Think
new things
Make
new connections

Terms of Reference

Artificial intelligence and national security in 2025

27 February-1 March 2025

DITCHLEY

Terms of Reference

Summary

As we enter 2025, artificial intelligence (AI) stands as perhaps the most transformative force in global security, potentially reshaping how nations protect their interests and respond to threats. The past two years have seen an unprecedented acceleration in AI capabilities, particularly in multimodal systems that integrate text, vision, speech, and, to a gradually increasing degree, more complex decision-making. This evolution marks a shift from AI as a useful productivity tool to a strategic enabler that could fundamentally alter the character of national security challenges. Democratic nations find themselves at a crossroads where AI presents both extraordinary opportunities and profound vulnerabilities. The ability of allied democracies to coordinate their efforts across national boundaries will be as critical as their technological capabilities in dealing with this rapidly changing, globally connected landscape.

This Ditchley conference will explore the dilemmas that AI poses for national security, bringing together leaders from diverse fields to assess emerging threats, explore responses, and map the infrastructure, capabilities and skills that will be essential for strategic advantage, while respecting democratic values. We will look at how states and the private sector will need to collaborate, within nations and across borders. The discussion will build on Ditchley's earlier AI and defence meetings but also the *Strategic Technologies Investors' Council* discussion that Ditchley staged in November 2024, which brought together officials and venture capital and private equity investors from the US and UK.

Detail

Emerging threats and opportunities

The threat landscape has grown increasingly complex as AI capabilities mature. Many commentators pointed to AI-enabled disinformation and manipulation of elections as a serious risk for 2024. Modern AI systems can generate persuasive, tailored content that adapts to different audiences and contexts. Emerging capabilities like OpenAI's Sora text to video generation suggest that soon we will struggle to distinguish truth from fiction. In theory that should enable hostile actors to conduct influence operations at unprecedented scale and precision. But so far AI does not seem to have had the major impact on the information landscape that many expected. Although there certainly were efforts to influence elections using AI, most do not seem to have been decisive, with the possible exception of Romania and the extraordinary annulment of the election. Some think this is because that is not how disinformation works: when people believe misleading information, they are applying their own biases and a crude cartoon can trigger those as easily as a hyper-realistic fake video. Equally, if a true high-quality video, image or story does not conform to someone's biases, then it is more likely to be dismissed. Will increasingly powerful AI capabilities change this picture? Are we understating the impact of AI in 2024?

AI's ability to analyse vast amounts of information in parallel and to join dots that human intuition cannot identify is already changing the nature of intelligence assessment and decision making. There is a risk that our adversaries will develop clearer maps of our critical systems and infrastructure and their interaction with the broader connected and open society we have created, than we have ourselves. They may come to understand the interdependencies and vulnerabilities of our critical systems better than we do. For example, energy grids, water systems, sewage systems, communication networks, transport networks, trade routes and city management now all depend heavily on digital systems that might be disrupted. But they also rely on human maintenance and troubleshooting teams to fix problems. People working on sensitive projects or classified work also have family lives online that could be exploited. How well do we understand our connected vulnerabilities? Democratic values are not consistent with a surveillance society that undermines freedom and privacy and so how can we map our domestic vulnerabilities without unacceptable intrusion?

What are the prospects for AI enabled intelligence and criminal operations, whether in the digital or human realms? Advanced language models are beginning to produce sophisticated automated persona going far beyond Chatbot interaction. What are the prospects for personalised digital engagement at scale of individuals working in sensitive positions? Might AI enable orchestrated personalised phishing attacks at scale to gain access to computer systems? Could AI-powered signals intelligence systems detect patterns in communications and vulnerabilities in computer networks at a pace and scale we have not yet seen? Is there a risk of a fundamental attackers' advantage emerging, where bad actors can take risks in the early deployment of such AI systems that good actors cannot for fear of breaching security through AI error (ie defensive AI own goals)? What does this mean for cyber defence and deterrence?

The ability automatically to identify and track targets across multiple data domains has compressed decision-making times and created new vulnerabilities in traditional security frameworks. An emerging example is the integration of AI powered analysis with earth observation imagery, now widely commercially available. How will global visibility, combined with trainable analytic systems[1], affect defence systems and national security? What will increased speed of analysis mean for responsible decision making?

AI enabled weapons systems are already in use. Evolving American military doctrine specifies a human "on the loop", rather than in the loop, meaning that a senior officer should take responsibility for launching an automated system. Autonomous weapons systems are technically possible now. Is international agreement to prohibit their deployment sustainable? Would such agreement survive a conflict where autonomous systems prevailed against human controlled systems? What are the implications for terrorism, whether by state or non-state actors? What does the development of AI-enabled weapons systems mean for deterrence?

[1] As an example the commercial Earth observation satellite imagery company, Planet, has shown how a trainable AI enabled analytical interface could be trained in a straightforward way to improve its recognition of Chinese weather (alleged spy) balloons as identified in the controversial course of one such balloon across the US. Once trained in spotting the balloon in imagery then analysts were able to follow the balloon back in time through the daily imagery to its launch balloon in China. This then allowed identification of other balloons and their routes that had been launched from the same site.

AI-enhanced biotechnological threats**,** such as precision bioweapons targeting specific populations, represent another growing concern at the overlap of AI and genetic engineering. Advances in personalised healthcare and genetically engineered vaccines could also lead to create new deadly viruses and toxins. How will biosecurity need to evolve?

The risk of strategic surprise has intensified through several mechanisms. "Silent" capability development makes it increasingly difficult to detect adversary advances until they are used. The democratisation of AI technology through open-source innovations further challenges conventional models of strategic stability, as significant capabilities can now be developed outside traditional state structures. The private sector is now leading in most areas of innovation and many commercial technologies are dual use. How can arms control be effective in this new environment?

Democratic Response and Essential Actions

Meeting these challenges requires a response from democratic nations while protecting our fundamental values, which we might think of as falling into four main areas.

- Technical Sovereignty: This will mean developing indigenous AI capabilities, secure supply chains, and resilient interdependence among trusted partners. Establishing strategic computing reserves and frameworks for resource sharing will mitigate the risks of dependency on adversarial states, ensuring that democracies retain strategic control over critical capabilities. Will the US and US companies remain willing to share AI capabilities with allies? How can allies contribute to the development of a capabilities in a way that wins and deserves a seat at the table?

- Strategic AI Governance: How can we develop governance frameworks that ensure ethical and accountable use of AI in national security contexts while keeping pace with adversaries? What channels of negotiation do we need to agree international guidelines? Do we need to control AI in the same way that the development of nuclear weapons has been restricted, with an AI equivalent of the Treaty on the Non-Proliferation of Nuclear Weapons (NPT) and inspections by the International Atomic Energy Agency (IAEA)?

- Public-Private Partnerships: The private sector is central to AI innovation. How can we build the partnerships we need? How can we manage competition and industrial policy whilst continuing to rely on market-based innovation? Traditional procurement cycles must evolve to accommodate the pace of AI innovation. How do we create a wider and more innovative supplier base for defence and national security? Democracies must prioritise interoperability among allied forces to avoid capability gaps that adversaries could exploit. How can NATO manage coordinated procurement of AI across allies, for example?

- Information and Societal Resilience: What do we need to do to strengthen societies against adversaries' manipulation of information and social networks? What is the right narrative for governments that sets awareness but does not sow panic? How can we continue to increase digital literacy across all ages?

Infrastructure and Capability Requirements

How can we develop the AI national security capabilities and supporting infrastructure that democracies will need to achieve technical sovereignty, as noted above?

- Energy Infrastructure: Sustainable and resilient energy infrastructure will be essential to power AI systems. This may require investment in next-generation nuclear and renewable technologies, distributed energy networks, and redundant power systems for critical installations. How can AI's energy requirements be met in a way that reconciles with our carbon reduction goals (and/or a focus on clean air and water)? What can the US' allies do to support at pace the development of energy for AI training and inference?

- Training AI: So far, the US and China have close to a monopoly on the increasingly vast capacity data centres and associated investments in GPUs and energy needed to train AI at the frontier level. How can we expand access to AI compute and useful, curated data sets to increase innovation across democratic allies? (Ditchley's recent Strategic Technologies Investors' Council meeting put forward the concept of "TRAIN", a transatlantic AI network combining GPU access and data sets, to enable greater innovation across allies.)

- Computing Resources: Scaling GPU and TPU capabilities is critical for AI training, while edge computing solutions enable tactical applications. Democracies must also develop secure AI model-sharing frameworks and strengthen supply chain security for key hardware components. How can we coordinate across allies? Quantum computing is also moving forward if not at the pace of large language model AI: what do we need to do to prepare?

- Cyber Resilience: Robust defences against adversarial manipulation are critical to ensuring AI systems remain secure and reliable. This includes protecting underlying data and infrastructure from sabotage, espionage, and cyberattacks. How can we protect the AI capabilities we build from digital attack?

For the middle part of the conference we will split into three groups so as to be able to discuss and work through some of these issues in more detail.

Group A will explore what AI capabilities and supporting infrastructure democratic nations and democratic alliances will need and how to the develop them?

Questions to address would include where and how we should build data centres and energy supply for AI training and interference for national security? What kinds of public and private partnership are needed? Should the pursuit of AGI become a nationally or internationally governed project? How can middle powers raise their game on innovation on AI for national security purposes? How can AI be integrated into NATO and other democratic projects such as AUKUS? What kinds of AI applications should we develop as democratic countries: cyberwarfare tools? Machine scale rapid analytical tools? Automated or autonomous weapon systems? Are there areas that should remain permanently off limits for AI development and deployment: bioweapons? Nuclear weapons? How will the development and deployment of AI systems change needs for talent and skills for national security? How can national security keep pace with the private sector, given its necessary focus on security and risk in its processes?

Group B will focus on defence and deterrence in the age of AI and the international management of the AI arms race and the establishment of norms and ethics

Questions could include the impact of the deployment on AI on deterrence. What AI capabilities could act as a form of deterrence? What AI capabilities could undermine traditional deterrence, for example identification of locations of strategic weapons systems, submerged submarines etc? How will defence and national security doctrine and practice need to evolve? What international backchannels, negotiations do we need to begin to address the implications of AI? Should we be working towards treaties that set norms and limits? As a technology being developed primarily in the private sector, to what extent should governments intervene as the technology becomes more powerful? Are some desirable international limits to the use of AI already clear, for example the development of bioweapons through genetic engineering? How can the international community police such limits?

Group C will look at societal resilience and cohesion

What is the right narrative for democratic governments on the threats and opportunities of AI for defence and national security that would increase resilience but not sow panic? Some citizens are already worried that AI will take their jobs. How can democratic governments assure citizens that they have a sensible plan to cope with the unintended consequences and risks of AI, as well as developing the AI capabilities required for defence and security? As systems become more powerful, then how much risk should private sector companies own and how much risk do they want to own? What are the roles for international organisations and civil society in campaigning for effective governance of AI? What digital tools, data sources and digital infrastructure does society need to be resilient in the age of AI?