Essay

Not peer-reviewed version

All-Source Intelligence Analysts as Confident Consumers of AI-Enabled Outputs: The Pressing Case for AI Literacy Training

George Ellison *

Posted Date: 18 June 2026

doi: 10.20944/preprints202606.1383.v1

Keywords: artificial intelligence; intelligence analysis; AI literacy; human-AI teaming; analytical training

Essay

All-Source Intelligence Analysts as Confident Consumers of AI-Enabled Outputs: The Pressing Case for AI Literacy Training

George Ellison

Centre for Intelligence Studies, Research & Knowledge Exchange Service, University of Lancashire, Preston, Lancashire PR1 9HE, United Kingdom; gthellison@lancashire.ac.uk; +44(0)7818 457 774 (M/W)

Abstract

Artificial intelligence (AI) is rapidly reshaping how information is collected, processed and analyzed by intelligence agencies. There is therefore a pressing need for foundational AI literacy training to ensure intelligence analysts become confident consumers of the outputs that AI-enabled tools can provide – but not for the specialised computational and statistical training required to become accomplished developers or operators of such tools. AI literacy training should distinguish between: AI-*facilitated* outputs, where automated collection, processing and analytical workflows can improve the consistency and efficiency of tasks that analysts themselves could perform; and AI-*generated* outputs, where semi-autonomous computational techniques can prompt unprecedented insights into hitherto, and in some cases previously unknowable, dataset features – insights that nonetheless entail dependencies and limitations requiring careful consideration and evaluation. Drawing on influential policy reports and academic articles exploring the potential utility of AI within intelligence analysis, this paper concludes that the transition to fully AI-enabled analytical capability demands a substantive expansion in AI literacy training; the conceptual and analytical foundations of which have been summarized in a companion piece to this article.

Keywords: artificial intelligence; intelligence analysis; AI literacy; human-AI teaming; analytical training

---

How AI Is Reshaping What We Know (And What We Need to Understand)

Ongoing developments across a broad range of automated and semi-autonomous computational procedures (popularly known, though unhelpfully lumped together, as 'artificial intelligence'; or simply 'AI'; Mitchell et al., 2019; Sanguinetti, 2025) continue to accelerate. The quickening pace of these innovations makes it challenging to keep abreast of the opportunities, risks and training needs involved. These span all areas of human experience and endeavor. Indeed, the scale, scope and reach of what has been dubbed the "Fourth Industrial Revolution" (4IR; Moll, 2021) can be bewitching (Thierry, 2025) and bewildering in equal measure (Marr, 2024).

Notwithstanding persistent uncertainty and an enduring degree of healthy skepticism regarding quite how "industrial" or "revolutionary" these developments might actually prove be (West & Allen, 2018; Marr, 2024), the potential for AI-enabled tools to challenge and upend much of how we currently view, understand and engage with the world appears unprecedented. It may yet prove transformational (West & Allen, 2018). This reflects the fact that AI-enabled outputs have both *epistemic* and *cognitive* impacts. These include those affecting *what* humans know – and *how* they know what they *think* they know – about the world and their place within it. As a result, the attendant and wide-ranging repercussions of AI seem likely to dwarf the earlier psychosocial and societal impacts of ostensibly comparable technological advances (Chantel, 2023; Volti & Croissant, 2024).

Moreover, the extent to which the outputs that AI-enabled tools provide, and the insights these outputs can support, might empower and enable us to learn more about the world than we ever thought possible, and make better decisions than we ever dared take, seems likely to substantively undermine the prevailing rationale for investing in our somewhat limited, and very human, sensory and cognitive capabilities. Yet AI-enabled outputs are neither infallible nor comprehensive; and they cannot (yet) ascribe the meaning(s) that ultimately determine their utility, value and importance to us. Humans will therefore need to keep their wits about them if they are to develop, maintain and apply a sound *foundational understanding* of the basis upon which AI-enabled tools operate – including the extent to which the outputs these tools offer, and the insights they support, are ultimately dependent upon, and constrained by the: datasets; computational techniques; and analytical/inferential competencies of the 'developers' and 'operators' of AI-enabled tools, and the 'consumers' and 'commissioners' of their outputs . This is because a sound foundational understanding – what Konishi (2015) dubbed 'AI literacy' – is both *fundamental* and *indispensable* for:

- differentiating between the questions and problem sets that different AI-enabled tools can usefully address (and those that they cannot); and
- critically evaluating the validity, practical utility and meaningfulness of the answers and solutions that AI-enabled outputs, and the insights these support, might provide.

In this regard it is especially important that all of us who can, make sure that we take personal and collective – that is, individual and institutional – responsibility for understanding the basis on which AI-enabled tools work. We cannot abrogate these responsibilities to third parties or specialist colleagues – even those who are proficient and trusted developers or operators of these tools. This is because third parties and specialist colleagues are both vulnerable to financial, commercial and/or reputational conflicts of interest. At best, they will be somewhat indifferent to the (accidental or intentional) misrepresentation of AI technology by others. At worst, they will have a vested interest in actively promulgating the deliberate obfuscation and mystification of AI-enabled tools – so as to enhance the perceived value and importance of these tools (and of those who can develop and operate them; Markelius et al., 2024). Indeed, the training received by a good many of those with substantial expertise and experience as developers and operators of these tools is often so narrowly focused on the technical practicalities involved that it pays little heed to the foundational understanding required to:

- meaningfully and judiciously select and apply only those tools to those questions and problem sets these tools are specifically designed (or equipped) to address;
- critically evaluate the outputs (and any insights these support) that context-dependent applications of these tools produce; and
- look beyond the *intended* outputs of these applications to consider any plausible alternative explanations or interpretations relevant to the analytical context involved.

The last of these is a particular concern when the outputs appear comfortingly and reassuringly consistent with prior expectations – and therefore attract less careful examination, less critical scrutiny, and less robust assessment than that required to attenuate the possibility that these might *actually* be:

- chance phenomena resulting from *non-systematic* errors (often unhelpfully labelled 'measurement errors', though these can also reside in the inherent instability and variability of the data concerned, and can occur at any stage of the analytical process, from design through to interpretation);
- methodological artefacts resulting from *systematic* errors (i.e. methodological biases) in analytical design and data sampling, measurement, collection, analysis or interpretation; or
- deliberate misrepresentation or falsification – comprising outputs that have been intentionally vitiated, manipulated, obfuscated or fabricated for the purposes of deception (Ellison, 2026c).

What the curricula of AI literacy courses need, and need not, include is worth stating plainly. So too is a clarification of what is meant here by 'AI literacy' which – in this instance – does not refer to the specialist competencies required to develop or operate AI-enabled tools, but to the conceptual

 doi:10.20944/preprints202606.1383.v1

and analytical understanding required to purposefully commission and judiciously consume, interrogate, understand and integrate their outputs, and the insights these support. The curricula of AI literacy courses therefore do not need to ensure that *all* analysts know how to develop or operate AI-enabled tools, since analysts do not need this skillset commission what tools they require, or to consume their outputs, by carefully and critically evaluating and interpreting these, and any resulting insights. Instead, what analysts need (and what AI literacy training must provide) is a sufficient level of conceptual and analytical understanding to answer three sets of critical questions when dealing with AI-enabled outputs, and the insights these support (Torres Torres et al., 2019; Mitchell et al., 2019). These concern:

- First, what kinds of questions any given AI-enabled tool is designed (or equipped) to address; and whether any such questions are relevant to the priority intelligence requirements (PIRs; DCDC, 2024) their decision-maker needs addressed;
- Second, the scope, scale, integrity and relevance of the dataset on which this tool was developed (i.e. 'trained'); and whether this was sufficiently comparable to the dataset(s) available in the context of the PIRs concerned; and
- Third, what the outputs of this type of tool actually represent (i.e. whether these are simply descriptions or predictive estimates useful for targeting or preparation, respectively; or reflect insightful classifications or tangible causal relationships amenable to exploitation or attenuation/augmentation, respectively); and what potential methodological dependencies and limitations might apply (and warrant evaluation), so as to assess the reliability, accuracy, precision and trustworthiness of these outputs (Ellison, 2026a).

These three sets of questions are not substantively *technical*; they are predominantly *conceptual* and *analytical*. They are also the sorts of questions that contemporary training in AI-relevant computational skills would not necessarily equip or encourage analysts – or indeed the developers and operators of AI-enabled tools – to ask (Mitchell et al., 2019; NASEM, 2022).

These considerations are particularly important within what we might call 'disciplines of uncertainty' – whose central task is to generate the best possible judgments within contexts where substantive evidence is incomplete, ambiguous, unreliable, or unavailable. In such contexts – and in the absence of comprehensive, dependable or definitive evidence – these disciplines must evaluate and combine whatever empirical information is available with less concrete theoretical understanding and plausible speculation (Ellison, 2026c) so as to determine: what might have happened; what might be happening; and what might yet happen. The natural and human sciences (Pollack, 2003), forecasting (Petropoulos et al., 2022) and intelligence analysis (Pili, 2024) are all 'disciplines of uncertainty'; and it is within these that the epistemic promise and cognitive impact of AI-enabled tools are likely to be most keenly felt.

Far from disinvesting in the training and skills required to painstakingly synthesise what we think we know to be true, and what we can only assess to be so, there is a pressing need to strengthen the conceptual and analytical training available within such disciplines, and the expertise, competencies and capabilities this sustains. This is especially important when analysts are required to interpret outputs *facilitated* or *generated* by AI-enabled tools without succumbing to overconfidence or hubris. Even optimistic assessments of AI integration acknowledge that its cognitive effects on analysts, and on their analytical capabilities, remain uncertain (Mitchell et al., 2019). Training of this kind will therefore remain crucial for enabling practitioners to evaluate and synthesise empirical evidence, theoretical understanding and plausible speculation from all available sources — including the outputs provided, and the insights enabled, by ongoing advances in AI, and by doing so in partnership with AI (NASEM, 2022). Indeed, ensuring that such training offers a sound understanding of how AI-enabled tools work will be essential for everyone, including those in non-technical roles who nonetheless need to be competent commissioners and consumers, of AI-enabled outputs. Including AI literacy as a core element of this analytical skillset will also provide valuable preparation for any later technical training as developers or operators of AI-enabled tools.

The aim of this article is to make the case for the additional AI literacy training needs of current and future cohorts of professionals and practitioners in one such 'discipline of uncertainty' – all-source intelligence analysis – to ensure its analysts can grasp the opportunities and meet the challenges that accompany AI-enabled outputs and the insights these support. Its arguments draw upon: Deloitte's somewhat ebullient and optimistic report from 2019 (which focused on both the necessity and potential benefits of integrating AI-enabled automation and analysis within intelligence workflows; Mitchell et al., 2019); a recent overview from 2024 (which explored contemporary open source reports describing the use of AI-enabled tools by intelligence agencies around the world; Sfetcu, 2024); and the US National Academies of Science, Engineering and Medicine's more circumspect report from 2022 (on "human-AI teaming"; NASEM, 2022).

These inform both 'why' and 'how' training for all-source intelligence analysts might need to evolve in order to support the 'what' that analysts will require in order to consume (and potentially commission) AI-enabled tools; and thereby work in partnership with AI to strengthen, deepen and broaden the assessments they are able to provide. In other words, training that can successfully provide the foundational understanding necessary to become 'AI aware' as a basis for further training to complete the transition to 'AI-enabled' analysts (MOD, 2025).

Why AI Literacy Matters Most Within 'Disciplines of Uncertainty'

Central to the development of AI literacy training is a clear understanding of the distinction between two broad categories of AI-enabled outputs: those that are AI-*facilitated*; and those that are AI-*generated*. The distinction can be difficult to make because the two are often used in combination, but it rests on whether AI-enabled tools are:

- simply *facilitating* the scale, scope and pace at which known – or at least knowable – patterns and associated outputs can be produced; or
- *generating* outputs that are entirely novel, from hitherto unknown patterns and signals hidden deep within the data.

This distinction matters because both automated data processing and semi-autonomous data analysis contribute to the step change in productivity — and the increasing variety, volume and velocity of outputs that advances in AI technology have brought about. Yet they do not pose the same epistemic or cognitive challenge.

AI-*facilitated* outputs tend to extend or accelerate analytical tasks that humans could in principle perform themselves. AI-*generated* outputs, by contrast, are more likely to support insights that analysts might not independently derive, and therefore demand more conceptually informed scrutiny.

The more prosaic AI-enabled tools – that is, those involving automated processing and analysis – seem likely to have the most immediate *practical* impact in applications where the mechanisms and processes involved are already: well understood; the result of deliberate human design; or subject to 'physical laws' or human control – and are therefore most amenable to reproducible replication and optimization over time and place (Ellison, – 2026c). In contrast, the more insightful AI-enabled tools (particularly those offering semi-autonomous output-generation) are likely to have a much deeper epistemic and cognitive impact wherever these novel techniques are able to resolve hitherto irreducible uncertainties. This is because a distinctive feature of this subset of AI-enabled tools (commonly known as "intelligent rules engines", and including 'machine learning' and 'deep learning'; Mitchell et al., 2019) is their ability to:

- generate credible estimates through interpolation or extrapolation (more simplistically described, and often mis-conceptualized, as 'predictions') of past and future 'known unknowns' and 'unknown unknowns' – i.e. known-but-as-yet-unmeasured or unobserved, and unknown-and-hitherto-unknowable dataset features, respectively (Luft & Ingham, 1955; Rumsfeld, 2002; Davies & Thomson, 2010); and

 doi:10.20944/preprints202606.1383.v1

- do so in ways that emulate (that is, match or *surpass*) the sensory, analytical and cognitive capabilities – and physical endurance – of even the most gifted, experienced and energetic human specialists (and even when using the most advanced *pre-AI* tools; De Cosmo, 2022; Karampelas, 2023).

While all citizens of technologically advanced societies might benefit from becoming competent and judicious commissioners and consumers of AI-enabled outputs, professionals and practitioners working within disciplines of uncertainty are likely to require a higher level of understanding as more expert commissioners and consumers of AI-*generated* outputs. Yet quite how current human scientists, forecasters and intelligence analysts might exploit either type of AI-enabled outputs, while retaining the critical and cognitive skills required to evaluate and interpret them, is likely to prove a challenging ask (Heuer, 1999; De Cosmo, 2022). Not least for 'analog natives', who may struggle to consume, let alone commission, AI-*generated* outputs (Karampelas, 2023); but also for tech-savvy 'digital natives', 'AI immigrants' and 'AI natives' (Gillespie et al., 2023; Tyson, 2024), whose familiarity with digital technology and computational analytics may make them more trusting and less critical of AI-*generated* outputs (Panthalookaran, 2024; Dolman, 2024).

Knowing What Questions AI Can Usefully Address

All three of the principal sources examined for this article (Mitchell et al., 2019; NASEM, 2022; Sfetcu, 2024) make clear that the versatility of programmable, automated and semi-autonomous computational procedures now encompassed by the term 'artificial intelligence' has meant that these tools have become critical enablers and force multipliers within two key steps of the intelligence cycle: information collection; and information processing and analysis (DCDC, 2024). Despite this, the operational rollout of this technology continues to require extensive preparation, planning and experimentation, as well as post-implementation monitoring so as to adapt existing doctrine in support of its alignment with novel TTPs (tactics, techniques and procedures) and SOPs (standard operating procedures). While a substantial proportion of these developments occur covertly, sufficient open-source reporting exists to reveal that innovative digital and computational technologies providing AI-enabled outputs (and the insights these support) are proliferating across intelligence agencies around the world (Sfetcu, 2024). They have also undergone well-publicized within-theatre testing, innovation and development during the ongoing conflicts in the Middle East (Abraham, 2024) and Ukraine (Bondar, 2025; though arguably these developments date back at least as far as the Vietnam War; Belcher, 2019; Young, 2021).

This means that – as Deloitte's report (Mitchell et al., 2019) would have it – "AI is not coming to intelligence work; it is already here." Indeed, most all-source intelligence analysts are already consuming information whose collection and/or processing has been *facilitated* by automated AI tools; and a growing number are also beginning to encounter outputs *generated* by tools operating semi-autonomously. What many analysts still lack, however, is the conceptual understanding needed to judge: what these tools are actually doing; what sorts of outputs they produce; and what kinds of tasks they are (and are not) well suited to address.

Knowing when, how, for what purposes, and on what processes and datasets these tools might be most effectively used – not least for novel or uncommon tasks that are specific to PIRs relevant to particular roles, teams, functions or disciplines – clearly constitutes a necessary enabler for many analysts, both now and in the future. At the same time, one of the most persuasive arguments presented in Deloitte's report (Mitchell et al., 2019) for the widespread adoption of AI-enabled automation and analysis systems within intelligence is the seemingly inexorable rise in the variety, volume and velocity of digitized information available for intelligence processing and analysis. Yet any efficiencies gained through such automation are unlikely to be realized without substantial investment in the knowledge, training and organizational support required to implement, maintain and evaluate these systems properly (Mitchell et al., 2019). If intelligence analysts are to effectively exploit all available sources of digitized information, then attaining the skills required to select or commission appropriate AI-enabled tools to automate information collection, processing and analysis procedures, and to judge what they (and others using these techniques) have done and found, is a pressing operational priority. This is especially important because a variable, and at times substantial, proportion of intelligence requirements concern singular phenomena, entities or processes that are historically unprecedented, or so spatio-temporally contingent as to be essentially (and quite literally) unique. In such cases, the absence of comparable or sufficiently relevant information from other places or other times may itself limit the utility of AI-enabled processing or analysis techniques. This means that in these instances, analysts will *still* need their traditional cognitive and inferential training to generate robust assessments in the absence of certainty.

Evaluating AI-*Generated* and AI-*Facilitated* Outputs

An important consequence of the differences outlined above is that the training required to evaluate AI-*generated* outputs is likely to prove substantially more demanding than that required to evaluate AI-*facilitated* ones. In part, this is because AI-*generated* outputs involve semi-autonomous analytical techniques that would often be too challenging, resource-intensive and time-consuming (or simply impracticable) for most human analysts to implement independently – and it is therefore not possible to compare findings generated by human analysts with those produced by these AI-enabled tools. But it is also because the complexity of the optimized algorithms such techniques construct – so as to generate their interpolative or extrapolative quasi-'predictive estimates' of hitherto unmeasured or unknown dataset features – makes their outputs difficult, and at times impossible, to interpret meaningfully (albeit as anything other than arithmetic and algebraic solutions to the tasks and questions at hand; and within the constraints imposed by the dataset[s] involved).

For these reasons, the evaluation of AI-*generated* outputs is likely to require substantial additional AI literacy training beyond that required: either to identify what questions or problems such tools might usefully address; or to judiciously commission or consume outputs from the most appropriate analytical tools for the task concerned. This is because the validity, generalizability and utility of the outputs these techniques generate depend on a number of critical dependencies determined by both the dataset(s) available and the knowledge, expertise and domain awareness of the analysts involved (NASEM, 2022). They are also subject to important limitations affecting both: which inferences *can* be reliably drawn (and those that *cannot*); and their subsequent practical utility. None of these dependencies or limitations are necessarily self-evident to non-specialist consumers of AI-*generated* outputs; and even proficient developers and operators of AI-enabled tools commonly overlook them (NASEM, 2022). Yet these are precisely the considerations that analysts need to be able to recognize and evaluate – and, for those willing to engage with them directly, these have been summarized in an accessible conceptual and analytical primer as a companion piece to this article (Ellison, – 2026b).

By contrast, the automation of information collection and processing using AI-enabled tools is often far more prosaic. This is particularly the case where such automation is applied to tasks that analysts would otherwise have completed (or been able to complete) themselves, had they the time and capacity to do so. In such cases, analysts are not only better placed to recognize, evaluate and comprehend the processes performed and outputs achieved, but are also often in a position to validate at least some of these outputs against those produced 'manually' by human analysts.

It therefore seems likely that – beyond the additional training required to select and apply suitable AI-enabled tools to automate existing (but time- and resource-intensive) collection, processing and analytical tasks – all-source intelligence analysts may not require substantial further training in order to evaluate the meaning and utility of AI-*facilitated* outputs. In many such cases, well-designed collection and processing protocols or SOPs that remind, or require, analysts to validate automated outputs against a manually produced subset of these may suffice. But a greater additional training burden exists in relation to AI-*generated* outputs, whose meaning, validity and utility depend on conceptual and analytical considerations that are less visible, less intuitive, and less readily checked against the practices (and outputs) of human analysts.

 doi:10.20944/preprints202606.1383.v1

The Pressing Case for AI Literacy Training for All-Source Intelligence Analysts

A compelling case can therefore be made for providing all-source intelligence analysts with additional AI literacy training in the knowledge and skills required to:

- select suitable AI-based tools to automate the collection and/or processing of information prior to analysis; and
- analyze, evaluate and interpret the outputs that automated and semi-autonomous AI tools can facilitate and generate, respectively.

This case is based not only on the efficiencies that might be achieved through the automation of information collection, processing and analysis within the intelligence cycle (Mitchell et al., 2019); but also on the distinctive epistemic and cognitive contributions that AI-*generated* outputs may make to all-source intelligence analysis. While “AI is not the solution to every problem” (Mitchell et al., 2019), whether in intelligence analysis or elsewhere; and while not all information sources and datasets are relevant, large, diverse or consistent enough to warrant or support the use of AI-based automation or analysis techniques – the inexorable rise in the variety, volume and velocity of information available to intelligence analysts makes the adoption of AI-based tools to automate processing and analysis an opportunity analysts should be equipped to exploit. Their adoption will also be essential if analysts are to offer the more comprehensive analyses and assessments in near real-time that the proliferation of digitized information sources should permit through human-AI teaming.

The training implications of these two broad categories of AI-enabled output are, however, distinct. For AI-*facilitated* outputs – where automation accelerates processes that analysts could in principle perform themselves – the AI literacy training required is relatively modest. This chiefly consists of the ability to select outputs produced by the most appropriate tools, and to validate automated outputs against those that are manually produced (albeit whenever possible). For AI-*generated* outputs – where semi-autonomous techniques produce outputs beyond the capabilities of analysts alone – the training required is substantially more demanding. In this instance the training must equip analysts to evaluate:

- whether the tool used was suited to the question at hand;
- whether the dataset on which the tool was trained is comparable to the data- and dataset-generating mechanisms available within the context at hand; and
- what the outputs actually represent, under what conditions they can be trusted, and how such trust ought to be evaluated.

These are the inherent dependencies and limitations of AI-*generated* outputs that non-specialist analysts are currently least equipped to assess – and which AI literacy training must therefore prioritize. As mentioned earlier, the conceptual and analytical foundations that this AI literacy training must contain – including the three critical dependencies and five substantive limitations of AI-*generated* ‘predictive estimates’ – have been summarized in more detail in a companion piece to this article (Ellison, –2026b). This has been crafted with a non-specialist audience in mind, who should find it accessible without prior technical knowledge. Together, the two articles argue that the transition from ‘AI-aware’ to fully ‘AI-enabled’ intelligence analysis is achievable – but only if the AI literacy training proposed here is treated as a professional baseline for all analysts; not simply a specialism reserved for the technically inclined.

References

Abraham, Y. (2024, April 3). ‘Lavender’: The AI machine directing Israel’s bombing spree in Gaza. *+972 Magazine [and Local Call].* web.archive.org

Belcher, O. (2019). Sensing, territory, population: Computation, embodied sensors, and hamlet control in the Vietnam War. *Security Dialogue, 50*(5), 416–436. doi.org

Bondar, K. (2025, March 6). Ukraine’s future vision and current capabilities for waging AI-enabled autonomous warfare. *Center for Strategic and International Studies.* web.archive.org

Chantel, J. (2023, April 25). Smartphone history: The timeline of a modern marvel. *Textedly Blog.* web.archive.org

Davies, A., & Thomson, M. (2010). *Known unknowns: Uncertainty about the future of the Asia-Pacific* (Special Report No. 35). Australian Strategic Policy Institute.

DCDC (Development, Concepts and Doctrine Centre). (2024). *Intelligence, counter-intelligence and security support to joint operations* [Bite-sized doctrine]. UK Ministry of Defence. web.archive.org

De Cosmo, L. (2022, July 12). Google engineer claims AI chatbot is sentient: Why that matters. *Scientific American.* web.archive.org

Dolman, J. (2024, September 11). Naïve or native? Do digital natives really understand their world? *The AI English Teacher.* web.archive.org

Ellison, G. (2026a). What 'typical insights' might 'core statistical analytical techniques'provide to intelligence analysts? *Preprints*, 2026060055. doi.org (Under review at: *Journal of Applied Mathematics, Sciences, and Technology for National Defense*)

Ellison G. (2026b). Understanding AI-enabled outputs: A conceptual and analytical primer for all-source intelligence analysts. *Preprints*, 2026060055. doi.org (Under review at: *Journal of Artificial Intelligence Research and Innovation.*)

Ellison, G. (2026c). Might 'uncertain military leadership' minimize the risks and optimize the opportunities of strategic uncertainty? *Contemporary Military Challenges*. in press.

Gillespie, N., Lockey, S., Curtis, C., Pool, J., & Akbari, A. (2023). *Trust in artificial intelligence: A global study.* The University of Queensland and KPMG Australia. doi.org

Heuer, R. J. (1999). *Psychology of intelligence analysis.* Center for the Study of Intelligence, Central Intelligence Agency.

Karampelas, A. (2023, August 25). The emergence of AI-natives. *Medium.* web.archive.org

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

Luft, J., & Ingham, H. (1955, August). The Johari window: A graphic model of interpersonal awareness. *Proceedings of the Western Training Laboratory in Group Development*, *1*(1), 1-7.

Markelius, A., Wright, C., Kuiper, J., Delille, N., & Kuo, Y.-T. (2024). The mechanisms of AI hype and its planetary and social costs. *AI Ethics*, *4*(3), 727–742. doi.org

Marr, B. (2024, April 15). AI: Overhyped fantasy or truly the next industrial revolution? *Forbes.* web.archive.org

MOD (UK Ministry of Defence). (2022, June). *Defence artificial intelligence strategy.* web.archive.org

Moll, I. (2021). The myth of the fourth industrial revolution. *Theoria*, *68*(167), 1–38. doi.org

Mitchell, K., Mariani, J., Routh, A., Keyal, A., & Mirkow, A. (2019, December 11). The future of intelligence analysis: A task-level view of the impact of artificial intelligence on intel analysis. *Deloitte Center for Government Insights.* web.archive.org

NASEM (National Academies of Sciences, Engineering, and Medicine). (2022). *Human-AI teaming: State-of-the-art and research needs.* The National Academies Press. doi.org

Panthalookaran, V. (2024). A revised taxonomy of educational objectives for AI-natives: Reimagining thinking skills in the age of AI. *Higher Education for the Future*, *12*(1), 51–64. doi.org

Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Ben Taieb, S., Bergmeir, C., Bessa, R. J., Bijak, K., Boylan, J. E., Browell, J., Carnevale, C., Castle, J. L., Cirillo, P., Clements, M. P., Cordeiro, C., Cyrino Oliveira, F. L., De Baets, S., Dokumentov, A., … Ziel, F. (2022). Forecasting: Theory and practice. *International Journal of Forecasting*, *38*(3), 705–871. doi.org

Pili, G. (2024). A new theory of surprise: Unraveling the logic of uncertainty and knowledge. *Intelligence and National Security*, *39*(4), 695–708. doi.org

Pollack, H. N. (2003). *Uncertain science… uncertain world.* Cambridge University Press.

Rumsfeld D. (2002, February 12) DoD News Briefing. United States Department of Defense 2002; 12 Feb. web.archive.org

Sanguinetti, P. (2025, January 22). Why the term 'artificial intelligence' is misleading. *IE Insights.* web.archive.org

Sfetcu, N. (2024). *Artificial intelligence in intelligence agencies, defense and national security.* MultiMedia Publishing House. doi.org

 doi:10.20944/preprints202606.1383.v1

Thierry, G. (2025, April 14). We need to stop pretending AI is intelligent – here's how. *The Conversation*. web.archive.org

Torres Torres, M., Hart, G., & Emery, T. (2019, July 23). *The Dstl biscuit book: Artificial intelligence, data science and (mostly) machine learning* (1st ed., v1_2). UK Ministry of Defence. web.archive.org

Tyson, J. (2024, October 17). Digital natives are now AI immigrants. *UX Magazine.* web.archive.org

Volti, R., & Croissant, J. (2024). *Society and technological change* (9th ed.). Waveland Press.

West, D. M., & Allen, J. R. (2018, April 24). How artificial intelligence is transforming the world. *Brookings.* web.archive.org

Young, D. (2021). Sensors, interpreters, analysts: Operating the 'electronic barrier' during the Vietnam War. *Digital War, 2*(1), 51–63. doi.org

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.