

Developments in Artificial Intelligence - Opportunities and Challenges for Military Modeling and Simulation

Andrew J. Fawkes

Consultant Thinke Company Ltd

UNITED KINGDOM

andy.fawkes@thinke.co.uk

ABSTRACT

One of the principal themes the NATO Science and Technology Organization (STO) is fostering in 2017 is "Military Decision Making using the tools of Big Data and Artificial Intelligence (AI)". Simulation might play a significant role to play in these developments as it can act as a testbed for such concepts and support the military decision makers in future operations that are enhanced by AI. Simulation is already making a significant impact in the development of AI outside of the defence sector. Companies such as DeepMind and Nvidia are using computer games and simulations to "train" AI and autonomous systems, analogous to humans training in simulations. The rate of progress is high, driven by increases in computing power, availability of data and improved algorithms. AI can now "beat" humans at many computer and board games and is moving towards tackling more strategic games that have parallels with military C2. If such developments translate into the defence sphere then we could foresee humans and autonomous systems training in the same simulation systems, both separately and together, and the AI in the autonomous system being the same as that in the simulation. As autonomous systems proliferate across the nations, M&S technology and techniques might be used to improve the interoperability of autonomous systems. To maximise such synergies, it will be essential that NATO embraces all communities that have an interest in AI. Assessing the risks of potential adversary's use of AI and commercial autonomous systems is also necessary. Despite recent advances, AI development still faces significant technological and ethical challenges and these must be monitored and addressed as necessary.

1.0 CONTEXT

Artificial Intelligence or AI is a technology or concept that has developed over many decades and periodically becomes mainstream news. In the latter half of the 2010s this is still very much the case with regular forecasts of its impact on society, jobs and the world economy. Some of these predictions appear to be nearing reality and AI devices are even entering the home. AI also has the potential to influence and sometimes disrupt the ways that companies and organisations operate and AI-based technology revolutions are anticipated. AI also has enduring impact on media and culture and like much technology it can be considered to have beneficial and harmful uses and its impact has and will have political and ethical implications.

Over the decades many AI predictions have come true to some degree but in some cases not at all or only partially. Good examples of this are to be found in transport where passenger aircraft have considerable levels of automation but society remains some way off from accepting pilotless passenger aircraft. For

railways, some are now fully automated whilst others continue to put high reliance on the human. Cars can now park themselves and have high levels of automation but are yet to be fully autonomous in all environments and applications. There is no doubt however, that there is a trend towards greater use of autonomous systems and AI. This is being driven by ever greater processing power together with the ability for very large data sets (“big data”) to be captured and used to help build more capable AI. Such resources can also be accessed online in the cloud, driving down the cost of developing and distributing AI programs.

The military have developed and deployed autonomous systems for a very long time, for example in the use of land and sea mines. In the 20th century proximity fuzes came into service that were semi-intelligent, sensing and exploding at the most appropriate time for the target. Analogue computers also assisted operators as part of fire control calculations and missiles and rockets in World War 2 become remotely piloted or fully autonomous. With the advent of digital computing, autonomy in military systems is commonplace, reducing the manpower requirement or in assisting the human, but there remains a significant ethical dimension in the use of fully autonomous systems. Developments in AI and autonomous systems outside of defence are of significant interest as they may provide answers how to better manage and interpret data within military command and control systems but also because they may enhance potential adversary’s capabilities. This was recognised in 2017 by the NATO Science and Technology Organization’s (STO) which made "Military Decision Making using the tools of Big Data and Artificial Intelligence (AI)" one of its principal themes.

The modelling and simulation (M&S) community has strived itself to develop AI, reducing or eliminating the need for human input. Sometimes termed Semi-Automated Forces (SAF) or Computer-Generated Forces (CGF) this has benefits in analysis and training, reducing the number of role players and improving consistency. Simulation itself is now being used to “train” AI/autonomous systems, as such environments are repeatable and controllable and can generate highly tailorable data output. However, reproducing credible and realistic behaviours in simulation remains a significant challenge and the M&S community continues to strive to enhance its AI. The computer games industry also sees AI as a challenge as games can easily lose their entertainment value if their AI is poorly implemented.

2.0 WHAT IS ARTIFICIAL INTELLIGENCE?

Artificial intelligence (AI) is a broad topic area as it depends on what nature of human intelligence are being replicated and that AI technology can take many different forms. The Oxford Dictionary definition is “*The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.*” The nature of the intelligence can range from “narrow” intelligence which is highly tailored or specialised through to artificial “general” intelligence which is flexible, adaptive and inventive, much like the human brain.

There are many approaches to AI for example, decision trees, fuzzy logic and neural nets with some approaches becoming synonymous with AI. For example, machine learning is an approach that gives “computers the ability to learn without being explicitly programmed” by learning from and making predictions from data. In broad terms AI is the ‘what’, machine learning is an approach to the ‘how’, and self-driving cars might be the ‘why’.

Machine learning methods are based on learning data representations, as opposed to task-specific algorithms. Learning can be supervised, partially supervised or unsupervised. Neural nets or networks are computer systems modelled on the human brain and nervous system with an interconnected group of nodes, akin to the vast network of neurons in a brain. Deep or reinforcement learning, which is inspired by the way animals seem to learn, has taken the neural nets approach and added layers of nodes taking

advantage of current day higher processing power and making significant advances in image recognition for example and is generally seen as at the current forefront of AI technology.

An autonomous system builds on the use of AI and extends it into the physical world, in for example a robot or vehicle, requiring an awareness of the world through sensors, a task(s) and minimal human intervention.

Some argue that some AI algorithms are not really showing intelligence but are a predetermined and limited set of responses to a predetermined and limited set of inputs. Professor Isbell of Georgia Tech suggests that systems should have two features before they can be considered AI. Firstly, they must learn over time as their environment changes. Secondly, their challenge must be demanding too for humans to learn, so a machine programmed to automate repetitive work would not be considered an AI system. Another example would be the AI in many computer video games; it may appear to represent human behaviour but this is pre-programmed and there is little or no learning over time.

3.0 HISTORICAL CONTEXT

An overview of the general history of AI including work on AI by the M&S community will be provided to give context to the progress currently being made and the possible trajectory of the field into the future.

3.1 Early Years

Perhaps the first consideration of computers acting like humans was made by Alan Turing in his 1950 "Computing Machinery and Intelligence" paper. He proposed a test of a machine's ability to exhibit intelligent behaviour equivalent to, or indistinguishable from, that of a human, and this so-called Turing Test remains influential to this day. The term "artificial intelligence" is generally considered to have been coined by John McCarthy in 1955, a professor at Dartmouth College (USA), ahead of a Dartmouth Summer Research Project on Artificial Intelligence in 1956. The intent was to study "the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it" and "find out how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves." Although the project delivered a shared vision it could not agree on the specifics of a way forward.

In 1960 the American psychologist and computer scientist, J.C.R. Licklider, published his paper "Man-Computer Symbiosis". Licklider foresaw significant advances in artificial intelligence but argued that such machines would not replace humans but be symbiotic: "Men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions in technical and scientific thinking".

In the subsequent decades there were early successes in areas such as mathematical problem solving, natural language, and robotics with some of the ideas being central to modern AI, such as those behind neural networks. Funding for AI research ebbed and flowed as progress was sometimes slower than expectations with hype cycles in the 1970s, 80s and 90s. Indeed, in 1984 the term "AI Winter" was coined at the "American Association of Artificial Intelligence" annual meeting, describing a period of reduced funding and interest in AI research.

Looking at the military dimension, upon becoming DARPA Director in 1975, Dr. G. Heilmeier, raised "very fundamental and pragmatic questions about the AI research field" as it transpired that such questions had not been asked in the past. "It wasn't that I was never a believer in AI, I just wanted them (the AI program leaders) to answer basic questions, and they couldn't." Nevertheless, a number of AI applications began to appear in the late 1970's, including some for military systems, largely based on technology and

technologists supported by DARPA. In 1983 AI technology was incorporated as a key component of DARPA's Strategic Computing Program.

3.2 AI and the Modelling and Simulation Community

The M&S community also had significant incentives to exploit AI approaches, for example, to better represent humans in operational analysis and to reduce the number of people required to support training. The SIMNET program of the 1980s helped to advance M&S AI significantly.

Through to the 1980s, interactive simulation equipment was very expensive, and reproducing training facilities was likewise expensive and time consuming. In 1983 DARPA hence funded the development of SIMNET (SIMulator NETworking), a prototype research system to investigate the feasibility of creating a more cost effective real-time distributed simulator for combat simulation, with armoured vehicles operators the principal training audience. After a series of successful prototypes, by February 1988 the SIMNET facility at Fort Knox had grown to a total of 71 interactively networked simulators on a LAN and included other players such as attack helicopter simulators. Distributed training could also be supported with links to Forts Benning and Rucker and Germany.

As larger and more complex exercises were conducted, many more vehicles were needed to populate flanking and supporting units to provide a realistic context. Even if enough manned simulators had existed to fill these roles, doing so would have required hundreds of troops who would be doing little more than playing the roles of “extras” and therefore costly and manpower inefficient. Thus by 1988 a semi-automated forces (SAF) capability had been developed. This capability permitted one person at a SIMNET “command module” to orchestrate interactively the battlefield movements of several units such as a platoon or company. A semi-automated approach was chosen over a fully autonomous one because the restriction to carefully planned scenarios was too limiting to achieve the kind of free-play man-in-the-loop environment DARPA was aiming for.

The term semi-automated forces (SAFOR) was originally used in DARPA's Advanced Simulation Technology Program. The goal of this effort was to provide large software driven forces (flanking, supporting, enemy) to interact with manned simulators on the SIMNET battlefield. SAFOR units were required to execute their orders in accordance with doctrine and in a manner realistic enough that, from an operational standpoint, simulated elements would be indistinguishable from other participants on the battlefield, much like the classic “Turing Test” of intelligence. These simulated forces were also referred to by a variety of terms other than SAFOR, including SAF, Computer Generated Forces (CGF), Intelligent Forces (IFOR), and Automated Forces (AFOR).

First released in 1993, ModSAF (Modular Semi-Automated Forces) was the third generation of SAF systems. Early SAF coding had evolved to be non-modular, making it complex to understand and reuse. ModSAF aimed to provide a modular open architecture that would allow the addition of a wide variety of battlefield entities, including by third party developers, and scale up to allow operators to control higher echelon units.

By the mid-1990s, US SAF development had become uncoordinated such that there were several different SAFs in service or in development funded by different organisations. For example: BBN SAF, BDS-D CGF, IFOR/WISSARD, SWEG/SUPPRESSOR, ModSAF 1.0, CCTT SAF, and Janus. In response to this, the OneSAF concept was introduced by the US Army in 1996. The OneSAF aim was that one modular software package would meet all the M&S needs of the R&D and Acquisition; Concepts and Requirements; and the Training, Exercises and Military Operations M&S domains.

Outside the military in the late 1990s AI was predominantly driven by increasing computational power, greater emphasis on solving specific problems, and new ties between AI and other fields. For example,

IBM Deep Blue's chess program was the first to beat a reigning world chess champion, Garry Kasparov on 11 May 1997.

Outside of the OneSAF project, other nations and companies developed their own SAF/CGF or constructive simulations, for example, JSAF, Abacus, GUPPIS, VR Forces, and Masa Sword. This is partly because different nations and customers have different doctrines and languages but also newer products have entered the market exploiting more modern software approaches. SAF/CGF have developed into two broad categories, aggregated and disaggregated. The former simulating aggregations of units more suited to support command training and exercises, and the latter simulating individual entities from a few to many thousands. These boundaries have blurred but the SAF/CGF selected depends very much on the activity being supported and often also it is down to legacy issues; operators may be only experienced in using one SAF/CGF, companies may wish to tie their customer to their technology, and the software may have been rewritten over time.

The majority of modern SAF/CGF are broadly interoperable such that they can interact through simulation protocols such as DIS/HLA and can link to live and virtual simulation systems, providing automated forces to enrich the training and reduce the manpower support. Virtual simulations may have their own built-in AI behaviours that can support a range of training activities that do not require a sophisticated constructive simulation.

3.2 Military Command and Control Systems

As military Command and Control (C2) systems have become more digitized there has been a need to use simulations to stimulate and be stimulated by the Command, Control, Communications, Computers, and Intelligence (C4I) systems, particularly for training. This is principally to reduce the number of humans required to provide the interface between the training unit and the simulation. In response, the Battle Management Language (BML) initiative was instigated in 2001. BML is an unambiguous language for the command and control of forces and equipment conducting military operations, providing situational awareness and a shared, common operational picture. It is a standard digitized representation of a commander's intent to be used for real troops, for simulated troops, and for robotic and automated forces.

Building on BML, from 2004 the Coalition Battle Management Language (C-BML) introduced a coalition perspective. Under SISO, C-BML (together with Military Scenario Definition Language (MSDL)) has been replaced by the Command and Control Systems - Simulation Systems Interoperation (C2SIM) standard. Such initiatives offer the potential for C4I systems to be able to easily switch between interacting with real people and systems, including robotics and autonomous systems, and simulated people and systems.

4.0 ARTIFICIAL INTELLIGENCE - CURRENT DEVELOPMENTS

We have seen that AI technology has developed over several decades, typically funded by the military, and that the M&S community has developed a number of AI technologies, principally to replicate human behaviour on the battlefield and reduce the number of real humans required to support training. The M&S community has also developed protocols that allows C4I systems to interact with both simulation systems and real robots and autonomous systems.

4.1 Driving Factors in Artificial Intelligence

In the 2010s the development of AI technology is very much in the realm of academia and commercial technology companies. Further, the computer games industry strives to enhance its AI technology to reduce production costs and maximise the entertainment value. Such progress has led some to consider that the military is lagging behind the technology industry in leveraging AI.

There are several forces that are currently driving AI, but these are generally considered the predominant factors:

1. Hardware - Processors are being developed to enable faster and more powerful AI, hosted locally or in the cloud. Lately, the synergies of AI programming (eg. solving matrices) and graphics cards architectures is being exploited.
2. Data - Through large data sets, machine learning algorithms learn by refining hypotheses iteratively. From real-time information discovery to the integration of algorithms and data, more tools for working with data are enabling analysis of an increasing number of publicly available datasets. Termed “data farming”, data can also be “grown” using M&S, giving more control over the data quality.
3. Algorithms - Algorithms such as deep learning and its hierarchical pattern recognition are a major force driving the adoption of AI.

There is debate also about whether the rate of technological progress is accelerating threatening to make radical changes to society. The Electronic Frontier Foundation, which campaigns to protect civil liberties from digital threats, has started its own effort to measure and contextualize progress in AI and to inform this debate.

4.2 Games and Simulation as a Testbed for the Development of Artificial Intelligence

Any discussion of AI current developments could be very extensive but perhaps the most relevant to the M&S community is the use of games and simulation as a testbed for AI. This relationship is not new. As noted earlier, the game of chess was used to test computer systems and algorithms in the 1990s, leading to IBM Deep Blue’s capabilities in 1997. More recently, (Google) Alphabet’s DeepMind has used board games to test and train its AI. Of note is their development of “AlphaGo” which used both humans and deep learning to train its AI to play the board game “Go” and beat the world’s leading Go players. This goal was achieved more quickly than anticipated by many Go players as Go has many alternatives to consider per move. Subsequently AlphaGo has helped Go players learn new strategies that had not been considered previously. This reinforces the concept of AI not replacing humans but helping them to learn new ways of working.

DeepMind has also tested its AI using computer games. As opposed to other AIs, such as IBM’s Deep Blue, which were developed for a pre-defined purpose and only function within its scope, DeepMind claimed that its system was not pre-programmed: it learnt from experience, using only raw pixels of Atari games as data input. In 2013, DeepMind published a paper describing an AI playing seven different Atari 2600 video games (Pong, Breakout, Space Invaders, Seaquest, Beamrider, Enduro, and Q*bert). Since then DeepMind have exploited and successfully played more complex games such Space Invaders, Breakout and 3D games such as Doom. Although impressive, these games are well bounded with full information availability. DeepMind together with the games company Blizzard are hence moving towards developing their AI using more strategic games, notably the StarCraft game.

StarCraft is a real-time strategy (RTS) game that combines fast paced micro-actions with the need for high-level planning and execution. It has millions of casual and highly competitive professional players. From a reinforcement learning perspective, StarCraft II also offers four principal challenges:

1. It is a multi-agent problem in which several players compete for influence and resources and each player controls hundreds of units, which need to collaborate to achieve a common goal.
2. It is an imperfect information game with a “fog-of-war”, obscuring the unvisited regions of the map
3. The action space is large and diverse and players select actions among a combinatorial space of approximately 108 possibilities, using a point-and-click interface.
4. Games typically last for many thousands of frames and actions, and the player must make early decisions with consequences that may not be seen until much later in the game.

This research commenced in 2017 and the results are not available. However, the company OpenAI has developed an AI agent that in August 2017 beat professional players in a similar strategy game Dota 2 (although this was in the tactical part of the game and human players have since learnt to beat the AI). This research has a number of parallels with military C2 systems and if it is transferrable it would have the potential to improve military decision support systems or automate some current human activities.

Beyond the computer realm and into the physical world, virtual worlds/simulations are being explored as a medium to train robotic systems. Nvidia has created the "Isaac simulator" in which systems can self-learn through carrying out tasks and interacting with the simulation. Isaac uses Nvidia's graphics processing technology and the Unreal Engine 4 to render a realistic virtual environment in which systems can self-learn through trial and error to carry out tasks and interact with the simulation. Repetition of this process enables the speedier iteration of AIs and their deep learning systems in an environment where testing is safer and more scalable than real-world testing in the early stages of development. At Computex 2017, Nvidia demonstrated how an AI controlled robot can be trained to play ice hockey from scratch using Isaac in simulation prior to the program being transferred to a real-world robot.

Nvidia is also exploiting its AI technology to develop Metropolis, an intelligent video analytics platform which applies deep learning technology to video streams to aid traffic management, public safety and other smart city tasks. The motive for this is that only a small fraction of captured video is properly analysed by humans, and that existing analytics systems are not capable enough to make sense of the subtleties and complexity of video feeds from busy cities. Such technology might find its way to military monitoring systems such as military base cameras and UAV video feeds.

AI has also been an important element of computer video games. It is used to generate responsive or intelligent behaviours primarily in non-player characters (NPCs), similar to human-like intelligence. Video game AI technology can draw on many of the AI techniques seen in computer science in general. However, they may not constitute "true AI" in that such techniques may only constitute "automated computation," or a predetermined and limited set of responses to a predetermined and limited set of inputs.

Beyond controlling NPC behaviour there are other uses for AI in games. AI may track the ability and emotional state of the player, to tailor the game appropriately. It may be used to procedurally generate game content such as environmental conditions or interactive stories. AI may also be used to analyse and learn from the data from games so that designers can discover how players use their games, including what causes them to stop playing games or what maximises monetization.

AI can be trained to play computer games but the Georgia Institute of Technology have for the first time had some success in training AI to recreate a computer game engine by only watching the game. This work could lead to automated testing of games and simulations, easier games production, and potentially AI that comprehends the real world through observation.

4.3 Artificial Intelligence and Autonomous Physical Systems

Moving towards more complex autonomous physical systems considerable resources are being put into the development of autonomous vehicles. This is a very challenging task as traffic environments are complex and variable and essentially built around humans, not autonomous systems. There are also many legal and transitional issues that must be resolved. Nevertheless, significant progress is being made such that semi-autonomous systems such as parking aids are already widespread use.

Alphabet's Waymo is developing autonomous cars through testing in both the real world and in simulation. In 2016 Waymo logged 3 million miles on real world public streets whilst 2.5 billion virtual miles were driven in their simulation systems. Waymo give three reasons to use simulation. Firstly, many more miles can be driven than would be possible with a physical fleet. Secondly, simulated miles focus on the

still-difficult interactions for the cars rather than uneventful miles. And thirdly, the development cycles for the software can be much faster.

4.4 The Military and Artificial Intelligence

Recent AI developments are not exclusive to the commercial world. The \$813m US Navy X-47B program developed an unmanned combat air vehicle (UCAV) that by 2014 had completed a significant number of aircraft carrier catapult launches, arrestments, and touch-and-go landings, with only human supervision (not through remote control). In April 2015, the X-47B successfully conducted the world's first fully autonomous aerial refuelling. The US Navy is also funding research into autonomous swarming boats and UAVs.

US AFRL has funded an alternative AI approach to the neural networking using fuzzy logic. Fuzzy logic is a form of many-valued logic in which the truth values of variables may be between 0 and 1. Fuzzy logic is not new but University of Cincinnati researchers have developed "genetic fuzzy trees" that are able to handle many inputs with relatively low processing demands. This research had led to the "ALPHA" system which in flight simulator tests has beaten an experienced combat pilot in a variety of air-to-air combat scenarios. This work may lead to improved UAV capabilities and AI being used as an "assistant" for pilots in manned aircraft.

A variation on autonomous system and human-controlled vehicle military operations is the manned-unmanned teaming (MUM-T) concept. This combines the strengths of each platform to increase situational awareness, allowing the armed forces to conduct operations that include combat support and intelligence, surveillance, and reconnaissance (ISR) missions. The US Army is making MUM-T an established part of its tactics, techniques, and procedures (TTPs) by outfitting its combat aviation brigades with Boeing's AH-64D/E Apache helicopters and Textron Systems RQ-7B Shadow Tactical Unmanned Aircraft Systems (TUAS). Unmanned systems are expected to extend the breadth of a human system's comprehension of their surrounding environment enabling enhanced decision making and more synchronized, responsive actions. This teaming theory also provides a level of safety for the manned platform, as it can remain in non-hostile areas whilst the unmanned system can provide the forward scout mission until targets are identified and enemy positions are known.

4.5 Autonomous System Interoperability

As autonomous systems proliferate there will be an increasing need for them to interoperate, much in the same way that humans must. In a NATO context, autonomous systems and their human controllers will need to act together coherently, effectively and efficiently to achieve Allied tactical, operational and strategic objectives. In such an interoperable system, interfaces will need to be understood, and heterogeneous systems, present or future, will need to work together with no or few restrictions.

Some related work is being done in this field, and in research published in 2016 at the International Workshop on Modelling and Simulation for Autonomous Systems (MESAS), an HLA federation was set up that included that hardware and software of a real Autonomous Underwater Vehicle (AUV) together with a virtual simulator. The M&S and robotic communities are expected to benefit from such an approach which allows the development of more complex and realistic simulated scenarios with hardware-and software-in-the-loop. It is perhaps feasible that in the future autonomous systems might interoperate using HLA technologies and architectures.

4.5 The Ethical Dimension of Artificial Intelligence

Ever since artificial intelligence came to the public consciousness there has been a debate and sometimes fears about how AI might progress and be controlled. The concept of a "technological singularity" has

been debated which is a point in time when AI exceeds human intelligence and develops beyond. Prior to any such hypothetical point there is also much debate as to how jobs will develop, disappear and be invented with the impact of AI.

How AI will develop taking into account human safety is very much a political and ethical issue. For autonomous vehicles, the German Government has accepted the following recommendations, which is the first time such fundamental principles have been articulated:

1. Automated/networked driving is ethically necessary if the systems cause fewer accidents than human drivers.
2. In the event of danger, the protection of human life always has top priority.
3. In the case of unavoidable accidents, any qualification of people according to personal characteristics (age, sex, physical or mental constitution) is not permitted.
4. In any driving situation, it is necessary to clearly define and who is responsible for the driving task: the human or the computer.
5. Anyone who drives must be documented (to clarify possible liability issues).
6. The driver must be able to decide themselves about the use of their vehicle data.

For the military, the UN has voted to begin formal discussions on the use of autonomous weapons including drones, tanks and automated machine guns. An open letter to the UN signed by 116 AI experts called for a ban on such weapons stating: “Once developed, lethal autonomous weapons will permit armed conflict to be fought at a scale greater than ever, and at timescales faster than humans can comprehend. These can be weapons of terror, weapons that despots and terrorists use against innocent populations, and weapons hacked to behave in undesirable ways.” In response, the UK Ministry of Defence (MoD) has taken a step towards developing a doctrine for the use of more autonomous systems stating: “It’s absolutely right that our weapons are operated by real people capable of making incredibly important decisions, and we are guaranteeing that vital oversight”; the MoD has not called for a pre-empt ban, however.

There is much discussion of the ethics of AI in warfare but in future our adversaries may not see these as constraints. Terrorist groups or adversaries engaged in asymmetric warfare may use AI to bring down both military and civilian IT infrastructures and attack the operation of vital utilities. We are already seeing AI, or at least sophisticated algorithms together with big data, being used to assess and sometimes influence public opinion. Autonomous systems and robotics may be used to physically attack installations and people, and these systems may well be commercially available such as autonomous vehicles and drones.

4.6 Artificial Intelligence Technology Challenges

Although there have recently been ground-breaking developments in AI, many experts highlight that there are still problems that must be overcome, particularly for neural net based approaches:

1. Large volumes of data are required,
2. AI programs typically can only do one task,
3. AI programs can be difficult to understand and verify after programming.

AI systems require more information than humans to understand concepts or recognize features, and this can be hundreds of thousands of times more. Internet companies such as Google and Amazon have access to large amounts of data but most AI researchers and companies do not and hence developments can tend to be centralised in the “tech giants”. Problems can also occur where the data is sensitive, for example in healthcare, acting as a barrier to developing AI. This is also a problem for the military where there may be security concerns and even for large militaries, the volume of data may be comparatively low.

Current systems can be trained to do one task very well but cannot do more than one task. In spite of their successes such as AlphaGo, DeepMind recognise "there is no neural network in the world, and no method right now that can be trained to identify objects and images, play Space Invaders, and listen to music." This can be further illustrated by the training of neural networks to play computer games. Systems can currently master one game but they must be retrained from scratch to play any new game; any knowledge or skills cannot be passed on as is the case for humans.

Techniques such as neural nets and deep learning can provide excellent results and exceed human performance at some tasks. However, such programs learn from data and can interpret the data patterns in alternative ways to humans. As an example, researchers at Virginia Tech trained a neural network using a large sample of photos to detect what cover a window had. It appeared to give good results. However, after further human analysis, the program had made the decision based on whether the photo included a bed; the program had made the incorrect deduction that all bedrooms would have a curtain. It had seen the bed not the curtain.

AI programs generating new insights may be advantageous but it may be an incorrect outcome if the context of the data is not fully understood. Further, neural nets would ideally have a good level of interpretability but this can be detrimental to accuracy as there is typically a trade-off between interpretability and accuracy. In response to this challenge, DARPA are funding an Explainable Artificial Intelligence (XAI) research program with the two goals: To develop more explainable models, while maintaining a high level of learning performance (prediction accuracy); and enabling human users to understand, appropriately trust, and manage their AI systems.

5.0 CONCLUSION

Artificial intelligence as a concept is 60 or so years old. It has made a significant impact on human culture, but as a technology it has had several cycles of progress followed by periods of reduced funding and significance. Through much of the 20th Century, funding for AI research was principally public sector, especially from the military. However, we see in the 21st Century significant contributions being made by the private sector, particularly the large Internet companies and others such as in the automotive field. There have been ever greater synergies between the use of games and simulation to train and test AI and the suitability of gaming graphics cards for processing AI algorithms and simulation systems are being used to train vehicle AI systems ahead of real world usage.

AI is now never far from the news and is being driven by significant improvements in computer processing speeds, availability of data, and enhanced AI algorithms. There is much debate as to how AI will affect society and jobs and the military, with sometimes grave predictions of the future should AI developments not be controlled. Assessing the risks of potential adversary's use of AI and commercial autonomous systems is also necessary.

The military M&S community itself has had a long involvement with the development of AI, for example to reduce manpower to support training and enhance operational analysis. Modern constructive simulations are modular, easier to use, and typically can interoperate with other constructive simulations and wider virtual and live simulation systems. The M&S community has also developed interoperability standards that allow C2 systems to interact with both simulations systems and real-world robotics and autonomous systems. The interoperability of autonomous systems across the NATO nations might be facilitated using M&S standards such as HLA.

Although great strides are being made in AI and autonomous systems there are still significant challenges. AI programs tend to be single function and cannot transfer learning and sometimes require very large datasets to train from. There can also be significant verification issues in terms of the AI drawing the

wrong conclusion or the human not understanding how an AI might reach a conclusion. Ethical and social issues can also be very challenging now and into the future.

AI and autonomous systems impact across many communities both inside and outside of the military. The military M&S community itself has much to contribute to NATO and the Nations given its long-standing development of AI in simulation systems and the development of interoperability standards such as C2SIM and HLA. Equally, there is potential for the community to learn from non-defence sphere in the use of simulations and games to train and test AI, both for its own uses and wider exploitation. Although the human is likely to always be part of the decision cycle, working towards the optimum human and AI teaming should be a clear objective for all communities to maximise the contribution of both. Simulation should have a key role in testing and training both humans and AI to achieve this optimum blend.

6.0 REFERENCES

- [1] MSCO email to NMSG Community; 18 May 2017
- [2] Machine "Some Studies in Machine Learning Using the Game of Checkers"; IBM Journal of Research and Development; Arthur Samuel; 1959
- [3] 'Artificial Intelligence' Has Become Meaningless; The Atlantic; 4 March 2017
- [4] "Computing Machinery and Intelligence"; Mind; Alan Turing; 1950
- [5] McCarthy, J., Minsky, M., Rochester, N., Shannon, C.E., A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence; 1955
- [6] J. C. R. Licklider. Man-Computer Symbiosis, IRE Transactions on Human Factors in Electronics, volume HFE-1, pages 4-11; March 1960
- [7] Wikipedia - AI winter
- [8] DARPA Technical Accomplishments An Historical Review of Selected DARPA ... I Sidney G, Reed Richard H. Van Ana Seymour J. Deitchman; February 1990
- [9] DMSO Survey of Semi-Automated Forces; 1993
- [10] DARPA Technical Accomplishments, Volume II: An Historical Review of Selected DARPA Projects", Institute for Defense Analyses; 1991
- [11] MODULAR SEMI-AUTOMATED FORCES; Andy Ceranowicz, Loral Advanced Distributed Simulation; 1994
- [12] SIMNET and Beyond: A History of the Development of Distributed Simulation; Duncan C. Miller; 2015
- [13] Engineering Principles of Combat Modeling and Distributed Simulation; edited by Andreas Tolk; 2012
- [14] Essens, P.J.M.D., Vogelaar, A.L.W., Mylle, J.J.C. Baranski, J.V., Goodwin, G., Van Buskirk, W. & Hof, T. (2010). CTEF 2.0 – Assessment and Improvement of Command Team Effectiveness: Verification of Model and Instrument. NATO RTO technical report. AC/323(HFM-127)TP/310; 2005
- [15] Adaptability in Coalition Teamwork; NATO RTO Meeting Proceedings; AC/323(HFM142)TP/212; 2008

-
- [16] Standardizing Battle Management Language - A Vital Move Towards the Army Transformation, Scott A. Carey, Martin S. Kleiner, Michael R. Hieb, Ph.D., Richard Brown; SISO 2001
 - [17] 7 Key Factors Driving the Artificial Intelligence Revolution, David J. Hill, SingularityHub; 29 Aug 2016
 - [18] “Do We Need a Speedometer for Artificial Intelligence?” Wired; 30 August 2017
 - [19] Playing Atari with Deep Reinforcement Learning; DeepMind: Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller; 19 Dec 2013
 - [20] StarCraft II: A New Challenge for Reinforcement Learning; DeepMind, Blizzard Games; 2017
 - [21] Game Engine Learning from Video; Matthew Guzdial, Boyang Li, Mark O. Riedl; School of Interactive Computing, Georgia Institute of Technology; 2017
 - [22] Business Insider; 11 Aug 2017
 - [23] “Inside Waymo's Secret World for Training Self-Driving Cars”; The Atlantic; 23 August 2017
 - [24] Wikipedia - Artificial Intelligence (video games)
 - [25] Genetic Fuzzy based Artificial Intelligence for Unmanned Combat Aerial Vehicle Control in Simulated Air Combat Missions; Journal of Defense Management; Psibernetix Inc.; Air Force Research Laboratory, USA; University of Cincinnati; 2016
 - [26] Silicon UK; Nvidia Isaac Is A Virtual Simulator For Training AI And Robots Speedily And Safely; 30 May 2017
 - [27] Silicon UK; Nvidia Metropolis Platform Brings AI Smarts To Smart City Video Feeds; 9 May 2017
 - [28] MUM-T operations on the U.S. Army's UAS roadmap; Military Embedded Systems; 23 August 2017
 - [29] Modelling and Simulation in the Autonomous Systems' Domain; MESAS; Current Status and Way Ahead Jan Hodicky; NATO Modelling and Simulation Centre of Excellence, Rome, Italy; 2015
 - [30] "The Coming Technological Singularity: How to Survive in the Post-Human Era"; Vision-21: Interdisciplinary Science and Engineering in the Era of Cyberspace; G. A. Landis, ed., NASA; 1993
 - [31] Report of the Ethics Commission Automated and networked driving (ethics rules for driving computers); German Federal Government; the 23 August 2017
 - [32] The Observer Newspaper; 10 September 17
 - [33] The Telegraph Newspaper; 21 August 17
 - [34] “These are three of the biggest problems facing today's AI”; The Verge; Oct 10, 2016
 - [35] Das, A., Agrawal, H., Zitnick, C.L., Parikh, D., & Batra, D. Human attention in visual question answering: Do humans and deep networks look at the same regions? Conference on Empirical Methods in Natural Language Processing; 2016
 - [36] Defense Advanced Research Projects Agency (DARPA); Explainable Artificial Intelligence (XAI); Mr. David Gunnin
-

