
Principles of Search Theory

By J. R. Frost

Principles of Search Theory

Part I: Detection

by J. R. Frost

Introduction

This is the first in a series of four articles that, as a group, attempts to present the fundamental principles of search theory in layman's terms (parts I and II appear in this issue; parts III and IV will appear in the next issue of Response). Collectively, these articles are intended to ground the reader in some of the basic principles and terminology of search theory in an easy-to-understand manner. While these articles may include some comments on how aspects of search theory relate to SAR, practical application (such as specific search procedures based on the theory) is beyond the scope of this particular series. Before we can begin, however, we must lay down some ground rules, express some caveats, and make some disclaimers.

The principles of search theory described in this series of articles have been established by the scientific community over the last 50 years, and may be found in various university-level textbooks and in scientific journals. Unfortunately, these sources express the principles of search theory in the language of higher mathematics, making them all but unreadable for nonmathematicians. The objective of these articles is to remove this impediment to understanding the basic concepts of search theory by translating the mathematics into analogies that are easier to grasp.

Most of the terms used in this series of articles are taken from either the scientific search theory literature itself or the *U.S. National Search and Rescue Manual*. (Where there is overlap, these two sources are consistent.) Terminology is important for understanding search theory's basic principles. However, many terms used herein have also been used elsewhere with different, and sometimes vague or even contradictory, meanings. Therefore, to gain a full appreciation of the material being presented, the reader may need to set aside familiar concepts and definitions from other informal discussions of "search theory."

Although the author has striven for clarity and simplicity, it should be no surprise if several careful re-readings and some computation are necessary to grasp the concepts involved. Search theory is not simple and intuitive; many of its concepts are difficult to understand after only an initial exposure. Readers who have been exposed to other treatments of this subject that may not have adhered as strictly to precise terminology or the underlying scientific research are likely to have the most difficulty.

Finally, the principles of search theory represented here are not the opinions (or theories) of this author. The author's role is merely that of translator and messenger. Some readers

may see challenges to cherished concepts that have come to be accepted as "conventional wisdom." Addressing these situations would require far more than the available space. With this in mind, comparisons with other informal descriptions of "search theory" will be deliberately avoided and we will confine ourselves to what may be obtained from the scientific literature.

History

The theory of how to search for lost, missing, hidden and even evasive objects has been a subject of serious scientific research for more than 50 years. It is a branch of the broader applied science known as operations research. The term operations research can be traced to work done during World War II in support of the war effort. At that time, operations research was an apt title since the objective was to find the most efficient and effective ways for conducting military operations. During the war, one important type of military operation was, in fact, searching. Searches were conducted to locate the enemy, and to locate and recover one's own lost or missing personnel or those of one's allies. In more recent years, the principles of operations research have been applied to a wide variety of problems that involve making good decisions in the face of uncertainty about many of the variables involved. These problems often do not involve "operations" in the classical sense, and so the term operations research has become an anachronism to some extent. However, the original meaning is very close to the subject we want to discuss—namely, effective, efficient ways of searching for lost or missing persons.

B.O. Koopman^{1,2} did the initial work on search theory during World War II for the U. S. Navy. The Navy's primary search objects were enemy ships and submarines, and its own downed fliers adrift on the ocean. Koopman had to first develop the general principles of search theory before he could get down to the specifics of naval problems. These fundamental principles, which apply to all types of searches for lost or missing objects, are the principles we will be discussing.

A Search Analogy

To avoid descending too deeply into the pit of mathematics, we will need to discuss a common, easily visualized activity that can be used as a model, or analogy, for searching. 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 search theory.

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 1 shows a schematic representation of the four different designs. We construct the brooms and label them as B1, B2, B3, and B4, respectively.

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 (the width of the broom head), as illustrated in Figure 2. It swept up 50 grams of sand—all the sand within the $0.5 \text{ m} \times 10 \text{ m}$ 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 limita-

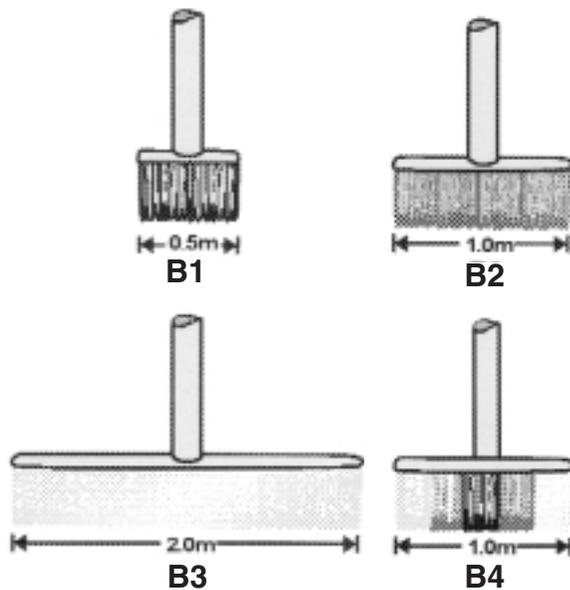


Figure 1



Figure 2

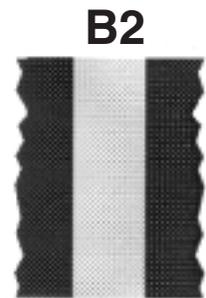


Figure 3

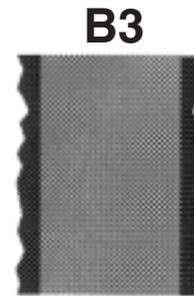


Figure 4

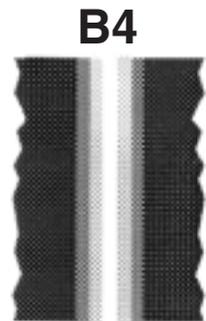


Figure 5

tion 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 3. 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 4. 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

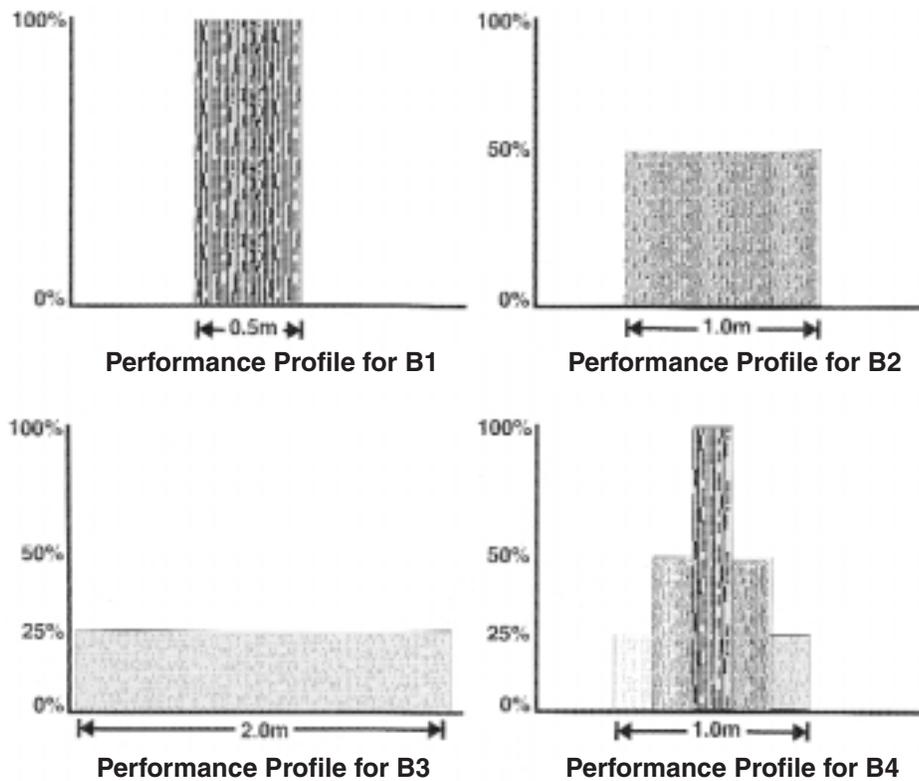


Figure 6

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 5 illustrates the uneven performance of broom 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 6.

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 when moving at 0.5 m/sec, 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 6 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. As we shall see, this is not a mere coincidence. The results of our experiments and some values of interest that may be computed from them are shown in the table below. Although the utility of some of the computed values may not be immediately apparent, their usefulness will become clear in the subsequent parts of this series.

The results tabulated below are valid only for situations that are exactly like our experiment. If we change the speed at which the brooms are pushed, use another surface (e.g., the asphalt in the parking lot), or use BBs instead of sand, we may or may not get the amount of sand (or BBs) found in a 50 cm swath along the tracks. Likewise, the four brooms may or may not continue to perform "equivalently" with respect to one another. We need a more general definition of effective sweep width for it to be useful.

We may define *effective sweep width* of a broom moving over the floor at a certain speed as the ratio of the amount of material swept up per unit time to the product of the density (amount per unit area) of material covering the floor and the broom's rate of travel. This definition is easier to grasp when written as an equation:

$$\text{Effective Sweep Width} = \frac{\text{Amount of Material Swept Up Per Unit Time}}{[(\text{Amount of Material Per Unit Area}) \times (\text{Broom Speed})]}$$

The term *amount of material* could mean any quantitative measure of the material, including grams of sand (as in our experiment), number of objects (such as number of BBs), volume of a liquid (e.g., for sponge mop evaluation), etc.

	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
Effective Sweep Width	0.5 m	0.5 m	0.5 m	0.5 m
Area Effectively Swept	0.5 m x 10 m			
Effective Sweep Rate	0.25 m ² /sec			

Also note that we are using "effectiveness" to mean "has the same effect as" according to some agreed-upon measurement (grams of sand swept up, in this case). We could have used the word "equivalent" in place of the word "effective" whose usage here is taken directly from the scientific literature. Readers are invited to substitute "equivalent" for "effective" if it makes the articles in this series easier to understand. The important thing to note is that the modifier "effective," as used here does not imply a broom is only, or even highly, effective over a swath having a physical width equal to the *effective sweep width*. When we say that all four brooms have an effective sweep width of 50 cm, we are saying that all four sweep up the amount of sand found on the floor in a swath 50 cm wide. Only broom B1 does this in a literal sense. All of the others sweep up the same amount of sand in one pass, but each removes the sand in its own way from a wider swath.

Effective Search (or Sweep) Width (W)

In his groundbreaking work on search theory, Koopman¹ defined the *effective search (or sweep) width* (often shortened to just *sweep width*) as follows: 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,

$$[1] \quad W = \frac{\text{Number of Objects Detected Per Unit Time}}{(\text{Number of Objects Per Unit Area}) \times (\text{Searcher Speed})}$$

where all values are averages over a statistically significant sampling period. If the *performance (or detection) profile* (called a *lateral range curve* in search theory) is known for a certain search situation, then the *area under the detection profile* equals the *sweep width*, W , for that situation. This *effective sweep width* is also twice the maximum detection range of an "equivalent" *definite range detection profile* (one that is 100% effective out to some definite lateral range either side of its track and completely ineffective beyond that range, like broom B1 in our floor-sweeping analogy). Here, "equivalent" means that the definite range detection profile

and the actual detection profile both detect, on average, the same number of objects per unit time under the same conditions of object density and searcher speed.

Some Sweep Width Examples

To see how Equation [1] works, suppose we devise an experiment where a large number of identical cardboard dummies, having about the same size, shape and color of a lost person, are uniformly, but randomly, distributed over a square test area in measuring one mile on a side. (A uniform random distribution is one where the

exact locations of the objects are chosen at random, but the number of objects per unit of area is about the same throughout the test area.) Since our test area is in hilly, forested terrain, we use a total of 1920 cardboard dummies for a density of 1920 objects per square mile. On average, that is 3 objects per acre or one object for every square patch of ground measuring 120.5 feet on a side. Now suppose we have several searchers each make a single straight pass through the test area and find that their average speed is 0.4 miles per hour with an average detection rate per searcher of 12.8 objects per hour. Using Equation [1],

$$W = \frac{12.8 \text{ objects per hour}}{1920 \text{ objects per square mile} \times 0.4 \text{ miles per hour}} = 0.0167 \text{ miles}$$

The *sweep width* is 0.0167 mile or about 88 feet for this particular search situation.

Suppose we conduct a similar experiment with several aircraft, each flying over the same area one time in a straight line at 100 miles per hour (about 90 knots). Suppose the average detection rate from the aircraft is 13.3 objects per minute. That converts to 800 objects per hour. (*Note: At 100 mph, each aircraft spent only 36 seconds crossing over our square test area and detected 8 objects on average.*) Using Equation [1],

$$W = \frac{800 \text{ objects per hour}}{1920 \text{ objects per square mile} \times 100 \text{ miles per hour}} = 0.0042 \text{ miles}$$

We find the sweep width for this situation is 0.0042 mile or about 22 feet. Note that despite the very high detection rate achieved by the aircraft due to their high rate of speed, the smaller sweep width we have computed indicates the cardboard dummies are significantly more difficult to detect from the air than from the ground. In fact, they are about four times as hard to detect from the air as from the ground. In some environments, detection from the air is so difficult that air search is considered ineffective. Nevertheless, since an aircraft can search an area several times in less time than it takes for a ground team to search the same area once, sometimes the handicap of a small sweep width may be more than

offset. (We will revisit this issue in Part II when we discuss something called coverage.) For many environments, it is not at all unusual for objects the size of a person to be much, much harder to detect from a light fixed wing aircraft flying a several hundred feet above the ground at 100 miles per hour than from the ground while walking at less than one-half mile per hour. On the other hand, sweep widths for aircraft operating over relatively flat, open terrain can be significantly larger than those for searchers on the ground when both are looking for the same object.

Note: The “experiments” and figures described above are completely hypothetical and are presented only for conceptual illustration. There is a great deal more involved in designing and conducting scientifically rigorous experiments, and collecting and analyzing the data from them, than the above paragraphs convey.

Importance of Sweep Width

As the reader may have already discerned, *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 with a given amount of searching.

The *sweep width* concept is extremely robust. It has stood the test of time and a great deal of scientific scrutiny. 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. *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, Equation [1] may be applied to any sensor looking for any object under any set of conditions. For visual search, note that Equation [1] 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, Equation [1] may be applied to any SAR search situation, although it makes more sense to apply it 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., experienced vs. inexperienced 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.

The concept of effective sweep width is extremely powerful and lies at the very core of search theory.

Factors Affecting Sweep Width

There are three classes of factors that affect detection and hence the *sweep width*.

1. The search object’s characteristics affecting detection. Examples include such things as the object’s size, color, contrast with surroundings, etc. One would not expect the sweep width for a discarded candy wrapper or a footprint in the summer forest to be nearly as large as that of a person wearing bright clothing.
2. The capabilities of the sensor(s) in use. Examples include sensor type (e.g., unaided human eye, infrared devices, air-scent dog, etc.), searcher/operator abilities (e.g., training, experience, fatigue), search platform (searcher on foot, all-terrain vehicle, boat, aircraft, etc.), speed of the searcher’s movement in relation to the search object, etc.
3. The environmental conditions at the place and time of the search. Examples include terrain, amount of ground cover, lighting conditions (e.g., sunny vs. overcast, deep forest vs. open meadow), visibility/weather (e.g., clear, foggy, rainy), etc.

All of these factors may interact with one another in complex ways. This leads to the single most important difficulty of the *sweep width* concept—there is no simple, easy way to directly measure *effective sweep width* in the field for each search situation. However, it can be estimated from factors that may be measured, or at least observed, directly. With the help of data from some scientifically designed and executed experiments covering a reasonably broad range of search situations, *effective sweep width* values can be estimated quickly and reliably based on the sensor(s) in use, the characteristics of the search object(s) and the environmental conditions in the search area(s). Note that although the maximum detection range is both measurable and affected by many of the same factors, this value alone does not reliably indicate how much detection will take place, whereas the effective sweep width does. Even when we know the maximum detection range, all we can say with certainty is that the sweep width can never be greater than twice its value.

The Coast Guard’s Experience

For more than 20 years, the U.S. Coast Guard has been conducting scientific experiments and analyses to develop tables of validated *effective sweep width* values for use in marine SAR. These experiments first identify the significant factors affecting detection and then go on to quantify the effects of the identified factors. It does not appear that any experiments of a similar sophistication or scope have been undertaken for the benefit of inland SAR. Coast Guard experience has shown that relatively few (but expensive, unfortunately) experiments, covering a representative cross-section of conditions typically encountered in SAR missions, produces useful results across a wide spectrum of SAR situations. Consequently, the *National SAR Manual*³ contains extensive

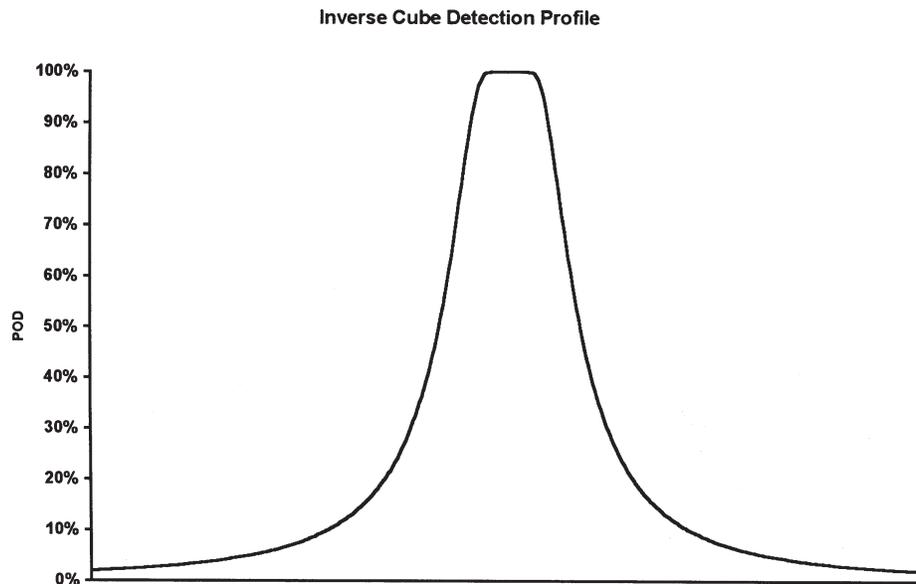


Figure 7

tables of *sweep width* values for a wide range of conditions encountered in marine SAR. From these, good estimates of *effective sweep widths* for marine situations other than those directly tabulated may be obtained quickly and easily by using correction factors, interpolation, limited extrapolation, etc. Maritime search planners the world over use these tables to good effect every day. Coast Guard experience has also shown the need for a significant level of continuing experimentation to keep pace with changing technology—both that which a distressed person might use and that which becomes available to the searchers.

Koopman's Model of Visual Detection

We will conclude our discussions with a brief look at a mathematical model of visual detection developed by Koopman during his initial work on search theory. This model is important for both historical reasons and because it is still used today in a SAR context.

Koopman had no empirical data, such as the results of controlled experiments, from which he could develop *detection profiles*. Since the primary objective of his research involved detection of enemy ships and surfaced submarines from patrol aircraft flying over open ocean, he analyzed the geometry of this situation and made two reasonable assumptions. The first assumption was that an observer in an aircraft first detects a vessel by sighting its wake. The second assumption was that the instantaneous (one glimpse) probability of detecting the vessel is proportional to the solid angle (like that at the apex of a pyramid) subtended at the observer's eye by the wake's area. Working through the geometry and associated mathematics, this is approximately equivalent to saying the instantaneous detection probability is inversely proportional to the cube of the distance from the observer to the ship's wake. Hence, this model came to be called the inverse cube law of visual detection. Continuing to develop this model, Koopman found that it produced a particular type of bell-shaped detection profile, shown in Figure 7.

Unfortunately, we will not be able to fully appreciate the importance of Koopman's model until near the end of Part II. However, it seemed most appropriate to introduce its detection profile here first.

Sweep Width and Speed

There is one important observation, however, that Koopman's notion of "glimpses" will help us understand before we move on from *effective sweep width* to the next topic. For visual search, the *sweep width* decreases as the speed of the searcher in relation to the search object increases. Using Koopman's approach, we can see at least one reason why this should be true. A searcher "glimpse rate," or number of glimpses per minute, is roughly constant. As his speed increases, the searcher has to scan more area with the same number of glimpses. This gives him less time to focus on each small patch of ground and decide whether the search object is there. Similarly, he has fewer opportunities to catch a glimpse of, and detect, a search object as he passes by. As Hill⁴ and many others have stated, seeing and detecting are not the same thing. Koopman^{1,2} observed, ". . . the act of [the search object's] *recognition* is essential: what the searcher perceives is a set of sensory impressions which he must *interpret* before he knows what is causing them. When the object is in plain view, its recognition is so immediate that this may hardly seem to take place; but in the typical problems of search, recognition can easily be a matter of real difficulty." In his presentation, Hill⁴ dramatically demonstrated, with photographic slides taken in a wooded area, just how difficult recognition can be. It takes a small, but finite, amount of time for a searcher to move his gaze to a new patch of ground (or water), focus on it, give himself enough time to recognize objects of interest if any are present, decide whether there are any such objects present, and move his gaze to a new patch. The faster a searcher moves, the more likely it becomes that he will fail to detect objects of interest even when they are present, and the more likely

it becomes that he will fail to even look at some patches of ground.

Increased speed can also produce other effects. To give an extreme example, consider how many objects a searcher would detect while running through a wooded test area having no trails to follow. The running searcher would have virtually all of his attention devoted to the problems of navigating the terrain at such a high speed. By concentrating on obstacle avoidance, as he must, such a searcher is unlikely to detect any but the closest and most obvious objects, making the number of detections per unit time quite low in comparison to his speed. In terms of Equation [1], the number of objects detected per unit time in the numerator would not be large enough to offset the large searcher speed in the denominator. If, as in paragraph 5, a searcher moving at 0.4 miles per hour detects 12.8 objects per hour, then a searcher running at 4 miles per hour would have to detect 128 objects per hour in order to maintain a computed *sweep width* of 88 feet.

The generalization that increased *speed* results in decreased *sweep width* does not necessarily hold when comparing two very different resources. For example, the advantages of a “bird’s-eye view” from an aircraft will often more than compensate for the detrimental effects of its high speed as compared to that of searchers on the ground. If the *sweep width* from the air is equal to or greater than that for searchers on the ground, the aircraft’s speed becomes a huge asset instead of a liability because it can do much more searching in much less time than ground teams can. Nevertheless, increasing the aircraft’s speed (e.g., doubling it from 100 knots to 200 knots) will have a detrimental effect on *sweep width*. Searching aircraft fly low and slow for the very good reason that it makes the search object below more detectable than it would be from higher altitudes and/or speeds.

Many important factors, such as search platform capabilities, nature of the terrain, searcher safety and fatigue, go into determining what search speed is appropriate. As a result, *sweep width* experiments are normally conducted using search speeds that seem to provide the best balance among the competing demands. Furthermore, the *sweep widths* so obtained are always used with their corresponding *search speeds* for both planning and operations. We will continue the practice of treating *sweep width* and *search speed* as an inseparable pair of values throughout the remainder of this series. In fact, we will combine these two quantities in Parts III and IV into a single variable called the *effective search (or sweep) rate*.

Coming Attractions

In the next article, we will see how the *effective sweep width* concept allows us to develop an objective relation-

ship among the amount of *effort* expended in searching an area, the size of the area, and the *probability of detecting (POD)* the search object if it is present in the searched area. (In fact, objective *POD* estimates are just not possible without a basic measure of *detectability*, i.e., the *effective sweep width*.) In the third article we will look at means for constructing *probability density distributions* that quantify the search manager’s estimate of where the search object is more likely and less likely to be. This will allow the *probability of the object being contained within a defined geographic area (POA or POC according to individual preference)* to be computed. In the fourth and final article we will see how the objective relationship among *POD*, *effort* and *area* can be applied to *probability density distributions* to produce *optimal search plans* that maximize the *probability of success (POS)* obtainable with the effort available.

Food for thought until next time: When we look at how our four broom designs perform when they are used to sweep four identical test *areas*, will they all still produce the same results for the same effort? Do not jump to any conclusions! The answer(s) may come as a surprise!

Acknowledgments

I wish to thank Don Cooper for taking the time from a very busy schedule to review several drafts of these articles. His many suggestions and comments were always enlightening and led to many substantial improvements. I also wish to thank Don for his patience and the insights he has given me into the nature of inland SAR over the past two years (and many hundreds of e-mails).

I wish to thank Martin Colwell for his helpful review and comments, and especially his suggestions regarding the illustrations.

I must also thank Dr. Lawrence D. Stone of Metron, Inc., Reston, Virginia, for patiently tolerating my questions 20 long years ago when I first became involved with CASP from an analyst’s perspective (he was one of the developers), more recently for taking time from his very busy schedule to tutor me through parts of his comprehensive book on optimal search theory, and most recently for reviewing these articles.

Dr. A. S. G. (Tony) Jones also provided many helpful comments and detected a number of important typographical errors, for which the author is more than grateful.

Finally, I must thank Hugh Dougher for suggesting that I write these articles, thereby providing the final bit of incentive needed to get started.

Principles of Search Theory

Part II: Effort, Coverage, and POD

by J. R. Frost

Sweeping Areas

In Part I: Detection we established the concept of *effective sweep width* (usually shortened to just *sweep width*) by using floor sweeping as an analogy for searching. We then found that *sweep width*, W , is used in search theory as a standard method for expressing *detectability*. In this article, we will first develop a definition for effort. Combining this definition with *sweep width*, we will define the amount of *area effectively swept* (called *effective search effort* in a SAR context). Next, we will see how to relate the *area effectively swept* to the actual physical size of the area being swept using a value called *effective coverage*. *Effective coverage* will represent the degree of “thoroughness” with which an area has been swept. We will then see how *effective coverage* may be used to determine what fraction of the material present in the area is swept up. In our floor sweeping analogy, this fraction will be represented as the *percentage of dirt (pod)* originally present that is swept up, although we are using sand as a substitute for dirt in our experiments. Finally, we will provide an example showing how these concepts are used to estimate *Probability of Detection (POD)*. An underlying assumption in all that follows is that effort is applied as uniformly as possible over the entire test area and is not confined to a single, small, portion. We also define a broom’s track as the line followed by the center of the broom head and the track spacing in a pattern of parallel sweeps as the distance between adjacent tracks.

Effort

In the floor sweeping experiments described in Part I, we covered a 10 meter square area uniformly with sand at a “density” of 10 grams per square meter. Then, we pushed each of four different brooms (B1, B2, B3 and B4), across the width of the square, recorded how much sand each swept up and analyzed each broom’s *performance profile*. Pushing a broom a distance of 10 meters represents a certain amount of *effort*. In fact, *effort* is defined as the distance traversed within the area of interest. *Effort* may be equivalently defined as the amount of time spent in the area of interest times the *average speed* of the broom. If multiple brooms are used simultaneously, then the *effort* will be multiplied accordingly.

[2] $Effort = Distance\ Traversed\ in\ the\ Area,$

or, equivalently,

[3] $Effort = (Time\ in\ the\ Area) \times (Average\ Speed).$

It should be easy to see that the amount of sand swept up depends on the amount of *effort* expended. The farther a broom moves, the greater the amount of sand we expect to sweep up. *Effort* has units of length (e.g., meters). Note that this is different from the more usual purely time-based definition used in the work place (i.e., labor hours), although the time-based definition is a component of search theory’s *effort*.

Area Effectively Swept

Just knowing how far a broom moved (i.e., the *effort*) in an experiment will not tell us how much floor sweeping was effectively done. However, if we also know the *effective sweep width* of the broom, then we can compute the amount of area that was *effectively swept*. This computation simply requires multiplying the *effective sweep width* by the *effort* (i.e., the distance the broom moved).

[4] $Area\ Effectively\ Swept = Effort \times (Effective\ Sweep\ Width).$

Provided no portions of the area are swept more than once, we can quickly compute the amount of sand swept up by multiplying the *area effectively swept* by the *density* (in grams per square meter) of sand on the floor. We will momentarily defer the important issues raised when the broom passes over parts of the floor more than once.

Effective Coverage

Effective coverage (usually shortened to just *coverage*) is defined as the ratio of the amount of *area effectively swept* to the actual physical size of the area where the sweeping was done.

$$[5] \quad Coverage = \frac{Area\ Effectively\ Swept}{Physical\ Size\ of\ the\ Area\ Where\ Sweeping\ was\ Done}$$

Thus, if the *area effectively swept*, computed using Equation [4] above, is one half the actual physical area involved, then the *coverage* is said to be 0.5. One can think of coverage as a measure of how “thoroughly” the floor has been swept. Note that for a given *sweep width*, the coverage is proportional to the *area effectively swept*, which in turn is proportional to the level of *effort*. How much sand is swept up, as we shall see, may also depend on the broom *performance profile*, or it may not, depending on the *coverage* and/or exactly how the broom is used. *Coverage* is the ratio of two areas and therefore has no units.

Parallel Sweeps

Most people given the job of sweeping a square area with a push broom would elect to do so using a pattern of parallel sweeps. When the sand is uniformly distributed over an open floor's area, this technique is the most efficient method. We will now describe a series of experiments using parallel sweeps with each of the brooms to see what we can learn. But first let us re-state the initial conditions: We have a square test area measuring 10 meters on a side for a total area of 100 square meters, and the sand is distributed uniformly over the area at a "density" of 10 grams per square meter making the total amount of sand in the test area 1,000 grams (1 kilogram).

Parallel Sweeps with Broom B1

Recall from Part I that B1 is 100% effective over a width of 50 cm and completely ineffective outside that width. Suppose we expend an *effort* of 50 m (or 100 seconds at 0.5 m/s) inside the test area. To satisfy our assumption of uniformly applying the effort to the degree possible, we divide the area into five strips, each two meters wide. We then push B1 down the center of each strip, as shown in Figure 8. (Note that when sweeping the rightmost and leftmost strips, the center of the broom will follow a track one meter from the nearest edge of the test area, while the spacing between broom tracks will be two meters. This is the standard method used in parallel track searches of rectangular areas.)

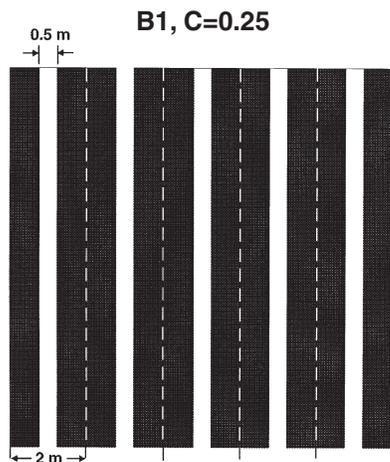


Figure 8

The physical area swept and the *area effectively swept* are both 25 square meters (50 m times 0.5 m), making the coverage 0.25 (25 m²/100 m²). The amount of sand swept up is 250 grams or 25% of the total in the area. Hence the fraction of the sand swept up is exactly equal to the coverage in this case. It is easy to see that if we divide identical "virgin" test areas into more strips and push B1 down the center of each, the equality between the fraction of sand swept up and coverage will continue until we reach a coverage of 1.0. At that point, we will have divided the area into 20 strips having a width of 0.5 m each. After sweeping each of these strips, broom B1 will have swept up all of the sand (100%). Any

further sweeping, i.e., any *coverage* greater than 1.0, would be pointless since it would not improve the results.

Before proceeding further, we need to make a few observations. Even though our first experiment involved five equally spaced parallel broom tracks, we could have achieved the same result using any non-overlapping sweeps as long as a total distance of 50 m is swept. Provided the broom remains within the test area and the sweeps do not overlap, the spacing between the broom tracks is irrelevant. The point to remember is that the fraction of sand swept up depends on *coverage*, not on track spacing. Also observe that we have the advantages of an open floor free of obstructions, precise navigation, and the ability to see exactly where we've already swept. Presently, we will investigate more realistic situations where we do not have these advantages.

Parallel Sweeps with Broom B2

Recall that broom B2 was uniformly 50% effective across a width of one meter and completely ineffective outside that width. Pushing B2 down the centers of each of five strips two meters in width produces the effect illustrated in Figure 9.

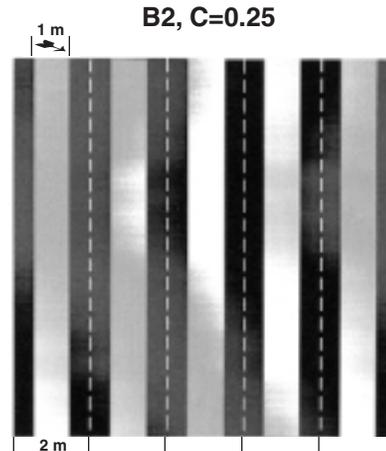


Figure 9

Although B2 passed over 50 square meters of physical floor area, the *area effectively swept* will still be 25 square meters because the *effective sweep width* is still 0.5 m and the *effort* is still 50 m. Hence, we will again have a coverage of 0.25 and we will again collect 250 grams of sand or 25% of the total in the area. It is easy to see that if we divide identical "virgin" test areas into more strips and push B2 down the center of each, this equality between the fraction of the sand swept up and the coverage will continue until we have divided the area into 10 strips, each one meter wide. At this point, broom B2 will pass over every point in the test area once but will achieve a coverage of only 0.5. To reach a coverage of 1.0 using parallel sweeps, we must move B2 a total distance of 200 meters along tracks that are only 50 cm apart, just as we did with B1. Because B2 is twice as wide as B1, it will pass over each part of the test area twice at this coverage. Since B2 removes 50% of the remaining sand with

each pass, it will remove 75% of the sand from the test area at a *coverage* of 1.0.

We should pause again to make a few observations. The most striking of these is the difference in performance between brooms B1 and B2 at higher *coverages*. While both brooms achieve identical results for *coverages* at or below 0.5, the performance of B2 at higher *coverages* falls farther and farther behind that of B1 up to a *coverage* of 1.0. At that point, B1's performance reaches the maximum possible value (100%), and since it can go no higher, broom B2 begins to catch up. At a *coverage* of 2.0 using a standard parallel track search pattern, B2 will sweep over the test area four times and will sweep up 93.75% of the sand in the area. For ease of comparison, the *percentage of dirt swept* up by each broom versus *coverage* will be graphed after we have examined brooms B3 and B4.

Parallel Sweeps with Broom B3

Recall that broom B3 is two meters wide but only 25% effective over that width. When it is pushed down the centers of five "virgin" two-meter strips, a *coverage* of 0.25 is again achieved and 25% of the sand present in the area is swept up. However, B3 has also passed over the entire area of 100 square meters once, as illustrated in Figure 10.

At a *coverage* of 0.5, it will pass over the area twice and will sweep up 43.75% of the sand. At a *coverage* of 1.0, about 68% of the sand will be swept up.

B3, C=0.25

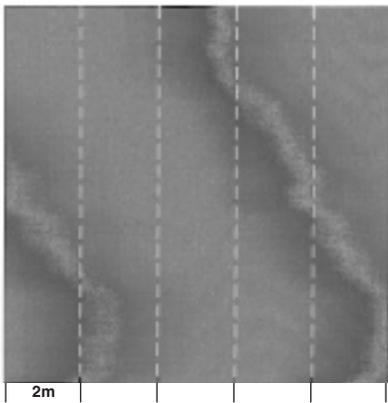


Figure 10

Parallel Sweeps with Broom B4

Recall the hybrid design and "stair-step" performance profile of broom B4 over its one-meter width. Figure 11 illustrates sweeping the area with B4 using a *coverage* of 0.25.

Like B1 and B2, the fraction of sand swept up by B4 equals the *coverage* up to a *coverage* of 0.5. However, at higher *coverages*, broom B4 outperforms B2, but does not do as well as broom B1. At a *coverage* of 1.0, broom B4 sweeps up about 80% of the sand in the test area.

B4, C=0.25

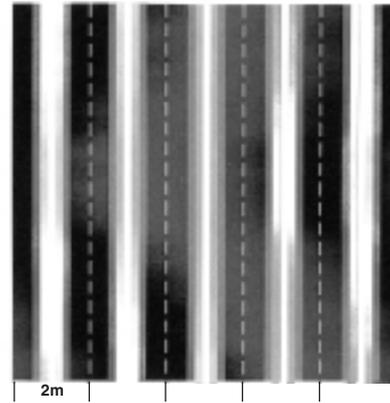


Figure 11

Graphing the Experimental Results

Figure 12 graphs each broom's performance when used to sweep the test area with equally spaced parallel sweeps. On this graph, the *percentage of dirt (pod)* swept up is plotted on the vertical axis against *coverage* on the horizontal axis.

The significance of the fifth, and lowest, curve is discussed below.

Randomizing Influences

Random variations are a fact of life in our activities and in our environment. Failure to account for them is frequently the reason mathematically precise solutions that look good on paper do not always work in the real world. There could be several sources for "randomness" in a real floor sweeping problem. There might be obstructions in the area that prevent the sweepers from moving along straight, equally spaced, parallel tracks or the sweepers might not choose to use such tracks for some reason even if they were possible. The density of the sand might not be uniformly distributed over the area, but randomly piled up in mounds of higher density in some places and lower density in other places. This would make it difficult to predict the amount of sand being swept up, particularly at lower *coverages* or for non-uniform detection profiles like that of broom B4. Finally, we could be trying to sweep up randomly scurrying ants instead of stationary grains of sand. Even if a swath was swept clean, some of the ants might run into it after the broom had passed. Although we cannot predict with certainty the fraction of sand (or ants) that will be swept up in any single floor-sweeping problem having random variations in one or more of its elements, it is possible to determine what the average results of many such sweepings would be. The graph of that answer is the curve labeled "Random Sweeping" in Figure 12.

Note that the lower, broader, and flatter the broom *performance profile*, the closer its performance over the area will be to the "random sweeping" curve, even though perfectly straight, parallel tracks are still being used. Although

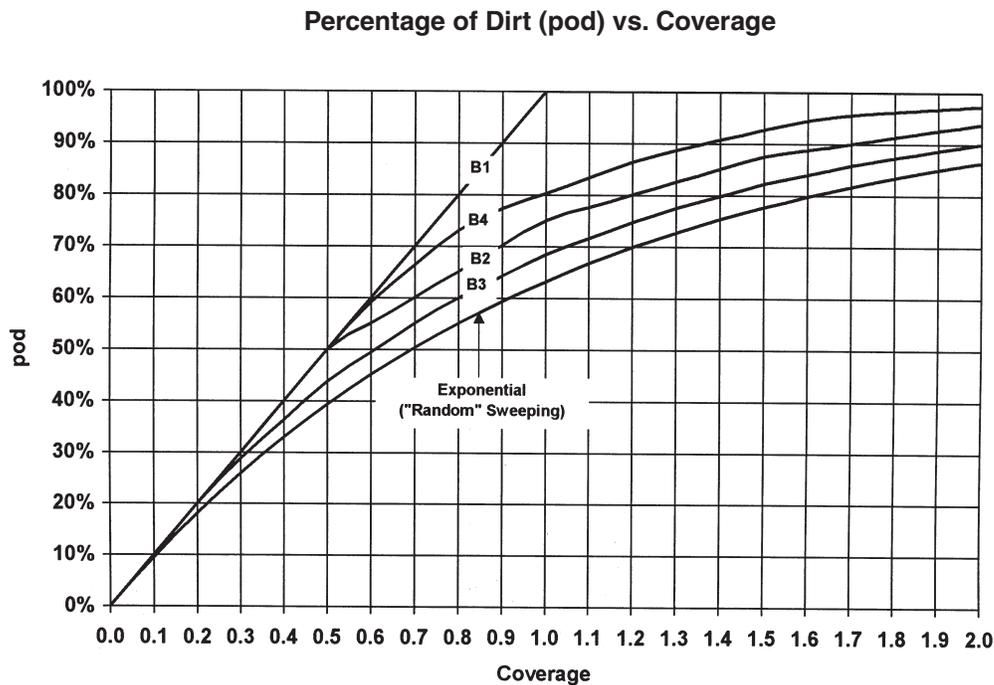


Figure 12

not obvious, using any of the brooms in a random fashion, rather than parallel sweeps, or having other significant random influences present, would place their average performances on the “random sweeping” curve. Therefore, when it is feasible to use straight, equally spaced, parallel tracks uniformly distributed over an area free of significant random variations, this tactic will make more efficient use of the available effort than “random” sweeping. How often such favorable conditions occur in real-world situations is another question.

We now have the answer to the question posed at the end of Part I, but the answer is not a simple one. When used in a pattern of perfectly parallel, equally spaced tracks to sweep a floor uniformly covered with sand, the four different broom designs perform equivalently at low *coverages* (less than or equal to 0.25) and very high *coverages* (greater than 4 where even the curve for B3 surpasses 99%). However, they do not perform equivalently at intermediate *coverages* when sweeping areas. The differences among the brooms in this intermediate, and most useful, range are significant. Alternatively, when the brooms are used in a somewhat “random” fashion or when significant randomizing influences are present, all four broom designs perform equivalently at all *coverages*, as shown by the so-called “random sweeping” curve.

Multiple Low-Coverage Sweepings vs. Single High-Coverage Sweepings

One might ask whether there is any advantage to sweeping the floor twice at a *coverage* of 0.5 as opposed to once at a *coverage* of 1.0 since both require the same *effort*. When

using broom B2, it should be easy to see that it doesn’t matter whether the floor is swept once at a *coverage* of 1.0 (track spacing = 0.5 m) or twice at a *coverage* of 0.5 (track spacing = 1.0 m). The broom is uniformly effective across its width. After one pass, the sand left behind is uniformly distributed over the swept area. Whether the area between two adjacent tracks is immediately swept over a second time as the broom moves down the adjacent track, or later as the result of a second complete sweeping of the area, makes no difference. Either technique will sweep up 75% of the sand. A similar argument may be applied to B3. However, let us examine the situations created by brooms B1 and B4. Unlike B2 and B3, which left the remaining sand uniformly distributed over the test area, broom B1 left half of the test area cleanly swept in the form of 10 “bare” 50 cm corridors separated by 10 untouched corridors also 50 cm wide. Only by a very precise placement of the second set of sweeps will we get the same results from two sweepings at a *coverage* of 0.5 that we got with one sweeping at a *coverage* of 1.0. Even a small error in the placement of broom B1 will produce a substantial decline in the amount of sand obtained from the second sweeping. Broom B4 presents a similar problem.

In general, two sweepings of an area at one-half of a given *coverage* can produce results no better than one sweeping at the given *coverage*, and the two *low-coverage* sweepings could easily do worse. Because B1 and B4 are very effective close to their tracks, accurate broom placement, relative to the first set of tracks, is necessary if the results of two sweepings are to match those of expending the same total effort in a single sweeping. Therefore, the outcome of two *low-coverage* sweepings is not as predictable as with B2 and B3 since the results of two sweepings are so

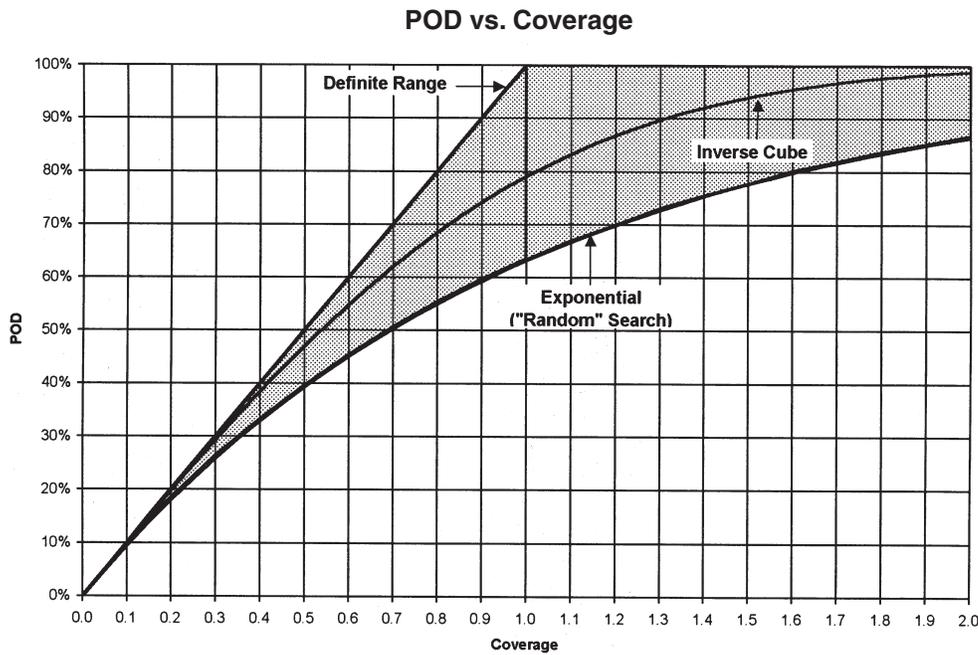


Figure 13

sensitive to where the second set of tracks is placed with respect to the first set. However, if the placement of the second set of parallel sweeps is independent, i.e., randomly offset, from the first set, then *on average*, over many sweepings, we will get substantially less sand for two sweepings at a *coverage* of 0.5 than for one sweeping at a coverage of 1.0. Situations often arise where close coordination between sweepings is not possible. For example, when we think of searching instead of sweeping floors, it seems unlikely that two independent parallel track searches of an area using 5-person teams would exactly replicate the tracks of a single 10-person team.

Returning to Searching

Let us restate Equations [2]–[5] in the terminology of searching and search theory. *Effort*, z , is defined as the distance traversed by the searcher(s) while operating in the area being searched. Stated as a formula,

$$[6] \quad z = L,$$

where L is the distance traveled by the searchers in the searched area. Note that this is the distance traveled along the searcher's actual path, not the lengths of the beelines connecting the searcher's position on one side of the searched area with his position when he reaches the other side. Equivalently,

$$[7] \quad z = v \times t,$$

where v is the average speed of the searcher(s) while searching and t is the time spent searching. The *area effectively searched* is called the *effective search effort* (or just *search effort* for short), Z , and is defined as the product of *effective sweep width* and *effort*. That is,

$$[8] \quad Z = W \times L = W \times z = W \times v \times t,$$

where W is the *effective sweep width*. (Note that *effort* has units of length while *search effort* has units of area.) Finally, *coverage*, C , is defined as the ratio of the *effective search effort* to the area being searched. Expressed as an equation,

$$[9] \quad C = \frac{Z}{A}$$

where A is the physical size of the area being searched. If multiple searchers or several similar resources are used simultaneously in an area, then the *effort* is multiplied accordingly, causing corresponding increases in *search effort* and *coverage*. (However, simultaneous searching by dissimilar resources, such as searchers on the ground and in an aircraft, should normally be treated as separate, independent searches.) For the very special case of perfectly straight, equally spaced, parallel tracks uniformly distributed over a rectangular area, we may take a short cut and compute *coverage* as the ratio of *sweep width* to *track spacing* (S), i.e.,

$$[10] \quad C = \frac{W}{S}$$

POD vs. Coverage

When we try to transfer the findings of the parallel track experiments described above to searching, we find that searching is rarely, in fact almost never, as clean and straightforward a proposition as sweeping floors. There are always all sorts of random influences on the searching and detection processes that are beyond the searchers' (or anyone else's) control. This is sometimes true even when searching for objects adrift on the open ocean, and probably even more frequently true for searches conducted on or over the ground. For this reason, a search model that accounts for random

variations is probably a very good unbiased estimator of actual organized search performance in the field.

The mathematical derivation of the so-called “random search formula” (also called the exponential detection function) is beyond the scope of this article, so we will simply state, without proof, the formula derived by Koopman:^{1,2}

$$[11] \quad POD = 1 - e^{-c}$$

In this formula, e is the base of the natural logarithms ($\cong 2.71828$) and C is the *coverage*. Koopman’s definition of “random” in this context contains the restriction that the *effort* be uniformly distributed over the area being searched. In that sense, the searchers’ movements cannot be completely random as they must fill the area to the same degree everywhere, like a liquid fills an open bucket, in order to satisfy Koopman’s restriction.

The exponential detection function has some interesting properties. Unlike perfect parallel track searching, the shape of the *detection profile* does not affect the *POD* computed by Equation [11], and perfectly straight, equally spaced, parallel tracks are not required. In other words, all *detection profiles* perform equally at equal *coverages* whenever significant random variations in the search parameters are present—as long as the *effort* is uniformly spread over the area. The exponential detection function also has the property that splitting a given amount of *effort* in two and performing two successive searches of an area always produces the same *cumulative POD* as using all the *effort* to do a single search of the area.

Figure 13 shows the graphs of *POD vs. Coverage* for three detection profiles that often arise in search theory.

The upper “curve” that is linear from the origin to (1, 100%) is the *POD* graph for a *definite range detection profile* (like the performance profile of B1 shown in Part I) following perfectly straight equally spaced parallel tracks. The middle curve is based on Koopman’s^{1,2} mathematical model of how warships underway are detected visually from the air—the so-called inverse cube law whose *detection profile* was shown near the end of Part I. This middle curve is the same one that appears in Figure 5-19 on page 5-29 of the U.S. *National SAR Manual* (1991)³ for estimating the *POD* of any single search. It is based on using Koopman’s inverse cube law of detection along perfectly straight equally spaced parallel tracks. The lower curve is the graph of the exponential detection function. As Koopman^{1,2} observed,

“At one extreme is the definite range law, at the other the case of random search. All actual situations can be regarded as leading to intermediate curves, those 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.”

In other words, provided there is no systematic error or bias affecting the searches, the average results of any search technique over many searches must fall on or between the

upper and lower curves shown in Figure 13. Under ideal search conditions where *sweep widths* are relatively large and the parallel tracks are accurately navigated, Coast Guard experiments show the middle curve in Figure 13 predicts *POD* remarkably well for experimentally determined sweep widths, even those of drifting objects that leave no wake. However, as search conditions deteriorate, not only do *sweep widths* decline, the actual *detection profiles* become much lower and flatter than that of Koopman’s inverse cube law. As with our broom experiments, this change in the *detection profiles* drives *POD* values toward the exponential detection function, even for perfect parallel tracks. In other words, a *coverage* 1.0 parallel sweep search for a boat under ideal conditions would have a *POD* of about 78% while a *coverage* 1.0 search for the same boat under poor or difficult search conditions would not only require more *effort* (due to the reduction in *sweep width*), it would have a *POD* of only about 63%.

The Case for Exponential Detection

Many would argue that most actual search results fall closest to those predicted by the exponential detection function in Figure 13. Searching is a difficult, demanding, and sometimes dangerous business. Under actual operational conditions, no searcher or sensor package can perform with mathematical precision free from random variations, nor can any search pattern be followed with absolute precision, nor is the search environment perfectly uniform. Even though GPS now makes it possible to come very close to perfect navigation, statistical modeling shows that it takes surprisingly little variation from perfectly straight equally spaced parallel tracks for expected *POD* values to be close to the exponential detection function. Furthermore, the smaller the *sweep width*, i.e., the harder detection becomes, the less variation from a mathematically perfect pattern it takes to make the exponential detection function the most reliable predictor of *POD*. Adding other sources of random variability or uncertainty about the search parameters will only reinforce the exponential detection function as the most reliable estimator of *POD*.

Note: We are discussing the average performance over many searches and using that information to predict or estimate the results of a single search. However, the actual *POD* achieved on any single search can be outside the shaded envelope of Figure 13. One could get lucky and find the search object very early in the search without expending much effort, creating a (statistically incorrect) temptation to claim a 100% *POD* for a coverage of less than 1.0. A more likely situation is one where there is some hidden bias or systematic error in the conduct of a search, such as avoiding difficult patches of terrain, making the actual *POD* for that search significantly less than the value predicted by the exponential detection curve. We can never know the actual *POD* of any search. We can only know the “found/not found” outcome. However, search managers and search team leaders need to be aware that actual *PODs* can be less than those predicted by the exponential detection function. These individuals need to be alert for and report any problems with

the execution of individual searches that might cause such a reduction in POD.

Characteristics of *POD* Coverage vs. Coverage Curves

All *POD* vs. Coverage curves have some important properties in common. All begin at the origin (0, 0%) with a slope of 1.0, i.e. the “rise” of the line tangent to the curve at the origin equals the “run.” All but that of the “perfect” definite range detection profile fall away from this 45° line at some point and approach a *POD* of 100% more and more gradually as coverage increases. This property produces a situation of diminishing returns. For example, suppose five searchers can search an area with a coverage of 0.5 during one sortie. Using the exponential detection function, this would produce a *POD* of about 39%. Adding five more searchers would give a coverage of 1.0 and a *POD* of about 63%—an increase of about 24%. Note that doubling the effort did not double the *POD*. Adding five more searchers would give a coverage of 1.5 and a *POD* of about 78%—an increase of only about 15%. Finally, adding five more searchers (for a total of 20) gives a coverage of 2.0 and a *POD* of about 86.5%—an increase of less than 9%. High coverages, and hence high levels of effort, are required to achieve high *PODs*. However, the *POD* per unit of effort expended goes down as the total amount of effort expended increases. Another lesson is that a very high coverage probably is not the wisest way to apply limited resources unless a very high *POD* is required after only one search, such as an evidence search where searchers are likely to disturb or destroy evidence not found on the first pass.

An Example

We’ve covered a lot of ground in this article. It may be helpful at this point to give an example of how to use what we’ve learned so far. Suppose we are searching for a lost child, age 9, in the woods near his home. Suppose we assign a three-person search team to a segment of the search area having a size of 0.25 square miles. Finally, suppose we have previously conducted sweep width experiments in similar terrain for search objects resembling a 9-year-old child, and those experiments produced a sweep width of 106 feet for comparably trained/skilled searchers moving at 0.5 mile per hour. We expect our search team to move at about 0.5 miles per hour and take about six hours to complete their search of the segment. What *POD* should we expect?

From Equation [7], we start by computing the effort, or total distance moved, for one searcher, as

$$z = v \times t = 0.5 \text{ mph} \times 6 \text{ hours} = 3 \text{ miles per searcher}$$

Since we have three searchers, the total effort is 3 x 3 or 9 miles for the team as a whole. We will assume the sweep width determined from earlier experiments is a good estimate for the sweep width in our current situation. Converting this value from feet to miles so our units of measure are consistent, we get

$$W = \frac{106 \text{ feet}}{5280 \text{ feet/mile}} = 0.02 \text{ miles}$$

Now we may compute the effective search effort, *Z*, as follows, using Equation [8]:

$$Z = W \times z = 0.02 \text{ mile} \times 9 \text{ miles} = 0.18 \text{ square miles}$$

We can now compute the coverage using Equation [9]:

$$C = \frac{Z}{A} = \frac{0.18 \text{ square miles}}{0.25 \text{ square miles}} = 0.72$$

Using Figure 13, we enter at the bottom with a coverage of 0.72, go up to the exponential detection curve, and read a *POD* of about 51% from the vertical axis to the left. If a scientific calculator is handy and has an exponential function button, we may also compute the *POD* from Equation [11] as follows:

$$POD = 1 - e^{-c} = 1 - e^{-0.72} = 1 - 0.48675 = 0.51325 \text{ or } 51.325\%$$

The added precision is, of course, completely superfluous.

Although some experienced searchers may not agree, maritime experience has shown that *PODs* estimated in this fashion are generally more reliable, and definitely more consistent, than direct estimates of *POD* provided by the searchers themselves. *POD* estimates from searchers often tend to be optimistic, i.e., too high. Based on the results of carefully designed and controlled scientific experiments, maritime search planners have adopted a philosophy of asking searchers to report those things affecting detection which they can actually observe. This includes such things as meteorological visibility, wind velocity, sea state, and crew fatigue, to name a few, along with any other on-scene observations they think may be important. (Ground teams would be expected to report things like the type of terrain and amount of ground cover they actually encountered, for example.) However, searchers are not asked to estimate *POD*. The search planner uses their reported observations in combination with tables of experimentally determined sweep width values and correction factors to estimate the actual sweep width, coverage and *POD*.

It is worth considering how to deal with the situation if our search team in the above example unexpectedly returned in only three hours instead of six. Let us examine a few of the many possibilities.

First, the searchers could have been moving at the desired search speed, traveling a total of only 4.5 miles instead of 9 miles. This means they expended only half as much effort as expected. If the team is sure they “covered” the entire segment, we can recompute the effective search effort and coverage using 4.5 miles of effort instead of 9 miles. The results will be 0.09 square miles and 0.36 respectively. Entering Figure 13 with a coverage of 0.36, we get a *POD* of 30% from the exponential detection function.

Second, the searchers could have “covered” only half of the segment at the desired search speed. If it is possible to determine which half of the segment was “covered,” we can recompute the coverage by applying the effective search effort of 0.09 square miles to an area of 0.25/2 or 0.125 square miles to get a coverage of 0.72. Using Figure 13 (exponential detection), we may assign a *POD* of 51% to the half of the segment that was searched, and a *POD* of 0% to the half that was not.

A third possibility is that they searched the segment at twice the expected search speed (1.0 mph). In this case, they

would have traveled the expected distance and expended 9 miles of *effort*. However, recalling the discussion in Part I about the detrimental effects of increased speed on *effective sweep width*, it is likely that the actual *sweep width* for the search was less than the 106 feet obtained in experiments where searchers were moving at only 0.5 mph. If *sweep width* data or an experimentally verified correction factor is available for the higher speed, then a *sweep width* value for the higher speed could be estimated and used to recompute *effective search effort*, *coverage* and *POD*. Otherwise, we will have to estimate the actual sweep width as best we can and use that value to recompute *effective search effort*, *coverage* and *POD*. A smaller *sweep width*, of course, will produce a smaller *effective search effort*, *coverage* and *POD*.

Finally, a fourth possibility is that the terrain was not as difficult to traverse or search as we had anticipated. This could mean the search object should have been more detectable than originally thought, making the *sweep width* greater than 106 feet. To find out, we would go back to the table of sweep width values for various environments and find the one for the terrain (and search speed, search object, etc.) that most closely corresponded with what the searchers actually encountered. Using that new *sweep width* value, we would then recompute the *effective search effort* and *coverage*. A larger *sweep width* would, of course, produce a larger *effective search effort*, *coverage* and *POD*.

Looking Ahead

Determining where to send a limited number of resources and how large their assigned areas should be is the ultimate

question we are striving to answer. One requirement for answering this question is having a reliable way to estimate *effective sweep widths (detectability)* so we can compute the *effective search effort* each resource can deliver. A second requirement is the availability of a reliable *POD vs. Coverage* curve so we have an objective method for estimating how changes in *coverage* resulting from different allocations of *effort* will affect *PODs*. That is, we need to know what *PODs* we can expect from concentrating our resources in small areas, and what *PODs* we can expect as we spread our resources out over progressively larger areas. A third requirement is an estimate of search object's location *probability density distribution* over the area. Part III: Probability Density Distributions will address this issue. In Part IV: Optimal Effort Allocation we will seek to answer the question of how to make the best use of the available *effort*.

Food for thought until next time: Suppose the area containing a lost or missing person is divided into 10 regions. Suppose two of these regions are each assessed as "very likely" to contain the person and all the others are assessed at varying, but lower, likelihoods. Finally, suppose one of the top two regions is significantly larger than the other. Does the assessment of "very likely" mean the top two regions each have about the same probability of containing (*POC* or *POA*) the person or does it mean they each have about the same *probability density* (probability per unit area or *POC/A*)? Warning: The answer could make a significant difference in how the available *effort* should be allocated when we get to Part IV.

This page intentionally left blank.

Principles of Search Theory

Part III: Probability Density Distributions

by J. R. Frost

Note: Readers are encouraged to re-read “Principles of Search Theory, Parts I and II,” printed in Volume 17, Number 2 of Response. We will be referencing them frequently. Readers are also encouraged to have pencils, a scratch pad, some graph paper, and a calculator handy.

Where to Search

Our first two articles dealt exclusively with the mechanics of searching. We developed the concept of *effective sweep width (detectability)* and examined how *sweep width, effort, coverage, and probability of detection* over an area are all related. However, we did not discuss any issues related to where the searching should be done. Our ultimate goal 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

$$[12] \quad Pden = \frac{POC}{A},$$

where *POC* (also called *POA*) is the probability that the search object is contained in some area and *A* is the size of that area.

Effective Search (or Sweep) Rate

Another important quantity we will need for our discussions here, and for the optimal effort allocation discussions in Part IV, is the effective search (or sweep) rate. It is simply the product of the effective sweep width, *W*, and the search speed for which that effective sweep width is valid. Stated as an equation,

$$[13] \quad \text{Effective Sweep Rate} = W \times (\text{Search Speed}).$$

The *effective sweep rate* has units of *area per unit time* (e.g., square miles per hour). As we will see in Part IV, the best placement of the available effort will depend on the interplay between *probability density* and *effective sweep rate* as one evaluates the search area looking for the best places to search during the next search cycle.

Probability Density Distributions and Probability Maps

A *probability density distribution* is usually represented by a *probability map* consisting of a regular grid. For the pur-

poses 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 14 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 14

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 distribu-*

tion. 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 the subjective scenario weights) of the cell probabilities from each scenario.

Unfortunately, formal search theory does not shed much light on how to go about turning an inconsistent body of evidence and data from a variety of sources into numbers on a probability map. As Stone,⁵ 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.”

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 of, at least partially, 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*; a topic discussed more fully in Part IV. 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 1. The values in the last

two columns were computed from the corresponding area and density values in the second and third columns. Figure 15 illustrates the situation.

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%

Table 1

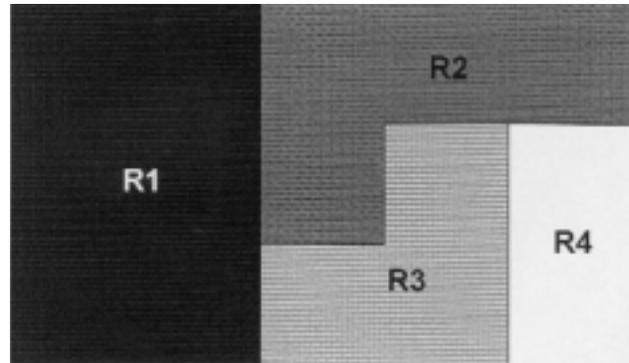


Figure 15

Suppose we have only one sweeper, whose broom is B2 from our *sweep width* experiments (see Part I) 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*, as defined in Part II, is 150 m. Since broom B2 is one meter in width, the sweeper could sweep an area of 150 m^2 . 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 Equation [4] from Part II,

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

the *area effectively swept* in five minutes is computed to be 150m x 0.5m or 75 square meters. From Equation [5] of Part II,

$$\text{Coverage} = \frac{\text{Area Effectively Swept}}{\text{Physical Size of the Swept Area}},$$

a *coverage* of 75 m^2 /150 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 12 from Part II shows B2 will sweep up 50% of the sand initially present in 150 m² 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

$$POC = 1 - e^{-\frac{R^2}{\sigma^2}}$$

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 16, 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.

1.42%	9.08%	1.42%
9.08%	57.91%	9.08%
1.42%	9.08%	1.42%

Figure 16

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 16 (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 17 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 17

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 18.

5%	12%	6%
14%	18%	16%
9%	2%	7%
6%	1%	4%

Figure 18

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 15. We now obtain an undistorted photograph of the entire floor from a point directly above its center and make three copies. Like Figure 15, 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 assignment. 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 1.

TOM’S ASSESSMENTS				
Region	Area (m ²)	Estimated poc	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

Table 2

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. In Part IV, we will see 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 suboptimal, 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 3.

DICK'S ASSESSMENTS						
Region	Letter Designation	Numeric Rank	Computed poc	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

Table 3

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

A	A	B	B	B
A	A	B	C	D
A	A	C	C	D

Figure 19

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 19.

Grouping the cells by region, she gets the results shown in Table 4.

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

MARY'S ASSESSMENTS						
Region	Letter Rank	Numeric Rank	Computed Cell poc	Computed Region poc	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	12,800/ 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

Table 4

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 As, 16 Bs, 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 4.

From Mary’s assessment, using a regular grid of cells, we may produce a “map” like that in Figure 20, 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 20

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.

DENSITY-BASED ASSESSMENTS				
Region	Area (m ²)	Relative Density	Relative Amount of Sand	Computed poc
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%

Table 5

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 5 shows how the *percentages of containment* could be computed from these relative *density* estimates.

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 21.

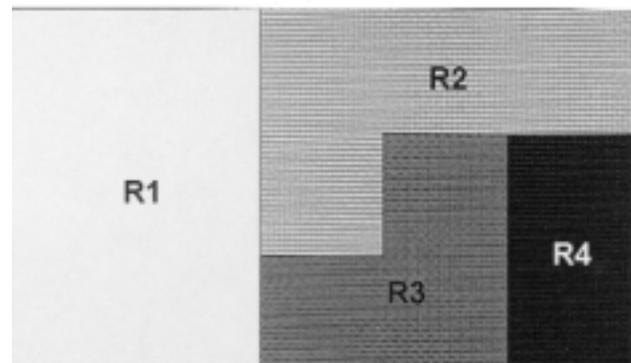


Figure 21

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 6.

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 6 as the basis for a simple ranking, the results will be inaccurate. Table 7 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 6 below.

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.

CELLULAR ASSESSMENT OF FIGURE 21						
Region	Letter Rank	Numeric Rank	Computed Cell poc	Computed Region poc	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

Table 6

SIMPLE RANKING ASSESSMENT OF FIGURE 21					
Region	Letter Designation	Numeric Rank	Computed poc	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

Table 7

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 21.

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 15 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 1. In fact, his results would have been identical to Tom's in Table 2, as shown in the table below.

PROPORTIONAL RATING ASSESSMENT OF FIGURE 15		
Region	Proportional Assessment	Computed poc
R1	10	10/20 = 50%
R2	6	6/20 = 30%
R3	3	3/20 = 15%
R4	1	1/20 = 5%
Totals	20	100%

For **Figure 21**, 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 6**. (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.

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 5** above illustrated how this could be done.

Containment vs. Density Estimates

It is important at this point to reconsider the question posed at the end of **Part II**: 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:

1. 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*.
2. 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.

Using the densities of sand on different parts of a floor as an analogy for *probability density*, readers are encouraged to develop a *probability map* for Figure 21 using Table 6. (The correct answer is contained in this article.) Readers are

Obtaining meaningful Probabilities of Containment REQUIRES the use of a Proportional Assessment Technique.

also encouraged to make up their own exercises, like those above, and develop the corresponding *probability maps*. This practice will provide a deeper understanding of the concepts involved.

Proper Use of Probability Density

We must pause again for a preview of things to come if we are to avoid leaving false impressions from our simple examples. So far, the only variable we have considered is *probability density* and the only problem we have really considered is where to place the first small increment of *effort*. (Recall that a searcher’s *effort* is defined as the distance traveled by the

searcher while searching, or, equivalently, the searcher’s speed times the amount of time spent searching.) We have looked at the effects of *density* differences while keeping *speed* and *sweep width* (i.e., the *effective sweep rate*) constant in order to give a simple demonstration of why *probability density* is important. We have not yet tried to show how variations in *probability density* should be used in the more complex, and more typical, effort allocation problems that also involve variations in *effective sweep rates* as well as the simultaneous placement of significant numbers of resources in different parts of the search area. Therefore, readers should not jump to the conclusion, for example, that regions, segments, cells, etc., should be searched in descending order of *probability density*. Unfortunately, the answer is not that simple.

The ultimate issue for the search planner is determining how the available effort should be apportioned among the various parts of the search area. A simple ranking might tell the planner where to send a single resource initially, but it does not tell him how to distribute a number of resources over the area as a whole. In other words, if we are to make the best use of our resources, we must know not only where to place effort, but how much of the available effort should be placed in each part of the overall search area. The *probability densities* in the various portions of the search area are an important factor to consider. Other equally important factors include the *effective search (or sweep) rates* in the different portions of the search area as well as the sizes (areas) of those portions. In Part IV—Optimal Effort Allocation, we will see how the combined effects of all these factors affect the choice of areas where effort should be placed and how much of the available effort should be assigned to each in order to maximize the probability of a successful outcome.

Creating Generalized Probability Density Distributions

Whenever search planners outline areas on a map or chart and assign probability values to them, they are creating a *probability density distribution*, regardless of whether they are thinking of their estimates in that way. A good estimate of how the *probability density* is distributed over the *possibility area* is an essential input when deciding how to deploy the available *effort* so we maximize our chances for finding

the lost person in the minimum amount of time. Therefore, we must be very careful about how subjective assessments are translated into *probabilities of containment (POCs)* with their corresponding *probability densities (Pdens)*. As we have just seen, even when there is evidence as good as a photograph showing how the *density* is distributed (something that will never be available for search planners), significant problems can arise from some techniques of turning subjective estimates into numeric values. All of the available evidence bearing on where the lost person might be during the next set of search sorties needs to be carefully evaluated in a way that will produce a valid estimate of how the *probability density* is distributed. Although very subjective assessments will always be necessary and current practices typically produce estimates of *probabilities of containment*, evaluators should be aware of the important role the *probability densities* computed from their *POC* estimates will play in finding the optimal distribution of *effort*.

Some Things to Consider

We will now offer some additional thoughts for SAR managers and search planners to consider.

a. **Pre-planned Searches.** If historical records are available for an area of responsibility (AOR) where SAR incidents are relatively frequent, consider analyzing those records for historical trends and insights into where lost persons are most often found. Consider building *historical probability* maps and developing optimal search plans for them (after reading Part IV) for use in initial searches while the evidence pertinent to a specific incident is being carefully evaluated. Consider working with a professional statistician or practitioner of operations research while doing this work.

b. **Assessing the Evidence and Other Data/Information.** When a SAR incident arises, one, or preferably more, persons should evaluate the evidence, clues, historical data, behavior profiles, etc., and develop estimates of where the lost person is likely to be. This topic alone could be the subject of many articles or even books. For now, it is important to emphasize the requirement for carefully evaluating the body of available, often conflicting, evidence to extract as much information as possible about the lost person's probable locations. Evaluators should develop a number of scenarios compatible with self-consistent subsets of the available data. The different scenarios should then be weighted according to their relative likelihood of representing the true situation. However, evaluators should strive to avoid "over-assessing" the evidence used in individual scenarios. When making estimates that will ultimately be used to produce *probability maps*, evaluators must resist the temptation to make distinctions between cells, regions, segments, etc., that are finer, or more detailed, than the available evidence will support. Recall the example above of the light air-

craft being used to survey a wilderness area. If there truly were no further information available from any source, it would be difficult to justify anything other than a uniform *probability density distribution* for the scenario of the plane going down while engaged in survey operations in the area.

c. **Assumptions vs. Facts.** Evaluators need to clearly document assumptions they make when developing possible scenarios and keep them separate from the known facts. An assumption, if repeated too often and questioned too seldom, gradually takes on the appearance of fact and can lead to something called "scenario lock." Scenario lock occurs when planners become fixated on a particular (and not always the most likely) scenario to the exclusion of all others. Such fixations may lead planners far astray and result in significant delays or even complete failures. This unsatisfactory situation can arise from basing an extended search on an initial cursory assessment that is never revisited. Therefore, it is important to conduct regular re-evaluations to account for new evidence (including negative search results), re-evaluate assumptions, and prevent scenario lock.

d. **Assessments and Planning vs. Operations.** Evaluators should concern themselves only with evaluating the available evidence. Such things as the logistical and management problems associated with the conduct of search operations should not be allowed to affect the evidence assessment process. The ultimate objective of evidence assessment is a *probability map* that correctly reflects the evaluators' assessment of the available data. In other words, it is the job of the evaluators to ascertain, to the best of their abilities, the *meaning* of the available evidence and data in the context of the current incident and quantify that meaning via a *probability map*. Once a probability map has been constructed from the assessment results, then search planners can proceed, via the methods to be described in Part IV, to determine how to distribute the available *effort* to the greatest advantage. Once this is done, search managers can work out the details of how subdivide regions into manageable segments, how to deploy and recover the search resources, etc.

e. **A Process of Elimination.** Like many other types of investigations, SAR cases are "solved" largely by a process of elimination. In SAR, the objective is to eliminate uncertainty about the lost person's location and condition. Searching is but one of many tools used in the process for eliminating this uncertainty. However, it is by far the most involved, expensive, and risky tool, and one that is used only when it is believed the lost person is in imminent danger. These are characteristics that require searching to be done in the most efficient, effective manner possible and justify significant investment in the assessment of evidence, planning of searches, and search planning aids, such as computer programs to compute optimal effort allocations and keep track of search results.

*Estimating
probabilities of
containment (POCs)
for SAR situations
remains a
very subjective and
a very difficult task for
which no simple
solutions exist.*

Principles of Search Theory

Part IV: Optimal Effort Allocation

by J. R. Frost

Note: Readers are encouraged to re-read Parts I, II, and III. We will be using all of the concepts they introduced. Readers are also encouraged to have pencils, a scratch pad, some graph paper, and a calculator handy.

Goal of Search Planning

The ultimate goal of any search planner is to develop a plan for applying the available resources to the search space in a way that maximizes the chances for finding the object of the search in the minimum amount of time. In other words, the search plan should represent the most effective and most efficient use of the available *effort*. Search effectiveness is measured by a quantity called the *probability of success*. The probability of success is defined as the product of the probability that the searched area contained the object at the time of the search (*POC* or *POA*) and the probability of detecting the object if it was there (*POD*). The general formula for computing the probability of success for a searched area is

$$[14] \quad POS = POD \times POC.$$

Equation 14 may also be used to predict the *probability of success* for a search of an area based on predicted *POC* and *POD* values. The overall effectiveness of all searching done to date is given by the *cumulative overall probability of success*. This value is the sum of all the un-normalized *POS* values over all segments of the search area for all searching done to date. It represents the chances that all searching done to date would have found the search object if it was anywhere in the *possibility (search) area(s)* of the *scenario(s)* under consideration, regardless of whether all parts of these *possibility area(s)* have actually been searched.

Search efficiency is measured by how quickly *POS* increases as the search progresses. The search plan that increases *POS* at the maximum possible rate for the *effort* that is available is said to be a *uniformly optimal search plan*. A search plan that achieves the same final *POS* for the same *effort*, but takes longer and/or does not increase *POS* at the maximum possible rate in the early stages, is said to be a *T-optimal search plan* where *T* is the time spent expending the available *effort*. A *T-optimal search plan* is the next best thing to a *uniformly optimal search plan*. As this article progresses, we will see how to develop *uniformly optimal search plans* in the absence of real-world operational con-

straints. Then, briefly, we will discuss how to modify such plans so that operational constraints may be addressed without decreasing the rate of *POS* growth or the final *POS* too badly.

General Review

In Part I, the concepts of *detection profile* and *effective sweep width* were introduced. In Part II, the concept of *area coverage* was introduced and it was defined in terms of *effort*, *effective sweep width*, and the amount of *area* over which the *effort* was being applied. The relationship between *coverage* and *probability of detection, POD*, was explored for several different detection profiles when used in parallel sweep search patterns. The relationship between *coverage* and *POD* for so-called “random” searching was also explored, producing the exponential detection function. Graphs depicting these relationships were constructed. An argument favoring the exponential detection function as the most realistic estimator of *POD* under actual operational search conditions was advanced. In Part III, the importance of *probability density* and having a good estimate of how it is distributed over the search area (the *probability density distribution*) were demonstrated using the *density* of sand on a floor as an analogy. The concept of *effective search (or sweep) rate* was introduced. The notion of a *probability map* was introduced by using a regular grid of cells and estimating the amount of probability contained in each (*probability of containment, or POC*). From these values, *probability densities* could be computed for each cell. It was shown that it was also possible to use the reverse process of first estimating the *probability densities* and then computing the *probabilities of containment*. We will now investigate how all of these concepts can be used together to make the most productive use of the available *effort*.

In the paragraphs that follow, we will revisit the concept of *effective search (or sweep) rate*, set up some floor-sweeping experiments and define a term called *productive sweeping rate (psr)*. After performing some experiments to get a “feel” for the general nature of the optimal *effort* allocation problem, we will define a new term called the *probable success rate (PSR)* using the productive sweeping rate from our experiments as an analogy.

Effective Search (or Sweep) Rate

We will begin by reviewing the notion of *effective search (or sweep) rate*. Recall from Part I that the *effective sweep width, W*, is a measure of detectability for a particular sensor “sweeping” an area at a particular rate of speed, *v*, looking for a particular search object under a particular set of environmental conditions. If objects were uniformly distributed over a large area, then *W* is the width of the swath along the sensor’s track that contains the same number of objects as the sensor detects in a single pass through the area in a straight line. This does not imply, however, that all the objects detected will be within that swath. In fact, the number of objects detected outside the swath will equal the number not detected within the swath. (This is yet another way to define *effective sweep width*.) The *effective search (or sweep) rate* is simply the product of the *effective sweep width* and the corresponding *search speed*. That is,

[15] $Effective\ Search\ (or\ Sweep)\ Rate = W \times v.$

The *effective search rate* has units of area per unit time (e.g., square miles per hour). Note that multiplying the *effective search rate* by the time, *t*, expended in an area produces the amount of *area effectively swept*.

[16] $Area\ Effectively\ Swept = (Effective\ Search\ Rate) \times t$

A Floor-Sweeping Experiment

Let us recall the second gymnasium floor problem used in Part III and depicted in Figure 22.

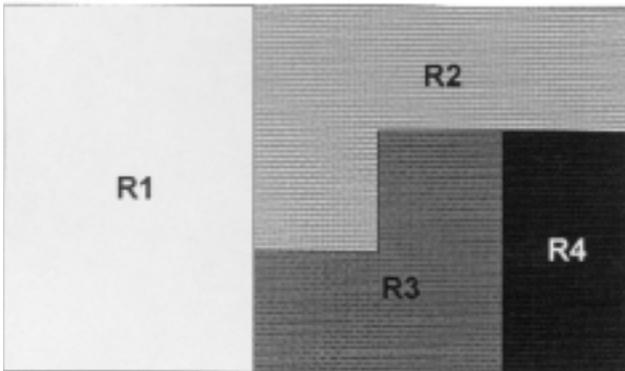


Figure 22

The area shown has dimensions of 30 meters by 50 meters. Each of the four regions can be formed from a whole number of square “cells” measuring 10 meters on a side. Recall that we used sand as an analog for probability, spreading the sand uniformly within the different regions at certain pre-determined densities. Table 8 lists the information on how the sand is distributed.

Figure 23 is a “probability map” corresponding to Figure 22 and Table 8 where the “probabilities” actually show what fraction of the total amount of sand initially present is contained in each cell.

Region	Area (m ²)	Density of Sand (g/m ²)	Amount of Sand Contained (kg)	Percentage of Sand Contained
R1	600	5	3.0	19.35%
R2	400	10	4.0	25.81%
R3	300	15	4.5	29.03%
R4	200	20	4.0	25.81%
Totals	1500	10.33	15.5	100.00%

Table 8

1	2	3	4	5
3.23%	3.23%	6.45%	6.45%	6.45%
6	7	8	9	10
3.23%	3.23%	6.45%	9.68%	12.90%
11	12	13	14	15
3.23%	3.23%	9.68%	9.68%	12.90%

Figure 23

Assume that we have obtained the services of five sweepers for a period of 10 minutes (600 seconds). We will assume they each have a broom like broom B2 from our earlier articles. Recall that B2 is uniformly 50% effective across its one-meter width, giving it an *effective sweep width* of 0.5 meters (m) at a speed of 0.5 meters per second (m/sec) or, equivalently, 30 meters per minute. Having five sweepers for 10 minutes gives us a total of 50 sweeper-minutes. Multiplying this value by the speed, *v*, we get an available *sweeping effort* of 50 minutes times 30 meters per minute or 1,500 meters. We now wish to determine how we should apply this *effort* so that the maximum amount of sand is swept up in the minimum time. Readers are encouraged to pause for a moment to consider how they would approach this problem.

Productive Sweeping Rate (psr)

Let us begin by defining the *productive sweeping rate* as the amount of sand a broom sweeps up per unit time. The *productive sweeping rate* depends on the *effective sweep rate* and the *density* of the sand in the area that is being swept. The *productive sweeping rate (psr)* may be computed using Equation [17].

[17] $psr = (Effective\ Sweep\ Rate) \times Density$

For example, if the *effective sweep width* is 0.5 m and the *sweeping speed* is 0.5 m/sec in all regions, then we see from Equation [15] that the *effective sweep rate* per broom in all regions is

$$\text{Effective Sweep Rate} = W \times v = 0.5 \times 0.5 = 0.25 \text{ square meters per second.}$$

If the density is 20 grams per square meter, then we see from Equation [17] that the *productive sweeping rate* per broom is

$$psr = 0.25 \times 20 = 5 \text{ grams per second.}$$

Computing the values for each region, we get the results shown in Table 9.

Using Productive Sweeping Rate

Sweeping region R4 will produce more sand per second than sweeping anywhere else, at least initially. Table 9 shows that 5 g/sec per broom can be swept up. Working together by spacing themselves one-meter apart in a line abreast formation, the five sweepers can sweep R4 completely one time with a uniform *coverage* of 0.5 in just 80 seconds. The total *effort* expended is 80 sec x 0.5 m/sec x 5 sweepers or 200 meters. This will remove half of the sand, or a total of 2 kg from the 4 kg originally present in R4. (That is 12.90% of the amount of sand initially present in the gym.) However, once this is done, R4 no longer has the highest density. In fact, its density and *psr* have been reduced by half, to those of R2, or 10 g/m² and 2.5 g/sec respectively. Updating Table 9 to reflect this change produces Table 10. Figure 24 shows the updated probability map. Note that *percentages of containment (poc)* and *percentages of sand swept up to date (cumulative pos)* will be computed with respect to the total amount of sand initially present. “Re-normalizing” these percentages to make them reflect the percentages of sand remaining

1	2	3	4	5
3.23%	3.23%	6.45%	6.45%	6.45%
6	7	8	9	10
3.23%	3.23%	6.45%	9.68%	6.45%
11	12	13	14	15
3.23%	3.23%	9.68%	9.68%	6.45%

Figure 24

would add considerably more computation but would not contribute to our purposes. In fact, re-normalizing would make keeping track of actual *densities*, *productive sweeping rates (psr)*, and the total amount of sand removed to date (*cumulative pos*) a much more difficult chore. Although omission of the re-normalization step may be somewhat disconcerting to those familiar with the *rule of Bayes* and *Bayes’ Theorem* from statistics, it will not affect the outcome of our effort allocation decisions, the amount of sand removed or the amount of sand remaining in any way. However, it will make the computations much simpler.

R3 now has the highest *psr* value at 3.75 g/sec. In another 120 seconds, the five sweepers can complete one *coverage* 0.5 sweeping of R3, removing half of its sand, or 2.25 kg. The *effort* required is 120 sec x 0.5 m/sec x 5 sweepers or

REGIONAL VALUES BEFORE ANY SWEEPING						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.25	10	2.50	0	0	0
R3	0.25	15	3.75	0	0	0
R4	0.25	20	5.00	0	0	0

Table 9

REGIONAL VALUES AFTER SWEEPING R4 ONCE						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.25	10	2.50	0	0	0
R3	0.25	15	3.75	0	0	0
R4	0.25	10	2.50	80	200	2
Totals				80	200	2

Table 10

REGIONAL VALUES AFTER SWEEPING R4 & R2 ONCE EACH						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.25	10	2.50	0	0	0
R3	0.25	7.5	1.875	120	300	2.25
R4	0.25	10	2.50	80	200	2.00
Totals				200	500	4.25

Table 11

1	2	3	4	5
3.23%	3.23%	6.45%	6.45%	6.45%
6	7	8	9	10
3.23%	3.23%	6.45%	4.84%	6.45%
11	12	13	14	15
3.23%	3.23%	4.84%	4.84%	6.45%

Figure 25

1	2	3	4	5
3.23%	3.23%	3.23%	3.23%	3.23%
6	7	8	9	10
3.23%	3.23%	3.23%	4.84%	3.23%
11	12	13	14	15
3.23%	3.23%	4.84%	4.84%	3.23%

Figure 26

300 meters. As with R4, R3's *density* and *psr* values are reduced by half, to 7.5 g/m² and 1.875 g/sec respectively. In just 200 seconds of elapsed time, 2 + 2.25 or 4.25 kg of sand out of the total of 15.5 kg on the gym floor have been removed (*cumulative pos* = 4.25/15.5 or 27.42%). After this sweeping, R3 is no longer has the highest *psr*. That honor is now shared between regions R4 and R2, as shown in Table 11 and Figure 25.

Between them, R4 and R2 contain 6 kg of sand. Sweeping these two regions one time at a *coverage* of 0.5 will require 240 seconds (6 minutes). One such sweeping will remove 3 kg of sand, reducing their *densities* to 5 g/m² – the same density as R1. R4 will be left with 1 kg of sand and R2 will have 3 kg remaining. So far, we have removed

2 + 2.25 + 3 or 7.25 kg of the sand in the gym for a *cumulative pos* of 46.77% (7.25/15.5). Note that R4 has been swept twice, R2 and R3 have each been swept once, and R1 has yet to be swept at all. Table 12 and Figure 26 show the state of affairs at this point.

Now, R3 again has the highest *psr* value. In another 120 seconds, the sweepers can remove half of its remaining sand, or 1.125 kg, reducing its *density* and *psr* values to 3.75 g/m² and 0.9375 g/sec respectively. At this point, we have removed 2 + 2.25 + 3 + 1.125 or 8.375 kg of sand (*cumulative pos* = 54.03%) and we have expended (80 + 120 + 240 + 120) sec x 0.5 m/s x 5 sweepers or 1400 meters of *effort*. Table 13 and Figure 27 apply.

REGIONAL VALUES AFTER SWEEPING R4 TWICE, R2 & R3 ONCE EACH						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.25	5	1.25	160	400	2.00
R3	0.25	7.5	1.875	120	300	2.25
R4	0.25	5	1.25	160	400	3.00
Totals				440	1100	7.25

Table 12

REGIONAL VALUES AFTER SWEEPING R4 & R3 TWICE, R2 ONCE						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.25	5	1.25	160	400	2.000
R3	0.25	3.75	0.9375	240	600	3.375
R4	0.25	5	1.25	160	400	3.000
Totals	560	1400	8.375			

Table 13

The honors for having the highest *psr* value are now shared among R1, R3, and R4. However, we have only 100 meters of *effort* left. That is just enough to sweep one cell on the probability map one time at a *coverage* of 0.5. Any of the 12 cells that still contain 3.23% of the original 15.5 kg of sand will do. Each contains 0.5 kg of sand. We choose to sweep the middle cell (8) for reasons that will become apparent later. Sweeping this cell with our remaining effort will remove 0.25 kg of sand, bringing the total amount of sand removed up to 8.625 kg for a *cumulative overall pos* of 55.65%. Figure 28 is the final probability map.

Examining the Results

The plan just described has some interesting characteristics. If the sweepers moved from region to region exactly as described (and we do not count transit times against the 50 available sweeper-minutes), then the amount of sand removed at any point in time during the process would be the maximum possible amount that could have been removed up to that point. Therefore, this is a *uniformly optimal sweeping plan*. Note that the order in which the regions were swept to achieve uniform optimality was neither initially nor intuitively obvious, even though this is a relatively simple problem. In addition to having regions with differing sizes and densities, we could have complicated the problem further by varying the *effective sweep widths* and their corresponding sweeping *speeds* from one region to another as well (and we will, shortly). Also note that one region was not swept at all even though it contained nearly one-fifth of the total amount of sand.

1	2	3	4	5
3.23%	3.23%	3.23%	3.23%	3.23%
6	7	8	9	10
3.23%	3.23%	3.23%	2.42%	3.23%
11	12	13	14	15
3.23%	3.23%	2.42%	2.42%	3.23%

Figure 27

Analogies with Searching

We should pause and consider how the floor-sweeping experiments we have just completed relate to searching. Recall that we are using sand on the floor as an analogy for *probability*. Therefore, searching an area may be thought of as sweeping up probability. The amount of probability swept up is the *probability of success (POS)*. The amount of sand left behind is analogous to the updated (post-search) *probability of containment (POC)*. Therefore,

$$[18] \quad POC_{new} = POC_{old} - POS,$$

where POC_{old} is the *POC* immediately prior to the last search and the *POS* is the value obtained from the last search. Equation [18] is exactly mathematically equivalent to the better-known formula,

$$[19] \quad POC_{new} = (1 - POD) \times POC_{old}$$

Expanding the expression on the right of the equals sign and substituting from Equation [14] easily proves this assertion. The total amount of sand swept up by all sweeping done to date is analogous to the total amount of probability swept up by all searching done to date. The latter is called the *cumulative overall probability of success*. Just as the goal in our floor-sweeping experiments was to sweep up sand at the greatest possible rate, the goal of the search planner is to maximize the increase in the *cumulative overall probability of success* by “sweeping up” probability the greatest possible

1	2	3	4	5
3.23%	3.23%	3.23%	3.23%	3.23%
6	7	8	9	10
3.23%	3.23%	1.61%	2.42%	3.23%
11	12	13	14	15
3.23%	3.23%	2.42%	2.42%	3.23%

Figure 28

rate. To do this, the search planner needs to know the *probable success rates (PSR)* in the different regions. The *probable success rate (PSR)* is exactly analogous to the *productive sweeping rate (psr)* used above. The probable success rate is computed by,

$$[20] \quad PSR = (\text{Effective Search (or Sweep) Rate}) \times Pden,$$

where *Pden* is defined as the amount of probability per unit area in a region or,

$$[21] \quad Pden = \frac{POC}{A},$$

where *A* is the area of the region. We may re-write Equation [20] as

$$[22] \quad PSR = W \times v \times Pden,$$

provided we are careful to remember that the *effective sweep width*, *W*, and *search speed*, *v*, come as together as a single package. (As a practical matter, small changes in *v* will not usually affect the value of *W* seriously.) *PSR* has units of *probability (of success) per unit time*.

Artificial "Features" of the Experiments

With the exception of the last 100 meters of *effort*, the optimal search proceeded by sweeping an entire region exactly once at a *coverage* of 0.5 before proceeding to the next region in the sequence. This is an artifact of our broom's *performance profile* (B2 is uniformly 50% effective for a width of one meter) and our method of sweeping (straight parallel tracks spaced exactly one meter apart). These physical features along with the dimensions of our regions meant that each time we started sweeping a region, we could keep the brooms operating at peak productivity until we had finished sweeping it one (more) time. Very few SAR situations give rise to *detection profiles* that are so uniform and sharply defined. Almost all detection profiles are at their highest close to the searcher's actual track and decline in some fashion as distance from the searcher's track increases. Therefore, we must not jump to the conclusion that all searching should be done at a *coverage* of 0.5. In fact, we can quickly show the fallacy of such a premature conclusion by simply switching brooms. If we use broom B1, (100% effective across a width of 0.5 m), the optimal spacing is 0.5 m, making the *coverage* 1.0. The optimal sweeping sequence becomes R4, R3, R2 and R1. In ten minutes, we will be able to sweep R4 once (160 sec, 4 kg), R2 once (240 sec, 4.5 kg) and 2.5 cells in R2 in the remaining time (200 sec, 2.5 kg) for a total of 11.0 kg of sand. This represents 11/15.5 or about 71% of the sand initially present. Similarly, for broom B4 (25% effective across a width of two meters), the optimal spacing is two meters, making the *coverage* 0.25. The optimal sweeping sequence becomes R4, R4 & R3, R4 & R3 (again), R2, R4 & R3, R2, R4 & R3. The total amount of sand removed will be about 7.877 kg or about 50.82% of the amount initially present. We will not attempt to develop a uniformly optimal allocation for the uneven performance

profile of broom B4 due to its computational complexity. However, we will observe that B4's *performance profile* is probably much more typical of actual SAR *detection profiles* than those of the other brooms.

Although we carefully avoided the overlapping of swaths during any single sweeping of a region, we should not jump to the conclusion that the overlapping of *detection profiles* from adjacent searcher tracks is to be avoided under all circumstances. This is another artifact of our broom's physical characteristics. With realistic *detection profiles*, some overlap is often required to achieve a practical approximation to the *optimal search plan*.

Alternative Strategies

Achieving uniform optimality (Strategy 1) during a single operational period is hard to do as it is often impractical to move searchers around in the manner just described for our sweepers. However, we can still take advantage of the data we have just computed to develop a nearly optimal, but more practical, sweeping (search) plan. If we add up the total amounts of time and *effort* expended in each region by all five sweepers, we get the results shown in Table 14 below.

Region	Effort (m)	Time (sweeper-sec)
R1	0	0
R2	500	1000
R3	600	1200
R4	400	800
Totals	1500	3000

Table 14

Strategy 2: We could assign two sweepers to R3 for the full 10 minutes (600 seconds), allowing them to sweep R3 exactly twice at 0.5 m/sec with their 600 meters of *effort*. We could also assign a third sweeper to R2 for the full 10 minutes (300 meters of *effort*), allowing him to sweep three of that region's four cells exactly once. Finally, we could assign the other two sweepers to R4 for 6 minutes and 40 seconds (400 seconds, 400 meters), allowing them to sweep it exactly twice. For the remaining 3 minutes and 20 seconds (200 seconds, 200 meters), these two sweepers would be assigned to the remaining unswept cell in R3 which they would be able to sweep exactly twice. This plan significantly reduces the need to move sweepers from one region to another as compared to the previous plan. It is also a *T-optimal sweeping plan* because it would remove the same total of 8.625 kg of sand in the allotted ten minutes. However, in the early stages, it would not remove sand as quickly as the uniformly optimal sweeping plan did.

Strategy 3: If we are not required to recognize regional boundaries for the purpose of sweeping the floor, we can develop another *T-optimal plan* that requires no inconvenient movements of sweepers from place to place. Going back to our initial "probability map," Figure 23, we could

assign two sweepers to cells 13, 14 and 15, two sweepers to cells 8, 9 and 10, and one sweeper to cells 3, 4 and 5 for the full 10 minutes in each case. We could have them sweep first from right to left across all three of their assigned cells, then back and forth until their *efforts* were uniformly spread over their assigned cells. Cells 8, 9, 10, 13, 14 and 15 would all be swept twice while cells 3, 4 and 5 would be swept once. We would still get a total of 8.625 kg of sand in ten minutes and the sweepers' brooms would never need to leave the floor.

Now let us look at three other strategies that might have suggested themselves before we performed our experiments above.

Strategy 4: One of these alternatives would be to sweep the regions one time each in descending order of how much sand each contained initially. The order of sweeping in this case would be R3, R4, R2 and R1, if we allow *density* to be the tiebreaker between R4 and R2. There happens to be exactly enough effort available to carry out this strategy. The end result would be removal of 7.75 kg of sand, or 50% of that initially present.

Strategy 5: Another strategy that might come to mind would be sweeping the regions once each in descending order of *density*. The sweeping order in this case would be R4, R3, R2 and R1. Again, 7.75 kg of sand, or 50% of that initially present, would be removed.

Strategy 6: Finally, since we have five sweepers and each sweeper can sweep three cells exactly once in ten minutes, we could assign one sweeper to each "column" of three cells in Figure 23.

Figure 29 graphs the results of our sweeping strategies, showing what percentage of the sand initially present each strategy sweeps up as a function of time as the sweepers pass back and forth over their assigned cells.

The upper curve, labeled "S1" for "Strategy 1," shows the results of applying the *uniformly optimal sweeping plan*. In our simple example, Strategy 5 is also uniformly optimal for the first six minutes. Strategy 4 catches up with S1 at three minutes, twenty seconds and manages to stay with it until six minutes have passed. Then both S4 and S5 become sub-optimal and depart sharply from the optimal curve. Strategies 2 and 3 do not catch up with S1 until five minutes have passed, but then they remain very close to the uniformly optimal curve for the remaining five minutes. Interestingly, the first five strategies are all *T-optimal* when *T* equals five minutes. Only the first three are *T-optimal* when *T* is ten minutes, while the other three are all sub-optimal at that point. Finally, the worst plan of all is clearly Strategy 6. Strategy 6 is the one where the available *effort* is spread uniformly over the floor for the entire ten minutes and is never concentrated anywhere. In other words, Strategy 6, if applied to searching, would seek to obtain the same *POD* everywhere at once. This is rarely the best plan.

Again, the author must caution the reader against jumping to conclusions. The excellent performance of Strategy 5 in the early stages does not imply searching regions in descending order of *probability density* is always a good way to start. In developing examples, the author has struggled mightily to balance the competing demands of making them simple enough to follow yet complex enough to reflect reality. It has not been easy and the author has not always been entirely successful. In the next paragraph, we will complicate matters enough to show why one should not jump to "obvious" conclusions.

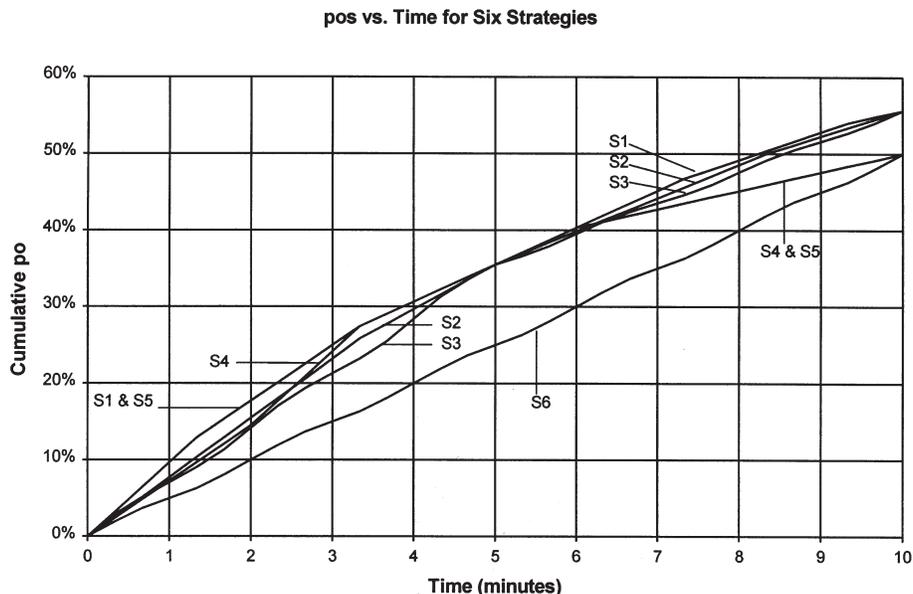


Figure 29

A More Complex Experiment

We will now move a step closer to simulating a search situation with our floor sweeping analogy. Suppose we return to our initial problem represented by Figures 22 and 23, and Table 8. However, suppose that one morning when we enter the gym to set up our experiments, we find the floor is being refurbished. It is no longer uniformly smooth everywhere but has varying degrees of roughness. This forces us to perform some additional experiments, like those of Part I, to determine new *effective sweep width* and *sweeping speed* values for the new conditions. We conduct our experiments using broom B2 and find that although it remains uniformly effective across its one-meter width, the level of that effectiveness (i.e., the *effective sweep width*) varies with the roughness of the floor. We also find that eventually the roughness of the floor also impacts the *speed* at which the sweepers move. The results of the experiments are shown in Table 15 below.

Region	Effective Sweep Width (m)	Sweeping Speed (m/sec)	Effective Sweep Rate (m ² /sec)
R1	0.5	0.50	0.25
R2	0.4	0.50	0.20
R3	0.3	0.40	0.12
R4	0.2	0.25	0.05

Table 15

Apparently region R1 is still in its original condition for our purposes since neither the *effective sweep width* nor the *sweeping speed* have changed. In R2, the *sweep width* suffers somewhat, but not the *speed*. Both R3 and R4 suffer increasingly lowered *sweep widths* and *speeds*. Using our new *effective sweep rate* values to compute new *productive sweeping rates* for sweepers using a B2-type of broom, we get the results shown in Table 16.

One thing is abundantly clear: Five sweepers will not be able to sweep up nearly as much sand in ten minutes as they did before. We also see that the *productive sweeping rate values* no longer parallel the *density* values. Because we are using the same brooms as before (all identical to B2), we will still want to use parallel tracks at a spacing (*S*) of one meter. However, our *effective coverage* values for each region will now be different at this spacing thanks to the differing *effec-*

tive sweep widths. Using the shortcut formula for *coverage* that is valid for parallel sweeps of rectangular areas,

$$[23] \quad C = \frac{W}{S},$$

where *C* is the *coverage*, *W* is the sweep width and *S* is the track spacing, we find we will be using *coverages* of 0.5 for R1, 0.4 for R2, 0.3 for R3 and 0.2 for R4. From Part II of this series, we recall that at these low *coverages*, the *percentage of remaining sand swept up* by B2 at each sweeping will equal the *coverage*. We begin with sweeping R2 one time. It takes 160 seconds for our five sweepers working together to “cover” R2 once at 0.5 m/sec. Having done so, they have removed 40% of the 4 kg of sand initially present or 1.6 kg. This leaves 2.4 kg behind in R2. Table 17 summarizes the situation following this sweeping.

Sweeping R3 one time at the reduced speed of 0.4 m/sec requires 150 seconds and a total expended effort of 300 m. Note that while the *efforts* required to sweep the regions does not change from our previous experiments, the *times* required to expend those *efforts* when sweeping *speeds* have been reduced must increase accordingly. We sweep up 30% of the 4.5 kg of sand initially present in R2, or 1.35 kg. This leaves 3.15 kg behind. Table 18 summarizes the situation following the sweeping of R3.

With more than half of our time used up, we have swept up less than 3 kg of sand. Since R3 still has the highest *productive sweeping rate*, we sweep it again. This time, we get 30% of the remaining 3.15 kg of sand or 0.945 kg. The results are summarized in Table 19.

We have now used 460 seconds of our original 600 seconds (ten minutes) of sweeper availability, leaving 140 seconds. The next region we want to sweep is R1 since it has the highest *psr* value. At 0.5 m/sec, our five sweepers can sweep 3.5 cells out of the six in R1, removing 50% of the sand present in the swept area. The amount of sand contained in 3.5 cells of R1 is 3.5 x 0.5 kg/cell or 1.75 kg and 50% of this value is 0.875 kg. Table 20 summarizes the results after 10 minutes of *uniformly optimal sweeping*. *Density* and *psr* values for R1 are given for both the swept and unswept portions.

Note that even with a uniformly optimal plan, we have managed to sweep up only 4.77/15.5 or 30.77% of the sand initially present. Figure 30 shows how the final “probability map” would look if the 3.5 cells in R1 chosen for sweeping were 2, 7, 12 and one-half of 11.

REGIONAL VALUES BEFORE ANY SWEEPING						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.20	10	2.00	0	0	0
R3	0.12	15	1.80	0	0	0
R4	0.05	20	1.00	0	0	0

Table 16

REGIONAL VALUES AFTER SWEEPING R2 ONCE						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.20	6	1.20	160	400	1.6
R3	0.12	15	1.80	0	0	0
R4	0.05	20	1.00	0	0	0
Totals				160	400	1.6

Table 17

REGIONAL VALUES AFTER SWEEPING R2 & R3 ONCE EACH						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.20	6	1.20	160	400	1.60
R3	0.12	10.5	1.26	150	300	1.35
R4	0.05	20	1.00	0	0	0
Totals				310	700	2.95

Table 18

REGIONAL VALUES AFTER SWEEPING R2 ONCE, R3 TWICE						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	5	1.25	0	0	0
R2	0.20	6	1.20	160	400	1.600
R3	0.12	7.35	0.882	300	600	2.295
R4	0.05	20	1.00	0	0	0
Totals				460	1000	3.895

Table 19

REGIONAL VALUES AFTER SWEEPING R2 & R1 ONCE EACH, R3 TWICE						
Region	Effective Sweep Rate (m ² /sec)	Density (g/m ²)	Productive Sweeping Rate (psr) (g/sec)	Expended To Date		Amount of Sand Swept Up To Date (kg)
				Time (sec)	Effort (m)	
R1	0.25	2.5 & 5	0.625 & 1.25	140	350	0.875
R2	0.20	6	1.20	160	400	1.600
R3	0.12	7.35	0.882	300	600	2.295
R4	0.05	20	1.00	0	0	0
Totals				600	1350	4.770

Table 20

1	2	3	4	5
3.23%	1.61%	2.58%	2.58%	2.58%
6	7	8	9	10
3.23%	1.61%	2.58%	4.74%	12.90%
11	12	13	14	15
2.42%	1.61%	4.74%	4.74%	12.90%

Figure 30

Note also that the optimal plan for this problem placed no effort at all in the region having the highest density (R4) whereas the optimal plan for the previous problem ignored the region with the lowest density (R1).

Examining Alternative Strategies

We will not take time to analyze all the different strategies examined above, but it will be instructive to look briefly at the three “obvious” alternatives; namely Strategies 4, 5 and 6. Recall that in Strategy 4, regions are swept in order of decreasing *percentages of containment (poc)* or, equivalently, decreasing amounts of sand. In Strategy 5, regions are swept in order of decreasing *densities*. In Strategy 6, one searcher is assigned to each “column” of three cells on the “probability map” and sweeps those “columns” for the entire ten minutes, starting at the “bottom” and moving “upward” on the first leg. We will not belabor the reader with the nec-

essary arithmetic. We will simply show the graph of the results as Figure 31.

Figure 31 makes it quite clear that none of the more simplistic strategies work very well for the more complex problem we have just examined. This time, sweeping the regions in descending order of *density* (Strategy S5) was the worst thing to do in the early stages. However, the results of Strategy S6 are similar to those of the previous experiments in that apportioning the resources, i. e. the sweepers, evenly over the entire area again turned out to be the worst strategy at the end of ten minutes time. Note that a uniform distribution of resources over the floor’s area does not produce a uniform distribution of *effort* this time, nor does it cause a uniform *percentage of dirt (pod)* to be swept up everywhere. In fact, Strategy S6 has left some portions of the floor unswept this time because three of our five sweepers could not complete their assigned “columns” of three cells each in the allotted time due to the speed reductions in regions R3 and R4.

Assessing the Costs of Sub-Optimal Planning

One may think that the maximum difference between the optimal *pos* value of 19% (2.95 kg of sand) and the 14% (2.15 kg) or 15% (2.32 kg) of the other strategies at 5 minutes, 10 seconds in Figure 31 is small. However, the perspective changes when one considers the additional *time* and/or *effort* required to obtain the optimal *pos* values. For example, it will take the same five sweepers 6 minutes, 30 seconds to reach the 19% mark based on the next best strategy of the three alternatives considered. That’s nearly 26% more time to get the same result. In a SAR mission, adding that much time to reach an early *POS* goal could be serious. Alternatively, shaving 20% or so off the time required to achieve a 19% *cumulative probability of success* by making

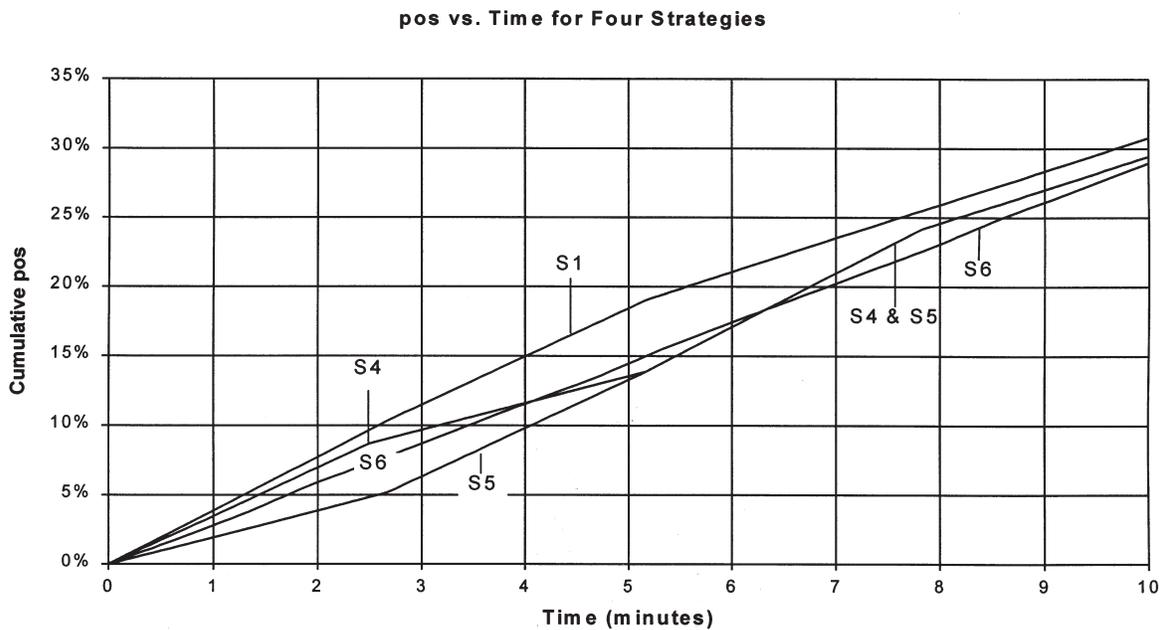


Figure 31

more efficient use of the available resources could contribute substantially to increasing the number of successful missions. Another way to view the issue is to consider how much more *effort* would be required to achieve the optimal result when using sub-optimal strategies. We would have to increase our available *effort* by about 26%. Roughly speaking, this means we need 26% more sweeper-minutes to make the sub-optimal strategies produce results as good as the uniformly optimal plan in the first 5 minutes. That is a substantial increase. (Of course, if we had the additional effort, we would want to set a new optimal *POS* target.) The point is that we can pay a significant price when search plans are sub-optimal.

The Charnes-Cooper Algorithm

One feature of both the above *optimal effort allocation* problems is particularly worth noting. As optimal sweeping progressed, the *productive sweeping rates* tended to become more and more nearly the same everywhere as sand was removed. In fact, an *optimal effort allocation* strategy does seek to “level the playing field,” as it were. In other words, the general idea is to search the region with the highest *probable success rate (PSR)* until enough probability is swept up to make the *PSR* there equal to the second highest *PSR* in the list. Then both regions are searched together so their *PSR* values are kept equal to one another as they decrease at the same rate to the *PSR* value of the third highest value in the list. This process continues until as many regional *PSR* values as possible are the same. Any remaining *effort* is then spread over these regions in a fashion that keeps the *PSR* values equal to one another as they all decrease together toward the next level. Note that the distribution of effort required to keep *PSR* values dropping at equal rates everywhere is not uniform. We must still compute an appropriate *coverage* and corresponding level of effort for each region using the *W*, *v*, and *Pden* values appropriate to that region and the resources available to search it.

Note: In the examples given above, physical constraints and the method of sweeping i.e., using brooms one meter in width moving along perfectly straight, parallel tracks one meter apart, forced us to push psr values down below the next highest value before moving on to the next region. We could not move on until completing the current region because if we did, we would leave some portion of the current region behind that still had high psr and move to an area with a lower psr. Again, the need to use simple, easy-to-visualize examples prevents us from having an exact analogy with the mathematical principles involved.

In 1958, A. Charnes and W. W. Cooper developed an algorithm for computing the optimal distribution of *effort* for situations where the *probability density distribution* was known and the exponential detection function applied. Stone [6] describes an adaptation of this algorithm in some detail. Conceptually, the algorithm works as described above, i.e., it “levels the playing field” in terms of *probable success rate*. However, there is a good deal of mathematical detail needed to make the concept work. That detail involves every equation presented in this series of articles as well as others. The algorithm is really practical only if a computer is available on

which it may be programmed and run. The good news is that the algorithm is not too difficult to program and it is very efficient. The bad news is that it computes what is known as an *unconstrained optimization*. It makes no allowances for the real-world limitations on how resources may be deployed. If the algorithm computes that the *effort* represented by one searcher searching for one hour is needed in an area of 160 acres, then it will assume, for computational purposes, that the searcher can somehow uniformly “search” the entire 160 acres in that hour, albeit with a very low *coverage* and *POD*. There are other assumptions that are equally unrealistic. However, if the algorithm’s intermediate iterative workings are not taken too literally, its final results can be quite useful in a practical sense. The algorithm can even be modified to accept “effort” defined in the more commonly used terms of searcher- or resource-hours instead of the classical search theory definition using distance. If the algorithm is run for the total amount of available “effort” (and the *sweep widths*, *search speeds*, *probabilities of containment*, *areas*, etc., have been entered for all the regions) the search planner can see how much of that “effort” (e.g., how many searcher-hours) were accumulated in each region. The results of such computations will provide a very useful guide to the search planner regarding where he should place the available resources during the next search cycle. In other words, the output of the Charnes-Cooper algorithm may be used in the same way we used the final values in Table 14 to develop alternative Strategies 2 and 3 for our first set of experiments.

Lessons

We have covered a great deal of ground in these articles. Some of the things we have learned along the way are:

- The goal of search planning is to maximize the *cumulative overall probability of success* and minimize the time required to achieve it within the constraints of the available resources.
- The *probability of detection (POD)* is an estimate of the chances that the object of the search would have been detected if it had been in the searched area during the search.
- The concept of *effective search (or sweep) width* as a quantitative measure of “detectability” is the key to objective, reliable and consistent *POD* estimates. Without this concept and supporting data from rigorous scientific field experiments, *POD* estimates must necessarily be regarded as highly subjective “guess-timates.”
- The concept of *effective search (or sweep) width* lies at the very core of search theory and is the key to planning effective, efficient searches of areas and evaluating search results.
- In search theory, *effort* is defined as the distance a searcher travels while searching within a defined area.
- The concept of *effective coverage* relates the *effective sweep width* and the amount of effort expended in an area to the size of that area.

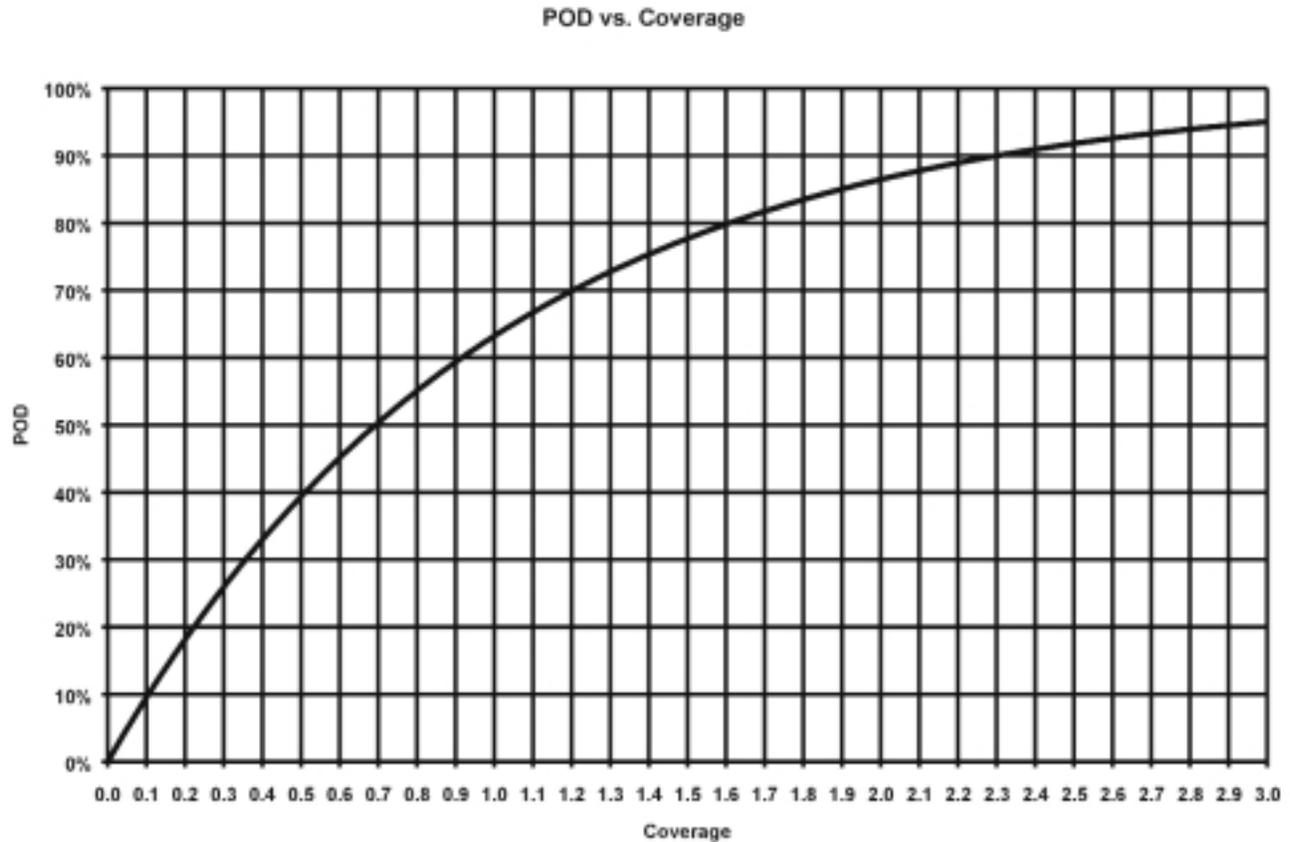


Figure 32

- *POD* is a function of *coverage*, as depicted in Figure 32. The graph shown there is that of the exponential detection function for *coverages* ranging from zero to 3.0.
- There is no theoretical basis for the claim that two successive *low-coverage* (i.e., low *POD*) searches of a region will produce a higher *cumulative POD* for the same effort as a single *higher-coverage* (i.e. higher *POD*) search would. In fact, search theory suggests the opposite effect is far more likely.
- The probability that a region, segment, or other geographically defined area contains the search object is called the *probability of containment (POC)* or, equivalently, the *probability of area (POA)*.
- Initial *POC* values are estimated subjectively by *scenario analysis* and *consensus*. For a particular scenario, *POC* values for sub-divisions of the *possibility area* may be estimated directly or by assignment of *proportional assessment* values that are then normalized to produce *probabilities of containment*. However, simplistic schemes that only assign ranks and thus do not keep the values in the correct proportions to one another can lead to *POCs* that are inconsistent with the evaluators' assessments of the available information. Such schemes should be avoided.
- When multiple scenarios are under consideration, they may be assigned different "weights" to reflect their relative likelihoods of representing the true situation.
- The *probability density (Pden)* of a region, segment, or other geographically defined area is the ratio of its current *POC* to its *area*.
- A *probability map* is a regular grid of cells where each cell is labeled with the amount of probability it contains. A *probability map* is constructed by laying a regular grid over a map labeled with the results of the scenario analysis and consensus processes. *POC* values for the cells are then computed from the regional *POC* and *Pden* values established by scenario analysis and consensus. A cell's *POC* is based on the product of its area and the *probability density* of the region it lies within. If the cell spans more than one region, then the areas of the fractions of the cell lying in each region are multiplied by the respective regional *densities* and the results added together to get the cell's *POC*. Because all cells are the same size, the *Pdens* of the cells are proportional to their *POCs*. Therefore, a *probability map* shows at a glance both the *containment probabilities* and where the *probability densities* are high and where they are low. The *probability map* is a representation of the search object's location *probability density distribution*.
- Developing *scenarios* and their corresponding *probability density distributions*, or *probability maps*, from evidence, clues, behavior profiles, historical records, and any other available information is not a simple task. *Scenario analysis* is an essential part of the search planning process.

deserving more dedicated time, resources and attention than it generally gets.

- The probability that a search of a region, segment, or other geographically defined area will, or should, locate the search object is called the *probability of success (POS)*. *POS* is a function of *POD* and *POC*—in fact, it is the product of the two.
- Searching an area is tantamount to “sweeping up” or “removing” probability from it. The amount removed is the *POS* while the amount remaining is the new post-search *POC*. The *POD* of the search determines both values.
- The *cumulative overall probability of success* is the sum of all individual *POS* values achieved to date. It measures the chances of having found the search object if it was anywhere within the *possibility areas* of the *scenarios* under consideration. Achieving a high *cumulative overall POS* value without locating the search object is an indication that further searching based on the scenarios currently under consideration is unlikely to be successful. It is also an indication that a thorough re-evaluation of all available data and information is needed to determine whether key facts have been overlooked, whether other plausible scenarios exist, etc.
- The *probable success rate (PSR)* for a region, segment, or other geographically defined area is an estimate of the rate at which *POS* can be increased by searching there. *PSR* is the product of the *effective sweep width*, the corresponding *search speed*, and the area’s *probability density (Pden)*.
- The *optimal allocation* of search resources is not a simple task—in theory nor in practice. Simplistic guidelines about placing most of the resources where the *probability of containment* is highest, or where the *probability density* is highest, are unreliable.
- In the most basic terms, the idea behind *optimal effort allocation* is to put search resources into the region(s) where probability can be swept up most quickly, moving them to other regions when and as necessary to ensure they are always searching where they can be the most productive.
- For large-scale searches involving significant amounts of area and requiring more than a few hours to resolve successfully, search theory, properly applied, can substantially improve success rates in most jurisdictions.

A Final Word

In these articles, we have not developed a practical search planning methodology based on search theory nor was it our intent to do so. However, perhaps we have at least raised the reader’s awareness of the potential benefits that development

of such a methodology would bring to SAR missions. A project to produce a set of scientifically valid yet practical search planning procedures would require a development team whose collective talents and knowledge covered the entire spectrum from the most mathematically esoteric aspects of search theory to the most practical aspects of planning and conducting search operations. Such an undertaking would also require significant amounts of time and resources. Tasks would include:

- Designing, developing and conducting sweep width experiments in inland environments.
- Testing and evaluating different search tactics.
- Integrating material and knowledge from many diverse areas of expertise into a clear, concise, coherent, practical, and scientifically valid set of guidelines for planning searches.

However, it is to be hoped that we have provided a rarely-seen scientific perspective on the nature of searching that will help search planners think about and approach search problems in new and improved ways even before these tasks can begin.

Consolidated References

1. Koopman, B.O., *Search and Screening*, OEG Report No. 56, The Summary Reports Group of the Columbia University Division of War Research, 1946. (Available from the Center for Naval Analyses.)
2. Koopman, B.O., *Search and Screening: General Principles with Historical Applications*, Pergamon Press, New York, 1980.
3. *National Search and Rescue Manual* (U.S.), 1991.
4. Hill, K., “Search, Detection, and the Visual Briefing,” presentation at *NASAR’s 27th Annual Conference and Exhibition*, Portland, OR, May 27-30, 1998.
5. Stone, L.D., “The Process of Search Planning: Current Approaches and Continuing Problems,” *Operations Research*, 31 (March-April 1983), 207-233.
6. Stone, L.D., *Theory of Optimal Search*, 2nd ed. Operations Research Society of America, Arlington, VA, 1989.