Essay

Not peer-reviewed version

Understanding AI-Enabled Outputs: A Conceptual and Technical Primer for All-Source Intelligence Analysts

George Ellison *

Posted Date: 18 June 2026

doi: 10.20944/preprints202606.1363.v1

Keywords: artificial intelligence; intelligence analysis; decision support; AI literacy; Human-AI Teaming; predictive estimation; mechanistic prediction

 doi:10.20944/preprints202606.1363.v1

Essay

Understanding AI-Enabled Outputs: A Conceptual and Technical Primer for All-Source Intelligence Analysts

George Ellison

Centre for Intelligence Studies, University of Lancashire, PR1 9HE, UK; gthellison@lancashire.ac.uk

Abstract

Intelligence analysts are increasingly able, and required, to consume and interpret outputs generated by artificial intelligence (AI) enabled tools — yet most receive little training in what these outputs actually represent, or how these might be robustly evaluated. This primer addresses that gap. It argues that analysts do not need to know how to develop or operate AI tools in order to use these tools' outputs critically and judiciously. But they do need sufficient conceptual understanding, and foundational technical knowledge, to evaluate these outputs competently. Three principles provide the framework for this understanding: First, the distinction between AI-*facilitated* outputs – where automation improves the pace, scale and fidelity of data collection, processing and analytical procedures that analysts could otherwise perform; and AI-*generated* outputs – where many of the novel outputs generated by the semi-autonomous techniques involved could not have been produced by analysts working independently; Second, the critical difference between interpolative and extrapolative estimation and mechanistic prediction; and Third, the critical dependencies and substantive limitations that govern the reproducibility and practical utility of all AI-*facilitated* and AI-*generated* outputs. Together these principles constitute the technical and conceptual foundations of the AI literacy training that all-source intelligence analysts should receive – the case for which is presented in a companion piece to this article.

Keywords: artificial intelligence; intelligence analysis; decision support; AI literacy; Human-AI Teaming; predictive estimation; mechanistic prediction

---

The epistemic and cognitive impact of 'artificial intelligence'

'Artificial intelligence' or 'AI' has become a troublesome, catch-all term (Sanguinetti, 2025) encompassing a growing array of computational techniques that use rules-based software. These include those that:

- simply automate existing or novel, analyst-defined procedures (known simply as "[un-intelligent] rules engines"); or
- can independently respond and adapt to the results of user-defined procedures (known as "intelligent rules engines"); or
- include components of, and contributions from, each of the above (Mitchell et al., 2019).

All of these techniques can appear 'intelligent' in their ability to perform tasks systematically and automatically, and thereby *facilitate* the production of what Donald Rumsveld (2002) famously called 'known unknowns' – i.e. "things that we know we don't know" – and to do so at a scale and pace, and with a level of fidelity that would otherwise require substantial time and effort on the part of experienced analysts. But the second (and, therefore, the third) of these techniques can also *generate* estimates of what Rumsveld (2002) called 'unknown unknowns' – i.e. "things that we don't know we don't know" – deduced from patterns hidden within the datasets analysed, and that reflect and reveal entities or features that would otherwise have been impossible for even the most skilled human analysts to detect (and even if unconstrained by the time and effort they would require).

Notwithstanding the unprecedented advance in analytical productivity, innovation and discovery that these computational techniques provide, the broader contemporary use of the term 'artificial intelligence' (or simply, 'AI') risks mistaking the outputs of AI-enabled tools as equivalent to human intelligence, intuition and sense-making by conflating:

- AI as a discipline with AI as an entity;
- AI as an aspiration with AI as a reality; and
- AI as a tool with AI as an independent agent.

As Sanguinetti (2025) points out, this conflation has led to "misunderstandings [that]… have spread throughout the lexicon of the discipline", and widespread misrepresentation and misunderstanding of what AI-enabled tools can *actually* achieve. Indeed, the widespread adoption of such terminology seems designed to imply that some AI techniques are *fully* autonomous or even sentient; even though, in reality, the most immediate advances that AI-enabled tools have achieved are simply a consequence of the automation of processes that humans themselves are able to perform – albeit at a far slower pace, a far lower scale, and with far less fidelity than AI-enabled tools commonly achieve. Terminology that favours the former representation of what AI-enabled tools are able to achieve include: 'machine *learning*' and 'deep *learning*' (two 'intelligent rules engines' that involve semi-autonomous, though user-designed/selected, model optimisation protocols; Mitchell et al., 2019); '*predictions*' (for the estimates of unmeasured, unknown or hitherto unknowable dataset features in the past, present or future, that AI-enabled tools can generate; Mitchell et al., 2019)); and 'AI *hallucinations*' (for outputs that have been: produced from patterns that result from non-systematic errors and systematic biases in data sampling, measurement and analysis; or fabricated by the tools themselves when programmed to impute, infer dataset features even when the information required to do so is incomplete or inconclusive; Sanguinetti, 2025).

Those who develop and operate the tools that produce AI-enabled outputs – as well as those who commission these tools and consume their outputs – should all be aware of the insidious effects these semantic sleights of hand can have on more prosaic yet more accurate and literal conceptualisations of: the tasks that AI-enabled techniques can (and cannot) perform; and the questions and problems for which AI-enabled techniques can (and cannot) provide meaningful answers, solutions and explanations (Thierry, 2025). Yet this can nonetheless be challenging, since the outputs produced by AI-tools have *both* epistemic *and* cognitive effects:

- 'Epistemic' as in the empirical and experiential evidence on which human knowledge of the world is based; and
- 'Cognitive' as in the variable and often imperfect information processing mechanisms, reasoning skills and heuristics that humans employ to evaluate and integrate such evidence into the corpus of knowledge and understanding they hold to be true.

Without sufficient conceptual understanding of how AI-enabled tools produce these outputs, those who consume or commission them risk misunderstanding both: what those outputs mean; and how they can (and should) be used. In the absence of this understanding, the insights that consumers and commissioners derive from the outputs of AI-enabled tools are vulnerable to inferential errors and biases. This is particularly the case within 'disciplines of uncertainty' – such as intelligence analysis (Pili, 2023) – in which the paucity of empirical information means their analysts are more open to, and reliant on, alternative sources of information, such as the theoretical and experiential understanding, and plausible speculation of the analysts themselves. In such disciplines, the allure of AI-enabled tools lies not only in the productivity gains these tools can offer, but also in their ability to reveal patterns hidden deep within the datasets available to them, that would otherwise be challenging, impracticable or even impossible for human analysts to identify *without* such tools.

 doi:10.20944/preprints202606.1363.v1

Yet the attraction of AI-enabled tools as a novel source of information (and the novel insights this information can support) wherever uncertainty is pervasive, also heightens the need for analysts to understand: what AI-enabled tools are *actually* doing; what their outputs *actually* represent; and how any inherent dependencies and limitations might be evaluated, attenuated or accommodated within the insights analysts derive from the outputs these tools produce.

The twin aims of this article are therefore to:

- provide an accessible conceptual and technical primer to accompany a companion piece in which the case for providing foundational 'AI literacy' training (Konishi, 2015) to all-source intelligence analysts is made (Ellison – submitted); and
- situate this against the backdrop of AI-related policies and practices being developed to support the integration of AI-enabled tools within intelligence workflows.

The three principles that follow are intended to address the first of these aims; and although none of these principles requires specialist technical training to understand or apply, each can help analysts navigate the analytical and inferential pitfalls that AI-enabled techniques entail. These principles mirror the more technical considerations that analysts using computer-enabled statistical techniques have needed to bear in mind when applying these to analyse and interrogate large and complex datasets – techniques that include early versions of 'intelligent rules engines' (such as: multivariable linear and logistic regression; and analysis of covariance, ANCOVA; Ellison 2026).

As for the second of these aims, this will be addressed by critically examining Deloitte's "task-level" review exploring the potential impact of AI on intelligence analysis – a review that offers a compelling, if somewhat partial, case for the integration of AI-enabled tools within intelligence workflows (Mitchell et al., 2019).

The Future of Intelligence Analysis in the Age of AI

Although the Deloitte review (Mitchell et al., 2019) review paid little heed to the conceptual and technical understanding that intelligence analysts require to critically and competently consume AI-enabled outputs, it enthusiastically endorses a growing role for AI within intelligence analysis. In particular, it summarises the key incentives for, and many benefits of, integrating AI-enabled tools into intelligence workflows – and emphasises the efficiencies that AI might bring by: automating lower-value tasks; supporting higher-value analytical and inferential tasks; and providing more timely, comprehensive and customer-focussed decision support (Mitchell et al., 2019). However, any such efficiencies are likely to be substantively attenuated by the time, resource and training costs required to implement, maintain and sustain this shift in analytical practice. Even were the review's projected efficiencies to be achievable, they would still depend on substantial organisational restructuring, sustained investment, and comprehensive workforce redesign to deliver the organisational systems and structures required, and recruit or develop staff who are competent to: consume and commission AI-enabled tools; and critically evaluate these tools' outputs. Indeed, Mitchell et al. (2019) acknowledged a number of limitations and uncertainties that undermined the findings of their review, and cite three specific caveats that are particularly relevant to the training needs of all-source intelligence analysts:

- First, that "intelligence work never ends" – meaning that prioritisation will need to remain a central and essential component of professional intelligence practice irrespective of any improvements in information processing and analytical capacity that AI might deliver;
- Second, that "AI is not the solution to every problem" – meaning that the utility of AI-enabled tools will depend in no small part on the nature of the task, the datasets available, and the costs of implementation; and
- Third, that the "impact of AI on analysts' cognitive biases remains uncertain" – meaning that intelligence agencies will need to carefully monitor the cognitive capabilities of analysts using or teaming with AI, and not least given the risk that familiarity with AI-enabled outputs can

make consumers more trusting and less critical of these (Heuer, 1999; Mitchell et al., 2019; Gillespie et al., 2025; Dolman, 2024).

Beyond these caveats, there are at least four additional considerations Mitchell et al. (2019) overlook that are also relevant to the realisable utility of AI-enabled outputs in all-source intelligence analysis; and each of these is worth examining in some detail:

1. The unprecedented (and essentially unique) nature of intelligence requirements

A variable (and potentially substantial) proportion of intelligence requirements concern phenomena, entities, processes, or features thereof, that are historically unprecedented, or so spatio-temporally contingent as to be essentially (and quite literally) unique. As a result, the data available for these may lack the volume or variety required to justify or support the use of AI-enabled information processing and analytical techniques. Under such circumstances – and where substantial uncertainty resides in the paucity of definitive evidence – analysts will need to retain the specialist cognitive and analytical skills required to generate more theorised and speculative, yet resolutely robust, assessments on which to provide decision support (Pili, 2023; Ellison, in press).

2. Data classification, storage and fusion constraints

Combining datasets and fusing information from different systems, often held at different classifications, poses a long-standing challenge to automated processing in intelligence analysis (Hanna et al., 2017). This becomes more difficult still when the information streams and datasets involved are large, require substantial secure storage, and must be hosted within systems that can support the necessary processing, automation and analytical software; and from which any resulting algorithms and outputs can then be extracted or released.

3. AI is evolving, and so are the analysts who are expected to use it

Much of Mitchell et al.'s (2019) enthusiasm for integrating AI-enabled automation and analysis within intelligence workflows reflects the state of AI technology in 2019 (when their review was published), and the contemporaneous capabilities of intelligence analysts whose exposure to analog vs. digital technologies now include those born in the 1960 through to the early 2000s. Although the pace, direction and character of developments in AI technology remain uncertain, it is clear that current and future cohorts of intelligence – including yesteryear's 'analog natives', today's 'digital natives', tomorrow's 'AI natives', and the future's 'quantum natives' (Dolman, 2024) – may require very different forms of training and support to strengthen and sustain their analytical judgment, assessment skills and capacity for human-AI teaming. And whatever the future brings, it is unlikely to reproduce the same discrete set of challenges that intelligence analysts face today. This means that the doctrinal, training-related and technical infrastructure required to integrate AI into intelligence workflows will need continual revision as both the technology, and the socio-technical environments that shape analysts' cognitive capabilities, continue to evolve.

4. The people who can do this well are scarce and in demand

Analysts trained to consume, commission or produce AI-enabled outputs – and to integrate such outputs judiciously within intelligence analysis and assessment for decision support across complex, unpredictable and uncertain multi-domain contexts — remain in short supply and in high demand. For the foreseeable future, intelligence agencies are therefore likely to face stiff competition to recruit and retain personnel with the competencies, motivation and aptitude required to develop and deploy these skills — not least from private-sector technology and commercial intelligence organisations, where experienced and technically competent analysts are highly sought after (and often command levels of remuneration far above those available within public-sector defence and security agencies).

Taken together, these additional considerations suggest that the integration of AI-enabled tools into intelligence workflows cannot be treated as simply facilitating the provision and integration of these more powerful tools. It will substantively depend upon an analyst workforce that can understand what these tools are doing, what sorts of outputs they produce, and what kinds of insights can be safely drawn (and what claims can be confidently made) on the basis of these outputs. The three principles that follow are central to these capability-based dependencies, and cover the most pertinent conceptual and technical consideration that analysts need to understand to align their

 doi:10.20944/preprints202606.1363.v1

practice with the system-wide transformation of their operational workplaces that the integration of AI-enabled tools will require.

AI-Enabled Outputs: A Conceptual and Technical Primer for Intelligence Analysts

Principle 1: The important distinction between AI-facilitated and AI-generated outputs

Much of the excitement surrounding AI focuses on the unique quasi-'predictive' capabilities of a subset of AI techniques that involve 'intelligent rules engines' — particularly so-called 'machine learning' and 'deep learning' (Mitchell et al., 2019). Yet it is important not to ignore how tasks performed by much simpler, and far less sophisticated, innovative or inscrutable '[un-intelligent] rules engines' can greatly facilitate analyst-generated outputs and the insight these can support. This is because whenever the automated tasks these '[un-intelligent] rules engines' perform – by, for example: observing and measuring; sourcing and locating; collating and combining; formatting and standardising; classifying and categorising; organising and arranging; and summarising and visualising information – enable the human analysts involved to undertake analyses, produce outputs and derive insights, then the uplift in (human-derived, yet AI-*facilitated*) insight can be just as profound, and potentially more relevant, meaningful and useful, as that *generated* solely by AI.

In this sense, the key strengths of AI-*facilitated* outputs are their (relative) simplicity, accessibility and interrogability, and the benefits that accrue when human analysts remain central to the analytical process (through 'human-AI teaming'; NASEM, 2023). These benefits include the broader contextual understanding that only humans can provide, such as the attribution of meaningfulness and utility, and the filtering of discernible and emergent patterns according to their actionability – thereby identifying outputs and deriving insights that not only help reduce a decision-maker's uncertainty, but can also offer them more practical (that is, relevant, meaningful and actionable) decision support. Added to these is the substantive benefit of being able to replicate whatever AI-*facilitated* analyses have produced, since these predominantly involve the automation of analyses that analysts themselves *could* perform (though not at such pace, at such scale or with such fidelity; Ellison, in press).

In contrast, 'AI-*generated* outputs' comprise those produced by a sub-set of semi-autonomous AI techniques. These techniques exploit the statistical-associational properties of relationships amongst variables within a dataset to identify optimal algorithmic (i.e. arithmetic and algebraic) solutions to user-specified computational tasks. This is an approach that was widely practised using within computer-based statistical packages decades before the high-powered computational capacity enhanced the pace and scope with which these could be performed semi-autonomously. These include:

- *'supervised machine learning'* – which involves comparing multiple permutations of alternative algorithms to identify (or 'train') the optimal algorithm that: most consistently or accurately generates similar values to those that are known (and available) for a pre-specified specified 'target variable' or 'target dataset property or parameter' within the training dataset concerned – algorithms that can then be used to estimate, impute or 'predict' (through interpolation or extrapolation) values for such 'targets' in subsequent, comparable datasets where these values are unattainable or have yet to have occurred, been measured, observed, or recorded; and
- *'unsupervised machine learning'* – which operates in a similar fashion, but in lieu of a 'target variable' or 'target dataset property or parameter' on which the optimal algorithm can be identified (or 'trained'), the 'target' involved is one of a number of unmeasured (i.e. 'latent') and *user-defined* dataset properties, or parameters – such as a *'latent variable'* (i.e. an unmeasured and therefore hidden phenomenon, entity, process, or characteristic thereof, that can be elucidated from user-specified statistical-associational properties observed within the dataset); or a discrete number of *'latent classes'* of cases (i.e. unmeasured and therefore hidden

groups or clusters of cases that share more similar statistical-associational features than cases more closely aligned with *other* groups or clusters of cases, or none). These techniques make a powerful contribution to the perceived utility of AI-enabled tools, because their ability to estimate the value of missing values, or hidden properties, parameters and features – and to do so (in a good many datasets) with considerable accuracy and precision – offers what appear to be hitherto unattainable 'predictive estimates' of two principal sources of uncertainty: 'known unknowns'; and 'unknown unknowns' (Luft& Ingham, 1955; Rumsveld, 2002; Davies, 2010; Ellison, in press).

Indeed, the ability of this subset of AI techniques to offer predictive estimates of hitherto 'unknown unknown' phenomena, entities, processes, and characteristics thereof, means these offer an unprecedented advance in measurement one that seems likely to eclipse even the most dramatic step-changes in 'knowing' that have occurred previously – such as the invention of optical instruments (microscopes and telescopes; Wall, 2018; Lopez, 2021) in the 16th and 17th centuries; and sensors for measuring the full extent of the electromagnetic spectrum (in the 19th and early 20th century; Ball, 2007). Moreover, 'intelligent rules engines' look likely to dramatically expand the capabilities of all preceding technologies – while capitalising on the data these technologies generate to reveal structures, processes and mechanisms that might well go some way beyond our wildest imagination.

For the analyst, the practical importance of this distinction is straightforward. AI-*facilitated* outputs can often be interrogated, replicated or checked against processes and judgements that trained analysts understand well, and are able to produce themselves. AI-*generated* outputs, by contrast, can offer genuinely novel insights that analysts could not independently derive, and therefore demand a different and more conceptually advanced form of scrutiny. Distinguishing between these two broad categories of output is therefore a necessary first step in knowing what questions to ask of AI-enabled tools, and what kinds of inferences and claims can safely be made on the basis of their outputs.

Principle 2: 'Predictive analytics' generate descriptive estimates, not bona fide 'predictions'

In common parlance, a 'prediction' reflects the ability to accurately determine something that has yet to occur, or an observation or measurement of a past, present or future entity, phenomenon, process, or characteristic thereof, that has yet to be made. As a concept it has particular resonance in contexts where there is little, if any, knowledge or understanding of the mechanisms involved in producing the discrete entities, phenomena, processes, or characteristics thereof, to which the prediction relates. Indeed, so-called 'intelligence failures' are often erroneously attributed to those predictive assessments and judgements that, despite a high absolute or relative 'estimative probability' (Dhami et al., 2021), do not subsequently occur (Jensen, 2012).

Under such circumstances, a prediction that subsequently turns out to be true is commonly taken as proof of specialist predictive expertise, or even of mystical or prophetic powers. This is not without justification if the basis on which the successful prediction was made is unclear or unknown to anyone else. Moreover, even predictions that would have later turned out to be wildly inaccurate can inspire sufficient confidence to take on a life of their own as 'self-fulfilling' or 'self-defeating' prophecies. In such cases, the prediction acts as an intervention because it persuades those who believe it to precipitate, pre-empt or prevent, whatever it was that had been predicted to occur.

Unsurprisingly, then, the use of the term 'prediction' to describe the statistical estimates that AI techniques known as 'intelligent rules engines' (such as machine learning or deep learning; Mitchell et al., 2019) can generate – whether of past or future, and hitherto unobserved or unmeasured, entities, phenomena, processes, or characteristics thereof — imbues such techniques, and the very notion of AI, with capabilities that far exceed anything most of us would be willing to ascribe to even the most proficient (and successful) human forecaster. At the same time, the complex and opaque algorithmic - that is, arithmetic and algebraic - solutions that these AI-enabled techniques construct to generate their optimal predictive estimates only serve to enhance their allure as something we know or believe to be true but cannot interrogate, and might never fully understand or explain.

In reality, the facts of the matter are somewhat more mundane. 'Intelligent rules engines' simply exploit the statistical, distributional and associational properties of the dataset(s) provided to generate optimal algorithmic solutions to user-specified computational tasks. Yet even consumers, commissioners and producers of AI-enabled outputs who fully grasp the basis on which these AI techniques generate their so-called predictions can struggle to resist the temptation to view AI as anything other than uncannily prescient – as evident in the case of the Google engineer (de Cosmo, 2022) and the celebrated evolutionary biologist (Moss and Kolin, 2026) who came to believe that the chatbots they had helped develop, or had worked with, were 'alive' and sentient. Both cases serve as vivid examples of the substantive cognitive impacts5 that AI-enabled tools (and their outputs and associated insights) can have.

For the analyst, the practical implication is straightforward. AI-enabled 'predictions' should be approached not as privileged glimpses of what is likely or going to happen, but as context-dependent statistical estimates whose meaning, utility and trustworthiness depend on: the data from which they were derived; the tools used in their analysis; the conditions under which these tools were applied; and the extent to which these conditions are comparable to those that currently prevail.

Principle 3: The three dependencies and five limitations of AI-enabled 'predictive estimation'

The three principal 'dependencies' of all currently available, semi-autonomous, AI-enabled analytical tools are that these can only generate quasi-'predictive' estimates of unmeasured, unknown or hitherto unknowable variables or dataset properties and parameters *if*:

- the dataset(s) concerned contains the 'informational perspectives' and 'statistical power' required to algorithmically identify the patterns that are necessary to generate interpolative or extrapolative estimates of unmeasured, unknown or hitherto unknowable variables or dataset properties/parameters;
- the analysts concerned have sufficient prior 'knowledge and understanding' (i.e. empirical and theoretical evidence, subject matter expertise and domain-specific awareness; NASEM, 2023) to interpret, evaluate and validate the meaning and utility of any patterns identified, and the 'predictions' these support; and
- the algorithmic identification of these patterns can subsequently be faithfully *and* usefully replicated in broadly 'comparable datasets' that have been derived using similar 'data- and dataset-generating mechanisms and procedures'.

The five principal 'limitations' of all 'predictions' generated by currently available, semi-autonomous, AI-enabled analytical tools are that – even when the dependencies listed above have been satisfied, these 'predictions':

- will be more accurate estimates of the *sample average* value of the 'predicted' phenomena, entities, processes, or characteristics thereof, than of their value for any *individual* case within the sample analysed;
- will not necessarily be an accurate estimate of the *population average* value of the predicted entities, phenomena, processes, or characteristics thereof, unless the *sample* of cases analysed is truly representative of the wider *population* of cases from which the sample was drawn;
- will not provide interpretable evidence of the direction, strength or precision of any (direct or indirect) causal effects amongst the variables retained in the optimal predictive algorithm;
- will rarely offer accurate estimates of *future* phenomena, entities, processes, or characteristics thereof, and only then if the data and dataset generating mechanisms involved are *not* sensitive to, dependent on, or subject to known, unknown, unpredicted, unpredictable or unprecedented spatio-temporal changes in *external* factors; and

 doi:10.20944/preprints202606.1363.v1

- will not be definitively validated (and cannot be considered 'proven') by subsequent phenomena, entities, processes, or characteristics thereof, that are concordant with those that had previously been predicted – since these may still have occurred as a result of chance; or as a result of 'self-fulfilling' or 'self-defeating' prophecies.

As a result of these dependencies and limitations, **evaluating the validity and practical utility** of any meaningful and actionable inferences that can reliably be drawn from AI-*generated* 'predictions' and their associated insights, will require substantial knowledge and understanding relevant to:

- the dataset(s) *on* which (and the spatio-temporal contexts *in* which) the 'predictive' algorithms were trained and subsequently applied; and
- the subject matter expertise (i.e. the empirical, experiential and theoretical evidence, and plausible speculation) and domain-specific awareness necessary to offer robust assessments of any credible alternative explanations for the insights concerned.

While some of the knowledge, understanding and skills required to evaluate the utility, meaning and validity of AI-*generated* outputs is already covered within the post-basic, specialist training, and subsequent in-service supervision and support, that many all-source intelligence analysts receive, additional augmented AI literacy training (Konishi, 2015) will be required to cover the more technical dependencies and limitations of AI-*generated* outputs if analysts are to be sufficiently competent and confident to apply these techniques (or use insights produced by others who have used these techniques) to strengthen their analyses and assessments.

As a result of these dependencies and limitations, evaluating the validity and practical utility of any meaningful and actionable inferences drawn from AI-generated 'predictions' and their associated insights requires both: knowledge of the dataset(s) on which the relevant algorithms were trained, and applied; and sufficient subject matter expertise needed to derive plausible alternative explanations. Some of the knowledge, understanding and skills required to do this may already be covered within the specialist training and subsequent in-service support received by many all-source intelligence analysts. However, additional AI literacy training will still be required if analysts are to become competent and confident in applying these techniques – or in using insights produced by others who have applied them – so as to strengthen rather than compromise their analyses and assessments. For the analyst, the practical implication is clear: AI-generated outputs should never be treated as self-explanatory or self-validating, but always as context-bounded estimates whose value depends on whether: their dependencies have been satisfied; and their limitations properly understood.

Conclusion

The utility of AI-enabled tools across a range of intelligence collection disciplines and analytical specialisms has meant they have become an integral component of strategic, operational and tactical intelligence practices from at least the mid-1960s onwards. Yet the growing use of AI-enabled tools within intelligence collection, processing and analysis has not been matched by a comparable growth in the conceptual and technical understanding required to critically interpret their outputs. This matters because the outputs of such tools are neither self-explanatory nor infallible, and because the confidence these tools can inspire may easily exceed the understanding with which their outputs are consumed.

The three principles set out in this primer are intended to address that gap. Taken together, they make a straightforward but important case, namely that:

- AI-*facilitated* and AI-*generated* outputs are not the same kind of thing and should not be treated as if they were;

 doi:10.20944/preprints202606.1363.v1

- so-called AI-enabled 'predictions' can be better understood as statistical estimates generated through algorithmic interpolation or extrapolation than as *bona fide* (or 'mechanistic') predictions in any firmer sense; and
- the validity and practical utility of AI-generated outputs are always governed by a critical dependencies and substantive limitations that analysts can, and should, learn to evaluate.

None of this requires most analysts to attain the specialist knowledge, experience and expertise required to become developers or operators of AI-enabled tools. Instead, it requires a level of conceptual and technical understanding that allows them to interrogate AI-enabled outputs with the same discipline these analysts would bring to any other source of information or intelligence. In this sense, the aim of this primer has not been to turn analysts into developers or operators of AI-enabled tools, but to help them become more confident, critical and competent consumers and commissioners of their outputs.

The broader training implications of this argument are developed in a companion piece, which makes the case for treating AI literacy training as a professional baseline for all-source intelligence analysts rather than as a specialism reserved for the technically inclined. Read together, the two articles argue that the effective integration of AI into intelligence analysis workflows depends not only on the complexity and sophistication of the tools available, but also on the quality of the judgment brought to bear on what these tools produce.

References

1. Ball, D.W. The electromagnetic spectrum: a history. *Spectroscopy* 2007; 22: 14-20. Archived here.
2. Davies A, Thomson M. Known unknowns: Uncertainty about the future of the Asia-Pacific. *Australian Strategic Policy Institute Special Report* 2010; 35: 1-16.
3. de Cosmo L. Google engineer claims AI chatbot is sentient: why that matters. *Scientific American* 2022; Jul 12. Archived here.
4. Dhami MK, Mandel DR. Words or numbers? Communicating probability in intelligence analysis. *American Psychologist* 2021; 76: 549–560. DOI: 10.1037/amp0000637
5. Dolman, J. Naïve or native? Do digital natives really understand their world? *The AI English Teacher* 2024; Sep 11. Archived here.
6. Ellison, G. What 'typical insights' might 'core statistical analytical techniques' provide to intelligence analysts? *Preprints* 2026; 02 Jun: 2026060055. Available here.
7. Ellison, G. (in press). Might 'uncertain military leadership' minimize the risks and optimize the opportunities of strategic uncertainty? Contemporary Military Challenges.
8. Ellison G. All-source intelligence analysts as confident consumers of AI-enabled outputs: The pressing case for AI literacy training. *Journal of Intelligence, Security and Resilience Education* – submitted.
9. Gillespie N, Lockey S, Curtis C, Pool J, Akbari A. *Trust in Artificial Intelligence: A Global Study*. The University of Queensland and KPMG Australia 2023; 1-79. DOI: 10.14264/00d3c94. Archived here.
10. Hanna M, Granzow D, Bolte B, Alvarado A. NATO intelligence and information sharing: Improving NATO strategy for stabilization and reconstruction operations. *Connections Quarterly Journal* 2017; 16(4): 5-33. DOI: 10.11610/Connections.16.4.01
11. Heuer RJ. Psychology of Intelligence Analysis. Center for the Study of Intelligence, Central Intelligence Agency 1999: 184pp. ISBN: 1929667000.
12. Jensen MA. Intelligence failures: What are they really and what do we do about them? *Intelligence and National Security* 2012; 27: 261–282. DOI: 10.1080/02684527.2012.661646
13. Konishi, Y. (2015, October 10). What is needed for AI literacy? Priorities for the Japanese economy in 2016. Research Institute of Economy, Trade and Industry. web.archive.org
14. Lopez SG. A brief history of optical microscopy. *John Innes Centre, Blog* 2021; May 12. Archived here.
15. Luft J. Ingham, H. The Johari window: A graphic model of interpersonal awareness. *Proceedings of the Western Training Laboratory in Group Development* 1955; Aug: 1-246.

16. Mitchell K, Mariani J, Routh A, Keyal A, Mirkow A. The future of intelligence analysis: a task-level view of the impact of artificial intelligence on intel analysis. *Deloitte Center for Government Insights* 2019; Dec 11. Archived here.
17. Moss MF, Koplin J. Is Richard Dawkins right about Claude? No. But it's not surprising AI chatbots feel conscious to us. *The Conversation* 2026; 7 May. Archived here.
18. NASEM (National Academies of Sciences, Engineering, and Medicine). *Human-AI Teaming: State-of-the-Art and Research Needs*. The National Academies Press, Washington DC 2022; 125pp. DOI: 10.17226/26355
19. Pili G. A new theory of surprise – unraveling the logic of uncertainty and knowledge. Intelligence and National Security 2023; 39: 695–708. DOI: 10.1080/02684527.2023.2278845
20. Rumsveld, D. (2002, February 12). DoD News Briefing. United States Department of Defense. Available here.
21. Sanguinetti P. Why the term 'artificial intelligence' is misleading. *IE Insights* 2025; Jan 22. Archived here.
22. Thierry G. We need to stop pretending AI is intelligent – here's how. *The Conversation* 2025; Apr 14. Archived here.
23. Wall WA. *History of Optical Telescopes in Astronomy*. Springer Nature, Switzerland AG; 2018. 173pp. ISBN: 978-3-319-99087-3.