

## The End of Human Intelligence Analysis—Better Start Preparing

*John F. Galascione*

---

**As a follow-on thought experiment, it might be useful to look beyond the future Gartin visualizes to a more distant and very different future.**

*Editor's Note: This article is a slightly updated version of the original published internally in 2020. It is reprised here as a complement to Dennis Gleeson's article "Artificial Intelligence and Analysis," beginning on page 11 of this edition.*

My reaction to reading Joseph Gartin's excellent "The Future of Analysis" (*Studies* 62, no. 2, June 2019) was that it described a step on the way to the future of analysis. It envisions human analysts using enhanced computer-based tools to produce finished products and insights for delivery to customers. This sounds familiar to me, as if the future will be just a more idealized version of the present—the same basic job but with better and more reliable tools. One can only hope.

As a follow-on thought experiment, however, it might be useful to look beyond the future Gartin visualizes to a more distant and very different future. In the spirit of certain structured analysis techniques, we can then step back to see what actions can and should be taken now and in the near future. The intention is not to debate the timeline of this more distant future but instead to assume it is approaching and to outline its implications. Some key assumptions about this more distant future are needed:

- Assume the paradigm that people will use computers still exists (and

not the other way around). For a distant future, this largely involves hope, but if such an assumption is wrong, then the rest of this discussion doesn't much matter.

- Assume computer-based systems will become—and from then on will always be—better than human beings at processing and analyzing information. By "better" I mean more complete, more accurate, faster, less biased—all the hallmarks we recognize from our own analytical training and all the qualities to which we analysts aspire.
- Assume that our mission persists, which in its basic form involves anticipating and warning to enable senior officials in our government to respond to and/or shape the course of events. This belief in mission is tied, somewhat, to the notion that we (human beings) can use what we do here to control or direct the course of history. A key assumption here is that this will remain the case into the not-so-distant future.

The 2030 that Gartin describes—a time in which the intelligence analysis-production-and-delivery system will be a more-or-less tweaked version of the current system—is different from what I would argue will be an inevitable future in which computer-based systems will be able to perform these functions:

---

The views, opinions, and findings of the author expressed in this article should not be construed as asserting or implying US government endorsement of its factual statements and interpretations or representing the official positions of any component of the United States government.

---

---

***In this future, which is probably not as distant as we think, it seems as if an argument would need to be made for how and why humans would still be included at all in the intelligence analysis profession.***

---

- track and synthesize information faster and better than any human or group of humans.
- better anticipate customer interests and needs by tracking an overall larger dataset of government and international activities than any analysts could and ever would follow.
- produce adequately “written” finished products that lay out assessments clearly, use language consistently from product to product, are relatively immune to review, style, and editing fads, and are clear about how each aspect of those assessments is changing and may be vulnerable to change.
- anticipate future intelligence needs, generate requirements, and more accurately and completely track the validity or suspiciousness of reporting streams.

The systems will also be truly multiculturally diverse, having ready access to worldwide patterns of thought, cultural histories, and priority systems and be able to apply them specifically to localized

intelligence-related assessments and they will be less prone to experience bias, confirmation bias, and cognitive dissonance.

---

***What’s a Person to Do?***

In this future, which is probably not as distant as we think, it seems as if an argument would need to be made for how and why humans would still be included at all in the intelligence analysis profession. Some arguments for continued human involvement arise automatically from traditional roles and functions of intelligence analysts, but these arguments are not without problems.

***Briefer Version 2.0***

If we assume that our key consumers will continue to be human beings, then perhaps analyst-like humans will be needed to act as interfaces between computer-based assessment-generation systems and those human consumers. This is unlikely to be practical, however, because the deep-learning-based systems will pump out assessments likely to defy the explanation of humans, even those working alongside them. Could

the programmers of AlphaGo<sup>a</sup> explain the infamous “move 37” in one of its historic games against one of the world’s best Go players? Not really, and neither could the world’s best Go player. Similarly, humans will almost certainly be inept at explaining all of the reasons why a certain assessment is being made, being made now, or is changing from its previous form.<sup>b</sup> This version of “briefer,” then, devolves to mechanically transferring information and attempting to translate between the customers and the algorithms. Or worse, it becomes the job of analyzing the computer-generated analysis.

***Data Scientist Version 2.0***

Some would also argue, as part of the assumption that humans will still be key actors in this future, that people will be needed to develop and tweak the algorithms that generate these assessments. Even this idea is becoming increasingly falsified, however, by how deep-learning systems are being developed to use only limited, if any, human insights from which they can quickly “outthink” humans. So far, most progress is on domain-specific tasks, as in the case of AlphaGo, but even this is changing as self-learning systems have been developed to conduct more generalized “learning” with little-to-no input from humans.<sup>c</sup> In terms of the

---

a. AlphaGo, a computer system developed by DeepMind (now part of Google), used a combination of human data and self-play to become proficient in the ancient game of Go, and then beat some of the world’s best players. DeepMind has since developed variations of AlphaGo—AlphaGoZero and AlphaZero—that developed domain-specific expertise solely through self-play and with no actual human input or human-supplied data.

b. In an unrelated but similar example, researchers in 2018 developed an AI machine learning tool to detect depression (in people) based on speech patterns during interview responses—without additional context. This research led to an accurate prediction tool but the opacity of its conclusions was seen to be a drawback—basically the tool was capable of diagnosing a person with depression but did so in a way that was too opaque to the psychiatrists themselves to be completely useful.

c. DeepMind, the same group behind the development of AlphaGo, had earlier developed a self-learning system that trained on a wide variety of Atari video games using only the pixels on the screen and game scores, from which it was able to develop the capability to play multiple games at professional human levels. See V. Minh, K. Kavakcuoglu, D. Silver, et al., “Human Level Control Through Deep Reinforcement Learning,” *Nature* 518 (2015): 529–33.

impact of these developments on the future of intelligence analysis, it's difficult to imagine how humans will be expected to make substantive changes to algorithms they were not involved in developing in the first place. The idea that humans will be somehow needed to "assist" computers seems increasingly more laughable.

Propagating machine-learning and computing trends makes it seem as if 20 to 30 years from now—and possibly sooner—the idea of humans doing intelligence analysis professionally will be of historical significance only. A debate about the future of AI and its development timelines is beyond the scope of this article, but the truth is that some advances in domain-specific AI and deep-learning techniques have surprised many, and experts have been somewhat poor in their ability to predict the future of AI-related progress.

Although skeptics continue to assess that "generalized" artificial intelligence is still a distant development, it's not clear that this argument falsifies the future of the analytical profession proposed here. The progress to date in domain-specific machine intelligence makes such a future seemingly inevitable, especially given that intelligence analysis is a domain-specific task. Just as no single person would ever have the generalized intelligence and expertise required to do all analytical jobs, it's not difficult to imagine a suite of separate artificial intelligent systems that handle different analytical "accounts."

---

### ***Between Now and Then . . .***

In keeping with the spirit of structured analytic techniques, we should

be asking what we can do to prepare for this future. Interestingly, I think some of the possible steps we can take and issues we'll confront will mirror those our predecessors faced. This time, however, we'll be tasked with ensuring that our machine-based successors are able to do this job at least as well as we think we do it now.

On the October 8, 2023, episode of the CBS news program *60 Minutes*, AI guru and pioneer Geoffrey Hinton underscored the extent to which AI tools may well leave humans out of the loop. While noting that the algorithms themselves are designed by humans, once these algorithms interact with data they produce "complicated neural networks," and the developers (much less the users) "don't really understand how they do what they do." Unpacking the AI-derived analytical conclusion may one day be beyond the grasp of human intelligence analysts.

### ***Technical Analysis Goes First***

Technical intelligence analysis, specifically weapons-related analysis, is likely to be the first to be transitioned into a completely humanless process. Technical analysis is about numbers, linear cause-and-effect relationships, predictability, and tracking of R&D activities (much of which is already computer based). Technical expertise—and by extension technical intelligence analysis—has always been severely stovepiped. Nuclear analysts view reporting through a strictly nuclear-technology filter, biological weapons analysts see developments as they fit into that specific technology, and so forth. This is how human expertise has worked for generations and there's no good reason to believe that the future includes a

different type of more generalized human technical expert.

Initially, technical analysis will be most amenable to domain-specific AI techniques, but then—as in the case of other AI-related advances—the "domain" is likely to be expandable to include subjects across multiple technical fields. Having a much deeper background in a much wider range of technologies, the computer-based technical "DA analyst" of the future is less likely to automatically view developments as fitting into a narrow field in the way most human analysts would. This machine-based analyst is also more likely to identify connections in technologies, materials, and people across multiple fields—something human analysts are practically incapable of doing.

As a constructive step toward this future, technical analysts now and for the next decade should be establishing robust knowledge-management efforts and standards with which to more efficiently enable ingestion by their machine-based successors. It sounds as if we're handing the bank robber the combination to the vault, but in the future we're imagining here, such an approach makes more sense than the alternative. We speak of our legacy, and this has always meant passing along knowledge to the next generation of analysts. Wouldn't this still be important even though the next generation of analysts are not people? We also speak of mission focus and we assume the mission will persist. Wouldn't this mean that it's actually our duty to ensure that the next generation of analysts—human or otherwise—serve that mission well?

---

***Finally, analysts of all stripes probably should be finding ways to become comfortable with the inevitable and be developing strategies for transitioning from what is now an almost completely human process to what is likely to become a nearly completely humanless process.***

---

***Documenting How We Played the Game***

A key aspect of teaching and training the next generation of analysts has always involved documenting and explaining our past successes, failures, and everything in between. If we truly care about the health of our enduring mission after we're gone, then documenting what is right and wrong about our analysis should become a large, or larger, part of our job over the next decade. AlphaGo learned partly from about 100,000 human-played games of Go, which served as the foundation upon which it built even better and more perplexing strategies and learned how to teach itself. The same groundwork needs to be put in place for the DA, even if it's almost certainly more difficult to characterize and document our results than the moves of a tightly constrained game whose rules never change.

***Be the Solution, Not the Resistance***

Finally, analysts of all stripes probably should be finding ways to become comfortable with the inevitable and be developing strategies for transitioning from what is now an almost completely human process to what is likely to become a nearly completely humanless process. We are not the first industry to face this, and we have the advantage of having studied our own jobs and our own history extensively. The self-reflection and recursive self-analysis that

have long been hallmarks of CIA analysis probably should be brought to bear on establishing a graceful transition from humancentric intelligence analysts. Indeed, CIA and other intelligence agencies are in the business of delivering finished analysis; they are not necessarily in the business of employing analysts.

As present-day intelligence analysts staring down a very different future, we are also responsible for ensuring that the transition to this future is relatively painless for everyone. As with any revolutionary change, resistance is inevitable and false arguments will abound. Replacing all automobiles with the current self-driving technology—even right now when this technology is arguably only in its infancy—probably would reduce the number of accidents overall. Nonetheless, every accident involving a self-driving car will enable skeptics to compare machine-based results to an absolute (impossible) standard instead of seeing how they are actually better than humans despite being imperfect. It's not difficult to see that such false comparisons will continue to be employed to undermine assessments of the actual progress and inevitability of advances, including in the field of intelligence analysis. A nonsensical result from early machine-based intelligence analysts would be met with a similar response, even though such events only help those focusing on resisting

this inevitable future—slowing, but not preventing, its arrival.

---

***All Is Not Lost (Yet)***

The future posited here is one in which many if not all the people who do this job today are likely to be replaced. One could even imagine a time when the term “intelligence” would refer implicitly to the machine-based type and “human intelligence” (HUMINT!?) would refer to our particular human flavor of the stuff. Nonetheless, to the extent that we remain committed to the mission of the Intelligence Community, we should be transitioning—mechanically, mentally, and emotionally—into enablers and facilitators of this inevitable future—if not its most ardent supporters, developers, and champions. One could argue that this is actually our job.

Our job also involves assessing what our professional futures as intelligence analysts will be, regardless of what that future might look like. We can argue about how quickly this future arrives, but we probably shouldn't argue about whether some form of it will arrive. The machine-learning-related advances to date suggest that those who disagree with the inevitability of this future have the responsibility of generating an alternative future that is somehow more likely than this one.

It would be ironic, would it not, for one of the last major (human) intelligence failures to be that we failed to identify our own extinction by failing to predict the future of our own profession?



*The author:* John Galascione is a 35-year veteran in the intelligence analysis field, having focused his entire career on analysis of foreign nuclear weapon-related development and testing activity.