and k is an arbitrary positive constant.

Proof:
Note: Axioms (1) - (3) imply $\delta(\mathbf{p},\mathbf{1})$ is, up to a positive constant, the entropy of the distribution **p**, $\delta(\mathbf{p},\mathbf{1}) = -\Sigma p_i\ln(p_i)$, by Shannon's results for message sources [10]. For completeness, we give an expanded version of the argument here, in the context of our notation. Axiom (4) is then used to extend the result to information gain.

Use mathematical induction to establish
$$\delta(\mathbf{s^m},\mathbf{1}) = m\delta(\mathbf{s},\mathbf{1});\ s > 1 \text{ an integer}, \tag{1}$$
as follows. The proposition is obviously true for m = 1. Let us illustrate the induction step by examining the case for m = 2:

---

[4] We extend the notation to allow random variables having given distributions to represent those distributions as arguments in δ, and we let "E_I" denote expected value with respect to the distribution of I. By the notation conventions at the beginning of this section, **1** is a mass function assingning mass unity to a single point of S.

Let Z represent an experiment consisting of choosing a point at random from a set S containing s^2 points. Suppose S is partitioned into s subsets each containing s points. We may view an outcome on Z as being determined through a compound experiment I;X, where I indicates a randomly selected partition subset from which a point X will be chosen at random (see our discussion of Axiom (3) below). It follows that I has uniform distribution **s** and, conditionally upon the outcome on I, X also has uniform distribution **s**. By Axiom (3), it follows that

$$\delta(\mathbf{s^2},\mathbf{1}) = \delta(\mathbf{s},\mathbf{1}) + E_I\, \delta(\mathbf{s},\mathbf{1}) = 2\delta(\mathbf{s},\mathbf{1}).$$

Assume $\delta(\mathbf{s^{m-1}},\mathbf{1}) = (m-1)\delta(\mathbf{s},\mathbf{1})$. By Axiom (3), with an indicator variable I showing which of s sets each having s^{m-1} points from which to choose a point at random, we have $\delta(\mathbf{s^m},\mathbf{1}) = \delta(\mathbf{s},\mathbf{1}) + E_I\, \delta(\mathbf{s^{m-1}},\mathbf{1}) = \delta(\mathbf{s},\mathbf{1}) + (m-1)\delta(\mathbf{s},\mathbf{1}) = m\delta(\mathbf{s},\mathbf{1})$, and equation (1) follows for integer $s > 1$.

Let t be an arbitrary (fixed) integer greater than one, as is s. Then by the above, $\delta(\mathbf{t^n},\mathbf{1}) = n\delta(\mathbf{t},\mathbf{1})$. For any positive integer n there exists a positive integer m such that $s^m \le t^n < s^{m+1}$, so $m \ln(s) \le n \ln(t) < (m+1)\ln(s)$, thus

$$\frac{m}{n} \le \frac{\ln(t)}{\ln(s)} < \frac{m+1}{n} = \frac{m}{n} + \frac{1}{n},$$

so (2)

$$\left|\frac{\ln(t)}{\ln(s)} - \frac{m}{n}\right| < \frac{1}{n} \quad .$$

By Axiom (2), $\delta(\mathbf{s^m},\mathbf{1}) \le \delta(\mathbf{t^n},\mathbf{1}) < \delta(\mathbf{s^{m+1}},\mathbf{1})$, that is, by equation (1),

$$m\delta(\mathbf{s},\mathbf{1}) \le n\delta(\mathbf{t},\mathbf{1}) < (m+1)\delta(\mathbf{s^{m+1}},\mathbf{1})$$

or

$$\frac{m}{n} \le \frac{\delta(\mathbf{t},\mathbf{1})}{\delta(\mathbf{s},\mathbf{1})} < \frac{m}{n} + \frac{1}{n},$$

so

$$\left|\frac{m}{n} - \frac{\delta(\mathbf{t},\mathbf{1})}{\delta(\mathbf{s},\mathbf{1})}\right| < \frac{1}{n} \quad .$$

Combining this with equation (2), by the triangle inequality it follows that

$$\left|\frac{\ln(t)}{\ln(s)} - \frac{\delta(\mathbf{t},\mathbf{1})}{\delta(\mathbf{s},\mathbf{1})}\right| < \frac{2}{n},$$

and, since s and t are arbitrary integers greater than 1, and n may be chosen arbitrarily large, it follows that

$$\delta(\mathbf{s},\mathbf{1}) = k \ln(s);\ s > 1. \qquad (3)$$

By Axiom (2), it follows that $k > 0$.

If $s = 1$, then by Axiom (4), $\delta(\mathbf{1},\mathbf{1}) = \delta(\mathbf{1},\mathbf{1}) + \delta(\mathbf{1},\mathbf{1})$, so $\delta(\mathbf{1},\mathbf{1}) = 0$. Thus equation (3) holds for the case $\mathbf{s} = \mathbf{1}$ as well.

We now extend the result of equation (3) from uniform distributions to distributions having rational masses. If the t components of **p** are positive rational numbers, they may be expressed with common denominator, in the form $p_i = n_i/\Sigma_j n_j$; $i = 1, 2, \ldots, t$, where the n's are positive integers. Let an experiment Z have a uniform distribution $\Sigma \mathbf{n_j}$. Partition S with Σn_j points into t subsets with $n_1, n_2, \ldots, n_t$ points, respectively. We can realize an outcome on Z by performing the compound experiment I;X, where I indicates selection of a subset from the partition, and X indicates (conditional) random selection of a point from the subset indicated by I. The random variable I has distribution **p** and, given I = i, X has uniform distribution $\mathbf{n}_i$. By Axiom (3) it follows that

$$\delta(\Sigma \mathbf{n}_j, \mathbf{1}) = \delta(\mathbf{p},\mathbf{1}) + E_I\, \delta(\mathbf{n}_I, \mathbf{1}).$$

By equation (3), it thus follows that $k \cdot \ln(\Sigma n_j) = \delta(\mathbf{p},\mathbf{1}) + \Sigma_j k \cdot \ln(n_j) \cdot p_j$, so

$$\begin{aligned}\delta(\mathbf{p},\mathbf{1}) &= k[\ln(\sum_j n_j) - \sum_i p_i \ln(n_i)] \\ &= k[\sum_i p_i \cdot \ln(\sum_j n_j) - \sum_i p_i \ln(n_i)] \\ &= -k \sum_i p_i \ln\left(\frac{n_i}{\sum_j n_j}\right) \\ &= -k \sum_i p_i \ln(p_i)\,.\end{aligned} \tag{4}$$

We invoke Axiom (1) to extend equation (4) to mass functions **p** with arbitrary positive mass values (i.e., where components of **p** may be irrational). For any mass function **p** there exists a sequence $\{\mathbf{p}^m\}$ of mass functions with rational masses converging to **p**, so continuity of δ implies

$$\lim_{m\to\infty} \delta(\mathbf{p}^m, \mathbf{1}) = \delta(\lim_{m\to\infty} \mathbf{p}^m, \mathbf{1}) = \delta(\mathbf{p}, \mathbf{1}).$$

But, by equation (4) and continuity of the function x ln(x),

$$\lim_{m\to\infty} \delta(\mathbf{p}^m, \mathbf{1}) = \lim_{m\to\infty}(-k \sum_i p_i^m \ln(p_i^m)) = -k \sum_i p_i \ln(p_i)\,.$$

Finally, we extend the foregoing to the information gain function $\delta(\mathbf{p},\mathbf{p}^*)$. By axiom (4), taking $\mathbf{q} = \mathbf{p} = \mathbf{p}^*$, we have

$$\delta(\mathbf{p}^*,\mathbf{p}^*) = \delta(\mathbf{p}^*,\mathbf{p}^*) + \delta(\mathbf{p}^*,\mathbf{p}^*),$$

so it follows that $\delta(\mathbf{p}^*,\mathbf{p}^*) = 0$. This, with another application of axiom (4), implies that $\delta(\mathbf{1},\mathbf{p}^*) = -\delta(\mathbf{p}^*,\mathbf{1})$. With yet another application of axiom (4),

$$\delta(\mathbf{p},\mathbf{p}^*) = \delta(\mathbf{p},\mathbf{1}) + \delta(\mathbf{1},\mathbf{p}^*) = \delta(\mathbf{p},\mathbf{1}) - \delta(\mathbf{p}^*, \mathbf{1}) = -\Sigma p_i \ln(p_i) + \Sigma p^*_i \ln(p^*_i),$$

up to a positive multiplicative constant. Q.E.D.

Remarks:

- We shall assume the multiplicative constant is unity. One might regard the constant to be related to the choice of base for the logarithm function, which we assume to be the natural logarithm.
- $\delta(\mathbf{p},\mathbf{1})$ is the information gain realized in totally resolving uncertainty inherent in an experiment with distribution **p**, since **1** is degenerate at a point of S.

Corollaries:

(1) If **p** and **p*** have the same set of masses, $\delta(\mathbf{p},\mathbf{p}^*) = 0$.
Note **p** and **p*** could be different mass functions, however.
In particular, for any $\mathbf{p} \in \Omega$, $\delta(\mathbf{p}, \mathbf{p}) = 0$.

(2) For fixed $\mathbf{p^*}$, among distributions $\mathbf{p}$ with exactly n positive masses, $\delta(\mathbf{p},\mathbf{p^*})$ is maximal for $\mathbf{p} = \mathbf{n}$.

The distribution over a given set with the most randomness is the uniform distribution over that set.

(3) For fixed $\mathbf{p}$, $\delta(\mathbf{p},\mathbf{p^*})$ is maximal for $\mathbf{p^*} = \mathbf{1}$.

A distribution over S with the minimal randomness is one degenerate at a point of S.

(4) For two compound experiments I;X and J;Y,

$$\delta(I;X, J;Y) = \delta(I, J) + E_I\delta(X, \mathbf{1} \mid I) - E_J\delta(Y, \mathbf{1} \mid J).$$

In particular, for compound experiments I;X and I;Y,

$$\delta(I;X, I;Y) = E_I\delta(X, Y \mid I).$$

This gives a result for information gain with compound experiments. The second expression is interesting in that the information gain depends solely on the expected gain in resolving X to Y, and not upon the entropy associated with the "mixing" distribution, I.

(5) For any positive integers k, s, and m, $\delta(\mathbf{km}, \mathbf{sm}) = \delta(\mathbf{k}, \mathbf{s}) = \ln(k/s)$.

For uniform distributions, the information gained in finding sub-regions to be certainly target free is dependent only on the ratio of prior-to-posterior area, and not on the individual sizes of these regions. That is, there is one nit of information gain in narrowing from a region of size e to one of size 1, just as there is one nit of information in narrowing from a region of size 1 to one of size 1/e. In both cases, one nit of information would be required to inform which region (the smaller region or its complement) the target is in.

(6) $\delta(\mathbf{q}, \mathbf{p}) = -\delta(\mathbf{p}, \mathbf{q})$.

With time reversal, one obtains information loss (negative of information gain).

Discussion of the axioms:

The first axiom is a technical condition used to extend the theorem from distributions with rational masses to distributions with real masses. It requires that the function δ has no jumps (i.e., $\delta(\mathbf{q}, \mathbf{q^*})$ tends to the value $\delta(\mathbf{p},\mathbf{p^*})$ as $\mathbf{q}$ tends to $\mathbf{p}$ and $\mathbf{q^*}$ tends to $\mathbf{p^*}$). This "smoothness" axiom seems entirely reasonable.

The second axiom asserts that, for two cases with uniform prior distributions, where uncertainty is resolved to a given posterior distribution $\mathbf{p^*}$ in both cases, the gain in information is greatest for the case with uniform prior over the most mass points. That is, a uniform prior distribution with more mass points has less "starting" information than has one with fewer mass points (hence giving more information gain upon resolving uncertainty to the posterior $\mathbf{p^*}$).

For example, suppose we are calculating the information gains related to narrowing down the location of a given enemy unit to be within a fixed subset U of cells within the original area of concern. Considering two possible prior distributions, say one uniform over a large superset of U and another uniform over a smaller superset of U, the axiom asserts that narrowing the possible location to one within U in the first case gives more information gain than that in the second case.

The third axiom concerns the behavior of δ when the prior distribution is regarded as a compound experiment. Consider **p** having masses {1/4, 1/4, 1/3, 1/6} where these masses represent the respective probability of four distinct outcomes of an experiment. Imagine the random variable Z represents the outcome of the experiment. We may, without changing the stochastic features of the experiment, break the experiment of observing an outcome on Z into an initial outcome on I which denotes whether the outcome is in the set associated with the first two masses or the second set of two masses, followed by a conditional experiment X giving the outcome within whichever of the two sets occur. Clearly, I has probability 1/2 of indicating the first set, and similarly for the second set. Given I indicates the first set, (conditionally) either of the two outcomes in the first set are equally likely, so the conditional distribution of X has probability 1/2 at each of these points. If I indicates the second set occurred, X has conditional probability 2/3 of giving the first outcome in the set, and probability 1/3 of giving the second value. Thus the original experiment Z and the compound experiment I;X have the same overall probabilities of giving each of the four original outcomes, where we use standard conditional probability calculations in the second case. For example, the probability Z gives the third value in the original set, 1/3, is equal to the probability I indicates the second set occurred (probability 1/2) and conditionally X gives the first outcome in this set (conditional probability 2/3), so the overall probability of this outcome with I;X is (1/2)(2/3) = 1/3, as before with the experiment Z.

The point is, we may, if we wish, view the experiment Z as a compound experiment I;X, where first I is observed, then conditionally a value of X is observed. Since $\delta(Z,\mathbf{1})$ may be viewed as the gain in totally resolving the uncertainty in the experiment Z, axiom (3) asserts the information gain in totally resolving the uncertainty in I;X (or Z) can be determined by adding the gain in totally resolving the uncertainty in I, and the average (over possible values of I) of the conditional gains in totally resolving the uncertainty in X, given each value of I.

For the numerical example above, note

$$\delta(Z, \mathbf{1}) = 2(1/4)\ln(1/4) + (1/3)\ln(1/3) + (1/6)\ln(1/6),$$

which may be seen to be equal to

$$\delta(I,\mathbf{1}) + E_I\delta(X,\mathbf{1} \mid I) = [2(1/2)\ln(1/2)] + (1/2)[2(1/2)\ln(1/2) + (2/3)\ln(2/3) + (1/3)\ln(1/3)].$$

The fourth axiom asserts one may view resolving the uncertainty with a prior distribution **p** to the uncertainty with a posterior distribution **p*** in terms of two information gain steps involving an intermediate distribution **q**. From the point of view of information gain from **p** to **p***, it does not matter what intermediate distribution might have been attained through some portion of the information that resulted in resolving the uncertainty in **p** to that in **p***. In our applications, the information gain is usually computed over each fixed time increment, Δt. Axiom (4) asserts we could obtain the information gain over Δt by adding the gains computed for the two increments of length $\Delta t/2$ making up Δt, for example. Thus, for example, the information gain over the period of a battle, t, may be computed by accumulating the incremental gains over time segments Δt_i making up the battle period.

Operational illustrations of the third and fourth axioms are contained in sections I-1, (a') and I-1, (b').

III-2. Assumption of Independence Between Targets

We note in the combinatorial setting that, in the case where targets *are* distinguishable and may occupy common areas, the sample space changes from that considered in Section I-4. If there is no constraint on how many Red targets can occupy a cell, Red could deploy his forces in any of R^N ways, where R is the number of cells in the area of concern. With a uniform prior, the initial entropy is then $ln(R^N) = N\,ln(R)$. But if there was only a single target to find, the initial entropy would be $ln(R)$, as a special case of the first equation in the preceding example. Thus it appears that for N targets, with this sample space, entropy for all N targets is the <u>sum</u> of the entropies for the individual targets. This suggests a more general relationship exists between individual targets within the R possible areas and the combined array of N targets over the appropriate sample space.

A sufficient condition for this additivity property is the independence of the positions of the N individual targets. To see this, suppose the joint density function $p(\cdot)$ of the N target positions factors into the product of marginals:

$$p(t_1, t_2, \ldots, t_N) = p_1(t_1)p_2(t_2)\ldots p_N(t_N).$$

For simplicity of notation, denote vectors of t_i-values by $\underline{t}$ which can range over a sample space S. Then

$$\begin{aligned}-\Sigma_{t \in S}\, p(\underline{t})ln(p(\underline{t})) &= -\Sigma p(\underline{t})[ln(p_1(t_1)) + \ldots + ln(p_N(t_N))] \\ &= -\Sigma p_1(t_1)ln(p_1(t_1)) - \ldots - \Sigma p_N(t_N)ln(p_N(t_N)),\end{aligned}$$

so $e_{1,2,\ldots,N} = e_1 + e_2 + \ldots + e_N$, where "$e_i$" denotes entropy with respect to target i. That is, the joint entropy of the N targets is the sum of marginal entropies of the respective individual targets. In applications, it might be possible to gain independence (at least roughly) by carefully defining what constitutes "areas" and "targets". For example, Blue might know Red deploys tanks in platoons of four tanks, so finding a single tank actually provides information about three additional tanks in the same vicinity. In this case, one might want to model the "targets" as tank platoons, rather than individual tanks. Locating a tank would then indicate presence of a platoon within some appropriate area, rather than precisely locating the platoon.

We have investigated several numerical examples using bivariate distributions with varying levels of correlation, and it appears the correlation in target locations must be fairly large before entropy calculated by summing marginal entropies differs appreciably from the exact joint entropy. The following example illustrates this for a case with correlation about 0.2.

Example: A Discrete Bivariate Distribution.

Suppose there are $R=2$ areas, labeled "0" and "1", and suppose there are two targets, $T1$ and $T2$. Imagine Red deploys the targets such that, from Blue's point of view, the joint distribution of the location of $(T1, T2)$ is in accordance with the following bivariate mass function:

$p(t_1, t_2)$	0	1	$p_1(t_1)$
0	.2	.1	.3
1	.3	.4	.7
$p_2(t_2)$	.5	.5	1.0

The joint entropy is the sum over the four cells of the joint table of terms of the form $-p \ln(p)$, which gives 1.28. The marginal entropies for *T1* and *T2* are 0.611 and 0.693, respectively, which sum to 1.30. Note there is significant correlation between *T1* and *T2* ($\rho=0.22$), as evidenced by the relatively larger mass values on the main diagonal of the table, yet the joint entropy is not greatly different from the sum of the marginal entropies.

Now suppose intelligence is obtained indicating T2 may not be in position 0, such that the prior marginal probabilities are updated from (.5, .5) to (.4, .6). Then the posterior joint distribution becomes

$p(t_1, t_2)$	0	1	$p_1(t_1)$
0	.16	.12	.28
1	.24	.48	.72
$p_2(t_2)$	.4	.6	1.0

where we note the marginal distribution of T1 has also changed, due to the correlation in T1 and T2 locations. The correlation between T1 and T2 remains 0.22. Now the joint entropy is 1.24 and the sum of marginal entropies is 0.593 + 0.673 = 1.27, which differs from the joint value by 0.03. However, the information gain, using the joint entropies in both cases is 1.28 - 1.24 - .04, whereas the value obtained using the sum of marginal entropies in both cases is 1.30 - 1.27 = 0.03. We note the error in information gain associated with assuming independence, 0.01, is smaller than the errors in either of the entropy calculations (0.02 and 0.03, respectively).

It should be noted the argument given at the beginning of this section is valid for a finite discrete joint distribution. In Section III-4 we give an example of a bivariate continuous distribution over a disc in the plane. One can define a uniform distribution over such a domain, assumed to be centered at the origin, in terms of independent random variables representing polar coordinates of the outcome. However the joint entropy, which is the sum of the marginal entropies of the random variables, is not the entropy of the original distribution uniform over the disc. This example shows caution must be exercised in exploiting independence in entropy computations with continuous random variables.

Conditioning approach

If the individual target positions are not independent, one can compute entropy using the joint distribution of the target locations, or one can sum entropies of conditional distributions instead of marginal distributions. It is always the case that the joint mass function factors into a product of conditional mass functions, so one can express the joint entropy as a sum of terms related to corresponding conditional entropies. The joint entropy is then given as the sum of expected values of these conditional entropies. For example, if X, Y, and Z are jointly distributed random variables, the joint entropy can be given by $e_{X,Y,Z} = e_X + E_X e_{Y|X} + E_{X,Y} e_{Z|X,Y}$, where E_X denotes expectation with respect to the marginal distribution of X, $e_{Y|X}$ denotes entropy of the conditional distribution of Y given X, and similarly for the other expression. (An outline of a proof of this is given in Section III-3.) Thus, in principle, we need only be concerned with entropy computations for univariate marginals and conditionals. Shannon's basic inequality asserts that a conditional entropy cannot exceed the corresponding unconditional entropy [7], so approximating the joint entropy by the sum of marginal entropies is conservative. That is, it will over-state the true entropy and hence we will tend to over-state the degree of apparent randomness remaining in Red's deployment. This was seen in the example above, where the sum of marginal entropies exceeded the joint entropy in both cases. Since we are concentrating on *decreases* in entropy, the amount of error involved in summing the marginal entropies for both the prior and posterior distributions may be negligible for practical purposes. Again, in the example above, we see assuming independence gave error in information gain that was much smaller than the respective errors in the individual entropies.

III-3. Combining Entropy Measures in a Compound Experiment

An implication of the third axiom of information gain given in Section III-1 is that in order to combine entropies in a compound experiment, one cannot simply sum the marginal entropies. To further illustrate this, imagine drawing a target type at random from a total set {1, 2, .., m}of target types, then determining the location entropy of a target of the selected type. Note targets of the various types may have different location distributions (obstacles and tanks aren't distributed over a piece of terrain in the same way, for example). Consider an indicator random variable I with possible values 1, 2, ..., m, representing the outcome on drawing the type of target. Let $e_{T,I}$ denote the entropy of the compound outcome on I and location of the selected target, T, that is, with respect to the joint distribution of I and T. Similarly let e_I, and $e_{T|I}$ denote the entropy of the distribution of I and the conditional entropy of target location, given the outcome on I, respectively. Then $e_{T,I} = e_I + E_I(e_{T|I})$, where "$E_I$" denotes expected value with respect to the distribution of I. This can be motivated by a conditioning argument with the definition of entropy, as follows:

$$
\begin{aligned}
e_{T,I} &= -\sum_{t,i} p_{T,I}(t,i)\ln(p_{T,I}(t,i)) = -\sum_{t,i} p_{T,I}(t,i)\ln(p_{T|I}(t|i)p_I(i)) \\
&= -\sum_i \sum_t p_{T,I}(t,i)\ln(p_I(i)) - \sum_i \sum_t p_I(i)p_{T|I}(t|i)\ln(p_{T|I}(t|i)) \\
&= -\sum_i p_I(i)\ln(p_I(i)) - \sum_i \left(\sum_t p_{T|I}(t|i)\ln(p_{T|I}(t|i))\right)p_I(i) \\
&= e_I + E_I(e_{T|I})
\end{aligned}
$$

This argument can be repeated for higher-dimensional cases.

III-4. Extension from Discrete to Continuous Distributions

Univariate distributions

The extension of the entropy concept to continuous distributions is not entirely straight-forward. Many authors have defined the entropy of a continuous distribution with density f to be

$$- \int_{\{x|f(x)>0\}} f(x) \cdot \ln(f(x))dx,$$

which is closely analogous to the earlier definition for discrete distributions [4, 9]. For discrete distributions, e is a measure of the dispersion of probability mass over points, *without regard to what those points are*. Thus, if we form a sequence of increasingly fine discrete approximations of a continuous distribution, the sequence of corresponding entropies will increase without bound!

To illustrate, consider a continuous uniform distribution over the interval [a,b], so $f(x) = 1/(b - a)$, for x between a and b. Then the above integral gives

$$-\int_a^b f(x)\ln(f(x))dx = \frac{\ln(b-a)}{b-a}\int_a^b dx = \ln(b-a)$$

(which, we note, is negative when $0 < b - a < 1$). Now suppose we form a sequence of discrete approximations of this distribution, based on partitioning the interval [a,b] into n sub-intervals Δx of length $(b - a)/n$. Let us consider the approximating mass function that takes values $p_i = f(x_i)\cdot\Delta x_i = 1/n$, where x_i is the center of the i^{th} sub-interval and Δx_i is its width. The entropy of this discrete approximation is the maximal value attained with a discrete uniform distribution over n points,

$$\ln(n) = e = -\sum_{i=1}^{n} f(x_i)\Delta x_i \ln(f(x_i)\Delta x_i) = -\sum_{i=1}^{n} f(x_i)\ln(f(x_i))\Delta x_i - \sum_{i=1}^{n} f(x_i)\Delta x_i \ln(\Delta x_i)$$

Now consider refining the partition and taking the limit of the terms on the right as $n \to \infty$. The first term converges to $-\int_a^b f(x)\ln(f(x))dx$, the "continuous analogy" expression for entropy mentioned above. The limit of the second term,

$$\lim_{n\to\infty} -\sum f(x_i)\Delta x_i \ln(\Delta x_i) = -\lim_{n\to\infty} n\cdot\frac{1}{n}\ln(\Delta x) = \lim_{n\to\infty}[-\ln(\Delta x)]$$

diverges to $+\infty$. Therefore, $-\int f(x)\ln(f(x))dx$ is only part of the limit as we form finer and finer discrete approximations to f. Indeed, in the present case, the sequence of "approximating entropies" does not converge.

Thus, from the point of view of extending the definition of entropy for discrete distributions, the expression $-\int f(x)\ln(f(x))dx$ *may not be appropriate for measuring information.* However, since the term $-\ln(\Delta x)$ adds out in the computation of information gain, *integrals of this form may be employed in computing* δ *for continuous distributions, with an interpretation identical to that for discrete distributions.*

Some interesting conclusions result from such models. For example, suppose the prior distribution of location of a certain target is continuous uniform over an interval

(1, 5) and information is received that shows the target cannot be in the interval (3. 5), so it must be uniformly distributed over the interval (1, 3). Then the information gain is ln(4) - ln(2) = ln(2). If subsequent search clears the interval (2, 3), the target distribution is updated to uniform over the interval (1, 2). The information gain this time is ln(2) - ln(1) = ln(2), same as in the first step. However, note in the first step a region of length 2 was cleared, while in the second step a region of length only one was cleared. (This is related to the result in the mobile target example, in Section III-5, that the information gain curve is independent of the target movement rate, r, and is a direct analog of the result for discrete uniform distributions discussed in Corollary (5), Section III-1.)

Exponential example

Let us illustrate the foregoing with a second example. Consider an exponential distribution with parameter λ and note

$$-\int_0^\infty f(x)\ln(f(x))dx = -\int_0^\infty \lambda e^{-\lambda x}[\ln(\lambda) - \lambda x]dx = -\ln(\lambda) + 1 = 1 + \frac{1}{2}\ln(\sigma^2),$$

which is negative for sufficiently large λ (i.e., small variance, σ^2). Suppose a change in information status can be represented by a change in the parameter to λ^*. Then the information gain is

$\delta(\lambda,\lambda^*) = ln(\lambda^*) - ln(\lambda) = ln(\lambda^*/\lambda) = (1/2)ln(\sigma^2/\sigma^{*2}) > 0$ if and only if $\sigma^{*2} < \sigma^2$, so information gain is positive in this case exactly when the posterior distribution has smaller variance than had the prior. The fact that the integral above can be negative does not affect the interpretation of the information gain. The comment on positive gain remains valid even when the variances involved are both smaller than $1/e^2$, so that the individual integrals are negative.

Multivariate considerations

Caution must be exercised in computation of entropy for two- and higher-dimensional continuous distributions. If X and Y are independent jointly uniform over a rectangle in the plane defined by opposite corners (0,0) and (a,b), the double integral defining the entropy of (X,Y) gives

$e_{X,Y} = \ln[(a - 0)(b - 0)] = \ln(\text{area of rectangle}) = e_X + e_Y$,

as expected. For a uniform distribution over non-rectangular region, R, of two dimensions, one may argue as follows. Approximate the distribution by forming a "partition" of the region into $\Delta\times\Delta$ squares, such that a limiting process (as Δ gets small) will provide exact coverage of the region. There are approximately [area(R) / Δ^2] of these squares. Now let I be an indicator variable in a compound experiment, where I chooses which of the squares to sample, and, given the value of I, (X,Y) is an outcome distributed uniform over the chosen $\Delta\times\Delta$ square. Then by Axiom (3) of Section III-1,

$e_R = e_I + Ee_{X,Y|I} = \ln(\text{area}(R) / \Delta^2) + E_I(\ln \Delta^2) = \ln(\text{area}(R))$,

again as expected.

Now consider a uniform distribution over a circle C with center at the origin and radius r_0. The entropy of this distribution, by the above argument, is $\ln(\pi r_0^2)$. Let R and Θ be jointly distributed random variables such that (R,Θ) has a joint density function which is uniform over C. Then R has the "triangular" density $f(r) = 2r/r_0^2 = kr$ (where

$k = 2/r_0^2$), for $0 < r < r_0$, and Θ has density $g(\theta) = 1/2\pi$; $0 < \theta < 2\pi$. It follows the "marginal" entropies of R and Θ are, respectively,

$$e_R = r_0^8/16[\ln(r_0^3/2) - 1/2] \quad \text{and} \quad e_\Theta = \ln(2\pi).$$

It is immediately obvious that the joint entropy over the region C, e_C, is not equal to $e_R + e_\Theta$, in spite of the fact that the random variables R and Θ are independent! Now, by Axiom (3) of Section III-1, it is true that $e_C = e_R + E_R e_{\Theta|R}$, regardless of independence. To compute $E_R e_{\Theta|R}$, first consider the region within C associated with a given value, r, of R. As Θ varies over its range, a circle with circumference $2\pi r$ is swept out. Thus, for R = r, the conditional entropy associated with Θ is $e_{\Theta|R} = \ln(2\pi r)$, and the expected value is

$$E_R e_{\Theta|R} = \int_0^{r_0} \ln(2\pi r) f(r) dr = \ln(2\pi) + \int_0^{r_0} kr \cdot \ln(kr) dr - \ln(k)$$

$$= \ln(\frac{2\pi}{2/r_0^2}) - e_R = \ln(area(C)) - e_R.$$

Then $e_C = e_R + E_R e_{\Theta|R} = \text{area}(C)$, as should be the case.

A caution about continuous random variables

The forgoing example illustrates the fact that even though R and Θ are independent, it may not follow that $e_{\Theta|R} = e_\Theta$ nor does it follow that $e_C = e_R + e_\Theta$, so in particular, $e_C \neq e_{R,\Theta}$. Dilemmas such as this illustrate the fact that, in using continuous random variables, we inherently establish a coordinate system. For example, in the exponential (with parameter λ) example above, we saw the entropy is $1 - \ln(\lambda)$. But λ is a *scaling* parameter; choosing a value of this parameter is equivalent to choosing a scale for the coordinate system used to represent outcomes on the exponential experiment. This is not consistent with the observation that, with discrete distributions, entropy depends only on probabilities. Note, however, in the exponential example, the information gain, $\ln(\lambda/\lambda^*)$, is unaffected by changes in scale. In Section III-6, we show this is true in general.

The point is, with continuous models, caution must be exercised lest *values* of outcomes of experiments enter into evaluations of entropy (which should depend only upon *probabilities* of outcomes). The same is true for calculations of information gain, although in some cases effects of an implicit selection of coordinate system add out in this case. In general, for computations of information gain in applications, it is good insurance to form discrete approximations of any continuous distributions involved. This should help eliminate the potential for gross errors such as might well have occurred in the preceding example involving a continuous distribution over a disc in the plane. The probability contour maps in the application discussed in Section II-4 defined continuous bivariate distributions that were subsequently converted to discrete mass functions precisely for this reason.

III-5. An Example Application to Mobile Targets

It seems reasonable to take into account the age of location information. If a mobile target is located at some point in time, one does not know its location at a later time, assuming no information about its location has been received in the meantime.

From a Bayesian updating point of view, the "spike" of probability at the target's location when it is located begins to "melt" with the passage of time. Based on the target's ability to move in the neighborhood of its location, the probability the target is in neighboring cells begins to grow as time passes. If we imagine a plot of uncertainty about the target's location versus time, we would expect the curve to increase in some fashion. Alternatively, cumulative information gain should become more negative as time increases. But what might be the shapes of these curves?

A simple model

We illustrate how a relationship between information gain and time might be established, using as an example a simple model of how the target might move. Suppose a certain target located at the point (0, 0) on a certain terrain could move at average rate r in any direction, and could change directions at random times. We imagine the location (X, Y) of the target after t minutes is distributed bivariate normal, with mean (0, 0) and covariance matrix $\sigma^2 I$ (as would be expected by a diffusion approximation argument). We link σ^2 to the average rate r and time t as follows: the target must be within a circle of radius rt; suppose rt/2 is the median distance of the target from (0, 0). Since $(X^2 + Y^2)/\sigma^2$ is chi-square distributed with two degrees of freedom (that is, exponential with parameter 1/2), it follows the median of $(X^2 + Y^2)/\sigma^2$ is 2ln(2), so the median distance of the target from (0, 0) is $\sigma(2\ln(2))^{1/2}$. If we set this equal to rt/2, it follows that $\sigma^2 = (rt)^2/8\ln(2)$.

For a (univariate) normal distribution with parameters μ and σ^2 the entropy is

$$-\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \left(-\ln(\sqrt{2\pi}\sigma) - \frac{1}{2\sigma^2}(x-\mu)^2\right) dx = \frac{1}{2}(1 + \ln(2\pi)) + \ln(\sigma).$$

(We note, in a normal distribution, entropy increases linearly with the logarithm of variance and does not depend on the mean.) For a bivariate normal with independent components (inherent in the assumed form of the covariance matrix), by the result in Section III-2, the joint entropy of X and Y is thus two times the marginal value shown above, or $e_{X,Y} = 1 + \ln(2\pi) + \ln(\sigma^2)$. Setting $\sigma^2(t) = (rt)^2/8\ln(2)$ gives entropy of target location at time t, $e(t) = 1 + \ln(2\pi) + 2\ln(rt) - \ln(8\ln(2))$. The cumulative information gain, say from initial time $t_0 > 0$ to time $t > t_0$, $\delta(t_0, t)$, is then given by

$$\delta(t_0, t) = 2\ln(rt_0) - 2\ln(rt) = 2\ln(t_0/t), \text{ for } t > t_0.$$

We see the shape of the information gain curve is therefore like $-\ln(t^2)$, *independent of r*! The information gain rate at time t is given by

$$\lim_{\Delta t\to 0} \delta(t, t + \Delta t)/\Delta t = 2\lim_{\Delta t\to 0} [\ln(rt) - \ln(r[t + \Delta t])]/\Delta t = -2/t, \text{ for } t > t_0.$$

Of course, if the initial time t_0 is chosen so the "area of uncertainty" is of fixed radius, then the information gain from time t_0 to time t *does* depend on target average movement rate. For example, if we choose $t_0 = 1/r$, then $\delta(1/r, t) = -2\ln(rt) = \ln(1/r^2t^2)$, for $t > 1/r$.

With other models of how the distribution of target location expands with time we obtain similar results. For example, if the location of the target t hours after it is located is assumed to be uniform over a circle of radius rt, then the entropy is proportional to the area of the circle, so information gain from t_0 to t is the logarithm of the ratio of the area containing the target at time t_0 to that containing the target at time t. That is, in this case also, $\delta(t_0, t) = 2\ln(t_0/t)$, again independent of r. This result may seem counter-intuitive, at first, but it is consistent with a similar result stated as Corollary 5 in Section III-1. It can be motivated and illustrated by considering the number of nits of information required to narrow the location area to some fraction of that area, thinking in the reverse direction.

III-6. Functions of Random Variables

The information gain from **p** to **p*** can be viewed in terms of prior and posterior random variables X and X* having these respective distributions, in some cases. In certain applications, it could be of interest to consider some function, g, of these random variables. In principle, one could find the distributions of g(X) and g(X*) and proceed as usual, but in many cases this step is not necessary; one may obtain information gain directly.

Consider an example related to the mobile target model described in Section III-5, where we know the entropies of X and X* but want the information gain going from X^2 to X^{*2}. If the distribution of X is continuous, with density function f having support $(0,\infty)$ or $(-\infty,\infty)$, the density of X^2 is given by $f_2(r) = f(\sqrt{r})/2\sqrt{r}$; $r>0$. Then the entropy of X^2, e_2, is given by

$$e_2 = -\int \frac{1}{2\sqrt{r}} f(\sqrt{r}) \ln[\frac{1}{2\sqrt{r}} f(\sqrt{r})]dr$$
$$= \int \ln(2\sqrt{r}) f(\sqrt{r}) d\sqrt{r} - \int f(\sqrt{r}) \ln(f(\sqrt{r})) d\sqrt{r}$$
$$= E_X \ln(2X) + e_X$$

More generally, suppose $g' > 0$ on the support of f (still assumed to be $(0,\infty)$ or $(-\infty,\infty)$). Then

$$f_{g(X)}(r) = f(g^{-1}(r)) / g'(g^{-1}(r))$$

so

$$e_{g(X)} = -\int f(g^{-1}(r)) \ln\left(\frac{f(g^{-1}(r))}{g'(g^{-1}(r))}\right) dg^{-1}(r) \qquad (5)$$
$$= E_X \ln g'(X) + e_X$$

Examples

1. g(X) = F_X(X). As a verification of expression (5), consider the case where g is the CDF, F, of X. We know, by the probability integral transformation, that F(X) is distributed uniform over the interval (0,1), so by the expression obtained in Section III-4, it follows that $e_{F(X)} = \ln(1) = 0$. By the preceding expression, this is equal to

$$E_X(\ln(F'(x))) + e_X = E_x \ln(f(x)) + e_X,$$

so it follows that

$$e_X = -E_x \ln(f(x)) = -\int (\ln(f(x))\, f(x)\, dx;$$

consistent with the definition of entropy for continuous distributions.

2. $g(X) = e^X$ or $g(X) = \ln(X)$. If X has continuous distribution with mean μ and $Y = e^X = g(X)$, then $\ln(g'(X)) = X$, so the entropy of e^X is $\mu + e_X$. Similarly, if X has support contained in $(0,\infty)$, $e_{\ln(X)} = e_X - E_X\ln(X)$.

3. $\delta(X,Y)$ is invariant under location and scale changes. If g is a linear transformation, say $g(z) = az + b$, then $e_{aX+b} = \ln(a) + e_X$ and $e_{aY+b} = \ln(a) + e_Y$, so $\delta(aX+b,aY+b) = \delta(X,Y)$. Thus, while the re-scaling part, a, of the transformation (but not the re-locating part, b) effects entropy, it does not effect information gain.

The preceding can be extended to 1-1 transformations of jointly distributed random variables, using standard methods with the Jacobian of the transformation. For example, for a bivariate case, suppose X_1 and X_2 are jointly continuous with density f over the real plane, and suppose $Y_1 = g_1(X_1, X_2)$ and $Y_2 = g_2(X_1, X_2)$ is a 1-1 transformation whose inverse is given by $x_1 = h_1(y_1, y_2)$; $x_2 = h_2(y_1, y_2)$, valid over some two-dimensional domain D contained in the real plane. A standard result of probability theory states the joint distribution of Y_1 and Y_2 is given directly by

$$f_{Y_1,Y_2}(y_1,y_2) = f(h_1(y_1,y_2),h_2(y_1,y_2))|J(y_1,y_2)|;\ (y_1 y_2) \in D$$

where $|J|$ is the absolute value of the 2×2 determinant $J = |\,\partial h_i/\partial y_j\,|$, which is the Jacobian of the transformation. The joint entropy of Y_1 and Y_2 can thus be expressed as follows:

$$\begin{aligned}
e_{Y_1,Y_2} &= -\iint_D f_{Y_1,Y_2}(y_1,y_2)\cdot\ln(f_{Y_1,Y_2}(y_1,y_2))dy_1dy_2 \\
&= -\iint f(h_1(y_1,y_2),h_2(y_1,y_2))|J(y_1,y_2)|\cdot[\ln(f(h_1(y_1,y_2),h_2(y_1,y_2))) + \ln(|J(y_1,y_2)|)]dy_1dy_2 \\
&= -\iint \ln(|J(y_1,y_2)|)f(h_1(y_1,y_2),h_2(y_1,y_2))|J(y_1,y_2)|dy_1dy_2 - \\
&\qquad \iint \ln(f(h_1(y_1,y_2),h_2(y_1,y_2)))f(h_1(y_1,y_2),h_2(y_1,y_2))|J(y_1,y_2)|dy_1dy_2 \\
&= -E_{Y_1,Y_2}\ln(|J(Y_1,Y_2)|) + e_{X_1,X_2} \\
&= E_{X_1X_2}\ln(|J^{-1}(X_1,X_2)|) + e_{X_1,X_2},
\end{aligned}$$

where J^{-1} is the Jacobian of the inverse transformation, $-|\,\partial g_i/\partial x_j\,|$.

Example: transformation from cartesian to polar coordinates

Suppose X_1 and X_2 are independent and identically distributed as N(0,1). Then by results shown in the preceding section, the joint entropy of X_1 and X_2 is $1+\ln(2\pi)$.

Now suppose $Y_1 = X_1^2 + X_2^2$, the squared radial "distance" from (0,0), and $Y_2 = \mathrm{Tan}^{-1}(X_2/X_1)$, the "radial angle," so the region D is $[0,\infty)\times[0,2\pi)$. We note Y_1 has entropy $\ln(2) + 1$, by the fact that Y_1 is a sum of squares of independent standard normal random variables, hence has a chi-square distribution with two degrees of freedom, which is an exponential distribution with parameter (1/2), whose entropy was computed in Section III-4. We note further that the radial angle Y_2 is uniformly distributed over the interval $[0, 2\pi)$, so its entropy, $\ln(2\pi\text{-}0)$, is also given by an expression in Section III-4. Finally, we observe that Y_1 and Y_2 are independent, so we expect the joint entropy of Y_1 and Y_2 will be of the form $[\ln(2) + 1] + \ln(2\pi)$. Let us illustrate how this value can be obtained directly from the transformation and the joint entropy of X_1 and X_2, using the result above.

Note the inverse of the above transformation is

$$X_1 = \sqrt{Y_1}\cos(Y_2);$$
$$X_2 = \sqrt{Y_1}\sin(Y_2),$$

so the Jacobian is

$$J = \begin{vmatrix} \frac{1}{2\sqrt{y_1}}\cos(y_2) & -\sqrt{y_1}\sin(y_2) \\ \frac{1}{2\sqrt{y_1}}\sin(y_2) & \sqrt{y_1}\cos(y_2) \end{vmatrix} = \frac{1}{2}.$$

Then it follows by the result above that the joint entropy of Y_1 and Y_2 is given directly by

$$\begin{aligned} e_{Y_1,Y_2} &= E_{X_1X_2}\ln(|J^{-1}(X_1,X_2)|) + e_{X_1,X_2} \\ &= E_{X_1,X_2}\ln(2) + (1 + \ln(2\pi)) \\ &= (\ln(2)+1) \;+\; \ln(2\pi) \\ &= e_{Y_1} \;+\; e_{Y_2}, \end{aligned}$$

in agreement with the anticipated value.

These ideas can, of course, be extended to higher dimension spaces, and transformations that are piece-wise 1-1 over regions forming a partition of the plane. In some cases, the resulting integrals do not converge, as can be seen by considering $g(X) = e^X$, where X is distributed as t with one degree of freedom. In other cases the results can be somewhat novel, as is the case for the entropy of the radial miss distance, $\sqrt{Y_1}$ in the above example. In this case, by the result at the beginning of this section,

$$\begin{aligned} e_{\sqrt{Y_1}} &= E_{Y_1}\ln(\frac{1}{2\sqrt{Y_1}}) + e_{Y_1} \\ &= -\ln(2\sqrt{2}) - \int_0^\infty (\ln(t))e^{-t}\,dt + e_{Y_1} \\ &= -\ln(2\sqrt{2}) + \gamma + \ln(2) + 1 \\ &= 1 - \ln(\sqrt{2}) + \gamma, \end{aligned}$$

where γ is Euler's constant ($\gamma \approx .5772$).

III-7. Information Gain Rate, δ′, with a Fixed Prior

The fourth axiom of information gain and the sixth corollary (Section III-1) combine to simplify the expression for the derivative of $\delta(\mathbf{p},\mathbf{p}^*(t))$, where we imagine a fixed prior $\mathbf{p}(t_0)$ at a fixed time t_0, and a posterior $\mathbf{p}^*(t)$ that changes with time $t > t_0$. In this case we may regard δ to be a function of t, written as $\delta(t_0,t)$, and the derivative can be written as

$$\delta'(t_0,t) = \lim_{\Delta t \to 0} \frac{\delta(\mathbf{p}(t_0), \mathbf{p}^*(t+\Delta t)) - \delta(\mathbf{p}(t_0), \mathbf{p}^*(t))}{\Delta t}$$
$$= -\lim_{\Delta t \to 0} \frac{\delta(\mathbf{p}(t_0), \mathbf{p}^*(t)) - \delta(\mathbf{p}(t_0), \mathbf{p}^*(t+\Delta t))}{\Delta t}$$
$$= -\lim_{\Delta t \to 0} \frac{-\delta(\mathbf{p}^*(t), \mathbf{p}(t_0)) - \delta(\mathbf{p}(t_0), \mathbf{p}^*(t+\Delta t))}{\Delta t}$$
$$= \lim_{\Delta t \to 0} \frac{\delta(\mathbf{p}^*(t), \mathbf{p}^*(t+\Delta t))}{\Delta t}$$

We note the last expression above does not depend on the prior, $\mathbf{p}(t_0)$, as should be the case. To find the information gain rate, one need only measure the gain from t to $t + \Delta t$, using only $\mathbf{p}^*$.

Example for a discrete case.

As a specific example, if the distribution $\mathbf{p}^*(t)$ is discrete, so the entropy of the posterior distribution is $-\Sigma p^*_i(t)\ln(p^*_i(t))$, we obtain from the preceding expression, $\delta'(t_0,t)=$

$$= \lim_{\Delta t \to 0} \frac{-\sum p^*_i(t)\ln(p^*_i(t)) + \sum p^*_i(t+\Delta t)\ln(p^*_i(t+\Delta t))}{\Delta t}$$
$$= \sum \lim_{\Delta t \to 0} \frac{p^*_i(t+\Delta t)\ln(p^*_i(t+\Delta t)) - p^*_i(t)\ln(p_i(t))}{\Delta t}$$
$$= \sum \frac{d}{dt} p^*_i(t)\ln(p^*_i(t))$$
$$= \sum p_i^{*\prime}(t)[1 + \ln(p^*_i(t))]$$

which can be verified directly by taking the derivative of $-\Sigma p_i(t_0)\ln(p_i(t_0)) + \Sigma p^*_i(t)\ln(p^*_i(t))$, exploiting the fact that the prior entropy is constant.

REFERENCES

[1] Barr, D.R., and T. Sherrill, "Estimating the Operational Value of Tactical Information," Department of Systems Engineering Technical Report, U.S. Military Academy, 1995.

[2] Barr, D.R., M. Tillman and S. Strukel, "Entropy Measures of Reconnaissance," ORCEN Technical Report, U.S. Military Academy, 1994.

[3] Barr, D.R., J. Armstrong, S. Strukel, and M. Tillman, "A Methodology for Determining the Value-Added of Reconnaissance, Proceedings of the NATO Helicopter Symposium, SHAPE Technical Center, 1993.

[4] Barr, D.R., M. Weir and J. Hoffman, "An Indicator of Combat Success," *Naval Research Logistics*, Vol. 40, 1993.

[5] Carroll, N., W. Glaser and M. Webb, "A Study in Recon," Department of Systems Engineering, U.S. Military Academy, SE489 Course Report, inst. M. Tillman, 1993.

[6] Champion, M., *et al*, "Intelligent Minefield Evaluation System," Department of Systems Engineering Technical Report, U.S. Military Academy, West Point, NY, 1996.

[7] Khinchin, A. I., *Mathematical Foundations of Information Theory*, translated by R. Silverman and M Friedman, Dover Publications, Inc., 1957.

[8] MacWillie, S., "Contributions of Reconnaissance: MOE/MOP Guide," Directorate of Combat Developments, U.S. Army Aviation Center, Combat Dev. Pamphlet, 1992.

[9] Moraved, H. and M. Blackwell, "Learning Sensor Models for Evidence Grids," *1991 Annual Review*, The Robotics Institute, ed. Melvin W. Siegal, 1992.

[10] Shannon, C., "A Mathematical Theory of Communication," The Bell System Technical Journal, **27**, 1948.

[11] Sherrill, T. and D. Barr, "Measuring Information Gain in Janus," Department of Systems Engineering Technical Report, U.S. Military Academy, West Point, NY, 1996.

[11] Sonntag, R. E. and G. J. VanWylen, *Fundamentals of Statistical Thermodynamics*, John Wiley & Sons, Inc., 1966.

[12] Stone, L. D., "Search for the *SS Central America*: Mathematical Treasure Hunting," *Interfaces*, Vol. 22, No.1, 1992.

[13] Zehna, Peter W., *Probability Distributions and Statistics*, Allyn and Bacon, Inc., Boston, 1970.