# Training Analysts in the Proper Use of Ignorance in Intelligence

**Martin Hill, John Salt**

**Center for Electronic Warfare, Information & Cyber**

**Cranfield University**

**Defence Academy of the United Kingdom**

**Shrivenham**

**SN6 8LA**

## Abstract

The right information must be balanced by the right ignorance to avoid overloading intelligence work. Analysts are therefore trained in mechanical simplifications - such as hatching out areas on maps - that are appropriate to particular circumstances but fail catastrophically in others. The solution is not to throw out the simplifications but to train in awareness of their limitations and the ability to adapt as required.

**Keywords: Ignorance, Intelligence Assessments, Knowledge**

## Introduction

Considerable effort has been made to understand and convey uncertainty in intelligence analysis since at least Kent's 'Words of Estimative Probability' [1]. More recent guides such as Heuer's "Structured Analytic Techniques", the DIA's "Quick Wins for Busy Analysts" and the Canadian Forces Intelligence Command's "Aide Memoire on Intelligence Tradecraft" [2]–[4] all cover uncertainties, theory and scenario generation and comparative assessments, and provide guidelines for conveying uncertainty.

The useful role and communication of Ignorance remains less well studied and is rarely included in analytic training except as identifying gaps in knowledge that need to be filled. Ignorance is seen as the antithesis of knowledge - knowledge and ignorance vary in inverse proportion to each other - and so the analyst's task is to reduce or eliminate ignorance.

We suggest that there is more to ignorance than simply not knowing; that beneficial and harmful ignorance can be distinguished from each other; and that the costs of reducing ignorance can outweigh the benefits. Identifying those conditions and contrasting what can be deliberately ignored with what needs to be known can help intelligence teams focus their effort on useful assessments.

We will describe two aspects of ignorance in intelligence: 'depth' of ignorance and 'flavours' of ignorance. We will suggest some beneficial and harmful motivations of the ignorant and ignorance-providers. We combine these to suggest training approaches that reinforce the values of deliberate ignorance in intelligence work.

# Depths of Ignorance

Robinson suggested classifications for data to be fed into simulation models which also seems relevant here:

> Type A data has been collected and is immediately available. (We may not know something but can do so quickly and easily).
>
> Type B data has not been collected, but could be if the time and effort were spent to do so. (We do not know something, but know how to know).
>
> Type C data is not collectable. (We cannot know something).

- [5]

Types A and B can obviously be expanded to describe how immediately the data is available (within reach, in another office, in another building, in another language, etc) and the costs and delays in acquiring it. Type C (eg past events that have not been observed) reminds us that some ignorance simply cannot be reduced

Later Phillip Armour offered the following more introspective five levels of ignorance in developing software:

> 0: Lack of Ignorance: You know something and can demonstrate it
>
> 1: Lack of Knowledge: You do not know something and can readily identify it
>
> 2: Lack of Awareness: You are not aware that you do not know something relevant, but you have ways of finding out. This is typical of software engineers starting out on a project.
>
> 3: Lack of Process: You are not aware that you do not know something relevant, and you do not have a means of finding out.
>
> 4: Meta Ignorance: You are not aware of the five orders of ignorance; you are not aware that you are not aware.

- [6]

Armour's second and third levels describe being ignorant of ignorance; not knowing what is not known. This is what Donald Rumsfeld later famously called "unknown unknowns":

> "There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns – there are things we do not know we don't know. "

- [7]

While Rumsfeld received some notoriety for that statement, it is a fairly succinct summary of knowledge scope. The implied four-way combination of 'known' and 'unknown' can be completed with 'unknown knowns' as the 'tacit knowledge' suggested by Polyani [8]:

| Known Knowns (Explicit knowledge) | Unknown Knowns (Tacit Knowledge) |
| --- | --- |
| The things we know.<br>*My bicycle is red* | Assumptions that are true; things in our worldview that we have forgotten we know.<br>*Roadcraft* |
| **Known Unknowns (Explicit Ignorance)**<br>Things we know we do not know<br>*Is this bridge strong enough for a tank?* | **Unknown Unknowns (Tacit Ignorance?)**<br>The things we do not even realise we do not know; hidden assumptions |

These suggest a range of 'depths' of ignorance:

**Not ignorant**: Knowing.

**Knowledge Ignorance:** Not realising what is known (tacit knowledge).

**Irrevocable Ignorance** Not being able to know; can be applied to any of the following:

**Simple ignorance**: Knowing what is not known - have the question and methods to answer it.

**Investigatable Ignorance**: Not knowing what is not known but have exploratory methods to discover the questions.

**Stumped ignorance**: Not knowing what is not known and not knowing how to know: Aware of ignorance in general but ignorant on how to deal with it.

**Invisible Ignorance**: Not realising that one is ignorant at all of a thing; blind spots, hidden assumptions, etc.

...each of which can be applied to ignorance of facts, theory, expertise and predictions.

# Flavours of Ignorance

## Ignorance of Information

*Ignorance of what, who and when*: You don't know something because you haven't been told; this is mere non-possession of a piece of information about the world.

Where the depth of this ignorance is limited to Simple or Investigatable Ignorance we should be able to find this information out by expending some effort and resources. The value of this ignorance therefore is the cost of the effort and resources to find the answer in time; if the value of the answer is less than the cost of answering it then this ignorance has net value.

A specific form of this is the practice of ‘hatching out’ areas on the map as ‘no go’ or ‘slow go’ areas: areas where the enemy is unlikely to operate at full mobility. These areas are essentially then deliberately ignored, which has value as attention and other resources can be focussed on areas more likely to be used by the enemy. However it requires careful procedures and training to decide what areas are hatched out, and how often they should be reviewed. It could be argued that the German surprise attack through the Ardennes in their invasion of France in 1940 surprised the French because they had effectively ‘hatched out’ the Ardennes as a ‘no go’ area.

Ignorance is not always well distinguished from ‘knowing there is not’ (*lack of evidence* vs *evidence of lack*). The empty space on the enemy overlay of a military map indicates that either nothing is known about the area or that the area has been searched and nothing found. This is obviously a crucial distinction but rarely made explicit. Even when a search is made, the analyst should beware that a search that finds nothing should not be taken as solid evidence that there is nothing there. Few searches are 100% effective, not least because the enemy hides. A particular example of this arose is an experiment in Australia when a UAV overflew and reconnoitred an area in which an armoured brigade was concealed (Paul Syms, personal communication). The sensor magnification was chosen to rapidly cover the search area but was too low for the targets to be registered; it was thus possible to look straight at a target and not see it. As a result the exercise analysts were not only ignorant of relevant information but had also denied themselves the means of resolving that ignorance. One of the authors (Salt) produced a simulation program to help find the best balance of sensor resolution and sweep rate to avoid this sort of problem. It is possible to imagine similar "Goldilocks problems" arising wherever rapid coverage needs to be balanced with close scrutiny.

## Ignorance of Theory

*Ignorance of why and how*. You don't know how things work, what tends to cause what, or why things are the way they are.

As with 'Ignorance of Information', Simple and Investigatable Ignorance of Theory can be resolved with suitable time and effort.

An alternative to understanding is to capture activities as processes, called 'routines' by Levitt & March [9],and 'rules' by Kieser & Koch [10], and so avoid the attention and resources required to teach and learn theory. An example would be the scripts used by first-line telephone support operators, written by technical experts, but that require no technical knowledge to follow. Military intelligence operators are encouraged to create 'products' such as SWOT tables and ASCOPE/PMESII walkthroughs which are essentially checklists of viewpoints that can be semi-mechanically completed.

Such processes may be learned by experience rather than supplied by experts. For example in this popular story of blindly learned process:

1) Place five gorillas in an outdoor cage in a cold, windswept location.
2) Suspend a banana in the cage above a ladder.
3) When any gorilla attempts to use the ladder, soak all five gorillas with a fire-hose. Repeat until they avoid the ladder.
4) After a short period, replace one gorilla. When the new gorilla attempts to use the ladder, the other gorillas beat him up.

5) Remove the fire hose.
6) After a short period replace another gorilla, and repeat until all have been replaced.
7) At this point no gorillas use the ladder; any that try are beaten up and none of them knows why.

This mocks the gorillas for being foolish in following useless process without knowing why. However this story also tells how groups can capture knowledge as process without being subject to community-wide pain. If the hose had been left in place in step 5, the gorillas have ‘learned’ to avoid being soaked without knowing why.

As well as defects in knowledge, one may be unaware of one's defects in skills. The awkward fact that unskilled people are the least likely to have an accurate perception of their own lack of skill has long been acknowledged. Recently much attention has been given to Dunning and Kruger's Ig Nobel-winning research "Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments" and the phenomenon is known as “The Dunning-Kruger effect” [11].

It seems obvious that expert analysts are to be preferred to incompetent ones. However, expertise can bring its own highly specific kind of ignorance.

## Ignorant Expertise

In his work on Recognition-Primed Decision-making (RPD), Gary Klein [12] discovered that a lot of expert decisions are typically made in a very short space of time, apparently without time to think. This is because experts can immediately recognise the essential elements of a problem in their field of expertise from a “pattern-book” they have accumulated from long experience, rather than working the problem out from first principles. They are using what Kahneman [13] in “Thinking Fast and Slow” calls “System 1” thinking – fast, automatic, subconscious and emotional. This produces quick decisions, whereas “System 2” thinking is deliberative, conscious, logical, and slow. Most of us get through life without too much “System 2” thinking, and experts are no exception.

This does not make expert decisions any less valuable, but it does mean that experts, their expertise having become subconscious, can find it difficult or even impossible to explain how the decision was reached. It is not that the expert does not know, but rather that they do not know how they know.

This may explain why so many military commanders, in the British army at least, are quite poor at specifying their CCIRs (Commander's Critical Information Requirements). Their military expertise is subconscious (in token of which the British army has for many years now used a tactical estimate procedure originally based on Klein's RPD), so they find it difficult to articulate; they will know what to do when critical information is presented to them and it triggers a decision, but they cannot so easily say in advance just what those critical information items will be.

One author (Salt) observes from experience as an Operational Research Analyst that the role often requires one to act as a sort of “expert ignoramus”.  An ignorant (but enquiring) newcomer looking with fresh eyes at a problem that has baffled the experts is forced to resort to “System 2” thinking, simply by their own lack of experience, and may thereby spot something the expert missed. Ignorance in this case may also circumvent the constraints imposed by conventional thinking among the people who own the problem, in which case it is important not just to bring a fresh mind to the problem but to bring an empty – ignorant – one.

## Ignorance of the Future

Given the stochastic nature of the material universe, one must accept that many phenomena can be characterised only in probabilistic terms. It is never possible to know, with certainty, what numbers will show when a pair of dice are rolled unless one is using some highly suspect dice. However knowing the statistically likely spread of results is useful, if it is possible to gain this. In some cases – such as drawing packages from a lucky dip – it may not be.

Predicting future events in the complex real world is usually at least as uncertain as statistical unknowability, with the added problem that the uncertainty is difficult or impossible to quantify. Any sufficiently complex system entails this kind of unpredictability, as occurs for example in weather forecasting or human behaviours such as economics or sporting contests. Perhaps surprisingly, a sufficiently complex system can be unpredictable without involving any stochastic elements. The complexity required to produce such deterministic unpredictability can be quite low, as for example in John Horton Conway's “Game of Life” [14], Craig Reynolds' “Boids” [15] or indeed the few lines of programming code required to implement a pseudo-random number generator (PRNG). In such cases we can only hope that there are recognisable patterns that arise from this complexity, as with Mandelbrot factors.

This deterministic complexity is in principle knowable only if the exact computations required are executed; the screen patterns from a game of life or a boids program, or the output from a PRNG, are simply the results of (sometimes quite simple) computation. However the results of such computations are strongly dependent on the precise starting conditions. Outside a digital computer, in real life, the limited precision with which these starting conditions are measurable mean that in practice they might as well be unknowable. But knowing their unknowability is at least knowledge of a sort; an acknowledgement of ignorance.

Where uncertainty is understood and expected, this kind of ignorance can be clearly labelled. As an example, in anti-submarine warfare, contacts are classified as NONSUB, POSSUB (LOW), POSSUB (HIGH), PROBSUB and CERTSUB according to the assessed probability that the contact is really a submarine.

# Choosing Ignorance

There are many epistemic, technical and social barriers to knowledge, such as not having sufficient resources to be able to obtain all the information required in the first place. Here we will introduce some good and bad motivations to *choose* to be ignorant.

## Ignoring Clutter to Avoid Overload

This class of ignorance is one that is deliberately chosen. Because human minds – even if not operating under stress – have a fairly low hrair limit on the amount of information they can hold at one time, it is necessary in all forms of analysis to simplify. This requires deliberately ignoring certain details; it is a choice to either forget it or not to find out in the first place. One of the key skills of the successful analyst is successfully recognising which details can safely be ignored in this way, and which are important.

While ‘breadth’ is already scoped geographically for military intelligence operators by defined Areas of Responsibility and surrounding Areas of Interest, the ‘depth’ is not well defined. For example standard military intelligence practice for land operations includes marking maps with overlays that marry the terrain to the enemy’s mobility, and so simplifies it. This might include marking terrain as ‘no go’ - the enemy is not able to move at significant speed through such terrain - which allows attention and sensors and the deployed forces to focus elsewhere. However, as described above in the example of Germany using the Ardennes in 1940, this ignorance is also the space that the enemy will look to operate in.

It might be argued that an analyst should always report the maximum amount of information available, and leave it to the recipient to decide which bits they think are important. This however just moves the cognitive load rather than solving it, and furthermore adds to the communication load. Similar issues face data scientists in the commercial world, where the ‘big data’ forced on analysts when ‘drinking from the data firehose’ needs to be transformed into ‘small data’ that can be understood by the executives. As the late Graham Mathieson used to remind people (quoting Claude Shannon): we should not only ask what consumes information, but also what information consumes; and information consumes the attention of the recipient (if it doesn't, it doesn't inform them). All extraneous information not required to inform a decision is not just harmless stodge, it is a waste of attention and thus of time, and risks distracting the recipient from the points that actually matter. If the decision-maker wants to dig deeper and have more information in a particular area, then certainly the analyst should go ahead and provide it, but they should not pre-emptively pack in everything “just in case”.

This seems especially likely to happen in two cases. The analyst might not really understand exactly why the decision-maker needs this information, and so provides everything that might possibly be useful in the hope of getting some hits with a scattergun approach. Worse still, in a low-trust working environment, analysts might include everything in an attempt to cover themselves from later accusations of overlooking some vital piece of information. One of the authors (Salt) recalls from his time in Saudi Arabia seeing the security briefings from the US Embassy which circulated occasionally. These typically warned that unspecified terrorists might be planning to strike American or other western military or civil targets in the Middle East or elsewhere at some time in the next month or later. By warning against everything, they warned against nothing, and everybody was lamentably unprepared when the al-Khobar Towers bomb went off.

The advanced analyst might also suspend judgement. This is not so much deciding not to know, as deciding not to decide. In general the military tends to admire prompt decision-making, but there is much to be said for the old Civil Service principle that “when it is not necessary to make a decision, it is necessary not to make a decision”. People find it harder to accept information that contradicts a decision they have already made [16] so not making a decision should help the analyst be less vulnerable to the effect. It does, however, require the analyst to be quite strong-minded in resisting the usual human need for psychological closure.

## Ignorance in Learning

Pavese [17] suggests that there might be no effective way of acquiring new scientific information without strong prior knowledge such as which experiments are valid, which variables must be controlled, which can be ignored, and so on. This is a familiar problem for educators; it may not be possible to teach everything at once and so Wood et al suggest that a ‘scaffolding’ of temporary ideas could be used [18].

The typical military problem is that land field commanders and intelligence operators are relatively young, and are trained and experienced in the importance of terrain as the significant factor in combat. Few have significant social or psychology scientific training. In typical peacekeeping or insurgency operations they are therefore ignorant of the frameworks of ideas needed to understand complex socio-political situations and suitable ways to shape them effectively. Adding such education to all the relevant staff would be a considerable training burden; instead a suitable ‘scaffolding’ of metaphors and models such as “Human Terrain” are used that make sense to people who work with maps. This ignorance becomes more costly when dealing with new operational environments that are not geographic such as ‘information operations’ and ‘cyber warfare.’

## Security/Secrecy

Protective security relies on building walls of ignorance based on the "need to know". Secrecy is ignorance deliberately enforced, whether on the enemy or on friends. Of course the problem with the “need to know” principle is that it can often be difficult to know whether one does or does not need to know something without knowing that something in the first place.

Current fashions tend to stress the “duty to share” much more, perhaps acknowledging that in the past excessive secrecy has damaged the quality of analysis to an extent not justified by the security benefits. Nonetheless, a certain amount of this kind of ignorance seems necessary.

## Preserving Status

Moore and Tumin [19] identify the role of ignorance in enforcing social hierarchies of class, power, and expertise.  The phrase “knowledge is power” is well known, and in authoritarian management hierarchies the senior managers try to preserve their status by keeping the underlings ignorant.  Anyone with much experience of working with government departments will probably be able to think of occasions when protective markings or other security measures were applied for reasons of social control (to assert status or avoid embarrassment) rather than genuine security. None of these practices tend to make for good analysis.

Shannon's idea that information content is measured by “surprisal” means that data that is exactly what you expected carries very little information (merely “nothing has changed”), whereas surprising data contains a good deal of information. It seems safe to say that information or intelligence is not much use if it does not tell you something new and interesting to inform your decisions. Therefore it must be something you were previously, to some extent, ignorant of. It seems likely that a status-conscious decision-maker who tries to enforce ignorance in others to bolster their own status is going to have some difficulty acknowledging their own ignorance by accepting the new information. In this way concern for social status inhibits successful communication, tending to justify Hagbard Celine's principle that there can be no communication except between equals [20]. Here the problem is not one of producing good analysis, but of having the results of good analysis accepted.

# Conclusion

We have described choices of ignorance that can sometimes be good and sometimes bad (decluttering, staged learning), choices that can be both good and bad depending on who you are (security), and choices that are usually bad (preserving status). These choices can be applied to various described depths and flavours of ignorance, some of which may also be directly useful (‘expert ignoramus’). This list may not be exhaustive; we don’t know. However it does demonstrate that ignorance is not, as one might naively assume, entirely negative. We hope that it will prove useful to intelligence trainers to support ignorance-structuring, ignorance management, and ignorance exploitation.

It follows that analysts should be trained in the uses and abuses of ignorance. In order to understand the value of ignoring clutter, intelligence training exercises should supply enough information and demand assessments in sufficient time that analysts are overloaded. The information supplied should include considerable clutter - irrelevant detail that nevertheless looks useful - encouraging choices to deliberately ignore distracting ‘rabbit holes.’  Examples of ‘surprise’ from the past - such as the German invasion of France through the Ardennes - can provide counter-lessons so that analysts are aware of their ignorance and review it. Analysts should be taught to distinguish between ‘ignorant’ and ‘empty’ when asked about the blank areas of the map or the edges of ORBATs.

Intelligence exercises from other domains - such as criminal investigations for military staff – can help train analysts to imagine new ways of discovering what they need to know and what they do not.

Finally common practices such as red teams and devil advocates should be exercised to expose “invisible ignorance”. In all cases trainers should be encouraged to echo Vroomfondel and demand of their analysts “rigidly defined areas of uncertainty and doubt.”

# References


[1] S. Kent, *Words of Estimative Probability*. CIA 2007, 1964.

[2] R. H. Pherson and R. Heuer, *Structured Analytic Techniques for Intelligence Analysis*. CQ Press, 2010.

[3] DIA, *Quick Wins for Busy Analysts*. 2013.

[4] CFIC, *Aide Memoire on Intelligence Analysis Tradecraft*, no. August. 2015.

[5] S. Robinson, *Successful Simulation: A Practical Approach to Simulation Projects*. Wiley Subscription Services, Inc., A Wiley Company, 1994.

[6] P. G. Armour, “The five orders of ignorance,” *Commun. ACM*, vol. 43, no. 10, pp. 17–20, 2000.

[7] D. Rumsfeld, “News Transcript: DoD News Briefing - Secretary Rumsfeld and Gen. Myers,” 2002. [Online]. Available: http://www.defense.gov/transcripts/transcript.aspx?transcriptid=2636. [Accessed: 13-Apr-2012].

[8] M. Polyani, *The Tacit Dimension*. 1966.

[9] B. Levitt and J. G. March, “Organizational Learning,” *Annu. Rev. Sociol.*, vol. 14, no. 1988, pp. 319–340, 1988.

[10] A. Kieser and U. Koch, “Bounded Rationality and Organizational Learning Based on Rule Changes,” *Manag. Learn.*, vol. 39, pp. 329–347, 2008.

[11] J. Kruger and D. Dunning, “Unskilled and unaware of it: How difficulties in recognizing one’s own incompetence lead to inflated self-assessments.,” *Journal of Personality and Social Psychology*, vol. 77, no. 6. pp. 1121–1134, 1999.

[12] G. Klein, “Sources of Power: how people make decisions.” 1998.

[13] D. Kahneman, *Thinking Fast and Slow*. 2011.

[14] M. Gardner, “Mathematical games: The fantastic combinations of John Conway’s new solitaire game ‘life,’” *Sci. Am.*, vol. 223, no. October, p. 120–123., 1970.

[15] C. W. Reynolds, “Flocks, Herds, and Schools: A Distributed Behavioral Model,” in *ACM SIGGRAPH*, 1987, vol. 21, pp. 25–34.

[16] S. Palminteri, G. Lefebvre, E. J. Kilford, and S.-J. Blakemore, “Confirmation bias in human reinforcement learning: Evidence from counterfactual feedback processing,” *PLOS Comput. Biol.*, vol. 13, no. 8, p. e1005684, Aug. 2017.

[17] C. Pavese, “Skill in epistemology I: Skill and knowledge,” *Philos. Compass*, vol. 11, no. 11, pp. 642–649, 2016.

[18] D. Wood, J. S. Bruner, and G. Ross, “the Role of Tutoring in Problem Solving,” *J. Child Psychol. Psychiatry*, vol. 17, no. 2, pp. 89–100, 1976.

[19] W. E. Moore and M. M. Tumin, “Some Social Functions of Ignorance,” *Am. Sociol. Rev.*, vol. 14, no. 6, p. 787, Dec. 1949.

[20] R. Shea, R. A. Wilson, R. Shea, R. Shea, and R. Shea, *The Illuminatus! trilogy*. Raven, 1998.