Open Peer Review on Qeios

Notes on the Implications of Ignoring Bayes' Rule in Search and Rescue Practice in the UK

Andre Clark[1]

1 University of South Wales

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.

Abstract

Thomas Bayes (1702-1761) has had a big influence on the science of inference since he discovered the mathematically correct way of adjusting probabilities to account for new evidence. Nonetheless, it is still the case that in practice it is not always clear where and when to apply the rule he derived, or the consequences of not doing so. In this note, the effects of not doing so when searching an area of ground for a missing person (misper), where the chances of finding them depends both on whether they are there and how well the ground is searched, is investigated. This investigation suggests that within the range of probabilities that generally apply to search operations in rural settings in the UK, the widespread failure to apply Bayes' rule may incline search managers to widen search areas more than is warranted by the evidence and may thereby reduce overall search effectiveness (ceteris paribus).

Andre Clark (andre.clark@southwales.ac.uk)

The author is an active member of a Mountain Rescue team based in Wales.

Keywords: Search & Rescue, Bayes, Search Theory, Missing Person Behaviour.

Discussion

When dealing with conditional probabilities, where the probability of an event or outcome occurring is based on the occurrence of a previous event or outcome, applying Bayes' rule (theorem) gives us the correct numerical result. The rule is that for events A and B the Pr(A|B) = Pr(B|A) Pr(A)/Pr(B), where Pr(A) is the probability of A occurring, Pr(B) the probability of B occurring, Pr(A | B), the probability of A given B, and Pr(B | A), the probability of B given A, and where / means division, and adjacent variables means multiplication (although the letter x is also used for this in more textual parts of the following discussion).

At an intuitive level we know that this rule must apply to the case of finding someone in a field since this depends both on how well the area is searched (the Probability of Detection, POD) and the likelihood of the misper having walked into the field at some point in the past (the Probability of their being in the Area, POA). In SAR practice in the UK, however, what we do is to simply multiply POA and POD together (Cooper, 2022, Charnes and Cooper, 1958), to determine the probability of finding a misper (Probability of Success, POS). This calculation is standard practice whether it is done manually, or using search management programmes such as ‘SarMan,’ and in this note we are simply comparing the results of doing it this way compared to using the alternative Bayesian calculation in order to check if there are any practical implications arising from this. That it could have implications rather depends on what is done with the POS information, but typically search managers, in considering which segment of a search area to task a search party to, will look for the highest POA in deciding where to send them. If for example a POA of 20% is attached to a segment of land in the West and 10% to one in the East, then the party will be sent West regardless of whether Bayes’ rule is applied or not. The difference between the traditional approach and the Bayesian alternative comes when a search party subsequently returns to ‘control’ (usually a control vehicle from which the search manager is organising the search) having not found the misper. If, for example, a party is believed to have covered the segment with a POD of 50% the search manager has a difficult decision to make since this means that in effect only half the segment was searched and half of 20% is 10%. This means that the same probability now applies to sending the party East and sending them back to the West since in effect there is a remaining POA of 10%, equal to the previous POA x (1-POD). If, instead, we apply Bayes’ rule to the estimation of these new POAs we may get a different result and quite possibly a different plan, which is better in theory, but which might also be better in practice -although this depends on a number of other factors, some of which are addressed in what follows.

In order to maintain some consistency with standard notation (as suggested by Taylor 2024, in response to an earlier draft of this paper), if we define POA as Pr(A) and POD as Pr(D) and use ‘not-A’ and ‘not-D’ to denote their counterparts, then having searched a segment of ground ‘i’ the conditional probability of the misper still being there, undetected, is Pr(A_i|not-D_i), which following Bayes Rule is;

Pr(A_i|not-D_i) = Pr(not-D_i|A_i) Pr(A_i) / Pr(not-D_i)

= [1 - Pr(not-D_i|A_i)] Pr(A_i) / [Pr(not-D_i|A_i) Pr(A_i) + Pr(not-D_i|not-A_i) Pr(not-A_i)]

Given that if someone is not in segment ‘i’ then they can't be found there Pr(not-D_i|not-A_i) = 1, this can be simplified to;

Pr(A_i|not-D_i) = [1 - Pr(not-D_i|A_i)] Pr(A_i) / [Pr(not-D_i|A_i) Pr(A_i) + Pr(not-A_i)]

and further to;

Pr(A_i|not-D_i) = [1 - Pr(not-D_i|A_i)] Pr(A_i) / [(1 - Pr(D_i|A_i)) Pr(A_i) + 1 - Pr(A_i)]

and ultimately to;

Pr(A_i|not-D_i) = [1 - Pr(not-D_i|A_i)] Pr(A_i) / [1 - Pr(D_i|A_i) Pr(A_i)]

In terms of activities for a search manager this means calculating a new POA for the searched area segment ‘i’ in time ‘t’ that takes the POA prior to the search in t-1 to give a new POA equal to POA_{t-1} x ((1-POD)/(1-POA_{t-1} x POD)). The chances of the misper still being in the search segment despite our efforts to find them will go down in line with the formula, while the POAs for other segments will go up, with new POAs = POA_{t-1} x (1/(1-POA_{t-1} x POD)), although it is not entirely clear how best to operationalise this, (as discussed in more detail below). The size of the adjustment in the POAs matters since repeatedly allocating returning search parties according to the highest new POAs will enable a search manager to maximise the chances of finding the misper given the resources available (Colwell, 1994), both by shifting resources between segments in a given search area and by comparing the new POAs to those that have yet to be constructed for segments outside the current search area. On the other hand, in practice in the UK one repetition tends to be the maximum and this is therefore less of a problem in reality than talk of repeated allocations would suggest.

In addition to the issue of what maths to apply, there are a couple of additional practical challenges to directing a search through adjusted POAs in the manner described above. One is that the aforementioned construction of original POAs are a bit of a guess in the first place: Following Mattson (1976), search managers have to combine a lot of different types of information together. There is some widely available statistics to help them with this in the UK, (such as can be found in Perkins et al, 2011), but they will also be expected to combine this with a lot of unquantifiable information about the misper's habits provided by police, friends, and family, as well as all manner of local geographic knowledge from search team members, (since in the UK Mountain Rescue and Lowland Search teams are made up of local volunteers). A second problem is that there is disagreement on how exactly any POA adjustment should be done, and although one might assume that in a paper about Bayes' rule the author would recommend following Syrotuck (1975) and apply Bayes' rule every time, to do so would entail overlooking some significant practical difficulties. The problem is that to correctly maintain the relativities between each new POA (given that it is the relativities that will determine which areas to prioritise), every POA segment need to be adjusted each time a search party returns to control and although modern technology means this is not the burden it once was, it is still the case that search manager may struggle to keep their intuitive feel for numbers that are continually moving about. In an attempt to help, some have suggested ensuring that all the POAs are always constrained to sum to one, but such normalisation is not without its problems (Cooper et al, 2003) as although this stops scales moving along with the numbers themselves, it can still leave search managers scratching their heads over what the numbers are telling them, regardless of the fact that they sum to the same number each time (Cooper, 2000). Rules of thumb could be developed to simplify matters of course, one obvious one is to simply apply rank orderings, but decisions on whether to re-search a segment are not made lightly, particularly as this implies some sort of failing on the first search undertaken, and some means of retaining quantification of scales is likely to be needed.

Little wonder then that rigorous POA adjusting is seldom done in practice and the simple alternative of keeping PODs high (70% or more) has become the norm (on the assumption that this will ensure that each searched segment falls below the POA of any newly created segment), with searches therefore tending to evolve by expansion rather than by re-working. Consequently, in what follows the focus is on the first round results and the question of whether to repeat or expand the search at that point, and the issue of what happens in subsequent iterations when POAs (with the exception of any vague unsearchable catchall Rest of the World segment) are likely to be lower, is ignored. This is consistent with the author's own experience of over 20 years in Mountain Rescue in the UK, in which frequent repetitions have seldom been observed, except occasionally where a crime is indicated.

Before attempting to illustrate the effect of adopting Bayes' rule in such circumstances, it is worth mentioning that although calculating PODs is not the focus here, this is also a problematic area: Often PODs are calculated by a search party simply walking towards, or tangentially, to an object such as a rucksack to see at what distance they gain sight of it. This ‘rain dance’ indicates how far apart (on average) searchers should walk when searching, and although the most diligent parties will adjust such ‘sweep widths’ as their estimates of how much ground they can effectively search is hit by changes in terrain, failing light, tiredness, and so on, it is inevitably a bit of a guess, (although ways to improve it are frequently being suggested, see Koester et al, 2014 for example).

An illustration of the effect of adopting Bayes’ rule.

Although many of the assumptions regarding the relationship between searching and finding can be traced back to the need to find things in the relatively undifferentiated marine environment (Koopman, 1946), landscapes add a lot of complexity and consequently in this illustrative scenario a simple location with very few search segments (A1-A5 superimposed on the map below) is used. The point at which the missing person was last seen (PLS) was just to the east of the Parking sign shown on the map where she parked her car since the official car park was full of troops (who soon departed). Given that a car park is the ideal place to locate the control vehicle and to assemble search parties, this became the designated central initial planning point (IPP) for the search.

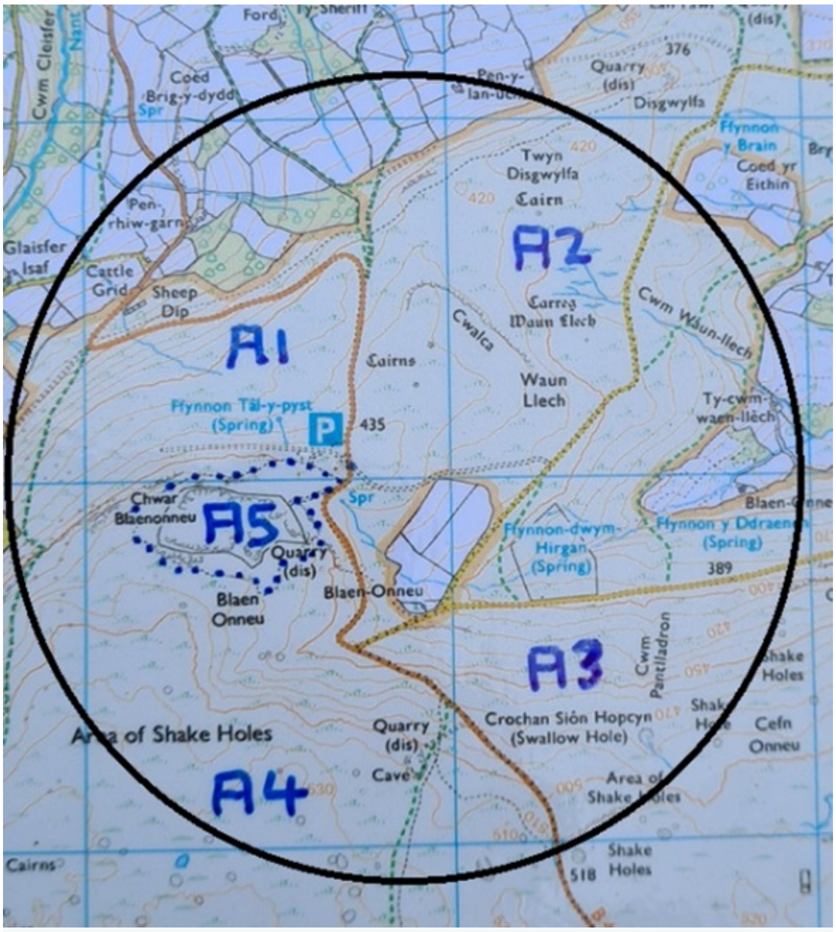

Figure 1. Map showing 4 roughly equal segments and a smaller, more complex, A5 segment.

In the absence of additional information, such as a route favoured by the misper, it was decided to conduct searches of the 5 segments A1-A5, which lie, more or less, within the circle drawn on the map. Such circles are often included (either by hand or by programmes such as SarMan), to represent data collected on missing person behaviour over the years (Syrotuck 1976, Koester 2008). In this case the circle represent the distance from the PLS in which 50% of 'solo female walkers' are found according to Mountain Rescue England and Wales (Perkins et al 2011). With modern GPS and map-enabled mobile phones it is possible for searchers to actually trace out the 1km square grids delineated by the blue lines, giving four equal squares for A1 to A4 at least (which if the terrains are similar might result in similar search timings), or to actually skirt the edge of the circle (which is neat in terms of ticking off similar POAs across segments). In practice, however, neatness on a map is not a priority and generally search managers will look for features on the ground that search party members can work within, such as field boundaries, to minimise the amount of time wasted checking proximity to circles or squares. Modern technology makes this easier than it has ever been before, with apps such as SarMan enabling search managers in control to draw the complex shapes that natural boundaries tend to take on a computer that within seconds appear on the mobile phones of each and every member of a search party.

In theory a search manager would follow the rules of 'constrained optimisation' and maximise the overall POS given the resource constraints (Koopman 1980, Stone, 2007), but in addition to the practical problems mentioned previously with such maximisation they may also be unclear on what the constraints are; how many searchers will turn up, how quickly they will cover the ground, and whether this is quick enough to find the casualty alive (given that the likelihood of survival tends to go down with time, as Adams et al, (2007) show). This is not to imply zero progress on the searching side, with the advent of such things as 'watershed' maps that help to show the most likely routes mispers take when confronted by gullies and rivers (Doke, 2012), and of course the use of drones in the searching process.

In order to keep it simple in this illustration such details are omitted and the only complicating factors that stop us from simply dividing the area into four segments is that one of the areas has a quarry which requires more effort to search because its relative topographical complexity narrows the required sweep widths to get a party into the desired 70% POD range. In effect, adding this fifth area allows us to allocate resources unevenly to match the unevenness of the 'probability densities' on the ground, while keeping each search party group to a sensible number of 4 personnel, given that on the day 20 searchers turned up.

Results

Although the principle is unaffected, by effectively eliminating the differences between segments this illustration allows us to look at the overall POA alone, which, with the overall original POA of 50% and POD of 70%, gives us a new POA figure of 15% using the traditional multiplication approach, and 23% using Bayes' rule. The 8% difference between the two is shown in bold alongside several others for a range of POAs and PODs in the table below:

Table 1. The extent to which Bayes' rule exceeds traditional approach in indicating need to re-do a previously searched area.

	POD					
POA	0.4	0.5	0.6	0.7	0.8	0.9
0.3	2%	3%	3%	2%	2%	1%
0.5	8%	8%	9%	**8%**	7%	4%
0.7	16%	19%	20%	20%	18%	12%
0.9	30%	37%	42%	46%	46%	38%

These results show that in considering where to search next, the likelihood of the misper being outside the search area is overstated by the traditional approach. A search manager may consequently find themselves a little more inclined (depending on what both the POAs and PODs are) to expand the search rather than re-do previously searched segments than is warranted by the data, unless we are looking at small POAs to begin with (as we may be in looking at segments), in which case the difference is small.

References

- Adams, A. L., Schmidt, T. A., Newgard, C. D., Federiuk, C. S., Christie, M., Scorvo, S. and DeFreest, M. (2007). Search is a time-critical event: When search and rescue missions may becomefutile. *Wilderness & Environmental Medicine.* 18: pp. 95-101.
- Charnes, A and Cooper, W. W. (1958) The Theory of Search: Optimum Distribution of Search Effort.*Management Science* 5(1): pp. 44-50.
- Colwell, M. (1994). *Search priority: an integrated approach to search planning* British Columbia, Canada: Private publication.
- Cooper, D. C. (2000). *The application of search theory to land search: The adjustment of probability of area*Private Publication, Cuyahoga Falls, Ohio, U.S.A.
- Cooper D. C. (2022). *Search Actions Outline; a guide for planning and managing a land search* Mountain Rescue England and Wales.
- Cooper, D. C., Frost J.R and Quincy Robe, R. (2003).*Compatibility of Land SAR Procedures with Search Theory.* Prepared for U.S. Department of Homeland Security United States Coast Guard Operations (G-OPR) Washington, D.C.
- Doke, J. (2012). *Analysis of search incidents and lost person behavior in Yosemite National Park.*Doctoral dissertation, University of Kansas.
- Koester, R. J. (2008). *Lost Person Behavior: A Search and Rescue Guide on Where to Look—for Land, Air, and Water.* Charlottesville, VA: dbS Productions.

- Koester, J., Chiacchia, K. B., Twardy, C. R., Cooper, D. C., Frost, J. R., and Robe, R. Q. (2014). Use of the visual range of detection to estimate effective sweep width for land search and rescue based on 10 detection experiments in North America. *Wilderness Environmental Medicine*. 25(2): pp. 132-147.
- Koopman, B. (1946). Search and screening. *Technical report OEG report No 56*. Operation Evaluation Group Office of the Chief of Naval Operations. US Navy.
- Mattson, R. J. (1976). Establishing search areas. *Search and Rescue Magazine*, Spring, 7-8.
- Perkins, D. Roberts, P and Fenney, G. (2011). *The U.K. Missing Person Behaviour Study* Mountain Rescue (England & Wales).
- Syrotuck, W. G. (1975). *An Introduction to Land Search Probabilities and Calculations*. Arner Publications, Westmoreland, New York, USA.
- Taylor, D. (2024). Comment on: Clark, A. Notes on the implications of ignoring Bayes' rule in Search and Rescue practice in the UK. *Qeios*. Comment posted on Jan 19, 2024.