

Second Draft
For Review & Comment

Controversial Topics In Inland SAR Planning

A NEWSAR White Paper
February 2004

By

Daniel O'Connor
NEWSAR

Copyright 2004 by Daniel O'Connor
for Northeast Wilderness Search And Rescue,
POB 920362 Needham, MA 02492

Table of Contents

Overview 3

Introduction 4

Open Systems vs. Closed Systems 7

Grid Searching & Manpower Requirements 9

POA vs. POD 11

Tactical Decision Aids. 13

Mattson Consensus vs. Initial Distribution 14

References 15

Appendix A 16

Overview

Inland Search & Rescue Planning is, in some quarters, experiencing a crisis of confidence. Strident voices in the extended SAR community are calling for Inland SAR planners to abandon traditional techniques and adopt wholesale, the methods of Maritime SAR.

The argument for this abrupt transition generally goes as follows: Maritime or Marine SAR (MSAR) is based on sound mathematical footing, dating back to Bernard Koopman's virtual invention of Search Theory during World War II. MSAR is a sophisticated sub-field of the broader discipline of Operations Research, while Inland SAR (ISAR) planning, as traditionally taught, is a makeshift assembly that is "unscientific."

While it is not in the scope of this brief paper to recount the evolution of ISAR methods, we argue that far from being "makeshift," classic ISAR is a creature in the family tree spawned by Koopman, an adaptation, filling a "niche" as it were, to situations and conditions not encountered in the marine environment. This is not to say that ISAR cannot learn from its more sophisticated cousin. It can indeed, and should embrace MSAR methods when they are appropriate.

Two of the early pioneers of ISAR, William Syrotuck and Robert Mattson, were field grade officers in the U.S. Air Force, undoubtedly familiar with the SAR techniques developed in the 1940's. Variations from those techniques, and later additions by others (Bownds et al), were not makeshift at all, but rather adaptations to the practical realities of moving from the two-dimensional plane of the ocean to the multi-dimensional land search.

The "controversial topics" of the title are areas where we contend that specific MSAR methods or concepts are either flatly inappropriate for ISAR, or where their use in ISAR should be treated as an option and not a mandate. Because the application of search theory directly affects the deployment of resources and the ability to save lives, criticism and counterarguments will always be carefully considered. We would encourage feedback from all interested parties.

This white paper is meant to serve a two-fold purpose: first, to reassure practitioners of classic Inland SAR planning that they need not immediately abandon all of its techniques developed over the last 30-odd years, and secondly, to formalize NEWSAR's position on these increasingly controversial topics.

Introduction

Before tackling the areas of controversy, it will be helpful to review commonalities shared by ISAR and MSAR, and some of the differences between them.

The basis of modern search theory can be summed up in the equation:

$$\text{POA} \times \text{POD} = \text{POS}$$

POA or Probability Of Area, is the likelihood that the subject (target) of a search is in a particular geographic sub-set of the Search Area (SA). In the MSAR world this is called POC, Probability Of Containment. POD or Probability Of Detection, is a measure of a resource's (sensor's) ability to detect the subject if it were actually in the area of interest.

Multiplying POA times POD creates an interaction between locational likelihood and the ability of a resource, or set of resources, to detect objects ranging from small clues, like the stub of a cigarette, to people, downed aircraft and even large vessels. This mathematical interaction creates the value of POS, or Probability Of Success, a measure of how much of the SA was effectively swept, covered or searched.

If a resource with a POD of 80% searches an area with a POA of 40%, then the resulting POS ($0.8 \times 0.4 = 0.32$), assuming all inputs are accurate, is equal to 32%. This means that 32% of the defined SA has been cleared and is not likely to contain the subject. However, the remaining 8% of the starting 40% might still contain the subject. This is due to the fact that there was a 20% chance for the resource, with a POD of 80%, to miss the target

Use of POS in this way assumes that the subject is either stationary or moving in a uniform way, and does not re-enter the searched areas. If the subject is not found and the sum of these POA x POD interactions approaches unity, the incident manager gains confidence, in the case of a maritime search, that the distressed vessel has sunk, or in the case of an inland search, that the subject is not in the SA. OPOS or Overall Probability Of Success, is the sum of all computed POS interactions; it approaches unity when the search has included near exhaustive effort within the defined SA.

Special care must be taken in land searches to ensure that the multi-dimensionality of the SA is considered. Unlike the 2-dimensional surface of the ocean, inland searches, in addition to the normal variations in terrain, can include a myriad of 3-dimensional sub-areas like lakes, ponds, rivers, trees, tree hollows, hollow logs, caves, wells, graves, mine shafts and buildings, within which are basements, attics, closets, air ducts, chimneys, spaces between walls, crawl spaces, etc., all of which may need consideration as a search segment. With the exception of some specialized searches for high value targets like a lost nuclear submarine, MSAR does not typically search below the surface of the ocean.

As Richard J. Toman has shown, inland bodies of water need to be sub-divided into three segments [1]. Buoyancy characteristics of a corpse can vary substantially, placing the recovery subject on the surface, on the bottom of a pond, lake or river, or suspended somewhere in-between.

Both MSAR and ISAR agree on the mathematical formulation $POA \times POD = POS$. Regarding this formula, Alan R. Washburn notes, “Thus the fundamental Bayesian paradigm is the same in both maritime and inland search. However, the two differ in that the idea of a motion model seems to be missing in inland search; that is, POAs do not ‘shift’ merely on account of the passage of time.” [2]

For ISAR, initial POA is generally established by a set of best guesses from a panel of experts. For MSAR, initial POA is developed through mathematical models driven by known or forecast environmental conditions like winds and currents. POD is a highly calibrated number in MSAR with its cadre of professional searchers. In ISAR, POD is more problematic and often simply an informed estimate. The lack of highly calibrated sensors in ISAR is not by choice but rather the nature of the beast. Calibration of sensor coverage and derivative POD requires time, money and a unifying command authority not presently available to the ISAR community.

As Washburn perceptively notes, ISAR and MSAR differ in other fundamental ways: “The inland SAR situation is essentially different from the marine situation in that many more organizations are involved. . . . inland SAR is conducted by a host of organizations, most of them voluntary. . . . The diffuse nature of the responsibility for inland search leaves it without a sponsor for development of either theory or software.” The differences are not just organizational, as Washburn continues, “Inland search must cope with the effects of terrain, a complication that is missing in maritime search.” [3]

We would add one other salient difference: the potential for foul play. Outside of combat situations, the underlying assumptions for finding a missing vessel in MSAR are driven by measurable physical factors like wind and sea currents. While maximum theoretical distances can be computed for ISAR searches, statistical distributions of lost person behavior are rudimentary where they exist at all. In ISAR, initial and subsequent assumptions are not confined to physical limits but are subject to the vagaries of the human psyche. Human agents can and do fake their own disappearance, and the subjects themselves vary broadly: hunters, hikers, elderly Alzheimer victims, children, generic “missing persons,” etc., all of which exhibit behavioral differences across classes and as individuals within broad classes. Some incidents are searches for abductees or other victims of foul play. Because of the highly volatile human elements, every ISAR Search Area, unlike MSAR target spaces, must always be treated as a potential crime scene.

ISAR has also developed the concept of the “Rest Of the World” (ROW), missing in the closed systems of MSAR. ROW acts as a shadow segment in ISAR, accumulating POA as segments are swept and the missing subject is not found. ROW is a characteristic of an open system, one that allows the possibility that the subject is not in the defined search area. ROW represents the space, i.e., the remainder of the planet, outside of the Search Area.

In MSAR, there is typically no ROW. The probabilities associated with swept areas recycle through the SA, with total POA equal to the sum of the initial POA after an infinite number of updates. MSAR techniques generally assume that the SA, as defined, always contains the missing subject within the closed system.

Proponents of MSAR techniques for ISAR insist that ROW be dropped in favor of a closed system, the subject of Topic 1. This unyielding attitude is curious, since the field of Operations Research (OR) makes no such demand. In OR, a probability distribution which does not sum to 100% is called a “defective” one. Except for special cases, defective target distributions are usually allowed. When a distribution sums to less than 1, “this simply means that the target has some probability of being outside the region in which (the) search is to be conducted.” [4] In ISAR, we call that area outside the region where the search is conducted the Rest of the World.

Topic 1: Open Systems vs. Closed Systems

In our view, the second-most misplaced idea from proponents of MSAR techniques is the insistence that ISAR operations be confined within closed systems. Certainly, successful searches can be conducted inland with closed search areas. However, a closed area severely limits a number of adaptive qualities characteristic of open systems.

Among the advantages of an open system are the following:

1. The ability to start a Search with a few segments, without having to divide up the entire theoretical Search Area, easing the initial computational burdens and mapping tasks by eliminating areas that are not immediately relevant.
2. The ability to “follow the clues,” i.e., to expand the Search toward the subject by adding additional segments based on new information without the requirement of developing a new consensus.
3. Allow the inclusion of non-contiguous segments. Did the despondent subject travel to a summer home out of the primary search area?
4. Monitor R.O.W. probability, the likelihood that the subject is no longer in the Search Area, or in the case of Bastard Searches, never was.
5. Deal with new information just outside the defined Search Area by expanding the SA out of R.O.W. and adding the Influence of Clue to shift POA toward the newly created segment.
6. Provides the potential to Shift POA without POD. POA is not recycled internally, but can be reintroduced into the defined Search Area from the accumulated total in ROW. The search effort of resources that usually do not produce POD, like tracking dogs, hasty teams, man-trackers, etc. can be quantified as clues and used to shift POA.

The list is not meant to be comprehensive, but representative of useful adaptations when employing open systems. Table 1 summarizes some of the differences between open systems in ISAR and closed systems in MSAR. The abbreviation “TDA” under TDA Support refers to Tactical Decision Aids, or computer programs used to manage the Bayesian mathematics and other associated search tasks, discussed in more detail as the fourth topic.

Closed systems require the SAR planner to define a search area that guarantees containment of the subject (target). This makes sense in MSAR where lack of a find indicates a vessel has sunk. In ISAR, this causes the Search Area to be over-defined, much larger than it need be if ROW was used. With ROW, the ISAR planner can grow the search toward the subject, on the basis of clues, as new information arises.

Table 1. Characteristics of ISAR Open & MSAR Closed Searches

Characteristic	Open-Inland	Closed-Maritime
Search Area is 3-Dimensional	Yes	No
Includes R.O.W.	Yes	No
All Segments must be pre-defined	No	Yes
Initial POA Distribution must encompass entire SA (no R.O.W.)	No	Yes
Initial POA Distribution derived from Physical Properties of Target and Environment	Sometimes	Always
Initial POA Distribution from Behavioral Properties of Humans	Usually	Never
Conducive to the discovery and assessment of Clues	Yes	No
POA changes over time, even if no search	No	Yes
Can accommodate non-contiguous segments	Yes	No
Can be readily applied in criminal investigations	Yes	No
TDA Support	CASIE3 (DOS) - Free CASIE4 (WIN) - Free Search Manager(WIN)-\$\$	CASP (USCG) NODESTAR (USN) both proprietary

Topic 2: Grid Searching and Manpower Requirements

A closed search area generally indicates grid searching as the primary detection tool. In ISAR, a grid search is often considered a last resort, a manpower intensive technique destructive of clues. Since the maritime environment is not clue rich (no footprints left in the water), and distances are vast, grid searching for the subject (lost person or vessel) via aircraft is the primary technique in MSAR.

The following example is offered to demonstrate the feasibility of grid searching in Maritime searches and its inappropriateness at the beginning of Inland searches.

Imagine a 10-by-10 mile target space, a perfect 100 square miles, where a search is initiated for a lost vessel. This can also be thought of as a checkerboard of 100 1-by-1 mile squares. The crew of a maritime search aircraft decides to use a sweep width of 1/8th of a mile (660 feet) with a lateral range of 330 feet on either side of the aircraft. At some chosen altitude, this provides a reasonably high POD or coverage. Dividing ten miles by 1/8th results in a requirement of flying 80 sensor tracks through the area. To be conservative, the crew adds 2 tracks to cover entry and exit at both ends of the 100 square mile box. Flying 82 ten mile tracks requires 820 flight miles. Adding 10 more miles, to account for the 80 eighth-of-a-mile shifts between tracks, results in a total of 830 miles. A turboprop or turbojet aircraft traveling at a stall-safe approach speed of 165 knots could search this 100 square mile area in about five hours. At faster speeds, in less than five hours.

Compare this to planning an Inland grid search of 1 square mile. Using standard ISAR planning formulas [5] with twenty foot spacing (a lateral range of ten feet on either side of the observer), a reasonably high POD or coverage could be expected. Allowing for two sweeps per operational period, one mile up and one mile down, 132 trained responders would need to deploy for seven hours, to grid search a 1-by-1 mile square.

To grid search 100 square miles, like the maritime example, in a similar timeframe at an estimated speed of one mile every 3.5 hours, the manpower requirement is astounding: 13,200 grid-searchers for seven hours! [6] Even spreading out these deployments over a week, 1,886 trained and rested grid searchers would need to be activated everyday. Clearly, ISAR units cannot provide this absurd level of manpower for extensive grid searching. ISAR needs to rely on non-grid search paradigms to compensate for its lack of speed in its human sensors and its inability to field a small army of trained responders at every emergency.

Is it unfair to compare a 100 square mile Search Area in MSAR to one in ISAR? We think not. Consider how long it takes a missing hiker to create a 100 square mile theoretical search area. A standard planning speed for a lost person is two miles per hour. Even at 25% of this figure, a lost hiker traveling at one-half mile an hour, for twelve hours on a summer day, could easily wander six miles away in a half-day. Using

the formula for the area of a circle, pi times the radius squared, a six mile radius from the subject's Last Known Point or Point Last Seen creates a circular theoretical search area of about 113 square miles, in just half a day! A thorough grid-search of 113 square miles would require something on the order of 105,000 searcher-hours (15,000 searcher deployments times seven hours each). An MSAR aircraft over the ocean can do, in less than a normal workday, what it would take an inordinate number of inland grid searchers a week to do at a 24/7 operational intensity. Wandering off at two mph, this same hiker could create a theoretical search area of over 1,800 square miles in that half day, increasing the manpower requirement by a factor of sixteen.

This reinforces our conviction that it is inadvisable to encourage ISAR to adopt grid-searching as the primary technique for finding the living. While painstakingly thorough, it is both highly inefficient and virtually impossible to man at levels that will provide a quick rescue rather than the drawn out recovery of a cadaver.

Over the years, ISAR has developed a host of methods that emphasize efficiency over thoroughness to compensate for low searcher (sensor) speeds. Techniques like sign-cutting, binary sectioning, K9 scent detection, are part of an ISAR toolbox designed to use the least amount of manpower (and dog-power) in the most efficient way. Because speeds are slow and manpower is scarce, shifting to grid-searching within a closed search area, as the primary detection method, would be an unfortunate step backward.

Topic 3: POA vs. POD

Some SAR practitioners have shown a bias toward POD over POA. Their reasoning centers around the idea that the efficiency of the sensing resource, expressed as POD, takes precedence over the locational likelihood of the subject, expressed as POA, in driving up POS.

Table 2 below shows the true distribution of POA in a 4-segment closed search area. Regardless of any initial consensus or empirical distribution, the subject, marked by an "X" can only be in one of the segments with a true Probability of Area of 100%. All of the other segments have a true POA of zero.

Table 2: True POA in a 4-Segment Search Area

X 100%	0%
0%	0%

A poor resource with a low POD still has a chance of finding the subject if deployed in the segment where the subject is located. A perfect resource deployed in an area where the subject is not located has a zero chance of a successful find. Consider the following imaginary problem: a family member is missing and the Incident Commander asks for your preference. He can deploy resources to the Moon with 100% POD, or he can deploy a very poor resource into the area where the family member was last seen. Obviously, we would all intuitively prefer the second choice.

While the choice may seem obvious, certain search priority schemes create a bias toward the resource POD and away from the POA. Is it better to search in an area with high POA or with a resource that has high POD? To look at the problem from another angle, Table 3 shows an example of two POA x POD interactions that produce the same POS. Which would you choose? Assume that the same resource is used but with a different POD per segment.

Table 3: Mathematically Equivalent POS

Segment	POA	POD	POS
1	90%	10%	9%
2	10%	90%	9%

In the first example, the likelihood of the subject being in the segment is high, while the resource only has a 1-in-10 chance of making a find. In the second example, the probability that the subject is in the area is relatively low (given only the two segments), but the resource is highly efficient, expected to have a find nine out of ten times. Note that these two situations produce the same POS. Mathematically they are equivalent, but is one to be preferred over the other?

We would argue that the first case should be preferred. Although the resource is poor in Segment 1 and excellent in Segment 2, it should nonetheless be deployed into Segment 1. The rationale for this is twofold: First, the purpose of a search is to find the subject as quickly as possible, not sweep the entire search area as quickly as possible. These goals are not the same and may have different starting points. The subject is more likely to be found by any resource in Segment 1.

Secondly, the POD of the resource, whenever it is in the segment where the subject resides, is likely to be higher than its calibrated value. This is due to serendipitous elements of chance that work in the resource's favor: the resource may simply get lucky and stumble onto the subject; the subject may have left strong clues or signs, or the subject may be actively looking for rescuers.

For these reasons, we would discourage the use of search priority schemes that bias deployment toward areas of high resource POD but low POA.

Topic 4: TDA's: Tactical Decision Aids

Tactical Decision Aids are computer programs that assist with search management tasks related to probabilistic assessments and the associated Bayesian mathematics. The Coast Guard uses Computer-Assisted Search Planning (CASP), a proprietary application that starts by running a Monte Carlo simulation of 10,000 iterations to create an initial probability distribution for cells in a closed search area. The US Navy uses a proprietary multi-target package called NODESTAR, capable of generating up to a half-million 6-dimensional cells to account for characteristics like depth, velocity and target type. [7] These applications are not available for use by Inland SAR planners.

Inland SAR currently has a handful of TDA packages available. The two best known IBM-PC compatible applications are CASIE (Computer-Aided Search Information Exchange), which is free, and Search Manager which costs in the range of \$750 plus fees for additional seats. Of the two, CASIE is the simpler to use. Search Manager, which does a host of SAR planning and management tasks, including the ability to connect across an agency's entire network, is not typically run by a single user and is a complex system to learn.

The controversial nature of this topic lies in the fact that proponents of MSAR techniques in ISAR do not recommend (or even mention) either CASIE or Search Manager as viable TDA's. We suspect that this is because both include ROW in their POA distributions. MSAR proponents seem wedded to the idea that ISAR must revert to closed systems with no ROW, in spite of the fact that the use of ROW is a mathematically acceptable practice, if not widely implemented, in Operations Research.

Oddly enough, the impact of this is to take computing power out of the hands of ISAR planners, as MSAR proponents now recommend the retrograde use of paper forms for tracking and calculating POA updates. MSAR proponents have bemoaned the "computational burden" imposed by ISAR calculations, never overtly mentioning that CASIE has been around almost a quarter century to do them. As Alan Washburn notes, "CASIE3 is a user-friendly TDA that has been well received in the Inland SAR community, and which reliably shifts the POAs to account for past search effort." [8]

It is as though MSAR proponents are insisting that ISAR borrow its techniques but not the computing power required to properly implement them.

Topic 5: Mattson Consensus vs. Initial Distribution

In our view, the most misplaced idea from proponents of MSAR techniques is the equating of a Mattson Consensus, as developed in ISAR, with an Initial Distribution, as developed in MSAR. Similar to imposing on ISAR the grid search as a primary detection technique, using a Mattson Consensus improperly is a mathematical blunder.

Each formal search, where POA is going to be tracked and updated, must start with either an Initial Distribution or a Consensus. In maritime searches, the US Coast Guard employs the CASP software to create an initial probability distribution. This is done by combining a sample of empirical data, how lost vessels behaved under various environmental conditions, and using Monte Carlo analysis to create a set of probabilities based on the sample inputs. As previously noted, CASP runs 10,000 randomly seeded iterations to produce the probability distribution needed to create POA values for every cell in the closed search area. This is a rigorous method for creating POA assignments. Nothing like this currently exists in the Inland SAR world.

ISAR planners use a Mattson Consensus, named after its inventor, Colonel Robert Mattson, USAF, to create a pseudo probability distribution. Since a rigorous method for creating an initial distribution was not, and still is not, available to ISAR, Col. Mattson devised a weighting scheme that would rely upon the opinions of individual experts, people familiar with aspects of a search like local terrain and typical behavior of the category of lost person. There was never any pretension that these were true, empirically derived probabilities. But they acted as a starting point to avoid paralysis and get the search going. Without a Mattson Consensus, all segment probabilities would be equal, leaving no clear indication where to send resources first.

In MSAR's closed search areas, which are effectively devoid of clues (remember, no footprints left on the water), the initial distribution is a critical metric that often guides the search from start to finish. In ISAR, the initial consensus is just a qualitative estimate that can be shifted quickly in the presence of a strong clue. Indeed, we would assert that in ISAR, a strong clue is far more important than the initial consensus. After all, a Mattson Consensus is only the weighted average of a series of best guesses. Unlike a rigorously derived probability distribution, the quality of the experts' opinions cannot be controlled from search to search.

Proponents of MSAR techniques in ISAR mistakenly give a Mattson Consensus the same weight as a rigorously derived Initial Distribution. If a closed system is used, the initial cell or segment probabilities have to be locked down at the start of the search. In a clue rich environment like ISAR, this means that an entirely new consensus would have to be developed every time a strong clue was found within or without the search area.

Furthermore, the substitution of a Mattson Consensus for an MSAR-type Initial Distribution, provides a false sense of precision. A Mattson Consensus is neither empirically nor probabilistically derived and should not be used within a closed system. At the moment when a Mattson Consensus is developed, ISAR planners have, relative to the search's timeline, virtually the least amount of information about the lost subject's whereabouts. In MSAR, the Initial Distribution, based on the physical properties of a vessel and forecasts for the surrounding environment, is quantitatively derived and may constitute most or all of the pertinent information available to planners for the duration of the search.

We assert that the use of a qualitatively derived Mattson Consensus as a surrogate for a quantitatively derived Initial Probability Distribution is a serious flaw in the effort to standardize ISAR with MSAR techniques. The use of a closed system bestows a misleading level of accuracy on Mattson derived initial POA distributions, by equating them with mathematically rigorous ones.

Please address comments on this white paper to dano@newsar.org

REFERENCES

1. Toman, Richard J., (1999) Unpublished MLPI lecture notes, Brandon, VT.
2. Washburn, Alan R., (2002) Search and Detection, 4th ed., Topics in Operations Research Series, INFORMS, Linthicum, MD
3. Washburn (2002)
4. Stone, L.D., (1989) Theory of Optimal Search, Operations Research Society of America, Arlington, VA
5. Hill, Ken et al, (1997) Managing the Lost Person Incident (MLPI), Chapter 15, Probability of Detection, NASAR, Chantilly, VA
6. Appendix A, Calculations
7. Washburn (2002)

APPENDIX A

Grid Search Calculations

1. For determining the number of searchers required, 1-Square Mile

Assumptions:

SA = 1 square mile

Searcher Speed = 1 mile per 3.5 hours

Searcher Spacing = 20 feet (lateral range of 10 feet either side)

Duration = 7 hours (allows 1 mile up & 1 mile back)

Searchers = (Area sq mi x 5280 x 3.5) / (Spacing in Ft x Hours)

$$= (1 \text{ sq mi} \times 5280 \text{ ft/mi} \times 3.5 \text{ hrs/mi}) / (20 \text{ ft} \times 7 \text{ hrs})$$

$$= (18,480 / 140)$$

$$= 132$$

2. For determining the number of searchers required, 100-Square Miles

$$(132 \text{ searchers/1 sq mi}) * 100 \text{ sq mi} = 13,200 \text{ searchers}$$

3. For determining the number of searchers required Daily for a Week

$$13,200 \text{ searchers} / 7 \text{ days} = 1,886 \text{ searchers per day}$$

4. For determining the Theoretical Search Area at 2 mph for 12 hours

$$2 \text{ mph} * 12 \text{ hours} = 24 \text{ mile radius}$$

$$\pi * (24 \text{ mile radius Squared}) = 3.14 * 576 = 1,809 \text{ Square Miles}$$