Notes to accompany “Intro to Search Theory and its Application” presentation slides

The slides are not intended to be read as a stand-alone document. These slides were intended as a visual aid to be used in a live presentation at the NM SAR Council's annual ESCAPE event in April 2015. Because of this, most of the content of the talk was delivered verbally and is not present in the slides themselves. These notes are meant to provide someone who missed the talk but asked for the slides a very rough sense of how the presentation went, but won't be the whole experience.

Slide 1

“The fact is, most searches are over pretty quick.”

Slide 2

Graphic from Koester's Lost Person Behavior class: histogram of number of searches of given length from the 12,900 search incidents in the ISRID database.

Because such a huge fraction of searches are completed in less than a day (93%) it is important to make sure that early search efforts are not impeded by an overuse of formal search planning methods – these searches are best solved by rapid tasking methods.

But neither is it correct to dismiss formal search methods just because they are rarely used, and neither is it correct to neglect teaching formal search methods just because of previous experience that folks who learned them overapplied them when inappropriate. It is precisely because long searches are so rare and that these rare searches are the ones that are difficult that we really *must* learn formal search planning and be prepared to use them *when necessary.* And formal search planning means using search theory.

Slide 3

The class will start with a quick overview of how search theory was developed and how it was first introduced into the land SAR world.

We will quickly proceed to discuss the concept of “Objective” probability of detection (POD), to be distinct from “subjective” POD, what it means, what it depends on, and how to calculate it.

From there we'll discuss the “Probability of Success,” its use in formulating a theoretically optimal effort allocation, and the development of a practical search plan.

The class will then go on to describe how these abstract mathematical tools would be used in real application in search operations. Using search theory properly requires careful attention in briefing and debriefing resources.

We'll conclude with a quick review and statement of how we need to proceed to bring these methods into common knowledge and standard use.

Slide 4

Search theory, like most modern sciences, began with a military application: figuring out how to find things to blow up.

Navy wanted to find and destroy German submarines more efficiently. “Operations Evaluation Group” (OEG) established to apply “Operations Research” to problem.

Operations research is a branch of applied mathematics that applies advanced mathematical analysis to aid decision making. Used throughout military and commercial world to find better processes to optimize some desired output (blowing up submarines, producing right number of widgets and gizmos to maximize output, minimize cost and waste, etc.).

Example of OEG success: prediction that changing fuses of aerially-delivered depth charges so they detonate at 25 feet below surface instead of 100ft would increase kill rate. Kill rate went from 1% to 7% – because if you spot a sub and head towards it to drop charges and they spot you as you approach, they'll evade. If they have time to get down to 100ft, they also have time to move farther from the spot where you saw them – and your depth charge will be very unlikely to be close enough (within 20 feet) to the sub by the time it goes off for the charge do destroy. But if they only got down to 5-45 feet, they'll have had less time to evade and your charge is more likely to be dropped near them.

Similar successful prediction: painting search/destroy aircraft white will increase kill rate (because the aircraft will be harder to spot from the sub, meaning it can get closer before the sub is alerted and can evade).

Bernard Koopman is the founder of the operations research sub-field “search theory.” He developed a theory that calculates a “detection index” from which a probability of detection can be objectively calculated. The theory provides the means to aid resource allocation decisions: searches are always “big” problems, resources for a search are always insufficient to guarantee success, and allocation of the limited resources should be chosen to maximize the likelihood of success.

Original report was classified (“confidential”) information, but was declassified and published in 1958. A second edition, revised and extended, was published in 1980, and can still be purchased from the Military Operations Research Society. It is no page-turner, but someone with a solid mathematics background will find it an entertaining, if soporific, read.

Slide 5

Unfortunately, search theory is not really taught well (if at all) in most land SAR classes. There are a lot of reasons for this.

While search theory required a staff of mathematicians to develop, it doesn't require mathematicians to **use.** The work of the Operations Evaluation Group could not stay in the hands of academic mathematicians – it was necessary to develop procedures that could actually be used by the people planning and executing search operations. This required development of manuals, tables, charts and tactical decision aids that could be used quickly by people who were *not* mathematicians. And finally, people needed to be trained to use these tools – without requiring them to learn the deep background that allowed search theory to be developed.

In the 1970s, the state-of-the-art tool for land search was to throw together a large group of searchers and have them comb areas in very tight “line abreast” formations (walking practically shoulder-to-shoulder). This was extraordinarily inefficient and slow, and people began to look for a better way. In the process of this investigation, some people with familiarity with the Navy's (and Coast Guard's, by then) procedures brought the terminology and some of the concepts of search theory into the land SAR world.

The manuals and procedures had been developed for maritime search, and therefore all of the charts and tables were inapplicable to land SAR – and these people recognized that. Not being experts in search theory (but knowing how to use the end-result procedures), and not having the sort of budget the military had had to develop the charts and tables, these folks wound up leaving out large chunks of the math and procedures. This led to a logically incomplete system that could not work without filling in the gaps. Again, not being mathematicians and experts in search theory, their attempts to fill the gaps left a logically inconsistent (and sometimes flat-out *wrong*) system. Land search "theory" then developed on its own without ever making contact with real search theory again until the very late 1990s, at which time researchers in the search theory world became aware of the problems that had grown in land SAR. Only since then has there been an organized effort to bring land SAR back into the larger fold and get search theory taught (and applied) properly.

The quotations on this page are meant to get two additional points across:

- Search theory is actually a serious scientific field, not a realm of speculation and opinion. As such, it may be possible to make it simple enough to use, but it is also possible to oversimplify it to the point of meaninglessness.
- Most of the terminology and concepts I will refer to in the course of the next few hours are terms and concepts the audience has certainly heard of before – but heard of out of context or with "baggage" from outdated, oversimplified teaching. I ask that people let go of these associations during the class and keep an open mind.

Slide 6

Slide basically speaks for itself.

Slide 7

This slide is meant to clear up a misconception about the meaning of "POD".

"Success" in Koopman's quotation on the previous slide means "finding what you're looking for."

In order for the mathematical relation "POS=POC*POD" to be meaningful, the three probabilities **MUST** refer to exactly the same thing. So probability of success is the probability of finding the thing you're looking for. POC is the probability that the thing you're looking for is contained in the area where you're looking. And POD is the probability that the technique you're using *would* find the thing you're looking for if it were where you were looking.

This is meant to counter the complete misunderstanding that POD means "probability of finding the subject or small clues" or "fraction of 100 balls of white socks you might find" that are continously bandied about by instructors in land SAR. Unless "the subject or small clues" are the things that define success, and unless "the subject or small clues" are the things you've developed a probability of containment distribution for, POD can't be the probability of finding "the subject or small clues." All three need to refer to precisely the same thing or the math is ***meaningless.***

Slide 8

The big beef about how search theory was introduced into land SAR is that people who did so left big chunks of the background out, resulting in a logically incomplete and/or inconsistent system. In teaching this stuff I simply don't have time to provide all of the background for every concept – if I did, I'd be teaching a graduate level course in applied math, not a two hour familiarization class for search managers and searchers.

In some places in this talk I have to fill in some details just to make sure that things don't seem like they're pulled from nowhere. But I have to make sure at least someone in the room stays awake, so I have to make some leaps of faith.

As it turned out, there were not quite so many nerds (i.e. scientists, engineers, engineering students, etc.) in the room (that I could tell) and in the end I did NOT make very many side remarks about more advanced math that I could have made on the whiteboard.. In these notes, I've made these asides in footnotes, which can safely be ignored by anyone not interested in them – the rest of the talk can be understood without any reference to the asides in the footnotes.

Slide 9

A brief discussion of how information needs to be compiled into a probability distribution map in order to use search theory. The click-animated graphic (visible only in slide show mode) is used to illustrate the points.

- If we have absolutely no information on which to base searching decisions other than "the subject is somewhere in here" we have no choices to make. The very best solution to the problem is to spread search effort uniformly over the whole search area and hope for the best. In the graphic, the large greenish square represents a 25 square mile search area with absolutely nothing to distinguish any part from any other part. The optimum solution to this search is "send whatever you have to search this entire area uniformly in the time they have." This may have a very low chance of success, but the chance of succes of this plan is the highest we can hope for.
- We usually don't have as little information as that. We usually have at least the spot where the subject was last seen (or last known to be). That is, the IPP.
- We usually have at least some statistics about how far away from the IPP subjects like this are found.
- We often have some information about the area that helps us divide it up into regions of different likelihood – and in this case we have roads, trails, highways and streams that create natural boundaries that represent decision points for a lost person. Our assessment of the likelihood of any bounded region having the subject in it is based on our analysis of this set of information *in some way that I do not explain here.*
- I have created a set of probabilities for the region in this figure from whole cloth, for the sole purpose of having numbers with which to work in the subsequent slides. ("A complex mathematical procedure I refer to as 'rectal extraction.'"). Because the numbers themselves are not really important to the discussion here (though naturally they are critical in the real application), the final build of the slide shows the regions of the original map color coded by probability of containment, with brightest red representing the highest value, and lower values being represented by colors with less red and more blue as they get smaller.

- In all color-coded slides for the rest of the talk, whatever quantity is displayed is displayed with the same sort of color scheme, but red always indicates the highest value *of that quantity in that figure*. The scale is relative to this highest value in the figure, not an absolute scale.

Important things to note about this search problem that will be relevant later:

- Areas 0 and 1 are heavily wooded pinon/juniper forest heavily overgrown with oak – difficult to see or move through
- Areas 7 and 8 are ponderosa parkland through which a fire has run a few years ago, clearing all the understory but leaving the largest trees. Very easy to move through and very easy to see through.
- The hill in the center is steep with cliffs.

Slide 10

POC distributions are hard to get at (I referred to it as a black art while talking on this slide), and I do NOT discuss the methods of obtaining them in this class (that's a four-hour class of its own). But this slide tells a few of the ways that they can be developed. This is the last we see of the question – for the rest of the talk it is simply assumed that we have obtained this POC map in some way other than picking random numbers.

Slide 11

Probability of detection is an example of a "conditional probability."[1] It refers to the probability that a detection would be made *assuming* that the search subject is contained in the area searched. As an example, I stated that if I were to poke my head in the door of the adjacent meeting room, it was reasonable to say that there would be a 99.999% probability that I would detect an elephant there if there were an elephant in the room. This is a statement about how observant I would be and how easy it would be to spot an elephant that was present. This is very different from the answer to the question "If you looked in room 3, what is the probability that you would find an elephant?" (i.e. not "assuming there's one in there") That would be very low, because it is very unlikely that there is an elephant in there.

Probability of detection depends, naturally, on how "detectable" a search object is by the "sensor" looking for it, and on how much effort is expended doing the looking. "Objective" POD refers to *computing* the probability of detection from a quantified "detectability" and "effort" in some means, as opposed to subjectively estimating it. To do this requires that we quantify both "detectability" and "effort."

1 In SAR, we refer to this thing as POD, which notation obscures its conditional nature. A mathematician would use a more precise notation to keep things clear. The conditional probability of detection assuming containment would be written as $P(D|C)$ which reads "The probability of D given C" where D is detection and C is containment. The probability of success is actually "the probability of containment and detection", $P(C\cap D)$, and so the "fundamental relation" of search probabilities in a more formal notation would be:

$$p(C\cap D)=P(C)P(D|C)$$

which reads as "The probability of C and D is equal to the probability of C times the probability of D given C." Depending on which source you read, this is actually the *definition* of conditional probability in terms of the joint probability, or the definition of the joint probability in terms of the conditional probability.

Slide 12

Mimicking the search theory literature, we start the subject of quantifying “detectability” by considering a “definite detector. This detector always finds a particular type of search object, so long as that object ever comes within a distance “W/2” on either side of the track along which it moves. An object that never comes within that distance is not detected. The number “W” is called the “Sweep Width”, and the distance of a search object from the sensor track is called the “Lateral Range.”

If you graph the probability of finding an object at a given lateral range, you create what is called a “detection profile” or “lateral range curve.” For a perfect detector, this curve looks like the graph on this slide. There is a region of width W where the probability is 100%, and outside that region the probability is zero.

We now want to know the probability of detection that would result from moving this sensor through some search area.

Slide 13

In this slide we describe moving a definite detector through a perfect “parallel track line search” pattern.[2] As the detector completes each pass, it carves out a region where it will certainly detect the search object if it were there. When the entire search of the area is complete, the sensor has moved a “total track length” of L. This track length is also referred to as the “Effort.” Effort may be expressed either as distance travelled or time spent (related, obviously, by the sensor's speed).

The total area over which the sensor could detect objects is its sweep width times its track length, the “Area Effectively Swept,” represented as Z.

The ratio of the area effectively swept to the total area of the search region is called the “Coverage.”

Slide 14

For the definite detector applied in this manner, the probability of detecting an object that is in the area is precisely equal to the probability that the object is ever within the sweep width of the sensor. In this case, that probability is exactly equal to the coverage. A graph of POD vs. coverage rises linearly from 0 to 1 as coverge changes from 0 to 1, and remains at 1 (100% POD) for all coverages above 1, meaning that getting coverge higher than 1 with this type of sensor in this type of search is a waste of time.

This represents the best possible case, using a perfect detector in a perfect pattern. Any other detector would produce a different graph. And even the perfect detector, applied imperfectly, would produce a different graph.

Slide 15

While seemingly abstract and useless, the analysis of the definite detector leads to some interesting and useful consequences.

Imagine the sensor flying through a “swarm” of identical search objects that have been distributed uniformly at some density of objects per unit area (e.g. 6400 objects/square mile or equivalently 10 per acre). All objects that fall inside the detector's sweep width anywhere on its track will be found. The

2 A huge “Thank you” to Tony Davis of White Mountain SAR for my new $10 word, “boustrophedonic,” which describes the parallel arrangement of search tracks here, and is derived from Greek words meaning, roughly, “as an ox plows a field.”

number of those objects must be the density times the area effectively swept. Rearranging gives us a relation between the sweep width and the rate of object detection, sensor speed, and object density. It is also the case just by looking at the lateral range curve that the sweep width is the area under the curve. These two things turn out to be generally true for *any* detector.[3]

Slide 16

This slide introduces the notion of the definite detector being applied to a search area in a completely random manner instead of a perfectly ordered manner. The animation builds up a simple argument for why the "takeaway message" is true, but relies on the use of limits and the definition of the exponential function. The build is strictly there just in case someone in the audience is a nerd. The answer appears first in case nobody wants to grok the derivation.

Slide 17

This slide is pretty self-explanatory.

Slide 18

The two "definitions" of effective sweep width here are actually not independent. In formal search theory texts, W is actually defined as the area under the lateral range curve. The other is strictly derivable from that definition.[4]

Even though random search law derived using the definite detector, turns out to be applicable to any detector whenever the detector's application is subject to human fallibility and other random influences.

Reber (1958) actually showed that even a definite detector in a parallel track search will tend to produce PODs more like the random search law if there are navigation errors that make the tracks anything other than perfectly spaced. The larger the navigation error standard deviation is relative to the sweep width (e.g. 3*W, 4*W, etc.) the closer the POD is to the random search law, and the farther it is from the definite range law.

None of which helps us if we have no means of obtaining the effective sweep width.

Slide 19

This slide is a "leap of faith" slide. The result here cannot be derived without calculus.[5]

---

3 Well, technically, one of them is used as the *definition* of the "effective sweep width," and the other is a straight consequence of that definition.

4 If p(x) is the lateral range function that represents the probability of finding an object at lateral range x, then there are $\rho v t \, dx$ objects at that lateral range, and from the definition of the expectation value of a random variable $\langle N \rangle = \int_{-\infty}^{+\infty} \rho v t \, p(x) \, dx = \rho v t \int_{-\infty}^{+\infty} p(x) \, dx$. The intergral is just the area under the curve, so it is W by the normal definition. A simple rearrangement gets the relationship on line 3 of the slide.

5 But is trivial *with* calculus. Continuing the argument of the previous footnote:

$$\langle N_{missed\ inside} \rangle = \int_{-W/2}^{W/2} \rho v t (1 - p(x)) \, dx \quad \text{and} \quad \langle N_{found\ outside} \rangle = \int_{-\infty}^{-W/2} \rho v t \, p(x) \, dx + \int_{W/2}^{\infty} \rho v t \, p(x) \, dx \quad .$$

Anyone who has gotten this far through this footnote without a vague feeling of panic will have no trouble showing the equivalence of these two expressions in a line or two of rearrangement.

Slide 20

This slide shows the results of carrying out an experiment of the form on the previous slide. Blue curve is "detections at this distance or further" and red curve is "misses at this distance or closer." Where the two curves cross is the lateral range corresponding to half the sweep width – the distance at which the number of objects missed closer than that distance equals the number of objects found farther than that distance.

Doing this experiment with a large number of searchers (pooling data into a single crossover graph) gives you one sweep width value for one particular object type in one type of place at one particular time of year in one set of environmental conditions. To generate a complete table to be used to look up sweep widths would be an enormous effort. This needs to happen, but will take many years, or an enormous grant and dedicated researchers with a huge pool of testers. In the meantime, we need a shortcut.

Slide 21

This slide describes the shortcut. It introduces the concept of the "average range of detection" and how to measure it. An object of the same kind as the search subject is placed in the environment of the search area, and a searcher walks away from it until it can't be seen. The object is then approached again and the distance at which it becomes visible is recorded. This process is repeated 8 times around the object and the distances averaged.

There is absolutely no reason to believe without verifying that this number has any relationship to the sweep width: it is the distance at which an object – known to be present and whose location is known to the searcher – can be seen. The sweep width is a measure of how well a given detector can detect objects without being alerted to their presence.

But in the cited paper, the researchers actually looked to see if there was a correlation between range of detection and sweep width for that type of searcher and object, and found a statistically significant correlation. Their work shows that if real sweep width values are not available, the range of detection can be used if a correction factor is applied. The slide gives the correction factors proposed in the paper. The QR code or link get you to the text of the paper itself.

Now we have something we can work with.

Slide 22

Recap provided mainly to break up the talk into digestible segments and set up the next section.

Slide 23

High POD is not itself an important thing. High POD must be achieved *in the right place.*

Slide 24

POD must be paid for. The slide shows a demonstration calculation of how much a given amount of POD "costs" in one particular combination of conditions. The point is that the cost grows very quickly at the high end. We have to determine the appropriate amount of effort to expend, because we don't have an infinite supply.

Slide 25

The real measure of quality of a search plan is its probability of success. Through the objective POD and the POC distribution, we have a knob we can turn in each region (how much effort we allocate there) and a payoff function (the POS). We can use this function to find the best allocation of effort.

Slide 26

It is often said that searching helps us "rule out areas where the subject isn't." But searching doesn't really do that, because searching isn't guaranteed to work – there is only a *probability* of detecting something even if it's there. What we are really doing by searching is providing evidence that is either consistent or inconsistent with our original hypothesis (the original POC distribution, the "prior" probabilities). Through Bayesian updating we can fold the results of our search (the "evidence", which presumably is "we did not find the subject") into our knowledge base and update our "state of belief" (the POC distribution). The result is a new POC distribution known as the "posterior" probabilities.[6]

Slide 27

This just shows the results of applying the formula of the previous slide to one of the areas of the original toy problem.

Slide 28

Mostly self-explanatory, but contradicts one of the common misconceptions perpetuated by many land SAR books and classes: POS does indeed have value after a search, and is the KEY quantity we aim to maximize with each search, whereas it is often taught that the quantity is useless.

Slide 29

Sets up the main resource allocation example used for the rest of the class. An extremely limited resource (200 searcher-hours) needs to be distributed in the 25 sq-mi area from the POC slides. What is the best plan?

Hint: It is ***NOT*** "search the highest POC areas first."

Slide 30

Region 0 is not the place to search first because while it does have a 15% POC, it is also enormous – so the probability per unit area is actually quite small. If consider probability *density* instead, the hot

---

6 These are called Bayesian updates because they are applications of Bayes' Theorem:

$P(C|D')=\frac{P(C)P(D'|C)}{P(D')}$. In words, "The probability of containment given non-detection is equal to the prior probability of containment times the probability of non-detection given containment, divided by the probability of non-detection. If one works it out one gets the formula for the normalized posterior POC. However, since all the formulae that are used in effort allocation only use *ratios* of area POCs, the common denominator P(D') is not really helpful. Instead, we can look at the *joint* probability of "containment and non-detection":

$P(C\cap D')=P(C|D')P(D')=P(C)P(D'|C)$. Making the associations that $POC_{new}=P(C\cap D')$, $POC_{old}=P(C)$ and $(1-POD)=P(D'|C)$ leads to the update formula on this slide. Then the total non-detection probability P(D') never appears, and the joint probabilities are used as the new POCs. Their ratios remain correct, and all the algorithms for effort allocation work.

spots change significantly. But still can't allocate, because not all areas are equally searchable.

Slide 31

Now we put numbers into the regions representing their sweep widths and searcher speeds, which are intended to reflect the vegetation and terrain differences described on slide 9.

Slide 32

Introduces the probable success rate, PSR. Mathematically, this is the derivative of POS with respect to time, evaluated at zero time. Because it depends on POC, doing search in an area reduces PSR just as it reduces POC.

The optimum way to apply an increment of search effort is to put it in the area of highest PSR. But putting a small increment there decreases the PSR, and eventually that same region is no longer highest, requiring distribution of effort to a different area. This is how the algorithms that exist ultimately find an optimum.

Slide 33

This slide simply shows the same areas, but color-coded by PSR value instead. The hot spots are very different than early intuition would have suggested.

Slide 34

Describes one algorithm for finding an optimum. This algorithm only applies if there is a single pool of identical resources to draw from (i.e. all searchers are ground troops). If that is not true (i.e. have a pool of ground searchers, a pool of ATVs, a pool of 4x4, a pool of mounted searchers, etc.) a different algorithm is required.

Charles Twardy of George Mason University had a package called SORAL that could do these allocation computations (written as part of the SARBayes project when he was at Monash University). This library is hard to find now but some of us still have copies – including Don Ferguson, who is reimplementing the algorithms in Python so they can be applied by IGT4SAR.

Slide 35

Another view of how Charnes-Cooper works. At each step of the algorithm, sufficient effort is applied to the area of highest PSR to reduce the PSR down to the level of the next highest. As areas are made equal, the algorithm assigns effort to all of the equivalent areas to bring them all down to the next level. If ever all areas are equal, the remaining effort is applied to all areas to keep them equal.

Slide 36

The results of running Charnes-Cooper on the toy problem. The overall POS is only 27.6%, we cannot do any better than this. But this allocation is not practical: it has some areas with ridiculously small allocations that cannot possibly be performed in the real world. This allocation needs to be “tweaked” into an operationally viable plan. The tweaked plan will have a lower POS than the optimum, but if done carefully the reduction will not be very significant.

Slide 37

This slide shows the allocation and post-search PSR distribution obtained by a tweaked allocation. The allocation is obtained from the optimum by removing all allocation to areas that got less than 0.5 coverage, and reallocating the effort to areas that got more than 0.5 coverage. The redistribution is done so that each area receives an equal increase in coverage.

This allocation does have lower POS than the optimum, but it is not much worse. 26.1% is close enough to 27.6%, especially if the difference is due to making the plan something we can actually implement.

Slides 38-40

Self-explanatory.

Slides 41-44

Mostly self-explanatory, but walk through how the methods presented so far are applied to subsequent search periods (i.e., pretty much the same as for the first period, but using the updated POC map for all calculations). The overall cumulative POS is just the sum of the POS values of all searches to date.

Slide 45

This slide present mainly because a lot of energy is expended in NM search planning training talking about the cumulative POD. This is basically an extension of the overemphasis on POD that pervades other discussions. I spend a little time saying that.

If we have been keeping track of updated POCs as we should have been, obtaining cumulative POD is very easy, and can be done in a single step with the formula on this slide. There is no need to resort to the much more tedious cascading formula usually taught in SAR classes. The reason that formula is taught is because it was also taught not to bother tracking updated POCs! Either way, though, the two methods are equivalent. This one's just simpler.

But even if there is an overemphasis on its value, cumulative POD does have its use. If information comes along that we wish we had had access to at the time we generated our initial POC distribution, we do have a use for cumulative POD. We compute the cumulative POD of each region usng the formula on this page. Then, we go back and generate a new initial POC map from our initial information, but including the new information *without considering any of the searching we've already done.* The total effect of all searching done to date can be taken into account by performing a single Bayesian update using the cumulative POD as if it were a single search.

If the new information indicates that the subject has been moving, however, our prior searching is pretty much rendered meaningless – the subject may have moved into an area we have already searched. If that is the case, we cannot apply the procedure of the last paragraph, we just have to start again.

Slide 46

Shows the optimum allocation for the second wave of 200 searcher-hours, assuming we have applied our first 200 searcher-hours according to the theoretically optimum allocation. Day 2's optimum is just as impractical as day 1's, but in a different way: now, all the areas we searched yesterday are getting the same low POD, because they all came out of yesterday with the same PSR --- and that PSR was still the highest. Now we have a huge area of equivalent regions, and must distribute coverage uniformly to keep accessing the peak.

Slide 47

Shows the result of applying the “operationally viable” plan for day 2 assuming we used the “operationally viable” plan on day 1. That is, we take the POCs from the day 1 “viable” plan, compute the theoretical optimum, and again reallocate effort from any area with low (<0.5) coverage . It is very viable, and is very different from the theoretical optimum of slide 46. That's because day 1's effort did not make all the high PSR regions equal – the ones we searched yesterday are now much lower PSR than the ones we skipped. The day 2 allocation therefore ignores them and puts effort in to what were high PSR regions yesterday that didn't get searched.

Even though neither day's search was the theoretical optimum, the reduction in cumulative overall POS is not very large – and both days are now allocations that actually make sense.

“Search theory doesn't tell us what to do. It gives us information that helps aid our decision making.”

Slides 48-53

These are mostly self-explanatory, but try to show how all of this stuff would flow in an incident. We need to collect information to use in planning. Since this formal planning is not part of the first few operational periods, the time to collect that information is actually IN the first few periods. We might never use it if the search doesn't go long, but if it does we'll have it. If it doesn't, we can save that kind of information (in a GIS layer, even?) for use in the future in some other search in this area.

To make sure we get this information, we need to be briefing resources to collect it! And we need to debrief resources to get it out of them.

Slides 54-56

Subjective POD is a vessel of fertilizer, and none can abide the odor thereof. Knock it off.

Slide 55

“Critical separation” is the theory that was proposed to compute POD from the ratio of searcher spacing to range-of-detection. The formula proposed was pretty much just postulated and was not done with proper use of search theory. That formula matches up with observations that doing a search spaced out at twice the range of detection gives a 50% POD, but does not match up with our “worst case” random search law at any other value. It gives nonsensical results like “pod is zero at coverage of 0.5”, and underpredicts POD at higher coverages. Critical separation is an idea whose time should never have come, and whose time has long passed.

Further, it is often said that it is better to perform “critical separation” in conjunction with “purposeful wandering,” and that better POD is achieved this way than with parallel tracks at critical separation. Search theory could have told you that – purposeful wandering means lengthening the total track, which means applying more effort, which means increasing POD (and of course taking longer). But whereas the “critical separation with purposeful wandering” thing just says “you'll get higher POD,” search theory tells us the actual improvement numerically.

Search theory with sweep width, effort, and objective POD from the random search law is the way to be working.

Slide 56

Segmentation doesn't belong in this section of “Myths” but there is a lot of talk in planning section chief classes that imply that segmentation is an important early step in the formal planning process. “Segments” are supposed to be small areas that can be searched “in 4-6 hours.” But that is an *operational* consideration that involves *how searchers can do a job.* For *planning* purposes we need to be dividing the search area into regions based solely on *what the subject would have done.*

For purposes of POC determination, search area needs to be divided into *regions* of distinct probability of containment. Areas that have no boundaries that represent a reason for POC to be different should not be divided up.

Segmentation needs to happen when tactical assignments are being chosen in order to achieve the PODs determined in the planning process.

And there are some articles out there suggesting that there are ways to get around segmenting even at the operational stage, but these are still up in the air and require very careful record keeping *in a GIS.*

Slides 57-61

These are self-explanatory.