
Compatibility of Land SAR Procedures with Search Theory

Prepared for

U.S. Department of Homeland Security
United States Coast Guard
Operations (G-OPR)
Washington, D.C. 20593-0001

Contract DTCG39-00-D-R00009
Task Order DTCG32-02-F-000032

Prepared
By

D. C. Cooper
J. R. Frost
R. Quincy Robe

Potomac Management Group, Inc.
30 December 2003

Abstract

The widely-accepted science of search theory as described by Koopman (1946, 1980), Stone (1989) and others was incorporated into the first edition of the *National Search and Rescue Manual* in 1959 after the U. S. Coast Guard provided the first comprehensive application to civil SAR in the 1950s. Applied search theory quickly gained acceptance by maritime SAR agencies worldwide and has remained in global use ever since. Various practical improvements and modifications to search planning techniques and data have been made over the years, but the application of the underlying theory remains unchanged, as shown in the *International Aeronautical and Maritime Search and Rescue Manual* (IAMSAR Manual, 1999) and recognized globally as the standard text on aeronautical and maritime SAR operations and methods.

After a preliminary review of the available land search planning literature at a special meeting of the National Search and Rescue Committee (NSARC) Research and Development (R&D) Working Group in March of 2001, it was determined that the results of scientific operations research as it relates to searching may not have been effectively applied by those working in land search. In partial response to this, the NSARC R&D Working Group tasked Potomac Management Group, Inc., with reviewing current published methods used for searching areas for lost, missing, or distressed persons on land who are in need of assistance. The purpose was to gain familiarity with current terminology and procedures, and to identify which procedures are compatible with the application of formal search theory to land search, which could become compatible with practical revisions, and which cannot be revised in a practical manner to achieve compatibility. The findings of this review are included in this report.

This report concludes that it does not appear there has ever been a comprehensive attempt to apply the science of search theory to the development of land search planning and techniques. The report finds that various individuals at various times have attempted to apply bits and pieces of what they believed to be search theory to the problem. There is clearly a great deal of room for improvement as search theory can make substantial contributions if properly applied. There is also a critical need to rectify some of the more crucial misunderstandings that could have a significantly detrimental effect on future inland search operations.

This report recommends the following:

1. Developing a standard methodology for land search planning.
2. Refining and validating the procedures for establishing land sweep width values.
3. Performing sweep width experiments for the land SAR environment.
3. Developing computer-based search planning decision support tools for land SAR.
4. Developing improved resources allocation guidance for area land searches.
5. Improving procedures for estimating POD on land.

Table of Contents

Abstract	ii
Table of Contents	iii
List of Tables	1
List of Figures	2
Acknowledgements	3
1. Introduction	4
1.1 Statement of the Problem	4
1.1.1 Formal search theory and its practical application	4
1.1.2 Applicability of Search Theory to Land SAR	8
1.1.3 Necessity and Benefits of Applying Search Theory to Land SAR	9
1.1.4 Development of current land search methods	13
1.2 Purpose of This Report	16
1.3 Definitions	16
2. Search Theory: Overview of Planning Aspects	17
2.1 Objectives of Search Planning	17
2.2 Elements of Search Planning	17
2.2.1 Sweep Width, Effort, Search Effort, Coverage and Probabilities	18
2.2.1.1 Sweep Width	18
2.2.1.2 Effort and Search Effort	24
2.2.1.3 Coverage	25
2.2.1.4 Probability of Detection (POD)	26
2.2.1.5 Probability of Containment (POC) or Area (POA)	27
2.2.1.6 Probability of Success (POS)	28
2.2.1.7 Probability Density Distribution (so POA can be estimated)	28
2.2.2 The Exponential Detection Function	30
2.2.3 Optimization Criterion of Maximizing POS Subject to Effort Constraint	30
2.2.3.1 Optimal Search Density vs. Optimal Searcher Path	30
2.2.3.2 T-Optimal and Uniformly Optimal Search Plans	31
2.3 Search Planning Process	31
3. Selected Land Search Procedures	33
3.1 Land Search References to Search Theory	33
3.1.1 May (1973)	33
3.1.2 Kelley (1973)	33
3.1.3 Wartes (1974-75)	34
3.1.4 Syrotuck (1974-75)	35
3.1.5 Bownds et al. (1981, 1991a-c, 1992); and Bownds & Harlan et al. (1991)	35
3.1.6 LaValla et al. (1997)	36
3.1.7 Colwell (1992, 1994)	37
3.1.8 Hill (1997)	38
3.1.9 Stoffel (2001)	39
3.1.10 Dougher (2001) & Dougher et al. (2001)	39
3.2 Probability Density Distributions	40
3.2.1 Probability and Segments	40
3.2.2 Methods of Developing Initial POA/POC Values	41

3.2.2.1	Mattson	41
3.2.2.2	Relative Ranking with Letters	42
3.2.2.3	Relative Ranking with Numerals	44
3.2.2.4	Ranking by Descending “Priority”	46
3.2.3	Normalizing Adjusted POA Values	46
3.2.3.1	Methodological History: Syrotuck, 1975	47
3.2.3.2	Methodological History: Bownds, c1982	48
3.2.3.3	Methodological History: Shea, 1988	49
3.2.3.4	The Normalization Issue	50
3.2.3.5	An Improved Method	50
3.2.4	Conclusions	51
3.3	Detection Functions	52
3.3.1	Koopman’s Approach to Detection Functions	52
3.3.2	References to Land Search Detection	54
3.3.2.1	Kelley, 1973	54
3.3.2.2	Wartes, 1974	56
3.3.2.3	Syrotuck, 1974, 1975	60
3.3.2.4	Bownds et al. (1981) and Bownds & Harlan et al. (1991)	62
3.3.2.5	Perkins (1989)	64
3.3.2.6	Colwell (1992)	65
3.3.3	Multiple Searches vs. Single Searches	68
3.3.4	One POD for Multiple Objects	71
3.3.5	Conclusions	72
3.4	Effort Allocation	73
3.4.1	“Non-thorough” Effort Allocation	73
3.4.2	“Sector Ladder”	75
3.4.3	Information Theory	76
3.4.4	POD as a Goal Used in Lieu of POS	78
3.4.5	“Rest of the World” (ROW)	81
3.4.6	Conclusions	85
3.5	Land Search Planning “Methodology”	86
3.5.1	Hill, 1997	86
3.5.1.1	Initial Planning Actions	87
3.5.1.2	Developing the Initial Probability Density Distribution	88
3.5.1.3	Establishing Incident Objectives and Evaluating Resources	88
3.5.2	Stoffel, 2001	89
3.5.2.1	Initial Actions	90
3.5.2.2	Developing the Initial Probability Density Distribution	91
3.5.2.3	Establishing Search Objectives and Evaluating Resources	92
3.5.3	Dougher, et al., 2001, & Dougher, 2001	93
3.5.4	Cooper, 2002	94
3.5.5	Conclusions	95
4.	Summary, Conclusions, and Recommendations	96
4.1	Summary	96
4.2	Conclusions	97
4.2.1	Needed Changes	97

4.2.2	Required New Empirical Data and Methods of Acquiring It	99
4.2.2.1	Sweep Width Data	99
4.2.2.2	Lost Person Behavior	100
4.2.3	Inland Search Planning Methodology.....	100
4.2.4	Land Search Literature Reviews and Critiques	101
4.3	Recommendations.....	102
4.3.1	Develop a Standard Methodology for Inland Search Planning	102
4.3.2	Perform Sweep Width Experiments for the Inland SAR Environment	103
4.3.3	Develop Computer-based Tools for Land SAR.....	103
4.3.4	Develop Resource Allocation Guidance for Area Searches	104
4.3.5	Objectives for Future Work	105
4.3.6	Steps to Accomplish the Objectives for Future Work	105
4.3.6.1	Validate and Refine Procedures for Establishing Sweep Width Values.....	105
4.3.6.2	Extend and Modify to Include Det. Data for Aerial Search from Aircraft	106
4.3.6.3	Develop Improved Procedures for Estimating POD.....	106
4.3.6.4	Develop Outline for Practical Search Planning Methodology.....	107
4.3.6.5	Describe Functional Requirements for Software Tools.....	107
4.3.6.6	Survey Existing Software	108
4.3.6.7	Develop Complete Land Search Planning Support Software	108
5.	References.....	109
Appendix A.....	Selected Inland Search Definitions	119
Introduction.....	Notation.....	119
Selected Definitions		119
Appendix B		134
Analysis of Bayesian Updates and ROW		134
Generalization of Postulate 3 [Addition Rule].....		134
General Rule of Multiplication		134
The Rule of Elimination.....		134
The Rule of Bayes.....		135
Non-Detection “Events”		136
Experiment 1		137
Experiment 2 – A Non-search Example		139
Experiment 3 – Optimal Effort Allocation		141
The Role of ROW (“Rest of the World”)		145
“Completeness” of ROW		146
ROW and Effort Allocation.....		148
ROW and Decision Support.....		148
Comparison of Bayesian-adjusted ROW “POA” and $OPOS_{CUM}$		150
Experiment 4 – An Easter Egg Hunt.....		150
Experiment 4-A.....		150
Experiment 4-B.....		151
Search Planning vs. Incident Co-ordination		152
Further Analysis.....		153
Utility of Un-normalised POA Values.....		153

Answering Search Planning/Management Questions.....	154
Convenient Features of the “Random” Search Curve.....	155
The Ultimate Objective.....	156
Appendix C.....	157
Simplified Explanation of Sweep Width.....	157
An Analogy.....	157
Importance of <i>Sweep Width</i>	162
Appendix D.....	163
Probability Density Distributions and Probability Maps.....	163
Probability Density and Its Importance.....	164
Creating Probability Density Distributions.....	166
Circular Normal Probability Density Distributions.....	166
Uniform Probability Density Distributions.....	167
Generalized Probability Density Distributions.....	168
Estimating Probability Densities.....	169
Estimating Containment Percentages Directly.....	169
Ranking the Regions.....	170
Ranking the Regions – Again.....	171
An Assessment Based on Density Estimates.....	173
Another Short Exercise.....	173
Analysis of Results.....	175
Proportional Assessment.....	175
Containment vs. Density Estimates.....	176

List of Tables

Table 1-1	VPCAS (OASIS) Operational Results.....	10
Table 1-2	U.S. Coast Guard Statistics.....	11
Table B1	Probability and SAR Notation Comparison.....	135
Table B2	Probability and SAR Notation Comparison, Non-detection Events.....	136
Table B3	Initial Values for Figure B2.....	142
Table B4	Values After One Charnes-Cooper (1958) Iteration.....	144
Table B5	Values After Two Charnes-Cooper (1958) Iterations.....	145
Table B6	Initial Values for Figure B3.....	148
Table C1	Broom Experimental Results.....	161
Table D1	Gymnasium Floor Values.....	165
Table D2	Tom’s Assessment.....	170
Table D3	Dick’s Assessment.....	171
Table D4	Mary’s Assessment.....	172
Table D5	Density-based Assessment.....	173
Table D6	Cellular Assessment of Figure D8.....	174
Table D7	Simple Ranking Assessment of Figure D8.....	175
Table D8	Proportional Rating Assessment of Figure D2.....	176

List of Figures

Figure 1-1	Comparison of SARPlan and CSAD (from Abi-Zeid, 2001)	12
Figure 2-1	A Lateral Range Curve	21
Figure 2-2	A Uniform Random Distribution of Search Objects.....	22
Figure 2-3	Effective Sweep Width for a Clean Sweep.....	22
Figure 2-4	Effective Sweep Width.	23
Figure 2-5	POD vs. Coverage (Koopman, 1946)	26
Figure 2-6	Probability map.....	29
Figure 3-1	POD versus Coverage curves for three detection functions.	54
Figure 3-2	POD versus Spacing (Wartes, 1974).	59
Figure 3-3	POD versus Spacing in “Critical Separations” (Perkins, 1989).	65
Figure 3-4	POD versus Spacing in meters (Colwell, 1992).	67
Figure 3-5	POS versus Time (exponential detection).....	75
Figure A1	POD vs. Coverage (Koopman, 1946)	120
Figure A2	A search area with four regions and their POA values after a consensus.....	123
Figure A3	egments within regions and their associated POA values.	124
Figure A4	A search area showing regions and a grid overlay.	125
Figure B1	Simple initial search situation.....	137
Figure B2	Another search problem.....	141
Figure B3	A third search problem.....	146
Figure C1	Four brooms (B1, B2, B3, and B4).....	157
Figure C2	Broom 1 (B1)	158
Figure C3	Broom 2 (B2).....	158
Figure C4	Broom 3 (B3).....	159
Figure C5	Broom 4 (B4).....	159
Figure C6	Broom performance profiles.....	160
Figure D1	Probability map.....	163
Figure D2	Gymnasium floor illustration.....	165
Figure D3	Probability map showing probable error.....	167
Figure D4	Probability map for a uniform probability density distribution.	168
Figure D5	Probability map after accounting for multiple probable errors.....	169
Figure D6	Probability map using four-letter ranking scale.....	171
Figure D7	Probability map based on Mary’s assessment.	172
Figure D8	Gymnasium floor illustration; new experiment.	174

Acknowledgements

The material reviewed for this report was obtained from several sources. The U. S. Coast Guard provided a copy of the *International Aeronautical and Maritime Search and Rescue (IAMSAR) Manual* (IMO/ICAO, 1999) and a copy of the U. S. *National Search and Rescue Supplement to the IAMSAR Manual* (NSARC, 2000). The National Association for Search and Rescue (NASAR) graciously provided copies of Hill (1997), Cooper, et al. (1996), and other relevant publications they distribute and with which NASAR-certified instructors teach search operations, planning and management. The Mountain Rescue Council of England and Wales provided several documents used in the United Kingdom of Great Britain and from which they teach search planning and management. Two private for-profit concerns (whose owners were until recent years partners in a single concern) provided copies of selected materials they have used to teach the management of land search operations in the U.S., Canada, and several countries abroad (e.g., Australia, Ireland, Iceland, New Zealand, South Africa, Sweden, the U.K., and others). The primary publications used in these courses currently include LaValla, et al. (1997), Dougher, et al. (2001), Dougher (2001), and Stoffel (2001). A number of miscellaneous documents, including privately published, limited distribution, and historical papers, were also obtained from the private libraries of various authors and instructors in the land search and rescue (SAR) community. A comprehensive list of these is provided in the references section.

1. Introduction

1.1 Statement of the Problem

1.1.1 Formal search theory and its practical application

The mathematical basis of searching for objects has been a subject of serious study since World War II, when the main impetus was detecting enemy submarines. B. O. Koopman (1946, 1980) established the basis for a rigorous study of search theory and practice with his pioneering work for the U. S. Navy during WWII. Koopman was a member of the Navy's *Operations Evaluation Group* (OEG). An important characteristic of this group was that its members were required to spend several years in the field working directly with operations personnel. All work produced by this group had to be both scientifically sound and practical enough for operational use by Navy personnel without requiring them to have any special scientific training. It also had to show practical results. The work initially done by the OEG made a significant contribution to winning the Battle of the Atlantic against the German U-boats. Although this kind of application may seem far removed from searching for lost persons on land, the basic theory of search Koopman established applies to all types of searching, including that of inland search (Koopman, 1980).

Koopman's (1946) work was published shortly after the war in classified form as Operations Evaluation Group Report 56, *Search and Screening*. This report was later declassified and the author eventually expanded and republished this work in his book of the same title (Koopman, 1980). Although search theory was applied to military SAR operations during and after WWII, the U. S. Coast Guard provided the first comprehensive application to civil SAR in the 1950s. The methodology was incorporated into the first edition of the *National Search and Rescue Manual* in 1959 and it quickly gained acceptance by maritime SAR agencies worldwide. It has remained in global use ever since. Various practical improvements and modifications to search planning techniques and data have been made over the years, but the application of the underlying theory remains unchanged, as shown in the *International Aeronautical and Maritime Search and Rescue Manual (IAMSAR Manual, 1999)* published jointly by the *International Maritime Organization* and the *International Civil Aviation Organization* and recognized globally as the standard text on aeronautical and maritime SAR operations and methods.

An essential part of Koopman's work was developing the concept of effective search (or sweep) width—an objective numeric measure of how easy or hard it is for a given sensor to detect a given object in a given operating environment. Whenever Koopman's basic theory has been applied, substantial improvements in search success rates and reductions in the average times and resources required to achieve success have been realized.

Over the years, several surveys, indexes and bibliographies of search theory literature have been published (Enslow, 1966; Pollock, 1967; Dobbie, 1968; Moore, 1970; Strumpfer, 1980; and Benkoski, Monticino, & Weisinger, 1991). The hundreds of documents included in these surveys offer numerous examples of the extensive study that has gone into the science of operations research and search theory. Because summarizing all these works would not be useful here, only selected significant references will be described.

In his seminal work, Koopman (1946) analyzed the mechanics and geometry of encounters between moving sensor platforms and both moving and stationary targets (search objects). He went on to analyze the instantaneous detection probabilities during such encounters as a function of the true range between the sensor and the search object—a function whose exact form and values were subject to three classes of variables, namely the characteristics of the sensor, the search object and the environment that affect detection. Koopman (1946) extended the analysis to integrate the instantaneous detection function along the sensor's track relative to the search object to develop detection probability as a function of lateral range—the distance between the sensor and the search object at the closest point of approach, assuming the sensor approached the target along a straight track from some large distance away, passed by the target, and continued on for some significant distance. From this analysis, the author developed the concept of effective search (or sweep) width—a single value quantifying the ability of a specific sensor to detect a particular object while both are operating in a given environment. He then analyzed the performance of several hypothetical detection models, characterized by different lateral range curves, when used to cover an area with a pattern consisting of long straight equally spaced parallel tracks on the one hand, and when used along short randomly placed, but uniformly distributed, tracks on the other. For the latter case, Koopman (1946) derived the so-called “random search formula” and showed that it applied to all types of sensors when used in the “random” fashion he described. For parallel track search patterns, he showed that the detection probability curve as a function of “sweep density” (also known as “effort density” or “coverage”) varied with differences in the lateral range curves. In particular, Koopman (1946) devised a hypothetical mathematical model of visual detection based on the geometry of sighting moving vessels from patrolling aircraft (the so called “inverse cube law of sighting” [p. 22]), and computed its corresponding detection probability as a function of “sweep density.” He found that it fell roughly midway between that of a sensor having perfect detection out to some definite range and no detection beyond that range, and the random search curve. As a result, Koopman (1946) observed,

At one extreme is the case of the definite range law, at the other the case of random search. All actual situations can be regarded as leading to intermediate curves, i.e., lying in the shaded region. The inverse cube law is close to a middle case, a circumstance which indicates its frequent empirical use, even in cases where the special assumptions upon which its derivation was based are largely rejected (p. 31).

It is this “inverse cube law” curve relating probability of detection to coverage that has been the standard for the U. S. Coast Guard over the past 40 years. Finally, Koopman (1946) was the first to articulate, in formal mathematical terms, the general problem of finding the optimal distribution of searching effort (also known as “optimal effort allocation”) and it was he who made the first forays into its solution. In short, it was Koopman (1946) who first established the basic principles of search theory.

Later, Koopman (1980) updated and extended his seminal 1946 report, but the fundamental concepts of his original work remained unaltered.

Using Koopman's (1946) random search (exponential) detection function for a stationary target, Charnes and Cooper (1958) described the optimal allocation of effort as a convex programming problem and developed an algorithm to obtain it. Extending the concept to a larger class of detection functions, Stone (1989) described optimal effort allocation for all “regular” detection

functions. “A detection function is regular if its first derivative is continuous, positive, and strictly decreasing” (Benkoski, et al., 1991, p. 473). Wegener (1981) extended these results further to include non-concave detection functions and provides a procedure for constructing an optimal plan.

Benkoski et al. (1991), the latest of the search theory surveys, developed a list of some 239 search theory references from the scientific literature. Of these, ten were bibliographies or other surveys of the search theory field, twelve were texts, and the remaining 217 were articles from various scientific journals, proceedings, etc. This survey, as the authors painstakingly acknowledge in their introductory paragraphs, is not all-inclusive. Their introduction bears recording here, not only to let the authors speak for themselves, but for the insights they provide into the types of search problems that are analogous to searching for lost or missing objects or persons and the types that are not. Benkoski et al. (1991, pp. 469-470) states the following in its introduction:

In 1980, B. O. Koopman published an updated and extended version (Koopman [1980]) of his classic 1946 report, *Search and Screening* (Koopman [1946]). In the preface to this work he states that the original work ‘contained no references to scientific literature on search: none existed.’ It is always easier to survey the literature when you are the first one in the field.

The aim of this article is to provide a comprehensive review of the existing, published search theory literature. However, the four and one-half decades since Koopman’s initial report have seen the growth of a substantial body of work in search theory. This work would easily fill many library shelves. The sheer bulk of these references makes it difficult to compile a complete list and raises doubts about the usefulness of producing merely an unannotated list of publications. As a result, any survey of the search theory literature must be selective – and hence, to some extent, subjective. This survey is no exception.

In approaching this task, we have made two basic assumptions. The first has been to restrict our attention to books and articles that can be found in a typical university library. Our second assumption has been our choice of a definition for ‘search theory.’ Informally, we might include in search theory any mathematical framework involving the search for lost or hidden objects. However, we have employed a much narrower definition. In particular, we have restricted our review to publications which

- Emphasize search planning, not search modeling
- Take a tactical, not strategic, viewpoint
- Assume inspections have uncertainty
- Aim more at obtaining initial detections, than at fusing multiple detections
- Involve a noncooperative target.

Our first restriction means that we are interested in techniques for planning effective search. Generally, we have excluded articles which deal with the construction of specific models for target motion, sensor effectiveness, or the search environment. The second restriction indicates that we are concerned with the details of allocating search effort. So, for example, high-level predator/prey

models which consider search only as a renewal process with detections at exponentially distributed times are excluded. The third restriction excludes the binary search trees, sorting problems, and weighting problems often discussed in the computer science literature. Our fourth restriction means that we view search theory in a low-data-rate environment. Techniques such as Kalman filtering which can fuse multiple contacts to refine the estimate of the target location are excluded. Finally, the fifth restriction indicates that our targets are passive (one-sided search) or evading (search games). We rule out cooperative games such as the rendezvous problem in which the participants search for each other. Also, we exclude problems from coding theory in which one attempts to search for the meaning in a friend's message that has arrived via a noisy channel.

The restrictions noted above are important and generally appropriate to the SAR search problem. Certainly target motion, sensor effectiveness and the search environment are important factors in practical search planning and modeling. Nevertheless, they are also largely separable from the *theoretical* effort allocation problem or they can have general representations. These factors are separable in the sense that search theory generally assumes the search object's location probability density distribution and the geographic area it covers have already been estimated or assumed. Operationally, all three factors impact the precise nature of this distribution, but not necessarily the general methods for optimally deploying effort against it. However, since optimal effort allocation techniques based on search theory must have a probability density distribution to work against, and since some operational techniques for estimating such distributions in SAR are highly questionable, we will need to include motion, sensors, and the environment in this study. As an example of generalization, a particular sensor's effectiveness against a particular search object under a particular set of environmental conditions is generalized by the concept of *sweep width*. The sweep width concept is central to search theory, as discussed later, and its omission from land search planning methods is the single most important factor rendering their present effort allocation and search evaluation techniques ineffective. Because we will be dealing with the practical application of search theory to real-world search problems, we will not have the luxury of ignoring specific models of search object motion, sensor effectiveness, the search environment, or methods for estimating the search object location probability density distribution in this study.

The appropriateness of the second restriction is obvious – SAR searches are clearly tactical in nature. The third restriction simply means that searching does not guarantee detection, even if the search object is in the area when it is searched. SAR search planners the world over have seen this property of SAR searches demonstrated all too often. The fourth restriction may be somewhat less obvious to those who have not actively participated in SAR missions. Nevertheless, it should be reasonably clear after a moment's reflection that SAR is necessarily performed in a low-data-rate environment and is clearly aimed at obtaining initial detections. Obtaining sufficient information to effectively carry out a SAR response requiring a search to be undertaken is almost always a difficult task. The fifth restriction is also reasonably typical of SAR situations, even if signaling devices are present. Generally, the distressed person does not seek out the searchers and the effects of signaling devices on detectability can be handled by adjustments to the sweep width. Evasion is also not unheard of in SAR, especially when searching for a lost child who has been taught to avoid contact with strangers, or a hunter with a delicate ego who would rather follow the searchers covertly until he can find his own way out, than be found and have to admit he was lost. The final sentence in the above quotation will be recalled later in

this study when discussing a publication that claims a valid analog exists between coding theory and searching for missing persons.

One other observation recorded by Stone (1989) regarding the current status of search theory should be included here: "...[T]he development of cheap and powerful microcomputers with high resolution color graphics has allowed the development of real-time interactive search planning systems. Because of these computer developments, the trend in search theory is toward algorithms for computers and away from the theorem-proof style of presentation given in this book" (p. xv).

It is clear that a wealth of pertinent search theory literature exists. It will be shown later in this paper that the existing land SAR search planning methods have not taken full advantage of this research.

1.1.2 Applicability of Search Theory to Land SAR

The main objective of applied search theory is finding the allocation of available search assets, under the conditions prevailing at the time and place of the search, that maximizes the probability of successfully locating the search object and minimizes the time required to find it. These goals are important ones for most searches, but particularly so for SAR searches. Maximizing the probability of success at the greatest possible rate with the available search resources will save more lives by finding and assisting persons in need more quickly. Time is an important factor in saving lives, especially when the distressed person is injured or the weather has become extreme enough to threaten continued survival. Locating distressed persons sooner also reduces costs of SAR operations and reduces risks to searchers because they are exposed to the hazards of searching for less time.

The principles of search theory apply to any situation where the objective is to find a person or object contained in some geographic region in the most efficient manner. In this context, "efficient" means minimizing the time required to find the search object while maximizing the chances for finding it with the available resources. For stationary objects, it can be shown that maximizing the probability of success (POS) also minimizes the mean time to find the object. Since the goal of searching in the context of land SAR is to find the distressed person(s) as quickly as possible, the optimal effort allocation features of search theory clearly apply.

Koopman (1980) stated that the history and development of operations research applies to, "...various fields, including search" (p. 12) and that what has come to be known as search theory is applicable to all types of searches including those involving missing persons. Although Koopman's (1946, 1980) early work was clearly aimed at naval interests, the general theory of search established is applicable to virtually any type of search problem. Stone (1989) also claims that, "Searches for persons lost on land also fit into the framework of [search] theory" (p. 1).

Washburn (2002) includes an entire section (7.3) on "Inland Search" and concludes:

Whatever the explanation for the current lack of sophistication in inland search theory and TDAs [tactical decision aids], and whatever the organizational difficulties of

achieving a remedy, it is probably true that better TDAs in the hands of well trained inland search managers would find more targets. There is evidence of such improvements for both the Navy and the Coast Guard in the maritime case. Inland search, with a larger incident base and additional complicating features such as the effects of terrain, should be an even better candidate for improvement (p. 7-10).

The application of search theory requires a quantitative description of the detection function(s) that correspond to the sensors being used and their method of employment, the search object being sought, and the nature of the operating environment in which the searching is being done. A detection function in this context is a function that relates the probability of detecting an object if it is in a given area to the amount of searching effort expended in that area. No such functions are described in the literature on which current land search planning procedures are based.

In the maritime case, two such detection functions do exist. One is based on Koopman's (1946) "inverse cube law of sighting" when employed with search patterns consisting of perfectly straight, parallel, equally spaced tracks relative to the search object. The other is based on Koopman's (1946) "law of random search" (p. 71). The effects that the characteristics of the sensor, search object, and environment have on detection are quantified in a single value called the "*effective search (or sweep) width*" (p. 65). The average density of the searching effort over the searched area is quantified by *coverage*, the ratio of the *area effectively swept* to the physical *area* over which the effort was approximately uniformly spread. It is important to note that coverage is proportional to the level of *effort* required to achieve it. These key concepts are missing from current land search planning methods. They are explained further in section 2.2.1.

1.1.3 Necessity and Benefits of Applying Search Theory to Land SAR

In the editor's preface to Stone (1989), J. D. Kettle observes,

The search process is inherently a nervous one. Either you will find the 'target' or you won't. This involves more stress than the continuous penalties or payoffs associated with dullness or brilliance in dealing with problems such as scheduling or logistics. This discontinuity makes search a little like litigation. During an actual 'case' there is a sense of urgency and emergency. This stress can trigger a major, sometimes frantic, effort. Experts can be mobilized. Armies (or navies) can be sent scurrying around. A nervous principal or client can make intuitive decisions that are painfully wrong. In short, search theory is a delightful challenge for operations research (p. i).

Two corollaries to search theory's primary objective of optimal effort allocation include (a) preventing "intuitive [effort allocation] decisions that are painfully wrong," and (b) avoiding disorganized, inefficient, and ineffective "HS" (helter-skelter) searches. These corollaries are particularly crucial in SAR where lives are often at stake. The only way to avoid making unfortunate effort allocation decisions that waste effort and delay detection is to develop more objective decision-making methods. Search theory provides the tools for developing and applying such objective methods to SAR searches.

As shown below, when the existing guidance for planning searches is replaced or enhanced by that derived from search theory, mathematical models and actual experience show that substantial increases in the probability of success (POS) will be realized for the same expenditure of effort. These models also show correspondingly significant reductions in the average time needed to find the search object.

The U.S. Navy (USN) has used theory-based search planning techniques extensively in many, if not most, of its search problems—including many classified applications. Although computerized implementations of search theory are standard practice today (U.S. Coast Guard Research and Development Center [USCG R&D], 2001), with the USN it began in the 1970’s as a means of fulfilling their anti-submarine warfare (ASW) mission. To produce their own computer search planning program, the USN developed the concepts behind the Coast Guard’s Computer Assisted Search Planning (CASP) system. This USN system was called VPCAS (Aviation Patrol Computer-Assisted Search system, formerly Operational ASW Search Information System [OASIS]). As it was introduced to the personnel who planned searches, it was used in parallel with the manual methods that preceded the VPCAS system (Benkoski, 1978, and McCoy, 1978). Analysts compared the effectiveness of the two systems and found that the probability of success was twice as high for the VPCAS system as it was for the manual method (USCG R&D, 2001)(Table 1-1).

Table 1-1
VPCAS (OASIS) Operational Results
(USCG R&D, 2001)

CASE	Manual Success %	VPCAS Success %
I-A	32% (157)	73% (48)
I-B	20% (157)	56% (48)
II-A	43% (65)	82% (17)
II-B	25% (65)	65% (17)
III-A	32% (65)	71% (17)
III-B	17% (65)	53% (17)

Note 1. This is an unclassified version of the analysis performed by Benkoski (1978) and McCoy (1978).

Note 2. The numbers in parentheses indicate the number of searches over which the percentage was calculated.

Note 3. VPCAS used the optimal-effort allocation for moving targets algorithm developed by Brown (1980) to maximize POS.

Note 4. VPCAS tended to be used on harder searches because it was less familiar to users than the manual method. Operators used it only when they felt it was needed, which was on the more difficult searches.

If the average effort expended on VPCAS and manually planned searches was the same, then Table 1-1 shows that VPCAS plans produced over twice as much success, on average, for the same amount of effort. Table 1-1 demonstrates the significant potential for improvement the scientific search theory-based methods may offer over the published inland methods that apparently lack a rigorous scientific basis—especially if computerized search planning tools can be developed to aid the inland search planning function.

The ultimate goal of any SAR system is to reduce the number of lives lost and the USCG’s SAR Program is no different. Table 1-2 compares the number of lives saved in the four years prior to the implementation of their Computer-Assisted Search Program (CASP) to the number of lives saved subsequent to its implementation.

Table 1-2
U.S. Coast Guard Statistics
(from USCG, 2002)

Period	Average Number of Lives at Risk Per Year	Average Number of Lives Saved Per Year	Average Number of Lives Lost Per Year
Prior to change in search planning methodology (1971 – 1974)	4105	2681 (65%)	1424 (35%)
After change in search planning methodology (1975 – 1978)	4977	3632 (73%)	1345 (27%)
Difference (%)	+ 872 (21.2%)	+ 951 (26.1%)	- 79 (5.5%)

Note 1. Four years on either side of the implementation of a new search planning methodology are shown.

When a comparison is made of data collected before and after the modification of search planning systems in the U.S. Coast Guard (four years before and after 1975; reporting methods changed in 1979), statistics show a 5.5 percent reduction in the average number of lives lost per year (United States Coast Guard [USCG], 2002; see Table 1-2). This increase in lives saved is particularly impressive when viewed in light of the fact that over the same period there was a significant (21.2%) increase in the number of lives at risk (potential lives to be saved) and lives saved (26.1%)(USCG, 2002). This represents an average increase of 951 lives saved per year and appears to support the findings illustrated in Table 1-1. Although many other factors likely had an influence on USCG improvements, there is no doubt that CASP played a significant role.

In further corroboration of these data, Dr. Irene Abi-Zeid, in a presentation to the United States National Search and Rescue Committee (NSARC) Research and Development Working Group (22 May 2001), illustrated how superior automated search planning methods with integrated op-

timal search effort allocation algorithms were when compared to manual methods that used no such algorithms. Abi-Zeid showed that, given varying levels of effort, SARPlan could produce a significant improvement in success when compared to the older, manual Canadian Search Area Definition (CSAD) method (Frost, 2001; see Table 1-1). SARPlan is software developed by the Canadian Defence Research Establishment Valcartier (DREV) that included a near-optimal effort allocation algorithm. As with the USCG example, when the cost of conducting search operations planned by the CSAD manual method is compared to those planned with the SARPlan software, the potential to save a great deal of time and money by automating and optimizing search effort allocation is obvious.

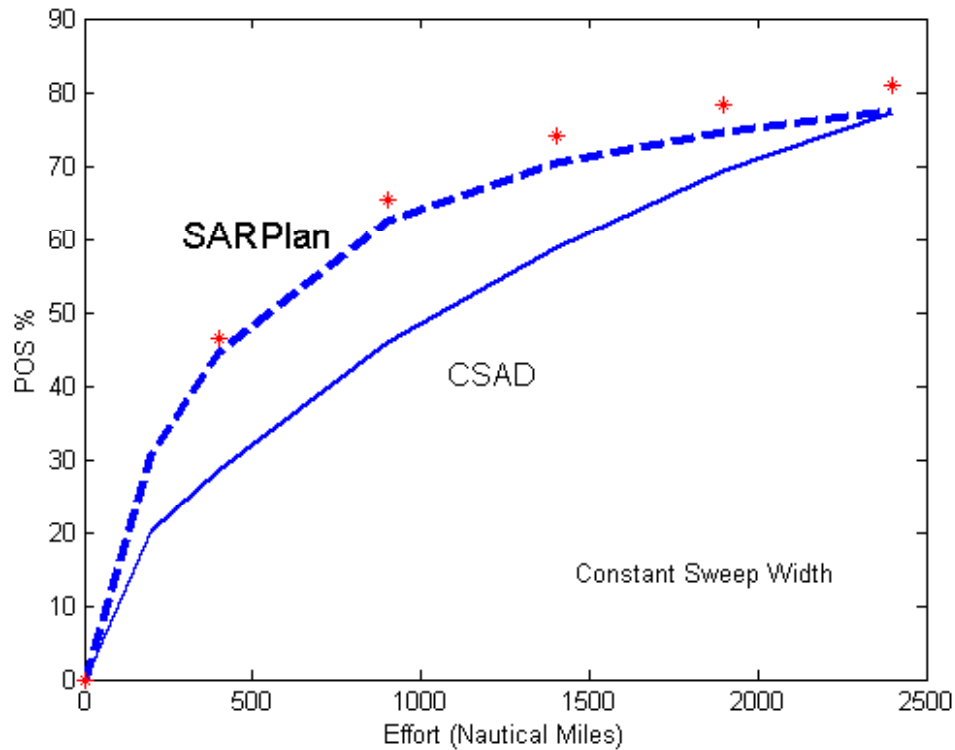


Figure 1-1. Comparison of SARPlan and CSAD (from Abi-Zeid, 2001; used with permission)

Examples of the successful use of scientific search theory-based methods abound. Richardson (1967) described his successful application of search theory to find a Hydrogen bomb on the ocean floor in 1964. The same principles were applied to the impressively successful search for the wreck of the submarine USS Scorpion that was found within 260 yards of the highest probability cell in the distribution (Richardson & Stone, 1971).

In 1985, Stone (1992) took on the task of developing a probability distribution and search plan for locating the SS Central America, a side-wheel steamer that sank in 1857 in 8000 feet of water nearly 200 miles from land. On board were some 425 people and 400 million (US) dollars worth of gold bars and coins. According to Stone (1999, p. 20), “This wreck had been the target of many previous searches, but...none were successful.” In 1989, the Columbus-America Discovery Group recovered one ton of gold bars and coins from the sunken wreck. The brief three years

spent on searching for the SS Central America, according to Stone (1999, p. 20), was, "...a sharp contrast to typical treasure hunting operations where individuals spend many years, or even entire lifetimes, in unsuccessful efforts involving far less difficult search conditions." Up-front investment in the application of the scientific principles of search theory apparently saved the Columbus-America Discovery Group anywhere from several to many years of expensive alternative search operations similar to those undertaken by competing groups. A group from the Lamont-Doherty Geological Institute of Columbia University performed an unsuccessful search for the Central America in the early 1980's. According to Stone (1992), "Their failure was...due to the lack of careful analysis and planning that the search theory supports" (p. 54). Around the same time, another commercial vessel attempted to locate the SS Central America: the Cameron Sea Horse chartered by Wally Kreisle. Their failure may have also been attributed to the fact that, although like many other treasure hunters they were willing to conduct search operations in the open sea, they did not for whatever reason apply the principles of search theory (Kinder, 1998).

The examples above show that search plans based on search theory often succeed where less scientific methods fail. It is abundantly clear that more scientific search planning using the levels of computing power now commonly available at low cost offers significant benefits. Decreasing the mean time to detect search objects saves resource hours and that translates directly into monetary savings (USCG R&D Center, 2001).

1.1.4 Development of current land search methods

For many years, the inland SAR techniques were advanced by individuals trying different methods of search and rescue (mostly rescue) and simply writing articles or presenting papers at SAR conferences about what seemed to work well and what did not. With the fielding of CASP circa 1974, the maritime SAR operations personnel were given their first introduction to actually quantifying the probability of the search object being in one place versus another, the probability of success concept, optimal search plans, etc. At about the same time, three members of the inland SAR community attempted to further the state of the art for inland search planning along similar lines. Circa 1974, Dennis Kelley (1973), Jon Wartes (1974, 1975a-b) and William G. Syrotuck (1974, 1975) were writing and publishing papers aimed at introducing similar improvements to land searches and putting inland search planning on a more "scientific" footing. Although the motives of these authors were undoubtedly true, and their desire to improve land search methods unquestionably genuine, the writing of these authors seem to indicate a less-than-complete understanding of search theory.

This incomplete understanding of the wealth of search theory by the developers of early land search methods is illustrated in what is likely the first published description of search theory in the inland literature: Kelley (1973). In it, the author initially describes how the theory of search was born out of Operations Research in WWII aimed at anti-submarine warfare, but makes no further use of the available information. This is in spite of the fact that the author described "coverage" (although quite differently than the 'coverage' used in search theory) and addressed the issues of search strategy, tactics, and planning.

William Syrotuck, another of the first land search authors to write about elements of search theory, referred several times to USCG search planning methods in his early work (Syrotuck, 1975,

pp. 21, 25, 34). However, the author's apparent misunderstanding of the concept of sweep width, along with his attempts to integrate it into land search planning, eventually led to a significant deviation from the accepted tenets of search theory.

The work of Jon Wartes was referenced several times in Syrotuck's (1975) writing. In short, Wartes (1974), performed limited detection experiments in an attempt to develop the concept of efficiency as a method to, "...produce more results or...produce the same results in less time" (p. 4). Unfortunately, the Wartes' (1974) study was fundamentally flawed as a detection experiment because only some aspects of the searcher/search object interaction were examined. In addition, search objects with widely differing characteristics were lumped together and treated in bulk. Furthermore, the results were analyzed as a function of searcher separation rather searcher/search object separation. The consequence of this was that the author deduced a relationship between searcher spacing and probability of detection (POD) that conflicts with the scientifically verified principles of search theory.

Later land SAR authors and experimenters generally continued in the same direction as the authors in the 1970s. In 1992, Martin Colwell undertook his own series of experiments in southwestern British Columbia, Canada. Again, the experiments were focused on the inappropriate goal of relating POD to between searcher spacing without using an objective measure of search object detectability. In addition, the data analysis was not grounded in search theory. As a result, questionable conclusions and recommendations on planning searches were published, as Colwell has since acknowledged.

Although search theory as it stands today is based on over a half century of formal, detailed and well-documented scientific research (Benkoski et al., 1991), current land search planning methods seem to be based primarily on ad hoc, informal research and opinion that evidently started in the early 1970s (Kelley, 1973; Syrotuck, 1974, 1975; Wartes, 1974, 1975a-b). Land search planning methods, and the associated literature, seem to depend little on the extensive formal research conducted by the maritime, aeronautical and business communities. Rather, conventional land search planning and management techniques have evolved independently of the scientific research over the past thirty years or so, primarily from operational trial and error and attempts to include the consideration and implementation of selected business management and operational field techniques into the planning of a search (LaValla, et al., 1997, Foreword).

As an example, LaValla, et al. (1997) contains an extensive bibliography but includes no entries from the wealth of scientifically established search theory literature. It appears that much of the research that is cited in this and similar texts came from individuals involved in land search operations who had little or no background in the science of search theory. Similarly, and not surprisingly, Benkoski et al. (1991), contains no references found in the bibliography included in LaValla et al (1997).

The land search literature appears to combine three different elements of search into one idea usually called "search management." This term is frequently used in the inland literature to describe varying combinations of incident response management, search planning, and search operations. Almost universally, the principal publications currently available in the land search literature in the United States (LaValla et al., 1997; Hill, 1997; Dougher et al., 2000; Dougher, 2001; and Stoffel, 2001) describe the use of the Incident Command System (ICS) to coordinate search resources and functions into a single, flexible incident management model. ICS is derived

from military and fire service models and is used by many agencies and organizations in the U.S. (e.g., federal agencies, emergency service organizations, NASAR, etc). Although ICS describes very well the general functions needed for organizing and managing an effective response to an incident (emergency or non-emergency), the ICS literature does not provide search-specific planning or operations guidance (NWCG, 1994; FEMA, 1998). Regardless, most of the literature that does provide specific guidance on the planning and conduct of land searches uses ICS as the axis around which the development and implementation of “search management” revolves.

Another universal characteristic of the principal management publications currently available in the land search literature in the United States is that they all provide some guidance on search techniques in the field in spite of the fact that they all purport to be directed toward managers of search operations (LaValla et al., 1997; Hill, 1997; Dougher et al., 2000; Dougher, 2001; and Stoffel, 2001). That is, they all offer some detailed suggestions on how the searching itself should be performed (tactics such as scanning techniques, formations, etc). This is probably due to the fact that when some of these materials were being developed specific guidance on land search strategies and tactics was limited and the authors were trying to fill the void as best they could.

The land search literature also contains limited information regarding the planning of searches and allocation of resources. However, it includes only incidental uses of selected elements of search theory. In contrast to the *National SAR Supplement* (NSARC, 2000) or *IAMSAR Manual* (ICAO/IMO, 1999a-c), most inland SAR manuals also do not contain a comprehensive, coherent search planning methodology. An exception to this rule was Hill (1997) where a plan of action for developing a search plan was outlined, and Stoffel (2001) where “a logical sequence for planning a search effort” (p. 191) was described but with few details and many missing elements. Instead, most contain a plethora of ideas, opinions, and, from search theory perspective, questionable effort allocation procedures. However, the inland manuals are like the *National SAR Supplement* (NSARC, 2000) and the *IAMSAR Manual* (ICAO/IMO, 1999a-c) in that they cover many aspects of SAR operations besides the planning of searches and contain a great deal of valuable practical advice in those areas. Except for a single short chapter in the *National SAR Supplement* (NSARC, 2000), the inland SAR methods and guidance in their use are not published by a governmental agency or international organizations whose delegates represent the governments of the member States. Instead, they are published by private organizations that use them in training courses provided for a fee.

After such an assessment, we must understand that all the above authors (and more yet to be described) were doing the best they could with the limited resources they had available. Volunteers form the backbone of land SAR operations. Funding for land SAR operations is extremely limited. Until recently funding for research has been virtually non-existent. The expense of the research that was done was borne almost entirely by the individual researcher and volunteers who assisted with the occasional help of a local business. These researchers are to be commended for having had the initiative, resourcefulness and persistence to do their research at significant personal sacrifice.

Land search planning methodology development has passed the point where informal “experiments” and anecdotal evidence about what worked or did not work are good enough. Unfortunately, it has taken until now to realize the full implications of this critical step, and this realization is still far from universal within the inland SAR community. The lack of general knowledge

within the community regarding search theory then, or even today, should not be surprising. Books on the subject can hardly be considered popular reading and they can be very difficult to find, even when using all the standard publication data to search for them. But, before a person can even begin such a search, he must first know of the subject's existence. Few people who participate in SAR know that search theory is a subject of serious scientific study.

If inland search planning is to be put on a scientific footing, it will require the use of scientifically rigorous methods. This is not to say the *operational* methods or software developed for field use will require users to have a rigorous scientific background. However, it will require both scientists and field personnel working together to develop truly useful, scientifically sound, and practical search planning tools.

1.2 Purpose of This Report

After reviewing the available land search planning literature, it is clear that the results of scientific operations research as it relates to searching has not been effectively applied by those working in land search. In response to this, the National Search and Rescue Committee (NSARC) tasked the NSARC Research and Development Working Group with hosting a meeting of land search experts. This was accomplished on March 24, 2001 in Laurel, Maryland and involved representatives from the U.S. Air Force, U.S. Coast Guard, U.S. National SAR School, U.S. National Park Service, National Association for Search and Rescue (NASAR), National Aeronautics and Space Administration (NASA), Federal Aviation Authority (FAA), Department of Transportation (DOT), Department of Defense (DOD) and private industry (represented by one of the world's leading authorities in search theory and its practical application, Dr. L. D. Stone). The meeting provided a unique opportunity for a small group of selected experts and other key persons from Canada and the United States to develop a preliminary assessment of land search planning needs. More specifically, the group was asked to identify efforts that could be undertaken by the federal government or others to develop practical improvements in land search planning through the application of science and technology. This study is one result of this tasking.

The Research and Development Working Group of the National Search and Rescue Committee tasked Potomac Management Group, Inc., with reviewing current published methods used for searching areas for lost, missing, or distressed persons on land who are in need of assistance. The purpose of this review was to gain familiarity with current terminology and procedures, and to identify which procedures are compatible with the application of formal search theory to land search, which could become compatible with practical revisions, and which cannot be revised in a practical manner to achieve compatibility. The findings are provided in this report.

1.3 Definitions

Due to the great quantity of special terms used herein, definitions have been listed in their own appendix (Appendix A).

2. Search Theory: Overview of Planning Aspects

2.1 Objectives of Search Planning

Searching is a truly ancient and ubiquitous activity. For this reason it is often taken for granted by the layman that searching is simply a matter of just looking around for the lost or missing object being sought. However, when a life may hang in the balance, such a simplistic approach is inappropriate, especially given the current state of knowledge about searching as a process.

During the Second World War a formal scientific discipline called search theory was established. The original work as well as all subsequent work has shown the "...operation of search as an organic whole having a structure of its own—more than the sum of its parts" (Koopman, 1980, p. 2). Although most would regard the mathematics of search theory as complex, it can be reduced for practical use to a few simple concepts and organizing principles. Implementing these concepts and principles in a manner appropriate to the type of search mission, operating environment and available search resources has repeatedly demonstrated its value. For the search and rescue (SAR) mission, the objective of all search planning is to deploy the available resources in a fashion that achieves maximum probability of success (POS) in the minimum time.

2.2 Elements of Search Planning

In 1957, the U.S. Coast Guard published its first search and rescue manual. This manual was to become the basis for the *National SAR Manual* that replaced it just two years later when the *National SAR Plan* was first adopted. Like its successors ever since, a large portion of this first manual was concerned with search planning. The search planning methodology contained in this manual was clearly based on the earlier work of Koopman (1946, 1956a-b, 1957). The fact that the SAR manual was published just after Koopman's (1956a-b, 1957) unclassified articles appeared is unlikely to have been mere coincidence. Koopman's (1946) original work was later declassified and updated (Koopman, 1980). Prior to the work of Stone (1989) no unified and comprehensive presentation on optimal search had been made. So, Stone (1989) further developed the work of Koopman and others to present a basic optimization technique, "...using Lagrange multipliers and maximization of Lagrangians...to solve most of the problems of optimal [effort] allocation..." (p. 2). In this work, Stone (1989, p. 3) defined the elements of the basic problem of optimal search. They are (paraphrased):

- A probability density distribution on search object location and state (so the probability of containment, POC (a.k.a. POA for "probability of area"), for any subset of the possible locations and states can be estimated),
- A detection function relating the probability of detecting (POD) the object if it is in a searched area to the density of the searching effort expended there,
- A known finite amount of available searching effort, and
- An optimization criterion of maximizing probability of finding the object in a desirable state (probability of success or POS) subject to the constraint on effort availability.

Given these elements, it is possible to develop an optimal search plan. However, translating the facts and assumptions about the case into a probability density distribution on search object location and state may be a quite complex undertaking. Such distributions are the combined result of many, often interrelated, factors. The available data often span a range of accuracies and degrees of relevance, contain inconsistencies, and come from sources of varying reliability. Consequently, considerable analytical skill is required to develop scenarios on which to base search plans that are logically consistent with substantial subsets of the available data.

The following material is presented here to set the stage for later discussions. The terms *optimal search density*, *optimal searcher path*, *T-optimal search plan*, and *uniformly optimal search plan* are taken from Stone (1989). The definitions given here are very loose paraphrases plus our own commentary.

2.2.1 Sweep Width, Effort, Search Effort, Coverage and Probabilities

The concept of *sweep width* is central to search theory as it relates to finding lost or missing persons or objects. Other terms related to optimal allocation of search resources include *effort*, *area effectively swept*, and *coverage*. Since we will have need of these terms later in this discussion, it is best that they be defined and described at this point.

2.2.1.1 Sweep Width

Sweep width is one of the central concepts of search theory and its application to SAR. The term *sweep width* has a specific mathematical definition different from what one might infer from the usual meanings of its component words. Therefore, we should discuss the term at least briefly before proceeding further and provide at least one or more informal definitions. References to more complete and mathematically rigorous discussions will be provided.

Sweep width is a single number characterizing the average ability of a given sensor to detect a particular search object under a specific set of environmental conditions. Thus each combination of sensor, search object, and set of environmental conditions will have a particular associated sweep width. In the vernacular, sweep width might be called a measure of “raw detection power.” Loosely paraphrasing Koopman (1980), sweep width may be described as follows:

Consider a sensor moving with constant velocity through (or over) a swarm of uniformly distributed, identical, stationary search objects under constant environmental conditions. If the average number of objects detected per unit time is divided by the object density (average number of objects per unit area), the resulting value is called the *effective search or sweep rate*. It is easy to see that the effective sweep rate has dimensions of area over time (e.g. square miles per hour). Dividing the effective sweep rate by the speed of the sensor gives the *effective search or sweep width*, which has units of length.

Notice that the above description does not imply that every object in the “swept area” is detected. Indeed, the meaning of “swept area” itself is not clear. To clarify how the term *sweep width* got its name, we will give an alternative description (also loosely derived from Koopman, 1980):

Consider an omnidirectional sensor that is “perfect” (i.e. 100% effective) within some definite range and completely ineffective beyond that range. That is, detection is guaranteed for any object the sensor approaches more closely than the definite detection range, and the sensor never detects any object beyond that range. This idea is analogous to setting a lawn mower’s blade to a height of zero and then pushing it into tall grass. The lawn mower would leave behind it a swath of bare earth having a definite width (twice the definite detection range), while blades of grass outside this width would be untouched. Inserting this particular sensor into the previous description, it is easily seen that in this special case (and this special case alone), the sweep width is literally the width of the swept area where the detections took place, i.e. twice the definite detection range. The concept is generalized by defining the *effective sweep width* of any sensor as equal to the sweep width of a definite range sensor that detects the same number of objects per unit time as the given sensor does under identical circumstances (i.e., same sensor speed, same object density, same environmental conditions). Generally the word *effective* is dropped, shortening the term to just *sweep width*. This is sometimes a source of confusion to new students of search theory and also to search planners in the field.

We see that in only one situation, namely definite range detection, does the sweep width actually correspond to a physical, geometric width measurement. Otherwise, it is a more abstract concept, but nevertheless one of great value and utility on both the theoretical and operational fronts. Additional treatments of the sweep width concept, some with illustrations, may be found in Koopman (1980), Stone (1989), and Frost (1998c, 1999b).

Unfortunately, sweep width cannot be measured directly for cases other than definite range detection. This is one reason why it is difficult to explain. Another reason is the ease with which the term “sweep width” is confused with other, sometimes similar, terms that have quite different meanings and uses. We will now rectify this problem by giving several different, but equivalent, descriptions of what sweep width represents.

For all of the following descriptions, assume that search objects are uniformly, but randomly, spread over an area. A uniform random distribution means that the search object locations occur at random so their positions cannot be predicted, but the number of objects per unit of area is about the same everywhere. Also assume that the area covered with objects is very large compared to the maximum detection range.

Suppose an experiment was done where every searcher detected every object within a given lateral range, say 50 feet either side of the searcher’s track, and detected no objects outside that range. That is, the searchers were 100% effective within 50 feet on either side of their track, and completely ineffective for objects farther from the searcher’s track. This would constitute a “clean sweep” of a swath 100 feet wide with no detections outside that swath. The effective sweep width in this case would be 100 feet. In this “ideal” but unrealistic example, the effective sweep width is the same as the width of the swath where objects were detected.

Now suppose another experiment is done in another venue using the same number of objects per unit of area. Further suppose that the searchers in this experiment find objects that are up to 100 feet either side of their tracks, but they detect, on average, only half the objects located in that swath of 200 feet. Note that there will be twice as many objects in a 200-foot swath as in a 100-

foot swath of the same length. Therefore, even though the searchers detect only half of those present in the 200-foot swath, they will detect just as many objects in one pass as the searchers in the previous experiment did. In this sense the two groups of searchers performed equivalently despite any differences in terrain, vegetation, searcher training, etc. So, for purposes of estimating how many objects will be detected in one pass, we would say the *effective sweep width* in both cases was 100 feet. That is, both groups of searchers detected the same number of objects as lay in a swath 100 feet wide even though only the first group did this in a literal sense.

This illustrates the difference between effective sweep width and maximum detection range. While it is possible to say that the width of the swath where searchers *can* detect objects will normally be about twice the maximum detection range, there is no way to predict from that information alone how many of the objects present in that swath *will* be detected, even if the number of objects present per unit of area is known. The effective sweep width, on the other hand, does allow us to estimate how many detections we should expect provided we also know the number of objects present per unit of area. Simply multiply the effective sweep width by the length of the searcher's track to get the area effectively swept then multiply this value by the number of objects per unit of area to get the number of detections that should be expected. Note that this value does not depend in any way on the maximum detection range and there is no known mathematical relationship between the two. Having a maximum detection range in one situation that is twice that of another situation does *not* mean objects in the first situation are twice as detectable, on average, as objects in the second situation. In fact, it is actually possible that a small, high-contrast object might have a very large maximum detection range in a given environment under just the right circumstances but be less detectable on average in that environment than a larger object with less contrast and a smaller maximum detection range. Knowing the maximum detection range does not help with POD estimation. Also note that just as knowing the maximum detection range does not tell us the effective sweep width, knowing the effective sweep width provides no information about the maximum detection range. However, knowing the effective sweep width gives us a way to reliably estimate POD since it is a measure of expected detection performance.

The effective sweep width may be thought of as the width of the swath where the number of objects *NOT detected inside* the swath are equal to the number of objects that *ARE detected outside* the swath. That is, when one gets to the point where the number of objects missed within a certain distance either side of track (areas B above the curve in Figure 2-1) equals the number that are detected at greater distances from the searcher's track (areas A below the curve in Figure 2-1), then one has found the effective sweep width.




Figure 2-1. A Lateral Range Curve. The number of missed detections (B) inside the effective sweep width equals the number of detections (A) that occur outside the sweep width.

For the more mathematically inclined who are familiar with calculus, the effective sweep width is also numerically equal to the total area under the lateral range curve down to the horizontal axis of the graph. One way to estimate effective sweep width from experimental data is to analyze the detection/non-detection results to first get an estimate of the lateral range curve and then compute the area under that curve. However, this is significantly more difficult than some other data analysis methods.

Finally, if detection were perfect (100% POD) within a swath of width W and completely ineffective (0% POD) outside that swath, then the effective sweep width would be W . That is, if a “clean sweep” were possible with no detections outside the swept swath, the width of the swath would be, by definition, the effective sweep width. Sensors with perfect detection within some definite maximum detection range and perfectly sharp cutoffs at that definite maximum detection range do not exist. However, this perspective on sweep width reveals another important property: The effective sweep width can never exceed twice the maximum detection range. It is almost always considerably less than that value, but just how much less depends on the search situation and all the factors affecting detection. It is not possible to establish any general mathematical relationship between maximum detection range and effective sweep width.

Figures 2-2, 2-3, and 2-4 below illustrate the concept of effective sweep width in another way. The black dots in Figure 2-2 represent identical search objects that have been scattered randomly but approximately uniformly over an area. The distribution is “uniform” because in any reasonably large fraction of the area there are about the same number of objects as in any other fraction of the same size. The distribution is “random” because the exact location of each object was chosen at random to avoid producing either a predictable pattern or a bias favoring one portion of the area over another.




Figure 2-2. A Uniform Random Distribution of Search Objects

Figure 2-3 shows the effect of a “clean sweep” where all of the objects within a swath are detected and no objects outside the swath are sighted. In this case the effective sweep width is literally the width of the swept swath. A total of 40 objects lay within the sweep width and all 40 were detected, as indicated by the empty circles. A “clean sweep” where the searcher/sensor is 100% effective out to some definite range either side of the track is unrealistic, but it serves to illustrate the sweep width principle.




Figure 2-3. Effective Sweep Width for a Clean Sweep.
 Dotted line represents searcher's track. Number missed within sweep width = 0.
 Number detected outside sweep width = 0.

Figure 2-4 represents a more realistic situation where objects are detected over a wider swath, but not all the objects within that swath are found. In this case, the total number of objects detected was also 40 but instead of making a “clean sweep,” the detections are more widely distributed. However, because in both cases 40 objects were detected over the same length of searcher track when the number of objects per unit of area was also the same, we say the *effective* sweep widths for both cases are equal.

Effective sweep width is a measure of detectability because, in a hypothetical situation where the average number of objects per unit of area is known, if we know the sweep width we can accurately predict how many of the objects will be found, on average, by single searchers on one pass through the area. As we will show later in this report, knowing the sweep width for a given combination of sensor (e.g., visual search), search object (e.g., a person) and environment (weather, terrain, vegetation, etc.) will allow us to accurately predict the probability of detection for any search conducted under those or similar conditions.




Figure 2-4. Effective Sweep Width.
 Dotted line represents searcher's track. Number missed within sweep width = 11.
 Number detected outside sweep width = 11.

Figure 2-4 also illustrates the property of effective sweep width where the number of undetected objects inside the swath equals the number of objects detected outside that swath.

Appendix C contains further clarification of the sweep width concept. An analogy is drawn between searching and sweeping floors. This analogy is used to provide a simplified non-technical explanation of effective sweep width.

To summarize: Sweep width is the metric used for estimating an object's detectability for a given search scenario. It is a single number having the dimensions of length. It may be derived from the lateral range curve that is produced from detection/non-detection data of an experiment that is appropriately designed and performed. It has the property that, on average, the number of search objects detected outside the effective sweep width is numerically equal to the number of search objects not detected within the effective sweep width (Figure 2-1 and Figure 2-4). It is used together with the amount of effort expended in a given area (e.g., a search segment) and the size of the area to get an objective, reliable, and accurate estimate of POD.

As a practical matter, it is not possible to directly "measure" sweep width at the place and time of a search. It is also impossible to develop sweep width values for the infinitely many possible combinations of sensor, search object, and environmental conditions. The Coast Guard has addressed these problems by designing and conducting numerous experiments to gather empirical data from which operationally useful sweep width estimates may be inferred. The Coast Guard's Research and Development Center has been conducting such experiments for more than twenty years, identifying the significant variables affecting operational sweep widths in the marine environment and producing extensive sweep width tables indexed to these variables. These tables are published in the U. S. *National SAR Supplement* (National Search and Rescue Committee [NSARC], 2000) and in a simplified derivative form in the *International Aeronautical and Maritime Search and Rescue Manual* (ICAO/IMO, 1999a-c).

Recently the National Search and Rescue Committee (NSARC) sponsored research to develop an experimental procedure for sweep width estimation on land that would be suitable for distribution to SAR agencies and groups around the country (Robe and Frost, 2002). These entities could then perform such experiments in their respective areas of responsibility, obtain sweep width estimates for their own use, and report their results for others to use and/or compare with their own experimental results. A notable feature of these procedures that is absent from all previously published work on land SAR POD "experiments" is that a detailed record of every detection opportunity and its outcome (both detection and non-detection) is made on an individual basis and the results are carefully analyzed. A detailed treatment of a method for determining sweep widths in land searches (procedures for conducting detection experiments) is available in Robe & Frost (2002).

2.2.1.2 *Effort and Search Effort*

Effort is a measure of resource expenditure and may be defined as the amount of distance covered by the searcher(s) in a search segment while searching. It could be measured in several ways, but the usual metric for search theory purposes is the distance a sensor platform travels while in the search segment. A search segment is defined as some bounded geographic area that a particular resource, such as a team of searchers, has been assigned to search. The distance a searcher covers while searching may be estimated by either estimating or recording the amounts of time spent searching (exclusive of rest or meal breaks, transit times to and from the assigned segment, etc.) and multiplying that value by the estimated average search speed using the familiar formula

$$d = rt$$

for *d*istance equals *r*ate times *t*ime. When a team of searchers is assigned a given segment, the total distance traveled by all members of the team will be needed. This value may be found by summing all the individual team member distances or, if all members moved at about the same speeds for about the same amounts of time while searching, then the distance covered by one searcher could be multiplied by the number of persons in the team to get the total distance covered in the segment. That is,

$$Effort = \sum_{i=1}^n d_i \text{ or } Effort = nd$$

where *n* is the number of searchers on the search team.

Search effort is a measure of how much “effective” searching is done by the sensor as it moves through the search area. *Search effort* is simply the product of the sweep width and the distance the sensor travels while in the search area or:

$$Area \text{ Effectively Swept} = Effort \times Effective \text{ Sweep Width}$$

It is easy to see that *search effort* has units of area. It is often called *area effectively swept*.

2.2.1.3 Coverage

Coverage (sometimes called *coverage factor*) is a relative measure of how thoroughly an area has been searched, or “covered.” *Coverage* is defined as the ratio of the area effectively swept to the physical area of the segment that was searched:

$$Coverage = \frac{Area \text{ Effectively Swept}}{Segment's \text{ Area}}$$

Searching an area and achieving a *coverage* of 1.0 therefore means that the *area effectively swept* equals the area searched. Note that this does not necessarily mean that every piece of ground was scanned nor does it mean that the POD of a coverage 1.0 search is at or near 100%. *Coverage* is a measure of how “thoroughly” the segment was searched. The higher the coverage, the higher the POD will be. However, the relationship is not linear. That is, doubling the coverage does not double the POD. Figure 2-5 (POD versus Coverage curve) shows the relationship between coverage and POD as derived by Koopman (1946, 1980) for situations where searchers do not move along a set of long, perfectly straight, parallel, equally spaced tracks but instead follow more irregular paths.

It is important to always remember that coverage and the corresponding level of effort are proportional. To double the coverage it is necessary to double the level of effort and doubling the level of effort doubles the coverage. In other words, although the relationship between POD and coverage is not linear, the relationship between coverage and effort is. This means, by extension, that the relationship between effort and POD is not linear, either. Doubling the effort assigned to a segment will not generally double the POD.

Since terrain and vegetation often prevent ground searchers from following a mathematically precise pattern of parallel tracks, and since ground searchers frequently alter their tracks to investigate possible sightings, look behind major obstructions, etc., the exponential detection function, as the curve in Figure 2-5 is called, seems to be the most appropriate for estimating ground search POD. This curve also works well when other “random” influences are present, such as uneven terrain and vegetation, even when the searcher tracks are perfectly straight, parallel, and equally spaced. The equation of this curve is

$$POD = 1 - e^{-Coverage}$$

where e is the base of the natural logarithms (approximately 2.718282). The function e^x or EXP is available with most handheld scientific calculators and electronic spreadsheet programs.

It can be seen that *coverage* is proportional to *search effort density*, the constant of proportionality being the *sweep width*. Therefore, any solution to the optimal search density problem is also a solution to the optimal coverage problem. In this sense, the two terms may be used interchangeably when discussing optimal search plans.

2.2.1.4 Probability of Detection (POD)

The probability of detection (POD) is defined as the *conditional* probability that the search object will be detected during a single sortie *if* the search object is present in the area searched during the sortie. Cumulative POD (POD_{cum}) is the cumulative probability of detecting the search object given that it was in the searched area on each of several successive searches of that area. Like coverage, it is a measure of how thoroughly an area was searched. The relationship between coverage and POD is usually plotted on a graph of POD vs. Coverage. Such a graph appears in Figure 2-5.




Figure 2-5. POD vs. Coverage (Koopman, 1946)

POD in itself is not the goal of search planning as some of the land search literature has suggested. POD is merely one part of a larger system.

2.2.1.5 *Probability of Containment (POC) or Area (POA)*

The probability of containment (POC), or the probability of area (POA) as it is called in the land search community, is the probability that a geographically bounded area contains the search object. POC and POA are exact synonyms and usage generally just depends on the author and/or the intended audience. Since the topic of this review is land search planning methods, the term “POA” will be used henceforth. If the probability density distribution function is known, the POA may be found by multiplying the mean probability density over the area by the size of the area. Conversely, if the POA is known, then the mean probability density may be computed by dividing the POA by the size of the area.

The POA is not a quantity that can be determined from objective experiment. In the maritime world, the distribution is estimated by giving the possible positions of a vessel a bivariate normal distribution with the reported distress position as the mean. The usual method for quantifying the amount of error about the mean is to estimate the probable error—the ellipse centered on the mean that contains 50% of the distribution. Post incident motion of the search object is estimated by using the mean wind (direction and speed), mean sea current (direction and speed), and mean leeway speed to compute an overall mean drift velocity. This velocity is multiplied by time lapse between the incident and the first search to obtain the mean drift displacement (direction and distance). This displacement vector is added to the incident position to estimate the mean datum position, and the probable error of this position is computed from the probable errors of the incident position and each of the factors affecting the object’s post-incident drift. After incorporating any safety factors, a distribution of probable search object locations (probability map) is developed that illustrates the POA of each of many equally sized cells, as shown in Figure 2-6.

In a land search situation, an objective analysis of historical data and lost person behavior can play a large part in the estimation of POA. In the land search community, POA is established subjectively by search planners based on experience; local conditions of terrain, vegetation, and weather; and the behavior of persons of similar status and condition.

Since POA is largely subjective the procedures for arriving at the value should be as consistent and rigorous as possible. In the recent versions of land SAR instructive material there is a good deal of agreement on exactly how to arrive at an estimate of POA for each search segment.

Although various search planning publications have approached the problem of estimating POA in various ways, there is general agreement that a form of consensus should be used. Any consensus method is likely acceptable as long as it retains proportionality between regions of probability (Wagner, 1989).

This is an area of research that shows great promise for reducing the mean time to locate lost or missing persons in need of assistance. However, there was insufficient time to look into this matter very deeply for this study. Based on presentations at SAR conferences by K. Hill, D. Heth and R. Koester (all behavioral scientists), research done so far seems to indicate somewhat predictable behavior patterns among members belonging to certain groups. For example, a lost hiker will most likely behave in one way (they seem to favor continued forward motion as opposed to retracing their steps) while an Alzheimer’s patient will most likely behave in another. Age also seems to play a role, as lost children do not seem to do the same things or take the same routes a lost adult would. Typical behaviors may also be dependent on locale, and the list goes

on. The benefits of “profiling” the lost person and using data obtained from many similar cases should be obvious. The resulting estimate of the probability density distribution of subject locations should be both much more accurate and much smaller than if no such profiling was used, especially in situations where individual traits or circumstances do not provide strong clues or information about them is not readily available. Such research may even be applicable in the marine environment since survivors often can, and sometimes do, significantly affect the movement of their survival craft in an attempt to save themselves, thus invalidating drift estimates. At the present time, much, perhaps all, of the research into lost person behavior seems to be unfunded.

2.2.1.6 *Probability of Success (POS)*

The probability of success (POS) is the probability that expending a certain level of effort in an area will successfully locate the search object. POS is simply the product of POD and POA for the area in question. That is, the probability of finding the object is simply the probability that it was in the area at the time of the search (POA) times the probability that it would have been detected had it been there (POD). POS is the value optimal search planning techniques seek to maximize.

2.2.1.7 *Probability Density Distribution (so POA can be estimated)*

A primary goal of search planning is to determine not only where to search in general, but also how to deploy the available effort in the most efficient manner. An essential factor in deciding how much effort to place in each portion of the search area is an estimate of how the probability density is distributed over the search area. Probability density (Pden) is simply defined as

$$Pden = \frac{POC}{A}$$

where *POA* is the probability that the search object is contained in some area and *A* is the size of that area. If the probability density distribution function is known, the POA may be found by multiplying the mean probability density over the area by the size of the area. Conversely, if the POA is known, then the mean probability density may be computed by dividing the POA by the size of the area.

A probability density distribution is usually represented by a probability map consisting of a regular grid of cells. Each cell is then labeled with its POA value. Since all cells are equal in size, a cell’s POA value is proportional to its Pden value. This type of display has the dual advantages of showing at a glance both how much probability each cell contains and where the highest probability densities lie. Although the POA and Pden values are not numerically equal, a cell with twice the POA value of another cell also has twice the Pden value of that other cell when a regular grid is used. Figure 2-6 is an example of a probability map.

3.23%	3.23%	6.45%	6.45%	6.45%
3.23%	3.23%	6.45%	9.68%	12.90%
3.23%	3.23%	9.68%	9.68%	12.90%

Figure 2-6. Probability map.

To determine where to search, we must first estimate where the lost or missing person could be. This requires a careful, deliberate, thoughtful assessment of all the available information as well as the continual seeking of additional information from all possible sources. “Available information” is an all-inclusive term referring to every scrap of evidence and data that might shed some light on the lost person’s probable locations. In addition to data about a specific incident, statistical data from similar situations, such as lost person behavior profiles, can be very useful. Historical data can also be useful, especially in popular recreational areas.

In SAR situations, data is frequently obtained from a variety of sources and is often inconsistent. However, such data also tends to form a number of self-consistent sets that each tell a “story” about what might have happened and where the lost person might be. These “stories” are called scenarios. Careful analysis of each scenario is then required to estimate the lost person’s probable locations if that scenario is true. These estimates are then quantified as probability maps, thus defining that scenario’s probability density distribution. The different scenarios are then subjectively “weighted” according to the search planner’s perceptions of their relative accuracy, reliability, importance, etc. and their probability maps are then combined appropriately. Probability maps for different scenarios are generally combined by computing, for each cell in an area large enough to include all scenarios, the weighted average (using subjective scenario weights) of the cell probabilities from each scenario.

Unfortunately, formal search theory does not shed much light on how go about turning an inconsistent body of evidence and data from a variety of sources into numbers on a probability map. As Stone (1983), a leading authority on search theory and its practical application, observes, “One of the greatest difficulties in generating prior [to searching] probability maps is the lack of systematic, proven techniques for eliciting subjective inputs for search scenarios.” He goes on to say, “In addition to obtaining subjective probabilities, we also have the problem of obtaining subjective estimates of uncertainties, times, and other quantitative information needed to form scenarios” (p. 213).

Scenario development and analysis is a complex, difficult, mentally demanding task requiring a good deal of concentration, attention to detail, and mental discipline. Appropriate resources should be dedicated to this task and insulated from the often frenetic, and always distracting, op-

erational activities. This frequently seems difficult to do in SAR situations. The first impulse is to get as much search effort as possible into the field as soon as possible because statistics show that a lost person's chances for survival decrease rapidly as time passes. While there is nothing wrong with mounting a large initial effort (provided more effort is on the way) based on only a cursory evaluation of the situation, too often this is not followed up with a more deliberate evaluation and planning effort for subsequent searching should the initial efforts fail. In a few publicized cases, it appears that lost persons who could have, and should have, been saved were not found in time – sometimes in spite of huge expenditures of effort in relatively limited areas. This appears to have been a result, at least partially, of poor analysis and planning.

2.2.2 The Exponential Detection Function

The detection function most commonly assumed in search theory is the exponential detection function, also known as the random search detection function. This function is used so often for several reasons. First, it is operationally realistic and applicable to a wide variety of real-world situations. Second, it has the property that it depends only on applying the available effort uniformly over the area being searched; i.e., it depends only on the *search effort density*. Third, it also has the property that expending a given amount of effort in an area always produces the same chances of finding the object, regardless of whether all the effort is expended in a single search or expended piecemeal in a series of searches. In short, the exponential detection function is both operationally viable and easy to work with mathematically. Many algorithms used to produce optimal search plans are based on this detection function.

2.2.3 Optimization Criterion of Maximizing POS Subject to Effort Constraint

2.2.3.1 *Optimal Search Density vs. Optimal Searcher Path*

There are two (at least) distinct ways to think of the optimal effort allocation problem. One is to determine the optimal allocation of *search effort density* over the distribution of probability density describing the possible search object locations. From this point of view, there are no constraints on sensor movement. It is assumed that search effort may be placed wherever it is needed, whenever it is needed and in whatever quantity it is needed, regardless of where or whether search effort is already being applied. In addition, the theory assumes there are no “overhead” or logistics costs associated with dividing or applying the available effort in such a flexible fashion. It is assumed that the available search effort may be instantaneously applied piecemeal in increments of any size less than its total value, i.e., that search effort is infinitely divisible. Despite this fine disregard for real-world operational constraints and costs, computing optimal search effort densities can still produce extremely useful results for real-world search planners. Such results can also be efficiently computed even for quite large problems involving objects that both move and undergo changes in state that affect both their motion and detectability characteristics.

Note: Methods for determining optimal search plans given certain types of constraints (*constrained optimization*) have been found. For example, a common operational constraint is the requirement to apply search effort uniformly over a large area, as opposed to varying the search

density from place to place within such an area. However, the intent here is to contrast optimal search density with optimal searcher path problems. The distinction is one of great importance.

The other approach to optimal search planning is to determine the optimal *searcher path* the sensor should follow as it travels through the distribution. Although this sounds like it is oriented along more operational lines, it is a very difficult problem to solve. In fact, it is tantamount to solving the classical NP-complete “traveling salesman” problem. (The traveling salesman problem, or “TSP” for short, is this: given a finite number of “cities” along with the cost of travel between each pair of them, find the cheapest way of visiting all the cities and returning to your starting point. *NP* stands for Non-deterministic Polynomial and *complete* means no solutions exist for which the upper bound on the number of required steps may be expressed as a polynomial in terms of the problem’s size or complexity where the order of the polynomial remains fixed.) This means that the difficulty of solving the problem (and the time required on a computer) does not increase as the square of the problem’s size or complexity, or the cube, or any other fixed power. Instead, the number of steps required for a solution increases exponentially with the number of “cities” (or probable search object positions) on the “tour.”

Thanks to search theory research, the optimal search density problem is now relatively easy to solve (using a computer) and provides useful guidance on how much of the available effort should be invested in each of various regions of the probability density distribution. Since the ideal solution, optimal searcher path, is very difficult (often impossible in practice) to solve, the optimal searcher path problem will not be considered further. Unless otherwise stated, the term *optimal search plan* will refer to solutions of the optimal search density problem.

2.2.3.2 *T-Optimal and Uniformly Optimal Search Plans*

A search plan is said to be “*T*-optimal” if it maximizes the probability of finding the search object by time *T*. A search plan is said to be “uniformly optimal” if it is not only *T*-optimal but is also *t*-optimal for all values of *t* less than *T*. Most of the theorems, techniques and algorithms found in Stone (1989) deal with uniformly optimal search plans. Because there are significant real-world constraints on how sensors can be deployed and moved about, actually attaining uniform optimality during a single search sortie is generally not a realizable goal. However, *T*-optimal search plans are often both attainable and operationally acceptable on a per-sortie basis. Since uniformly optimal search plans are also *T*-optimal for the time at which the available effort becomes exhausted, uniformly optimal search plans are still valid and valuable for planning purposes.

2.3 Search Planning Process

The first step in defining the requirements of a search planning tool is understanding the problem that the tool is supposed to help the planner solve. Like most problems, a thorough understanding requires that it be viewed from several perspectives and levels of detail. One of these perspectives is a “process view.” At a high level of abstraction, the search planning process consists of the following steps:

- 1) Create a Case when Alerted.

- 2) Gather Data and enter it into the Case folder or database (investigation). Revise as needed.
- 3) Enter Assumptions about needed but unavailable data elements into the Case folder or database. Revise as needed.
- 4) Analyze the Data and Assumptions from steps 2) and 3) to Define and Weight Scenarios. Review frequently and Revise as needed.
- 5) Estimate “initial” Probability Density Distribution(s) (PDD) based on Scenario definitions. Re-compute as needed based on revisions in previous steps.
- 6) Estimate PDD for next search (based on “initial” PDDs from Scenarios, probable post-incident State changes, previous Searching, etc.) This may require re-evaluating all activity to date, depending on revisions, if any, to data from previous steps. Often steps 2) – 4) will provide sufficient information the first time and will not be significantly revised in ways that affect the remaining steps as the search progresses. In other cases, proper revision of the “initial conditions” may be crucial. The “Review” process of step 4) is vital and necessary to prevent the condition known as “scenario lock” where the search planner pursues one scenario to the exclusion of others that are also plausible.
- 7) Estimate resource availability and capability (available Effort) for the next search.
- 8) Plan the next Search so that POS is maximized (optimal effort/resource allocation).
- 9) Promulgate the Search Plan.
- 10) Execute the Search Plan.
- 11) Evaluate the completed Search based on actual search activity and search conditions.
- 12) Repeat steps 2) – 11) until all Survivors are found and rescued or until active search is suspended pending further developments.
- 13) Close or “suspend” the Case.

Note that these steps involve all the essential elements described by Stone (1989). Although a “detection function” is not explicitly listed, it is implicitly included in step 8) since POD is required to compute POS, and the relationship between POD and effort is required to determine which allocation of the available effort will produce the highest POS. Even though Hill (1997) and Stoffel (2001) also include “lists” of actions, they are less complete and neither is specifically aimed at maximizing the POS.

3. Selected Land Search Procedures

In the paragraphs that follow, selected procedures, practices, and widely accepted doctrine described in land search planning documents will be reviewed and summarized in terms of their compatibility with scientific search theory. The source of these materials consists of those documents that have been widely referenced by, and available to, the inland SAR community over the past three decades. These documents range from textbooks to privately published papers with only a limited distribution.

3.1 Land Search References to Search Theory

In this section we will describe, in chronological order, those references on which the land search planning guidance contained in the most current widely used sources is based. Brief descriptions of the main features of these references will be included.

3.1.1 May (1973)

May (1973) included a brief chapter on search and suggested that when a line search was, "...not intended to cover every bit of the ground," a "loose line" search technique may be used in order to gain speed and, "...rapidly search only the most likely locations..." (p. 106).

Although search theory was not specifically mentioned, May's (1973) suggestion that five "man-hours" per square kilometer would produce a, "...very low probability of finding small clues" (p. 107) when compared to "saturation" searching that requires 250 man-hours per square kilometer (p. 104) expresses a direct relationship between man-hours (roughly, effort) and probability of detection. This relationship was consistent with the scientific search theory literature even though there is no evidence that the author was familiar with the concepts and the term "search theory" was never mentioned in the book.

3.1.2 Kelley (1973)

The first published mention of search theory to the land SAR community was Kelley (1973). Kelley's book covers a wide range of topics from search operations and strategy to training and base camp management. It also appears to be the first land search publication to mention the phrase—later to become quite popular with the land search community—"search is an emergency" (Kelley, 1973, p. 213).

The data on which Kelley (1973) bases many of the recommendations in his book—and indeed much of the entire contents of the book—are from a survey of 167 case studies from one law enforcement agency in Southern California from 1964-1971 (p. 261).

Kelley (1973) summarizes his perspective of the usefulness of "search theory" in one brief paragraph:

The essence of search theory is that preliminary search activities have the greatest effect on the search outcome. Specifically, initial searcher response time, the ability to confine victim activity, and the early detection of clues to the victim's whereabouts greatly improve the chances of finding a victim (p. 3).

The bibliography of Kelley's (1973) book contains sixty-two entries. Five of them are from the scientific operations research community (Morse & Kimball, 1946; Koopman, 1956a-b, 1957; Morse, 1970). Two of these references are not cited in the text at all (Morse & Kimball, 1946; and Morse, 1970). The other three (Koopman, 1956a-b, 1957) are cited only once in Appendix III. This appendix (by Rex Farquhar and Dennis Kelley) is titled "Probability of Success" (pp. 263-267). In it, they use the vernacular and attempt to describe a few of the concepts of Operations Research, probability theory and search theory in terms more familiar to persons involved in land SAR. They even reference Koopman (1956a-b, 1957) and attempt to define "probability of success," "coverage ratio," "random search," and "searcher rate." With the exception of "probability of success," the definitions offered do not match the definitions found in the scientific search theory literature up to that time or since, as examination of the literature surveyed by Benkoski, et al. (1991) shows. However, the definitions given are all incorrect for the same reason: the concept and definition of "effective sweep" width is completely missing even though this term is integral to the definitions of "searcher rate" and "coverage ratio," both qualitatively and quantitatively.

Although their representation of Koopman's (1956a-b, 1957) "random search" (exponential) detection function is correct in form, the incorrect definition of "coverage ratio" negates this positive aspect. (See discussion in section 3.3.1 below.) Farquhar and Kelley do observe that in the normal situations encountered in land searches, the exponential detection function is a "better model" for estimating POD. However, all these observations are restricted to Appendix III where they are offered only incidentally and do not impact the more intuitive, subjective methods found in the body of the text. In other words, the author mentioned some of the concepts derived from the Operations Research studies, but did not develop them nor suggest how they might be applied in the context of land search operations.

The problem of referencing the science of search theory but neither translating it into everyday operational terms nor using it correctly extends to the contemporary land search literature as well.

3.1.3 Wartes (1974-75)

Jon Wartes was a contemporary of Kelley (1973). At about the same time Kelley was writing his book, Wartes (1974) was conducting a detection experiment over a two-day period in the dense Pacific Northwest temperate rain forests of Washington State. The experiment did not address the central issue of effective sweep width ("detectability") but instead tried to relate POD estimates directly to searcher spacing. He produced a fifty-page report on the experiment that has since been widely referenced and used by the inland SAR community. Since this report has been so influential on the land SAR community, it will be discussed in some detail in section 3.3, *Detection Function*.

In Wartes' (1974) report the author states that, "Over 3400 man-hours of work went into the preparation, administration, and evaluation of the research..." (Wartes, 1974, p. 50). Although over one hundred people are included in the acknowledgements with many of the entries listing their relationship to the experiments, none are listed as having backgrounds or credentials in operations research, search theory, or the design of scientific experiments.

The study that served as the basis for Wartes' (1974) report was fundamentally flawed as a detection experiment for a number of significant reasons (see section 3.3 for more details). The consequence of this was that the development of sweep width values was not done and in fact could not be done from the data that was gathered.

Later, Wartes (1975b) made a case for "non-thorough" search methods. This report was also very influential in the land SAR community and is discussed further in section 3.4, *Effort Allocation*.

3.1.4 Syrotuck (1974-75)

In 1974, William Syrotuck published a paper that used Wartes' (1974) experiment and conclusions as its basis. The author was and is highly regarded for his scientific treatments of several issues important to land search including the behavior of lost persons. Because of this, many of his papers have been widely distributed and universally accepted in the land SAR community.

According to the author, the paper was intended to, "...examine the grid line, it's effectiveness, and to put forth some suggestions to increase efficiency" (Syrotuck, 1974, p. 1). Although the document did include some excellent practical advice for land searches, the author's POD conclusions revolved around the spacing-based POD findings of Wartes' (1974) that are examined in section 3.3, *Detection Functions*.

3.1.5 Bownds et al. (1981, 1991a-c, 1992); and Bownds & Harlan et al. (1991)

In 1991, a small group of SAR practitioners and mathematicians published a paper titled "Searchbusters" in four parts over multiple issues of the land SAR periodical *Response: Journal of the National Association for Search and Rescue* (Bownds et al., 1991a-c, 1992). This document presented several concepts that were widely accepted by the land SAR community. Some of these concepts included:

1. POS has no value after a search
2. POS only has value in prediction
3. "Success" in POS has "misleading overtones"
4. POA of ROW ("Rest of the world") is needed
5. Maximizing ROW POA as an optimization method
6. O'Connor (non-numerical, non-proportional) consensus method
7. POD has various definitions
8. POD is a measure of efficiency of a single resource after a search
9. Numerical influence of clue

Interestingly, the view that POS had no value as a post-search evaluation criterion was first expressed in Bownds et al. (1991c) even though the companion software package, Computer-Aided Search Information Exchange (CASIE), actually uses cumulative POS, which accounts for the post-search POS values of all previous searching, as an optimization criterion (Bownds et al., 1985). The main problem with CASIE is that it does not allocate the available effort over the probability density distribution. Instead, the user must decide ahead of time what resources may be used in each segment and give a (highly subjective) POD value for each segment-resource combination. This limits CASIE to a finite set of possible effort allocations, none of which may be optimal. All CASIE actually does is iterate through all possible effort allocations in this finite set to find the one that gives the highest cumulative POS. However, CASIE does not display the cumulative POS value it computes. Instead, it displays something called ROW POA (discussed in section 3.4.3, “*Rest of the World [ROW]*”).

Bownds et al. (1981) published a report from several detection experiments conducted with aircraft over the Sonoran desert in Arizona. A very similar series of experiments was conducted ten years later in Arizona by many of the same individuals but the detection scenarios were flown over mountainous terrain (Bownds & Harlan et al., 1991). Both of these reports are discussed in section 3.3.2, *References to Land Search Detection*.

3.1.6 LaValla et al. (1997)

There was a great deal of activity in the land SAR community between 1973 and 1975. Kelley (1973), Wartes (1974), and Syrotuck (1974, 1975) all published papers or books based on their understanding of “best practices” in land SAR and search theory. The National Park Service (NPS) budgeted some money to present a five-day course at the Grand Canyon National Park (Autumn, 1974) called “Managing the Search Function” (MSF). A “first cut” of a student manual for the course was compiled in 1975 from, “...articles, professional papers, books and other assorted SAR resources that were available from around the country” (LaValla et al., 1997, p. 2) before an instructor’s manual was eventually developed by Green and LaValla (1978).

At some point in the early 1980’s, several individuals working with the National Association for Search and Rescue (NASAR), a not-for-profit organization, developed a student manual for MSF that was distributed with classes (Brady et al., c1981). Soon thereafter, NASAR and the Emergency Response Institute (ERI), a for-profit partnership, collectively worked to improve the MSF course information and became the collective keepers of the material. At some point in the late 1980’s, NASAR and ERI each took separate paths with the MSF information. NASAR continued to manage, use and distribute it while ERI repackaged the materials and called their derivative book “Search is an Emergency: Managing Search Operations” (MSO) to differentiate it from MSF. In the mid-1990’s, the two principals of ERI dissolved their partnership and established their own separate commercial interests. One now operates “ERI International, Inc.” while the other venture is now called “Emergency Response International, Inc.”

LaValla et al. (1997) evolved from the early MSF and MSO student texts. In its current form, the book is essentially a collection of often dated ideas and concepts from various land search operators, managers and researchers. Most of the current contents of the book can be traced back to

the original MSF course; but, new, occasionally conflicting, ideas, procedures and practices have been added over the years. No specific search planning guidance from the many ideas and concepts in the book is offered and material later incorporated did not resolve conflicts in the differing approaches. Thus, readers are left to determine how to use the various concepts and resolve conflicts on their own. In addition, there is no evidence that any scholarly reviews or rigorous evaluations of the specific concepts and ideas included in the book have been conducted.

3.1.7 Colwell (1992, 1994)

In Colwell (1992), the author described an experiment he and his colleagues conducted in 1990 in four “search zones” covering 0.26 square kilometers in the slightly sloping mature coniferous forest near Vancouver, B.C., Canada. Two hundred life-sized, three-dimensional, cardboard mannequins—16 equipped with AM radios to produce sound for searchers to hear—were placed into the search area in three configurations (sitting, lying, and standing) and in both high- and low-visibility color configurations. With this report, the author intended to, “...explain, in detail, the procedures necessary to gain significant improvement in grid-searching techniques” (Colwell, 1992, p. 1). This was a serious attempt to calibrate field POD values but fell into the same trap as Wartes (1974) and Syrotuck (1974, 1975) by focusing on the inappropriate goal of relating POD to searcher spacing.

Colwell (1992) appears to be the first attempt to study sound as a “detection” method in the land search literature. However, the focus was on relating POD to searcher spacing and was completely devoid of the concepts of sweep width and coverage. “Search theory” is mentioned only in the context of “Grid Search Theory,” which is not a synonym for scientific search theory. POS, POA, sweep width, and coverage as defined by Koopman (1946, 1980) are not mentioned in the document.

The author claims to have “further refined” (Colwell, 1992, p. 2) the method of determining visibility distance that Perkins (1989) called “critical separation.” This refinement involves a measurement technique called the “visibility petal” (p. 68). Although Colwell generally agreed with the POD findings of Perkins, he also came to the conclusion that, “...visibility distance measurements, even when ‘improved’ by the authors refinements, are unreliable, anomalous and prone to wide variations” (p. 36).

Colwell (1992) concluded that pursuing “search efficiency” was a worthwhile goal and that searches with wide spacing between searchers were the most efficient types of sweeps. The author developed a “Recommended Sweep Search Conditions Table” based on his “findings.” In it, specific between-searcher spacing is recommended based on the type of sweep desired (e.g., high visibility, low visibility, standard, or body sweep).

Colwell (1992) defined “efficiency” as a, “...measure of what you get out of a search...for the effort you put in...” (p. 23). Mathematically, the author described search efficiency (he also later termed it “segment priority”) as the percent POD divided by searcher-hours required to achieve that POD, where searcher-hours served as his metric for effort. This description and its emphasis on POD as the desired outcome of a search was indicative of the state of the land search planning paradigm at the time. Unfortunately, this erroneous use of POD has survived and remains in

much of the current land search literature in spite of the fact that Colwell (1994) later revised his concept to incorporate POS (POA x POD) instead of just POD in the numerator of his metric. The fundamental truth that Colwell eventually comprehended—maximizing POS as the objective of all search planning—was not, and still has not been, similarly embraced in the bulk of the current published land search planning guidance.

Colwell (1994) attempted to, “...show how [POA, shifting POA, POD, single-pass POD calibration, POS, probability density, and search efficiency], along with the new concept of Search Priority, can be applied to provide a useful and workable, integrated search planning tool” (p. 1). Although several elements of search theory were discussed in this work, the concepts and application of sweep width and coverage were conspicuously absent. Because of this, the full application of POS also eluded the author causing him to write, “The concept of probability of success has probably been overstated” (p. 2) and, “The problem with POS is that it...does not take into account such real field problems as the size of the search area and the manpower required to actually perform the search” (p. 2). This latter statement is a direct result of not having a detection function that relates the POD in a segment to the effort required to achieve it.

Colwell (1994) suggested that, “...search areas, once defined, should usually be searched in order of the highest probability of area first” (p. 3). However, the author went on to suggest a better method that involved first establishing a probability distribution (called an “area-based search priority”) then multiplying this by the POD in each segment to come up with what the author called a “manpower-based search priority.” This term actually ranked segments in an order that would produce the most POS per unit of time (including the time it took to access and return from the segment)—a reasonable suggestion in terms of search theory. Unfortunately, the POD estimate for this effort allocation method is completely subjective and does not include the use of sweep width or coverage (effort). This ends up being a fatal flaw in the otherwise valid optimization method.

In short, Colwell (1994) was an attempt to develop a comprehensive approach to the application of search theory with only some of the necessary ingredients.

3.1.8 Hill (1997)

Much like LaValla et al. (1997), NASAR’s book, “Managing the Lost Person Incident” or MLPI (Hill, 1997), also evolved from the early MSF materials. In the mid-1990’s, NASAR decided to significantly improve and rename its MSF book. According to the Editor, “The most obvious changes pertain to the application of technology to SAR, search theory, stress management, and research on lost person behavior...Indeed 75% of the text is new” (Hill, 1997, p. ii).

Unlike LaValla et al. (1997), Hill (1997) did not include a comprehensive description of all land search ideas and concepts, but many remnants remained. The author also chose to provide more guidance than LaValla et al., but much of it was still based on the earlier work of Kelley (1973), Wartes (1974), and Syrotuck (1974, 1975). Nevertheless, a “logical sequence of planning actions” for search was provided—the closest thing to a search planning methodology published in the land search literature up to that point. This sequence is discussed in some detail in section 3.5, *Land Search Planning “Methodology.”*

3.1.9 Stoffel (2001)

The principal author of Stoffel (2001) is a former partner in ERI. Thus, the basis for much of the information in this book also comes from the original MSF and MSO material. However, this more contemporary book attempts to merge selected vernacular and principles of scientific search theory with the historic land search methods and thinking. The result is confusing at times. For example, the author briefly defines probability density (Pden) and describes its use as, "...[when] all other factors are equal, search parties assigned to segments with the greater Pden would likely produce results more rapidly" (p. 156). He then goes on to recommend that segments be searched in order of their POA values (high to low) when he states, "In effect, this [consensus] process ranks the segments in the order of priority that each should be searched" (p. 163). As another example, although the author correctly defines "sweep width" (p. 173), "track line length," "area effectively swept" (p. 174), and "coverage," and even includes a POD versus Coverage (exponential or random search) curve, these definitions are quoted from Cooper & Frost (1999) and go completely unused in the remainder of the document.

Although not as complete and easy to follow as Hill's (1997) approach, Stoffel (2001) also provides a "logical sequence for planning a search effort" (p. 191). This methodology is discussed in more detail in section 3.5, *Land Search Planning "Methodology."*

3.1.10 Dougher (2001) & Dougher et al. (2001)

Although the layout and presentations differ, the authors, contributors and content are essentially the same in both Dougher et al. (2001) and Dougher (2001). So, both of these publications will be discussed together. Both of the documents provide a list of recommendations (called "recommended actions" in one, and "action checklist" in the other) but neither really describes a search planning methodology. Rather, they both offer numerous practical recommendations for incident organization and planning, but little specific information regarding search planning.

Both documents employ a "six-step process" to solving operational problems as the centerpiece and basis of their recommended actions. The same six-step process is included in Stoffel (2001, p. 99) and termed "Recommended management steps," but its use is not developed or described in any detail.

Neither Dougher et al. (2001) nor Dougher (2001) describe area searching at all. Both limit their operational descriptions to "hasty searches" (quick searches of trails, roads, and likely areas without specific boundaries), "passive" search techniques ("...pure observation, leaving the target and its environment unaltered..." [Koopman, 1980, p. 16]) and practical advice on how to organize, procure, manage, and employ land search resources. However, both mention and/or define selected terms and concepts related to search theory. They just do not develop or describe their use.

In two brief pages in the appendix, Dougher et al. (2001) outline selected elements of search theory—interestingly titled "An Introduction to Probability Theory" (p. 34). These authors defined

POS as, “the Probability of Success (projected) of finding the subject or clues, using specific POA and POD values (a measurement of a future action)” (p. 34). They go on to describe POS as “having no value” after a search because, “After a search has taken place, the subject was either not found (success = 0%) or was found (success = 100%)” (p. 35). This of course ignores the considerable value of POS as an evaluative tool but is indicative of the land search planning paradigm.

Dougher (2001) includes POA and POD in the glossary, but does not use them in search planning. The same author does not even mention POS, but describes “POD targets” (p. 7-2) and POD as a value that resources estimate directly—issues that are addressed later in this review.

3.2 Probability Density Distributions

3.2.1 Probability and Segments

The development of a probability density distribution requires some means of estimating probable search object locations. In a land search, a committee of search planners accomplishes this through the development of a consensus. The details of this process are described in section 3.2.2, *Methods of Developing Initial POA/POC Values*. The process requires that the search area be subdivided into regions and that a POA value be assigned to each region. Although a definition of “region” (see Appendix A for definition) compatible with Koopman’s (1946) description of “region of probability” was published by Cooper & Frost (1999a), the use of the concept is not currently reflected in the land search literature. Rather, the subdivision of the search area to which a probability is assigned is the “segment”—an area sized to fit the resource assigned to search it. Thus, according to the land search literature the boundaries of the segments of the search area to which probabilities are assigned are not based on probability, they are based on searchability.

It may be useful to explain here what the land search literature describes as “segmenting the search area” and how they assign POA values. For management and logistics purposes, land search managers find it necessary to divide the search area into some number of searchable segments. Guidelines for this process have been published and appear consistent throughout the land search literature. As an example, some rules for segmentation include the following (Hill, 1997, pp. 113-114):

- Segment boundaries should be easy to identify, both in the field and on a map.
- Interior barriers within segments should be avoided.
- A search team should be able to cover its assigned segment in 4-6 hours of actual searching and be able to complete the entire sortie, including transit time to and from the assigned segment, in one “operational period” of about 8 hours.

In terms of search theory, the important point is that most often segmentation is not driven by considerations of where the search object is more or less likely to be or where search effort should be placed. Yet, segments are the smallest units described in the land search literature to which POA is assigned. Segmentation is most often driven by logistical and operational constraints unrelated to the probability density distribution on search object location. Note that the

level of coverage to be attained in the 4-6 hour interval is not specified, nor is the size of the search team. This is just one example of the vexing lack of specificity characteristic of critical inland search planning procedures where optimal effort allocation requires quantifiable estimates and measures.

A fundamental problem with the present land search practice of defining segments prior to developing a probability map—described by Hill (1997) and others—is that two completely independent aspects of the search planning process are being intermingled in a way that virtually forces them to influence one another when they should not. Where the search object is more likely or less likely to be located has absolutely nothing to do with the search manager's general management and logistics problems, and vice versa.

3.2.2 Methods of Developing Initial POA/POC Values

When vessels or aircraft become distressed and call for help giving their position, a probability density distribution on their location can be developed by using accepted models of the distribution of position errors associated with the method and mode of navigation used by the distressed craft. Even in the event of an overdue or unreported craft, the probability density distribution about the last known or reported position can be estimated in the same way, and the subsequent movement of the craft and the distribution of possible distress positions and times can be estimated with simple motion models. Post-distress motion is not an issue for downed aircraft on land, at least not until the forced landing site is found and there is evidence that survivors moved away from that location. On the ocean, post-distress motion is a function of winds and currents beyond the survivor's control.

However, the situation is quite different in the case of land search for lost or missing persons. Such persons rarely have the means to call for help. They often do not realize they are lost until a substantial time after making the navigational error that caused the condition. Post-distress motion is largely a function of the lost person's behavior and the environmental conditions (terrain, vegetation, weather, etc.) in the immediate vicinity of their location. Simple position error and motion models like those used in aeronautical and maritime cases do not apply. Hence land search planners have no choice but to find some other way for developing a probability density distribution on which to base search plans.

3.2.2.1 *Mattson*

The earliest written description in the land SAR community of a method to establish initial POA values for lost or missing persons was probably Mattson (1976) in an article for *Search and Rescue Magazine*. In this article, the author described a consensus method for establishing POA values from which to plan a search. This method, described numerous times in virtually every subsequent land SAR text, involves a group that assigns POA values to segments via a consensus-building process. Each of several individuals independently reviews the available information and evidence specifically related to the case at hand along with other general information or statistics compiled from similar incidents. These individuals assign POA values in percentages to the segments and then reconvene to compare notes and arrive at a consensus. The total of each

evaluator's percentages has to sum to 100%, and the mean value across evaluators for each segment becomes the initial POA value for that segment.

This method came to be known in the land SAR community as the "Mattson Method." It was easiest to use when the number of segments was small (especially in the days before computing equipment was readily available), and it allowed evaluators to maintain proportionality from one segment to another. That is, if an evaluator entered 40% under segment A1 and 20% under segment A2, it was obvious that the evaluator meant that there was twice as much probability that the subject was contained in segment A1 as there was in segment A2. At the time, proportionality was not highlighted as a benefit of the method. It was only later when other authors attempted to improve on the method that the importance of this characteristic was recognized. See Appendix D for examples and a more detailed description of the importance of using a proportional assessment technique.

The simplicity and field expediency of the Mattson Method quickly caused it to gain popularity. However, some difficulties in its use were eventually identified. One of these was the perceived difficulty of assigning POA values to a potentially large number of segments that were in the proper proportion to one another and summed to 100%. The possibility of "remainder bias" was raised. The hypothesis was that evaluators would tend to assign the few highest POA values first, and then be faced with splitting the small remaining probability among a large number of segments. It was feared that the first few POA values assigned would be overrated while the remainder would be underrated even though, in theory, the method should have preserved the proportionalities among the segments if properly done.

3.2.2.2 *Relative Ranking with Letters*

In response to these perceived problems, Bownds et al. (1991a) suggested an alternative that used letters (in lieu of numerals) as values on a relative scale, such as one ranging from "very unlikely" down to "very likely." Choices are converted to numerical values; usually consecutive integers covering the range of choices (e.g. a scale of 1 to 9) and several "rules" for use of the system were suggested by the authors.

Although the land SAR community readily accepted this method at the time, the method fell short of allowing evaluators to maintain proportionality between segments. In short, there was no way for an evaluator to determine the proportion of "very likely" to "very unlikely" or "A" to "D." In addition, the system permitted the range and scale on which likelihood was measured to vary in size (e.g., A-D or A-I) depending on the subjective ideas of the evaluator (Bownds et al., 1991a). When using this system, it is very likely that the evaluators using only letters and corresponding phrases are completely unaware, in any quantitative sense, of just how much less likely an E is than an A or how much more likely a D than a G. This fatal flaw in the concept may have been due to the unfortunate fact that the authors were unfamiliar with the science of search theory (Dr. David Lovelock, personal communication, 3 August 1998). Regardless of the reasons, the method suggested by Bownds et al. (1991a) was intended to improve on the shortcomings of the Mattson Method and ultimately fell short because it did not meet an essential criteria of probability density distributions: proportionality (Wagner, 1989).

Probability densities (Pdens) are computed from the assigned POA values and estimates of segment areas. Determining the area of an irregularly shaped segment from a map is not an easy

task, especially in steep terrain. For this reason, computation of probability densities is often omitted in practice, leaving only POA values on which to base effort allocation decisions. Nevertheless, once all of the segment POAs are assigned a probability density distribution has been defined. However, there are several reasons to question the validity of a distribution constructed in this fashion.

The consensus process is obviously very subjective and therefore prone to all the usual pitfalls of subjective analysis. This cannot be helped. However, it is also ambiguous in at least one important respect. This ambiguity is most obvious when segments are rated on a relative scale (e.g. very likely down to very unlikely) regarding the chances of the subject being contained within their boundaries. For example, consider two segments of unequal sizes. If both are assessed as “very likely” to contain the subject, does this mean their probability densities are equal or that their POAs are equal? The answer to this question will clearly affect how the effort should be allocated to obtain an optimal search plan. Current procedures in land search are strictly POA-based and just let the probability densities fall where they may. That is, all “very likely” segments will be assigned the same initial POA value regardless of the segment’s size in terms of area.

Although the objective of the consensus process is clearly the production of a probability density distribution according to the consensus of the evaluation committee, there is a dearth of guidance regarding how POA values (absolute or relative) should be assigned to segments. This in turn indicates the members of the consensus committee are probably unaware of their true mission. It is at least possible, and probably likely, that the intuitive processes of some individuals lean toward probability density assessments while those of others actually do lean toward POA assessments, even though all report their opinions in terms of POA or relative values on some scale. In any case, POAs are almost certainly assigned in the absence of any conscious, explicit awareness of how those assignments will affect the resulting probability density distribution or the allocation of search effort. This in turn can lead to a number of anomalies.

For example, it could easily happen that a region rated as only “likely” would have a higher computed probability density than one rated as “very likely” if the former was sufficiently smaller than the latter. All other things being equal, an optimal effort allocation algorithm would place effort in the more “likely” segment ahead of the “very likely” segment. This would probably surprise and concern most search managers.

As another example, suppose a test was conducted where the consensus group was asked to divide the search area into probability “regions” using only criteria and information bearing on where the subject was more and less likely to be located. They would be instructed to ignore all of the usual segmentation criteria. Also included in these instructions would be a directive to subdivide the search area into as many, *but only as many*, regions as they could justify from the available data, assigning a POA to each. (Ideally, we would want probability density values assigned, but this is probably asking too much.) Suppose further that the result was a relatively small number of regions that did not necessarily meet the segmentation criteria (e.g., too large or too small). These regions would have to be further subdivided, or gathered, into segments for the usual search management and logistics reasons. However, the question of how to assign the segment POA values remains unanswered. Two methods suggest themselves. In the first, a region’s POA is apportioned among its segments according to their respective areas so the probability density remains constant everywhere within the region. This is in keeping with the notion

that if the committee could have made a justifiable distinction about whether the probability density was higher in one place than another, they would have done so by creating another region rather than release the given region for segmentation. The other method is to go through the consensus process again, with the constraint that the sum of the segment POAs (where the segments are subsets of the region) in any region equal the previously assigned regional POA. This is a contradiction of the original directive to assign POAs only where justified by the available data. We would now be asking the group to assign POAs at a level of detail well beyond what the available data would support. Besides, the averaging process used to resolve differences in opinion could (and actually should) produce a nearly uniform density across the segments in a given region.

It may be possible to modify the Bownds et al. (1991a-c) method of establishing initial POA values to be compatible with scientific search theory. But, the likelihood range would have to be fixed (always using the same number of letters) and modified to allow each choice to be directly translated into a quantitative equivalent.

3.2.2.3 *Relative Ranking with Numerals*

Colwell (1996, 1998) modified the method described by Bownds et al. (1991a) by using integers (1-9) in lieu of letters and qualitative terms. This allows, but does not require, evaluators to maintain proper proportional relationships among the segments. The easiest way to ensure the segment POA values are in the relative proportions that reflect the evaluators' views is to require that all evaluators quantitatively *rate* or "score" (rather than rank) all segments against the same standard. One way to do this was suggested by Wagner (1989) where probability maps were produced that showed each cell's (segment's) "score" relative to the most probable cell (segment) on a scale of 1 to 10. Thus a "9" meant the cell (segment) was 90% as likely to contain the search object as the most probable cell (segment). Once all segments are "scored" in this fashion, the results may be normalized into a set of corresponding POA values that reflect these proportions and sum to 100%. This may be all that is necessary to introduce proportionality into this method, but additional modifications to the overall approach would also likely be required.

Dougher (2001) later described the same modification to the numeric-based consensus method first suggested by Colwell (1996, 1998). In it, a qualitative scale ranging from "very likely" to "very unlikely," exactly like the system suggested by Bownds et al. (1991a) was used. However, the consensus form provided included numerals (1 through 9) that were associated with each of the qualitative values on the "Probability Estimate Scale" (Dougher, 2001, p. 24-7). Although the author likely did not intend this, and it is not described in the instructions, the numerical values associated with each of the qualitative terms would allow an evaluator involved in a consensus to assess proportionality when comparing one region's value to another (e.g., 3 to 6). Unfortunately, the author does not include an explanation of the significance of proportionality and so evaluators have no way of knowing about its importance. The described system also suffers from some of the same shortcomings seen in other non-proportional methods of establishing initial POA values. For example, initial POA values are assigned to searchable segments or divisions (division: "a portion of the search area designated as being under the control of a supervisor," Dougher, 2001, p. 24-1) and not regions of probability (see next paragraph). That is, how to search is considered before establishing a probability density distribution. Interestingly, "region" is specifically defined in Dougher (2001) as, "A portion of the search area established to facilitate POA determinations" (p. 24-1), but the term is not used anywhere in the document outside of

its definition. This seems to imply that the author allows its use but does not recommend it, not unlike Stoffel (2001)(see section 3.5, *Land Search Planning “Methodology”*).

Most land search methods for establishing initial POA values mix “probability” with “searchability” when they attempt to use segments—areas sized to be searchable and based on the resource expected to search it—instead of regions of probability to develop a probability map. Although there is no logical connection between where a subject is likely to be and how a segment would be searched, the land search literature universally advocates establishing searchable segments before developing an initial distribution of probability. This is undoubtedly based in the land SAR community’s historical focus on operational issues and limited application of any search planning methodology. Fortunately, the solution may be as simple as separating the development of a probability map from operations (searching) when planning a search. It is very likely that this problem can be overcome by simply adopting a search planning methodology that isolates and addresses the development of a probability map prior to determining how searching is to be carried out.

Prior to developing CASP for the Coast Guard, D.H. Wagner had assisted the Navy in several deep ocean searches for sunken objects. Somewhat like a land search planning consensus group, the Wagner team carefully analyzed all the available data with a view toward developing a probability density distribution on which effort allocation could be based. However, there were several important differences. The Wagner team had orders of magnitude more time to do their analysis (but probably also had orders of magnitude more data to sift through). They were intimately familiar with search theory and were comfortable working with probability density distributions. Like the inland search planners, they determined a search area and then divided it into pieces of convenient size. However, unlike the inland search planners, these pieces were cells of equal size arranged in a regular grid (a luxury not generally available for practical inland search planning). Assigning a relatively high value to a cell meant the cell had both a relatively high POA and a relatively high probability density. There was no ambiguity. The Wagner team members were keenly aware of this fact and of the distinction between the values. As a result, their methods did not suffer from the anomalies just discussed. They knew that solving the optimal search problem required probability density distributions that faithfully represented the available data. They also had the tools and capabilities to produce such distributions without outside assistance. They were then able to apply optimal effort allocation algorithms to great advantage.

The searches where Wagner supported the Navy involved very costly search platforms and techniques. The more well-known search objects were items of great national importance (two sunken nuclear submarines (Scorpion and Thresher) and an accidentally dropped (unarmed, fortunately) hydrogen bomb). These characteristics are hardly typical of SAR searches and we can scarcely expect to have such high-powered teams of mathematicians and search theorists routinely available for SAR responses. Besides, they would be as much out of their element in trying to deal with the practical exigencies of SAR search management as most search managers would be in the world of mathematical theorems and proofs. The point to this observation is that each group, plus an important third group, has a vital role to play in the development, implementation and use of sound search planning practices. Theorists are needed to develop provably correct techniques, formulae, algorithms, etc. But before these can do the search managers any good, they must be translated into practical procedures that also account for the major, at least, operational constraints imposed by the real world. This is the job of the third group. They are

the ones who are able translate the theorems and algorithms into operationally useful techniques, bridging the gap between theory and practice (and sometimes literally acting as translators between the theorists and practitioners). Finally, no matter how well developed the theories or how faithful the translation from theory to practice, many, in fact most, important decisions must still rely on the judgment of experienced search managers. A methodology derived from search theory can be no more than one more tool in the search manager's toolkit, even if it is a quite valuable one.

3.2.2.4 *Ranking by Descending "Priority"*

Perkins and Roberts (1994) described a completely subjective method of prioritizing search segments (sectors) known as the "Sector Ladder." The Sector Ladder method falls short of producing a probability density distribution on search object location because it is merely a simple ranking of the segments by "relative priority." The authors claim that such a ranking is no different from one based on POA, "...except that PoA does it by means of numbers and calculations, whereas the Sector Ladder does it by means of writing down the sectors in the form of a list" (Perkins & Roberts, 1994, p. 6). There is no expectation that POA, POD, and POS can be used quantitatively to track search progress and to aid in the allocation of search resources. Further, it actually suggests that quantification should be avoided because, "...Search Managers...cannot do the maths needed to calculate POA's" (Perkins & Roberts, 1994, p. 3).

Although this qualitative and completely subjective method of prioritizing segments may have some limited operational use in representing where the search planner intuitively feels searching would be most productive, it does not produce a probability density distribution on which objective effort allocation decisions can be made. If used in the prescribed fashion, the Sector Ladder will almost always provide sub-optimal results and is not currently compatible with the tenets of scientific search theory. Further analysis is provided in section 3.4, *Effort Allocation*.

Koopman (1980), in his Preface, recognized the potential for, "...inappropriate handling of the mathematics" (p. 2) not unlike what Perkins and Roberts have attempted. Although, "...there is often an impulse to leave it out as such," Koopman continued, "...to leave out the mathematics is to leave out the essential reasoning" (Koopman, 1980, page 3).

3.2.3 Normalizing Adjusted POA Values

The benefits of the use of POS as a measure of search effectiveness are available to the search planner only after POA is "adjusted" for each segment searched. Thus, adjusting POA is an essential step in the use of POS and probability maps as search planning tools. What follows is a brief review of how adjusting POA has been previously described in the inland SAR literature. But, before proceeding, it must be noted that the same literature that seems to be overly concerned with adjusting POA—the only purpose of which is to aid in making effort allocation decision that ultimately require one to compute POS—does not describe what to do with POA after it is adjusted. Worse, the land search literature is almost completely void of any references or use of POS once POA is adjusted.

3.2.3.1 Methodological History: Syrotuck, 1975

Syrotuck (1975) stated that, “When one area is searched with a certain efficiency [POD] and the victim is not found, this increases the [relative] probability of the victim being in another area” (p. 9). This statement is based on the well-known rules of inference first articulated by Bayes. Put simply, when a particular segment is searched but the search object is not found, two things happen:

1. The relative probability of the object being somewhere else increases, and
2. The relative probability of the object being in the segment just searched decreases.

Mathematically, both of these requirements are satisfied if the POA value of the segment just searched is reduced by an amount proportional to the POD applied to that segment. Syrotuck (1975) called this concept “shifting” POA. Other land search authors have called it “adjusting” or “updating” POA. In this context, these terms are synonymous.

When Syrotuck (1975) first attempted to quantify the concept of shifting POA for the inland community, he suggested that POA be adjusted through an application of Bayes’ Theorem in its original form (equation [3-1]).

$$[3-1] \quad \text{Shifted POA} = \frac{Pa \times Pm}{Pa \times Pm + Pn}$$

- Where:
- Shifted POA* is the modified POA of a segment after an unsuccessful search in that segment.
 - Pa* is the probability that the victim is there.
 - Pm* is the probability that the victim was missed.
 - Pn* is the probability that the victim was not there.

The “shifted” POA produced by equation [3-1] is based on the assumption that only one segment was searched and that all the remaining segment POA values will be “shifted” in the following way: After the “shifted” POA value for a single searched segment is computed, the POA values of the remaining segments are normalized so the sum of all shifted POAs returns to the original value of 100%. The shifted POA computed by equation [3-1] is the correctly normalized value assuming only that segment was searched. In short, Syrotuck (1975) first subtracted the shifted POA value of the searched segment under consideration from 100% (1.0). The author then distributed this amount of probability over the remaining segments in proportion to their “unshifted” POA values to get new shifted POA values. In the end, the author’s method of normalizing the remaining segments to properly match up with the “shift” in the POA of a searched segment proved to be cumbersome and removed all possibility of applying POS in any useful way by manual calculation. While the premise of the author’s suggestion was valid, its use, and the normalization process specifically, imposed a considerable computational burden on the search planner and removed the use of cumulative POS—the primary purpose of adjusting POA in the first place.

3.2.3.2 Methodological History: Bownds, c1982

Shortly after Syrotuck's (1975) method was published, Bownds (c1982) advanced an alternative, and mathematically equivalent, method for manually adjusting POA values. This method involved a simpler two-step approach to adjusting POA values but still did not address POS and its use.

In the first step of Bownds' (c1982) modified approach, the adjusted POA value of the segment just searched (say, segment j) was calculated using equation [3-2].

$$[3-2] \quad P(A_j | D_j') = \frac{(1-D_j) \times A_j}{1-D_j A_j} \quad \{\text{Segment searched}\}$$

Where: $P(A_j | D_j')$ is the probability that the subject is in segment j (shifted POA) given that segment j was searched without detecting the subject.
 A_j is the POA of segment j prior to this update.
 D_j is the POD for this search in segment j .

Like Syrotuck's (1975) technique, equation [3-2] computes the correctly normalized value assuming only that segment was searched. In Bownds' (c1982) second step, the adjusted, normalized, POA values for each of the remaining segments is calculated by using equation [3-3] for all $i \neq j$. As the POA for each searched segment is adjusted, this calculation is performed once for every other segment, including those already adjusted.

$$[3-3] \quad P(A_i | D_j') = \frac{A_i}{1-D_j A_j} \quad \{\text{Segment not searched}\}$$

Where: $P(A_i | D_j')$ is the probability that the subject is in segment i ($i \neq j$), given that segment j was searched without results.
 A_i is the POA of segment i prior to this update.
 A_j is the POA of segment j prior to this update.
 D_j is the POD for this search in segment j .

This method ensured the sum of all POA values was equal to 100% at the conclusion of each computational cycle just as Syrotuck's (1975) technique did. In fact, both methods *require* the sum of the POA values be 100% at the end of each computational cycle, just as both methods require as many computational cycles as there are searched segments for a complete update of the POA values.

Although Bownds' (c1982) two-step method appeared considerably simpler than the style of Bayesian update suggested by Syrotuck (1975), it was only a modest improvement in terms of the computational burden imposed on the search planner. Bownds' (c1982) technique also required several pages of calculations (or a computer) for even relatively simple cases. In fact, both authors presented examples of shifting the POA values of a search area having only four segments, all of which had been searched, and each author required two full pages of calculations to obtain the final shifted POA values.

Neither Syrotuck (1975) nor Bownds (c1982) specified whether their terms for POA (i.e., Pa and Ai, respectively) meant shifted POA from previous searches, or whether their terms for POD (i.e., Pm and Di, respectively) represented cumulative POD. If both shifted POA from previous searches and cumulative POD were used in the calculations, the new shifted POA values would erroneously account for the search results twice—once in adjusting the POA and once in accumulating POD. Neither Syrotuck’s (1975) nor Bownds’ (c1982) notation and/or description specifically excluded the possibility of this type of error.

3.2.3.3 Methodological History: Shea, 1988

Shea (1988) published another method for adjusting POA that also required two steps and normalization. However, the author implied that the calculation of a cumulative POD was required, not optional as in earlier approaches. Further, he suggested that the use of cumulative POD would allow one to, “...wait until multiple resources have searched an area before calculating the shifted POA” (Shea, 1988, p. 24). The method allowed the adjustment of POA values after the application of multiple resources. This was a capability not available in previous methods. Indeed, the method could also adjust the POA values for any number of searched segments simultaneously before normalizing. In contrast, the methods of both Syrotuck (1975) and Bownds (c1982) required the adjustment and normalization of the POA value of a single searched segment, followed by the normalization of all other segments, before accounting for additional searched segments in their computations. These older methods took an extraordinary amount of calculation and erroneously implied that normalization was required. In the end, Shea’s (1988) method achieved the same mathematical result as Bownds (c1982) and Syrotuck (1975), but required much less computation. Unfortunately, Shea (1988) also did not address the use, value or application of POS in spite of the fact that, in terms of the application of search theory, this was the whole reason for adjusting POA.

$$[3-4] \quad POA^* = (1 - POD_{cum}) \times POA_{old} \quad \{ \text{Step 1} \}$$

Where: POA^* is an interim term needed in step 2 (equation [3-5])
 POA_{old} is the POA of the segment prior to a search being conducted.
 POD_{cum} is the cumulative POD for the segment.

$$[3-5] \quad POA_{shifted} = \frac{POA^*}{S} \quad \{ \text{Step 2} \}$$

Where: $POA_{shifted}$ is the new POA for each segment.
 POA^* is the figure determined from step 1 (equation [3-4]).
 S is the normalization factor determined from the sum of all POA^* .

One thing Shea (1988) did not point out was that the POD_{cum} used in equation [3-4] had to be based on only those searches conducted since the last “shifted” POA value was computed for the given segment. If the usual definition of POD_{cum} (cumulative POD for all searching done to date in the segment) is used with a POA_{old} value from an earlier adjustment, the effects of searching done prior to that adjustment will be counted multiple times and errors will result. (Changing POA_{old} to $POA_{initial}$, will allow equation [3-4] to work correctly with the usual definition of

POD_{cum} . Also, changing POD_{cum} to POD would clarify how the formula works when only one search of the segment(s) has been done since the last POA adjustment.)

Shea (1988) did point out that when using his first equation, the “POA*” value of any segment not searched would be the same as its “POA_{old}.” Therefore, performing his step 1 calculation on unsearched segments could be skipped. Further, as justification for his second step (equation [3-5]) the Author stated that, “We know that the sum of all POA must be equal to one, so we divide each POA* by S, giving the [normalized] shifted POA” (Shea, 1988, p. 24). This also implied that normalization was a requirement of the method even though this requirement was based on the false assumption that, “...the sum of all POA must be equal to one” (p. 24), and ignored the possibility of using a defective distribution (a distribution that does not sum to 100%) (Stone, 1989) or POS in any way. The reasons normalization is not required, and may be detrimental, are discussed in section 3.2.3.4, *The Normalization Issue*, below.

Shea’s (1988) step 1 was a simple algebraic variation of the conventional notation of search probability ($POS = POA \times POD$). This was a delightfully simple application of the notation, but the author, like Syrotuck (1975) and Bownds (c1982), did not completely define his terms. This caused confusion, the potential for miscalculations, and caused doubt about what was otherwise a valid method.

3.2.3.4 *The Normalization Issue*

What Syrotuck (1975), Bownds (c1982) and Shea (1988) did not explicitly state was that the process of normalization did not change the relative proportions of the figures derived by removing the denominator from the first step in the first two of these methods and letting Shea’s (1988) first step stand unaltered. Normalization may present the figures in a more visually acceptable form (i.e., they add up to 100%), but it is not required if one is only interested in adjusting POA (and computing cumulative POS) after searching. The additional computational burden of normalizing the figures takes time, is fraught with potential for errors, is unnecessary, actually destroys valuable information about the search, and removes the usefulness of POS—a primary purpose of the use of POA in the first place. Regarding the shifting of POA, Hill (1997) quite accurately stated, “...trying to compute the shifted POA by hand can be an extremely onerous and error-prone task...” (p. 145). These are powerful arguments against normalizing yet all of these authors seemed to believe that it was necessary.

3.2.3.5 *An Improved Method*

The U.S. Coast Guard’s Computer Assisted Search Planning (CASP) software uses a very simple algorithm to adjust POA after searches have been conducted. This method was published to the land search community by Cooper (2000, p. 21) in the following form:

[3-6]
$$POA_{s,n} = POA_{s,n-1} \times (1 - POD_{s,n})$$

Where:

$POA_{s,n}$ is the adjusted POA value in segment s (based on the initial POA value established in consensus c for region or segment s) after all searches (hereafter, c shall be presumed to be 1 unless otherwise specified). This value accounts for all searching done to date in segment s .

$POA_{s,n-1}$ is the adjusted POA value in segment s (based on the initial POA value established in consensus c for segment s) for the specific search just before n (e.g., n minus 1). This value accounts for all searching done in segment s prior to search n .

$POD_{s,n}$ is the probability of detection for search n in segment s . This value is not a cumulative value and indicates the POD for search n only.

So that the same confusion that surrounded Shea's (1988) notation does not resurface, Cooper (2000) more precisely defined his notation and terms. In addition, this method accounts for searches of a segment mathematically (e.g., adjusted POA and POS computed) one-at-a-time as POD values become available rather than requiring the computation of cumulative POD—another improvement on Shea's method. Although this method requires fewer computations by several magnitudes, and is by any measure far less complex than previously published methods, there is little evidence that the land search community has yet apprehended or applied the benefits of this improved approach.

In the USCG's CASP software, normalized versions of the non-normalized adjusted POA values are computed within the display modules for presentation purposes using the same normalization technique Shea (1988) used in equation [3-5]. However, the CASP software actually uses non-normalized POA values for all internal calculations and never uses normalized values for anything other than the visual presentation of probability maps. The accuracy of the calculations and the proportionality of the POA values are not affected by the use of non-normalized data. And, if values are normalized, the usefulness and potential of POS are removed. If the inland community used a similar approach, the computational burden of applying search theory to inland searches would be reduced substantially and the full value and usefulness of POS could finally be realized.

3.2.4 Conclusions

Only those methods of estimating POA values that are based only on subject behavior, terrain, weather and other factors that might affect where the subject might be located, and which also require evaluators to consciously assign POA values in a proportional manner should be used. Other methods that do not meet these requirements should be discarded.

POA adjustments to account for unsuccessful searching should be done in an un-normalized manner. This vastly reduces the computational burden while providing a clear cue to the search planner regarding how much of the information on which the search is based has been exhausted by unsuccessful searching. As POA values become very small, it becomes clear that continued

searching based on the same scenario(s) is unlikely to be successful. Either the search planning team needs to re-think the problem or consideration needs to be given to suspending active searching.

3.3 Detection Functions

3.3.1 Koopman's Approach to Detection Functions

Koopman (1946, 1980) provided the definitive analysis of the general detection problem. He first developed the concept of effective sweep width and defined it as follows (pp. 65-66):

Suppose a sensor moves at speed v miles per hour through a uniform (random) distribution of identical objects with an average density of N objects per square mile in a reasonably homogeneous environment and detects on average n objects per hour. Then the effective search (or sweep) width, W , for that combination of sensor, search object and environment is given by:

$$W = \frac{n}{vN}$$

Performing a “units analysis” we have:

$$W = \frac{\# \text{ objects found / hr}}{(mi / hr)(\# \text{ objects / mi}^2)} = mi$$

We see that effective sweep width has units of length (miles in this case). The effective sweep width is a measure of how detectable a given object is by a given sensor operating in a given environment. It is interesting to note that while Wartes' (1974) approach involved some of the same variables, the quantities “thoroughness” and “efficiency” were not meaningful. A similar “units analysis” shows that “thoroughness” is a vaguely defined POD of some sort, and “efficiency” has units of $\text{POD} \times \text{mi}^2/\text{man-hour}$. Wartes' approach is discussed further in section 3.3.2.2 below.

Using Koopman's (1980, p. 66) definition of W , the effective search (or sweep) *rate* is given by vW and has units of square miles per hour, for example.

Koopman's (1980, p. 74) description of the “length of the observer's path” is defined as *effort*, z , (L in Koopman's original notation) or the distance traveled in the search area while searching. If the average search speed, v , and time, t , spent searching are known, then the effort, z , is given by the familiar formula:

$$z = vt$$

If the average search speed, v , is known, then effort may be expressed in time-based units, such as flight-hours or searcher-hours as long as v is left in the equation as a constant. Using these quantities, the area effectively swept, Z , may be computed as follows:

$$Z = vWt \quad \text{or} \\ Z = Wz$$

Koopman (1980, p. 140) then reasoned that the probability of detecting an object, given that it was in the searched area, was dependent on the *effort density*, or the distance the sensor travels per unit of area covered. If effort is expressed in time-based units this becomes the amount of time spent searching per unit of area covered. In either case, the effort density can be normalized with respect to the size of the area covered by defining the *coverage*, C , as the ratio of the area effectively swept, Z , to the size of the area, A , over which the effort was more or less uniformly spread. That is,

$$C = \frac{Z}{A} = \frac{Wz}{A} = \frac{vWt}{A}$$

Koopman then derived the relationship between coverage and probability of detection for several situations. If the effort is expended in an area using a large number of uniformly randomly placed sensor tracks within the area that are short in relation to the dimensions of the area but reasonably long in relation to the maximum range at which the object could be reliably detected, Koopman (1980, p. 72) showed that:

$$POD = 1 - e^{-C} = 1 - e^{-\frac{vWt}{A}}$$

This is known as the exponential or “random” detection function. It is frequently used both operationally in the field and for theoretical work.

For an idealized definite range detector (one that detects every object within some definite range of the sensor and detects no objects beyond that definite range) that completely covers an area with a set of perfectly straight, parallel, equally spaced tracks, the POD is given by:

$$POD = C = \frac{vWt}{A}, \quad C \leq 1.0 \\ POD = 1.0, \quad C > 1.0$$

Finally, Koopman (1980, p. 77) developed a hypothetical model of visual detection based on patrol aircraft flying over the ocean and searching for (enemy) vessels underway. Based on the geometry of sighting opportunities, Koopman postulated that the instantaneous or one-glimpse probability of detection was inversely proportional to the cube of the range from the sensor to the search object. Thus it became known as the “inverse cube” model of visual detection. When used to cover an area with a set of perfectly straight, parallel, equally spaced tracks, this sensor produces a POD given by:

$$POD = \operatorname{erf}\left(\frac{\sqrt{\pi}}{2} C\right) = \operatorname{erf}\left(\frac{vWt\sqrt{\pi}}{2A}\right)$$

where “erf” is the “error function” from statistics. The POD vs. Coverage curves for these three detection functions are shown in Koopman’s Figure 3-12. Figure 3-1 below shows the graphs of these three detection functions. Note that the upper two curves require that the segment be completely covered with searcher tracks that are perfectly straight, parallel, and equally spaced. When these conditions are not met, then both definite range and inverse cube sensors, along with all other “regular” sensors, revert to the “exponential” detection function.




Figure 3-1. POD versus Coverage curves for three detection functions.

Koopman (1980) made a very important observation about this graph, “At one extreme is the case of the definite range law, at the other the case of random [exponential] search. All actual situations can be regarded as leading to intermediate curves...” (p. 79). We will want to recall Koopman’s observation as we examine the methods used to estimate POD for land SAR searches.

3.3.2 References to Land Search Detection

3.3.2.1 Kelley, 1973

Kelley (1973) discussed detection in the context of “coverage,” which he defined as follows:

Coverage – “Coverage was invented to give a simple measure to the amount of searching that will or has been applied to a given search area. Coverage can be construed as a level-of-effort or search rate” (p. 81).

The author's definition of "coverage" does not match the definition reported in the scientific literature. He uses the term coverage as a way of describing various combinations of two factors: different types of search techniques (e.g., various types of reconnaissance, lookouts, sweeps, and confinement methods), and how often they are applied. The author specifies that there are twelve types of coverage (1 – 12), each associated with a specific set of search techniques and frequencies (Kelley, 1973, pp. 84-85).

The science of search theory defines coverage as the ratio of the area effectively swept to the physical area of the searched segment. Coverage is a measure of how "thoroughly" the segment was searched: the higher the coverage, the higher the POD. Interestingly, each incremental increase in Kelley's (1973) "coverage" also describes a search technique/frequency combination that would tend to increase the level of thoroughness with which the segment was searched. However, categorizing search tactics by "coverage type" is not equivalent to the definition of "coverage" that is quantitatively useful for computing POD estimates, as in the well known formula:

$$POD = 1 - e^{-C}$$

In his Appendix III, Kelley's (1973) rendition of this equation is:

$$POD = 1 - e^{-\frac{a_t}{A}}$$

where a_t is the amount of area searched in time t . This implies several things:

- that the only way to get a "coverage ratio" of less than one in the exponent is to compare the area "covered" up to time t to the size (A) of the segment that is being searched,
- that a completed search always has a "coverage ratio" of 1.0, and
- the only way to obtain a "coverage ratio" greater than 1.0 is to do multiple searches of the same segment.

None of these "coverage ratios" are related to detection probabilities since the density of searching effort is never specified. Only if a_t represented the area effectively swept using Koopman's definitions of effective sweep width and effort would Kelley's version of the equation be correct. However, Kelley is not using those definitions but is instead using a definition of "area searched" that is unrelated to any detection parameters.

Kelley (1973) defined "search rate" as the amount of area "covered" divided by the number of searcher days (or hours) required to "cover" it. This means that the "search rate" is maximized by minimizing the effort density or coverage. Note that this does not address the rate of detection. Most importantly, note that although the units are the same (e.g., square miles per hour), Kelley's definition is fundamentally different from Koopman's (1980) "effective search (or sweep) rate." With Koopman's definition, if a number of identical objects were uniformly scattered over an area with a known density (number of objects per unit area, ρ) and a searcher was sent through the area, the number of objects that would be found per unit of time could be computed as:

$$\text{detection rate} = vW\rho$$

Kelley's "search rate" has no connection with detection, detection rate, or detection probabilities whereas Koopman's (1980) definition incorporates the measure of detectability known as the effective search (or sweep) width, W . Hence, such a computation cannot be done using Kelley's definition of "search rate."

Kelley (1973) also defined three types of sweep searches by the way searchers orient themselves with each other on the line: hand-in-hand, visual contact, and voice contact. The author suggested that, "A hand-in-hand line search will provide the maximum coverage achievable..., visual contact between searchers provides a lesser coverage, ...[and] voice contact between searchers provides the least coverage..." (Kelley, 1973, p. 111). Again, it is interesting that this one element of the author's use of coverage seems to parallel, but not capture, the scientific definition.

3.3.2.2 Wartes, 1974

The first person who attempted to quantify the relationship between POD and searcher activity for the land search community appears to have been Wartes (1974). His initial effort was to conduct an experiment to determine how POD related to searcher spacing in line abreast formations where the searchers followed straight, parallel, equally spaced tracks to cover an area in the shape of a parallelogram.

In the Wartes (1974) study, three search object types were deployed for the searchers: (a) 200-½ pint milk cartons, (b) 23 "unconscious" persons (motionless and silent), and (c) 22 conscious persons able to call out and respond verbally. A description of the method for placing these search objects in the area prior to the experiment was only given for the milk cartons. No information was presented on the placement of unconscious and conscious persons.

In the study, 100 milk cartons (half of those present) were placed so that they could be seen within 20 feet of the searcher's track. Therefore, the milk carton placements were certainly not a uniform random distribution and were made with a preconception of detectability. As one would expect from this distribution, the percentage of detections falls off rapidly as the searcher spacing increases since so many of the search objects were placed near the track. However, as the between searcher spacing increases, the term used by the author (size of area/number of items in the team's path) grows because the area rises much faster than the number of search objects in the searchers path due to the non-uniformity of the distribution. This, in combination with May's description of a "loose line" search (May, 1973, p. 106; see section 3.1.1 above), appear to have caused Wartes (1974) to infer that a type of efficiency was at work. However, the "efficiency" described by the author (see below) is a pure artifact of the way the data was collected and the non-uniform manner of search object placement.

A term called "Thoroughness" was defined by Wartes (1974) as, "The ratio of the items found to the number of items in the team's path" (p. 6). Mathematically, the author represented it as,

$$THOROUGHNESS = \frac{A}{B}$$

Where A = No. of items found,
 B = No. of items in the team's path

The term “in the searcher’s path” is never explicitly defined and the author notes that, “This definition allows thoroughness to exceed 100%” if search objects are detected outside of the searcher’s path (Wartes, 1974, p. 6).

Another term defined by Wartes (1974) was “Efficiency,” which was defined as, “The proportion of items found within a unit area per man-hour of search time” (p. 6). Mathematically, the author represented it as,

$$EFFICIENCY = \frac{A}{B} \times \frac{D}{C}$$

Where A = No. of items found
 B = No. of items in the team's path
 C = No. of man-hours of search time
 D = Size of area searched

Wartes’ (1974) “efficiency” is essentially his “thoroughness” multiplied by the area per hour covered by the searchers and was tied to searcher spacing. The author justified it as follows (capital letters refer to variables in his efficiency equation):

This definition came out of the reasoning that results achieved by teams (A) could only be meaningful if they were compared to potential results (B). Man-hours (C) is a good measure of effort expended to produce results but man-hours can be meaningfully compared only if area size (D) is equated (Wartes, 1974, p. 7).

This justification, and the author’s statement, “If method M covers twice the area as method N (everything else constant), it should be twice as efficient” (Wartes, 1974, p. 7), suggests that he was confusing the ability to detect the search object with the efficiency of the way the sensor was employed (tactic) and the efficiency of the search plan that dictated where the sensor was employed. In fact, there is no way the statement just quoted could ever be correct. Under Wartes’ definition, “twice as efficient” would mean covering twice the area with the same level of effort (man-hours) but getting the same POD. If the sensor, search object and environment remain unchanged, then so does the effective sweep width. If the level of effort remains unchanged, then the only way to cover twice the area is to cut the coverage in half. This will reduce the POD.

Wartes (1974) provided POD estimates based on his experimental results. These estimates were computed as the simple ratio of the number of objects detected divided by the number of objects present. The fatal flaw in this technique is that neither the total number of detections nor the total number of *detection opportunities* is obtained. Objects detected more than once (e.g., independent detections by two searchers during the same search) were recorded only as a single de-

tection. Similarly, objects that were within the maximum detection range of two or more searchers but were missed were counted as only a single non-detection. Finally, an object that could have been detected by two searchers but was detected by only one was counted as a detection. The non-detection by the other searcher was not counted. Such experiments do not examine the detection process at the individual searcher-search object level and produce unreliable estimates as a result.

Wartes (1974, p. 30) tabulated data for three spacings (20, 60, and 100 feet) for each of the three search object types (milk cartons, “conscious” and “unconscious” persons). The detection data for conscious and unconscious persons was combined to form a fourth search object category. Eventually, the POD values for this fourth category (92%, 71%, 53%) were rounded to the nearest 10% and the following formula for “POD vs. Spacing” was published in Green and LaValla (1978, p. 12-8):

$$[3-6] \quad \text{POD} = 100 - \frac{S}{2}$$

The graph of this function is shown in Figure 3-2 below. It would be very difficult to conceive of a detection process that would produce such a graph. The credibility of the POD values given by Wartes (1974, p. 30) can be examined in light of Koopman’s (1946, 1980) observation quoted above. Since Wartes’ daylight experiments took place over a very short period of time in the same location with the same searchers and search objects, it is reasonable to assume that the effective sweep width for each object remained constant. When parallel track search patterns are used in areas that are parallelograms with the first and last tracks one-half track space inside the boundaries, it is permissible to compute coverage as the ratio of sweep width to track spacing:

$$[3-7] \quad C = \frac{W}{S}$$

Under the stated conditions, this method produces the same computed coverage as the ratio of area effectively swept to the area of the parallelogram. We can now assume that one of Wartes’ (1974) POD results is correct, choose one of the curves in Figure 3-1 and use it in conjunction with the above formula to estimate sweep width. We can then see whether the other two PODs are consistent with these assumptions. Substantial differences will indicate the POD data estimated by Wartes are unreliable.

Wartes’ (1974, p. 30) shows a POD of 94% for “unconscious subjects” when the spacing is 20 feet. For the inverse cube model, a coverage of 1.5 produces this POD. Substituting 1.5 for C and 20 feet for S , and solving for W we get a sweep width of 30 feet. This means a spacing of 60 feet should produce a coverage of 0.5 and a POD of about 47%, not the 67% value Wartes shows. Similarly, a spacing of 100 feet corresponds to a coverage of 0.2 for which the POD for any of Koopman’s (1946, 1980) three detection functions is 20% or a little less, not the 51% Wartes gives. Looking at this the other way, the three apparent sweep widths derived from Wartes’ data are 30, 46.8, and 55 feet respectively. Using the average of these three values (44 feet) works out somewhat better with PODs of 99%, 64%, and 42%. In addition, the results of USCG research on visual detection during maritime searches show that the unaided human eye

does not produce the linear relationship with track spacing given in equation [3-6] but comes much closer to Koopman’s inverse cube model. Similar consistency problems exist with the other detection functions as well, so the choice of the inverse cube model was not the cause of the inconsistencies observed in this analysis. Therefore, it is very clear that conclusions drawn from Wartes data, especially when relating POD values to the effort levels required for each of the corresponding tabulated spacings, would be unreliable at best and completely misleading at worst.




Figure 3-2. POD versus Spacing (Wartes, 1974).

It is instructive to compare Wartes’ (1974) approach with that of Koopman (1946, 1980). As illustrated by the upper two curves in Figure 3-1, Koopman’s analysis showed that when the probability of detecting an object close to a sensor’s track is significantly higher than the probability of detecting it at longer lateral ranges, there is an advantage to using a pattern of straight parallel tracks. Such patterns will produce significantly higher POD values than covering the same area with the same effort while doing a “random walk.” However, Wartes’ stated objective was to find a “most efficient spacing” to determine where the maximum POD per “man-hour” occurred. Since effort is proportional to coverage, Koopman’s curves show that for *definite range detection*, POD is exactly proportional to effort expended, up to a coverage of 1.0 when any additional effort would be wasted. This supports Wartes’ (1974) conclusion that , “...there is no ‘most efficient’ spacing.” However, *in all other cases*, Koopman’s (lower) curves show the phenomenon of “diminishing returns.” The rate of POD increase (slope of the POD vs. Coverage (effort) curve) has its maximum value at the origin. Thereafter POD increases more and more slowly as the effort expended increases. This means that the most “efficient” thing to do under Wartes’ approach is to invest the minimum effort in the maximum area (maximize the ra-

tio D/C [Wartes, 1974, p. 6]). Regardless of how “efficient” this might be, it is unlikely to produce useful results since POD values will be at or near zero.

Wartes’ (1974) approach was doomed from the start. As Koopman (1946, 1980) showed, POD is related to effort density or coverage. For any given set of conditions, attaining a specific POD will require that a certain amount of effort be expended. In some situations a parallel track search pattern may yield higher POD’s than “random search” using the same level of effort. However, Wartes was not comparing parallel track patterns to “random search” but was trying to find some optimal spacing for parallel tracks—a goal that has no meaning in terms of POD yield. In the final analysis, Wartes’ approach used some of the same quantities that Koopman’s approach used, but Wartes’ was unable to combine them into a meaningful, coherent theory or method.

As Wartes’ (1975) later realized, the real metric of interest is the probability of success (POS) per unit of effort invested. Maximizing this quantity leads to an optimal search plan, as discussed in section 3.4, *Effort Allocation*.

3.3.2.3 Syrotuck, 1974, 1975

Syrotuck’s (1974) paper begins with a description of Wartes’ (1974) findings and goes on to relate that author’s POD findings to various line search configurations. Interestingly, Syrotuck works (1974, 1975) both include Wartes (1974) in their list of references, and Wartes (1974) includes Syrotuck (1974) in his list of references. In Syrotuck (1974), tables were provided that describe how many “man-hours” and how much time it took Wartes to achieve various levels of POD. Syrotuck did not analyze Wartes’ results but took them at face value. In fact, Syrotuck gives a number of examples based on Wartes’ POD vs. Spacing data to relate POD to the level of effort required to attain it via the corresponding spacing. As observed above, Wartes’ data are far too unreliable to support such conclusions because the experimental, data collection and data analysis techniques were all ill-conceived from the search theory and design-of-experiments perspectives. In addition, the results were shown to be unreliable since the resulting PODs indicate the basic detectability of the object changed when in fact there were no changes in the sensor, search object, or environment during the experiments.

Early in the paper, Syrotuck (1974) briefly stated that the, “United States Coast Guard has devised tables for probability of detection (POD) in searching for lost vessels at sea” (p. 2). He suggested that the tables were the results of experiments where, “...objects were set out on the surface of the water...a great number of sorties were flown over the area, and the number of objects sighted were tabulated.” He goes on to state that, “...the tabulations became a guide in estimating the chances of the object being sighted...” (Syrotuck, 1974, p. 2). By these statements, the author appears to conclude that (a) the USCG developed POD tables, and (b) these tables were directly and only related to the number of objects that were found in the USCG experiments. Unfortunately, neither was true.

The USCG did not develop or possess “POD tables.” The USCG tables to which the author was most likely referring were sweep width tables included in the *National SAR Manual* at that time. Those tables did contain sweep width values deduced from experiments like those Syrotuck described. Those sweep width values, along with the distance flown in the search area while searching (effort) could be used to determine coverage and then POD from a POD vs. Coverage

graph. They were not tables from which POD could be derived directly. In actuality, POD is a function of three variables: effort, effective sweep width, and size of the search segment. Unfortunately, Syrotuck's (1974) paper was primarily based on the premise that the USCG tables described POD directly—a false assumption. This coupled with the author's heavy reliance on Wartes' (1974) questionable conclusions about POD, searcher-spacing, and the relationship between the two, served as a poor foundation for accurate information, at least in terms of scientific search theory.

Syrotuck's (1974) descriptions of his incorrect impression of what USCG detection experiments were actually measuring, and what sweep width tables actually represented, caused many subsequent land SAR researchers (e.g., Wartes, 1974; Colwell, 1992; Bownds et al., 1981, 1991) to presume that POD could be “measured” through simple experiments where a number of objects are distributed throughout a search area and the POD is determined by the ratio of the number of objects detected to the number present. However, the USCG experiments Syrotuck described were sweep width experiments, not POD experiments. POD is a function of effort, effective sweep width, and size of the search segment and cannot be determined by experiments as simple as those of Wartes (1974).

A graph included in the appendix (Graph 1) of Syrotuck (1974) labelled the “chances of success versus number of sweeps” (Syrotuck, 1974, p. 21) appears to show cumulative POD. In terms of search theory, the only concept addressed in the document is POD, but the author's description and discussion of the concept are not based on the numerous scientific references available (even at that time). Instead he appears to be drawing only on Wartes' (1974) conclusions.

Syrotuck (1975) mentioned U. S. Coast Guard's search planning aids, specifically SARP and CASP, and even offered an example of a search plan that the author implied was developed using CASP-generated data (p. 34). The author also made a direct comparison between his methods and the methods used by the USCG when he said, “Readers may recognize some similarities to the US Coast Guard ‘SARPS’ [sic] system that is used to locate lost ships at sea” (Syrotuck, 1975, p. 21). He also mentioned the “repeated expansion” tactic, developed and used by the USCG for locating distressed mariners at sea, as an acceptable operational technique as if to suggest that maritime methods would also work in the land environment (Syrotuck, 1975, p. 25). The author seemed to hold the USCG methods up as examples of how search planning and operations should be conducted, at least in the marine environment.

Essentially no terms related to scientific search theory are used in Syrotuck's (1974) paper beyond those related to “probability of detection.” Unfortunately, every time these terms were used in this paper, it was describing some derivation of Wartes' (1974) findings. Although the terms “cover” and “coverage” were used a few times throughout the document, there is no evidence that the author intended them as they are defined in scientific search theory. He only used the terms as the dictionary defines them.

Syrotuck (1975) defined “sweep width” as follows:

Suppose we have 10 searchers and we space them 10 feet apart. They will form a line that has a width of 100 feet. This we will call “sweep width” (p. 27).

Syrotuck (1975) also defined “optimum sweep width” as being the same as the optimum size of area (p. 29), although this definition clearly does not match the definition described in the search theory literature. He goes on to say that, “We can vary track spacing by any amount...however, it is more convenient to make it some ratio of the sweep width” (p. 27). The author also gave only one description of “Coverage Factor:” “...a ratio of sweep width divided by track space or $C = W / S$ ” (p. 27). When “W” represents Koopman’s (1946, 1980) definition of sweep width that relates to the detectability of the search object, the searcher tracks are perfectly straight, parallel and equally spaced, and the area covered is a rectangle or parallelogram one-half track space larger than the search pattern on all sides, this equation is correct and has been used by the maritime SAR community for many years. However, when Syrotuck’s definition of sweep width above is used, the equation has no meaning. Placing 10 searchers at equal intervals along a 100-foot base line would always result in a coverage of 100/10 or 10, regardless of the values of any of the variables that affect detection.

3.3.2.4 *Bownds et al. (1981) and Bownds & Harlan et al. (1991)*

Bownds et al. (1981) published a report regarding several detection “experiments.” In these experiments, aircraft were used to search for people in the desert near Tucson, Arizona. The experiments were conducted in conjunction with the Pima County Sheriff’s Department and Detachment 1, 37th Air Rescue and Recovery Squadron, of the U.S. Air Force Air Rescue and Recovery Service (AARS). The stated purpose of the experiment was to, “...establish a conservative measurement of the probability of detection (POD) of the United States Air Force helicopter rescue teams searching Sonoran desert terrain for lost persons” (Bownds et al., 1981, p. 2).

The authors claimed that, “...almost no information is available on efforts to measure the PODs of SAR resources in the inland search for lost persons,...[but] the one exception is...Wartes [1974]...” (Bownds et al., 1981, p. 2). As in many of the inland references, the authors also cited Wartes (1974) as suggesting that, “...‘nonthorough’ searching with multiple passes produces a higher probability of success per hour spent in the search...” (Bownds et al., 1981, p. 31).

Five documents were included in the “References” section of the document: Wartes (1974), the 1973 U.S. National SAR Manual (NSM), and three statistics references. In the 1973 version of the NSM, no land search information was included and no sweep width values were provided for search objects on land. The few sweep width tables that were included were limited to the marine environment and typical maritime search objects. Compared to sweep width tables included in the much more recent *U.S. National SAR Supplement* (2000), they were very limited (less than a tenth as many entries) and were based on “survey data” of reported detection ranges during normal operations and searches, not on formal experiments.

In their experiments, Bownds et al. (1981) used a six square mile search area (2 miles east to west by 3 miles north to south). Five separate sorties were flown, each described as an “experiment.” Three were flown on bright, sunny days, and two were flown on overcast days. No other specific meteorological information (e.g., visibility, cloud levels, etc.) was provided. In addition, the authors reported that, “The speed, altitude, and spacing between creeping lines were decided on by each helicopter crew. (The average speed was about 60 knots, the altitude 175 feet, and the track spacing ¼ mile)” (Bownds et al., 1981, p. 7).

In the experiments, four crew members (two pilots, two “scanners” [observers]) in a “Huey” helicopter searched from the air for human subjects on the ground. The subjects were described as wearing “everyday” clothing; that is, clothing without large areas of highly visible colors. Because “...all victims who were assigned to move about in the open and attempt to attract attention were found” (Bownds et al., 1981, p. 10), the authors chose to include in their numerical analysis only data related to, “...victims under cover and simulating unconsciousness” (p. 10). Thus, the data in the report excludes all subjects who moved into the open, attracted attention, and were found.

The authors concluded that on bright, sunny days 7 out of a possible 24 victims were found (29%), and on overcast days 11 out of a possible 16 victims were found (69%). The results were displayed in two “single-pass” POD tables: one for bright, sunny conditions and the other for overcast, subdued light conditions. The authors suggested that the use of these tables allowed one to determine a range of POD values for each single, complete “pass” over the search area for various levels of confidence. From each of these two tables, another table was extrapolated that described how many passes would be required to achieve selected cumulative POD targets for a wide range of confidence levels. These secondary tables also allowed a user to estimate how many flight hours would be required to achieve the target POD values.

In Bownds et al. (1981), the following equation was used to compute cumulative POD (a curve representing this equation was also provided):

$$POD_{cum} = 1 - (1 - p)^m$$

where p is the “single pass” POD and m is the number of complete passes over the search area.

The use of this equation required that all successive m searches have the same POD value. Since this is how the authors intended to achieve higher target POD values (e.g., repeatedly reproducing the same search), this equation is adequate. However, when the POD values of successive searches differ, this more general form for computing cumulative POD should be used, but is not described in the document:

$$POD_{cum} = 1 - (1 - POD_a)(1 - POD_b)(1 - POD_c)...(1 - POD_m)$$

Where POD_a represents the POD values for each search up to m searches (complete passes over the search area).

The authors provided definitions of “POA,” “POD,” and “POS” (pp. 36-37) that match Koopman’s (1946, 1980) definitions of the same terms. “Coverage” was used several times in the document, but only as the dictionary defines it, not as defined by Koopman.

Bownds & Harlan et al. (1991) conducted experiments in southern Arizona similar to those conducted by Bownds et al. (1981). However, the latter experiment involved nine “scenarios” (sorties) and was conducted over mountainous terrain with elevations ranging from 6000 to 7904 feet. Three of the principle authors of the 1991 report were also authors of the earlier Bownds et al. (1981) report; and, the experimental design, procedures and participants were also quite similar: only “Hueys” were used as search aircraft, subjects wore “everyday” clothing, and the stated

purpose of the experiment was to, "...measure the effectiveness of United States Air Force Air Rescue crews searching for lost persons in a rugged mountainous environment of Southern Arizona" (Bownds & Harlan, 1991, p. 1).

Eight documents were included in the "References" section of this report: Wartes (1974), Bownds et al. (1981), a 1978 version of LaValla et al. (1997), Richardson & Corwin (1980) and four other less relevant references. The Richardson & Corwin citation was used to support the authors' statement that, "The United States Coast Guard has conducted research into the effectiveness of resources in searches of the maritime environment" (Bownds & Harlan et al., 1991, p. 1). This statement implied that the authors were simply doing for land search what the Coast Guard had done for maritime search. This was the only citation for this reference in the document. The authors also claimed that, "...little research exists that has quantitatively measured the effectiveness of search resources in finding lost persons in the inland environment" (Bownds & Harlan et al., 1991, p. 1). According to the authors, only two previous land detection experiments had been conducted prior to their work in 1991: Wartes (1974) and Bownds et al. (1981).

Nine scenarios were flown; all on sunny days but in various wind conditions. Some low-ceiling searches were scheduled but all were eventually aborted for safety reasons. The size of the search area and segments were not reported.

The locations for the subjects involved in the scenarios were selected from previous cases of the Pima County Sheriff's Department that had occurred between 1980 and 1983 ($n = 25$). All subjects were placed only in locations where previous subjects had been found; however, "The actual location of a victim and whether he/she was on a trail or in a drainage was determined randomly (with a random number generator)" (Bownds & Harlan et al., 1991, p. 11).

"Mission 9" in the report involved seven subjects that simulated an unconscious subject under cover. None of these subjects were found and none were included in the POD calculations reported in the "combined results" section of the report.

Several very practical suggestions for survivors were recommended in the report; but, unlike the earlier experiment (e.g., Bownds et al., 1981), only the resultant number of finds was reported for each mission. No confidence intervals were reported and no POD tables were included.

3.3.2.5 Perkins (1989)

"Critical Separation" is defined as that spacing between searchers that equals twice the maximum detection range of an object used to represent the spectrum of clues being sought. According to Perkins (1989) this yields a 50% POD. A graph showing POD values for other spacings was also provided in Perkins (1989, p. 6). The mathematical function that produces the graph was not given nor were any data that support an empirically derived graph provided. The curve appeared to be linear for spacings from zero up to about one critical separation where it then began a gradually lessening rate of decline as spacings continued to grow, apparently approaching zero as an asymptote from above. The linear portion seemed to obey the following equation, which was based on Wartes' (1974) findings and similar to an equation first seen in Green & LaValla (1978):

$$[3-8] \quad POD = 1 - \frac{S}{2}$$

where S is the searcher spacing expressed as a multiple (or fraction) of the critical separation. Following this formula for larger spacings leads to the absurd results of zero POD at a spacing of twice the critical separation and negative PODs for larger spacings. Apparently Perkins (1989) realized this and drew a curve that shows a more gradual decline in POD as spacing increases. The graph shown in Figure 3-3 approximates that shown in Perkins (1989). In Figure 3-3 it is assumed that the non-linear portion of the curve for values of S greater than one (critical separation) is given by:

$$[3-9] \quad POD = \frac{1}{2S}$$

This matches the last data point on the linear portion and has the approximate shape shown in Perkins (1989) for the remainder of the curve. It also corresponds to the demonstrable notion that when a finite maximum detection range exists, then the POD in a pattern of equally spaced parallel tracks will be inversely proportional to the spacing once that spacing exceeds twice the maximum detection range. However, there is no evidence that Perkins used this relation.

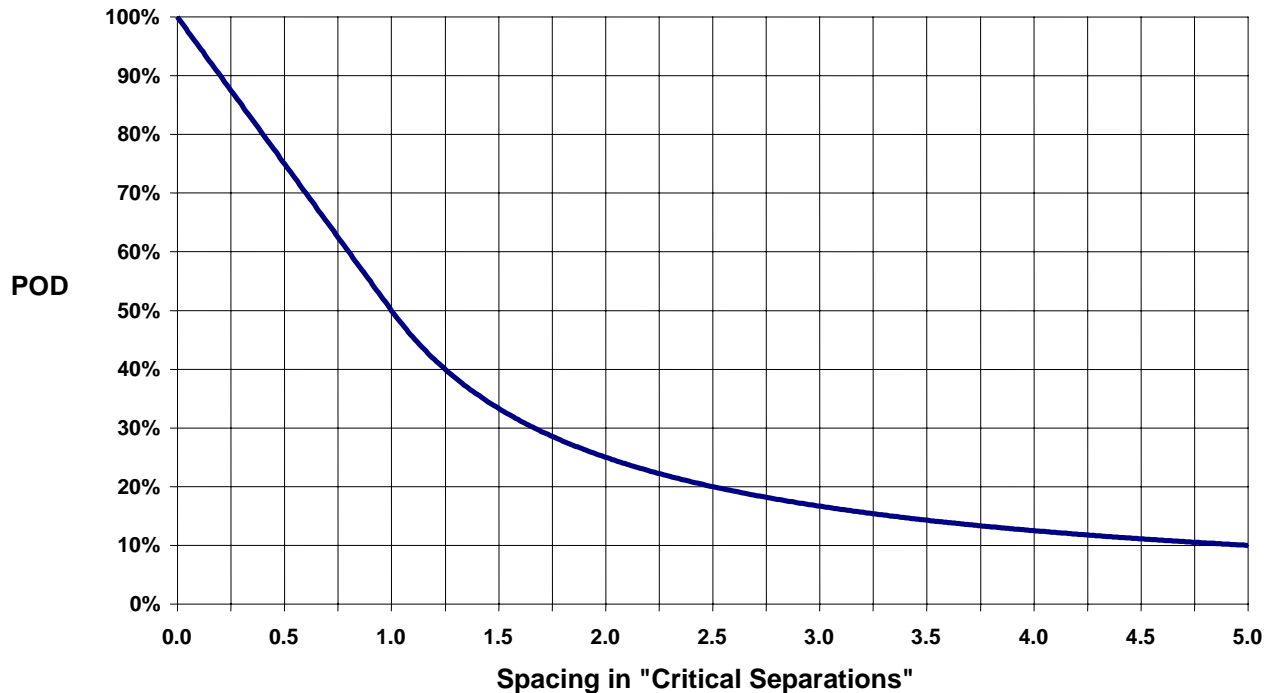


Figure 3-3. POD versus Spacing in "Critical Separations" (Perkins, 1989).

3.3.2.6 Colwell (1992)

In the early 1990s, Colwell (1992) attempted to refine the relationship between POD and searcher spacing. He conducted a number of experiments in British Columbia in an environment similar to that where Wartes (1974) had conducted his experiments. Although Colwell's experiments were somewhat more extensive, they suffered from the same fatal flaws as Wartes' original experiments. In addition, Colwell's analysis of the data was also fatally flawed, allowing Colwell (1992, p. 29) to reach the startling conclusion that at wider spacings, the spacing could be doubled without affecting the POD.

The flaw in Colwell's (1992) analysis was due to a fixation on searcher spacing, as opposed to level of effort. Colwell observed that once the spacing between searchers was so large that only two would "fit" between the sides of the search area he was using for his experiment, the number of objects detected remained roughly constant as the searchers became more and more separated until they reached the maximum separation determined by the width of the search area. The reason for the constancy in the number of detections is obvious: at each spacing there were two searchers moving the same distance across the search area. That is, the level of effort did not change and so one would not expect the number of detections to change either, assuming a uniform distribution of objects throughout the area being searched. This allowed many different spacings to be associated with the same level of effort and it was a very serious error. Nevertheless, Colwell (1992) concluded that, "The horizontal plateau region between approximately 100 m. and 200 m. spacing indicates that no benefit is obtained by placing the searchers at the closer 100 m. spacing. The 200 m. spacing permits a much larger area to be searched at approximately the same POD level" (p. 29). In other words, Colwell claimed that on a one kilometer baseline, ten searchers spaced 100 meters apart will have no better chance of finding a person known to be somewhere in the area than five searchers spaced at 200 meter intervals. This simply cannot be true.

Figure 3-4 below shows an example of one of Colwell's POD vs. Spacing curves. It is a third-order polynomial "fitted" to his detection data. The "horizontal plateau region" is simply the region in the vicinity of the inflection point. That is, the "horizontal plateau region" is an artifact of the mathematical function that was chosen for "fitting" the data and is not reflective of any "real-world" phenomenon any more than the negative POD values this function produces for spacings above about 250 meters. Other functions could have been chosen that would fit the data just as well and would also have reflected the obvious fact that POD should decline toward, but never reach, zero as the level of effort declines toward, but never reaches zero (because that would require an infinite spacing), at least in theory.




Figure 3-4. POD versus Spacing in meters (Colwell, 1992).

As observed above, Colwell (1992) allowed many spacings to be associated with the same level of effort. He also made the opposite error with his “baseline offset” technique (Colwell, 1992, p. 19). In the author’s Figure 13 (p. 51), the author shows how data from a search with a 10-meter searcher spacing can be grouped in subsets to produce data for a 30-meter spacing. In that figure, the data from ten searchers is divided into three groups—one group of four and two groups of three—all at the same 30-meter spacing. So, the number of detections for four searchers was grouped with the detections made by only three searchers as if the 33% additional amount of effort were irrelevant. In the example, it is interesting to note that the number of detections were apparently the same for the four-searcher grouping and one of the three-searcher groupings. The number of detections was significantly less for the other three-searcher grouping. This indicates significant “noise” or “scatter” in the data that is only exacerbated by grouping data from differing levels of effort under the same spacing. In any case, it was a very serious error to allow substantially different levels of effort to be associated with the same spacing, merely as a result of how the data from a single search was grouped.

Colwell (1992, p. 23) defined “efficiency” as the ratio of POD to the corresponding number of searcher-hours. Recall that the above analyses basically ignored level of effort when obtaining POD estimates. It is a logical non-sequitur of the first order to first claim, in effect, that the POD of a search is unrelated to level of effort by grouping the results from different effort levels under the same spacing, and then claim that the “efficiency” of that same search depends on a relationship between POD and level of effort. The fact of the matter is that POD is related to the level of effort expended and that in any particular circumstance, a given POD will require a corresponding level of effort. The ratio of POD to effort expended has no meaning.

Like Wartes (1974), Colwell (1994) later realized that the ratio of interest was POS/effort. The search plan that produced the highest overall probability of success with the available effort was the best, most “efficient” plan. The problem of deciding where to apply effort and in what concentrations (densities) given the probability density distribution on search object location is the optimization problem that needs to be solved.

3.3.3 Multiple Searches vs. Single Searches

The experiments and analyses described above had several fundamental flaws:

- The detection process itself was not examined (e.g., the characteristics of how a searcher detects objects as he or she approaches, passes, and recedes from them).
- Numbers of detections were not compared to numbers of detection opportunities, leading to false POD “readings.”
- Analyses of the data did not even attempt to ascertain the relationship between POD and effort expended. In fact, the connection between POD and effort was overlooked during the data analysis process, as exemplified by Colwell’s (1992) procedure that allowed many spacings to be associated with the same level of effort and multiple levels of effort to be associated with the same spacing. This made the resulting relationship between POD and spacing not credible. It also made any later attempt to use such data to deduce a relationship between POD and effort not credible.

This combination of oversights led to some false conclusions about the relationship between effort expenditure and POD when later attempts were made to relate the two on the basis of experiments and analyses that ignored this essential relationship initially. One of the most popular misconceptions is that a series of low POD (i.e., low coverage) searches can produce a higher cumulative POD for the same total amount of effort than one would expect from a single search. The following quotes exemplify this misconception:

- “Segments are frequently searched more than once, as it is usually better to search an area twice under low-POD conditions than to search it once more thoroughly” (Hill, 1997, p. 132).
- “Repeated sweeps of the same area with a wide spacing will be more efficient than a single sweep with close spacing” (LaValla et al., 1997, p. 164). (attributed to Wartes [1974] by the authors.)

The land search literature (e.g., Stoffel, 2001; Hill, 1997; LaValla et al., 1997) is in almost universal agreement that multiple low POD searches of a segment are better than searching the same segment once more thoroughly for the same total number of searcher-hours (roughly, effort) expended. This notion is false and seems to have several sources. First, it is a misinterpretation of a statement published by Wartes (1975b) and based on his earlier experimental results. The primary impetus for Wartes’ (1974) experiments was a growing body of anecdotal evidence indicating untrained people (usually friends and relatives of the missing person) searching randomly often found the subject before the trained, organized search teams did. Wartes’ (1975b) actual statements read, “Repeated sweeps of the same area with wide spacing will yield better results

than a single sweep with close spacing” (p. 6) and, “Non-thorough methods produce better results in less time” (p. 7). The “results” the author was speaking of involved maximizing the POS while minimizing the mean time required to achieve it when using line abreast formations (grid searching). They were not related to the final POD or POS results after all the effort had been expended.

To illustrate Wartes’ (1975b) actual finding, we will use one of his examples as a basis (Wartes, 1975b, p. 7). Assume 15 searchers are available to search a rectangular area. Suppose further that the search object is known to be somewhere in the rectangle and probability density (even if Wartes [1975b] was not thinking in these terms) is uniform over the rectangle. Finally, suppose it takes a searcher 3 hours to traverse the full length of the rectangle. Using all 15 searchers in a single line abreast on a 30-foot spacing, at the end of the first hour $1/3$ of the area would have been searched. Using an assumed sweep width of 90 feet and the exponential detection function (more concepts foreign to Wartes in 1975b), a coverage of 3.0 is computed making the POD over the area searched so far 95%. This produces a POS at the end of the first hour of 0.3333×0.95 or 31.67%.

An alternate tactic would be to divide the rectangle into thirds and create three lines abreast of five searchers each on a 90-foot spacing. In this way, the entire rectangle can be covered in one hour at a 90-foot spacing (or a coverage of 1.0) with a resulting POD (and POS) of 63.21%. Clearly the second tactic has a much better chance of locating the search object in just one hour than the first tactic does. In fact, the second tactic makes a successful outcome nearly twice as likely in that first critical hour. However, by the time all the effort has been expended at the end of three hours (assuming the five-person search teams offset their search legs appropriately in the second and third hours to ensure uniform coverage), the POD and POS would come out to 95% for both tactics.

Nevertheless, Wartes (1974, 1975b) did apprehend a fundamental truth, even if it was later misinterpreted. Apparently the standard method of grid searching at that time placed the searchers very close together. An analysis of Wartes’ (1974) description estimates that coverages as high as 8 were being routinely applied using line abreast formations. The time and resources required to cover any significant area at such high coverages approach astronomical magnitudes. It is worth noting that neither the POD vs. Coverage graph then found in the *National SAR Manual* or today found in the *IAMSAR Manual* or the *National SAR Supplement*, go above a coverage of 2.0. Most maritime search planners try to achieve coverages in the neighborhood of 1.0. Clearly the excessively high coverages apparently being used on the ground circa 1974 were well beyond the point of “diminishing returns” for adding more effort. It is little wonder that Wartes (1974, 1975b) was able to make a strong case for “non-thorough methods” in SAR searches. Evidence searches, on the other hand, might require a “thorough” search since the search activity could alter or destroy evidence not found on the first pass. However, time is usually a much less important factor in such cases than it is in SAR.

Search theory shows that the optimal search plan for uniform distributions is to apply the available search effort uniformly over the entire distribution in the shortest possible time (everywhere simultaneously, if possible), and this was the truth apprehended by Wartes (1974, 1975b). However, it is widely believed by many inland search managers that it is significantly better to completely search the same segment multiple times, using small numbers of searchers over several

operational periods, than to expend the same total amount of effort in the first operational period searching it once at a higher coverage. That is not the conclusion Wartes (1974, 1975b) reached, and it is not supported by search theory. Wartes (1975b) was clearly describing an *alternative tactic* for applying the *same effort* to the *same segment* in the *same operational period*. Nevertheless, statements such as Hill's (1997), "Segments are frequently searched more than once, as it is usually better to search an area twice under low-POD conditions than to search it once more thoroughly" (p. 132), are to be found throughout the land SAR literature. No alternative tactics like those of Wartes (1975b) are given and the clear implication, based on the "4-6 hour rule" of segment sizing, is that two low-POD searches completed over two operational periods are usually better than one higher POD search completed in one operational period with the same total expenditure of effort. This notion could easily cause search managers to hold effort in "reserve" or send it to segments where the "return on investment" (ROI), i.e. increase in cumulative POS, will be much lower than the ROI that could be obtained by using that effort to increase the coverage of a more "profitable" segment.

This widespread belief that the primary characteristic of searcher deployment that affects POD is the spacing between searchers has led to the few experiments actually described in the land search literature being aimed at establishing and developing the relationship between POD and spacing. Wartes' (1974) results were interpreted by some as indicating a linear relationship between POD and searcher spacing. A linear formula describing this apparent relationship has even been widely published. Not satisfied with this formula, Colwell (1992) tried to do some more sophisticated experiments and analysis, but was still ensnared in the POD vs. Spacing trap. This misconception about the relationship between POD and spacing may be borne out of a misunderstanding of the short-cut formula that computes coverage as the ratio of sweep width to track spacing. In maritime search planning, there is a definite relationship between the number of tracks in a parallel sweep (PS) search pattern, their spacing, and the size of the area searched. Uniform coverage is also a requirement. Thus, there is a definite relationship between spacing and effort (coverage) given these constraints. For the same area, doubling the spacing halves the effort required; but it also halves the coverage and reduces the POD accordingly.

Other problems with the notion that POD is purely spacing-dependent include the biases in Wartes' (1974) experimental POD results, and the work performed by Colwell (1992). The latter actually implied that, at wider spacings, increasing the spacing still more would not decrease, and might even slightly increase, the POD. Colwell concluded that it is better to search a segment twice at the wider spacing than once at the narrower spacing. Even a simple search model will demonstrate this is impossible. Other problems with Colwell's work are described in Frost (1998a).

Even the *National SAR Supplement* (National SAR Committee, 2000) can be cited, if one is not careful, as supporting the contention that two poor searches are better than one good one. The land SAR POD tables given there as a function of track spacing for searches conducted from the air seem to support this notion. However, if one observes that all POD values are multiples of 5, it is possible in most cases to explain the problem as one of rounding error. The same is true for the cumulative POD table given in the same chapter. However, these "rounding errors," if that is what they are, can produce some large anomalies due to large-scale error propagation when the two tables are used together. For example, it is possible to conclude that two searches done at 700 feet AGL over heavy tree cover using a two-mile track spacing can produce a cumulative

POD of 15% while a single search with the same effort using a one mile track spacing will produce a POD of only 5%. This is a mathematically impossible result.

It should be reiterated at this point that, when using an exponential detection function, POD is solely dependent on the amount of effort applied uniformly to the segment regardless of whether it is applied all at once or incrementally—i.e., the cumulative POD comes out the same either way (Koopman, 1946, 1980). For any other detection function, i.e. one that is “better” than exponential in the sense of being closer to the ultimate of definite range detection, incremental application of search effort will always produce *lower* PODs than using all the effort in a single search. In other words, for detection functions that are “better” than exponential, the cumulative POD over several searches will always be less than the POD of a single search using the same total amount of effort. This is the precise opposite of the guidance given in the land SAR literature.

3.3.4 One POD for Multiple Objects

It is widely believed that POD measures the likelihood of finding anything associated with the subject, whether it is a small clue or the actual subject. This notion is exemplified by a typical question asked of search teams during the debriefing following their return from a search: “If there were 10 clues of varying size [and detectability] in the area you were assigned to search, how many would you have found?” (Cooper et al., 1996, p. 296). Essentially all of the land search authors represent POD this same way: it represents the chances of finding any clue or evidence in a segment, and completely disregards the (usually wide) range of possible clue types and their respective detectabilities.

Hill (1997) defines POD as, “...the percentage of clues that a search resource would be expected to find,...” (p. 131). As an example, the author goes on to state that, “...if there are 100 clues available in a segment, a POD of 50% means that the segment is searched in such a way that approximately 50 clues should be found” (p. 131). Dougher et al. (2000) defines POD as the, “...chances that the subject or clues will be detected by a search of a designated area” (p. 34). The use of one POD to represent the chances of finding any object or objects, regardless of the range of possible object characteristics (size, color, contrast, etc.), sensor and environment, is a universal error contained in the land search literature.

In fact, POD is a value that applies to one “clue” in one environment, under one set of conditions, being sought with one particular sensor. So, although the answer to the questions regarding multiple items asked by Hill (1997) and Cooper et al. (1996) are interpreted to be POD, in reality they are not.

At the present time, inland search procedures have no way to estimate “detectability” or sweep width for search parties on the ground. (Some limited land sweep width tables for airborne searchers do exist, but their accuracy, origins and validity are unknown.) Although some authors have used average maximum detection range (the average maximum distance from a searcher that an object can be visually detected in a representative environment) as an estimate for sweep width as a stopgap method in the absence of experimental data, sweep width data is just not yet available for the various inland environments. Thus, the primary method currently used in land

search to estimate POD values is purely subjective (e.g., answering the question “If there were 10 clues of varying size [and detectability] in the area you were assigned to search, how many would you have found?” [Cooper et al., 1996, p. 296]). The logical conundrum presented by asking a searcher to come up with a single number that represents the probability he would have found an unspecified “something” if there had been such a “something” in the segment by asking him to estimate a count of how many such “somethings” he would have found out of a possible 10 or 100, should be obvious. Notice that, although not stated, the question being asked implicitly assumes a uniform distribution of clues over the segment’s area – something that is unlikely to be true except in experiments or training sessions. It is not the same question as, “If there was only one item, and that one item was equally likely to be anywhere in the segment, what were your chances of finding it?” Yet, the answer to the latter question, if based on some objective criteria, is the one that truly reflects POD for that item.

3.3.5 Conclusions

When Koopman’s (1946, 1980) analysis is compared to those of Kelley (1973), Wartes (1974, 1975a-b), Colwell (1992, 1994), and all others that form the basis for the search planning guidance in today’s land SAR manuals, it becomes clear that the latter analyses were severely hampered by omission of the sweep width concept. In its absence, terms and metrics, like “thoroughness” and “efficiency,” were created that are not meaningful. Finally, all tried to relate POD directly to searcher spacing, an approach that lacks the mathematical correctness and generality of Koopman’s approach and eventually leads to the incorrect conclusion that POD is related only to the distance between adjacent searchers. The fixation on relating POD to searcher spacing hampered all attempts to address probability of detection and effort allocation issues in land SAR up until 2002 when the first demonstration of a procedure to determine effective sweep width for ground search was performed under the auspices of NSARC.

The sweep width concept in its pure form is far more flexible and makes more sense mathematically. Expressing search effort as sweep width multiplied by distance gives it units of area that make sense in the computation of coverage. Coverage defined as the ratio of area effectively swept to the area of the segment certainly expresses how “thoroughly” the segment was covered. Finally, expressing POD as a function of the coverage is more intuitive. Zero coverage means zero search effort and that certainly produces a zero POD. The POD graph based on spacings does not show this because to reach zero effort, the spacing must be infinitely large.

Similarly, as coverage increases, 100% POD is approached asymptotically from below – another intuitive feature of graphing POD vs. Coverage. This feature clearly illustrates the phenomenon of “diminishing returns” where adding more effort to already high levels produces little benefit in terms of POD. The POD vs. Spacing graphs, on the other hand, leave the opposite impression because the curve rises most steeply where the decreases (reading the graph from right to left) in spacing are smallest. It is not obvious that decreasing the spacing from 20 feet to 10 feet, for example, requires doubling an already high level of effort. Taking this to its logical extreme, the graph shows 100% POD at zero spacing. This might be intuitive were it not for the fact that zero spacing requires an infinite amount of search effort to cover any area at all—an impossible situation.

POD vs. Spacing graphs also leave the false impression that very low levels of effort are almost as good as low levels of effort, leading in turn to the false notions of “efficiency” described in the sections above. Note that on Koopman’s graph (Figure 2-5), a coverage of 0.1 produces about one-half the POD of a coverage 0.2 search. That is, at low coverages, cutting the effort in half cuts the POD in half also. Note that the graph in Figure 3-3, assuming our guess about the non-linear portion is approximately correct, shows the same kind of result in that a spacing of four “critical separations” produces half the POD (12.5%) as a spacing of two “critical separations” (25%). However, the visual impression is entirely different from that of Koopman’s representation and it is not nearly so obvious how rapidly zero POD is being approached as effort (coverage) decreases.

There may be a temptation to compare effective sweep width with critical separation and try to find some connection between the two. This is not possible. Effective sweep width is a measure of how much detection can be expected from a given sensor when searching for a given object in a given environment, when the object’s location is not known. Critical separation is merely a measure of the maximum distance at which a given sensor can detect a given object in a given environment when the object’s location is known in advance of making the measurement. Note that Koopman’s definition of effective sweep width contains no quantities related to maximum detection range. In fact, the maximum detection range could be anything from microns to light-years without affecting Koopman’s definition. This is not to say that determining maximum detection ranges in the field are without value. They might be useful for determining sweep width correction factors, for example, but any such relationship would have to be determined empirically by analyzing data from experiments in the field.

In short, none of the POD estimation procedures found in the land SAR literature are compatible with search theory and none can be modified to make them compliant with search theory. Only the new approach that is based on established search theory principles shows promise. This work is sponsored by NSARC and endorsed by NASAR and the AFRCC.

3.4 Effort Allocation

The goal of search planning is to allocate the available searching effort (resources) so that the probability of successfully finding the search object is maximized in the minimum time. Stone (1989) describes the necessary and sufficient conditions for computing such an optimal search plan and provides methods for doing so. The land SAR literature also contains methods for allocating the available resources but they do not generally result in optimal search plans and often result in substantially sub-optimal plans. Not surprisingly, there is significant disagreement among the various methods presented in the land SAR literature, and none agree with or approximate the methods presented in the scientific literature.

3.4.1 “Non-thorough” Effort Allocation

A curious and erroneous result came out of Wartes’ (1974) report that has been repeated in nearly all of the subsequent land search literature. In Wartes (1974), the author suggested that there is an advantage to conducting a search using a wider search spacing and this has led to the

assertion that two low effort searches are better than one high effort search, given the same total level of search effort (see section 3.3.3, *Multiple Searches vs. Single Searches*, for a more detailed discussion of this concept). This is of course not true—the same level of effort applied in the same segment produces the same coverage, regardless of whether it is applied in one or multiple installments— but the myth has survived nearly thirty years in the land search literature. In fact, examination of Koopman’s (1946, 1980) curves shows that applying effort in multiple installments can do no better than equal (as with the exponential detection function) and can easily do worse than (as with the “better than exponential” detection functions) the POD of a single search that consumes the same total amount of effort.

In a way, Wartes’ (1974) experiment was actually aimed at the correct objective, uniform optimality, because the author was looking for ways to locate the search object more quickly. However, the stated objective was to find the most “efficient” spacing for searchers in a line abreast formation to use when searching a *single*, contiguous segment. The author correctly concluded that the then-current methods for applying search effort (i.e., traversing the length of the area only one time using a single line abreast, with very narrow spacing) were sub-optimal in terms of mean time to detect. He did not provide a definitive method for optimally allocating effort and used some vague and inconsistent metrics for “thoroughness” and “efficiency.” In the end he concluded, “There is no ‘most efficient’ method” (Wartes, 1974, p. 37). The author also did not look at the problem of how to allocate effort among a number of different segments. His findings and methods were widely misinterpreted and these misinterpretations were subsequently used by others to reach a number of false conclusions.

Interestingly, Wartes later used POS to justify his “efficiency” principle. Wartes (1975b) suggests that attaining a higher POS, when comparing one search plan to another, was a desirable goal. By virtue of the author’s, “...assumption that the lost person is within the area” (p. 5), a uniform probability distribution is implied when he provides measured search areas in his POS examples.

The author even recognized the importance of time (and searcher speed by inference) when he stated, “It is essential to find the lost person quickly...” (Wartes, 1975b, p. 7), and when he graphed POS as a function of time in his examples. This is also why the author started on the right track for effort allocation that maximizes POS in the minimum time when he generally stated, “Repeated sweeps of the same area with wide spacing will yield better results than a single sweep with close spacing” (Wartes, 1975b, p. 6). This statement, when read in context, meant that when the probability on search object location was uniformly distributed over the area, the POS could be increased more quickly in the early hours by doing low-coverage sweeps that together covered the entire area in a short time and then repeating them until the available effort was exhausted, than by doing one high coverage sweep that took much longer to cover the area just once, as shown in Figure 3-5. However, if the total effort expended and total elapsed time were the same for both methods, the final POD and POS and the time required to achieve them would be the same as well. This is illustrated in Figure 3-5 by the concordance of the last data point at the upper right of each graph.

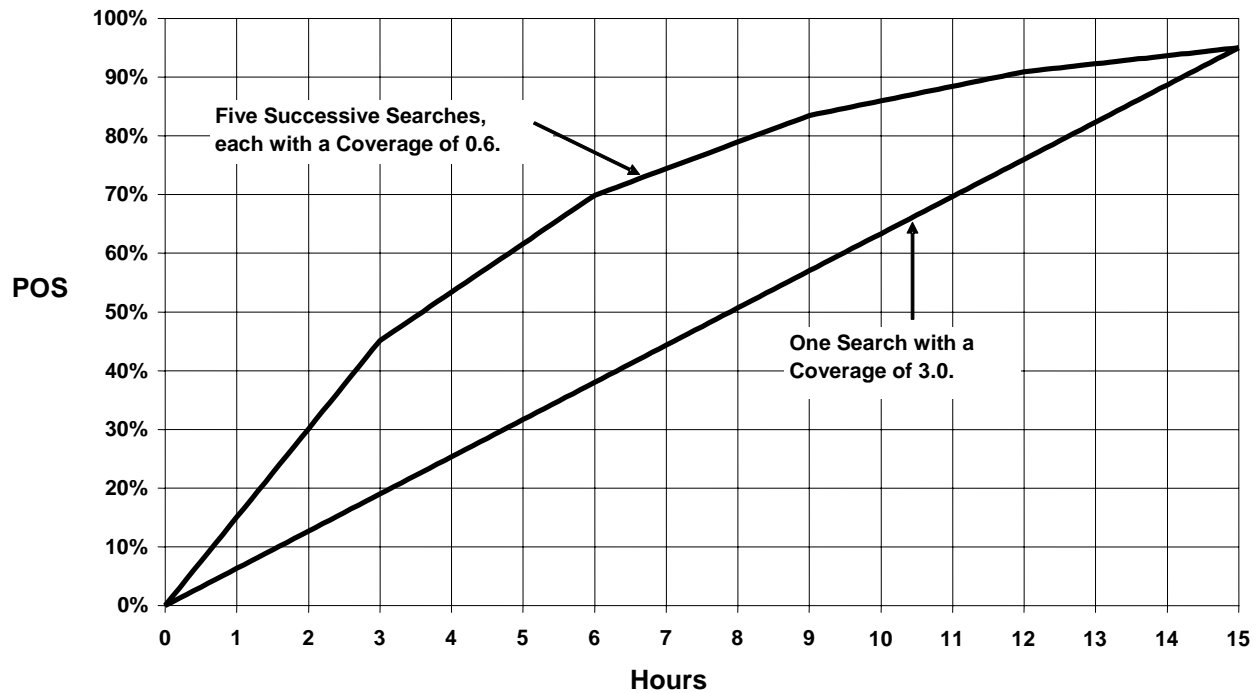


Figure 3-5. POS versus Time (exponential detection).

Unfortunately, later authors incorrectly took Wartes' (1975b) statement to mean that all segments, even high priority ones, should only get a cursory low-coverage search initially and then be re-searched at a later time, again with a low coverage. Even if the search planner does not do this deliberately, the potential enlargement of the search area with the passage of time can become the focal point and cause resources to be spread thinly over large areas while those originally thought to be worth a larger investment of effort languish following an initial cursory search. Very often the tendency to cover more area once before concentrating effort in more promising locales, whatever the motivation, is not an optimal use of the available resources and delays in locating and assisting the subject can result. Returning to Wartes (1975b), the treatment of POS was still based solely on his earlier POD "experiments" and analyses that did not include the concept of sweep width. Without this essential element of search theory, Wartes (1974, 1975b) was never able to achieve a set of consistent conclusions and methods compatible with the scientific search theory literature.

3.4.2 "Sector Ladder"

In brief, the "Sector Ladder" approach ranks search segments in order of "priority" or "relative importance," which the authors considered, "...no different than POA..." (Perkins & Roberts, 1994, pp. 3 & 6). Resources are then assigned to these segments starting at the top of the ladder (e.g., the highest priority). However, no guidance was provided regarding the resource levels assigned to individual segments. When the highest priority segment has been searched, it goes to the bottom of the list and is not searched again until all other segments have been searched. However, the authors stated, "...it is possible to allocate searchers to a sector which is not at the

top of the ladder, but you should only do this if you have a very good reason for doing so...” (Perkins & Roberts, 1994, p. 9). In short, this process is completely subjective. Since computation is being avoided, no updating of segment POA values is done. The result is that one cursory search of a high priority segment without success may prevent that segment from being re-searched for a considerable period of time, even if its Bayesian adjusted post-search POA would have left it at the top of a POA-ranked list.

3.4.3 Information Theory

Since the mid-twentieth century when the scientific theories of information and search first became subjects of interest, attempts have been made to apply information theory to problems of search. However, the fundamental differences in the subject matter of information and search theories make them essentially incompatible (Koopman, 1967). In addressing this specific issue, the man widely considered the father of search theory, B.O. Koopman, said that attempts to apply information theory to problems of search, “...have proved disappointing; neither the formulas nor the concepts of [information] theory have found a place in clarifying the problems of [search theory]” (Koopman, 1967, p. 1). In describing why this was true, the author said:

It has seemed to the present author that this fact is a natural consequence of a fundamental difference in the subject-matter of the two theories: in search, geometry (in the sense of positions, distance, areas, etc.) is an essential factor of the operation—in the elementary act of detection is to select a position and look near it. In the classical theory of information, on the other hand, no attention is paid to such metric matters, the ideas being confined to dichotomies: the elementary act is to ascertain in which of two subsets of a given set (e.g., of states of a system) the actual object (or state) belongs; and the geometrical shape or extent of the subsets has no necessary connection with the operation (Koopman, 1967, p. 1).

Recall also that Benkoski, et al (1991) chose to exclude both coding (information) theory and binary search tree types of problems from their survey as being inappropriate for their purposes. The present authors found their purposes completely consistent with SAR searching and agree that binary search and coding/information theory problems are not valid analogs for SAR searches.

In spite of these widely accepted conclusions, Hill (1995) chose to assert that information theory can be applied to search problems. In later work reiterating the same idea, Hill (1997) suggested that:

- Lost persons often leave detectable clues along their path of travel.
- The science of information theory tells us that if several clues are present in an area, then even a low POD search has a very good chance of finding at least one of them.
- *“Conclusion: Clue-sensitive resources, even when working a segment under conditions yielding a low POD, should be able to come up with at least one or two good clues if any clues exist”* (Hill, 1997, p. 132).

Hill (1997) also states, “If the resource is potentially able to detect **ten** clues existing in a segment, rather than merely one, there is only the remotest chance that *all* ten clues will be missed (see Figure 15.1)” (p. 131). Hill’s (1997) Figure 15.1 is a table of probabilities based on the binomial distribution that purports to represent the “probabilities of detecting clues as a function of the number of detectable clues” (p. 132). According to this table, if an area is searched with a 50% POD for finding one clue if it is present, then the probability of finding at least one of ten such clues, if there are ten in the segment, rises to 99%. In his earlier article, Hill (1995) states,

By interpreting the search for a lost person as a problem in reducing uncertainty, information theory suggests that ‘the best way to find someone’ (in addition to applying binary search methods, as Kelley suggested) is to search *each* plausible segment once with clue-sensitive resources employing relatively low POD tactics before searching any segment a second time. That is, rather than “boosting” the cumulative POD of a favorite high-priority segment by re-searching it, we can reduce much more uncertainty by searching another plausible segment for the first time. Indeed, not having found clues in a segment we originally thought should have contained plenty of clues is extremely important information, the gist of which is to ‘go look somewhere else!’ (p. 6).

There are several problems with Hill’s (1995, 1997) concept of applying information theory results to SAR searches, and all stem from the fact that it is based on several false premises. First of all, it is not possible to estimate, in any objective way, how many “detectable clues” might be present in a segment. So, the utility of the concept comes immediately into question. Second, the premise from information theory is that “errors” in a digital “message” consisting of some known number of bits being transmitted through a “noisy” channel will be completely independent of one another and uniformly distributed throughout the “message.” Clues, if present, will not be independent of one another as the same “agent” (the subject) left them only along a one-dimensional path of travel passing through a two-dimensional area. It is also highly unlikely that the clues will be uniformly distributed over the segment’s area, even in the unlikely event that the subject’s path of travel meandered uniformly throughout the segment. The only chance for the information theory analogy to have any validity at all is if the searcher happens to precisely follow the subject’s path – a possibility with a truly infinitesimal probability unless the “searcher” is actually a tracker. However, tracking is a different kind of problem more related to synthesizing data from “multiple contacts” to localize the search object than to making the initial detection. Third, clues that are likely to be the most numerous because they are natural consequences of the subject’s passage (e.g. footprints, broken twigs, etc.) are also among the most difficult to detect. A 50% POD for finding one such clue if only one exists is, in all probability, extremely optimistic. Finally there is Koopman’s (1967) observation quoted above.

It may be a good tactic for pre-deployment briefings to tell searchers, as an incentive, that they should not come back empty-handed but should find at least some clue left by the subject. However, as a way to either plan searches or evaluate negative search results, Hill’s table of PODs as a function of the number of detectable clues is dangerous. Hill (1998) acknowledged that clues are rarely found in practice, even when later information indicates they were almost certainly present during the search. He even demonstrated, with photographic slides, the ease with which clues can be missed.

3.4.4 POD as a Goal Used in Lieu of POS

Many land search management texts incorrectly recommend or imply that the goal search managers should strive for is maximizing the cumulative probability of detection (POD_{cum}) throughout the entire search area. They do not recognize that maximizing the cumulative probability of success (POS_{cum}) in the minimum amount of time is the valid optimization criterion.

The land search literature reflects a widespread belief that POD is the correct measure of search effectiveness, not POS. This has led to a variety of attempts to find tactics that would increase the POD achieved by a given level of effort, without actually having a function that related POD to effort expenditure. Since misinterpretations of Wartes (1974) statements and Colwell's (1992) data seemed to show that two successive poor searches produced significantly higher cumulative PODs for the same expenditure of effort than one good one, they were seized upon and implemented. As we have seen, this notion is incorrect. However, the goal of attaining an 80% cumulative POD in all segments, even those that are very unlikely to contain the subject, is still firmly entrenched in the present culture and is used to allocate effort via successive low-coverage searches (Stoffel, 2001; Hill, 1997; LaValla, 1997).

Dougher, et al. (2001) suggest that POD targets be used to, "search an area to a particular POD in a given time frame" and that POD be used to, "...plan a search by helping to estimate manpower requirements" (p. 35). With these statements the authors are suggesting that POD goals be set prior to allocating effort (resources) and that POD be used in place of POS in some type of effort allocation scheme. Without the use of POS and a probability density distribution, the results of this method are likely to be significantly sub-optimal. This approach is also contrary to search theory as described by Koopman (1980) and Stone (1989) where the development of an optimal search plan maximizes the cumulative POS and requires an estimate of the probability density distribution. In short, the allocation of effort implied by the stated POD goals do not produce an optimal search plan in the general case.

Dougher (2001), a book with similar content and contributors as Dougher et al. (2001), also suggests that POD goals be set by search planners. The author's statement, "A resource stating they searched a unit to a POD of 50%..." (p. 7-2), suggests that searchers are being asked directly for an estimate of POD. This method unfairly requires searchers to provide a purely subjective analysis and measure of their own detection performance when the only actual detection information searchers can reliably report is whether they found something or not.

This exclusively subjective manner of POD estimation in the land search community is demonstrated by the single question asked of searchers after a search: "If there were 10 clues of varying size [and detectability] in the area you were assigned to search, how many would you have found?" (Cooper et al., 1996, p. 296). A response of "5" would be converted directly to a POD of 50%. Unfortunately, nearly all of the land search authors suggest POD be estimated in a similar fashion. The use of one POD to represent the chances of finding any object or objects, regardless of their characteristics or the characteristics of the sensor and environment, is also a universal problem in the land search literature.

Robe & Frost (2002) addressed this widespread problem when they stated,

POD estimates should be based on objective measures and observations rather than on intuitive and therefore highly subjective assessments by either the search planner or the searchers. ... A searcher is generally a reliable source of information on the search environment experienced during the search and his/her physical condition, fatigue, level of training and experience that bear on the searcher's capabilities, etc. However, at the end of the day, the only direct detection information the searcher can reliably report is what objects, if any, they detected and approximately where and when they were detected. Searchers should be required to report only what they can observe; search planners and managers should estimate POD values based on those observations and the results of [sweep width] experiments... (pp. 2-3).

Dougher (2001), Dougher et al. (2001), Hill (1997), Stoffel (2001), and LaValla et al. (1997), among many others, define and describe what appears to be a widely accepted concept in the land search literature: "critical separation." This concept appears to be the quintessential derivation of the "POD is a function of searcher spacing" idea originated by Wartes (1974); and, contrary to the science of search theory, describes POD as having a simple, linear relationship to the distance a search object is from the searcher (sensor). Although the author (Perkins, 1989) claims that the concept is based on "research," only anecdotal evidence and personal "experience" were included in the concept's defining document (Perkins, 1989). No documentation of experimental techniques, no data analysis methods to support the author's assertions, and no serious studies have been found to substantiate the claims made regarding the concept. "Critical separation" also seems to be based on a fear of wasting effort by way of "visual overlap." Given that even Perkins (1989) estimate of a linear lateral range function produces decreasing POD values as lateral distance from the searcher's track increases, if reasonable POD levels near the region midway between adjacent searchers is to reach any worthwhile level, some visual overlap is necessary. This is not accounted for nor is the danger of leaving uncovered strips between searchers that may have to be covered later.

In contrast, there is a large body of published and critically reviewed (in the sense of rigorous scientific/academic review) knowledge about how sensors perform, including the unaided human eye, that contradicts the "critical separation" concept. Although the concept may have some practical operational use, what little has been written on the subject describes an overall concept that is mathematically incompatible with the science of search theory.

Dougher (2001) includes the terms "POA" and "POD" in the glossary, but does not use them in search planning. The same author does not even mention the term "POS," but describes "POD targets" (p. 7-2) and POD as a value that resources estimate directly.

Other criteria used to allocate effort include POA and Pden. Some authors recommend prioritizing the segments in order of POA, and allocating effort to them in that order until the effort runs out. Others recommend prioritizing based on Pden. Colwell (1994) comes closest to a solution by prioritizing segments according to the ratio of the estimated achievable POS to the number of searcher-hours required, including transit times. A significant problem with Colwell's method is that it requires the search manager to set rather arbitrary POD goals and use extremely subjective estimates on how much effort will be required to meet those goals in each segment. None of the methods objectively deal with situations where detection of the subject is much more difficult in one place than another, or situations where searchers can move through one segment more

quickly than another and still be effective. For the simpler problems (stationary search objects, no state changes), it turns out that the three factors determining where the effort should be placed are probability density, sweep width and searcher speed, or, combining the latter two factors, probability density and effective search rate. Frost (1999d) contains some alternative effort allocations for two different hypothetical situations and compares them with the optimal allocation computed using the Charnes-Cooper (1958) algorithm as described in Stone (1989). In the examples given (reproduced, here, in Appendix 4), none of the allocations driven solely by POA or Pden perform optimally. In fact, they do not even consistently perform as well or better than simply splitting the effort equally among the regions without regard to size, POA or Pden.

Colwell (1994) tries to include the logistics costs of transits to and from different segments in determining where to place the effort. This is indeed a real-world consideration that cannot be ignored in any practical search planning method. The time searchers spend in transit is time they cannot spend searching. Therefore, for the same number of available searcher-hours, more effort is available for searching nearby segments than distant ones. This significantly adds to the search manager's dilemma concerning effort allocation.

It is widely known that, while involved in a search, most land search managers do not perform a consensus, devise a probability distribution, adjust POA values, or compute POS. Anecdotal evidence suggests that this is likely because (a) they do not know how to do these things or use the results, (b) they are too much of a burden, or (c) they are successful in their search before such things are done or needed. Some land search management courses are actually built around the philosophy that most searches (>90%) take less than 24 hours to resolve and so more advanced techniques (e.g., adjusting POA, computing POS, etc.) do not need to be used or taught (Dougher, 2001; Dougher et al., 2001). While teaching people to resolve the more common types of searches is understandable, there is little evidence that courses and methods are being developed to address the remainder of the land searches where success is more difficult to achieve. It seems reasonable that these longer, less frequent searches (<10%?) are precisely the ones that would cause the most problems for everyone involved (e.g., fatigue, injuries, deaths, legal challenges, expense, etc.). Perhaps this is a good reason to somehow include them—or at least not ignore them—in any comprehensive guidance for managing land searches.

Another misconception about cumulative POS stems from the way segment POA values are updated to account for negative search results. It was initially recommended that search managers update segment POAs immediately after every search sortie (Syrotuck, 1975). Computing re-normalized Bayesian updates of segment POAs is very cumbersome and unnecessary for practical purposes. This was initially done by the usual method of multiplying each POA value by one minus the estimated POD for the corresponding segment, summing all the resulting POAs to get the normalization factor, and then dividing all these “un-normalized” POA values by this factor to get back to a total sum over the search area of 1.0 (100%). For any significant number of segments, this was a burdensome task many search managers felt they could (and often did) do without. Shea (1988) published an article describing a way to lighten the computational burden somewhat. He simply recommended that instead of updating the POAs after every sortie, just track each segment's cumulative POD over several sorties (something some search managers felt they had to do anyway) and update and re-normalize the POA values less frequently. Unfortunately, he did not make it clear in his article that the “cumulative POD” he was describing was only for the searches done since the last POA update and *not* the cumulative POD since the searching began. The latter definition was the one familiar to most search managers. Further-

more, a critical table in his article was printed with some incorrect values and an incorrect column label. The final Bayesian-adjusted and re-normalized POA (he called it “shifted” POA) column was mislabeled as POS. Errata were subsequently published. A more detailed discussion of adjusting POA issues in land search literature is included below in section 3.2.3, *Normalizing Adjusted POA Values*.

Finally, in spite of the fact that the land search literature addressed the issue of adjusting POA, no method of computing cumulative POS was provided since this cannot be done using updated, re-normalized POA values. The cumulative POS has to be tracked separately using either un-normalized values or by applying cumulative PODs for all searching done to date to the original, initial POA values. Alternatively, the re-normalization step can simply be removed from the process, greatly reducing the computational burden and making cumulative POS more directly available. This technique was only recently brought to the attention of the land search community and described by Cooper (2000). In any case, Bownds et al. (1991a-c, 1992) felt compelled to denounce POS because they thought it was being computed as the product of an updated, re-normalized POA and the cumulative POD to date—clearly an incorrect procedure. Besides, they reasoned,

What is the point of computing the “probability that the last conducted search was a success?” If the victim was found, it was a success, and, if the victim was not found, it was not a success no matter what the “after search” POS turns out to be! Thus POS is of no use after the search (Bownds et al., 1991c, p. 14).

Since the authors were computing individual “operational period” POS values using re-normalized POA values as input, the numbers were indeed good only for planning purposes, as they go on to explain. However, this unfortunate statement by someone with academic credentials (Ph.D. in mathematics), gave the false impression that even a properly computed cumulative POS value was useless as an evaluation and decision-making tool. So, Bownds et al. (1991a-c, 1992) developed a new method for evaluating the cumulative effectiveness (see section 3.4.5, “*Rest of the World [ROW]*,” below).

3.4.5 “Rest of the World” (ROW)

Bownds et al. (1991b-c) invented an extra “segment” called “rest of the world” or ROW rather than compute cumulative POS (Dr. David Lovelock, personal communication, 3 August 1998). This “segment” has no boundaries and cannot be searched. However, in order to perform its function, it has to be “primed” with at least a small initial POA value. As the physical segments on the ground are searched and their POAs updated, and all segments, including ROW, are re-normalized, the ROW’s POA gradually grows (but never shrinks). Hence, it has come to be believed that the most efficient search is the one that produces the largest growth in the ROW’s POA. In a sense this is true. Bownds et al. (1991b) even flatly state that maximizing the POS for a given operational period (based on updated, re-normalized POAs and estimated PODs) is equivalent to maximizing the increase in the ROW’s POA for that period.

Employing ROW is called “open system” search planning as opposed to a “closed system” where all 100% of the initial POA is tied up in physical segments (Bownds et al., 1991b). Stone

(1989) referred to an open system as a “defective distribution” or a distribution where the, “...probabilities sum to a number less than 1” (p. 19). In a closed system, if the POA values are updated and re-normalized at every step, the apparent total POA will remain at 100%. That is, this process is conditioned on the certainty that the search object is somewhere in the search area and nowhere else. As searching continues, segment POA values may rise and fall depending on which ones were searched, but without cumulative POS, there is no single, quantitative, criterion for determining how effective the searching has been to date. There is no indicator of when further searching in the area will almost certainly prove fruitless and it is time to either suspend search operations or find another place to look. Bownds et al. apparently thought that opening a small “leak” in the closed system’s POA “balloon” and pumping out a little probability at a time using the re-normalization process as the “pump” (primed with the initial ROW POA) was the way to break the endless cycle of the “closed system.” That the process just described is *a* way to break the cycle is beyond doubt. Whether it is the best way is open to question. When it came to actually solving the optimal effort allocation problem with a computer, Bownds, et al. (1991b-c) took an interesting turn.

One of the outcomes of the work of Bownds et al. (1991a-c, 1992) was a program called CASIE (Computer Aided Search Information Exchange). The current version is called CASIE III (www.math.arizona.edu/~dsl/casie/casie.htm). CASIE III was originally modeled after how the Pima County Sheriff’s Office (Arizona) practiced SAR at the time—a local solution to a general problem. So, the developers had to match to a large degree the agency’s intuitive process or they would not use the tool (Dr. David Lovelock, personal communication, 3 August 1998).

CASIE III accepts as input a list of segments and their POA values, a list of available search resources, and, for each resource, a list of estimated (subjective) POD values that the resource would be expected to achieve in each segment if applied there during the next operational period. CASIE III then exhaustively evaluates all possible combinations of segment POA values and resource-segment POD values to determine which will produce the highest POS. This is a cumbersome and inefficient method for solving the problem. In fact, it is an exponential solution to a problem that can be solved in linear time (e.g., by using Charnes-Cooper, 1958). Therefore, the numbers of segments and resources it can handle in a reasonable amount of running time are very limited. However, it is guaranteed to produce the correct result if given correct POD inputs and enough time.

Not only does CASIE III use POS as the optimization criterion, the source code shows it explicitly looks for the maximum (and correctly computed using cumulative POD applied to initial POA values) *cumulative POS* even though this information is never provided to the user. Instead, only the cumulative POD and updated, re-normalized POA values, including ROW, are provided after the actual search results are entered.

CASIE III allows the user to distribute the initial ROW POA among new segments and ROW. This means that if the initial ROW POA was 10%, there is no more than 10% of the original POA available for distribution. It would seem more prudent to re-evaluate all the available evidence and data and possibly re-assign initial POA values as a result (i.e., do a new consensus) if expansion of the search area is under consideration, rather than just apportioning a small amount of initial ROW POA over new segments. In this context, CASIE III does not appear to allow the user to do a new consensus without starting over and re-entering all of the data entered so far. On the other hand, as long as un-normalized POA values are being tracked and boundaries of

previously established segments remain the same, it is very easy to jump from a new consensus to the new current state in just one step. It does not even take a computer and can be done efficiently with a stubby pencil if necessary (although having a computer do this work would certainly be convenient). In short, it is recommended that expansion of the search area be accompanied by a complete re-evaluation of all available information and data obtained to date; and, an appropriate re-distribution of “initial” POA values be computed if, and as, indicated by such re-evaluation.

ROW seems to be capable of having at least three different definitions:

1. The probability that the subject is lost and in need of assistance in the vicinity of, but not inside, the current search area in some location where he could be found and aided by local search resources if the search area were appropriately expanded.
2. The probability that the subject is not lost in the vicinity of the search area and is not in need of assistance (at least not the kind of assistance SAR personnel can provide) and is in fact in a location where it is impossible or inappropriate to send local search resources.
3. Both 1. and 2. above.

The search planner, as opposed to the incident coordinator (even though the same person might be “wearing both hats”) is concerned only with option 1. He or she really does not need ROW, certainly not as defined in either 2 or 3 above, especially if indicators like fading un-normalized POA values and properly computed overall cumulative POS values are available. These indicators will motivate the planner and incident coordinator to consider other alternatives, expanding the search area, or suspending active search pending further developments.

There is also a fundamental statistical flaw in CASIE’s use of ROW. Consider the following:

Suppose the incident manager is told that the subject has been found but is not told where. How can the chances that the subject was found in ROW be computed? This is a perfectly valid textbook problem to pose (even if it has no practical value) and it is a classic example of how Bayesian inference is often employed. If the procedure for answering this question is not obvious, then consider the following statement of Bayes’ Rule:

If there are m possible mutually exclusive events e_1, e_2, \dots, e_m , exactly one of which must occur, and there are n possible mutually exclusive outcomes o_1, o_2, \dots, o_n , exactly one of which must occur, then given that the j^{th} outcome o_j has occurred, the probability that the i^{th} event e_i was the antecedent of outcome o_j is given by Bayes’ Rule:

$$P(e_i | o_j) = \frac{P(e_i) \times P(o_j | e_i)}{\sum_{k=1}^m [P(e_k) \times P(o_j | e_k)]}$$

Our simple textbook problem cannot be answered because we are missing one of the possible outcomes – that of finding the subject somewhere outside the search area. While we may have covered all possible antecedent events by inventing ROW and assigning it a POA, we have not

covered all possible outcomes. As a result, we do not have a conditional probability for “detecting” the subject given that the subject is in ROW. Hence, the thing that is missing is a “POD” for ROW. There is usually a non-trivial probability of locating the subject if he is in ROW through investigation, public service announcements, etc. If the statistical picture were complete, it would be quite easy to compute the answer to our simple textbook problem. However, if we had a way of estimating “POD” for ROW, then most of the justification for its use in CASIE would evaporate.

The ROW concept as represented in Bownds (1991b-c) and in CASIE III does not admit the possibility of finding the subject in ROW. The use of ROW POA as a one-way “safety valve” to bleed off probability, get out of the endless Bayesian update loop of a “closed” system, and avoid using POS-related values is actually a bit of statistical sleight of hand because the necessary conditions for a true Bayesian update have not been met since at least one possible outcome is not covered. A more detailed treatment of Bayesian updating and its relationship to ROW is included in Appendix B.

As a concept, ROW has a number of shortcomings. First of all, it is not clear exactly what ROW is intended to represent. It could be interpreted as part of the “search” scenario where the subject is lost and in need of assistance, but the search area may not be quite large enough to contain the subject’s actual location. Or, it could just as easily be interpreted to include all possible alternative scenarios, including those where the subject is not actually in distress but just had a change of plans and did not bother to inform anyone. Or, it could be interpreted as both at the same time. Second, cumulative POS answers the question, “If the subject is in the search area, what are the chances that all the searching to date would have found him or her?” This is a very useful thing to know. On the other hand, the value of ROW’s POA does not seem to measure any useful quantity. Third, ROW is starting to cause some real confusion. One group in Alberta, Canada, is recommending a starting value of 20% for the ROW POA. This is based on research of historical records that indicates subjects are found outside the initial search area 20% of the time (Dr. Edward H. Cornell, personal communication, 11 December 1997). It does not seem to have occurred to these researchers that perhaps the method being used to initially establish the search area may be flawed, or in need of an empirically determined safety factor to increase the size of the initial estimate. Fourth, the description of the concept in Bownds (1991b-c) overlooked a fundamental statistical flaw because the necessary conditions for a true Bayesian update have not been met.

In terms of its compatibility with search theory, the general concept of ROW may be useful during pre-operational search planning (e.g., determining if and where to conduct a search). But, once operations commence, ROW has no quantitative value. Although some have suggested that the idea of “maximizing ROW” might be used in lieu of POS, there seems to be no justification for doing so. The overall objective of search planning is to maximize POS, which will lead to the same effort allocation as maximizing ROW plus it will provide a valid measure of search effectiveness to date.

3.4.6 Conclusions

The ultimate problem, and all others pale in comparison, is that the land SAR community has had no standard method for relating probability of detection to effort density. That is, POD estimates are not based on any estimates of “detectability” (sweep width) or any detection function that relates POD to the level of effort and size of the area over which it was expended. A land search manager could easily assign two equally matched search teams for equal times to two segments equivalent in all respects except that they were significantly different in size. Then, according to current (highly subjective) procedures, the same POD could be assigned to both without anyone even being conscious of the inconsistency. Similarly, substantially different levels of effort could be assigned to two identical segments with similar results—same POD value despite the difference in the two effort levels without any conscious realization of the inconsistency. This situation is the reason all previous efforts by land SAR practitioners to develop valid optimal effort allocation procedures have failed. It has also led to a plethora of misconceptions about search tactics and some highly questionable suggestions about how searches should be conducted.

The term “POS” is described in much of the land search literature but it is not used in any meaningful way in either the literature or in actual land search planning. There may be several reasons for this. First, there appears to be a limited understanding of POS and its use in search planning. This is illustrated by the development of the concept of “maximizing ROW” as a replacement for POS. Those who believe POS is not a valid method of allocating effort invented this concept. Unfortunately, much of the land search literature parrots this view and thus POS has been generally relegated to the category of “interesting but not useful.” Second, many land search authors believe that maximizing POS is not a valid method of allocating effort. Instead, they attempt to maximize POD as if it were the correct measure. This misinterpretation of the science of search theory is likely due to the complete absence of a relationship between POD and effort density in the land search literature. Without the concept of search effort, the concept of effort allocation in general is difficult to comprehend, and even more difficult to apply.

Simple effort allocation schemes that recommend the order in which segments should be searched based on POA, Pden, or position in a purely subjective list like the “sector ladder” should be discarded. The question is not a matter of the order in which the segments should be searched. The question is how the available effort should be divided among the regions of probability during the next operational period. Even simple schemes for apportioning the effort in proportion to POA or Pden are not sufficient. The optimal effort allocation problem is not one that can be reliably addressed by simple intuitive means. In fact, the results of proven optimal effort allocation methods are often quite non-intuitive.

There is currently no effort allocation guidance present in the land SAR literature that is compatible with search theory, but such guidance built on a sound scientific foundation is sorely needed.

It has been clear for decades that information theory does not apply to the search problem for at least two reasons: there is no way to establish an objective estimate of the number of detectable clues in a segment; and unlike errors in a digital message, clues will not be independent of one

another or uniformly distributed. The notion that information theory results may be productively applied to SAR search planning should be discarded.

3.5 Land Search Planning “Methodology”

There is no single, coherent, consistent, and accepted search planning methodology published in the land search literature. However, there are several examples of suggested search action and planning sequences that warrant mention and discussion.

3.5.1 Hill, 1997

The following guidance is quoted from Hill (1997, pp. 66-67):

What follows is a recommended sequence of actions which will normally ensure the development of an efficient search plan. ...

- Develop the subject profile.
- Locate the PLS [point last seen] or LKP [last known position].
- Define the [overall] search area and define its boundaries. Consider theoretical, statistical, deductive and subjective methods, search objectives, subject profile and lost person behavior data.
- Confine the subject. That is, search the edges of the search area and take steps to ensure the subject does not cross the boundaries [undetected].
- Segment the search area into manageable units.
- Determine POA [probability of area or containment] for each segment.
- Determine the approximate size of each segment.
- Determine the Probability Density. Probability density (PDEN) is the probability that the lost person is in a given segment, divided by the size of that segment ($PDEN = POA/Size$). ... With all other factors equal (e.g. terrain, segment accessibility), higher priority should be given to segments with the highest PDEN.
- Establish incident objectives (see above)
- Evaluate resources. Work up an estimate of the total resources needed to achieve the incident objectives.
- Hold a planning meeting.
- Outline and review the Incident Action Plan.
- Allocate and brief resources.
- Debrief searchers as taskings are completed.
- Calculate changing probabilities as segments are searched.
- Brief General Staff personnel for the next operational period.

The “theoretical ... subjective methods” cited in the third step are described in nearly all of the land search references including LaValla et al. (1997), Stoffel (2001), Hill (1997), and Dougher

et al. (2001) as methods of establishing the search area. The following is quoted from LaValla et al. (1997, p. 119):

1. **Theoretical.** Distance that the subject could have traveled in the time elapsed.
2. **Statistical.** Information which reflects the distances other subjects have traveled given similar circumstances.
3. **Subjective.** Evaluation by the Search Manager of the limiting factors that exist for the specific incident and geographic location.
4. **Deductive Reasoning.** Methodical step by step analysis of circumstances surrounding the loss of the subject. Going from the general to the specific.

The example of “incident objectives” given by Hill (1997) and referenced above is quoted below (p. 64).

In this fictional example, the subject was reported missing at 2200 hrs on New Year’s Eve (yes, these things do happen!). Assume that the search team arrived at the scene within an hour and that the objectives are formulated at 2300 hrs.

1. Find John Doe by 0600, Jan 1. [Within 8 hours of the incident.]
2. Obtain reliable searching and planning data (clothing, equipment, trip plans, circumstances of loss [incident]) by 2400 hrs, Dec 31.
3. Establish his direction of travel by 2400 hrs, Dec 31.
4. Confine him to an area within 3 miles of the PLS, by 0200 hrs, Jan 1.
5. Obtain POD’s of 50% for all trails, drainages, and water edges in the search area by 0300 hrs, Jan 1.
6. Identify and obtain POD’s of 90% for all cabins, camps, and other likely spots, by 0300 hrs, Jan 1.
7. Obtain POD’s of 50% in segments 1, 2, 3, and 4, by 0500 hrs, Jan 1.

The “incident action plan” referenced above is a term from the Incident Command System (ICS) used to describe a general outline or strategy that covers how a response to an incident will be conducted. “Incident objectives” are often a part of the incident action plan. The ICS is an all risk incident management system designed to organize and manage all the functions involved in any type of emergency incident. The ICS is used across the United States, many parts of North America, and some other parts of the world to manage a wide range of civilian and military resources and events.

We will now go through the steps listed in Hill’s (1997) recommended sequence of planning actions, examining them both individually and as they relate to other actions in the sequence. Many of the details of this examination apply to other land search references and in some cases will carry us rather far afield. The reader’s indulgence is requested.

3.5.1.1 Initial Planning Actions

Obtaining pertinent information about the lost person, determining when and where the person was last seen, and defining the approximate size and location of the area where the subject could

be, assuming he or she is in fact lost and in need of assistance, are all important early steps. Reducing the likelihood that the subject can leave the search area undetected, is also appropriate.

3.5.1.2 Developing the Initial Probability Density Distribution

The fifth, sixth, seventh and eighth steps segment the search area, determine segment POA values, estimate segment sizes (areas) and compute segment probability density values. Although not explicitly stated, the ultimate goal of these actions is the establishment of a probability density distribution that reflects and quantifies all the available information about the subject's possible and probable locations.

Since the development of segments is based on logistical and operational constraints, and a probability density distribution is developed based on where the search object is more or less likely to be, there is a fundamental disconnect between what Hill (1997) is trying to do (develop a probability density distribution) and what he is trying to do it with (segments). Because segment boundaries have nothing to do with the location of the search object, using them as a basis for the distribution of probability cannot work. See section 3.2, *Probability Density Distributions*, for more information.

To solve an optimal search problem, one must first obtain, create, compute, or assume a probability density distribution representing where the search object is more likely and less likely to be (Koopman, 1980). Various algorithms may then be used to determine the optimal allocation of the available effort. If the probability density distribution is not constructed with great care so that it faithfully represents the implications of the available data, no effort allocation process can consistently provide optimal search plans. This is an example of the infamous "GIGO" principle – garbage in, garbage out.

3.5.1.3 Establishing Incident Objectives and Evaluating Resources

In this section we will review the objectives of the example quoted above to see how search theory may be applied to them and whether the problem should be restated with different objectives. We will also discover why optimal effort allocation in land search, particularly for searches employing teams on the ground (as opposed to searches using aircraft), has been such an elusive goal.

In short, a survey of the existing manuals and supporting papers indicates search theory has never actually been applied to the land SAR search problem. Even though some actually claim to be discussing "search theory," they have little in common with the texts and papers on the subject found in the scientific literature. Instead, the inland SAR publications generally record the attempts of authors not familiar with the science of search theory to re-create search theory in their own vernacular. These attempts have created a substantial number of now widespread misconceptions and have a serious potential for encouraging poor search management decisions.

3.5.1.3.1 *Setting "Deadlines"*

All of the objectives listed in the example have associated completion time goals or deadlines. These have obvious general management utility. However, search theory also directly addresses some of them. The best example is the very first objective. Stated in search theoretical terms,

this objective consists of developing a T -optimal search plan where T is 0600, January 1st. That is, the search manager has decided to develop a search plan that maximizes the probability of successfully finding the subject (POS) by 0600, January 1st. The next objective in the example is a deadline the search planner has set for obtaining pertinent case-specific information needed to plan the search. The third objective could be a combined data collection and optimal search goal since some searching for identifiable (as belonging to the subject) signs of the subject's passage may be needed to establish the direction of travel. The fourth objective may or may not be feasible. By 0200, the subject will have been missing for at least four hours and may well have traveled more than 3 miles. Only if the signs of the subject crossing the search area perimeter are unavoidable and so obvious that detection and identification are guaranteed (such as footprints in virgin snow with a unique tread mark or other identifiable feature belonging to the subject) will this goal be attainable by a perimeter search.

3.5.1.3.2 *Setting POD Objectives*

The first objective, discussed above, was very clear. Even if resources were insufficient to guarantee detection by the stated time, the obvious intent was to maximize the cumulative probability of success (POS_{cum}) attained by 0600 on January 1st with the effort that was available. Stone (1989) and other researchers have shown what conditions are necessary and sufficient for the existence of T -optimal and uniformly optimal search plans meeting this type of objective. Furthermore, these researchers have demonstrated a number of methods for actually developing such optimal search plans.

The problem with setting POD objectives like those in the example is that they may be inconsistent with the first, overall, objective. The previous search planning action from the list given by Hill (1997) should have established the initial probability density distribution. An optimal search plan would apportion the available effort among the segments so that POS_{cum} was maximized. Applying a certain effort density or coverage in a segment, as per the optimal search plan, implies a certain POD will be achieved. In other words, for a given combination of sensor, search object and environment, there is a one-to-one relationship between POD and effort density (or coverage). The POD values given in the last three objectives may or may not be consistent with those obtained from the optimal effort allocations. In other words, the allocation of effort implied by the stated POD goals may be significantly sub-optimal. The penalty for allocating effort according to the POD objectives rather than using the optimal search plan could be substantial, significantly delaying the location and rescue of the subject.

Because the land SAR community has no standard measure of "detectability" such as sweep width or any detection function that relates POD to the level of effort density (coverage), POD in the land search literature is a purely subjective value. This ultimately is the greatest weakness in any use of POD in the land search literature. Without sweep width or some similar measure of detectability and a valid detection function, "POD" as the land search literature defines it cannot be used in any scientifically acceptable application of search theory.

3.5.2 Stoffel, 2001

As previously noted, there is no single, coherent, accepted search planning methodology in the inland arena. However, the following guidance is quoted in Stoffel (2001, pp. 191-196) and is as close as this author comes to establishing a search planning methodology:

A Logical Sequence for Planning a Search Effort

1. Initial response
2. If you are the Search Manager, begin planning immediately or assign the function early.
3. Assign someone early to track information about the status of the situation as a whole.
4. Assign someone to keep track of resource status.
5. Initially develop some plausible scenarios that could explain the missing person's disappearance.
6. Develop a subject profile and begin investigative efforts to collect additional data from the community at large.
7. Establish search objectives.
8. Determine regions on the map if probability distribution mapping and Regions are being used. These relate to victim location probability. This is a depiction of general regions on a map where the subject might be. (20%, 25%, 40%, etc.) This is not segmentation.
9. Segmentation and assignment of probability of area and probability density designations for each segment and/or region.
10. Maintaining proportionality.
11. Estimate total number of resources needed to achieve the search objectives.
12. Encourage input!
13. Prepare assignments.
14. Commit resources to the field and plan for operational periods.
15. Debrief.
16. Crunch numbers.
17. Brief.

This sequence is a mix of operational, organizational, and planning elements. Because of this, it is difficult to discuss the list as a search planning methodology. So, the search planning elements will be discussed separately and in order of their appearing on the list. They will also be examined individually, as they relate to other actions in the sequence, and as they compare to Hill's (1997) recommendations.

3.5.2.1 Initial Actions

Initiating a response, beginning the planning function, and tracking the status of the situation and resources are all important early organizational and operational steps.

Although Hill's (1997) recommended sequence of actions included operational instructions early in his list (e.g., locating or identifying where the subject was last seen or known to be [PLS or LKP], and confining the movement of the lost subject), Stoffel (2001) does not. Although any sequence may be arguable and some details will always be left out of any summary, what is clear is that there is little to no consistency between the recommendations of the two authors in spite

of the fact that both claim to be describing the same thing: a sequence of actions for search planning.

3.5.2.2 Developing the Initial Probability Density Distribution

Stoffel's (2001) step five describes the development of "plausible scenarios" (p. 192). Although Hill (1997) did not include it in his recommended sequence, the author did separately describe a similar action he called "scenario analysis" (Hill, 1997, pp. 116-117). Although not explicitly stated, the ultimate goal of both of these steps is to help determine where the subject of the search is likely to be and thus impact an initial probability density distribution. The development of this probability density distribution is specifically what Stoffel (2001) is addressing—albeit vaguely—in his step eight. In this step, the author even goes as far as to differentiate "probability distribution mapping" from "segmentation," the latter of which is also used by Hill (1997) to describe how a larger search area is divided into smaller parts for ease of searching. Unfortunately, both authors require that probability be distributed to segments that have been specifically sized for searching. This is another example of the land search literature inappropriately suggesting—in contrast to Stone's (1989) recommendations—that the search area be segmented (broken into searchable segments) before a probability distribution is established.

Attempting to use segments—areas sized to be searchable and based on the resources expected to search them, instead of regions of probability—when developing a probability map is putting the "cart before the horse." Stoffel (2001) specifically suggests that, "After the search area has been segmented, a value must be assigned to each of the segments that represents the probability (POA) that the subject is in that piece of terrain" (p. 163). Koopman (1980) clearly intended the development of a probability density distribution to occur well before and separately from effort allocation (read resource application). However, the land search literature almost universally describes initial probability distribution be carried out on segments, the boundaries of which require knowledge of resource manageability and area searchability but not the probability density on search object location (Stoffel, 2001; Hill, 1997; Dougher, 2001; LaValla et al., 1997).

Although in his step nine the author implies the development of POA in regions, Stoffel (2001) also states, "At the time of this printing, the Region concept is still very controversial" (p. 164) and does not mention or describe the use of regions again. Thus, the author is affirming that the use of regions (e.g., developing a probability map based on Koopman's, 1980, concept of "Regions of Probability") is not consistently, if ever, used in land searches. Even more confusing is that the author also allows for either method to be used when he says, "It should be noted here that many prefer not to use the Regional probability mapping concept. In that case, probabilities are assigned directly to segments..." (Stoffel, 2001, p. 194).

Step ten of the list advocates the use of the Proportional Method of developing initial POA values. However, the author unfortunately insists that this be done only after, "...the search area has been segmented..." (Stoffel, 2001, p. 194). This is contrary to how the Proportional Method was first described for use in the land search literature (Cooper & Frost, 1999b), or how described by Wagner (1989), but follows how the inland literature in general has chosen not to describe probability (e.g., regions of probability) separate from, and prior to, segmentation. (See section 3.2, *Probability Density Distributions*, for more information)

Stoffel (2001) states, “POD tables can be used to plan a search by helping to estimate manpower requirements to search an area to a particular POD in a given time frame” (p. 153). The same significant limitations exist with setting POD objectives as described earlier for Hill (1997). The fact that land SAR community has no standard method for relating effort density to probability of detection has led to a plethora of misconceptions about search tactics and some highly questionable suggestions about how searches should be conducted.

3.5.2.3 Establishing Search Objectives and Evaluating Resources

When explaining his statement regarding “Establish[ing] Incident Objectives,” Hill (1997) asked, “What probability of detection can you accept for each of the segments in the area? (That is, how *thoroughly* will you search them?)” (p. 64). (Emphasis is that of the author’s.) This is a clear statement that the author suggests a POD target is one type of objective. Stoffel (2001) essentially quotes Hill when he asks a similar question regarding “Establish[ing] Search Objectives” (Stoffel’s step 7): “Which final Overall POS can we accept (e.g. *How thoroughly will we search*)?” (Stoffel, 2001, p. 193). (Emphasis is that of the author’s.) First, the similarity in the language is striking. Stoffel (2001), seems to have only replaced Hill’s “probability of detection” with “Overall POS” without offering further explanation. Thus, what Stoffel is suggesting is unclear because of the confusing mixture of thoroughness (POD) and POS. In the same short paragraph, the author also states, “...objectives usually prescribe a level of coverage for searching both the overall search area and the individual segments” (Stoffel, 2001, p. 193). With the addition of this second statement, the text appears to be suggesting a confusing combination of the development of POD targets (a poor method as previously described) and acceptable level of cumulative POS for both segments and the entire search area (a reasonable and encouraging suggestion). Unfortunately, no more explanation is provided anywhere in the text to describe what the author intended. He appears to use some of the right language but then does not follow up with an explanation or instructions. In terms of the science of search theory, this is a common theme throughout this document.

The author’s choice to “Establish Search Objectives” (step seven) before developing a probability distribution (step nine) or knowing the available level of effort (step eleven) differs from the order of actions described by Hill (1997). When he described “Establish[ing] Incident Objectives,” Hill (1997, p. 66) did so only after he listed determining initial POA values and computing probability density. Although the shortcomings of some of Hill’s steps have been described, the order of this particular section of his list appears to make sense in terms of search theory: determine POA, measure segment size, determine probability density, and establish objectives. The same cannot be said for the order of Stoffel’s recommendations (establish objectives, establish probability distribution if regions are being used, segment and assign POA) which seem to be contrary to Koopman’s (1980) recommendation that the development of a probability density distribution occur well before and separately from effort allocation.

Later in the document, Stoffel (2001) does briefly define probability density (Pden). However, the author appears to recommend probability density as the sole criterion for optimization when he describes its use as, “...all other factors are equal, search parties assigned to segments with the greater Pden would likely produce results more rapidly” (p. 156). Outside of the definition, this is the only place in the book where probability density is mentioned so the author’s intentions are difficult to infer. Regardless, ranking probability density as a scheme to allocate effort has been shown to produce significantly sub-optimal results (Frost, 1999d).

Stoffel (2001) also recommends that segments be searched in order of their POA values (high to low) when he states, “In effect, this [consensus] process ranks the segments in the order of priority that each should be searched” (p. 163). The description of POA values serving as a “ranking system” is clearly not how POA was intended to be used and when used as a search effort allocation scheme has been shown to produce significantly sub-optimal results (Frost, 1999d). Interestingly, Stoffel (2001) also stated that, “These calculations can be used to allocate resources in such a way to maximize the increase in overall POS (OPOS)” (p. 153). Unfortunately, the author does not describe how this is done and clearly contradicts his statements about ranking POA values.

Stoffel’s (2001) step 11 (“Estimate total number of resources...”) is nearly identical to Hill’s (1997) tenth step (“Evaluate Resources”). A few words have been altered in the former, but most are identical. Unfortunately, both authors describe the estimation of a POD target for each resource. They differ, however, in that Stoffel includes his “resource estimation” step much earlier in the list than does Hill. Although one may argue their respective positions in their respective lists, it is clear that no single search planning methodology is in use, and that Stoffel at least considered Hill’s information, and its order, before developing a very similar list.

3.5.3 Dougher, et al., 2001, & Dougher, 2001

When anything relating to search theory is mentioned in either of these documents (some specifics were described earlier in this review), the misinterpretation and misuse of search theory is consistent with the other land search literature. For example, both documents:

- Regard POS as having no value—“you either found them or you didn’t”
- Suggest that POD goals are acceptable
- Describe how POD can be directly estimated by questioning field searchers
- Describe and recommend “critical separation” as an acceptable method/concept from which POD can be directly derived
- Suggest that search tasks and scenarios be simply “prioritized” (the closest either gets to describing an optimal effort allocation scheme)
- Describe ROW and the use of an “open” system, and recommend them both
- Describe POD as a function of between-searcher spacing (a la Wartes, 1974)

Dougher et al. (2000) do not describe segmenting in any way. They describe “establishing the search area” but nothing further related to area searching. Although Dougher et al. (2001) stopped short of describing a method of establishing initial POA values, Dougher (2001) did not. The method described in the latter uses a qualitative scale ranging from “very likely” to “very unlikely,” exactly like the system suggested by Bowns et al. (1991a). However, the consensus form provided also included numerals (1 through 9) that were associated with each of the qualitative values on the “Probability Estimate Scale” (Dougher, 2001, p. 24-7). See 3.2, *Probability Density Distributions*, for more information on this author’s approach to developing initial POA values.

Even though Dougher et al. (2000) defines POD as, "...the probability or chance, usually expressed as a percentage, that a clue or Subject will be (predictive), or would have been detected (in retrospect) by the search action if a clue of the Subject was in the search area" (p. 96), the authors later suggest that POD represents "clues and/or the subject" (p. 96). All land search authors use this multiple object definition of POD in spite of the fact that POD values must necessarily vary with the detectabilities of the various objects that may be present. See section 3.3.4, *One POD for Multiple Objects*, for more details on this issue.

Although Dougher (2001) includes little detail about search theory and/or its use, the document does include a short description of "Bayes' Theorem of Subjective Probability" (p. 24-2). In this description, the author discusses the successful search for the USS Scorpion but unfortunately does not give a source of the information. In this section, "Bayes' theorem of subjective probability" is described as a, "...process to quantify the opinions of a group of knowledgeable persons" (Dougher, 2001, p. 24-2). This is offered, it seems, as evidence of the validity of the described consensus method in that the author also states, "The search consensus method described...is similar to the technique used [to find the Scorpion], and is also based on Bayes' theorem" (Dougher, 2001, p. 24-2). These claims appear to conflict with the more conventional definitions of Bayes' Theorem that describe, "...the relationships that exist within an array of simple and conditional probabilities" (Papoulis, 1984, p. 4), and with the fact that it is often used to reflect the relationships between conditional probabilities but has nothing to do with establishing the initial probability values. Although this may not be a significant point regarding search theory, it may serve as a reasonable indicator of the relative level of scientific sophistication of this particular work.

3.5.4 Cooper, 2002

Cooper (2002) developed a "Search Actions Outline" for the Mountain Rescue Council (MRC) of England and Wales. It was designed for search managers and incorporates, in the form of a brief, ten-page outline, practical recommendations regarding the organization, management, operations, and planning of a land search. The document clearly states its intention which is, "...to provide a standard methodology for land search operations" (Cooper, 2002, p. ii). The document does not include much detail; it is an outline only. It contains many of the features available in other similar lists of recommended actions such as Dougher (2001) and Dougher et al., (2001). But, in terms of search theory, it goes further to specifically describe the use of regions of probability (before segmenting) (p. 5), proportional consensus (pp. 6, 7), the development of a probability map (p. 9), the calculation (not subjective estimation) of POD and POS (p. 9), and the use of sweep width (p. 9). Although it is a brief summary of some larger body of information, it is notable that, unlike the previously described land search literature, the document does not include anything about subjectively estimating POD directly from field personnel, setting POD goals, critical separation, ROW, or POD as a function of between-searcher spacing. These distinctions alone uniquely separate this document from other land search literature in terms of its use of search theory. The basis for the outline is, currently, a privately published collection of presentation materials and handouts used in the MRC's Search Management and Planning Course. According to Dr. Anthony S.G. Jones, M.B.E., Vice Chairman of the MRC, the materials are currently under development and were developed from the existing land search literature,

practical experience, and the large body of scientific search theory literature introduced to the group less than a decade ago (personal communication, 5 September 2002).

There is nothing evident in the outline provided that contradicts the science of search theory. This appears to be the only land search document in this review that attempts to offer a search planning methodology. Unfortunately, there is also currently no evidence that it has been accepted in the United States. Based on communications with the principle developers of this material, it appears that they are aware of the many shortcomings of the current state of the land search literature, and they have been attempting to remedy the situation by better aligning their procedures and materials with the science of search theory (Dr. Anthony S.G. Jones, M.B.E., personal communications, 5 September 2002).

3.5.5 Conclusions

The land search literature contains many practical suggestions for improved search operations and management. However, land search planning methods, in the rare situation where they are described in the literature, do not follow the scientifically derived guidelines of Koopman (1946, 1980), Stone (1989), and others. Although a few recommended action sequences have been published in the land search literature, no comprehensive search planning methodology could be identified. Of the limited search planning advice offered, it is universally inconsistent, incomplete, and addresses mostly organizational and resource management issues. Where search planning is involved in this advice, it is based solely on other land search literature and universally excludes the volumes of information available on the topic of scientific search theory.

Some progress is being made to bring the science of search theory to the land search community. Robe & Frost (2002), Cooper (2002), and a number of presentations at national and international conferences (e.g., annual SARSCENE in Canada, annual NASAR conference in the U.S., biannual MRC conference in the U.K.) are attempting to integrate the science of search theory as described by Koopman (1946, 1980) and others into land search planning methods. But there are many obstacles, including the inertia of already-accepted, but not scientifically valid, concepts and methods that continue to be perpetuated through the current land SAR literature and course offerings.

4. Summary, Conclusions, and Recommendations

4.1 Summary

It is clear that much of the material being published in various inland SAR manuals, journals, conference proceedings, and private papers could be substantially improved. A new start is needed to develop a comprehensive, coherent, scientifically correct, but still practical and teachable, search planning methodology for land SAR search managers. At the heart of this effort must be research to identify the significant variables affecting *sweep width* in the inland environment, both for ground search parties and airborne searchers, and the development of aids for estimating sweep widths based on estimates or measurements of these variables in the field. While manual planning methods should be developed to the extent they can be, software tools to aid in effort allocation decisions should also be developed. Ideally, these tools would be integrated with other tools used to aid in managing the considerable logistics burden associated with search operations. Similarly, commercial off the shelf (COTS) packages for displaying maps and charts are commonly available and should be integrated into any search planning application. As the cost of GPS receivers continues to decrease and they become more widely available, search management software must have the ability to accept and display actual searcher tracks and properly evaluate the effectiveness of searches given this very valuable information.

In some ways, land SAR may be able to obtain even more benefits from the application of search theory than maritime SAR. Observe the operational constraint on aircraft searching over the ocean to fly long, straight, parallel search legs that prevents them from putting a little effort into one cell, then another, then another as an unconstrained uniformly optimal plan would have them do. In contrast, ground search managers assign separate assets to each segment and have considerable flexibility in the amount of effort allocated to any one segment at any one time. Even searching from the air has more flexibility over the ground than over water since light aircraft and small helicopters of limited endurance are assigned to a larger number of small segments than in the marine environment. In other words, the premise of infinitely divisible search effort used for unconstrained optimal search density solutions is more closely approached in the land environment than in the marine environment. Therefore the land search manager, if given the right tools, will probably be able to come much closer to actually realizing a uniformly optimal search plan than the maritime search planner. In addition, a proper application of search theory will help everyone better understand the land search problem and provide better paradigms for thinking about it. In the end, this may be the most valuable contribution made by bringing scientific search theory to bear on the land search problem.

No IAMSAR-like search planning methodology exists in the land search community. Although several authors suggest “logical” or “recommended sequences of action,” there is no consistency among the recommendations because no specific methodology or paradigm exists. In fact, the land search literature is mostly a collection of various authors’ ideas, thoughts and advice that, if followed, lead to as many different answers as there are authors. An effective analogy might be that of baking a cake. The IAMSAR Manual provides a “recipe” for combining the ingredients, including instructions on how to bake it. The land search community, on the other hand, seems to only have ingredients—and not all that are necessary—without specific instructions on how to combine them. Although there is some chance that certain combinations of the land search in-

redients will be edible, most combinations will not be from an optimal effort allocation perspective, unlike the IAMSAR recipe.

For the first time, a standard, simple, practical, and low-cost method for conducting detection experiments for ground searches was successfully developed, demonstrated and described in Robe & Frost (2002). Land SAR organizations will now be able to conduct detection experiments in their own respective areas of responsibility using their own resources to produce effective sweep width values for their own use and the use of others in similar search situations. This work constitutes a major breakthrough for improving land search planning and evaluation methods by replacing subjective estimates for POD with objective ones that are more reliable, repeatable, and accurate than current subjective techniques. This work will also make it possible to bring known and proven methods for the optimal allocation of search resources to each situation that requires areas to be searched, leading to multiple benefits. These benefits include finding survivors sooner on average and thereby saving more lives, reducing risks to searchers through reduced search times, reducing costs, reducing the time volunteers must take from their normal lives, and making resources more available for other missions if needed.

The inland SAR community, being made up of many diverse, very localized groups, is simply too fragmented and fiscally limited to mount or fund an effective research and development effort. It has long been a traditional and widely accepted role of the federal government to conduct research and provide practical advice and methodologies for activities that are in the public interest. Certainly SAR falls into that category. Fortunately NSARC has stepped into this void with funding provided by its members, most notably the Department of Defense and the Coast Guard.

4.2 Conclusions

It does not appear there has ever been a comprehensive attempt to apply the science of search theory to the development of land search techniques. Various individuals at various times have attempted to apply bits and pieces of what they believed to be search theory to the problem. There is clearly a great deal of room for improvement as search theory can make substantial contributions if properly applied. There is also a critical need to rectify some of the more crucial misunderstandings that could have a significantly detrimental effect on future inland search operations.

4.2.1 Needed Changes

When reading the portions of the land SAR literature that bear on search planning, it is clear that even though a few of the early investigators were vaguely aware of Koopman's early efforts, certain key elements of Koopman's work were overlooked. Chief among these was the concept of "effective search (or sweep) width" (Koopman 1946). This omission set in motion a chain of events whereby persons who were not professional mathematicians, analysts, or operations researchers tried to re-invent the search theory wheel without this key element. A few isolated spokes were found but the rest remained missing. Some faulty spokes were manufactured. However, a complete set of correct spokes was never assembled and no coherent rim to tie them together was ever discovered. This resulted in the problems identified in this review. The following bullets summarize the findings:

- Methods for developing initial probability density distributions should be modified in two important ways:
 - Regions of probability should be developed prior to and independently from the development of specific search assignments (segments). Regions of probability should be based only on subject behavior, terrain, weather and other factors that might affect where the subject might be located.
 - Assignment procedures for POA (POC) values must allow and require evaluators to assign values that are, in the evaluator's best judgment, in the correct relative proportions.
- POA adjustments to account for unsuccessful searching should be done without re-normalization.
- The sweep width concept should be incorporated into land SAR search planning and data to support it should be obtained through scientific sweep width experiments using procedures like those developed by Robe & Frost (2002) or refinements of them.
- The idea that POD is a function of sweep width ("detectability"), effort, and area searched (i.e., a function of Coverage) should be vigorously promoted. The notion that POD is dependent only on the spacing between adjacent searchers should be discarded. POD vs. Spacing tables, formulas, and graphs should likewise be discarded and replaced with the appropriate POD vs. Coverage function and graph.
- An appropriate detection function that objectively and quantitatively relates POD to Coverage needs to be adopted for use in land SAR search planning and evaluation. The exponential detection function $POD = 1 - e^{-C}$ is recommended. Subjective POD estimation techniques have no demonstrable validity and should be discarded.
- The notion that two successive poor searches of a segment will produce a significantly higher POD than a single search when both tactics employ the same total amount of effort needs to be discarded. It is incorrect and leads to poor effort allocation decisions.
- The view that POD is a valid measure of search effectiveness and the quantity that search planners need to maximize is incorrect and should be discarded. It should be replaced with the correct view that POS is the quantity to be maximized through the optimal allocation of the available effort.
- Hill's (1995, 1997) assumption that searching segments where multiple "clues" may be present is equivalent to the information theory problem involving detection of "errors" created by a uniformly noisy channel ("white noise") is incorrect and should be discarded.
- The correct computation and use of POS and cumulative POS for planning (effort allocation) purposes and as valuable evaluation and decision-making tools should be vigorously promoted. Effort allocation schemes based on overly simplistic rankings, such as ranking

by POA or Pden, should be discarded. Even the notion of prorating effort on the basis of such values should be discarded.

- As a tool for helping search managers keep an open mind regarding scenarios different from the one(s) on which the search is based, the ROW concept may be useful. However, as a planning tool, it has only limited utility. As a quantitative evaluation and decision-making tool it has little or no validity and cannot take the place of cumulative POS. Such uses of ROW should be discontinued and replaced with a proper utilization of POS and cumulative POS.
- There is no clear, consistent, comprehensive and scientifically sound search planning methodology available to planners of land searches at a practical level. Such a methodology needs to be developed and made widely available.

4.2.2 Required New Empirical Data and Methods of Acquiring It

There are several areas where new or improved empirical data is sorely needed. They are the areas of sweep width and lost person behavior. New empirical sweep width data is needed for inland SAR since none currently exists for ground parties and that which does exist for airborne searchers is of unknown origin and reliability. Although some research into lost person behavior has been done, it is far from enough. Research in this area should be expanded since it can make initial probability density distribution estimates much more accurate and can often significantly limit the area that must be searched.

4.2.2.1 *Sweep Width Data*

Clearly the most urgent need is for the development of meaningful sweep width values for inland SAR, especially for ground parties. These include visual sweep widths for human searchers for both subjects and common types and sizes of clues under a variety of typical environmental conditions. The need extends to obtaining, if possible, sweep widths for search dog teams and certain types of electronic sensors and sensing aids. This latter group includes both ground-borne and airborne infrared sensors, night vision goggles, thermal imaging equipment, etc. It may even be appropriate to include certain types of satellite-borne sensors since image resolutions suitable for SAR are becoming available.

As the Coast Guard has shown, extremely useful and valuable empirical sweep width data may be obtained from carefully designed experiments and equally careful analyses of the data collected in these experiments. The same concept should be applied to inland SAR, and experiments specifically tailored to obtain sweep width estimates for the various sensors, search objects, and environmental conditions encountered in inland SAR should be undertaken without delay. Robe & Frost (2002) have described a valid experimental model. Now all that is required is the application of the model and the collection and analysis of data.

Although the Coast Guard has assumed, for many years, that Koopman's (1946, 1980) hypothetical model of visual detection is valid and has also assumed they were realizing the increased POD benefits of perfectly executed parallel sweep search patterns, these assumptions are proba-

bly not appropriate for inland SAR. However, the exponential detection function, given reasonably valid sweep width data, will probably prove to be an excellent and objective estimator of POD for inland searches. In other words, if valid sweep width values can be obtained, further research into the details of how detections occur from instant to instant as a sensor approaches the search object is probably unnecessary. The exponential detection function should perform adequately without further refinement.

Empirical research into sweep widths needs to continue indefinitely, as in the maritime arena, to accommodate new sensors as they become available, changes in search object characteristics affecting detection (e.g. high efficiency laser reflective tape if it comes into widespread use along with corresponding sensor packages), etc. In short, sweep width research should not be regarded as strictly a one-time effort. Even though a substantial initial effort should be undertaken to overcome the current almost complete lack of inland sweep width data, there should be a continuing effort to maintain currency with the available technology.

4.2.2.2 *Lost Person Behavior*

This is an area of research that shows great promise for reducing the mean time to locate lost or missing persons in need of assistance. However, this topic was outside the scope of this study. Based on research by Syrotuck (1976) and presentations at SAR conferences by Hill, Heth, Cornell, and Koester (all behavioral scientists), research done so far seems to indicate somewhat predictable behavior patterns among members belonging to certain groups. For example, a lost hiker will most likely behave in one way (they seem to favor continued forward motion as opposed to retracing their steps) while an Alzheimer's patient will most likely behave in another. Age also seems to play a role, as lost children do not seem to do the same things or take the same routes a lost adult would. Typical behaviors may also be dependent on locale; and the list goes on. The benefits of "profiling" the lost person and using data obtained from many similar cases should be obvious. The resulting estimate of the probability density distribution of subject locations should be both much more accurate and much smaller than if no such profiling was used, especially in situations where individual traits or circumstances do not provide strong clues or information about them is not readily available. Such research may even be applicable in the marine environment since survivors often can, and sometimes do, significantly affect the movement of their survival craft in an attempt to save themselves, thus invalidating drift estimates. At the present time, much, perhaps all, of the research into lost person behavior seems to be unfunded.

4.2.3 Inland Search Planning Methodology

Perhaps the single most important finding in this study is the critical need for a clear, concise, coherent, comprehensive, scientifically based but still practical search planning methodology for inland SAR. While some attempts have been and are being made to integrate scientific search theory into land search planning methods (e.g., Cooper, 2000, 2002), these methods do not appear to be widely known or used in the United States in spite of the fact that the Mountain Rescue Council of England and Wales has been integrating them into its curricula for teaching land search planning since 1996 (Cooper, 2002). Research into developing an accepted methodology for inland SAR situations and environments may and probably will suggest new, improved tac-

tics for maximizing the chances of finding the subject in minimum time with the available resources.

Once such a methodology, including improved tactics, is developed, it can then become the basis for step-by-step manual “cookbook” methods and computerized search planning aids which help the search planner make the best possible effort allocation decisions. The methodology needs to be flexible enough that it can be tailored to local conditions, but rigorous enough in its approach to withstand scientific scrutiny tempered by recognition of unavoidable operational exigencies. Most of all, it needs to help search managers consistently make good effort allocation decisions. While it is impossible to compete with the flexibility of complete subjectivity, it seems clear from the material presented earlier in this study that the present methods do not satisfy the more important latter criteria. If there is any remaining doubt about the need for such a methodology, and search managers trained to use it consistently and effectively, the reader is invited to peruse Hill’s (1997) description of the Andy Warburton case, reproduced in Hill (1997, pp. 5-8).

4.2.4 Land Search Literature Reviews and Critiques

It was found that, on subjects where search theory bears, pre-publication reviews and critiques within the inland SAR community are insufficient. One does not want to discourage anyone from trying to better understand a problem and suggest better solutions. After all, it is impossible to predict where the next great idea will come from. On the other hand, allowing ideas and opinions not supported by an existing body of pertinent scientific research to be published and taken as if they were, even when the authors themselves make no such claims (although they often do), can only produce confusion, chaos, and a loss of credibility – not to mention possible loss of a life that could have been saved.

In response to National Search and Rescue Committee (NSARC) tasking, the NSARC Research and Development Working Group hosted a meeting of land search experts on March 24, 2001 in Laurel, Maryland. The meeting provided a unique opportunity for a small group of selected experts and other key persons to develop a preliminary assessment of land search planning needs. The Group quickly agreed that existing land search planning procedures, reference documents, graphs, formulas, and data should be rigorously reviewed and either validated or given valid replacements. This conclusion of the group was one reason for this report. However, the inland SAR community does not provide an ongoing method of reviewing and validating new information on land SAR. No refereed or peer-reviewed scholarly journals or publications are readily available to the land SAR community. Often the only consideration given material prior to publication is whether it sounds plausible to the editors. Much of the land search literature reflects this culture and seems to have promoted quantity over quality. Many of the references reviewed in this report describe a collection of opinions and ideas that have never been scientifically scrutinized or validated. Too often one person’s apparent good idea has quickly become land SAR gospel and taught everywhere as “state-of-the-art.” This trend should be reversed and sound methods of review and validation should be integrated into the land search development processes.

4.3 Recommendations

4.3.1 Develop a Standard Methodology for Inland Search Planning

While most of the land SAR literature describes the elements of search theory concepts (e.g., notions of assigning POA values to sub-regions of the possibility area, the notion of POD, and the fact that POS is the product of POA and POD), none of them present a comprehensive methodology for planning a search based on search theory. Guidance for how to evaluate the available data about an incident and how to combine local knowledge and general statistics (e.g., lost person behavior profiles) with the case-specific data is sparse. POA values are estimated in an almost purely subjective manner by surveying several individuals for their opinions on where they think the most likely locations are. Some of the survey techniques do not address the critical issue of ensuring the probability density values generated properly reflect the proportions that were in the evaluator's mind.

The report of the 24 March 2001 meeting of NSARC R&D Working Group (NSARC, 2001) suggested that a basic land search planning methodology should be developed, published and distributed for national, state and local authorities and volunteer groups to consider incorporating into their search planning procedures, if they deem it to be applicable to their needs. It was suggested that a guide describing the methodology could be made available through electronic means such as the Internet and CD-ROMs, and be included in the land SAR chapter of the *National Search and Rescue Supplement* to the *International Aeronautical and Maritime Search and Rescue Manual*. This will require carefully assembling a team consisting of researchers from several pertinent fields of science, a representative cross-section of the inland search operations community, and a group to manage the project and act, when and if needed, as a bridge between the "researchers" and the "operators." The guide developed from this methodology might include but not be limited to the following sections (NSARC, 2001):

- Management Advance Planning
 - Pre-search organization and coordination
 - Prevention education
 - Resource agreements
 - Resource training
 - Communication needs
 - Vulnerability
 - Hazard assessment for the local area

- Search Operations
 - Notification and resource recall
 - Case investigation/data collection
 - Resource management and deployment
 - "Hasty Search" employment
 - Extended searches
 - ✓ Using lost person behavior statistics
 - ✓ Using scenario analysis

- ✓ Using techniques for estimating effective sweep widths
 - ✓ Using guidance for allocating available resources to maximize the probability of success (POS) as quickly as possible.
- Post Search Operations
 - Search suspension guidance
 - Demobilization
 - Critique
 - Documentation/data collection and reports
 - Lessons learned

Since implementation of such a methodology will require sweep width estimates, the second recommendation below needs to be addressed concurrently.

4.3.2 Perform Sweep Width Experiments for the Inland SAR Environment

There is no detection function in use in land SAR that relates POD to the density of searching effort expended. In fact, the very concept of relating POD to effort expenditure is conspicuously absent as is the entire concept of a measure of detectability (e.g., sweep width). Estimated POD values are purely subjective and based on highly speculative answers to questions like, “If there were ten objects in your assigned search area (a.k.a. “segment”), how many do you think you would have found?” (Cooper et al., 1996).

Before any truly objective methods for estimating PODs, POSs or allocating effort in the inland environment can be implemented, some objective measure of the basic “detectability” of the search object(s) we are seeking by the sensor(s) being used under the environmental conditions prevailing at the time is needed. Such a measure does not currently exist for searches conducted on the ground. However, the basic requirements and characteristics of such experiments are discussed in great detail in Robe & Frost (2002). Refinement of these techniques and development of a “cookbook” approach for use by SAR personnel to establish sweep width values for their own situations should be vigorously pursued. The next step in this process has already been funded under the auspices of NSARC using funds provided by DoD.

4.3.3 Develop Computer-based Tools for Land SAR

The potential value of computers applied to the search problem is certainly nothing new. Nearly three decades ago, Syrotuck (1975) realized the potential of the devices when he made this statement on their use:

Detailed search plans could easily be called from the computer. Such as specific areas to search and by which resource. The time it would take and the probability of success.

A centralized Computer Search Planning System that was used by many agencies, in a short time, would gain far more “experience” than any individual contributor. However, each contributor would gain by the collective experience of all the others. ... The cost of the entire system may be more than its’ ultimate value. However, what is the value of a “life”? (Syrotuck, 1975, p. 35)

Although some tools exist, and at least one would come close to correctly performing the optimal effort allocation function if given the correct data and enough running time, it does not appear that any adequate computer-based search planning aids exist. Even for a tool like CASIE III that could find the most nearly optimal effort allocation in a large but finite set of possible resource-segment assignments, sweep width data and some significant off-line work would be required to correctly generate the needed inputs. Computerized tools based on a correct implementation of search theory principles are needed just as badly as the basic methodology itself and must be fully integrated with this improved approach. Furthermore, it should be possible to develop such tools, suitable for use on laptop computers, for only a relatively modest investment.

4.3.4 Develop Resource Allocation Guidance for Area Searches

In the land search literature, the optimal allocation of effort to maximize POS is not addressed in any concrete, useful way. This is not surprising since this cannot be done without a detection function that relates POD to effort. However, there exists a dire need for search resource allocation guidance that should be used to plan searches following an unsuccessful “hasty search.” The goal of search planning is to, “Utilize the available search resources in a way that maximizes the probability of locating the distressed person(s) in the minimum amount of time” (NSARC, 2001, p. 2). Land search planning methods and procedures must be restructured to work toward this objective and minimize the length of successful searches.

This report can be used to highlight the existing material used by the land search community that is, and is not, compatible with the science of search theory. From this, modifications should be made to the land search procedures that are incompatible with search theory and new methods and procedures developed.

Practical, technically correct procedures for using effective sweep width values with other data already generally available to produce the needed resource allocation guidance for land searches must be developed. The objective of this work would be validated procedures for using effective sweep widths and other data to maximize POS and minimize average search duration by optimally allocating the available resources. This information would be made available to the land search and rescue community and potentially become additions to publications such as the *National Search and Rescue Supplement* (NSARC, 2000) and eventually the *International Aeronautical and Maritime Search and Rescue Manual* (ICAO/IMO, 1999a-c).

4.3.5 Objectives for Future Work

1. Provide a practical probability of detection (POD) estimation procedure with worksheets, graphs and/or other appropriate job aids that is suitable for land SAR and based on proven scientific concepts. The procedure will produce objective, accurate, consistent, and reliable POD estimates, replacing current subjective techniques.
2. Based on proven scientific principles, provide a practical search planning (as opposed to search management) methodology that will allow land search planners to maximize the effectiveness of their available resources during every operational period of every search. The advantages of such optimal resource allocation are:
 - a. More successful searches.
 - b. Reduced average time and resources required for finding the search object.
 - c. More lives saved as a result of reduced average search time
 - d. Reduced risk to searchers as a result of reduced exposure time to the risks of searching.

4.3.6 Steps to Accomplish the Objectives for Future Work

4.3.6.1 Validate and Refine Data Collection and Analysis Procedures for Establishing Sweep Width Values

Refine and validate the preliminary detection data collection and analysis procedures developed and demonstrated in Robe & Frost (2002) for ground searchers.

- a. Conduct additional test demonstrations in different terrain and with different SAR groups.
- b. Develop data analysis software using readily available commercial off-the-shelf software to automate the data analysis procedures and reduce the opportunity for human-induced error.
- c. Conclude with a set of experiments in one venue to produce actual detectability (sweep width) data for realistic search objects in that venue.
- d. Final procedures and software for ground search detection experiments are to be suitable for use by SAR organizations without assistance from professional analysts or special scientific training.

The effort outlined in Robe & Frost (2002) could do no more than demonstrate whether the procedure under development shows promise of being practical for general use by SAR teams/agencies to develop search object detectability (sweep width) estimates in their respective geographic areas of responsibility (AORs). Although it has shown such promise, the procedure still requires further development and refinement. It must also be shown to be practical for use in a variety of environments by a variety of personnel. Therefore, it should be tried in several

geographic areas with differing environments. This will also provide the opportunity to involve and train more SAR personnel and prove the utility of the procedure.

Robe & Frost (2002) was of insufficient scope to actually produce reliable detectability (sweep width) data. This was known from the outset. Valid values for use in actual searches was not, and in fact could not be, a goal of that first demonstration. The experiment procedure should be exercised sufficiently to convincingly demonstrate practicality, usefulness and reliability of results.

A (full) set of experiments should be done in at least one location that will be sufficient to produce enough data for obtaining an actual sweep width estimate for a given set of conditions.

4.3.6.2 *Extend and Modify to Include Detection Data for Aerial Search from Aircraft*

Existing detection data collection and analysis procedures should be extended and modified as needed for aerial searches over land for the use of the Civil Air Patrol (CAP) and other agencies that search from the air.

- a. Conduct additional test demonstrations in different terrain and with different CAP wings to adapt and refine the sweep width experiment procedures as necessary for searches conducted from aircraft.
- b. Extend data analysis software as needed to apply to aerial search, to automate the data analysis procedures and reduce the opportunity for human-induced error.
- c. Finalize procedures and software suitable for use by SAR organizations without special scientific training assistance from professional analysts.

The sweep width data currently published in the *IAMSAR Manual* (IMO/ICAO, 1999a-c) for aerial search over land are quite limited and of uncertain origin. No supporting studies for these data have been found to date. At a minimum, these data should be validated. The above procedures for ground searchers should be expanded to make them applicable to aerial search over land where there is an even greater potential for improving search effectiveness stemming from the natural advantages of aerial search. Due to their high speed, a procedure to obtain detectability (sweep width) data for search objects on the ground when an aircraft is performing the search will necessarily involve the need to populate much larger areas with objects. This is likely to introduce some unique procedural and logistics issues to be addressed. In addition, it will be necessary to accurately track the movements of aircraft during the procedure. This tracking can easily be accomplished by means of a GPS receiver that uses a laptop computer as a data-logging device. Substantial Civil Air Patrol involvement in both the adaptation of the ground detectability (sweep width) procedure and its implementation will be highly desirable.

4.3.6.3 *Develop Improved Procedures for Estimating POD*

Develop procedures for reliably estimating probability of detection (POD) that are suitable for use in all types of ground and aerial search of areas. These procedures would be based on the detectability index (sweep width) for a given search, the amount of effort expended in a given area, and the size of the area that was covered by the search resources.

Note: A sample procedure is provided in Part IV of Robe & Frost (2002) but it may require further explanation or changes before it will be suitable for field-level users. This type of procedure will be new to the land search community. Understanding and proper use will require development and explanation of a new (to the land SAR community, but standard elsewhere) paradigm for the search problem. The explanation must be written in clear terms that are appropriate to the land search problem. This will require discussions with and assistance from recognized authorities on land search.

POD is a function of the search object's detectability (sweep width), the amount of effort expended in an area, and the size of the area where the effort is expended. There are no authoritative and generally applicable procedures, graphs, job aids, etc. currently available to the land SAR community as a whole for estimating POD based on these parameters. Currently, ground search POD estimates are either purely subjective and based on how well searchers think they could have detected the search object, or they are based on data from "experiments" done many years ago that attempted to relate POD directly to searcher spacing without any estimate of object detectability, effort expenditure, or area searched. The forms of the POD vs. Spacing graphs that resulted from those experiments, and hence a significant portion of their values, can be shown to be inconsistent with the scientifically established principles of search theory.

Therefore, there is a need to develop practical procedures, graphs, job aids, etc. to estimate POD values from sweep width data, effort estimates and area estimates.

4.3.6.4 Develop Outline for Practical Search Planning Methodology

Develop an outline for a practical search planning methodology for use in land searches involving static search objects. This project will review existing published land search planning methods to determine which are already consistent with known search theory principles and best practices, which can be modified to become consistent with known principles and best practices, and which should be discarded. It will then go on to outline a practical search planning methodology that is consistent with search theory and will allow search planners/managers to make better resource allocation decisions.

4.3.6.5 Describe Functional Requirements for Software Tools

Develop a functional description for software tools to aid the land SAR search planner. Truly optimal resource allocation is a computationally complex problem that requires computer assistance. Tracking the progress of a search, updating probabilities based on search results and allocating the next amount of available search effort optimally based on these updates are all activities where computer assistance would be extremely useful and enabling.

The goal of any search planning system is to provide a method for optimally allocating the available search resources so that the probability of success (finding the object being sought) is maximized, the expected time required to find the object is minimized, the resources are used in the most efficient manner, and the risks to search personnel are reduced (primarily through reduced exposure times). The more quickly a distressed person can be found, the more quickly lifesaving assistance can be provided. The proper deployment of the available search resources is extremely important to achieving this goal. However, finding the optimal allocation for the

available search effort is a computationally intensive and complex process better left to an appropriately programmed computer.

Requirements for developing such software that will be suitable for use by the land SAR community need to be developed and documented.

4.3.6.6 Survey Existing Software

Survey existing search management software packages and evaluate the feasibility of integrating a search planning module meeting the requirements found in step 4 above with each.

There are a few software packages available now that aid land SAR users with the search management function but provide little to no search planning or effort allocation advice. The little that does exist is merely a codification of manual procedures that have no scientific, and on close inspection often have no sound empirical, foundation. A survey of search management software currently in use should be done with a view toward determining whether and how a search planning module meeting the requirements determined by the previous task may be integrated with them.

4.3.6.7 Develop Complete Land Search Planning Support Software

Develop a complete, practical, robust land SAR search planning methodology and supporting software similar in scope and level of detail to that already available for other SAR communities.

5. References

- Abi-Zeid, I. (2001). *SARPlan progress report*. Presentation to the United States National Search and Rescue Committee (NSARC) Research and Development Working Group, 22 May 2001.
- Allais, M. (1957). Method of appraising economic prospects of mining exploration over large territories. *Management Science* 3, pp. 285-347.
- Allais, M. (1957). Method of appraising economic prospects of mining exploration over large territories. *Management Science* 3, pp. 285-347.
- Benkoski, S.J., Monticino, M.G., & Weisinger, J.R. (1991). A survey of the search theory literature. *Naval Research Logistics*, 38, pp. 469-494.
- Benkoski, S.J. (1978). *CAS operational summary, COMPATWINGSPAC memorandum 15 may*. (Secret Noform).
- Blake, B., Baynes, R. & Dougher, H. (1983). Applying shifting probabilities of area (POA) in open system searching: a new concept in search management. In Bill Wade (Ed.), *NASAR Conference Papers* (pp. 21-35). Dallas, Texas: National Association for Search and Rescue.
- Bownds, J., Lovelock, D., McHugh, C., & Wright, A. (1981). *Desert searches: The effectiveness of helicopters*. Retrieved from www.math.arizona.edu/~dsl/casie/helicop.htm on 20 February 2001.
- Bownds, J.M. (c1982). *A simple version of the Bayesian information processing for updated POA estimates*. Private Publication, Desert SAR Group, Department of Mathematics, University of Arizona.
- Bownds, J., Ebersole, M., Lovelock, D., & O'Conner, D. (1991a). Reexamining the search management function, *Response*, 10(1), pp. 12-15, 28.
- Bownds, J., Ebersole, M., Lovelock, D., & O'Conner, D. (1991b). Reexamining the search management function—part II, *Response*, 10(2), pp. 16-20.
- Bownds, J., Ebersole, M., Lovelock, D., & O'Conner, D. (1991c). Reexamining the search management function—part III, *Response*, 10(3), pp. 12-15.
- Bownds, J., Harlan, A., Lovelock, D., & McHugh, C. (1991). *Mountain searches: Effectiveness of helicopters*. Fairfax, VA: National Association for Search and Rescue.
- Bownds, J., Ebersole, M., Lovelock, D., & O'Conner, D. (1992). Reexamining the search management function—part IV, *Response*, 11(3), pp. 10-14.

- Brady, J., LaValla, P., Stoffel, R., & Wade, W. (c1981). Managing the search function. Olympia, Washington: Emergency Response Institute.
- Brown, S.S. (1980). Optimal search for a moving target in discrete time and space. *Operations Research*, 28, pp. 1275-1289.
- Charnes, A., & Cooper, W.W. (1958). The theory of search: optimal distribution of search effort. *Management Science*, 5, pp. 44-50.
- Colwell, M. (1992). New concepts for grid searching. *Search and Rescue Technical Publications Series, 1(1)*, Olympia, Washington: International Search and Rescue Trade Association (INSARTA).
- Colwell, M. (1994). *Search priority: an integrated approach to search planning*. British Columbia, Canada: Private publication.
- Colwell, M. (1996). *Trail-based probability of area: A terrain-based approach to POA estimation*. British Columbia, Canada: Private publication, 6 August 1996.
- Colwell, M. (1998). *Unifying geographical search assignments: Coordinating and prioritizing both area and trail searches*. British Columbia, Canada: Private publication, 16 October 1998.
- Cooper, D.C. (2002). *Search actions outline: A guide for planning and managing a land search*. Bangor, North Wales: Mountain Rescue Council of England and Wales.
- Cooper, D.C. (2000). *The application of search theory to land search: The adjustment of probability of area*. Private Publication, 20 March 2000, Cuyahoga Falls, Ohio, U.S.A.
- Cooper, D.C., & Frost, J.R. (1999a). *Selected inland search definitions*. Private Publication, 24 June 1999, Cuyahoga Falls, Ohio, U.S.A.
- Cooper, D.C., & Frost, J.R. (1999b). *Proportion-based consensus: Establishing initial POA values*. Private Publication, 6 August 1999, Cuyahoga Falls, Ohio, U.S.A.
- Cooper, D.C. (1998). "An Outline of Selected Concepts of Land Search Theory." *NASAR Conference Proceedings*, Response '98, 27th Annual Conference and Exhibition of the National Association for Search and Rescue, 27-30 May 1998, pp 83-106.
- Cooper, D.C., & Jones, A.S.G. (1997a). An approach to a grand unified theory of search (G.U.T.S.). *NASAR Conference Papers*, pp. 77-88, Richmond, Virginia: National Association for Search and Rescue.
- Cooper, D., LaValla, P., & Stoffel, R. (1996). *Search and rescue fundamentals: Basic skills and knowledge to perform wilderness, inland, search and rescue*. Third edition, revised. Olympia, Washington: ERI, Inc., and D.C. Cooper.

- Cornell, E.H., & Heth, C.D. (1996). Distances traveled during urban and suburban walks led by 3- to 12-year olds: Tables for search managers. *Response*, 15, pp. 6-9.
- Cornell, E. H., Heth, C. D., Kneubuhler, Y., & Sehgal, S. (1996). Serial position effects in children's route reversal errors: Implications for police search operations. *Applied Cognitive Psychology*, 10, 301-326.
- Cornell, E. H., Heth, C. D., & Alberts, D. M. (1994). Place recognition and wayfinding by children and adults. *Memory & Cognition*, 22, 633-643.
- Cozzolino, J.M. (1972). Sequential search for an unknown number of objects of nonuniform size. *Operations Research* 20, pp. 293-308.
- Dobbie, J.M. (1968). A survey of search theory. *Operations Research*, 16, pp. 525-537.
- Dougher, H. (2001). *Search management systems: A practical approach to initiating and directing efforts to resolve missing persons reports*. Third Revision, June 2001. Olympia, Washington: ERI International, Inc.
- Dougher, H., Goodman, R., LaValla, P., Long, C., Perkins, D., & Roberts, P. (2001). *Search management for the initial response incident commander*. Second Edition. Olympia, Washington: ERI International, Inc.
- Enslow, P.H. (1966). A bibliography of search theory and reconnaissance theory literature. *Naval Research Logistics Quarterly*, 13, pp. 177-202.
- Federal Aviation Administration [FAA]. (2003). *Aeronautical information manual (AIM)*. Newcastle, Washington: Aviation Supplies and Academics, Inc.
- Federal Emergency Management Agency [FEMA]. (2001). *National warning system operations manual*. FEMA Manual 1550.2. Washington, D.C.: FEMA.
- Federal Emergency Management Agency [FEMA]. (1998). *ICS: Incident command system*. Independent Study Course, IS-195. Washington, D.C.: Emergency Management Institute.
- Freund, J.E. (1974). *Modern Elementary Statistics*. 3rd Ed. Englewood Cliffs, NJ: Prentice Hall.
- Frost, J.R. (2001). *Comparing SARPlan with manual methods and actual cases*. Prepared for Defence Research Establishment Valcartier [DREV], Contractor Report Number W7701-001992/001/QCA, Fairfax, Virginia: Soza & Company, Ltd.
- Frost, J.R. (1999a). Principles of search theory, part I: Detection. *Response*, 17(2), pp. 1-7.
- Frost, J.R. (1999b). Principles of search theory, part II: Effort, coverage, and POD. *Response*, 17(2), pp. 8-15.
- Frost, J.R. (1999c). Principles of search theory, part III: Probability density distributions. *Response*, 17(3), pp. 1-10.

- Frost, J.R. (1999d). Principles of search theory, part IV: Optimal effort allocation. *Response*, 17(3), pp. 11-23.
- Frost, J.R. (1998a). *What's missing from ground search theory*. Fairfax, Virginia: Soza and Company, Ltd.
- Frost, J.R. (1998b). *Search theory enhancement study: prepared for interagency committee on search and rescue research and development working group*. Fairfax, Virginia: Soza and Company, Ltd.
- Frost, J.R. (1998c). *The theory of search: a simplified explanation*. (Rev. ed.). Fairfax, Virginia: Soza and Company, Ltd. & U.S. Coast Guard.
- Green, P. & LaValla, R., Eds. (1978). *The national park service instructor manual: Managing the search function*. Washington, D.C.: National Park Service.
- Goodman, R. & Cowan, B. (1996). *The "POC" factor*. Privately published paper presented to the William G. Syrotuck Symposium on Search Theory and Practice, Denver, Colorado: NASAR, May 29-30.
- Gorry, G.A., Kassirer, J.P., Essig, A.E., & Schwartz, W.B. (1973). Decision analysis basis for computer-aided management of acute renal failure. *American Journal of Medicine*, 55, pp. 473-484.
- Haley, K.B., and Stone, L.D. (1980). *Search theory and applications*. Nato Conference series, Series II: System Science. Published in cooperation with NATO Scientific Affairs Division. New York and London: Plenum Press.
- Hanfland, K.A., Keppel, R.D., & Weiss, J.G. (1997). *Case management for missing children homicide investigation: Executive summary*. Olympia, Washington: Office of the Attorney General, State of Washington, and U.S. Department of Justice's Office of Juvenile Justice and Delinquency Prevention, May 1997.
- Harbaugh, J.W., Doveton, J.H., & Davis, J.C. (1977). *Probability methods in oil exploration*. New York, New York: John Wiley & Sons.
- Heth, C. D., & Cornell, E. H. (1998). Characteristics of travel by persons lost in Albertan wilderness areas. *Journal of Environmental Psychology*, 18, 223-235.
- Hill, K. (1999). *Lost person behavior*. Ottawa, Ontario, Canada: the National Search and Rescue Secretariat.
- Hill, K. (1998). *Search, detection, and the visual briefing*. Presentation at the National Association for Search and Rescue's 27th Annual Conference and Exhibition, Portland, Oregon, 27-30 May.

- Hill, K. (1997). *Managing lost person incidents*. Chantilly, Virginia: National Association for Search and Rescue.
- Hill, K.A. (1996). *Distances traveled and probability zones for lost persons in Nova Scotia*. Unpublished data: author.
- Hill, K.A. (1995). How many clues do you need? Applying information theory to land search. *Response*, 14(3), pp. 6-8.
- Hill, K.A. (1992). Analyzing lost person scenarios. *Response*, 11(1), pp. 23-27.
- Hoel, P.G. (1976). *Elementary Statistics*. 4th Edition. New York: John Wiley & Sons.
- International Maritime Organization and International Civil Aviation Organization (IMO/ICAO). (1999a). *International Aeronautical and Maritime Search and Rescue Manual: Vol. I. Organization and Management*. London/Montreal: the International Maritime Organization (IMO) and the International Civil Aviation Organization (ICAO).
- International Maritime Organization and International Civil Aviation Organization (IMO/ICAO). (1999b). *International Aeronautical and Maritime Search and Rescue Manual: Vol. II. Mission Co-ordination*. London/Montreal: the International Maritime Organization (IMO) and the International Civil Aviation Organization (ICAO).
- International Maritime Organization and International Civil Aviation Organization (IMO/ICAO). (1999c). *International Aeronautical and Maritime Search and Rescue Manual: Vol. III. Mobile Facilities*. London/Montreal: the International Maritime Organization (IMO) and the International Civil Aviation Organization (ICAO).
- Kelley, D. (1973). *Mountain search for the lost victim*. Marsh, H. (Ed.), Montrose, California: Dennis E. Kelley.
- Kinder, G. (1998). *Ship of gold in the deep blue sea*. New York, New York: The Atlantic Monthly Press.
- Koester, R.J. and Stooksbury, D.E. (1995). Behavioral profile of possible Alzheimer's disease patients in Virginia search and rescue incidents. In *Wilderness and Environmental Medicine*, 6, 34-43.
- Koopman, B.O. (1946). *Search and screening* (OEG Report No. 56, The Summary Reports Group of the Columbia University Division of War Research). Alexandria, Virginia: Center for Naval Analyses.
- Koopman, B.O. (1953). The optimum distribution of effort. *Operations Research*, 1, pp. 52-63.
- Koopman, B.O. (1956a). The theory of search, Part I, Kinematic bases. *Operations Research*, 4, pp. 324-346.

- Koopman, B.O. (1956b). The theory of search, Part II, Target detection. *Operations Research*, 4, pp. 503-531.
- Koopman, B.O. (1957). The theory of search, Part III, The optimum distribution of searching effort. *Operations Research*, 5, pp. 613-626.
- Koopman, B.O. (1967). *Search and information theory: Part of final report on stochastic processes in certain naval operations*. Presented in preliminary form before the Operations Research Society of America, Chicago meeting, 1 November 1967.
- Koopman, B.O. (1979). Search and its optimization. *American Mathematical Monthly*, 86(7), August-September, pp. 527-540.
- Koopman, B.O. (1980). *Search and screening: general principles with historical applications*. Revised. New York, NY: Pergamon Press.
- Larson, R. (1972). *Urban police patrol analysis*. Cambridge, Mass.: MIT Press.
- LaValla, P., Stoffel, R., & Jones, A.S.G. (1997). *Search is an emergency: a text for managing search operations* (Fourth ed. revised). Olympia, Washington: Emergency Response Institute, Inc.
- Lovelock, D. (c1991). *The influence of clues on the POAs*. Private publication.
- Mangel, M. & Clark, C.W. (1982). *Uncertainty, search and information in fisheries*. Technical Report No. 82-6, Department of Mathematics, University of California, Davis (to appear in *Journal du Conseil*).
- Mattson, R.J. (1976). Establishing search areas. *Search and Rescue Magazine*, Spring 1976, pp. 7-8.
- May, W.G. (1973). *Mountain search and rescue techniques*. Boulder, Colorado: Rocky Mountain Rescue Group, Inc.
- McCoy, K.F. (1978). *Operational ASW search information system (OASIS), COMPATWING-SPAC memorandum 7 march*. (confidential).
- Mitchell, B.L. (1985). *A summary of the national association for search and rescue data collection and analysis program for 1980 – 1985*. Washington, D.C.: National Association for Search and Rescue.
- Moore, M.L. (1970). *A review of search and reconnaissance theory literature: Department of industrial engineering technical report*, University of Michigan.
- Morse, P.M. (October, 1970). Search theory and browsing. *The Library Quarterly*, October 1970.

- Morse, P.M. (1982). In memoriam: Bernard Osgood Koopman, 1900-1981. *Operations Research*, 30, pp. 417-427.
- Morse, P.M. and Kimball, G.E. (1946). *Methods of operations research: OEG report 54*. Operations Evaluation Group, Office of the Chief of Naval Operations, U.S. Navy Department, Washington, D.C.: U.S. Superintendent of Documents.
- National Search and Rescue Committee [NSARC]. (1999). *United states national search and rescue plan*. Washington, D.C.: U.S. Superintendent of Documents.
- National Search and Rescue Committee [NSARC]. (2000). *United states national search and rescue supplement to the international aeronautical and maritime search and rescue manual*. Washington, D.C.: U.S. Superintendent of Documents.
- National Search and Rescue Committee [NSARC] Research and Development Working Group. (2001). *Report of meeting of land search planning experts*. Meeting held 24 March 2001 in Laurel, Maryland, U.S.A.
- National Wildfire Coordinating Group [NWCWG]. (1994). *The incident command system national training curriculum*. Modules 1 – 17. Developed by the Interagency Steering Group. Boise, ID: National Interagency Fire Center, Division of Training.
- Paloheimo, J.E. (1971). On a theory of search. *Biometrika*, 58, pp. 61-75.
- Papoulis, A. (1984). *Probability, random variables, and stochastic processes*. Second Edition. New York: McGraw-Hill, pp. 38-39, 78-81, and 112-114.
- Perkins, D. (1989). Critical separation: its role in search management in theory and practice. In LaValla, P. (Ed.), *Probability of detection (POD) research, and other concepts for search management*, Olympia, Washington: Emergency Response Institute (ERI), pp. 1-39.
- Perkins, D. & Roberts, P. (1994). *The sector ladder: an alternative approach to managing searches*. Northumberland, England: Northumberland National Park Search and Rescue Team.
- Pollock, S. (1967). *Search, detection, and reconnaissance selection bibliography*. Department of Operations Analysis, U.S. Naval Post Graduate School.
- Reaves, B.A., & Hickman, M.J. (2002). "Census of state and local law enforcement agencies, 2000" in *Bureau of Justice Statistics and Bulletin*, October 2002.
- Richardson, H.R., & Discenza, J.H. (1980). The united states coast guard computer-assisted search planning system (CASP). *Naval Research Logistics*, 27, pp. 659-680.
- Richardson, H.R., & Stone, L.D. (1971). Operations analysis during the underwater search for Scorpion. *Naval Research Logistics*, 18, pp. 141-157.

- Richardson, H.R. (1967). Operations analysis. Chapter V of Part 2 in *Aircraft Salvage Operation, Mediterranean: Report to Chief of Naval Operations from the Supervisor of Salvage and the Deep Submergence Systems Project*, Washington, D.C.: U.S. Navy.
- Robe, R.Q., & Frost, J.R. (2002). *A method for determining effective sweep widths for land searches: Procedures for conducting detection experiments*. Prepared for the National Search and Rescue Committee, USCG Contract Number DTCG39-00-D-R00009. Alexandria, VA: Potomac Management Group, Inc.
- Robert T. Stafford Disaster Relief and Emergency Assistance Act, 42 U.S.C. 5121, et seq., as amended by the Disaster Mitigation Act of 2000, Pub L. No. 106-390, 114 Stat. 1552 (2000) (the Stafford Act).
- Sanderson, B. & Wartes, J. (c1983). *Applying probability density to probability of area*. Private publication.
- Setnicka, T.J. (1980). *Wilderness search and rescue*. K. Andrasko, Ed. Boston, MA: Appalachian Mountain Club.
- Shea, G. (1988). Formula for the field: a simple method for shifting probability of area. *Response*, 7(3), pp. 20-24.
- Stoffel, R. (2001). *The handbook for managing land search operations*. Cashmere, Washington: Emergency Response International, Inc.
- Stone, L.D. (1999). *Advances in search theory*. Unpublished manuscript, Metron, Inc., Reston, Virginia. {this is integrated into USCG R&D, 2001 above and will be removed}
- Stone, L.D. (1992). Search for the SS central America: mathematical treasure hunting. *Interfaces*, 22, pp. 32-54.
- Stone, L.D. (1989). *Theory of optimal search* (Second ed.). Military Applications Section, Operations Research Society of America. Arlington, Virginia: ORSA Books.
- Stone, L.D. (1984). Generalized search optimization. In *Statistical Signal Processing*, E.J. Wegman & J.G. Smith, Eds. New York, New York: Marcel Dekker, Inc.
- Stone, L.D. (1983). The process of search planning: Current approaches and continuing problems. *Operations Research*, 31, No 2, pp. 207-233.
- Stromquist, W.R., & Stone, L.D. (1981). Constrained optimization of functionals with search theory applications. *Mathematics of Operations Research*, 6, pp. 518-529.
- Strumpfer, J. (1980). *Search theory index*. Technical Note No. TN-017-80, Simonstown, South Africa: Institute for Maritime Technology,
- Syrotuck, W.G. (1974). *Some grid search techniques for locating lost individuals in wilderness areas*. Westmoreland, New York: Arner Publications.

- Syrotuck, W.G. (1975). *An introduction to land search probabilities and calculations*. Westmoreland, New York: Arner Publications.
- Syrotuck, W.G. (1976). *Analysis of lost person behavior: an aid to search planning*. Syrotuck, J.A., (Ed.), Westmoreland, New York: Arner Publications.
- Taylor, A. & Cooper, D. (1990). *Fundamentals of mantracking: The step-by-step method*. Second edition. Olympia, Washington: ERI, International.
- United Kingdom of Great Britain and Northern Ireland. (2002). *Search and rescue framework for the United Kingdom of Great Britain and Northern Ireland*. London, England: Queen's Printer and Controller.
- United States Air Force. (1999). *AFRCC history and mission*. Retrieved from www.acc.af.mil on March 23, 1999.
- United States Coast Guard Research and Development [USCG R&D] Center. (2001). *Review of search theory: advances and applications to search and rescue decision support*. Prepared for U.S. Department of Transportation, U.S. Coast Guard (G-OPR), Report No. CG-D-15-01, Washington, D.C.: U.S. Coast Guard.
- United States Coast Guard [USCG]. (2002). *United States Coast Guard search and rescue summary statistics 1964 through 2001*. Retrieved from www.uscg.mil/hq/g-o/g-opr/U_S_%20Coast%20Guard%20SAR%20Statistics%20Introduction.htm on 23 December 2002.
- United States Coast Guard. (1999). *U.S. Coast Guard: History*. Retrieved from www.uscg.mil/hq/g-o/g-opr/ on May 23, 1999.
- United States Coast Guard Auxiliary. (1999). *U.S. Coast Guard auxiliary history and accomplishments*. Retrieved from www.cgaux.org on May 23, 1999.
- United States Department of Transportation [USDOT]. (1996). *Value of life and injuries for use in preparing economic evaluations*. Prepared by the Office of the Assistant Secretary for Transportation Policy, Washington, D.C.
- Wagner, D.H. (1989). *Naval tactical decision aids*. Monterey, CA: Naval Postgraduate School.
- Wartes, J. (1983). Probability density: an aid to search planning. In *NASAR Conference Papers*, pp. 36-47. Dallas, Texas: National Association for Search and Rescue.
- Wartes, J. (1975b). *New concepts in sweep searching: the evidence favoring non-thorough methods*. Private Publication.
- Wartes, J. (1975a). *Search area confinement: an overview*. Private Publication.

- Wartes, J. (1974). *An experimental analysis of grid sweep searching*. Sunnyvale, California: Western Region Explorer Search and Rescue (ESAR).
- Washburn, A.R. (2002). *Search and detection*. Fourth ed. Linthicum, Maryland: Institute for Operations Research and the Management Sciences.
- Washburn, A.R. (1983). Search for a moving target: The fab algorithm. *Operations Research*, 31, pp. 739-751.
- Washburn, A.R. (1981). An upper bound useful in optimizing search for a moving target. *Operations Research*, 29, pp. 1227-1230.
- Wegener, I. (1981). The construction of an optimal distribution of search effort. *Naval Research Logistics Quarterly*, 28, pp. 533-543.

Appendix A

Selected Inland Search Definitions

From Cooper & Frost, 1999a

Introduction

In reviewing the inland search literature, it quickly becomes apparent that confusion is likely when a term is defined differently in various locations or when two terms are used to mean the same thing. It is recognized that many of these terms are not currently in general use in the ground search and rescue (SAR) literature. It is the intention of the authors here to offer factual, scientifically based definitions for terms that may be used in ground SAR operations and planning. In the interest of standardizing this terminology and reducing confusion, the authors also suggest that the following list of definitions and terminology be accepted and used by the inland search and rescue community.

The origins of many of the terms contained herein vary widely but includes operations research literature, international SAR literature (e.g., *The International Aeronautical and Maritime Search and Rescue (IAMSAR) Manual*) as well as conventional probability and statistics references.

Notation

A descriptive and complete notation is required to insure that terms are not confused. The notations illustrated in Sidebar A1 will be used to insure accuracy and consistency.

Selected Definitions

Area Effectively Swept (Z) – A measure of the area that can be (or was) effectively searched by searchers within the limits of search speed, endurance, and effective sweep width (IMO/ICAO, 1999b). The *area effectively swept (Z)* equals the effective sweep width (W) times search speed (V) times hours spent searching (exclusive of transits, breaks, etc.) in the search area (T) ($Z = W \times V \times T$) for one searcher or one resource (such as a boat or aircraft and its crew). Alternately, $Z = W \times D$, where D is the linear distance traveled while searching. The *area effectively swept* is described in units of area (i.e., square miles, etc.). If multiple searchers simultaneously follow independent paths when searching and together achieve approximately uniform coverage of the segment, then the total *area effectively swept* is given by $Z = n \times W \times V \times T$ where n is the number of searchers. “*Area Effectively Swept*” is also referred to as “Search Effort” and the linear distance traveled is also referred to as “Resource Effort” or just “Effort.” Note: The amount of *area effectively swept* does **not** equal the amount of ground actually viewed by the searchers while searching. The amount of *area effectively swept* is the amount of area that would have been swept by a hypothetical sensor that was perfect (100% effective) over a swath as wide as the effective sweep width centered on each searcher’s track and completely ineffective (i.e., made no detections) outside that swath. No such sensor exists, of course, but the concept of “*area effectively swept*” is nevertheless valid and useful for computing coverage, and using coverage to estimate probability of detection (POD).

Coverage (C , also known as Coverage Factor, Normalized Effort Density) – The ratio of the area effectively swept (Z) to the area searched (A) or $C = Z/A$ (IMO/ICAO, 1999b). For parallel sweep searches where the searcher tracks are perfectly straight, parallel, equally spaced, and the area covered is a parallelogram one-half track space larger than the pattern of parallel tracks on all sides, *Coverage* may be computed as the ratio of effective sweep width (W) to track spacing (S) or $C = W/S$. “ A ” (area searched) and “ Z ” (area effectively swept) must be described in the same units of area. “ W ” (effective sweep width) and “ S ” (track spacing) must be expressed in the same units of length. *Coverage* may be thought of as a measure of “thoroughness.” The POD of a search is determined by the coverage, as shown in Figure A1 (Koopman 1946). Perfectly executed parallel sweep searches under ideal search conditions may achieve POD values somewhat higher than those shown in Figure A1. On the other hand, systematic errors or biases in the actual performance of a search that prevent uniform coverage may result in POD values below the curve shown in Figure A1.




Figure A1. POD vs. Coverage (Koopman, 1946)

Defective Distribution (also known as Defective Probability Density Distribution) – A probability density distribution that contains less than 100% of the search object’s possible locations under a given scenario or set of scenarios. (Stone, 1989).

Effective Sweep Width (W) – A measure of the effectiveness with which a particular sensor can detect a particular object under specific environmental conditions (IMO/ICAO, 1999b). A measure of “detectability.” *Effective sweep width* depends on the search object, the sensor and the environmental conditions prevailing at the time and place of the search. There is no truly simple or intuitive definition. Actual *effective sweep width* values for specific situations must be determined by rigorous scientific experiments. However, reasonably accurate esti-

mates may be made from tables of *effective sweep widths* that have been determined by rigorous experiments for various typical search situations by applying appropriate “correction factors” to accommodate other search situations. A less accurate method of estimation for visual search is to assume the *effective sweep width* equals the “visual distance,” or average maximum detection range (both of which are different ways of thinking of the same value). Since the relationship between *effective sweep width* and maximum detection range is not consistent across all search situations, this method may either over-estimate or under-estimate the correct value. Therefore, it should be used only until more accurate *effective sweep width* data is available. Robe and Frost (2002) describes a procedure for conducting detection experiments from which effective sweep width values may be estimated.

The *effective sweep width* may be thought of as the width of a swath centered on the sensor’s track such that the probability of failing to detect an object within that width equals the probability of detecting the same object if it lies outside that width, assuming the object is equally likely to be anywhere. Another equivalent definition is: If a searcher passes through a swarm of identical stationary objects uniformly distributed over a large area, then the *effective search (or sweep) width (W)* is defined by the equation,

$$W = \frac{\text{Number of Objects Detected Per Unit Time}}{(\text{Number of Objects Per Unit Area}) \times (\text{Searcher Speed})},$$

where all values are averages over a statistically significant sampling period (Koopman 1946). Note that *effective sweep width* values are at least partially dependent on search speed. Generally speaking, a significant increase in search speed will decrease the *effective sweep width*. *Sweep width (W)* is needed to compute the area effectively swept (search effort or *Z*), and *Z* is needed to compute the coverage (*C*) based on the amount of search effort expended in the segment relative to the segment’s physical area. The POD may then be derived from the POD vs. Coverage graph (Figure A1).

Effort – The linear distance traveled by searcher(s) or resource(s) while searching in a segment. For one searcher or resource, it is computed as ($V \times T$). For multiple searchers it is computed as the sum of the distances traveled by each searcher, or, if all searched for the same amount of time at the same speed, it may be computed as ($n \times V \times T$) where *n* is the number of searchers. Also known as *track line length (TLL)*. The unit of measure for *Effort* is in linear distance. Used in the calculation of *Area Effectively Swept*.

Estimated Position (EP) – Last computed or estimated position for a lost search object.

Last Known Position (LKP) – Last witnessed or reported position of a lost search object (IMO/ICAO, 1999b).

On Scene Endurance – The amount of time a facility (resource) may spend at the scene engaged in search and rescue activities (IMO/ICAO, 1999b).

Optimal Resource Allocation – The process of determining where to assign the available search resources so that they produce the maximum possible probability of success (POS) in the minimum time.

Optimal Search Plan – A plan that maximizes the probability of finding the search object in the minimum amount of time by using the results of the optimal resource allocation process.

Parallel Sweep Search – A search tactic where one or more sensors, searchers, or resources (e.g., a helicopter) search an area by following a pattern of straight equally-spaced parallel tracks. Primarily used by vessels and aircraft, and for very thorough ground searches (e.g., evidence searches in conjunction with police investigations). Advantages include more uniform coverage of open areas and often a somewhat higher POD in such areas for a given level of effort than other techniques. While it is always a good idea to search any area in an organized fashion with a uniform coverage (until sufficient evidence is discovered to suggest another technique, such as tracking), in many ground search situations the terrain and ground cover make strict maintenance of straight tracks and equal spacing both impractical and counter-productive. However, an approximation to a *parallel sweep search*, such as “purposeful wandering” in parallel corridors, is often useful to help assure reasonably uniform coverage. Care must be taken to ensure the level of effort (distance traveled) is accurately estimated when searches do not follow straight, parallel tracks, even when they remain in parallel corridors. The $C = W/S$ formula only works when the searcher tracks themselves are perfectly straight, parallel, and equally spaced, and the area covered is a parallelogram one-half track space larger than the pattern of parallel tracks on all sides. The formula should never be used under other circumstances. The formula $C = Z/A$ always works, however.

Possibility Area – (1) The smallest area containing all possible survivor or search object locations. (2) For a scenario, the *possibility area* is the smallest area containing all possible survivor or search object locations that are consistent with the facts and assumptions used to form the scenario (IMO/ICAO, 1999b).

Probability Density (Pden) – The ratio of a region’s or a segment’s probability of area (POA) to its physical area.

$$[1] \quad Pden = \frac{POA}{Area}$$

Probability Map – An illustration of the distribution of search object location probability over the possibility area where each cell or region is labeled with the probability of the search object being in that cell or region (IMO/ICAO, 1999b). Initially, *probability maps* are formed from a largely subjective analysis of the available information (LKP, terrain, evidence, clues, historical trends, lost person behavior profiles, etc.). This information is evaluated to determine regions (see “Region”) where the subject might be at the time of the search based on one or more scenarios (see “Scenario”). It quantifies the probability of the subject being in each region, as shown in Figure A2. (See “Initial POA” under “Probability of Area.”)




Figure A2. A search area with four regions and their POA values after a consensus.

If the regions are subdivided into searchable segments, segment POA values are determined from the regional POA in proportion to the segment areas. It is assumed that the probability density (Pden) is constant throughout any one region. That is, the ratio of segment POA to region POA is the same as the ratio of segment area to region area, as shown in Figure A3. If the Pden is not constant throughout any one region, the number of regions and choice of regional boundaries should be refined until it is no longer possible to distinguish parts of regions on the basis of Pden.

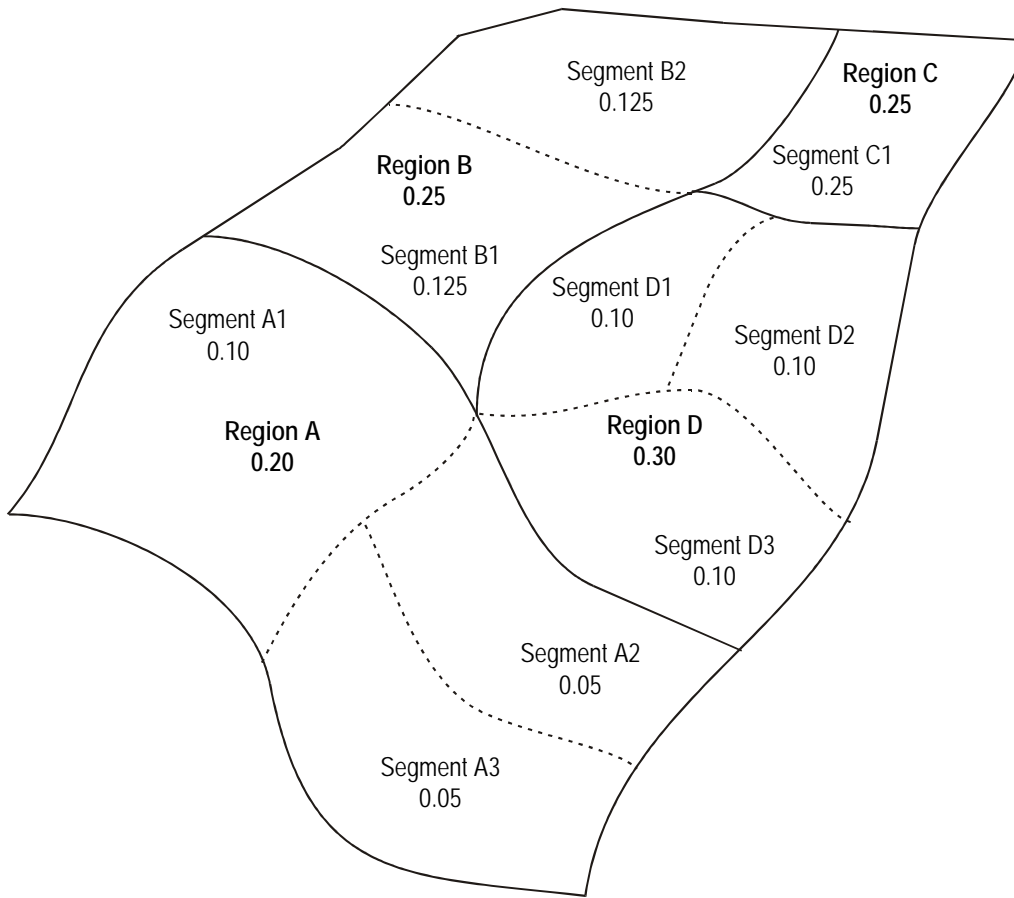


Figure A3. Segments within regions and their associated POA values.

In its purest mathematical form, a *probability map* consists of a regular grid of cells all of equal area as shown in Figure A4. Cellular probabilities are determined in the same way as segment probabilities. That is, each cell is assigned a fraction of the region's POA value in proportion to the fraction of the region's area contained in each cell. For cells that span regional boundaries, POA values are computed as the sum of the contributions from each region, pro-rated by the fractions of the regional areas contained in the cell. Although most useful in an open "unbounded" uniform environment (e.g., the ocean), this type of display may also be useful in mixed environments and has at least one advantage. When all the cells all have the same area, the POA values are proportional to the probability density (Pden) values so it is easy to tell at a glance where both POA and Pden values are high and where they are low. Note that by examining those cells that are completely contained within a region, it is clear that Region C has the highest density. It is also possible to tell that the Pden in Region C is nearly three times that of Region D. With segments or regions having unequal areas, it is possible to have a high POA and a low Pden and vice versa. Note that the POA of Region C is less than that of Region D. In general, Pden is more important to optimal resource allocation than POA.




Figure A4. A search area showing regions and a grid overlay.

A *probability map* may be made more readable by multiplying all the probabilities by some convenient constant. For example, if the cellular probabilities were all multiplied by 100, then 0.0129 would become 1.29%. Another technique (used in the original version of the U. S. Coast Guard’s Computer Assisted Search Planning (CASP) system) is to multiply all the cellular probabilities by 10,000 and record the results as whole numbers. In this case, 0.0129 would become 129.

Probability of Area (POA, also known as Probability of Containment or POC) – The probability that the search object is contained within the boundaries of a region, segment, or other geographic area. Regional *POA* values are generally determined by consensus and scenario analysis. Segment *POA* values may be computed from regional probability densities and segment areas.

Adjusted, Shifted or Updated POA (POA_{s,n}) – The modified POA of a segment after an unsuccessful search in that segment. Used to measure the decrease in the probability that the search object is in the segment after the segment has been searched. The following equations represent various methods of obtaining POA_{s,n}.

- [2] $POA_{s,n} = POA_{s,n-1} \times (1 - POD_{s,n})$
- [3] $POA_{s,n} = POA_{s,n-1} - (POD_{s,n} \times POA_{s,n-1})$
- [4] $POA_{s,n} = POA_{s,n-1} - POS_{s,n}$
- [5] $POA_{s,n} = POA_{s,0} \times (1 - POD_{cum_s})$
- [6] $POA_{s,n} = POA_{s,0} - (POD_{cum_s} \times POA_{s,0})$
- [7] $POA_{s,n} = POA_{s,0} - POS_{cum_s}$

Note: The *adjusted POA* values computed by the above formulas are not normalized. That is, the sum of the *adjusted POA* values will not equal the sum of the initial POA values. The omission of normalization is deliberate and necessary to the correctness of the formulas and definitions presented herein. Removal of the normalization computations does not violate the laws of probability and statistics in this context. Removal of normalization also substantially reduces the computational burden of maintaining *adjusted POA* values and preserves enough information about the search to make all other probability values of interest easily computable.

Initial POA ($POA_{initial}$) or *Consensus POA* ($POA_{s,0,c}$) – The initial POA assigned prior to any searching. *Initial POA* values must be based on a careful and thorough evaluation of all the available evidence, data, clues, etc., pertinent to the incident. *Initial POA* values, or the relative values used to compute them, must be in the correct proportions to one another. A region with an assessed value of “8” on a scale of 0 to 10 must be twice as likely, in the view of the evaluator making the assessment, to contain the search object as a region that is assigned a “4.” Similarly, a region with a POA of 20% must actually be viewed as twice as likely to contain the object as one with a POA of 10%. If, upon review before the evaluator submits his/her values the proportional relationships among the regional assessment values do not pass this test, then they should be revised until the evaluator feels they do correctly reflect his/her views in this regard. If the relative assessment values used are in the correct proportions, the POA percentages computed from them will also be in the correct proportions. The consensus POA values computed from the individual assessments should then be an accurate reflection of the collective views of the evaluators.

Ideally, the search area will be divided into some number of regions based on the available evidence, data, clues, etc., which bear on where the subject is more likely and less likely to be at the time of the first search. POA values would then be assigned to these regions. If necessary, these regions may be sub-divided into searchable segments. Segment POA values would be computed by prorating the region’s POA among the region’s segments by segment area. That is, a segment one-third as large as the region would get one-third of the region’s initial POA as its initial POA value. Stated as a formula:

$$[8] \quad POA_{s,0,c} = POA_{R,0,c} \times \frac{A_{s,c}}{A_{R,c}}$$

Where: $POA_{s,0,c}$ is the initial POA value for segment s in region R based on consensus c . Hereafter, it will be assumed that all values are based on the same consensus c if this subscript is omitted.

$POA_{R,0,c}$ is the initial POA value for region R based on consensus c .

$A_{R,c}$ represents the area of region R from consensus c .

$A_{s,c}$ represents the area of segment s in region R .

If a new consensus is necessary and new initial regional and segment POA values are established there is no need to discard all information about previous searching (i.e., segment POD values). Assuming that segment boundaries do not change, new adjusted segment POA values may be computed using the following procedure (the formulas show how to get from the adjusted POA values of the first consensus to those of the second consensus):

- Compute new initial segment POA values based on the new regional POA values from the new consensus using equation [7] above. (Note that $A_{s,2} = A_{s,1}$.)

$$POA_{s,0,2} = POA_{R,0,2} \times \frac{A_{s,2}}{A_{R,2}}$$

- Compute the cumulative POD for each segment (see *Cumulative Segment POD* ($POD_{cum_{s,n}}$) under Probability of Detection below) using equation [9] (preferred) or [10] or [11] below.

$$POD_{cum_{s,n}} = 1 - \left(\frac{POA_{s,n,1}}{POA_{s,0,1}} \right)$$

- Multiply the new initial segment POA by one minus the cumulative segment POD to get the new adjusted POA by using equation [5] above.

$$POA_{s,n,2} = POA_{s,0,2} \times (1 - POD_{cum_{s,n}})$$

Probability of Detection (POD , $POD_{s,n}$) – The probability of the search object being detected, assuming it was in the segment searched. $POD_{s,n}$ measures sensor effectiveness, thorough-

ness, and quality for search n of segment s . $POD_{s,n}$ is a function of the coverage (C) achieved in segment s by search n , as shown in Figure A1.

Cumulative Segment POD (POD_{cum_s}) - After the same segment is searched multiple times, the chances of having detected the search object, if it was present in the segment the whole time, are increased as compared to having searched the segment only once. This increasing probability of detecting a search object after multiple searches in the same segment is called *cumulative segment POD*.

$$[9] \quad POD_{cum_{s,n}} = 1 - \left(\frac{POA_{s,n,1}}{POA_{s,0,1}} \right)$$

$$[10] \quad POD_{cum_{s,n}} = \frac{POScum_{s,n,c}}{POA_{s,0,c}}$$

$$[11] \quad POD_{cum_s} = 1 - ((1 - POD_{s,1,c}) \times (1 - POD_{s,2,c}) \times \dots \times (1 - POD_{s,n,c}))$$

Predictive POD – estimated POD computed by search planners prior to the search of a segment based on predicted values for effective sweep width (W), area that will be effectively swept (Z), and coverage (C).

Retrospective POD – POD computed by using information obtained from debriefing the searchers to estimate the effective sweep width (W), area effectively swept (Z) and coverage (C) after the search of a segment.

Probability of Success (POS) - The probability of finding the search object with a particular search. POS measures search effectiveness.

Cumulative Probability of Success (POScum) - The accumulated probability of finding the search object with all the search effort expended over all searches to date (IMO/ICAO, 1999b). POS_{cum} may be computed for a segment, in which case it can never exceed the initial segment POA, or it may be computed for all searching in all segments to date (overall POS_{cum} or $OPOS_{cum}$ [see below]), in which case it can never exceed the total of all initial POA values (usually 1.0 or 100%).

Segment POS ($POS_{s,n}$) – The probability of finding the search object in the segment on a particular search (i.e., during a particular operational period).

$$[12] \quad POS_{s,n} = POA_{s,n-1} - POA_{s,n}$$

Segment POS_{cum} ($POScum_s$) – The sum of the POS values for each search in a particular segment. Used to measure the increasing probability that has been “extracted” from the segment by searching. This value can never exceed the initial POA value assigned to the segment. $POScum_s$ is a measure of search effectiveness to date in this segment.

$$[13] \quad \text{POScum}_s = \text{POS}_{s,1} + \text{POS}_{s,2} + \dots + \text{POS}_{s,n}$$

$$[14] \quad \text{POScum}_s = \text{POA}_{s,0} - \text{POA}_{s,n}$$

$$[15] \quad \text{POScum}_s = \text{POA}_{s,0} \times \text{PODcum}_s$$

Overall POS_{cum} (OPOS_{cum}) - The sum of all individual segment POScum values. Used to measure the increasing possibility that the search object is outside of the search area and the decreasing probability (1- OPOScum) that further searching based on the present scenario(s) will be successful. OPOS_{cum} is a measure of overall search effectiveness.

$$[16] \quad \text{OPOS}_{cum} = \sum \text{POA}_{s,0} - \sum \text{POA}_{s,n}$$

$$[17] \quad \text{OPOS}_{cum} = \sum_{s=1}^m \text{POScum}_s$$

Probable Success Rate (PSR) – The rate at which the probability of success (POS) is increased over time as the search progresses. An optimal search plan attains the maximum PSR possible from the available resources.

$$[18] \quad \text{PSR} = W \times V \times Pden$$

Where: *W* is the effective sweep width.
V is the search speed.
Pden is the probability density.

Resource Effort – See “Effort.”

Region (R) – A subset of the search area based only on factors that affect POA (*regions* may require segmentation prior to searching). *Regions* are based on probability of the search object’s location, not on suitability for assigning search resources. A *region* may contain searchable segments, or a *region*, itself, may be a searchable segment. A searchable segment may also contain one or more *regions* (based on probability) but rarely is the available data good enough to distinguish such small *regions* in ground search situations.

Scenario – A consistent set of known facts and assumptions describing what may have happened to the survivors (IMO/ICAO, 1999b). A description of what the subject(s) may have done and what the subject(s) may have experienced since last seen or known to be safe. A *scenario* should be consistent with a significant part of the available evidence and data. Normally, multiple *scenarios* should be considered, especially when not all the available pieces of evidence and data are consistent with all other pieces.

Search – An operation, normally coordinated, that uses available resources, personnel and facilities to find persons in distress or objects whose exact location is unknown (IMO/ICAO, 1999b).

Search Area – The area, determined by the search planner, that is to be searched. The *search area* may be divided into regions based on the probable scenarios and into segments for the purpose of assigning specific responsibilities to the available search resources (IMO/ICAO, 1999b).

Search and Rescue Facility – Any mobile resource, including designated search and rescue units, used to conduct search and rescue operations (IMO/ICAO, 1999b).

Search and Rescue Unit (SRU) – A unit comprised of trained personnel and provided with equipment suitable for the expeditious conduct of search and rescue operations (IMO/ICAO, 1999b).

Search Effort – See “Area Effectively Swept.”

Search Endurance – The amount of “productive” search time available at the scene. (IMO/ICAO, 1999b).

Search Object – A ship, aircraft or other craft missing or in distress or survivors or related search objects or evidence for which a search is being conducted (IMO/ICAO, 1999b). A generic term used to indicate evidence (clue) related to a lost subject or the lost subject. In the same segment, different search objects generally have different effective sweep widths (or “detectabilities”). This means that for any given search of a segment, different coverages, and hence different POD values, will be achieved for different *search objects*.

Search Speed (V) – The average rate of travel (speed over the ground) of searchers while engaged in search operations within a segment (IMO/ICAO, 1999b).

Segment (s) – A designated sub-area (subset of the search area) to be searched by one or more specifically assigned search resources. The search planner determines the size of a *segment*. The boundaries of a *segment* are identifiable both in the field and on a map and are based on searchability, not probability.

Sensor – Human senses (sight, hearing, touch, etc.), those of specially trained animals (such as dogs), or electronic devices used to detect the object of a search (IMO/ICAO, 1999b). A human, multi-sensor platform is often referred to as a “searcher.”

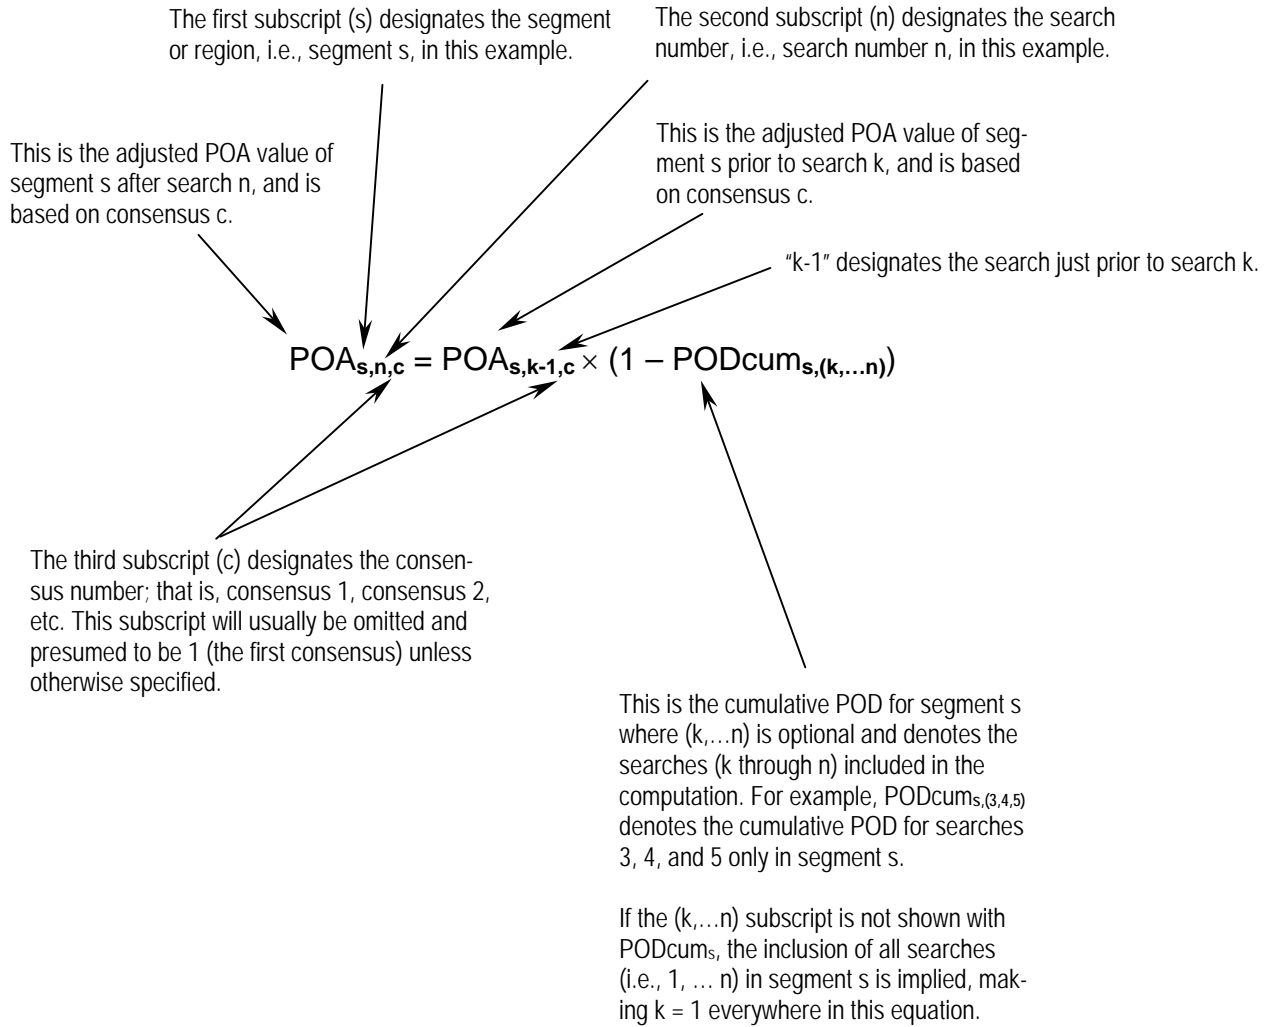
Sensor Track – The actual path followed by a sensor while engaged in searching. The length of this path is called *Effort*. For example, the actual path followed by a searcher carrying a GPS tracking device can be displayed on several computer-based mapping systems. Often these systems or the GPS receiver itself can compute and display the length of the path between any two recorded points.

Sortie – The individual movement of a resource in conducting a search or rendering assistance (IMO/ICAO, 1999b).

Sweep Width – See “Effective Sweep Width.”

Track Spacing – The perpendicular distance between adjacent tracks of a parallel sweep search pattern.

Sidebar A1 - Notation Used Herein



Sidebar A2 – Standard Symbols for Terms Defined Herein

A	Area
C	Coverage
c	Consensus (usually denotes the consensus number, e.g., first consensus, second consensus, etc.)
CASP	Computer Assisted Search Planning (US Coast Guard software)
cum	(as subscript) denotes cumulative value of associated term (e.g., POD_{cum} is cumulative POD)
EP	Estimated Position (usually computed)
LKP	Last Known Position
n	Search number
Pden	Probability Density
POA	Probability of Area
POC	Probability of Containment (identical to POA)
POD	Probability of Detection
POS	Probability of Success
PSR	Probable Success Rate
R	Region
S	(upper case) Track Spacing
s	(lower case) Segment
SRU	Search and Rescue Unit
T	Time
V	Velocity or Speed
W	Effective Sweep Width
Z	Area Effectively Swept (also known as Search Effort)

Appendix B

Analysis of Bayesian Updates and ROW

Ref (a): Freund (1974) {Note: Dr. Freund was Professor of Mathematics at Arizona State University when this book was published.}

Generalization of Postulate 3 [Addition Rule]

If $A_1, A_2, \dots,$ and A_k are mutually exclusive events, then

$$[1] \quad P(A_1 \cup A_2 \cup \dots \cup A_k) = P(A_1) + P(A_2) + \dots + P(A_k)$$

Freund (1974) p. 106.

General Rule of Multiplication

Given any events A and B

$$[2] \quad P(A \cap B) = P(B) \times P(A | B); \quad \text{provided } P(B) \neq 0$$

$$[3] \quad P(A \cap B) = P(A) \times P(B | A); \quad \text{provided } P(A) \neq 0$$

Freund (1974) p. 114.

The Rule of Elimination

If $B_1, B_2, \dots,$ and B_k are mutually exclusive events, of which none has zero probability and one must occur, then for any event A

$$[4] \quad P(A) = \sum_{i=1}^k P(B_i) \times P(A | B_i).$$

Freund (1974) p. 117.

Let A represent a detection event and let B_i represent the event of the subject being in segment i . Then the Equation [4] states, in words:

The probability of detecting the subject is equal to the sum of the products of the probability of the subject being in a given segment and the probability of detecting the subject if he is in the given segment.

In the notation used to express these quantities for use in SAR, $P(A) = OPOS$ (overall probability of success), $P(B_i) = POA_i$ (the probability of the subject being in segment i), and $P(A|B_i) = POD_i$ (probability of detecting the subject, given that the subject is in segment i). That is,

Table B1
Probability and SAR Notation Comparison

<i>Probability Notation</i>	<i>SAR Notation</i>
P(A)	OPOS
P(B _i)	POA _i
P(A B _i)	POD _i

Rewriting Equation [4] in SAR notation,

$$[5] \quad OPOS = \sum_{i=1}^k (POA_i \times POD_i).$$

The Rule of Bayes

If B₁, B₂, ..., and B_k are mutually exclusive events, of which none has zero probability and one must occur, then for any event A

$$[6] \quad P(B_i | A) = \frac{P(B_i) \times P(A | B_i)}{\sum_{j=1}^k [P(B_j) \times P(A | B_j)]}$$

for i = 1, 2, ... or k.

Freund (1974), p. 118.

Using the same meanings for *A* and *B_i* as before, Equation [3] may be stated in words as:

Given that a detection has occurred, the probability that the subject was found in segment *i* is equal to the *a priori* probability that the subject was in segment *i* times the probability of detecting the subject given he was in segment *i*, divided by sum of the products of all the *a priori* segment probabilities and the respective probabilities of detecting the subject given he was there.

We may rewrite the right side of Equation [6] in SAR notation as follows:

$$[7] \quad P(B_i | A) = \frac{POA_i \times POD_i}{\sum_{j=1}^k (POA_j \times POD_j)} = \frac{POS_i}{OPOS}$$

for i = 1, 2, ..., or k

and where *POS_i* denotes the probability of finding the subject in segment *i* based on the *a priori* probability of the subject being there.

What question does Equation [4] actually answer? Imagine the following situation. The subject is found at the end of an “operational period” and the event is duly reported to the search manager via radio. However, due to poor communications conditions, the search manager is unable to copy any other information. All he knows is that the subject was found somewhere. Given this additional piece of information, if the search manager wishes to amuse himself while waiting for a better radio communications link or a messenger, he might want to estimate, for each segment, the probability that the subject was found in that segment. Bayes’ Rule would apply and the search manager could simply compute all the $POS_j, j = 1, 2, \dots, k$, values and divide each by their total sum.

Non-Detection “Events”

The above problem is not a very useful application of Bayes’ Rule for SAR. The more usual application involves events that consist of searches that fail to detect the search object. To address this issue, let us define the event A' as a non-detection event, like a negative result from a medical test. Then our translation table from probability theory to SAR variables becomes

Table B2
Probability and SAR Notation Comparison
Non-detection Events

<i>Probability Notation</i>	<i>SAR Notation</i>
$P(A')$	$1 - OPOS$ (= $OPOA(\text{adjusted})$)
$P(B_i)$	POA_i
$P(A' B_i)$	$1 - POD_i$

We may then re-write Equation [7] as

$$[8] \quad P(B_i | A') = \frac{POA_i \times (1 - POD_i)}{\sum_{j=1}^k [POA_j \times (1 - POD_j)]} = \frac{POA_{i(\text{adjusted})}}{OPOA_{(\text{adjusted})}}$$

What question does Equation [5] actually answer? Imagine that a search was conducted without locating the subject. Given the knowledge that the subject was not located by the searching that was done, Equation [5] answers the following question for each segment i : “What are the chances that the failure to detect the subject was a result of the searchers failing to detect the subject even though the subject was actually there?” In other words, what are the chances that our “test” of the segment returned a false negative result? The answer to this question depends on two things: the *a priori* probability of the subject being in each segment and how well the searchers covered each segment (*coverage*, which leads directly to a POD_i value). Note that the question Equation [5] answers is not quite the same as, “Given the failure to detect the subject, what are the *a posteriori* chances that the subject is in segment i ?” although it may be given that interpretation.

At this point some concrete examples would be helpful. Figure B1 below shows a simple initial search situation with *a priori POA* values.




Figure B1. Simple initial search situation.

Experiment 1

Suppose segments A and B are both searched simultaneously with a coverage of 2.3. This gives a (“random” search) *POD* of 90% for each of these two segments. Segment C is not searched at all, giving a *coverage* and *POD* of zero. Treat segments A and B as a single segment (call it AB) having a *POA* of 90% for the moment. Assume that nothing was found, i.e., the search results were negative. We now ask the question, “Given that the subject was not found, what are the chances that this result was due to the searchers in segments A and B failing to find the subject even though he was there as compared to the chances that the subject was not in either of those two segments?”

If we simply compute the *POS* of the search of segment AB as $0.9 \times 0.9 = 0.81$ and subtract that value from the *a priori* probability of segment AB, we get a remaining *POA* of 0.09 or 9%. Alternatively, we could have computed $0.9 \times (1 - POD) = 0.9 \times 0.1 = 0.09$ to arrive at the same result. Note that this is the numerator of Equation [8] for segment AB. Also note that we can already answer our question. Assuming we have complete faith in the *POD* estimate and the *a priori POA* values, we may conclude that the chances that the search failed because the searchers missed the subject are slightly less than the chances that it failed because segment C was not searched (9% vs. 10%). Completing Bayes’ formula to compute the *relative* probabilities on a scale of 0 to 1, we find that, given the negative search result, there was about a 47% chance (0.09/0.19) that the search failed because the searchers missed the subject (i.e, that the negative result returned by the “test” of segment AB was false), and about a 53% chance (0.10/0.19) that the search failed because segment C was not searched.

This illustrates a basic characteristic of the use of Bayes' Rule: It is used to do a form of "backwards" reasoning from "effect" to "cause." Given a set of *a priori* probabilities for a complete set of events with known probabilities of occurring and the conditional probabilities for each possible outcome when it is known which event occurred, Bayes' Rule permits us, once a particular outcome is known but the *a priori* event leading to it is not known, to compute the probability that a given event was the antecedent of the known outcome.

We have specified and computed several probabilities. Let us see if we can state exactly what each one means in English.

1. $POD = 90\%$: In nine searches out of every 10, on average, we expect to find the subject if the subject is in the searched "segment." Alternatively, we expect to miss the subject 10% of the time even if he is in the searched "segment."
2. $POA = 90\%$: Nine times out of ten, on average, we would expect the subject to be in the "segment." Alternatively, we would expect the subject to be somewhere else 10% of the time.
3. $POS = 81\%$: On average, we would expect about four searches out of every five that use the strategy of searching 90% of the total POA with a coverage of 2.3 to succeed in finding the subject. That is, the chances of having found the subject by now using this strategy are about four out of five if our POD and *a priori* POA values are to be believed. Alternatively, we would expect that on average nearly one search in five using this strategy would fail to find the subject. While these results do not rule out the possibility that our POD and *a priori* POA values were correct, perhaps they do indicate we should consider other possible explanations for the search's failure despite the relatively high POS value.
4. Adjusted post-search $POAs$, $POA_{AB} = 9\%$, $POA_C = 10\%$: The failure of the search to locate the subject indicates that the subject is nearly as likely to be in the searched segment as out of it, based on our POD and *a priori* POA values. However, neither possibility is terribly likely; indicating there may be other likely alternatives (scenarios) not yet considered. Perhaps all the available information, data, and assumptions should be re-evaluated with a view toward other possible explanations for attaining such low adjusted POA values without finding anything.
5. Bayesian adjusted post-search $POAs$: Bayesian $POA_{AB} = 47\%$, Bayesian $POA_C = 53\%$: If we assume there are no other alternatives to consider and we believe our POD and *a priori* POA values are valid, we now "know" this search is among the 10% that were expected to fail even if the subject was in the searched segment and among the 19% that were expected to fail using this search strategy. If this is true, then the chances that the searching was unsuccessful because the searchers failed to detect the subject even if the subject was in the searched segments is just slightly less than the chances that it failed because the searchers were looking in the wrong place. Hence, if there are no other alternatives to consider and we believe our POD and *a priori* POA values are valid, the failure of the search indicates there is still nearly one chance in two that the subject is in the searched segments and a little better than one chance in two that he is in a segment that was not searched.

The normal technique for applying Bayes' Rule to negative search results uses the information in 1 and 2 above to provide only the information contained in the first line of 5, i.e., Bayesian $POA_{AB} = 47\%$ and Bayesian $POA_C = 53\%$. The information contained in 3 and 4 (which really amount to the same statistic since the $OPOS = 1 - (POA_{AB} \text{ (adjusted)} + POA_C)$), and the amplifying information in the body of 5 is not usually provided. It seems that one could easily argue that it is important for a search manager to know the odds for having obtained a successful outcome given the amount of effort expended so far and the strategy used to expend it. In fact, one could argue almost as easily that this statistic ($OPOS_{CUM}$) is just as important as knowing where to place the available effort during the next operational period and more important than simply knowing the relative $POAs$ of the segments on a scale of 0 to 1.

Experiment 2 – A Non-search Example

Ref (b): Hoel (1976) {Note: Dr. Hoel was Professor of Mathematics at the University of California, Los Angeles, when this book was published.}

It may be helpful at this point to examine an example of applying Bayes' Rule that is not related to search and rescue. The following example is based on an example given in Hoel (1976).

Suppose a test for detecting a certain rare disease has been perfected that is capable of discovering the disease in 97 percent of all afflicted individuals. Suppose further that when it is tried on healthy individuals, 5 percent of them are incorrectly diagnosed as having the disease. Finally, suppose that when it is tried on individuals who have certain milder diseases, 10 percent of them are incorrectly diagnosed. It is known that percentages of individuals of the three types being considered here in the population at large are 1 percent, 96 percent and 3 percent, respectively. The problem is to calculate the probability that an individual, selected at random from the population at large and tested for the rare disease, actually has the disease if the test indicates he is so afflicted (pp. 61-62).

We have a two-stage problem. The first stage consists of three types of events corresponding to the three types of individuals in the population at large. These events have the following probabilities:

- $P(e_1) = 0.01$
- $P(e_2) = 0.96$
- $P(e_3) = 0.03$

The second stage consists of two possible outcomes – either the test was positive, claiming that the tested individual has the rare disease, or it was negative, claiming that the individual tested does not have the rare disease. The conditional probabilities for the positive outcomes are:

- $P(o_1|e_1) = 0.97$
- $P(o_1|e_2) = 0.05$
- $P(o_1|e_3) = 0.10$

First, let us ask, “What is the probability that an individual chosen at random will test positive for the rare disease?” Using the multiplication and addition rules, we have

$$P(o_1) = P(e_1) \times P(o_1 | e_1) + P(e_2) \times P(o_1 | e_2) + P(e_3) \times P(o_1 | e_3)$$

$$P(o_1) = 0.01 \times 0.97 + 0.96 \times 0.05 + 0.03 \times 0.10 = 0.0607$$

Based on these computations, it is expected that only about 6 percent of the population will test positive for the rare disease. Now let us ask, “What is the probability that an individual selected at random from the population will have the rare disease and will test positive for it?” For this problem, we use the multiplication rule to compute

$$P(e_1 \cap o_1) = P(e_1) \times P(o_1 | e_1) = 0.01 \times 0.97 = 0.0097$$

That is, there is a little less than a one percent (0.97%) chance that an individual chosen at random will have the disease and will test positive for it. We may now answer our original question, “What is the probability that a person chosen at random from the population at large and who tests positive for the rare disease actually has it?” This is easily computed using Bayes’ Rule as

$$P(e_1 | o_1) = \frac{0.0097}{0.0607} = 0.1598$$

In other words, only about 16% of those who test positive for the disease will actually have it. Note that 84% of those who test positive (or about 5.1% of the population) will not have the disease despite the positive test result. Stated in yet another way, if a person is chosen at random from the population at large, there is about 1 chance in 100 that he has the disease, but if he is tested and tests positive, the chances are increased to about 1 in 6 that he has the disease.

We may also use Bayes’ Rule to compute the probability that a person chosen at random from the population at large will actually have the disease in spite of testing negative. We may infer from the previous computations that the probability that a person selected at random from the population at large will test negative for the rare disease is $1 - 0.0607 = 0.9393$. Using the multiplication rule, we may infer that the probability of a person actually having the disease and testing negative is $0.01 \times (1 - 0.97) = 0.01 \times 0.03 = 0.0003$. Now, Bayes’ Rule tells us that a person chosen at random from the population at large who tests negative for the disease has only about 1 chance in 3,000 (3,131 to be more accurate) of having the disease. Therefore, a person chosen at random from the population at large who tests negative has his chances of having the disease decreased from 1 in 100 to about 1 in 3000.

Let us pause and examine the final statements of the last two paragraphs. Are they true or false? Actually, they are both false. A particular individual's chances for having a particular disease have nothing to do with whether he or she has been tested for it. A person's chances for having a disease are usually linked to certain physical factors such as the person's genetic makeup, history of exposure to agents known to cause the disease or increase susceptibility to it, general condition of health, etc.

A test like the one examined here, although it is nearly 95% reliable (i.e., the negative and positive results returned are correct nearly 95% of the time), still returns positive results that are correct only 16% of the time when applied to the population at large. Therefore, it might not be applied to the population at large in practice but only applied to a subset consisting of individuals known to be at high risk based on family history, genetic analysis, medical history, work history, travel history, etc. In this case, we would expect the proportion of correct positive results to increase dramatically.

Returning to the "true-false" question that prompted this discussion, what the test results do is increase the individual's and his doctor's confidence levels about whether the individual has the disease from what these levels were prior to the testing. Again, the test results do not affect the individual's actual chances for having the disease, even though they would clearly affect the odds a bookmaker would be willing to give following disclosure of the test results.

Experiment 3 – Optimal Effort Allocation

Figure B2 below shows another search problem that is slightly different from Figure B1.




Figure B2. Another search problem.

The following table provides information about the three segments shown.

Table B3
Initial Values for Figure B2

Segment	A	B	C
Sweep Width (W)	0.075 km	0.050 km	0.100 km
Search Speed (V)	0.800 km/hr	0.500 km/hr	1.0 km/hr
Sweep Rate ($W \times V$)	0.060 km ² /hr	0.025 km ² /hr	0.100 km ² /hr
Area	0.667 km ²	1.000 km ²	9.000 km ²
Searcher Hours for $C = 1.0$	11.111 hrs	40.000 hrs	90.000 hrs
POA	0.300	0.600	0.100
Probability Density (Pden)	0.450/km ²	0.600/km ²	0.011/km ²
PSR ($W \times V \times Pden$)	0.027/hr	0.015/hr	0.001/hr

Note 1. The area of segment C does not include the areas of segments A and B.

Note 2. “Searcher Hours for $C = 1.0$ ” is the number of hours needed to search the given segment at a coverage of 1.0.

Note 3. PSR is “Probable Success Rate” and it is the product of sweep width, search speed and probability density. The sweep widths are probably unrealistically large.

Ref: (c) Charnes & Cooper (1958)

In 1958, Charnes and Cooper published a procedure for determining the optimum distribution of search effort over a group of cells (or segments) based on the information shown in Table 3 above and the assumption that the “random” search detection function applies. The procedure is actually simple in concept, but the mathematics can become somewhat involved.

The basic procedure Charnes and Cooper (1958) discovered has the following outline (paraphrased).

1. Determine the total amount of effort available for the next operational period.
2. Compute the *PSR* values for all segments.
3. Sort the segments in descending order of *PSR* values.
4. Determine how much search effort will be required to reduce the highest segment’s *POA*, and hence its *Pden*, so that its *PSR* value following the application of that effort will equal the *PSR* value of the second highest segment. If this exceeds the available effort, then stop this procedure and apply all of the available effort to the segment with the highest *PSR*. Otherwise, record the amount of effort needed and subtract it from the total available effort to get the amount of effort remaining to be allocated.
5. Determine how much search effort will be required in each of the two segments with the highest (and now equal) *PSR* values so that their post-search *PSR* values are reduced to that of the third highest segment. If this exceeds the remaining available effort, then apportion all of the remaining effort between the two highest segments so that each gets the same coverage with that effort. Add the effort from step 3 above to the first segment’s

portion from this step to get the total amount of effort that should be expended in that segment. Otherwise, add the efforts needed in each segment to the previous efforts allocated there and subtract them from the remaining effort to get a new value for effort remaining to be allocated.

6. Continue in this fashion until either the effort needed for the next step exceeds the amount remaining or the last segment has been reached (i.e., all segments have the same *PSR* value). In the former case, the remaining effort is apportioned across all segments that have already had effort allocated so that all receive the same coverage from that effort. These portions of effort are then added to each segment's effort total. In the latter case, the remaining effort is apportioned across all segments, including the one with the lowest initial *PSR* value, so that all receive the same coverage from that effort.
7. Apply the total effort accumulated for each segment to that segment in the next operational period.

We can go through the Charnes-Cooper (1958) procedure using the information in Table 3 as a starting point. For convenience, we will speak of available effort in terms of searcher-hours since that is how inland search managers tend to think of effort.

Suppose a search manager has 8 searchers available for 5 hours of searching, giving a total available effort of 40 searcher-hours. How should this effort be apportioned among the three segments to achieve the highest probability of success (*POS*)?

It is clear that segment A has the highest *PSR*. It can be shown that the *POD* needed in segment A to get its *a posteriori PSR* value down to that of segment B is given by

$$[9] \quad POD_A = 1 - \frac{PSR_B}{PSR_A} = 1 - \frac{0.015}{0.270} = 0.444 = 44.4\% .$$

It can also be shown that for the exponential detection ("random" search) function,

$$[10] \quad POD = 1 - e^{-C} ,$$

the coverage needed in segment A is given by

$$[11] \quad C_A = -\ln\left(\frac{PSR_B}{PSR_A}\right) = -\ln\left(\frac{0.015}{0.270}\right) = -\ln(0.556) = 0.588 .$$

The amount of *effort* (t_A , in searcher hours) is found by solving

$$[12] \quad C_A = \frac{W_A \times V_A \times t_A}{A_A}$$

for t_A to get

$$[13] \quad t_A = \frac{C_A \times A_A}{W_A \times V_A} = \frac{0.588 \times 0.667}{0.075 \times 0.800} = 6.531 \text{ hours} .$$

Subtracting this value from the 40 searcher hours originally available, we find we still have 33.469 searcher-hours left to allocate. Note that if we update the *POA* in segment **A** to reflect 6.531 hours of searching and create a new table like Table B3, we get Table B4 below.

Table B4
Values After One Charnes-Cooper (1958) Iteration

Segment	A	B	C
Sweep Width (<i>W</i>)	0.075 km	0.050 km	0.100 km
Search Speed (<i>V</i>)	0.800 km/hr	0.500 km/hr	1.0 km/hr
Sweep Rate (<i>W</i> × <i>V</i>)	0.060 km ² /hr	0.025 km ² /hr	0.100 km ² /hr
Area	0.667 km ²	1.000 km ²	9.000 km ²
Searcher Hours for C = 1.0	11.111 hrs	40.000 hrs	90.000 hrs
POA	0.167	0.600	0.100
Probability Density (Pden)	0.250/km ²	0.600/km ²	0.011/km ²
PSR (<i>W</i> × <i>V</i> ×Pden)	0.015/hr	0.015/hr	0.001/hr

Now we need to determine how much more effort is needed to drive the *PSR* values in segments A and B together down to that of segment C. Following the same procedure as before, we find that the coverage in segments A and B needs to be

$$[14] \quad C_{AB} = -\ln\left(\frac{PSR_C}{PSR_{AB}}\right) = -\ln\left(\frac{0.001}{0.015}\right) = -\ln(0.074) = 2.603 .$$

This is a very high coverage. As a shortcut, note that to search segment A to a coverage of 1.0 requires 11.111 searcher-hours while searching segment B to the same coverage requires 40.000 searcher-hours for a total of 51.111 searcher-hours to do both at a coverage of 1.0. To search both to a coverage of 2.603 would then require $51.111 \times 2.603 = 133.026$ searcher-hours. This is far more effort than we have remaining to allocate. Therefore, we must apportion the remaining effort between segments A and B. Since we already have the values for a coverage 1.0 search of each segment, it is easy to see that the fraction of the remaining effort that should go to segment A is simply $11.111/51.111 = 0.217$. So, $0.217 \times 33.469 = 7.276$ hours for a total of 13.807 hours allocated to segment A. The remaining 26.193 hours go to segment B. As a result, the optimal coverage for segment A is given by Equation [12] and found to be 1.243, producing a POD_A of 71.137% and POS_A of 21.341%. Similarly, the optimal coverage for segment B is 0.655, producing a POD_B of 48.047% and a POS_B of 28.828%. The *OPOS* for this search is found to be 50.169%. In short, the search has about a 50-50 chance of succeeding and that is the best that can be done with the available effort under the given circumstances.

As a practical matter, a search manager would no doubt approximate the above solution by assigning three searchers to segment A for 5 hours (a total of 15 searcher-hours) and the remaining

five searchers to segment B for five hours (a total of 25 searcher-hours). Note that this produces a coverage C_A of 1.35, a POD_A of 74.076%, and POS_A of 22.223% for segment A. For segment B, the results are a coverage C_B of 0.625, a POD_B of 46.474%, and POS_B of 27.884%. The $OPOS$ from this allocation is 50.107% – a very small decrease of no practical importance. The chances for success are still about 50-50.

The important lesson to learn from the last paragraph above is that when the theoretically optimal allocations are known, it is usually easy for the search manager to develop an operationally feasible plan that is very nearly optimal. However, without knowing what the optimal allocations are, the solving the optimal effort allocation problem is largely guesswork. Certainly the optimal solution is not intuitively obvious in the general case. {One might argue that in this case, the effort (in searcher-hours) was prorated in the closest feasible approximation to the POA values (3-to-5 effort ratio between segments with a 3-to-6 ratio of POAs) or Pdens (which are in the ratio of 3-to-4). To prove that this does not hold in the general case, exchange sweep width and/or search speed values between segments A and B and go through the Charnes-Cooper [1958] procedure again.} It is also worth noting that POS vs. effort curves are generally quite flat near the optimum value, thus giving the search manager quite a bit of latitude for altering the theoretically optimal allocation to accommodate practical exigencies.

Table B5 shows the results after the second (and final, in this case) Charnes-Cooper (1958) iteration.

Table B5
Values After Two Charnes-Cooper (1958) Iterations

Segment	A	B	C
Sweep Width (W)	0.075 km	0.050 km	0.100 km
Search Speed (V)	0.800 km/hr	0.500 km/hr	1.0 km/hr
Sweep Rate ($W \times V$)	0.060 km ² /hr	0.025 km ² /hr	0.100 km ² /hr
Area	0.667 km ²	1.000 km ²	9.000 km ²
Searcher Hours for C = 1.0	11.111 hrs	40.000 hrs	90.000 hrs
POA	0.087	0.312	0.100
Probability Density (Pden)	0.130/km ²	0.312/km ²	0.011/km ²
PSR ($W \times V \times Pden$)	0.008/hr	0.008/hr	0.001/hr

If we look back at Bayes' Rule, Equation [8] in particular, we see that a Bayesian update of the POA values in Tables [4] and [5] would simply divide all POA values now shown in the tables by their respective sums. Note that this would not change the PSR ratios in Equations [9], [11], or [14] on which the effort allocation computations are based. Therefore, the Bayesian normalization requires more computation but adds no value, at least not for effort allocation purposes.

The Role of ROW ("Rest of the World")

Several of the rules stated above began, "If $B_1, B_2, \dots, \text{ and } B_k$ are mutually exclusive events, of which none has zero probability and one must occur...." If this condition is to be met in the

global sense that the subject must be *somewhere*, it is necessary to acknowledge the possibility and some probability that the subject is not in the search area. One way to accomplish this is to create a fictitious “segment” and assign some amount of probability to it – presumably all the probability that cannot be attributed to the subject being in the search area. This fictitious “segment” is usually called the “rest of the world” or ROW.

Let us make a preliminary observation about ROW: It includes all possibilities except those of the subject being in the search area. This includes, for instance, the possibility that the subject is lost and in need of assistance, but not in the currently designated search area. In other words, the “*POA*” of ROW is not necessarily equal to one minus the probability that the subject is in need of assistance and located somewhere in the vicinity of the search area. ROW simply includes all scenarios that place the subject somewhere outside the search area. This group also contains everything from a simple change in the subject’s plans accompanied by a failure to inform key people (or maybe anyone), to abduction, runaway, victim of violence in some other area, and anything else imaginable (leading to the so-called “bastard” search).

“Completeness” of ROW

When developing *POAs* for a search area, search managers, or more appropriately Incident Coordinators, are (or should be) always mindful of the possibility that there may be some other explanation for the subject’s unknown whereabouts besides being lost somewhere in the search area. Accounting for this possibility by creating ROW and assigning probability to it gives the comfortable feeling of having “covered all the bases.” Unfortunately, this comfortable feeling is not entirely supportable.

Suppose we alter Figure B2 slightly to obtain Figure B3 below.

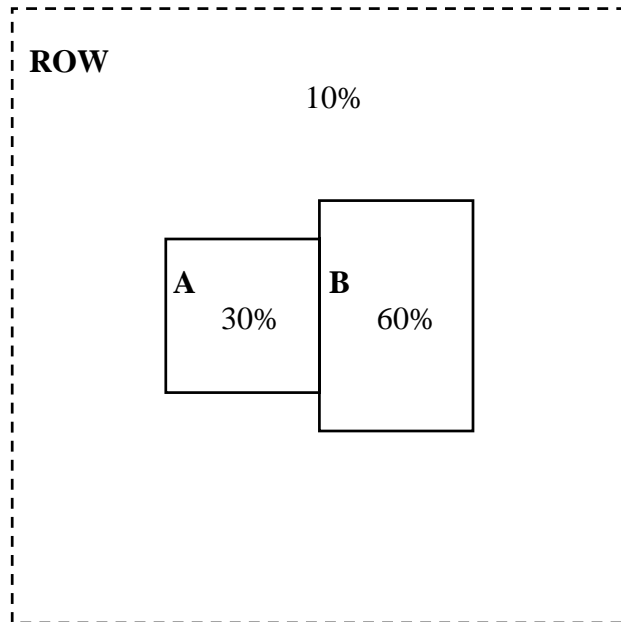


Figure B3. A third search problem.

Now suppose we return to our very first situation where the search manager was informed that the subject had been located, but was not told where. Suppose further that the *PODs* in segments A and B were both 90%. Then the probability that the subject was found in segment A may be computed using Equation [7] as 0.3333 and the probability that the subject was found in segment B may be similarly computed as 0.6667. Finding the subject in ROW is not admitted as a possibility. However in real life, subjects are found, and not infrequently, well away from the search area in situations quite different from the assumed situation on which the search was based. Something must be missing from our formulation of the problem.

To help us determine what is missing, Hoel [1976] provides a more complete statement of the necessary and sufficient conditions for applying Bayes' Rule. Paraphrasing his discussion on p. 60:

If there are m possible mutually exclusive events $e_1, e_2, \dots e_m$, exactly one of which must occur, and there are n possible mutually exclusive outcomes $o_1, o_2, \dots o_n$, exactly one of which must occur, then given that the j^{th} outcome o_j has occurred, the probability that the i^{th} event e_i was the antecedent of outcome o_j is given by Bayes' Rule:

$$[15] \quad P(e_i | o_j) = \frac{P(e_i) \times P(o_j | e_i)}{\sum_{k=1}^m [P(e_k) \times P(o_j | e_k)]}$$

We are missing one of the possible outcomes, that of finding the subject somewhere outside the search area. While we may have covered all possible antecedent events, we have not covered all possible outcomes. As a result, we do not have a conditional probability for "detecting" the subject given that the subject is in ROW. Hence, the thing that is missing is a "*POD*" for ROW.

Unfortunately, there is no reasonable way to estimate the probability that investigative and other efforts being applied to ROW will "detect" the subject or otherwise cause the subject's whereabouts and status to become known if the subject is not in the search area. However, it is known from experience that investigative and other efforts have a significant probability of locating the subject if he is not in the search area as well as a significant probability of turning up important clues and information even if the subject is in the search area. For these reasons, investigative and other efforts besides actual searching in the field should receive significant emphasis throughout the prosecution of the incident.

Now we have a conundrum. How can we justify increasing the "*POA*" of ROW through the Bayesian update process while at the same time denying the possibility of finding the subject there when such an outcome actually has a non-trivial probability of occurring?

Apparently, simply adding ROW to the problem does not actually "cover all the bases" since the probability of "detecting" the subject if he is in ROW seems to be a quantity beyond our powers of estimation.

ROW and Effort Allocation

If we construct a new table of values for Figure B3 using the information from Table B3 for segments A and B and updating the values in the last column to reflect the change from a third “real” segment to ROW, we would get the values shown in Table B6 below.

Table B6
Initial Values for Figure B3

Segment	A	B	ROW
Sweep Width (W)	0.075 km	0.050 km	—
Search Speed (V)	0.800 km/hr	0.500 km/hr	—
Sweep Rate ($W \times V$)	0.060 km ² /hr	0.025 km ² /hr	—
Area	0.667 km ²	1.000 km ²	∞
Searcher Hours for $C = 1.0$	11.111 hrs	40.000 hrs	∞
POA	0.300	0.600	0.100
Probability Density (Pden)	0.450/km ²	0.600/km ²	0.000
PSR ($W \times V \times Pden$)	0.027/hr	0.015/hr	0.000/hr

It is not possible (nor even sensible) to assign an effective sweep width, search speed or sweep rate to ROW. For all practical purposes, ROW has an unlimited area and, consequently, a zero probability density. Therefore, it will never be a factor in the optimal effort allocation process for the searching of areas, segments, etc. In other words, ROW is neither needed nor even useful for optimally allocating search effort.

Note: In some inland SAR manuals and papers, the statement is made that the “most efficient” search plan is the one that causes the largest increase in the [Bayesian] adjusted ROW “POA.” While this is certainly a true statement, provided one assumes the “POD” in ROW will always be zero, maximizing the [Bayesian] adjusted ROW “POA” is exactly equivalent to maximising OPOS (and OPOS_{CUM}), or minimising un-normalised adjusted OPOA, except that the latter concepts work regardless of whether ROW is acknowledged. Since full Bayesian updates are required for comparing one plan’s *a posteriori* ROW “POA” to another’s, this technique requires significantly more computation than that required for computing and comparing OPOS values.

ROW and Decision Support

As just noted, the Bayesian-adjusted ROW “POA” can have some *qualitative* value for comparing alternative search plans but it is not needed since there are easier ways to do such comparisons. We now ask whether the numeric value of ROW “POA” has any *quantitative* value for helping the search manager make important decisions. To answer this question, we must know exactly what ROW “POA” represents.

ROW “ $POA_{Bayesian-adjusted}$ ” – Assuming that the initial *a priori* POA values and cumulative POD values for subsequent searching are all correct, and assuming that the investigative and other efforts expended in ROW had no chance of locating the subject

even if he was not in the search area (i.e., ROW “*POD*” = 0.0), then ROW “*POA*_{Bayesian-adjusted}” is the probability that the unsuccessful outcome of all searching to date was “caused” by the subject being outside the search area.

If we assume that the cumulative ROW “*POD*” will always be small, then the ROW “*POA*_{Bayesian-adjusted}” may be a good approximation even if considerable investigative and other effort is expended in ROW. Given this assumption, if the search manager wants to be 95% confident that the reason the searchers are not finding anything is that the subject is not in the search area before he considers other alternatives, then he will have to continue searching until ROW “*POA*_{Bayesian-adjusted}” is 0.95. This can be a rather difficult goal to meet. If the initial *a priori* ROW “*POA*” was 0.1 (10%), then the un-normalised *OPOA*_{adjusted} for the search area would have to be driven down to only 0.00526 (0.526%). If the entire search area were covered uniformly, this would require a *POD* of 0.99415 (99.415%) in every segment of the search area. Obtaining such a *POD* would require a coverage of 5.14166 in each and every segment of the search area. In the search situation represented by Figure 3 and Table 6, nearly 263 searcher-hours would have to be expended in the 1.667 square kilometres of segments A and B to achieve this result. Even an optimal allocation of search effort requires just over 261 searcher-hours to achieve this result.

Relaxing the confidence level to 90% provides some relief, but not as much as one might hope. In this case, the un-normalised *OPOA*_{adjusted} for the search area would have to be driven down to 0.00870 (0.87%), assuming uniform coverage of the search area. This would require a *POD* of 0.98765 (98.765%), which in turn requires a coverage of 4.39445. For the situation represented by Figure 3 and Table 6, a total of 224.60 searcher-hours would be required. An optimal allocation would still require 222.94 searcher-hours to achieve the same result.

Note: Optimal effort allocation is most useful, and makes the biggest difference, when the area that can be effectively swept with the available resources is small compared to the size of the search area. The area that a searcher can effectively sweep (*Z*) is defined as the product of effective sweep width (*W*), search speed (*V*) and time available for searching (*t*). In other words, $Z = W \times V \times t$. If the total amount of area that all the available searchers together can effectively sweep is small compared to the size of the search area, and if the PSR values are not the same everywhere in the search area, then optimal effort allocation can significantly improve the chances for finding the subject earlier rather than later. (If the PSR values are the same everywhere, then the theoretically optimal solution is to place resources everywhere at once so that the coverage of the entire search area is uniform everywhere, no matter how low the resulting coverage. If this is impractical, then effort allocation decisions need to be based on other criteria, such as logistical constraints for example.) If the resource pool is large compared to the size of the search area, difficulty of detection of the subject and search speed, then optimal effort allocation can provide only marginal improvements. This is also true in the cumulative sense as the search progresses and more and more resources have been optimally expended. In other words, optimal effort allocation is generally most important in the early stages of a search and becomes progressively less important as the search progresses. However, if there is a “new start” because of significant changes in initial POA estimates based on late-arriving information or the search area is moved to a new location, then optimal effort allocation will again become important.

Comparison of Bayesian-adjusted ROW “POA” and $OPOS_{CUM}$

Let us begin with a preliminary observation. Examining Figure 3, it is immediately apparent that even a “perfect” 100%-effective search of segments A and B has only a 90% chance of being successful ($OPOS_{CUM} = 0.9$). This would be reflected in a Bayesian-adjusted ROW POA of 1.0 (100%), meaning that the search manager could state with absolute certainty that the search failed because the subject was not in the search area. However, this does not alter the fact that the search only had a 90% probability of finding the subject in the first place. Thus we see that overall search effectiveness ($OPOS_{CUM}$) and Bayesian-adjusted ROW POA are not the same thing.

Now let us make another interesting comparison. Suppose we treat segments A and B in Figure 3 as a single segment AB and forget about optimal effort allocation for the moment. Suppose we search segment AB at a coverage of 2.3, achieving a POD of 0.90 (90%). We may now state, at a 90% confidence level, that the subject was not in segment AB when it was searched since there was a 90% chance of detecting him if he was there. Note that this statement does not depend on the *a priori* POA_{AB} , ROW “ POA ,” or anything else. It stands cleanly on its own.

When we do consider the *a priori* POA values, the $OPOS$ for this search is 0.810 (81%) and the un-normalised adjusted POA_{AB} is 0.09 (9%). However, when we compute the Bayesian-adjusted POA values for segment AB and ROW, a quite different picture appears. For segment AB, the Bayesian-adjusted POA_{AB} is 0.47368 (47.368%). This means that the chances that the search’s failure was “caused” by the subject being in segment AB, undetected by the searchers, are about 47% or nearly one chance in two. The chances that the negative search result was “caused” by the subject being outside the search area, i.e., in ROW, are about 53% or slightly better than one chance in two.

We seem to have two quite dissimilar statistics for the same thing. The un-normalised adjusted POA_{AB} indicates there is only a 9% probability of the subject being in segment AB while the Bayesian-adjusted POA_{AB} seems to indicate there is a 47% probability of the subject being there.

Perhaps we should take a break and go on an Easter egg hunt.

Experiment 4 – An Easter Egg Hunt

Suppose a supplier of outdoor equipment is conducting a promotional “contest” where a check for \$10,000 is placed in one of 1,000 identical plastic “Easter eggs.” The eggs are then mixed together so it is not known which of the 1,000 contains the check. Nine hundred of the eggs are then selected at random and scattered in a uniformly random fashion within a wilderness area having known boundaries. The other hundred are retained by the promoter and are not available to be found. The team that becomes the finalist in the contest is given the opportunity to search the wilderness area and collect as many eggs as they can find.

Experiment 4-A

The initial conditions are that the *a priori* *POA* of the wilderness area is 0.9 (90%) and the *a priori* *ROW POA* is 0.1 (10%).

Suppose all the eggs are sealed and that the contest rules call for all recovered eggs to be returned unopened with their seals intact. Also assume there is no way to tell whether an egg contains the \$10,000 check without breaking the seal and opening it. Finally, suppose that the winning team searches the wilderness area and recovers 810 eggs, leaving 90 undetected eggs in the search area. The probability that they have recovered the egg containing the £10,000 check, if it was in the search area, is 0.9 (“*POD*” = 90%). The probability that they have recovered the egg containing the check is 0.81 (“*POS*” = 81%).

Suppose the team is given the opportunity to search the area again and recover as many of the remaining eggs as they can find. Based on their initial performance, they could expect to recover another 81 eggs for a total of 891 eggs. This would increase their chances of recovering the check, if it was in the search area, to 0.99 (“*POD*” = 99%). Their chances of having recovered the check, however, are only 0.891 (“*POS*” = 89.1%). If they do a second search, they can increase their chances for recovering the check by only 8.1 percentage points.

Note that since we do not yet know the outcome of the first search, there has been no opportunity to apply Bayes’ Rule. Also note that we are begging the question of whether a second search would be worthwhile. This will depend entirely on what “costs” and “risks” the team will incur by mounting a second search effort, and the value of the potential benefits (after taxes, etc.). Note that the cost of recovery per egg in the second search will be ten times that of the first search if the costs of both searches are the same. In real life, the incremental cost of the second search might be considerably less than that of the first search. Although there are statistical methods for assessing whether the potential costs are worth the potential benefits when the outcome is not guaranteed, we will defer this issue for the moment.

Experiment 4-B

Now suppose that the contest rules allow the searchers to know the outcome of the first search before deciding whether to mount a second search and they learn that the check was not recovered. Using Bayes’ Rule they can determine that the probability that they failed to find the egg containing the check because it was there, undetected, when they searched the area is 0.47368 (about 47%) while the probability that they failed to find the check because it was one of the 100 eggs held back by the promoter is 0.52632 (about 53%). It is clear that these are also the respective probabilities that, of the remaining possibilities, the check is in the search area or held by the promoter. Another interesting statistic is that since there are only 190 eggs of the original 1,000 left, each remaining egg is now a little more than five times as likely to be the one containing the check than it was prior to knowing the outcome of the search.

We may now ask, “How does knowing the outcome of the first search affect the decision about whether to mount a second search?” Clearly, if the outcome of the first search had been positive, there would have been no benefit at all to conducting a second search. Since the first search results were negative, it is just as clear that if the team is to have any chance of recovering the egg with the check, they will have to search again. However, does this knowledge affect the total probability that two searches, each having a 90% *POD*, will recover the egg with the check? At the end of the second search in either case, 891 eggs will have been recovered and the chances

that the one containing the check was among them will still be 891/1000 or $OPOS_{CUM} = 89.1\%$. The total “cost” of searching to get to this point will also remain unchanged.

If we used the Bayesian-adjusted *POAs* to estimate the chances of finding the check with a second search using the same amount of effort as the first (and presumably achieving the same “*POD*”), we would get a “*POS*” for the second search of $0.9 \times 0.47368 = 0.42632$ (about 43%). This is a little more than half the probability of success of the first search. In a cost/benefit type of analysis, this number might be useful for helping the team decide whether a second search is worthwhile.

Returning to the original question of whether the Bayesian-adjusted ROW *POA* has any real quantitative “meaning” or value for making SAR search planning decisions that are not typically based on a cost/benefit analysis, it appears that, standing alone, the Bayesian-adjusted ROW *POA* really does not provide much useful information. One question the Bayesian-adjusted ROW *POA* specifically does **not** answer is, “What is the probability that, given all the searching effort expended to date, the subject would have been found by that searching if all of our *a priori* *POAs* and subsequent *PODs* were correct?”

Search Planning vs. Incident Co-ordination

Perhaps part of the confusion surrounding the ROW issue stems from mixing search planning issues with those of overall incident management. The search planner is concerned with only a certain particular subset of the possibilities. That subset consists of those scenarios where the subject is lost and in need of assistance somewhere in an area where it will be possible for searchers to find him. The search planner’s job is to allocate the available effort so that if one of those scenarios is true, then the probability of finding the subject alive in the minimum time is maximised.

Note: The added condition “alive” makes for an optimal survivor search, something that is slightly more complicated to compute. It is useful when the hazards to the subject’s continued survival are not uniformly distributed over the area. This includes the *a priori* risks to the subject. For example, if it can be said of some region that the subject was more likely to suffer an injury than in other places, then the hazard to the subject’s continued survival are high because the survival vs. time curve for injured persons drops dramatically from that of un-injured persons.

In any case, the search planner is not concerned with the “rest of the world.” Mathematically, there is nothing that prohibits us from considering a smaller sample space, such as one conditioned on the hypothesis that the subject is in the area and in need of assistance. This in turn means that all of the possibilities of concern to the search planner are “in the area” and there is no need to consider ROW.

On the other hand, the Incident Commander (IC) must be concerned with “ROW” since searching addresses only a portion of the possible scenarios. While it is possible to give the IC a quantitative measure of search effectiveness (ratio of $OPOS_{CUM}$ to the initial total *POA* in the search area), it is very difficult if not impossible to assess the effectiveness of investigative and other efforts in an objective, quantitative sense. If the latter were possible, then perhaps the IC could

compare the effectiveness of one activity to the other and decide which needed more emphasis. This would not be done in an effort allocation scheme, however, since investigative resources and searching resources are generally disjoint sets by virtue of their differing highly specialised tasks – i.e., resources from one set are not generally suitable for the taskings of the other set. That is, one cannot generally turn searchers into investigators or vice versa. Therefore, while the IC must keep the “rest of the world” in mind, computing its Bayesian-adjusted *POA* (as if investigative efforts were always completely ineffective) does not seem to provide any useful quantitative information on which to base decisions.

Further Analysis

Utility of Un-normalised POA Values.

Tracking un-normalised *POA* values preserves all of the “negative” information about the searching done to date while still permitting reasonably easy access to the Bayesian-adjusted *POA* values. For example, to compute the Bayesian-adjusted *POA* values for the segments at any point, simply divide the individual un-normalised adjusted *POA* values by their sum. That is,

$$[16] \quad POA_{\text{Bayesian-adjusted}} = \frac{POA_{\text{adjusted}}}{\sum_{\text{all segments}} POA_{\text{adjusted}}},$$

where POA_{adjusted} denotes the un-normalised adjusted *POA* of a segment. Recall that for searched segments, this value is the previous POA_{adjusted} value times one minus the *POD* of the last search. For un-searched segments, the *POA* value remains unchanged, i.e., the “*POD*” is zero and the “adjusted” value equals the previous value.

The cumulative *POD* for any segment is simply

$$[17] \quad POD_{\text{CUM}} = 1 - \frac{POA_{\text{adjusted}}}{POA_{\text{initial}}}.$$

The cumulative *POS* for any segment is simply

$$[18] \quad POS_{\text{CUM}} = POA_{\text{initial}} - POA_{\text{adjusted}} = POA_{\text{initial}} \times POD_{\text{CUM}}.$$

The $OPOS_{\text{CUM}}$ at any time is given by

$$[19] \quad OPOS_{\text{CUM}} = 1 - \sum_{\text{all segments}} POA_{\text{adjusted}}.$$

There is another very important capability as well. If the initial *POA* values should change due to the effects of late-arriving information, it is very easy to use Equation [17] to first compute the

cumulative POD values to date for all the segments, and then compute in a single step all the new adjusted segment POA values to date from the new initial POA values using

$$[20] \quad POA_{adjusted(new)} = POA_{initial(new)} \times (1 - POD_{CUM}).$$

This means that not only is the effect of previous searching not lost as the result of a re-evaluation of initial $POAs$, the effect of that searching given the new initial POA data set is easy to compute. Once the new adjusted POA values are known, then all of the other values of interest may be found with Equations [16], [18], and [19].

In short, any statistic one would want about the searching done to date with respect to either an individual segment or all segments collectively is readily available if the un-normalised POA values are computed and tracked.

Answering Search Planning/Management Questions

The line between the concerns of the search planner and those of the Incident Co-ordinator is not always a clear one. However, we may roughly delineate their concerns with the following two questions.

1. If the subject is in the general area of the point last seen, last known position, etc., how should the available search resources be deployed in order to maximize the chances of finding the subject alive in the minimum amount of time?
2. Should active searching in the field continue, or should it be suspended pending further developments?

We have seen that the first question is of concern to the search planner. We have also seen that neither full Bayesian POA updates nor ROW really contribute anything toward helping the search planner answer these questions. If that is the case, then what should the search planner think of low adjusted POA values and consequently high $OPOS_{CUM}$ values? Several things, alone or in combination, can lead to this situation before finding the subject. These include

- a. All or some part of the information used to form the scenario(s) on which the searching was based was interpreted incorrectly, leading to any or all of the following: an inappropriate search area, a likely scenario that was not among those initially considered, inappropriate scenario weights, incorrect initial POA values.
- b. The estimated POD values were optimistic.
- c. There has been some systematic error in the conduct of the search, such as searchers avoiding places they found difficult to search, i.e., non-uniform coverage of segments.
- d. A computational error has been made.

All of these issues, and the many more that one can imagine, have one thing in common: The only practical way to discover them is to carefully review all the available information gathered to date, including all the search team debriefs, and carefully inspect the facts, assumptions, observations and reasoning that led to the POA and POD values used in the computations. Inspecting the computations themselves is a good idea also. A thorough review is very difficult, mentally demanding work that requires concentration and hence peace and quiet away from the fre-

netic operational activities. Someone not too closely connected with the incident should do such reviews to avoid having someone who will try to justify an earlier position. These reviews should be done as frequently as practicable.

Despite their importance, it is easy to forget or put off reviews. Unlike the problems found in elementary statistics texts and many found in business, medicine, industry, etc., search planners and Incident Co-ordinators rarely have any firm statistics on which to base their decisions. In the problem involving the test for a rare disease, the population statistics were given and presumably based on a large sampling of medical records and possibly death certificates. The statistics about the efficacy of the test were given and presumably based on a large number of trials whose results were confirmed or denied by other, more reliable means. In the case of the “Easter egg hunt,” the numbers of eggs in and out of the search area were precisely known as were the number actually recovered.

Search planners do not have such luxuries. Initial scenarios, scenario weights and *POA* values are almost entirely subjective and are often based as much upon assumptions as upon facts. In the past, inland SAR *POD* estimates were also entirely subjective, almost certainly optimistic, and often inconsistent. The way to put *POD* estimates on an objective basis is by developing an objective means for estimating *effective sweep widths*. A procedure for conducting detection experiments and developing sweep width estimates from the data they produce may be found in Robe and Frost (2002). Regardless, search planners need to be reminded to check their work and to avoid putting undue faith in previous estimates, especially subjective ones and especially in the face of a continued lack of success. One of the easiest and most obvious reminders is declining un-normalised adjusted *POA* values. As he sees the *POAs* fade away without bringing the search to a successful conclusion, the search planner (and the IC) should be prompted to re-evaluate all the evidence gathered and searching done to date in a careful, logical manner.

The second question is more within the purview of the IC. It is a very difficult question to answer and there exists no hard and fast formula that provides an answer, with or without full Bayesian *POA* updates or ROW. Many factors have to be considered, and the overall effectiveness of all searching done to date is but one of them. Another is the projected incremental improvement in the chances for success given the application of additional searching effort. The subject’s probability of survival to date in the given environment is another. Additional facts turned up by the investigative efforts are still another factor. Political pressure can be another factor. The list is nearly endless.

Convenient Features of the “Random” Search Curve

The so-called “random” search detection function (also known as the exponential detection function) has a couple of mathematically convenient features. These include the following:

1. Simplicity: The formula for computing *POD* from coverage (*C*) is

$$POD = 1 - e^{-C}$$

where *e* is the base of the natural logarithms.

2. Cumulative *POD*: The cumulative *POD* for two or more searches that achieve coverages C_1, C_2, \dots, C_n is equal to the *POD* of a single search using a coverage C_T equal to the sum of the individual coverages, i.e., $C_T = C_1 + C_2 + \dots + C_n$. (The Charnes-Cooper [1958] algorithm takes advantage of this feature. Adapting it to other detection functions, such as that of Koopman's (1946, 1980) inverse cube law of visual detection, requires changing the algorithm somewhat. One way to do this is to pick a target *PSR* plateau and determine whether there is enough effort available to search all segments with higher *PSR* values down to that plateau. If not, increase the target *PSR* value. If so, decrease the target *PSR* value. Using a binary search algorithm, this method should converge to a solution fairly quickly.)

The Ultimate Objective

The ultimate objective is, or should be, to provide inland search planners/managers with a scientifically-based, clear, concise, mathematically consistent, simple and practical procedure for planning searches, evaluating search results, and using negative results of previous searching to plan the next search in a sequence. Satisfying all these goals is challenging, to say the least. The minimum requirement from a scientific perspective is a reasonable method for estimating and using sweep widths, as described in Robe and Frost (2002). This will make possible a level of objectivity and consistency in *POD* estimation that is sorely needed in the inland SAR arena. The exponential detection function seems to be a reasonable assumption for an unbiased estimator of *POD* as a function of coverage in the inland search environment. The tracking and direct use of un-normalised *POA* values also seems to be a reasonable compromise for achieving simplicity and practicality.

One of the issues that make the application of statistical techniques to SAR difficult is the short time scale. It is generally not possible to wait until one has achieved a 95% confidence level that the reason the searchers are not finding anything is that the subject is not in the search area before seriously considering other alternatives. On the other hand, we do not want to promote knee-jerk reactions and unproductive "helter-skelter" searching where searchers barely have a chance to search one region before they are told to abandon it for another region. Reasonable *POD* estimates based on sweep widths from detection experiments, a basic search-theory-based methodology using coverage and the exponential detection function, and fading un-normalised *POA* values as the search progresses seems to be the best combination for meeting SAR goals.

Appendix C

Simplified Explanation of Sweep Width

An Analogy

Even though *effective sweep width* (usually shortened to just *sweep width*) is essentially a mathematical concept, it can be explained or at least illustrated in mostly non-mathematical terms. To avoid descending too deeply into the pit of mathematics, we will need a common, easily visualized activity that can be used as a model, or analogy, for detection. So, let us pick the mundane activity of sweeping floors as an analogy for “sweeping” an area in search of a lost or missing person. We will use this analogy to describe hypothetical experiments that illustrate the basic principles of *effective sweep width*.

Suppose we wish to compare the performance of four different push-broom designs. In the first design, the broom head is one-half meter (50 cm) in width and has fine, closely-set bristles. In the second design, the broom head is a full meter in width but the bristles are more coarse and not as dense as with the first broom. The third broom is two meters in width with bristles that are even coarser and less dense than those of the second design. The fourth broom is again one meter in width, but it is a hybrid design where the center 20 cm is identical to the first broom, the 20 cm sections to the right and left of the center section are identical to the second broom, and the outboard 20 cm sections at each end are identical to the third design. Figure C1 shows a schematic representation of the four different designs. We construct the brooms and label them as B1, B2, B3, and B4 respectively.




Figure C1. Four brooms (B1, B2, B3, and B4).

In our first experiment, we want to know how the brooms compare to one another on a single sweep through a previously unswept area. To perform this test, we choose a smooth floor and mark off a square test area measuring 10 meters on a side. Using sand to simulate dirt on the floor, we cover the test area lightly, and uniformly, so that the “density” of sand is 10 grams per

square meter (g/m^2) of floor surface. We then push each broom in a straight line from one side of the test area to the other at a constant speed of 0.5 m/sec (1.8 km/hr or a little over 1 mph), collect the sand that was swept up, and weigh it.

When B1 is pushed through the test area, it appears to do a very good job. In fact, it makes a “clean sweep” with a width of 0.5 meters (or width of the broom head), as illustrated in Figure C2.



Figure C2. Broom 1 (B1)

It swept up 50 grams of sand—all the sand within the 0.5 m x 10 m swept area. Thus we may say that B1 is 100% effective out to a distance of 25 cm either side of the center of its track, and, because of the physical limitation of the broom’s width, it is completely ineffective at greater distances. The maximum lateral (side-to-side) range of the broom is 0.25 meters from the center of its track. Finally, since it took 20 seconds to traverse the 10-meter “test course,” B1 swept up the sand at the average rate of 2.5 grams per second.

Broom B2 is not as thorough as B1, but it makes a swath twice as wide as illustrated in Figure C3.



Figure C3. Broom 2 (B2).

When the sand from B2 is weighed, it turns out that it too swept up 50 grams of sand. As a quick calculation will show, B2 swept up 50% of the sand in the one-meter-wide swath it made. Further analysis shows that all parts of the broom performed equally, and both the sand swept up and that left on the floor were uniformly distributed across the width of the swath. Thus B2 is 50% effective out to a distance of 0.5 meters on either side of the center of its track, and completely ineffective beyond that distance. The maximum lateral range of B2 is 0.5 meters from the center of its track. Just as with B1, broom B2 swept up the sand at the average rate of 2.5 grams per second.

Broom B3 is even less thorough than B2, but it makes a swath twice as wide as B2 and four times as wide as B1, as shown in Figure C4.



Figure C4. Broom 3 (B3).

Furthermore, it too sweeps up 50 grams of sand and is found to be uniformly 25% effective over the two-meter swath it makes. The maximum lateral range is one meter either side of track and it swept up sand at the same rate of 2.5 grams per second.

Finally we push B4 through an unswept portion of the test area. When the sand from B4 is weighed, again we find we have 50 grams. More detailed analysis shows the center section made a clean sweep 20 cm wide, getting 20 grams of sand in the process. The two adjacent 20-cm sections swept up 10 grams of sand each for another 20 grams. This amounts to 50% of the sand present in the two corresponding 20-cm strips on the floor. Finally, the two outboard 20-cm sections got only 5 grams of sand each, which means they were only 25% effective in their respective strips. Figure C5 illustrates the uneven performance of broom B4.




Figure C5. Broom 4 (B4).

Based on the physical size of B4, the maximum lateral range of B4 is 0.5 meters from the center of its track. Finally, just as with the other brooms, B4 swept up the sand at the average rate of 2.5 grams per second.

If we graph each broom's *performance profile* as the *proportion of dirt (pod)* lying in the broom's path that is swept up across the width of the broom head as it moves forward, we get the graphs shown in Figure C6.

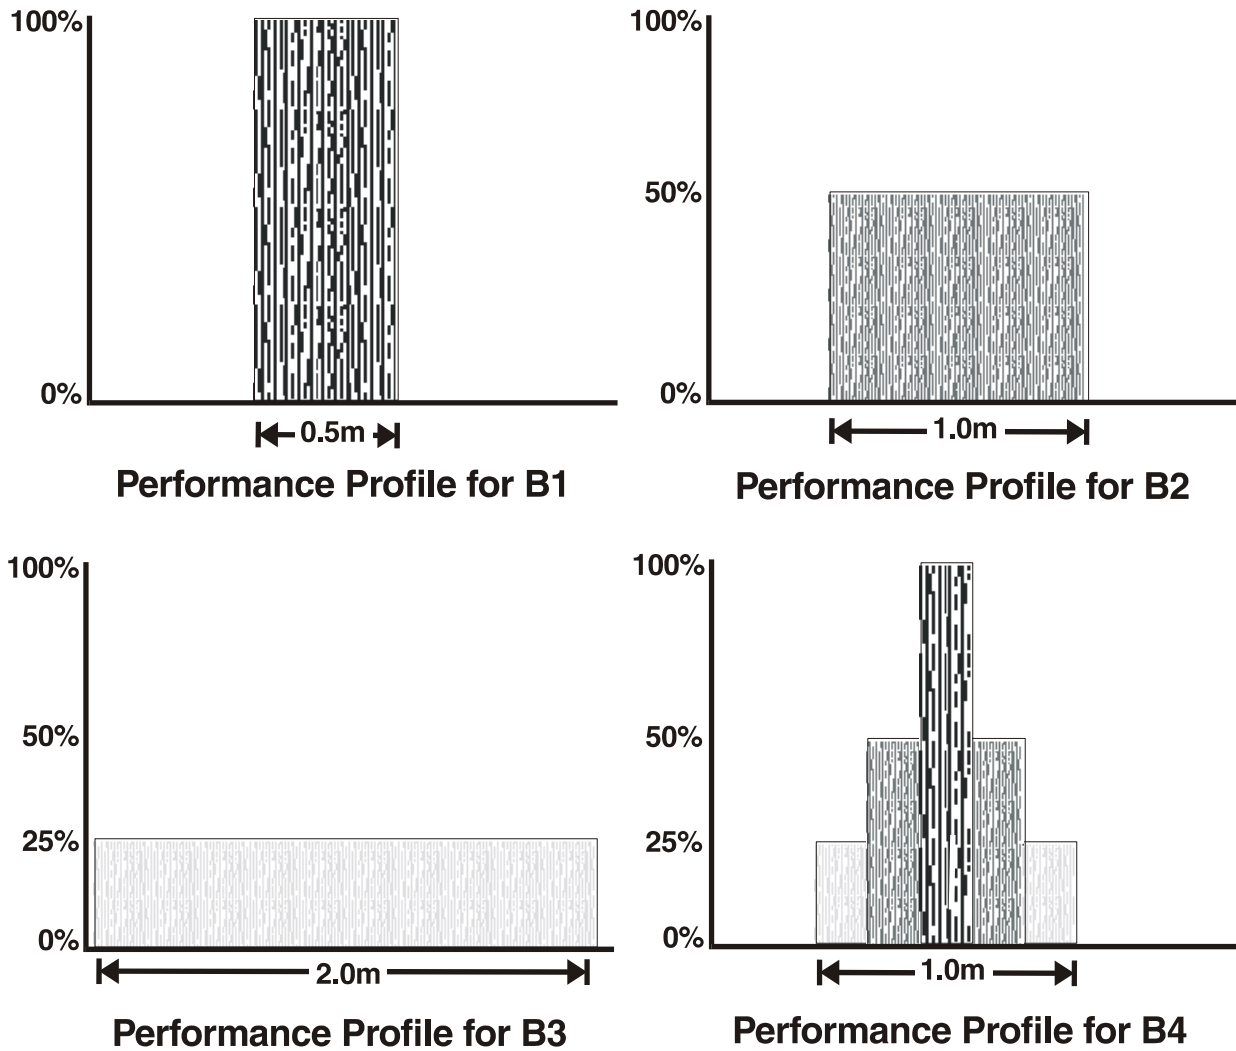


Figure C6. Broom performance profiles.

When looking at how the four brooms performed, we see that all four swept up the same amount of sand at the same rate under the conditions of our experiment, even if each broom did so in a different way. How can we characterize their “equivalent” performance? Note that the amount of sand swept up by each broom (50g) is exactly the amount found in a strip 50 cm wide and 10 m long. In fact, it is easy to show that no matter how far each broom is pushed under these same conditions, it will sweep up the amount of sand found in a strip 50 cm wide over the length of the broom’s movement. That is, we can say the *effective sweep width* of each broom, for the purposes of computing the amount of sand swept up, is 50 cm (or 0.5 m). If we convert the percentages on the vertical axes of Figure C6 to decimal values (e.g., 100% = 1.0), the amount of area “under the curve” (the shaded areas in the figure) is exactly equal to the *effective sweep width* in each case. This is **not** a mere coincidence. In fact, this is one of several equivalent definitions of *effective sweep width*.

One of the alternative, but equivalent, definitions is that the *effective sweep width* equals the width of the swath where the amount of sand left behind equals the amount swept up outside that swath in one pass over the floor. It is easy to confirm mentally without computation that this is the case for brooms B1 and B2. Now consider broom B3. In a central swath 50 cm wide, it

leaves behind 75% of the sand or 37.5 grams. Over the remaining 150 cm, consisting of two 75 cm swaths either side of the central 50 cm swath, it sweeps up 25% of the sand or $150 \text{ g} \times 0.25 = 37.5$ grams. It takes a little more computation, but a similar analysis of broom B4's performance will also agree with the result obtained by weighing the amount of sand swept up.

The results of our experiments and some values of interest that may be computed from them are shown in Table C1 below. Although the utility of some of the computed values may not be immediately apparent, their usefulness will become clear in the search planning process.

Table C1
Broom Experimental Results

	Broom B1	Broom B2	Broom B3	Broom B4
Broom Width	0.5 m	1.0 m	2.0 m	1.0 m
Maximum Lateral Range	0.25 m	0.5 m	1.0 m	0.5 m
Bristle Density	Dense	Less dense	Much less dense	Composite
Broom Effectiveness (avg.)	100 %	50 %	25 %	50%
Sand "Density"	10 g/m ²	10 g/m ²	10 g/m ²	10 g/m ²
Sweeping Speed	0.5 m/sec	0.5 m/sec	0.5 m/sec	0.5 m/sec
Time	20 sec	20 sec	20 sec	20 sec
Distance Moved	10 m	10 m	10 m	10 m
Area Swept	0.5 m x 10 m	1.0 m x 10 m	2.0 m x 10 m	1.0 m x 10 m
Amount of Sand Swept Up	50 g	50 g	50 g	50 g
Average Sand Removal Rate	2.5 g/sec	2.5 g/sec	2.5 g/sec	2.5 g/sec
<i>Effective Sweep Width</i>	0.5 m	0.5 m	0.5 m	0.5 m
<i>Area Effectively Swept</i>	0.5 m x 10 m	0.5 m x 10 m	0.5 m x 10 m	0.5 m x 10 m
<i>Effective Sweep Rate</i>	0.25 m ² /sec	0.25 m ² /sec	0.25 m ² /sec	0.25 m ² /sec

Although strictly speaking the results tabulated above are valid only for situations that are exactly like our experiment, *effective sweep width* tends to be relatively stable and not prone to sudden large variations as conditions change. A small change in the search situation produces only a small change in *sweep width*. Therefore, the results of tests performed for a typical search situation are useful for a fairly large range of similar situations. Furthermore, it is probably more practical and less error-prone for search planners to subjectively adjust the *sweep width* value determined by experiment for a known situation to a larger or smaller estimated value for a different situation than to subjectively estimate POD values directly based on no data at all.

In our floor-sweeping analogy of detection, the different brooms represented different sensors, the sand on the floor represented probability, the sweeping action represented the detection process, the amount of sand swept up represented the amount of probability "removed" by searching and the amount of sand left behind represented the probability that still remained after searching.

Importance of Sweep Width

Koopman (1946) defined the *effective search (or sweep) width* in his groundbreaking work on search theory. In the ensuing years right up to the present, it has withstood the tests of time, much scientific scrutiny, and a great deal of operational usage, especially in search and rescue.

Sweep width is a basic, objective, quantitative measure of *detectability*. Larger *sweep widths* are associated with situations where detection is easier while smaller *sweep widths* imply detection is more difficult. It should be clear that it must be important to know, in some quantitative way, how detectable the search object is in a particular search situation if we are to reliably estimate the probability of detecting that object, assuming it is present, with a given amount of searching.

The concept of *effective sweep width* is extremely powerful and lies at the very core of applied search theory.

The *sweep width* concept is extremely robust and extremely practical. An important property of *sweep width* is its relative independence from the details of the detection processes themselves, such as the exact shape of the detection profile, or exactly how the searcher's eyes and brain function to see and recognize the search object. In fact, *sweep width* integrates the effects of all the myriad factors affecting detection in a given situation into a single numeric value that is then easy for search planners to use. *Sweep width* is simply a measure (or estimate) of the average detection potential of a single specific "resource" (e.g., a person on the ground, an aircraft or vessel and its crew, etc.) while seeking a particular search object in a particular environment. Thus the concept may be applied to any sensor looking for any object under any set of conditions. For visual search, it will work for either relatively unobstructed views, such as searches conducted from aircraft over the ocean, or situations where obstructions are common, such as searching in or over forests. That is, *sweep width* may be applied to any SAR search situation, although it makes more sense to apply single *sweep width* values to situations where conditions are roughly uniform. Where there is a significant difference in environmental conditions (e.g., open fields vs. forests), sensor/searcher performance (e.g., trained vs. untrained searchers) and/or search objects (e.g., a person vs. "clues" like footprints or discarded objects), there will normally be a significant difference in *effective sweep width* as well. Where differences in these factors are small, the difference in sweep width will also be small.

Appendix D

Probability Density Distributions and Probability Maps

A *probability density distribution* is usually represented by a *probability map* consisting of a regular grid. For the purposes of this discussion, we define a regular grid as one that forms geometrically identical square cells. Each cell is then labeled with its *POC* value. Since all cells are equal in size, a cell's *POC* value is proportional to its *Pden* value. This type of display has the dual advantages of showing at a glance both how much probability each cell contains and where the highest probability densities lie. Although the *POC* and *Pden* values are not numerically equal, a cell with twice the *POC* value of another cell also has twice the *Pden* value of that other cell when a regular grid is used. Figure D1 is an example of a *probability map*.

3.23%	3.23%	6.45%	6.45%	6.45%
3.23%	3.23%	6.45%	9.68%	12.90%
3.23%	3.23%	9.68%	9.68%	12.90%

Figure D1. Probability map.

To determine where to search, we must first estimate where the lost or missing person could be. This requires a careful, deliberate, thoughtful assessment of all the available information as well as the continual seeking of additional information from all possible sources. “Available information” is an all-inclusive term referring to every scrap of evidence and data that might shed some light on the lost person’s probable locations. In addition to data about a specific incident, statistical data from similar situations, such as lost person behavior profiles, can be very useful. Historical data can also be useful, especially in popular recreational areas.

In SAR situations, data is frequently obtained from a variety of sources and is often inconsistent. However, such data also tends to form a number of self-consistent sets that each tell a “story” about what might have happened and where the lost person might be. These “stories” are called *scenarios*. Careful analysis of each scenario is then required to estimate the lost person’s probable locations if that scenario is true. These estimates are then quantified as *probability maps*, thus defining that scenario’s *probability density distribution*. The different scenarios are then subjectively “weighted” according to the search planner’s perceptions of their relative accuracy, reliability, importance, etc. and their *probability maps* are then combined appropriately. Probability maps for different scenarios are generally combined by computing, for each cell in an area large enough to include all scenarios, the weighted average (using subjective scenario weights) of the cell probabilities from each scenario.

Unfortunately, formal search theory does not shed much light on how go about turning an inconsistent body of evidence and data from a variety of sources into numbers on a probability map. As Stone (1983), one of the world's leading authorities on search theory and its practical application, observes, "One of the greatest difficulties in generating prior [to searching] probability maps is the lack of systematic, proven techniques for eliciting subjective inputs for search scenarios." He goes on to say, "In addition to obtaining subjective probabilities, we also have the problem of obtaining subjective estimates of uncertainties, times, and other quantitative information needed to form scenarios" (pp. 213-214).

Scenario development and analysis is a complex, difficult, mentally demanding task requiring a good deal of concentration, attention to detail, and mental discipline. Appropriate resources should be dedicated to this task and insulated from the often frenetic, and always distracting, operational activities. This frequently seems difficult to do in SAR situations. The first impulse is to get as much search effort as possible into the field as soon as possible because statistics show that a lost person's chances for survival decrease rapidly as time passes. While there is nothing wrong with mounting a large initial effort (provided more effort is on the way) based on only a cursory evaluation of the situation, too often this is not followed up with a more deliberate evaluation and planning effort for subsequent searching should the initial efforts fail. In a few publicized cases, it appears that lost persons who could have, and should have, been saved were not found in time – sometimes in spite of huge expenditures of effort in relatively limited areas. This appears to have been a result, at least partially, of poor analysis and planning.

Probability Density and Its Importance

To understand why *probability density* is important, we will return to our floor-sweeping analogy where the density of sand covering the floor is comparable to *probability density* in a search situation. We must also briefly jump ahead to *optimal effort allocation*. We will begin by extending our floor-sweeping analogy to a situation more complex than any we have discussed so far.

Consider a school gymnasium with a clear floor space measuring 50 meters by 30 meters for an area of 1,500 square meters (m^2). Suppose we divide the floor into four regions of unequal sizes so that region R1 covers 600 m^2 , R2 covers 400 m^2 , R3 covers 300 m^2 and R4 covers 200 m^2 . Suppose we cover each region uniformly with sand at the densities (in grams per square meter (g/m^2) of floor space) shown in the third column of Table D1. The values in the last two columns were computed from the corresponding area and density values in the second and third columns. Figure D2 illustrates the situation.

Table D1
Gymnasium Floor Values

Region	Area (m ²)	Density of Sand (g/m ²)	Amount of Sand Contained (kg)	Percentage of Sand Contained
R1	600	20	12	54.55%
R2	400	15	6	27.27%
R3	300	10	3	13.64%
R4	200	5	1	4.55%
Totals	1500	14.67	22	100.00%




Figure D2. Gymnasium floor illustration.

Suppose we have only one sweeper, whose broom is B2 from our *sweep width* experiments (see Appendix C, Figure C3) and whose rate of motion anywhere in the gym is 0.5 m/sec (30 m/min) regardless of the density of the sand. Finally, suppose our lone sweeper is available for only five minutes. If we wish for our sweeper to remove the greatest possible amount of sand in the time available, where should the sweeping be done?

In five minutes, the sweeper can move the broom a distance of 150 meters. In other words, the available *effort* is 150 m. Since broom B2 is one meter in width, the sweeper could sweep an area of 150 m². This is less than the area of any of the four regions. However, all other things being equal, the most productive place to sweep will be R1 because that is where the sand is most densely spread. Recall that broom B2 is uniformly 50% effective across its one-meter width and therefore has an *effective sweep width* of only 50 cm (0.5m). Recalling the equation from section 2.2.1.2,

$$Area\ Effectively\ Swept = Effort \times (Effective\ Sweep\ Width)$$

the *area effectively swept* in five minutes is computed to be 150m × 0.5m or 75 square meters. From the equation in section 2.2.1.3,

$$\text{Coverage} = \frac{\text{Area Effectively Swept}}{\text{Physical Size of the Swept Area}}$$

a *coverage* of $75\text{m}^2/150\text{m}^2$ or 0.5 is computed for the swept area. If the sweeper uses perfectly straight, parallel tracks at a spacing of one meter, Figure C3 in Appendix C shows B2 will sweep up 50% of the sand initially present in 150 m^2 of R1, or about 1.5 kg. Sweeping one-fourth of region R1 in this manner will sweep up more sand in less time than any other application of the same *effort* within the gymnasium. This is true because the *density* of the sand in R1 is higher than anywhere else, and it is tacitly assumed the *effective sweep width* and *speed* (i.e., the *effective sweep rate*) will be the same everywhere. The unwary could fall into a trap at this point by jumping to the conclusion that *density* is the only variable that needs to be considered. As we will see, the objective is to sweep up as much sand as possible in the least amount of time, taking into consideration any and all differences in **both** *density* and *effective sweep rate* from one region to another. It is the combined effect of these two variables that determines where sand can be swept up most quickly.

Note that although R1 also contained the most sand, it was the high *density*, not the high percentage of sand contained in the region, that caused sweeping there first to be more productive than anywhere else. In other words, when deciding where to place *effort*, the *density* of sand covering the floor in a region is far more important than the amount of sand contained there. Therefore, how the *density* of sand is distributed over the gymnasium floor will have a great deal to do with how the available *effort* should be distributed over the floor in order to sweep up the maximum amount of sand. Although *density* is not the only factor to consider when making effort allocation decisions, this brief example shows that it plays a major role.

Creating Probability Density Distributions

As mentioned previously, constructing a *probability density distribution* from the available information and evidence can be a difficult undertaking. In some cases, however, it is reasonable to assume a standard type of *probability density distribution*. We will briefly describe two such distributions and then return to the more general problem.

Circular Normal Probability Density Distributions

When a distressed aircraft flying over a remote area or a distressed vessel at sea reports its position, the known characteristics of navigation make it reasonable to assume the **actual position** may be some distance from the **reported position** (at least this was true before GPS receivers became so readily available). Analyses of these characteristics have shown that the actual positions often have a *circular normal probability density distribution* centered on the reported position. (Actually, the more general elliptical *bivariate normal distribution* is more correct, but the *circular normal* is a satisfactory example for this discussion.) For the mathematically inclined, the amount of probability contained (*POC*) in a circle drawn about the center of this type of distribution is given by

$$POC = 1 - e^{-\frac{R^2}{2}}$$

where e is the base of the natural logarithms ($\cong 2.71828$) and R is the radius of the circle in standard deviations (σ). (Note that for a circular normal distribution, the amount of probability contained within one standard deviation of the mean (center) is only about 39%, as compared to about 68% for the more familiar one-dimensional “bell curve.” Readers who want more information about the statistics of bivariate (two-dimensional) data are encouraged to consult a standard text on statistics.)

The radius for which the POC is 50% is defined by statisticians as the *probable error* of the position. The *probable error* defines the size of the circle where the chances of the actual position being inside the circle equal the chances of it being outside the circle. If we center a regular grid on the reported position and compute the amount of probability contained in each cell, we get a *probability map* like that shown in Figure D3, where the radius of the dashed circle is the *probable error*. The circle contains 50% of the probability. The other 7.91% contained in the center cell comes from the area that is outside the circle but inside the cell in the four corners.




Figure D3. Probability map showing probable error.

Although situations where this type of distribution would apply are relatively rare in inland SAR (e.g., the forced landing of an aircraft in a remote area), they are much more common in maritime SAR. Whenever it does apply, the search planner can estimate the *probable error* of a reported position and use Figure D3 (or a version with a finer grid) scaled to match the appropriate charts or maps, to plan the search. Of course, it might be necessary to adjust both the reported position and the size of the probable error based on such factors as the glide characteristics of the distressed aircraft or the drift characteristics of a life raft from a ship that sank.

Uniform Probability Density Distributions

Suppose the pilot of an aircraft issues a mayday call giving his tail number but no position. Assume checking the flight plan reveals that the aircraft was supposed to be engaging in a biological survey of a known wilderness area at the time, but no specific flight path was given. If no

other information is available, the search planner has little choice but to regard all parts of the area as equally likely to be the site of the distress. This means the *probability density* is uniformly distributed over the area. Figure D4 shows a *probability map* for a uniform *probability density distribution*.

5%	5%	5%	5%	5%
5%	5%	5%	5%	5%
5%	5%	5%	5%	5%

Figure D4. Probability map for a uniform probability density distribution.

Generalized Probability Density Distributions

Although resorting to a “standard” *probability density distribution* is the easiest way to generate a *probability map*, it is not always possible to find one that adequately describes what the available evidence indicates about where the search object may be located. This is a very common situation in inland SAR right from the start. Even in maritime cases, what may have started out as a “standard” distribution often becomes generalized rather quickly due to the vagaries and uncertainties of oceanic drift. The Coast Guard addresses this problem via its Computer Assisted Search Planning (CASP) system. CASP takes both the known variations in winds and current from one place and time to another and their respective *probable errors* into account. CASP then computes tens of thousands of independent drift trajectories using this data. The end result might look something like the *probability map* shown in Figure D5.

5%	12%	6%
14%	18%	16%
9%	2%	7%
6%	1%	4%

Figure D5. Probability map after accounting for multiple probable errors.

Estimating Probability Densities

Although formal search theory provides methods for optimally allocating effort once a *probability density distribution* has been defined, it does not shed much light on how to evaluate evidence, clues, historical data, lost person behavior profiles, etc., and use those evaluations to **create** a corresponding *probability density distribution*. While we cannot offer much guidance at this point about assessing the available information and data, we can examine some possible methods for assigning numeric values to those assessments.

Let us return to the gymnasium floor described above and shown in Figure D2. We now obtain an undistorted photograph of the entire floor from a point directly above its center and make three copies. Like Figure D2, there is enough contrast for a person to discern the four regions and the fact that the density in R1 is greater than that in R2 which is greater than that in R3 which is greater than that in R4. Finally, we arrange to have three floor sweepers, Tom, Dick and Mary, participate in some experiments.

Clearly, this is not a very realistic analogy for the kind of evidence a search planner would have to evaluate. Nevertheless, the examples that follow will provide some valuable insights into certain kinds of problems that can arise when attempting to translate assessments into *probability maps*.

Estimating Containment Percentages Directly

We begin by showing Tom (in isolation from the others) one of our photographs. We ask him to mark off the four regions and estimate what fraction of the sand is in each. We will call this fraction the *percentage of containment (poc)*. Tom will likely regard this as a difficult assign-

ment. It is clear that R1 covers a little less than half the floor's area but it is also clear that the sand is more dense there than anywhere else. Tom must weigh both factors when making his estimate. Table 2 summarizes Tom's estimates of how much sand, as a percentage of the total, each region contains. Compare the estimated percentages and the computed amounts and densities to the corresponding quantities in Table D1.

Table D2
Tom's Assessment

Region	Area (m ²)	Estimated <i>poc</i>	Computed Amount of Sand (kg)	Computed Density (g/m ²)
R1	600	50%	0.50 x 22 = 11.0	11,000/600 = 18.33
R2	400	30%	0.30 x 22 = 6.6	6,600/400 = 16.50
R3	300	15%	0.15 x 22 = 3.3	3,300/300 = 11.00
R4	200	5%	0.05 x 22 = 1.1	1,100/200 = 5.50
Totals	1500	100%	22.0	22,000/1500 = 14.67

The estimated *percentages of containment*, though imperfect, are actually very good, producing *densities* that are reasonably accurate and in about the correct relationship to one another. It could be shown that using these *densities* would cause a less-than-optimal level of *effort* to be assigned to region R1, and more-than-optimal amounts of *effort* to be assigned to the other three regions. (In this context, an "optimal" allocation of *effort* is the one that causes the greatest amount of sand to be swept up in the shortest amount of time.) Although the resulting sweeping (search) plan would be sub-optimal, it would not be dramatically so.

Ranking the Regions

We now call in Dick, give him one of our photographs, and ask him to mark off the four regions. We then ask him to rank the regions, using letters, by the amount of sand each one contains. Since there are four regions and it is pretty obvious all contain different amounts of sand, Dick chooses to use the letters A through D, with A denoting the region with the most sand. Dick finds this a very easy task, and his rankings, along with the percentages and densities they imply are shown in Table D3.

Table D3
Dick's Assessment

Region	Letter Designation	Numeric Rank	Computed <i>poc</i>	Computed Amount of Sand (kg)	Area (m ²)	Computed Density (g/m ²)
R1	A	4	4/10 = 40%	0.4 x 22 = 8.8	600	8,800/600 = 14.67
R2	B	3	3/10 = 30%	0.3 x 22 = 6.6	400	6,600/400 = 16.50
R3	C	2	2/10 = 20%	0.2 x 22 = 4.4	300	4,400/300 = 14.67
R4	D	1	1/10 = 10%	0.1 x 22 = 2.2	200	2,200/200 = 11.00
Totals		10	10/10 = 100%	1.0 x 22 = 22.0	1500	22,000/1500 = 14.67

Although the percentages reflect Dick's ranking, they are not very accurate. The computed *densities* are also inaccurate. As a result, the values computed from Dick's ranking fail to represent the photographic evidence and also fail to approximate the actual values as closely as Tom's estimates in three of the four regions. Although the simple ranking method was very easy in this case, we must conclude that it did not produce valid *densities* on which to base an optimal sweeping (search) plan. Clearly, there is something wrong with this technique.

Ranking the Regions – Again

We now call in Mary and present her with the same problem as Dick, (i.e., ranking by letters). We want to see if the difficulty we just experienced will repeat itself. She marks the boundaries of the four regions on the photograph but then goes a step further. She draws a grid on the photograph that is three cells wide by five cells long dividing the floor into 15 square cells of equal size. Conveniently, each region is comprised of a whole number of cells. She then ranks each cell using the same four-letter ranking scale Dick used. Each cell in region R1 is ranked as "A," each cell in R2 is ranked as "B," each cell in R3 is ranked as "C" and each cell in R4 is ranked as "D" as shown in Figure D6.

A	A	B	B	B
A	A	B	C	D
A	A	C	C	D

Figure D6. Probability map using four-letter ranking scale.

Grouping the cells by region, she gets the results shown in Table D4.

Table D4
Mary's Assessment

Region	Letter Rank	Numeric Rank	Computed Cell <i>poc</i>	Computed Region <i>poc</i>	Computed Amount of Sand (kg)	Computed Density (g/m ²)
R1	6 x A	6 x 4 = 24	4/44 = 9.09%	6 x 9.09 = 54.55%	0.5455 x 22 = 12	12,000/600 = 20
R2	4 x B	4 x 3 = 12	3/44 = 6.82%	4 x 6.82 = 27.27%	0.2727 x 22 = 6	6,000/400 = 15
R3	3 x C	3 x 2 = 6	2/44 = 4.55%	3 x 4.55 = 13.64%	0.1364 x 22 = 3	3,000/300 = 10
R4	2 x D	2 x 1 = 2	1/44 = 2.27%	2 x 2.27 = 4.55%	0.0455 x 22 = 1	1,000/200 = 5
Totals		44		100.00%	22	22,000/1500 = 14.67

At first glance, it appears Mary may have stumbled upon a perfect method since the regional *percentages of containment*, amounts of sand and *densities* computed from her assessments are all exactly correct! Further consideration may indicate that she was just lucky. The numeric values assigned to the letters in our ranking scale happen to be exactly proportional to the actual cellular *percentages of containment*. Multiplying each of the numeric ranking values (4, 3, 2, and 1) by 2.27 produces the actual cell *poc* values (9.09, 6.82, 4.55, and 2.27). From another, equivalent, point of view, we can say the numbers 4, 3, 2 and 1 are in the same relationship to one another as the different cell percentages (e.g. $9.09/6.82 = 4/3$).

It is worthwhile at this point to note the relationship of the ranking values to the *densities*. Multiplying each of the ranking values (4, 3, 2, and 1) by five produces the *density* values (20, 15, 10, and 5). This means these two sets of values are also proportional to one another, just as in the case of the cellular *percentages of containment*. This in turn means Mary could have used any smaller grid size she liked (e.g., one with 5 m x 5m cells), assigned letter values to each in the same way (e.g., 24 A's, 16 B's, etc.) and obtained the correct results for regional *percentages* and *densities*. She also could have dispensed with the grid altogether and used the areas of the regions in place of the number of cells in Table D4.

From Mary's assessment using a regular grid of cells, we may produce a "map," like that in Figure D7, showing how the sand is distributed. Note that on this "map", higher *percentages* imply proportionately higher *densities*.

9.09%	9.09%	6.82%	6.82%	6.82%
9.09%	9.09%	6.82%	4.55%	2.27%
9.09%	9.09%	4.55%	4.55%	2.27%

Figure D7. Probability map based on Mary's assessment.

Mary's good fortune illustrates an important lesson for search planning: Whenever an assessment value is assigned to a subdivision of the *possibility area*, that value must be proportional, in

a precise mathematical sense, to the subdivision's *probability of containing* the search object. Similarly, the assessment values must reflect the correct relationships among the subdivisions. If one subdivision is assessed as an "8" and another as a "4," the implication is that the first subdivision is twice as likely to contain the search object as the second. If the evaluator does not agree with this implication, then he has chosen one or both values incorrectly.

An Assessment Based on Density Estimates

It might have been an interesting exercise to ask the sweepers to estimate, from the photograph, the relative *densities* in the regions instead of *percentages of containment*. Such estimates could have been applied to the areas of the regions to get estimates of the relative amounts of sand contained in each. Then, these relative amounts could have been used to compute the *percentages of containment*. The results might have been both more accurate and more consistent if this had been tried. For example, suppose an evaluator had estimated from the photograph that the *density* in region R3 was twice that of region R4, the *density* in R2 was three times that of R4 and the *density* in R1 was four times that in R4. Table D5 shows how the *percentages of containment* could be computed from these relative *density* estimates.

Table D5
Density-based Assessment

Region	Area (m ²)	Relative Density	Relative Amount of Sand	Computed <i>poc</i>
R1	600	4	600 x 4 = 2400	2400/4400 = 54.55%
R2	400	3	400 x 3 = 1200	1200/4400 = 27.27%
R3	300	2	300 x 2 = 600	600/4400 = 13.64%
R4	200	1	200 x 1 = 200	200/4400 = 4.55%
Totals	1500		4400	4400/4400 = 100%

Another Short Exercise

To show that an assessment method works in general if the assessment values accurately represent the relative proportions of the *percentages of containment*, suppose we sweep the gymnasium floor clean and set up a new experiment as illustrated in Figure D8.




Figure D8. Gymnasium floor illustration; new experiment.

We will use the same regions and *densities* as before but distribute the sand as follows: 5 g/m² in R1, 10 g/m² in R2, 15 g/m² in R3 and 20 g/m² in R4. This means R1 will contain 3 kg of sand, R2 will have 4 kg, R3 will have 4.5 kg and R4 will have 4 kg for a total of 15.5 kg. Knowing the previous four-letter scale produces numbers that are in the correct proportions for these *densities* when using Mary’s cellular method, we can use these letters again with confidence to produce Table D6.

Table D6
Cellular Assessment of Figure D8

Region	Letter Rank	Numeric Rank	Computed Cell <i>poc</i>	Computed Region <i>poc</i>	Computed Amount of Sand (kg)	Computed Density (g/m ²)
R1	6 x D	6 x 1 = 6	1/31 = 3.23%	6 x 3.23 = 19.36%	0.1936 x 15.5 = 3.0	3,000/600 = 5
R2	4 x C	4 x 2 = 8	2/31 = 6.45%	4 x 6.45 = 25.81%	0.2581 x 15.5 = 4.0	4,000/400 = 10
R3	3 x B	3 x 3 = 9	3/31 = 9.68%	3 x 9.68 = 29.03%	0.2903 x 15.5 = 4.5	4,500/300 = 15
R4	2 x A	2 x 4 = 8	4/31 = 12.90%	2 x 12.90 = 25.81%	0.2581 x 15.5 = 4.0	4,000/200 = 20
Totals		31		100.00%	15.5	15,500/1500 = 10.33

Note that it would be more difficult to apply a simple ranking system to this distribution than the previous one because it is much less obvious which region contains the most sand and which contains the least. However, even if we use the correct regional *poc* values from Table D6 as the basis for a simple ranking, the results will be inaccurate. Table D7 shows the *percentages*, amounts of sand, and *densities* that would be computed from such a simple ranking. Compare these to the correct values in Table D6 above.

Table D7
Simple Ranking Assessment of Figure D8

Region	Letter Designation	Numeric Rank	Computed <i>poc</i>	Computed Amount of Sand (kg)	Computed Density (g/m ²)
R1	C	1	1/8 = 12.5%	0.125 x 15.5 = 1.9375	1,937.5/600 = 3.23
R2	B	2	2/8 = 25.0%	0.250 x 15.5 = 3.8750	3,875.0/400 = 9.69
R3	A	3	3/8 = 37.5%	0.375 x 15.5 = 5.8125	5,812.5/300 = 19.38
R4	B	2	2/8 = 25.0%	0.250 x 15.5 = 3.8750	3,875.0/200 = 19.38
Totals		8	100.0%	15.5000	15,500/1500 = 10.33

We must again emphasize that if assessment values are to produce accurate and valid *probability of containment* (*POC* or *POA*) estimates, the value assigned to each region, cell, segment, or any other subdivision of the search area, must be **mathematically proportional** to that subdivision's *probability of containment*. Stated another way, the assessment values assigned to the various subdivisions must be in the correct proportions to one another across the search area as a whole.

Analysis of Results

Tom had difficulty coming up with correct values because he had to mentally estimate *percentages of containment* by balancing the sizes of the regions against their apparent relative *densities*. Nevertheless, he was able to produce reasonably satisfactory results for this very simple problem. It is unlikely he would do as well with a more complex situation, such as that represented by Figure D8.

Dick's simple rankings produced unsatisfactory estimates of both *percentages of containment* and *densities*. A simple ranking does not address the essential proportionality relationships needed for estimating these values. Therefore, simple ranking systems should not be used since they produce inconsistent and misleading results.

Mary solved Tom's problem with unequal areas by using a regular grid. A grid worked well for this problem, but grids may not work as well in situations where irregular geographic features are a significant factor in assessing where the lost person is likely to be. Because Mary was also fortunate enough to be using assessment values that were in the same proportions as the actual *densities* (and cellular *percentages of containment*), her results were exactly correct. In a sense, Mary was not ranking the cells as much as she was *rating* them on a scale of 1 to 4 – a scale that happened to provide exactly the values she needed.

Proportional Assessment

Since correct proportionality is so important, we need a procedure for making proportional assessments that is more dependable than Mary's happy accident. One such procedure is for each evaluator to decide which region contains the most sand (probability) and then rate all other regions against this "standard." For example, suppose Dick had *rated* the *regions* of Figure D2 on

a scale of, say, 1 to 10 with R1 being assigned a value of 10. If he then decided that R2 contained a little more than half as much sand as R1, he might have rated it with a value of 6 (i.e., as containing about 60% as much sand as R1). Similarly, he might have rated R3 with a value of 3 (30% as much sand as R1) and R4 with a value of 1 (only 10% as much sand as R1). If Dick had chosen these *proportional assessment* values, his results would have been much closer to the actual values shown in Table D1. In fact, his results would have been identical to Tom’s in Table D2, as shown in the table below.

Table D8
Proportional Rating Assessment of Figure D2

Region	Area (m ²)	Proportional Assessment	Computed <i>poc</i>	Computed Amount of Sand (kg)	Computed Density (g/m ²)
R1	600	10	10/20 = 50%	0.50 x 22 = 11.0	11,000/600 = 18.33
R2	400	6	6/20 = 30%	0.30 x 22 = 6.6	6,600/400 = 16.50
R3	300	3	3/20 = 15%	0.15 x 22 = 3.3	3,300/300 = 11.00
R4	200	1	1/20 = 5%	0.05 x 22 = 1.1	1,100/200 = 5.50
Totals	1500	20	100%	22.0	22,000/1500 = 14.67

For Figure D8, using the same 10-point scale and *proportional assessments* of 6, 8, 10, and 8 for R1 – R4 respectively would have produced regional *poc* values of 18.75%, 25%, 31.25% and 25% respectively. These are very close to the correct values shown in Table D6. (The reader is encouraged to verify these figures and compute the amounts of sand and densities as an exercise.) It is important to understand that simply sorting the regions into a list in descending order of *percentage of containment* does not provide enough information to reliably estimate what those *percentages* are.

Obtaining meaningful *Probabilities of Containment* REQUIRES the use of a Proportional Assessment Technique.

Another way to solve the problem of unequal areas, from a mathematical standpoint at least, is to use a proportional assessment technique to estimate the relative *densities* and use them in conjunction with the regional areas to compute percentages of containment. Table D5 above illustrated how this could be done.

Containment vs. Density Estimates

It is important at this point to reconsider the question posed earlier: If two regions of different sizes are each assessed as being “very likely” to contain the search object, does it mean

- a) their *probabilities of containment* are both equally high
- or
- b) their *probability densities* are both equally high?

When an evaluator believes a particular portion of the search area is “very likely” to contain the search object he could mean one of two things:

- i) Considering all pertinent data, this portion of the search area is very likely to contain the search object **irrespective of its size** as compared to the other portions. In this case, he is estimating a *relative probability of containment*.
- ii) Considering all pertinent data, this portion of the search area is very likely, **relative to its size**, to contain the search object as compared to the other portions in relation to their sizes. In this case, he is estimating *relative probability density*.

When it comes to computing *probability densities* for use in the optimal allocation of effort, the distinction between these two interpretations is of paramount importance. A small portion of an area may have a high *probability density* and a low *probability of containment*. On the other hand, a large portion may have a low *probability density* but a high *probability of containment*. A small portion with a high *probability of containment* will necessarily have a high *probability density*. Similarly, a large portion with a high *probability density* will necessarily have a high *probability of containment*. It is easy to become confused, and it is necessary to take conscious steps to avoid such confusion. It all boils down to exactly how the evaluator accounts for differing sizes among the regions, segments, etc., comprising the search area. The evaluator's mode of thinking (*containment* vs. *density*) may in turn depend on the nature of the available information. When using a regular grid or other arrangement where all the basic subdivisions of the search area have the same size, the evaluator is freed from this potential point of confusion. In this situation, an estimate of the *relative probability densities* is also an estimate of the *relative probabilities of containment* and vice versa.