Putting AI into Air: What is Artificial Intelligence and What it Might Mean for the Air Environment

By Squadron Leader James Waller and Dr Phillip Morgan

Biography: Squadron Leader James Waller is an Engineer (Communications Electronics). He holds an MEng degree in Avionic Systems Engineering from the University of Bristol where he specialised in computational intelligence, pervasive and ubiquitous computing and human integration within semi-autonomous aircraft systems. He has also recently completed an MSc in Cyber Defence and Information Assurance with Cranfield University.

Biography: Dr Phillip Morgan is an Associate Professor in Cognitive and Human Factors Psychology, Director of Human Factors Excellence Research Group, and, Theme Leader within the Transport Futures Research Network at Cardiff University. Dr Morgan is Technical Lead for Cyber Psychology and Human Factors at Airbus. He holds a BSc (Hons) in Psychology, a PG Diploma in Social Science Research Methods, and PhD in Cognitive Science (supervised by Prof Dylan M Jones OBE), all from Cardiff University, alongside a PGCHE (University of Wales).

Abstract: Artificial Intelligence (AI) is heralded by some as the next Revolution in Military Affairs (RMA), a technology with the potential to offer decisive strategic advantage and revolutionise military tactics and force structures. Development of AI within the commercial sector continues at a galloping and, in some ways, alarming pace, and the academic community[1] has begun to consider the role AI could play within Air Power. This article begins with a brief history of AI, discussing some recent developments in the field and highlighting why there is currently so much excitement around the subject. We also discuss how our understanding of human sensory and cognitive processing limitations advocate a strong need for AI in the form of augmentation and decision support systems in several settings, whilst highlighting unique human abilities to adapt to and process information in different ways (that an AI currently cannot achieve). Based upon the areas where we believe AI will offer an edge to a human operator in the future, we propose a framework for assessing the key attributes of AI and discuss some of the implications for air warfare in the future.

Introduction

We can only see a short distance ahead, but we can see plenty there that needs to be done. [2]

Alan Turing wrote those words nearly 70 years ago, and yet they remain as prescient in 2019 as they did in 1949, perhaps more so. Within the subject of Computer Science, Turing is renowned for his eponymous 'Turing Test', where he first famously asked the question 'Can Machines Think?' and founded the discipline that is known today as Artificial Intelligence (AI).

Turing is also famed for his work at Britain's Second World War code-breaking centre at Bletchley Park and the contribution it made to the outcome of the War. Many consider that AI will have a comparable revolutionary effect on the future of warfighting, so much so that US policymakers have previously decreed that AI will form the bedrock of a '3rd offset'[3] strategy, insisting that AI is a technology with the potential to offer such decisive strategic advantage to its possessor that it could 'offset' the balance of power.[4] This momentum appears to be building: Artificial Intelligence appears over 60 times within the UK Ministry of Defence's analysis of global strategic trends[5] and is firmly considered one of 2019's key 'tech trends'.[6] It goes without saying that it is highly relevant to air power too: the US Air Force has announced that it will invest $100M in AI technology.[7]

Technology has always had a profound impact on the way in which wars are fought. However, it remains a truism that innovation is not confined to the technology alone, and current MOD thinking is clear that the innovation should not just be confined to the technology itself, but in how emerging technology is employed on the battlefield.[8] This is an idea neatly summarised within the introduction of the UK MOD Joint Concept Note[9] on Human/Machine Teaming:

> The winner of the robotics revolution will not be who develops this technology first or even who has the best technology, but who figures out how to best use it. [10]

'Figuring out how best to use AI' is at the heart of this article which is structured in two sections.

The first section will start with a discussion on the concept of AI, with a brief look at some key ideas and recent developments to establish what is and isn't hyperbole, but more importantly it will explain why AI is so important now. Using the examples of computer mastery of chess and language translation the article will illustrate how much the field of AI has achieved and how it is likely to be employed within the air domain. This section will conclude by recognising that AI is likely to complement human activity posing the question: 'What will a human be able to do better in the air war of the future?' The second section of the article will then examine what this means. By considering the likely limitations of the human brain and cognition, we will propose that AI possesses three core, interlinked, attributes which make its employment within air power desirable. These attributes are speed of cognition, ability to process large volumes of information, and ability to access physical spaces which are impossible for a human. The article will examine each of these attributes in turn and highlight possible implications for the air environment.

AI: A brief history and why the excitement now?

It is worth stating from the outset that this article is not about AI taking over the world, nor does it suggest that computers could immediately gain sentience and possess a decisive strategic advantage over humanity, a proposal which actually can be traced back to Alan Turing.[11] The phrase 'intelligence explosion'[12] has been coined to reflect Turing's original idea that once computers achieve a certain level of intelligence they will simply be able to redesign themselves iteratively, quickly surpassing any level of intelligence achievable by a human, referred to as Artificial General Intelligence (AGI).[13] There is credible academic literature which explores this fascinating idea,[14] however this article will not do so for two fundamental reasons. First, if humanity does develop AGI, it would be such a significant development that it could revolutionise virtually every aspect of life, not just warfighting. Debate on how this could affect air power would probably be low on humanity's priority list! Secondly, and more importantly, however, the academic community is divided on how 'soon' we will achieve AGI, though the broad consensus appears to be that it will take at least decades rather than years.[15] Instead this article focuses on the types of cognitive tasks within the air environment where AI is likely to perform as well as (or better than) a human in the near future.

The other significant area this article will not explore in detail is the ethical dimension of employing AI in warfighting, though it will recognise that this is a hugely important subject.[16] The reasons for this are twofold: first, like many novel technologies, there will always be the less scrupulous adversary who will employ the technology without regard to ethics. Secondly, many of the developments in AI are being pioneered by sectors outside of the defence industry; we will increasingly see AI technology being developed for one purpose not associated with Defence which will subsequently be employed for warfighting.[17] Therefore, regardless of ethical stance, it is important to consider the possibilities that AI could bring to air power. This article aims to explore the potential for AI to be employed in the air environment, agnostic of the ethical debate, though by doing so hopes to better inform it.

So, what can AI do now and what is it likely to be able to do within the next 10 years? To answer these questions, it is worth considering that, over the course of its short history (and like many novel technologies), the field of AI has been beset with 'overpromising and underdelivering' with several false starts, leading to 'AI Winters'.[18] A general theme is that computers are becoming increasingly proficient at activities provided they remain 'narrow', compared with humans who, despite numerous limitations (see examples later in this article), can process wider sets of tasks and are far more adaptable. For example, there have been significant developments in creating a machine which is comparable to a human in language translation or chess; however, at this time, there is not a machine that can do as well as a human at both. Humanity still retains 'general' intelligence i.e. a human can master both chess and language translation[19] or indeed any combination of other complex tasks. The crucial advantage a human possesses is that they can more readily be adapted to any complex tasks requiring a high level of cognition without requiring the level of extensive rebuilding and reprogramming in the way a computer currently requires.

However, this crucial distinction is changing, and a fundamental reason for this is that advances in hardware are now allowing AI programmers to experiment with AI techniques which previously were impractical due to hardware limitations. In spite of several research attempts to produce AI-capable hardware in the last century, it was ironically the video game market which aided this breakthrough.[20] During the early part of this millennium, market forces reduced the cost of multi-processor graphics cards known as Graphics Processing Units (GPUs) to support gaming. However, as a fringe benefit, GPUs are also well-suited to implementing some particular advanced types of Artificial Neural Networks, which had been previously just theoretical possibilities.[21]

An Artificial Neural Network mimics the human brain and, in its most basic form, consists of several processing elements (neurons) which have links between each other (synapses). The route from input through to output is referred to as a 'path'. When solving a problem, a neural network will process input data through possible 'paths' and if this produces a desirable output then the neural network will favour that path in the future, a system known as 'credit assignment'. Taking the analogy of the human brain, this could be considered learning.

Advances in machine learning are well demonstrated by the computer mastery of chess. A significant milestone was achieved in 1997 when IBM's Deep Blue beat the world champion Gary Kasparov,[22] considered by many to be a major development of AI. Some even considered it proof that a machine could pass the Turing Test.[23] Deep Blue was able to examine 200 million moves in a second, so, for each move, was comparing the positions on the board with a repository containing virtually every chess game and strategy played by every grandmaster through history. Playing Deep Blue at chess was akin to playing a human who, at the start of their turn, could stop time, and then take as long as they needed to consider their next move against the collective wealth of humanity's chess playing experience. Yet even some of the developers of Deep Blue did not consider their machine to be 'intelligent'.[24] Deep Blue was an example of a 'brute force' approach to AI, also known as 'Good Old-Fashioned AI' (GOFAI). The computer was not *thinking* about its next move; it was simply consulting an (admittedly vast) checklist for each move it made. Arguably the 'intelligent thought' was still being applied by the human programmers of Deep Blue, some of whom were chess grandmasters themselves.[25] Essentially Deep Blue had been programmed in advance with a playbook that could anticipate every move that its opponent could make.

Therefore, in this context, a much more significant milestone was reached when, in 2017 Google's Deepmind programme beat Stockfish, the 'reigning champion computer' at chess.[26]

Stockfish used AI technology which was fundamentally the same as that which Deep Blue used to beat Kasparov, a GOFAI style approach. However, Google's Deepmind had no prior knowledge of chess. Using its advanced neural network, it played millions of games against itself to 'learn' to play the game to expert standard in just 4 hours! It then proceeded to beat Stockfish. This was hugely significant. Up to that point, a computer could only beat a human by recalling so much data, and simulating permutations so quickly, that no real cognition was taking place. Google's Deepmind proved that, if a clear end state can be defined (in this case winning at chess), a computer can independently and quickly develop knowledge in a previously unknown subject area which can 'surpass not only one human, **but the collective wisdom of humanity**. In the case of this example, hundreds of years of human experience in playing chess was outmatched by Deepmind in four hours!'.[27]

We are now increasingly seeing the results of such advanced AI techniques in the field of language translation, and the developments in this field have introduced an additional nuance, which is that, increasingly, humans and computers will offer different and complementary abilities to solve complex problems; neither will be 'better' than the other. It is fair to say that currently *Google Translate* is more proficient than the average human with no linguistic ability at translating between languages. This is especially true if we consider a human versus a machine in a contest of learning two languages from scratch and then translating between them. Where many humans still retain the edge is when complex language and nuance must be interpreted perfectly by expertly trained analysts. However, the adage: 'good enough today rather than perfect tomorrow' often applies in warfare,[28] in this example AI could provide a complementary aid to linguists to provide a quick 'sense' of what is being said in an unfamiliar language, perhaps where slang or unknown dialect is used, potentially spotting patterns and linkages that a human mind may miss.

In a similar manner to learning chess, the AI agent behind Google Translate has 'learnt' to solve the problem of language translation by developing its own 'language'.[29] This 'language' is unintelligible to a human and is essentially a computational representation of the fundamental syntax of human speech, agnostic of native language. Google Translate's AI engine has done so largely independently of human input, and in doing so has devised a unique and efficient method of translating languages.

It is both fascinating and alarming that AI is developing to the point where it can solve problems completely independently and with apparently unique and innovative ideas, some of which are beyond the reach of a human mind. This is especially the case with a variant of Neural Networks called Genetic Algorithms, which solve problems by 'competing' different solutions to a problem in a manner reminiscent of natural selection. Genetic algorithms can produce creative solutions to a problem which may not be immediately apparent or even unimaginable. An example of this is 'Hackrod': a racing car chassis designed by AI which combines the optimum balance of tensile strength, aerodynamic performance and weight, acknowledged to be a better design than any professionally trained human engineer could have created.[30]

The above examples show that AI clearly is becoming more proficient at tasks which used to be the sole domain of humans, either as a direct replacement or in providing complementary abilities. Therefore, we must start asking: What are those things that 'only a human can do?', but set against the context of the previously mentioned 'false starts' within the field of AI, and the tendency for pessimists in AI technology to move the goalpost of what constitutes a breakthrough. Forty years ago, the academic community felt that once a computer could beat a human at chess, we would have achieved AI.[31] However, this has now easily been surpassed. Similarly, language translation was seen as a bulwark of human ability, yet arguably AI is now superior in some regards. Therefore, to understand how AI will be used in the air war of the near future, we will need to consider some of the limitations of the human brain where AI is likely to add significant value.

What does AI mean for the Air War?

The previous section has made clear that AI will not just give an advantage, but could become essential, either as a complete substitute or a subordinate to a human operator. But where might this be the case? Consider the following idea that compares humans against AI:

> On the other hand, [AI controlled devices] rely largely on pre-existing models and algorithms (although a degree of learning will be expected), and are less likely to show creativity, to detect unconventional opportunity in data collection, or data relevance, or a clever deception hidden in the data. [32]

Humans may continue to outmatch AI, certainly in the near future, within aspects of warfighting that require unconventional or creative thought. Where AI could add value in a complementary fashion alongside a human, is likely to be in areas such as augmentation and decision support. This article proposes that AI possesses three core attributes which make its employment in the air environment desirable:

- **Cognitive Speed (Speed)**: AI has the potential to process and respond to events far quicker than the human cognitive cycle can. Take the example of visual attention. When performing optimally, humans are capable of perceiving gist of scene within 20-ms at the neuronal level and the presence of a specific object within 150-ms, despite subsequent motor actions often taking longer.[33] Despite this, human perceptual and attention ability and accuracy can easily be degraded by, for example, increasing the number of distractor items in a visual scene,[34] increasing the presentation rate and decreasing the relative positioning of target events (attentional blink phenomenon).[35] Humans are also susceptible to failing to notice changes in scenes that occur within fairly rapid time delays (change blindness phenomenon)[36] and especially when cognitive processing is constrained and/or when changes are not expected. Thus, as AI becomes capable of handling more complex tasks at millisecond speeds, this could become a battle winning attribute.

- **Ability to process large volumes of information (Volume)**: Only hardware limits the amount of information an artificially intelligent agent can process. In comparison, there is a limit to the amount of information a human brain can process. For example, the classic view of short-term memory is that it can hold 7±2 items[37] for not much longer than 1/3 of a minute, and possibly less if rehearsal is limited, e.g., by interference from other tasks and/or items,[38] although, remembering information over the short (and sometimes subsequently longer) term can be improved through, for example, processing items to a deeper (and more meaningful) degree.[39] There is also evidence to suggest that short-term memory is adaptive, in that items will decay at a different rate depending on the rate of memory updating required due to factors such as the number and rate of distractor items.[40] An extremely important point here is that whilst human information processing abilities such as short-term memory are limited, the human brain is still incredibly impressive at processing particular types of information in a range of ways, and can adapt to task and environmental constraints.[41] At earlier sensory perceptual stages of processing, the retina can transmit up to ten million bits per second into the visual cortex of the brain, which even advanced AI struggles to match in terms of pattern recognition.[42] The key to leveraging this attribute, as previously highlighted, will be using AI in a complementary fashion to humans.

- **Access to physical space for making a decision (Access)**: Perhaps most prevalent in warfare, it may be impossible or too dangerous for a human to achieve the necessary access to the location where cognition is required.[43] This access could be physical or electronic; for example, while communications technology continues to advance, so do methods of intercepting or jamming communications. Furthermore, in the future we may not wish to place a human in the loop with recent evidence pointing to marked susceptibly of humans in detecting malevolent cues in computer communications masquerading as genuine computer updates.[44, 45]

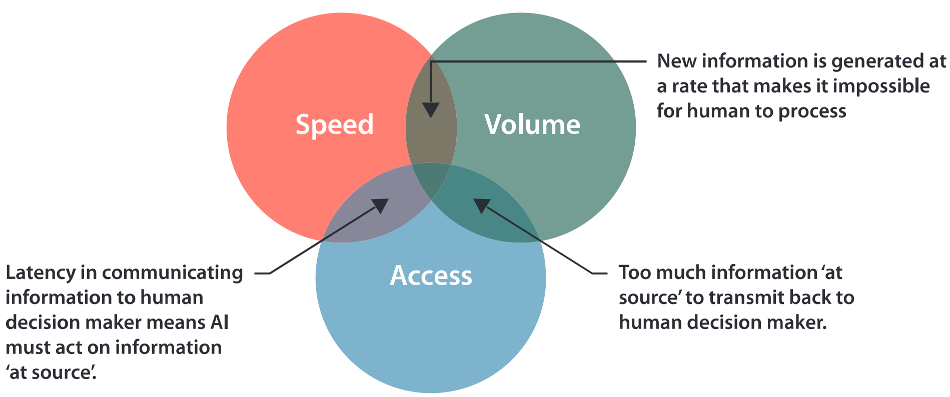

Figure 1: Interaction between Speed, Volume and Access for informing decisions.

All three of these issues could make a case for AI to be employed together with, or instead of a human. Furthermore, they all interlink in a complementary fashion. Practically, where two of the three factors apply this will further make the case *for* employing AI as part of the decision-making process. *Figure 1* illustrates this concept with concentric rings which represent the three proposed reasons for using AI. The areas of overlap present the secondary cases, which are:

- The speed at which new information is generated means the volume is too great for a human to process, or a large volume of information suddenly becomes available and cannot be processed quickly enough by a human to act upon.[46]

- The volume of information generated at the point of physical access is so great that it cannot be transmitted back to a human operator to act upon.[47]

- The transmission time itself from a point of physical access is large enough that it would affect the ability to make a timely decision.

As previously discussed, AI's effectiveness in solving problems diminishes when the problem 'broadens'. Current AI technology is particularly effective if the input variables can be limited and the problem defined in a series of 'values'. This is particularly relevant to air power as it has been suggested that within the air domain there is less complexity in the number of variables. Put simply 'air platforms do not have to dodge potholes in the ground'.[48] Air is likely to be an early adopter of AI technology, and a proving ground for innovative ideas. Consider some of the proposed attributes of AI and how they could affect air power in the next ten years.

Speed

Speed is an attribute of air power and enables the commander to control the tempo of the battle.[49] It is highly likely that AI will accelerate the speed of air warfare in the future and its use will potentially provide 'decision superiority' to commanders, should it be adopted correctly. This effect will be felt all the way from tactical control of aircraft through to the speed of operational and strategic decision making.[50]

AI taking tactical control of aircraft could be a consideration in the near future. In limited test cases, artificially intelligent autopilots have outperformed human pilots in air-to-air combat[51] and AI enabled, neural network-based autopilots are gaining 'superhuman' abilities for things typically considered the preserve of highly experienced pilots, such as judging landing speed.[52] AI's ability to process and act on multiple sources of information faster than a human are only likely to continue this trend, to the point where it may become, at best, an unnecessary extravagance to retain a human in the loop and, at worst, completely negligent to not allow AI to take over a manned aircraft to perform evasive manoeuvres and save the crew onboard, especially if fighting adversary air forces equipped with AI controlled aircraft.

However, perhaps more concerning will be the effect AI could have on decision cycles for entire aerial battles. To use an analogy based on the previous example, if air-to-air combat can be considered as a game of chess, with a series of defined inputs (i.e. sensor and position data of all aircraft) and desired end state (e.g. zero friendly force kills), then one could argue it is only a matter of time before AI will be better equipped to command and control a fast-paced aerial battle. Consider this alongside the fact that a computer could communicate with an aircraft electronically in an extremely precise manner. There would be no requirement to translate instructions or explain intent; an air force under the control of an artificially intelligent commander would possess a cognitive advantage over an adversary, moving and coordinating as a unified body in a way which would be virtually impossible for its human equivalent, as put succinctly by Scharre:

> forces [will] shift from fighting as a network to fighting as a swarm, with large numbers of highly autonomous uninhabited systems coordinating their actions on the battlefield. This will enable greater mass, coordination, intelligence and speed than would be possible with networks of human-inhabited or even remotely controlled uninhabited systems.[53]

In the first instance this is likely to be a continuum. Initially we are likely to see 'loyal wingmen': AI controlled aircraft which follow a human pilot. This could then evolve to 'flocking' whereby a command aircraft could be followed by a series of AI controlled aircraft, retaining the ability to intervene and control distinct elements of the formation. Finally, we may see 'swarming' where AI controls a fleet of air vehicles and a human can only control the group in aggregate, not individual elements.[54]

Volume

AI has an ability not only to cope with large volumes of information, which interlinks with the idea that it can improve cognitive speed, but will also enable correlation of data between vastly different raw formats in ways which would be impossible for the human mind. Recall the way in which Google Translate has developed its own 'language'. Now consider the application of this in the field of hyperspectral imaging, a technique which allows a machine to 'see' in both the visible and infra-red ranges of the electromagnetic spectrum and already has proven military application, for example, in discriminating different types of materials and thereby spotting camouflaged military units.[55] An artificially intelligent agent could easily compare and contrast data from a variety of hyperspectral electromagnetic sources, using neural networks to 'learn' what constitutes an enemy aircraft and communicating with other AI enabled analysis platforms using its own 'base language'. A likely subsequent trajectory will be that such AI will be developed to solve the same problems, but with less source data, or against more ambiguity.[56] The key implication for air forces will be that human operators will need to work in a complementary fashion with such AI, ensuring they understand how it works, and crucially where they can add unique value. In much the same way as the previous example, moving from loyal wingmen to swarming. We are also likely to see a continuum of AI being deployed to analyse data. Consider the idea that the future data analyst of tomorrow may spend less time analysing data themselves, and more time training and developing their 'subordinate' AI to do so.

Access

The final idea is that AI will be able to 'access' locations which a human cannot, and this idea is extremely important as, unlike the other examples it necessitates the deployment of AI, as opposed to being desirable. There are a variety of reasons for AI being preferable to a human operator. Risk is the obvious one within warfighting, but in the future other factors such as physical size may come into play. It may be that battle-winning advantage comes from deployment of aircraft which are simply too small to have a human operator onboard. An obvious point here is that the air war has already begun to encompass this type of technology in the form of Remotely Piloted Air Systems (RPAS). Use of AI to control warfighting aircraft is clearly a contentious proposal and one which goes right to the core of the ethical debate on use of AI and autonomous systems for warfighting. Leveringhaus makes a crucial distinction here, defining RPAS as 'Uninhabited' air systems[57] that are still controlled by a human. In the near future, as air battles take place in '*the most contested areas imaginable, where air, cyber, and the electromagnetic spectrum are all in play*'[58] it is likely to become necessary to design uninhabited air vehicles which can control themselves in the event that communications are lost for minutes at a time.[59] Consider in detail two potential cases where this would be required:

> **Loss of communications with an RPAS**. At the time of writing, Global Hawk is programmed to return to base if it loses communications with its controller,[60] using a basic level of computational intelligence which simply routes the RPAS back to its point of origin using a similar level of sophistication found in autopilots. However, this presents an obvious vulnerability, simply by jamming the communications to the RPAS an adversary could force it to return to base. Even if the jamming were only successful momentarily, it could have a disastrous effect on the prosecution of a mission, with the RPAS changing course to return to base and having to be re-corrected when the human operator regains control. If an adversary were to do this repeatedly, it could effectively deny the capability. Therefore, there is an obvious case for an artificially intelligent agent to be employed to ensure that the RPAS remains on course for momentary jamming, and such AI will likely increase in sophistication to be able to handle longer duration jamming or interception.
>
> **Cyber Attack on an RPAS**: Within cyber operations we will increasingly see artificially intelligent agents responsible for cyber defence. Alexander Kott from the US Army Battle Laboratory proposes that in the near future we will see AI enabled software which can autonomously be deployed within a network or potentially a weapons platform and be left to conduct cyber defence in the absence of an ability to remotely monitor the weapons platform.[61] Kott refers to this type of software as 'bonware' (i.e. the opposite of malware) and proposes that such software will reside in a semi dormant state within

systems, poised to activate in the event of a detected intrusion.[62] The question one must ask here is to what extent should such an artificially intelligent agent be allowed to take control of a weapons system? Clearly this is a question which will be fraught with ethical considerations, and far beyond the scope of this article. However, in considering this question, we should remember that there will be adversaries who will not hesitate to develop malware which will be able to maliciously control a lethal RPAS; AI, or at least AI to augment and support humans, could be the most effective counter to this.

Conclusions

As stated from the outset, the field of Artificial Intelligence has come a long way in the last 70 years, and while there have been several false starts, we appear to be on the cusp of an exciting, incredible and potentially daunting change. Taking the example of chess, during the 1990s computers had been able to leverage their ability to store large volumes of data and rapidly consider several million moves in a short space of time, however, arguably this wasn't true intelligence but more brute force: the Good Old-Fashioned AI. The accelerated development of Neural Networks has allowed computers to genuinely learn and it is clear that AI will be increasingly able to take on tasks which were typically the preserve of humans in the future. A key point this article has identified, however, is that AI will not necessarily directly replace a human, but will be able to offer complementary and potentially diverse capabilities.

To understand what this would mean for the air environment; this article has proposed three core attributes of AI which make its employment desirable when considered against a human. Each of these attributes have ramifications when considering the speed and ubiquity of air power.

The potentially superior cognitive speed of AI compared with humans could not only affect the tactical speed of engagements between air platforms, but may also affect the speed of C2 in the air environment, especially when considering that AI will enable air platforms to communicate electronically and in an unambiguous manner with each other. AI's ability to process volume of information will clearly revolutionise analysis of intelligence, and it is likely that intelligence staff will be some of the first to benefit from the AI revolution and will be at the forefront of genuine human machine teaming.

However, the final attribute: namely that AI offers a means to apply cognitive analysis and decision-making in a physical space, has perhaps the most significant ramifications for air power. While the idea of ceding a lethal air platform to the control of a computer is ethically sensitive, it is only a matter of time before technology will force the requirement to think about this critically upon us.

Humans will continue to bring unique skills including insight and creativity to warfare. However, those who understand what they will be within the age of AI and how best to harness them will have the decisive advantage. While the distance we can see into the future is probably shorter than it ever has been, there remains plenty to be done.

Notes

[1] For pertinent examples see: Andy Fawkes and Martin Menzel, "The Future Role of Artificial Intelligence - Military Opportunities and Challenges," *The Journal of the JAPCC*, no. 27 (2018): 70–80; Elsa B Kania, "Battlefield Singularity : Artificial Intelligence, Military Revolution, and China's Future Military Power," Center for a New American Security, 2017, **https://www.cnas.org/publications/reports/battlefield-singularity-artificial-intelligence-military-revolution-and-chinas-future-military-power**; William Gianetti, "Artificial Intelligence : Myths and Realities," *Air and Space Power Journal* 32, no. 3 (2018): 92–95.

[2] Alan Turing, "Computing Machinery and Intelligence," *Mind* 49 (1950): 460, **https://doi.org/10.1016/B978-0-12-386980-7.50023-X**.

[3] Nuclear weapons formed the 'first' offset while ISR and precision guided weapons formed the 'second'.

[4] Kenneth Payne, "Artificial Intelligence: A Revolution in Strategic Affairs?," *Survival* 60, no. 5 (2018): 7, **https://doi.org/10.1080/10.1080/00396338.2018.1518374**.

[5] Ministry of Defence, "Global Strategic Trends: The Future Starts Today," 2018.

[6] Matthew Wall, "Tech Trends 2019 : 'The End of Truth as We Know It?,'" BBC News, 2019, **https://www.bbc.com/news/business-46745742**.

[7] Justin Lynch, "Why the Air Force Is Investing $100M in AI," Fifth Domain, 2018, **https://www.fifthdomain.com/dod/2018/12/06/why-the-air-force-is-investing-100m-in-ai/?utm_source=Sailthru&utm_medium=email&utm_campaign=12%2F7%2F2018&utm_term=Editorial - Daily Brief&fbclid=IwAR2rkkBUa5JFMmyDJIV0XLMO_zFcg561EfO9ap_kfMOFhuPCzfSLa3KiT2M**.

[8] Nick Carter, "Annual Chief of the Defence Staff Lecture 2018" (London: RUSI, 2018).

[9] UK Ministry of Defence, "Joint Concept Note 1/18: Human-Machine Teaming," 2018, **https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/709359/20180517-concepts_uk_human_machine_teaming_jcn_1_18.pdf**.

[10] Paul Scharre, "Robotics on the Battlefield Part I: Range , Persistence and Daring," *Center for a New American Security*, 2014, 9.

[11] Turing, "Computing Machinery and Intelligence."

[12] I J Good, "Speculations Concerning the First Ultraintelligent Machine," *Advances in Computers* 6 (1965).

[13] Ben Goertzel and Cassio Pennachin, *Artificial General Intelligence* (Berlin: Springer, 2007).

[14] Two excellent books on the concept of Artificial General Intelligence are: Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014) and Max Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence* (New York: Penguin (Random House), 2017).

[15] Seth D. Baum, Ben Goertzel, and Ted G. Goertzel, "How Long until Human-Level AI? Results from an Expert Assessment," *Technological Forecasting and Social Change* 78, no. 1 (2011):

185–95, https://doi.org/10.1016/j.techfore.2010.09.006.
[16] For an excellent treatment of this subject see: Alex Leveringhaus, *Ethics and Autonomous Weapons* (London: Palgrave Macmillan UK, 2018); along with subsequent book review in this publication: Mark Phelps, 'Ethics and Autonomous Weapons: Book Review', *Air Power Review* 21, no. 3 (2018): pp.208–10.
[17] Linell A. Letendre, "Google. . . It Ain't Ford: Why the United States Needs a Better Approach to Leveraging the Robotics Industry," Air Force Law Review (AFLR) 77 (2017): 51–64.
[18] Jim Howe, "Artificial Intelligence At Edinburgh University : A Perspective," *June 2007*, no. June (2007), http://www.inf.ed.ac.uk/about/AIhistory.html.
[19] Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence*, 67.
[20] Jurgen Schmidhuber, "Deep Learning in Neural Networks: An Overview," *Acta Biochimica Polonica*, vol. 62, 2015, 8, https://doi.org/10.18388/abp.2015_1002.
[21] Schmidhuber, "Deep Learning in Neural Networks: An Overview."
[22] Gil Press, "The Brute Force Of IBM Deep Blue And Google DeepMind," Forbes, 2018, https://www.forbes.com/sites/gilpress/2018/02/07/the-brute-force-of-deep-blue-and-deep-learning/#7ba1ccda49e3.
[23] Press.
[24] Feng-Hsiung Hsu, *Behind Deep Blue: Building the Computer That Defeated the World Chess Champion* (New Jersey: Princeton University Press, 2004).
[25] Hsu.
[26] Mike Klein, "Google's AlphaZero Destroys Stockfish In 100-Game Match," Chess.com, 2017, https://www.chess.com/news/view/google-s-alphazero-destroys-stockfish-in-100-game-match.
[27] Yuval Noah Harari, *21 Lessons for the 21st Century* (New York: Vintage Digital, 2018).
[28] United States Air Force, "Air Force Future Operating Concept, A View of the Air Force in 2035," 2015, 9, https://doi.org/10.1080/17450918.2010.497851.
[29] Sam Wong, "Google Translate AI Invents Its Own Language to Translate With," *New Scientist*, 2016, https://www.newscientist.com/article/2114748-google-translate-ai-invents-its-own-language-to-translate-with/.
[30] Kumaran Akilan, "Artificially Intelligent Engineers — How AI Will Kill All Engineering Jobs .," The Forge, 2017, https://theforge.defence.gov.au/sites/default/files/ai_killing_engineering_jobs.pdf.
[31] Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence*.
[32] Alexander Kott and David S. Alberts, 'How Do You Command an Army of Intelligent Things?', *IEEE Computer* 50, no. 12 (2017): pp.96–100, doi:10.1109/MC.2017.4451205.
[33] Simon Thorpe, Denis Fize, and Catherine Marlot, "Speed of Processing in the Human Visual System," *Nature* 381, no. 6582 (1996): 520–22, https://doi.org/10.1038/381520a0.
[34] Anne M Treisman and Garry Gelade, "A Feature-Integration of Attention," *Cognitive Psychology* 12, no. 1 (1980): 97–136, https://doi.org/10.1016/0010-0285(80)90005-5.
[35] K. L. Shapiro, J. E. Raymond, and K. M. Arnell, "Attention to Visual Pattern Information Produces the Attentional Blink in Rapid Serial Visual Presentation.," *Journal of Experimental Psychology: Human Perception and Performance* 20, no. 2 (1994): 357–71.

[36] Ronald A. Rensink, J. Kevin O'Regan, and James J. Clark, "On the Failure to Detect Changes in Scenes across Brief Interruptions," *Visual Cognition* 7, no. 1–3 (2000): 127–45, https://doi.org/10.1080/135062800394720.
[37] George A. Miller, "The Magical Number Seven, Plus or Minus Two Some Limits on Our Capacity for Processing Information," *Psychological Review* 101, no. 2 (1956): 343–52, http://www.psych.utoronto.ca/users/peterson/psy430s2001/Miller GA Magical Seven Psych Review 1955.pdf.
[38] Nelson Cowan, "The Magical Mystery Four: How Is WMC Limited, and Why?," *Memory* 19, no. 1 (2010): 51–57, https://doi.org/10.1177/0963721409359277.
[39] Fergus I. M. Craik and Robert. S Lockhart, "Levels of Processing: A Framework for Memory Research," *Journal of Verbal Learning and Verbal Behavior* 11 (1972): 671–84, https://doi.org/10.1093/bja/45.12.1236.
[40] Erik M. Altmann and Wayne D. Gray, "Forgetting to Remember: The Functional Relationship of Decay and Interference," *Psychological Science* 13, no. 1 (2002): 27–33, https://doi.org/10.1111/1467-9280.00405.
[41] Wayne D. Gray et al., "The Soft Constraints Hypothesis: A Rational Analysis Approach to Resource Allocation for Interactive Behavior.," *Psychological Review* 113, no. 3 (2006): 461–82.
[42] Bostrom, *Superintelligence: Paths, Dangers, Strategies*, 45.
[43] Kott and Alberts, "How Do You Command an Army of Intelligent Things?"
[44] Morgan, Phillip L., Williams, Emma J., Zook, Nancy A. and Christopher, Gary 2018. Exploring older adult susceptibility to fraudulent computer pop-up interruptions. Presented at: AHFE 2018: International Conference on Applied Human Factors and Ergonomics, Orlando, FL, USA, 21-25 July 2018. Published in: Ahram, Tareq Z. and Nicholson, Denise eds. Advances in Human Factors in Cybersecurity: Proceedings of the AHFE 2018 International Conference on Human Factors in Cybersecurity, July 21-25, 2018, Loews Sapphire Falls Resort at Universal Studios, Orlando, Florida, USA. e Advances in Intelligent Systems and Computing Springer Verlag, pp. 56-68.
[45] Williams, E. J., Morgan, P. L. and Joinson, A. N. 2017. Press accept to update now: Individual differences in susceptibility to malevolent interruptions. Decision Support Systems 96, pp. 119-129.
[46] Sean A Atkins, "Multidomain Observing and Orienting: ISR to Meet the Emerging Battlespace," *Air and Space Power Journal* 32, no. 3 (2018): 37.
[47] John Mchale, "ISR Signal Processing Brings Performance to Sensors and Enables AI at the Edge," Military Embedded Systems, 2018, http://mil-embedded.com/articles/isr-processing-brings-performance-sensors-enables-at-edge/.
[48] Ben Jensen and Ryan Kendall, "*Waze For War: How The Army Can Integrate Artificial Intelligence*," War on the Rocks, 2016, https://warontherocks.com/2016/09/waze-for-war-how-the-army-can-integrate-artificial-intelligence/.
[49] UK Ministry of Defence, "JDP 0-30, UK Air and Space Power" (Shrivenham: UK Government, 2017), 25, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/668710/doctrine_uk_air_space_power_jdp_0_30.pdf.
[50] United States Air Force, "Air Force Future Operating Concept, A View of the Air Force in 2035."

[51] Payne, "Artificial Intelligence: A Revolution in Strategic Affairs?," 9.
[52] Chao Tong et al., "A Novel Deep Learning Method for Aircraft Landing Speed Prediction Based on Cloud-Based Sensor Data," *Future Generation Computer Systems* 88, no. 1 (2018): 552–58, https://doi.org/10.1016/j.future.2018.06.023.
[53] Scharre, "Robotics on the Battlefield Part I: Range , Persistence and Daring," 8.
[54] Daniel Wassmuth and Dave Blair, "Loyal Wingman , Flocking , and Swarming : New Models of Distributed Airpower," War on the Rocks, 2018, https://warontherocks.com/2018/02/loyal-wingman-flocking-swarming-new-models-distributed-airpower/.
[55] P WT Yuen and M Richardson, "An Introduction to Hyperspectral Imaging and Its Application for Security, Surveillance and Target Acquisition," *The Imaging Science Journal* 58, no. 5 (2010): 241–53, https://doi.org/10.1179/174313110X12771950995716.
[56] Kott and Alberts, "How Do You Command an Army of Intelligent Things?"
[57] Leveringhaus, *Ethics and Autonomous Weapons*.
[58] Wassmuth and Blair, "Loyal Wingman , Flocking , and Swarming : New Models of Distributed Airpower."
[59] Wassmuth and Blair.
[60] M. L Cummings, "Artificial Intelligence and the Future of Warfare," 2017, https://doi.org/10.1145/2063176.2063177.
[61] Alexander Kott, "Bonware to the Rescue: The Future Autonomous Cyber Defense Agents," in *Conference On Applied Machine Learning For Information Security* (Washington DC, 2018).
[62] Kott.

ROYAL AIR FORCE
Centre for Air and
Space Power Studies