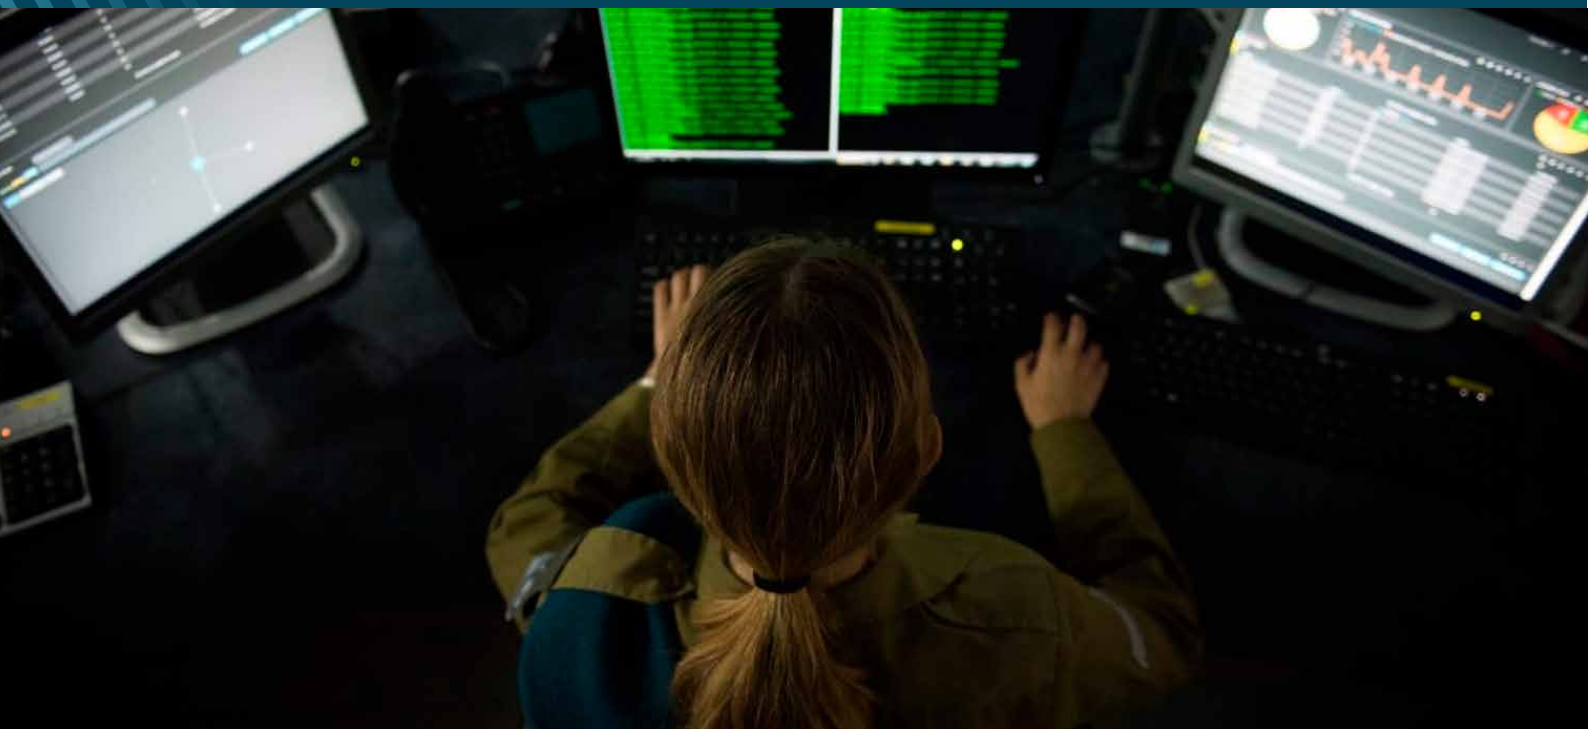

Using artificial intelligence in the IDF Computer Service Directorate. Photo: Bamahane website

Research in the Intelligence Community in the Age of Artificial Intelligence

Shmuel Even and David Siman-Tov

By their very nature, intelligence communities are ideal customers for new information technologies. This article addresses two questions: How do new information technologies, primarily artificial intelligence, contribute to the intelligence community's research activities? What challenges are posed by the assimilation of artificial intelligence in this field? The integration of artificial intelligence in strategic research may provide intelligence agencies with enhanced capabilities in helping leaders understand an emerging reality, detect changes of course early, and manage risks and opportunities. However, the path to achievement of these capabilities is replete with challenges.

Keywords: intelligence research, intelligence assessment, war alert, situation assessment, decision making, strategy, artificial intelligence, machine learning

Introduction

Artificial intelligence (AI) is a major evolutionary step in the broad field of information technologies. The elements enabling AI technology applications are the augmentation of computerization capabilities, growth in the volume of information, better algorithms, and growth in investments (Grimland, 2018). AI-based applications are increasingly integrated in daily life, and their impact can be expected to expand and intensify in many spheres ("How to Ensure Artificial Intelligence Benefits Society," 2020). The same is true with regard to AI-based applications in the security and intelligence systems. By their very nature, intelligence communities are ideal candidates for employing artificial intelligence, since the majority of their activities involve collecting enormous amounts of information from a variety of sources, processing data, conducting research, and formulating scenario-based assessments and predictions.

This article discusses information technologies in the age of artificial intelligence, in the context of intelligence research in intelligence communities, particularly strategic research. It will attempt to clarify how IT capabilities ("the technology" or "the machine") can contribute to intelligence research activity and to the formulation of intelligence assessments, and what challenges are posed by the assimilation of artificial intelligence in intelligence research.

The Artificial Intelligence Concept

There is no fully accepted definition of the term "artificial intelligence." The following section is based primarily on official United States national security texts. According to the definition in the John S. McCain National Defense Authorization Act for Fiscal Year 2019, artificial intelligence is:

1. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets.
2. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
3. An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
4. A set of techniques, including "machine learning," that is designed to approximate a cognitive task.
5. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

(US Congress, 2019)

These definitions are partially based on the similarity between the mode of operation and output of an AI system and human modes of thinking. In this article, the achievement required of the machine will be examined primarily according to the degree of its "intellectual capacity" to benefit humans in achieving their goals, and in this context, benefit research entities in the intelligence community, while exploiting the relative advantages of machines that are capable of performing specific tasks at levels of accuracy, speed, volume, and complexity that far exceed the capabilities of the human mind.

The brief description below of the nature of artificial intelligence is based on the article "A DARPA Perspective on Artificial Intelligence" (Defense Advanced Research Projects Agency, an agency of the United States Department of Defense responsible for the development of advanced technologies) by John Launchbury, the Director of the Information Innovation Office at DARPA (Launchbury, 2017). His article presents four characteristics of artificial intelligence capabilities:

1. The ability to grasp the environment outside the machine (perceiving): the machine's

ability to detect, analyze, and respond to information that it or systems connected to it collect from outside of the computer, such as a system that detects and analyzes vehicular traffic data, images in an environment, and so on.

2. The ability to learn (learning): the system's ability to learn from examples and apply the knowledge to new circumstances.
3. The ability to abstract (abstracting): for example, the ability to take knowledge discovered at a particular level and apply it at a higher level. This ability enables the creation of new meanings, but it requires a contextual capability.

According to Hallman, the main challenge is "taking the vast volumes of digital intelligence that the CIA receives from around the world and transforming them into a digital, dynamic and credible picture of the future." Hallman adds that "intelligence, in this context, becomes almost a super power."

4. The clarity of causality (reasoning): for example, to what extent can a human user of a machine understand the correlation between the raw data and the machine's conclusions. This is a critical ability for planning and decision making.

According to DARPA, the development of artificial intelligence may be differentiated by three waves:

The first wave is the programmed ability to process information. Experts take knowledge that they have of a particular subject, characterize it according to rules that are computer-compatible, and in turn, the machine processes the data according to algorithms that they wrote and generates output according to a defined pattern. Examples of this are logistics software, chess software, and tax computing software. At this stage, the capabilities of artificial intelligence are characterized by high "clarity of causality," in the sense that the researcher understands the logical cause-and-effect connections that are operating in the machine, and the source from which the machine took or received each item of data. However, this clarity is limited to a specific deterministic process that the person dictates through an algorithm and through the information that it feeds into the machine. The capabilities at this stage (which have been implemented in recent decades) are still of considerable value today. For example, in a DARPA project to augment cyber security, the first-wave systems succeeded in helping detect cyber security vulnerabilities.

The second wave is now at the center of the discourse and action in particular areas of artificial intelligence. The second wave includes applications of voice recognition, facial recognition, photo sorting, and more. The second-wave systems are highly capable of perceiving the environment outside the computer (using sensors and a link to big data). These systems are characterized by statistical learning and include the use of artificial neural networks that are characterized by deep learning. Using this technology, an AI system knows how to identify a phenomenon based on characteristics that it learned independently from examples (such as from a series of imaging of a disease) and not solely according to characteristics that the researcher input into the machine. DARPA uses second-wave technology, inter alia, to analyze the spread of cyberattacks and to gain autonomous tools. Despite all of these advantages, second-wave systems engage in specific tasks and have minimal capability of presenting reasoning. They do not have the ability to give an explanation (explainability).[1] Another limitation is that second-wave systems need an enormous amount of data in order to learn,[2] and they are not immune from errors.

The third wave is still at its initial research and development stages. It is supposed to overcome some of the limitations of the earlier waves and create additional capabilities. Sources in DARPA believe that the third wave

will be designed using contextual models that will enable, inter alia, the design of systems that know how to learn from a limited number of examples, how to provide explanations for their results, and how to create new meanings from data (the ability to abstract).

Thus, first-wave artificial intelligence is still relevant; the second wave provides high capabilities in particular areas; and the third wave reflects expectations of advances in the coming decade. It certainly will not be the last wave.

Table 1. Characteristics of AI technologies

	Situation	Perceiving the natural environment	Learning capability	Abstracting capability	Reasoning capability
First wave Logical data processing	Applied in recent decades and still evolving	Low	Nonexistent	Nonexistent	High
Second wave Statistical learning, neural networks	Applied, expanding, very effective in particular areas	High	High	Low	Low[3]
Third wave Contextual adaptation	R&D underway	High	High	Medium	High

Information Technologies in the Age of AI: Possible Contributions to Research in the Intelligence Community

In October 2015, early in his tenure as Deputy Director of the CIA for Digital Innovation, Andrew Hallman talked about the challenges in the field. According to Hallman, the main challenge is "taking the vast volumes of digital intelligence that the CIA receives from around the world and transforming them into a digital, dynamic and credible picture of the future." Hallman adds that "intelligence, in this context, becomes almost a super power" (Tucker, 2015). In June 2019, Hallman noted that artificial intelligence may help intelligence officers improve their ability to focus on their highest value activities, from the automation of routine tasks to the rapid exploitation of data, pattern recognition, and predictive analytics.

The American intelligence community is assimilating and advancing AI technologies—the CIA alone is developing about 140 projects that leverage AI technologies in order to streamline the performance of intelligence tasks such as photograph deciphering, analyses, and predictions. The Defense Department is taking action to form a headquarters called the Joint Artificial Intelligence Center, which will coordinate the efforts to develop and transfer AI technologies to operational uses (CRS, 2019).

There are several possible contributions to research and to intelligence assessments by the three waves of artificial intelligence technology, combined with other familiar information technologies, such as technologies of big data management, data merging, data fusion, and so forth. Some of these capabilities are not yet available at an adequate level for the variety of applications below:

Potential Contributions to Routine Research Activity

Facilitate processing and optimal use of big data: The information age provides researchers with more raw information from a variety of sources,

but the flow of information increases and accumulates at speeds and volumes that exceed their ability to even comprehend the information using existing means and integrate it into the intelligence analysis (Cruickshank, 2020). For example, the United States Army operates more than 11,000 UAVs, each of which takes pictures daily, and the volume of information that they acquire exceeds that of the high-resolution filming of three seasons of the National Football League, but the Defense Department does not have sufficient personnel or an adequate system to comb through this volume of data in order to derive actionable intelligence analysis (CRS, 2019). Artificial intelligence, in combination with big data management technologies, could help researchers and intelligence collection agencies contend with the massive volumes of information from distributed databases by processing, deciphering, and sorting information according to priorities. This technology can be particularly beneficial if there is a connection between the data system in the researcher's computer and those of other researchers and intelligence collection personnel, and other databases—including the information flow from UAVs, satellites, news websites, social networks, research institutes, and more.

AI technology may help present in quick, full, integrated, and ongoing fashion an intelligence analysis and intelligence assessment that will also include implications and predictions.

Improve researcher access to the original information: AI technology may help with language translation, so that researchers will be able to receive original information in various languages directly, thereby reducing their dependence on information collection and processing performed by the collection agencies (Recorded Future, 2019). Machine translations are not yet sufficiently accurate.

Help authenticate information and detect deception through AI's ability to ascertain whether information is authentic or edited, by checking the reliability of the information sources behind the text (checking their computer-documented history of inaccurate reporting of information), and cross-checking the information received with information from other sources.

Identify details, correlations, habits, patterns, and anomalies: We can expect to become able to identify events and behavioral patterns of an enemy swiftly, for example, through learning and monitoring phenomena that are identifiable from aerial photos and in the media, such that AI technology may help researchers assemble a richer and more reasoned intelligence analysis and even issue alerts about anomalous activities.

Facilitate data merging and data fusion: AI will improve these processes. For example, data on attacks on forces will be merged into the intelligence analysis; the system will display the ratio of the enemy's forces that were destroyed and its remaining capabilities (Buhbut, 2020).

Reduce human errors: AI technology may enable a reduction in cognitive errors and natural biases among intelligence personnel—an issue that the intelligence community grapples with constantly (CRS, 2019). For example, at issue are anchoring (when an individual depends too heavily on initial information), bias deriving from collective thinking, the use of partial information, and exaggerated importance of information that supports the researchers' position (Heuer, 2005). In all of these instances, AI technology may provide data and assessments that are immune to typical human bias (unless humans tendentiously introduce bias through the data that they choose to input into the machine), thereby serving as a control over human assessments. Second-wave AI technology may actually challenge researchers' logic, because it bases itself on statistical conclusions. For example, while researchers assess whether there will be a social uprising in a particular country based on the rationale of the situation, the machine

examines whether the data show an anomaly in the norm and whether there is any similarity to past uprisings, regardless of the rationale.

Improve the continuity of intelligence work: Digital systems can function fully and continuously 24/7 when human researchers are unavailable, apart from the skeleton crew whose abilities are inferior to those of the organization during routine workhours.

Help manage the research desk: Researchers in general and on-duty figures in particular may find value in using a virtual assistant (like "Siri" and "Alexa") for the purpose of managing the research desk.

Potential Contributions to Intelligence Research Products

Formulate an intelligence analysis and an intelligence assessment: AI technology may help present in quick, full, integrated, and ongoing fashion an intelligence analysis and intelligence assessment that will also include implications and predictions. This advantage is very significant considering the long and protracted process currently needed to produce a comprehensive strategic intelligence assessment (such as a national intelligence assessment)—which limits the number of these assessments—while leaders are looking to streamline the decision making process. Furthermore, it will be possible to generate reports of changes in the intelligence analysis at different cross-sections and strata. This technology already enables the presentation of force deployments on maps. Technological systems may notify researchers and collection units about locations where there are information gaps or a lack of updated data, so that they can supplement the intelligence analysis. Such systems may continuously integrate an enemy's intelligence analysis (the "red side") with that of one's own forces (the "blue side"), which will contribute to a dynamic situation report.

Provide intelligence and a digital representation of the intelligence to consumers: Consumers may benefit by receiving a rapid and ongoing supply of digital intelligence reports, including dynamic intelligence assessments, and the presentation of the intelligence using advanced visualizations, at any time and according to their needs.

The machine may learn from network activities about anomalous indications of preparations for violent and terrorist activities, and may even provide swift alerts about the start of an outbreak of a wave of violence and terrorist activities, like the second intifada.

Provide alerts about strategic course changes: The critical tests of strategic assessments by intelligence analysts are in identifying strategic course changes and issuing alerts about them in time, as well as in providing assistance with the formulation of strategic decisions and with the management of risks and opportunities (Even, 2017). Machines may help analysts by providing strategic alerts through "learning," the detection of anomalies and changes in the behavioral patterns of populations, leaders, and organizations. For example:

1. Alerts about war or other hostile activity initiated by an enemy: The machine will detect anomalies that may indicate abnormal activity by an enemy and help analysts understand their significance. It will present a picture of the anomalous indications, the correlations between them, and previous occurrences.
2. Alerts about domestic instability or social crises: The technology will enable early detection of a rise in social unrest in various countries or among population groups. This is one of the signs indicating an increased risk to regime stability, as occurred during the "Arab Spring" events (McKendrick, 2019). The CIA Deputy Director for Digital Innovation, Andrew Hallman, said that IARPA (Intelligence Advanced Research Projects Activity), which is subordinate to the

American intelligence community, launched a program in 2011 for the development of methods for continuous automated analyses of publicly available data in order to detect or anticipate major societal events, such as political crises, humanitarian crises, mass violence, riots, mass migrations, disease outbreaks, economic instability, resource shortages, and responses to natural disasters (Tucker, 2015).

3. Alerts about terrorist attacks or an intifada: The machine may learn from network activities about anomalous indications of preparations for violent and terrorist activities, and may even provide swift alerts about the start of an outbreak of a wave of violence and terrorist activities, like the second intifada. The machine can also identify connections among suspects, and between them and people who were not previously suspects (Eichner, 2017).
4. Exposes strategic forgeries or manipulative perceptual attacks: An example is the tactics that the American intelligence community attributes to Russia during the United States presidential elections in November 2016. However, artificial intelligence also enhances an enemy's ability to impersonate a person and launch an attack using fake news and deep fakes. This technological capability was demonstrated in a video clip showing how an actor impersonated former President Barack Obama, using voice and image processing software (Vincent, 2018).

Facilitate identification of security opportunities: AI technology may help with analyses of a broader spectrum of possible actions and propose unexpected decisions that appear to be irrational, but will surprise the enemy and give one's forces the upper hand (CRS, 2019).

Facilitate forecasting through scenarios and simulations: Significant improvement in the ability to present scenarios and simulations can be expected, and perhaps even to present estimates of the probability of scenarios. Such scenarios will enable the presentation of optimal courses of action from an enemy's perspective according to assumptions based on familiar and less familiar patterns of behavior. AI technology will contribute to the presentation of scenarios that integrate intelligence about the enemy with the conduct of one's forces. This may enable leaders to obtain higher quality intelligence assessments and situation assessments quickly, and even more relevant scenarios about wars and their outcomes—before making decisions. From this perspective, technological systems may reduce instances of miscalculation, and constitute a restraining factor against launching or continuing a war. On the other hand, opposing situations are also possible, in which leaders treat these systems as if they were crystal balls and are tempted to launch a military operation based on an optimistic simulation, or instances may occur in which the speed of the AI system's response will contribute to an escalation (see below, risk of losing control and of escalation).

Improve control: AI technology may assist with quality control of the intelligence assessments and situation assessments. The technology may be used to present the intelligence analysis according to the quality and updatedness of the information, and distribute it among intelligence collection agencies and sources, for the purpose of examining their contribution and the extent of dependence on them. Given good documentation, it will also be possible to use it for the purpose of monitoring the performance of the intelligence organization in the field of research and assessment, and to obtain information that will lead to improvement.

Potential Contributions to Improved Integration

Facilitate integration within and between research entities: Intelligence research is currently based on expert researchers working within hierarchical organizational frameworks (departments, divisions, and arenas), when not one has the comprehensive complete picture

at all times. A technological system can provide researchers with a shared, detailed, and up-to-date information and knowledge base. A technological system can help researchers identify at an early stage the emergence of phenomena that are already occurring in other countries, such as the phenomenon of the "Arab Spring," which began in Tunisia in December 2010 and spread throughout the Arab world.

Integration in the intelligence community: The technology will enable a significant increase in the integration between data collection and research systems and organizations within the intelligence community. For example, the data collection and research entities will constantly receive an up-to-date picture of the information gaps and of dynamic flagging of critical intelligence information, including ongoing guidance about areas where information gaps need to be closed. This integration may require interconnectivity between the databases in the community—which covert organizations are loath to do, whether due to considerations of data security or of competition.

Facilitate integration between intelligence agencies and planning and operational entities: AI technology will enable the integration of a dynamic intelligence assessment with a dynamic situation assessment, and will enable planning and operational entities to direct the intelligence agencies to accommodate their needs. This will contribute, inter alia, to an updated diagram of the line of contact during military operations and wars.

If the integration capabilities materialize, the technology may lead to a change in the organizational approach. For example, machines' ability to help with rapid, continuous, and compact integration in the above areas may have an influence on increasing the flexibility of the intelligence community's traditionally rigid structure, in which there is a sharp structural separation between information collection and research, and in security establishments that set up a sharp structural separation between the intelligence agency and the operational division and the planning entities.

Challenges when Assimilating Artificial Intelligence in Research

Along with the numerous advantages in AI technology are difficulties, risks, and requirements for optimal implementation of artificial intelligence in research. All these constitute technological, cultural, and organizational challenges, such as:

Difficulties in Applying Artificial Intelligence in Research

Machines have difficulty "understanding" complex human language: In order to investigate an enemy's intentions, information must be thoroughly understood, including speeches and dialogues by and between leaders speaking in their native languages, as well as cultural nuances. Prof. Yosef Grodzinsky, the head of the neurolinguistics laboratory at the Hebrew University's Center for Brain Sciences, states that "industry has not yet succeeded in engineering machines that understand language, or that even translate language properly." And indeed, it is difficult to find anyone who will rely on Google Translate (which is based on statistical machine translation) to translate a document for legal use (Grodzinsky, 2020). The existing machines are not yet capable of correctly understanding texts containing dual meanings, hints, subtexts, jokes, innuendos, and complex linguistic compositions. Similar to humans, the second-wave development of an AI machine is capable of learning from examples, but unlike humans, it currently needs an enormous quantity of examples, and it is incapable of understanding emotional situations in the same way that human researchers understand them, since they actually experience those emotions.[4] The development of translation software whose output will approximate the output of a human expert translator may herald progress in this direction for intelligence needs.

Difficulties in applying software that is based on statistical deductions in strategic research: Learning capabilities of second-wave machines are based on many examples and on statistical deduction tools. However, most of the major research issues in strategic intelligence cannot be learned from many examples that are difficult to generalize. For example: every war is a unique case and there is no representative sampling of information about wars that would enable a statistical prediction of the timing of the next outbreak of war between countries. This may explain why attempts to assimilate awareness and statistical tools in the work of strategic intelligence researchers in the American intelligence community have been unsuccessful (Tetlock & Gardner, 2017). Furthermore, systems that are based on statistical deduction can present a correlation between variables (for example: a positive correlation between a rise in arson attacks and seasonal temperatures), but they cannot explain why this correlation exists, i.e., what is the logic behind this correlation. The challenge is to find ways in which AI systems containing statistical software can be integrated in research work, and to develop technologies during the third wave that will succeed in learning from fewer examples and succeed in presenting reasoned explanations.

Difficulty of Intelligence researchers to rely on and control AI systems: One reason is that intelligence researchers need explanations and reasons, and they will have a hard time understanding how the machine reached its conclusions, both because the machine's conclusions can be an outcome of many complex actions, including the use of nonlinear functions, and because second-wave learning machines are incapable of providing logical reasons for their conclusions, since they are based on statistical deduction software. Another reason is that most analytical software currently in use are closed systems (black boxes) that do not enable researchers to fine tune them. Intelligence researchers need systems that they can command, reconfigure, or revise the informing algorithms according to the changing reality (Cruickshank, 2020). Furthermore, some researchers may object to machines, just like they object to any innovation that changes their world order, because they consider the possibility that machines might make them redundant or jeopardize their importance in the research enterprise.

Difficulty by consumers to accept AI outputs that cannot be directly explained. The final decisions are still in human hands and, for the most part, humans need to understand the logic behind the intelligence assessment and the recommendation to take action, which machines are still incapable of providing (CRS, 2019). This means that leaders and other intelligence consumers will have a hard time adopting and basing their decisions on intelligence assessments created by machines if they do not comprehend them (Vincent, 2019). Consequently, a significant amount of time will presumably be needed until leaders agree to base their decisions on assessments originating from machines without human mediation—at least until a series of successes of accurate strategic predictions of AI systems can be demonstrated, or until machines are developed that know how to respond to questions and provide detailed explanations that will satisfy decision makers.

Risks to Consider or Mitigate when Applying AI in Research

Overassessment by a machine: If intelligence analysts and decision makers have success with machine outputs, there is a risk that they might rely on them without understanding the logic behind the outcome and stop controlling them (CRS, 2019). The customary humanization of machine capabilities (capabilities of "learning," "drawing conclusions," and so on) and a machine's capability of simulating human dialogue are liable to mislead machine users into attributing extensive capabilities to a machine that it doesn't have. Sources at DARPA say that applications like voice recognition and facial recognition, which are based on

machine learning (from the second wave), have been so successful that people developed the misconception that the computer can just "learn things" (Launchbury, 2017).

Cybersecurity: For the purpose of applying artificial intelligence, the machine is supposed to have access to very large databases, some of which are also open and accessible to enemies that could perform manipulations and deceptions in them that will corrupt the machine's outputs. AI systems themselves are also exposed to cyber-hacking. Artificial intelligence can improve cybersecurity, but it is also liable to improve the attack capabilities of enemies in cyberspace. If an AI system's code is stolen, then within a very short timeframe the attacker will be able to use the system against that entity from which the code was stolen (CRS, 2019), thereby intensifying the magnitude and dispersion of the damage.

Risks of error: The machine itself does not guarantee certain identification, but rather only probable identification (a second-wave system may incorrectly recognize workers with tools as combatants). The machine may err due to errors in development, maintenance, and operation. The machine may encounter situations that its developers had not foreseen, due to a changing reality. Since artificial intelligence is liable to propose conclusions that may be incomprehensible or perceived as irrational, the machine's operators may perform erroneous corrective actions. In addition, it will be difficult to identify and repair the error in the machine and to investigate the source, as long as the machine is incapable of reporting its actions. Furthermore, unlike humans, an AI system is liable to repeat the same mistake a multitude of times at high speed in one or in several machines simultaneously, which could magnify the damage (CRS, 2019).

"Arms race": Artificial intelligence is merely another one of the new battlegrounds for a technology-based arms race (DNI, 2019). Colonel Avi Simon of the IDF's C4I and Cyber Defense Directorate noted that while in the past the technological power sources were held by militaries and governments, today they are in the hands of civilian companies like Facebook, Amazon, and Google, and such technological systems may be purchased online or used in civilian systems (Buhbut, 2020). In such a world, it is not sufficient to learn the behavioral patterns of a person or a population; it is also necessary to research and understand the "thought" pattern (the algorithm) of the enemy's computer, which will learn the thought pattern of the computer on one's own side and so on (the "double mirrors" effect). Consequently, the challenge is to develop, adapt, and implement artificial intelligence at a faster pace than one's enemies. This is also why supervision of exports of sensitive AI technologies is warranted.

Unlike humans, an AI system is liable to repeat the same mistake a multitude of times at high speed in one or in several machines simultaneously, which could magnify the damage.

Losing control and escalation by machines: At the techno-tactical level, the technology already enables autonomous tools to close circuits quickly, which includes: intelligence, decision making, and execution ("sensor to shooter")—such as detecting armed border infiltrators and opening fire, but the restriction on its application without human involvement is mainly an ethical restriction. Humans are liable to be tempted to rely increasingly on an intelligence analysis and on the "discretion" of machines due to their ability to respond rapidly to an enemy's activities, but this may create an increased risk that decisions to operate a strategic weapon might be made in the future without human involvement (Antebi & Dolinko, 2020; Johnson, 2020).

Infringement on privacy: The intelligence uses of artificial intelligence are sometimes at odds with the need to safeguard the public's rights to privacy, particularly in intelligence agencies operating in the domestic arena.

Therefore, on the one hand, a balanced regulatory framework that will improve the protection of civil rights against malign use of these technologies is warranted (Weinbaum & Shanahan, 2018). On the other hand, these systems may actually help safeguard privacy, because they can report phenomena and trends to the researcher without disclosing data on private individuals; or they can focus on exposing only specific people out of a large database, as opposed to the situation whereby human researchers comb through the entire database.

Increased dependence on the civilian sector: Since a significant share of the development of artificial intelligence comes from the civilian sector, the defense sector might develop dependence on it. Furthermore, the machine must undergo modifications before it enters the security system. This requirement, and the rapid development of artificial intelligence, could lead to security entities having to deal with an increase in the complexity of their procurement processes. In addition, civilian developers sometimes object to new uses of their software, due to ethical and/or commercial considerations. These considerations may deter some technology companies from agreeing to cooperative efforts with security entities (CRS, 2019).

It appears that the subject of artificial intelligence should be a new addition to the methodological knowledge required of intelligence researchers. The requirement of digital literacy, including AI literacy, may also affect the nature of the human resources in intelligence organizations.

Organizational Needs when Assimilating AI in Research Units

The need to adapt human resources to the artificial intelligence age: According to Dr. Yoel Mark, Vice President of Research at Amazon, intelligence researchers need to adapt themselves to the new age, to learn the concept of algorithms, to analyze meticulously and understand the outcomes obtained from machines, and to understand how their own actions contribute to the learning machine's optimization (Siman-Tov & Lt. Col. Z., 2018). A similar message is contained in a document from 2019 on behalf of the head of the American intelligence community called "the AIM Initiative." This document, which discusses the increased use of AI for intelligence purposes, states, inter alia, that investments are needed in programs that train and equip the workforce with essential skills for working in an AI environment. This does not mean that every researcher must be an AI expert, but it does mean that everyone must understand how artificial intelligence is integrated in their work and how it can contribute (DNI, 2019). It appears that the subject of artificial intelligence should be a new addition to the methodological knowledge required of intelligence researchers. The requirement of digital literacy, including AI literacy, may also affect the nature of the human resources in intelligence organizations.

In addition to the adjustment processes that investigators will undergo, a change in the composition of the human resources in research units is expected. On the one hand, as a result of the massive introduction of machines into intelligence work, a reduction in the number of intelligence personnel who will be needed in particular fields is expected. On the other hand, this technological evolution calls for the creation of new roles, so that a change in the job descriptions of many roles and an increase in demand for new roles, including experts in the development and operation of AI-based systems, are certain. In light of the considerable demand in the civilian market for suitable human resources, the intelligence organizations must compete on the work conditions in this market (Eichner, 2017). Although it is difficult to compete with the monetary fringe benefits that the business sector offers, the United States National Security Commission on Artificial Intelligence found that AI experts are willing to serve in the government sector in workplaces

offering a more compelling sense of purpose and a technical environment that will maximize their talents (CRS, 2019). The security system can also offer greater occupational stability than the civilian sector.

The need to adapt digital systems: Optimal communications must be created between databases and the machine. Even though AI is known for its excellent work in data processing from distributed databases, it would be advisable to plan and organize the architecture of the digital databases used by the intelligence community in order to fully tap the machine's capabilities (Cruickshank, 2020).

Conclusion

Integrating artificial intelligence in research work may provide intelligence researchers with ever-increasing benefits, in both their routine research activity and formulation and presentation of the research results, and in integration. In the coming decade at least, AI systems are not expected to replace researchers, but rather will continue to serve as tools, for example, that can differentiate between the important and the irrelevant in the information flow and can detect anomalies, correlations, and patterns that will generate important conclusions. The technologies may facilitate integrating the intelligence assessment in the assessment of the situation of forces, and enable intelligence organizations to present scenarios to leaders for predicting the behavior of human actors and high-level strategic course changes at a visual quality and speed that far exceed the capabilities of the past.

In terms of the challenges, the lack of high-quality language processing appears to be a significant constraint in a substantial share of the uses needed for strategic intelligence research; the same goes for the machine's inability to explain or rationalize its findings. These two inadequacies relate to the mode of operation of second-wave AI machines. Ideally a solution will be found in the third wave development. Furthermore, enemies are beginning to use artificial intelligence. The ability to understand the reality and identify risks and opportunities before enemies do may be a highly valuable strategic asset for leaders. The general question that will remain on the agenda is not whether machines will be capable of influencing the presentation of reality to humans and helping them make decisions, but rather, to what extent will humans allow machines to influence and control their world.

In order to advance the development and assimilation of artificial intelligence in intelligence research efficiently, a strategy and plan for integrating artificial intelligence in such research should be prepared. A good example is found in American intelligence. In January

Table 2. Expected contribution of AI technology to intelligence research

Facilitated research activity	Contribution to integration	Contribution to the research products
Processing of large volumes of information from distributed databases	Integration between the various research entities will enable the production of a consolidated intelligence analysis	Broader, more reliable, and continuous intelligence assessments
Information authentication and deception detection	Integration between research entities and collection agencies will facilitate the formulation of a full and dynamic intelligence analysis	More developed scenarios, the creation of various types of alerts
Detection of anomalies, correlations, and patterns	Integration between research entities and "our forces" will create an integrative and dynamic situation assessment	Higher quality research products that are more accessible by consumers

2019, the US Office of the Director of National Intelligence (ODNI) published the AIM strategy (A Strategy for Augmenting Intelligence Using Machines). This initiative includes four key AI objectives (DNI, 2019):

a. Immediate and ongoing: to create a digital foundation of information using artificial intelligence and automation processes, and to revamp the workforce in the intelligence community.
b. Short term: to adopt commercial technology solutions from the private market, particularly AI technologies, and capabilities of exploiting overt sources.
c. Medium term: to develop technological capabilities that will close the remaining gaps, so that the American intelligence community will have a strategic advantage over anyone engaging in intelligence.
d. Long term: to invest in the development of joint human–machine analysis capabilities.

*

The authors thank Prof. Yesha Sivan for his helpful comments on this article. Prof. Sivan is a visiting professor of digital, innovation, and venture capital at the Chinese University of Hong Kong Business School, and is the CEO of i8 Ventures.

Dr. (Col. ret.) Shmuel Even is a senior research fellow at INSS and a strategic advisor, and has extensive experience as a director of business companies.

Lt. Col. (res.) David Siman-Tov is a senior research fellow at INSS in the program National Security and Democracy in an Era of Post-Truth and Fake News, and Deputy Director of the Institute for the Research of the Methodology of Intelligence at the Israeli Intelligence Community Commemoration and Heritage Center.

References

Antebi, L., & Dolinko, I. (2020). Artificial intelligence and policy: A review at the outset of 2020. *Strategic Assessment, 23*(1), 94-100.

Buhbut, A. (2020, March 7). Artificial intelligence and creating targets in seconds: This is how the IDF prepares for the digital age. *Walla News!* https://news.walla.co.il/item/3341866 [in Hebrew].

Congressional Research Service (CRS). (2019, updated 2020). Artificial intelligence and national security. CRS. https://fas.org/sgp/crs/natsec/R45178.pdf

Cruickshank, I. J. (2020, February 14). The ABCs of AI-enabled Intelligence Analysis. *War on the Rocks*. https://bit.ly/366EXFF

Director of National Intelligence (DNI). (2019, January 16). The AIM initiative: A strategy for augmenting intelligence using machines. DNI. https://bit.ly/3hZO4ds

Eichner, A. (2017, January 18). The Google of the GSS. *Yediot Ahronot*. https://www.yediot.co.il/articles/0,7340,L-4909040,00.html [in Hebrew].

Even, S. (2017). From a national intelligence assessment to a national risk assessment. In S. Even and D. Siman-Tov (Eds.), *The challenges of Israel's intelligence community* (pp. 21-31). Tel Aviv: Institute for National Security Studies [in Hebrew].

Grimland, G. (2018, May 27). Intel explains: Six AI terms. *GeekTime*. https://www.geektime.co.il/private_channel/6-ai-topics-you-need-to-know/ [in Hebrew].

Grodzinsky, Y. (2020, February 28). Industry has not succeeded in engineering machines that understand language. *Haaretz* [in Hebrew].

Hallman, A. (2019). CIA's Andrew Hallman discusses digital futures at FedTalks 2019. CIA. https://bit.ly/3kJYKyN

Heuer, R. J., Jr. (2005). *Psychology of intelligence analysis*. Maarachot Publishing [in Hebrew].

Johnson, J. S. (2020). Artificial intelligence: A threat to strategic stability. *Strategic Studies Quarterly, 14*(1), 16-39.

Launchbury J. (2017). A DARPA perspective on artificial intelligence. *Technica Curiosa*. https://bit.ly/32N5u8P. See also YouTube, https://bit.ly/35TIqY1

McKendrick, K. (2019). Artificial intelligence prediction and counterterrorism. Royal Institute of International Affairs.

Recorded Future Team (2019, January 9). How artificial intelligence is shaping the future of open source intelligence. *Recorded Future*. https://bit.ly/2Ru0i4z

How to ensure artificial intelligence benefits society: A conversation with Stuart Russell and James Manyika. (2020, January 31). McKinsey Global Institute. https://mck.co/302Qn9H

Siman-Tov, D., & Lt. Col. Z. (2018). The big data revolution from the viewpoint of mega organizations: Interview with Dr. Yoel Mark, Vice President of Research at Amazon. *Intelligence in Theory and Practice, 3*, 33-35. https://bit.ly/362tYx1 [in Hebrew].

Tetlock, P. E., & Gardner, D. (2017). *Superforecasting: The art and science of prediction*. Kinneret, Zmora-Bitan, Dvir [in Hebrew].

Tucker, P. (2015, October 1). Meet the man reinventing CIA for the big data era. *Defense One*. https://bit.ly/37yGUce

US Congress (2019). John S. McCain National Defense Authorization Act (NDAA). Section 238, pp. 62-63. https://www.acq.osd.mil/eie/Downloads/IE/NDAA%2019%20BILLS-115hr5515enr.pdf

Vincent, B. (2019, May 31). How the CIA is working to ethically deploy artificial intelligence. *Nextgov*. https://bit.ly/35XO6Aq

Vincent, J. (2018, September 14). US Lawmakers say AI deep fakes "have the potential to disrupt every facet of our society." *The Verge*. https://bit.ly/2rtbrHU

Weinbaum, C., & Shanahan, J.N.T. (2018, July 3). Intelligence in a data-driven age. *Joint Force Quarterly, 90*. https://bit.ly/33mg0B2

Notes

1 For example: A second-wave system designed to classify photographs can identify a photograph of a particular tank. If the machine had the ability to speak, it would have said: after learning from a huge pool of examples and based on computations, it is highly likely that this is a T-72 tank (of Russian manufacture). But it cannot reason, because the identification process was not conducted according to tank features that researchers recognize: typical silhouette, particular continuous track system, 125 mm cannon, reactive defense, particular machine guns, and so on. That is, the machine has statistical recognition ability but it does not have the ability to explain the result logically.

2 For example: To train the system in handwriting recognition, you need to input into the system 50,000-100,000 examples. Sometimes it is difficult to provide relevant examples at such large quantities. By the way, humans can learn from a few examples, and the third-wave technology is striving to achieve this.

3 The reasoning of second-wave systems is lower than that of first-wave systems because first-wave systems generate output based on a logical algorithm that is usually understood by researchers. On the other hand, second-wave systems are based on statistical conclusions and on functions that are not familiar to the user.

4 In principle, although artificial intelligence will be able to detect signs indicating people's emotional situations, such as a demonstration of anger by learning examples of texts, vocal frequencies, facial expressions, and body language, it is highly unlikely that machines will be able to interpret all nuances of human cognitive and emotional behavior.