SUMMARY

22 April 2026

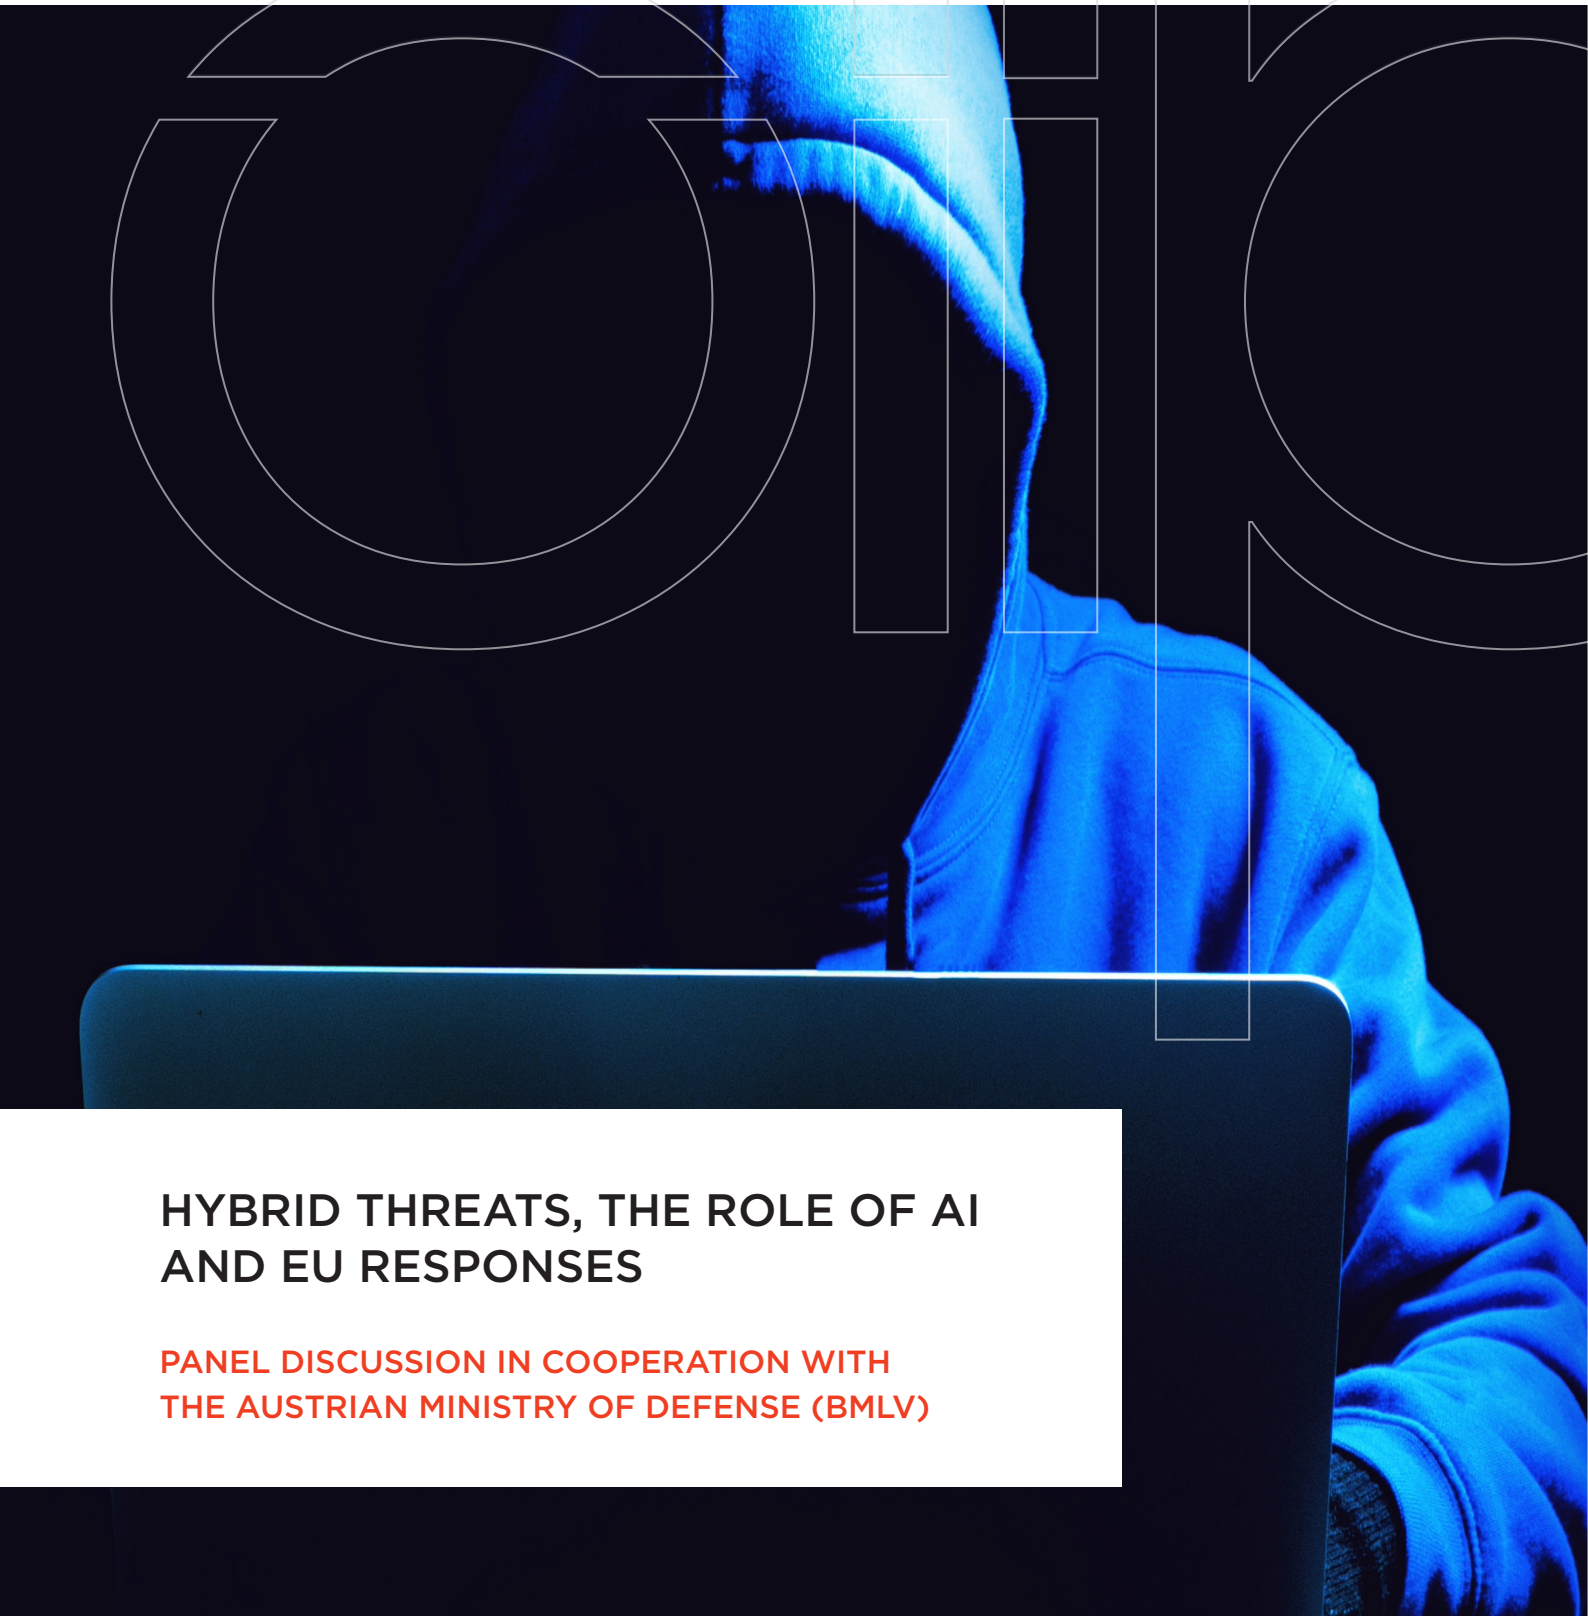

HYBRID THREATS, THE ROLE OF AI AND EU RESPONSES

PANEL DISCUSSION IN COOPERATION WITH THE AUSTRIAN MINISTRY OF DEFENSE (BMLV)

Österreichisches Institut für Internationale Politik
Austrian Institute for International Affairs

PANELISTS:

Miron Lakomy
Professor at the Institute of Political Sciences, the University of Silesia, Poland, and a non-resident fellow at the ITSTIME research centre in Italy
Presentation focus: terrorist exploitation of artificial intelligence, especially Salafi-jihadi actors, Islamic State/ISKP, al-Qaeda, and selected far-right uses of AI.

Ali Fisher
Explorer of Extreme Realms at Human Cognition Ltd., Lecturer at Università Cattolica del Sacro Cuore in Milan, and Fellow of the ITSTime research
Presentation focus: AI, hybrid threats, influence operations, online networks, and human-centred responses to adversarial communication strategies.

MODERATION:

Daniela Pisoiu
Senior Research Fellow at the Austrian Institute for International Affairs (oiip)

DATE:

April 22nd, 2026

VENUE:

oiip, Währinger Straße 3/12, 1090 Wien

Summary report author:

Carter Barnsback
Intern oiip

KEY TAKEAWAYS

1 AI is increasingly relevant to terrorism and hybrid threats, but it should not be treated as the core problem in isolation. Rather, AI is a tool used by human actors to pursue existing objectives: propaganda, recruitment, evasion, radicalisation, influence, and operational support. European institutions have begun to recognise the risks of malicious AI use, including through the Digital Services Act and the AI Act, but EU-level regulation alone cannot solve a global technological problem.

2 Terrorist organisations face several barriers when adopting AI: ideological limits, technical skill gaps, credibility concerns, platform security-by-design, and risks of exposure to law enforcement or intelligence agencies. In Salafi-jihadi circles, official media channels have been more cautious, partly because synthetic content may undermine their claim to authenticity and truth. Unaffiliated followers, however, appear more willing to use AI tools to generate, enhance, remix, and distribute extremist content.

3 Current terrorist use of AI is still uneven and often less revolutionary than public debate suggests. The most important uses are not necessarily fully synthetic propaganda, but AI-assisted enhancement of older human-made propaganda, production of social-media-friendly images and videos, evasion of automated moderation, and individual empowerment of supporters who previously lacked the skills or resources to contribute to "media jihad." Therefore, AI strengthens the persistent online presence of extremist movements without replacing older forms of propaganda.

4 Hybrid threats and influence operations remain fundamentally network-based. Content only matters if it reaches human audiences, moves through communities, and changes behaviour regardless of whether it is AI-generated or not. Ali Fisher argued against "AI tunnel vision": focusing only on whether content is AI-generated misses the more important question of how information travels through networks, how communities interpret it, and how adversaries exploit existing habits, grievances, and online clusters.

5 Effective responses should combine technical detection with a human-centred strategy. The speakers emphasised network analysis, security-by-design, cross-platform monitoring, and global cooperation. They also argued that resilience cannot be built only through counter-messaging. It requires listening to targeted communities, recognising their agency, reducing marginalisation, and building collaborative strategies in which the audience is not treated merely as a passive target of influence.

WICHTIGSTE ERKENNTNISSE

1 Künstliche Intelligenz (KI) gewinnt zunehmend an Bedeutung für Terrorismus und hybride Bedrohungen, sollte jedoch nicht isoliert als das zentrale Problem betrachtet werden. Vielmehr ist KI ein Werkzeug, das von menschlichen Akteuren genutzt wird, um bestehende Ziele zu verfolgen: Propaganda, Rekrutierung, Umgehung von Kontrollen, Radikalisierung, Einflussnahme und operative Unterstützung. Europäische Institutionen haben die Risiken des missbräuchlichen Einsatzes von KI bereits erkannt, unter anderem durch den Digital Services Act und den AI Act. Eine Regulierung auf EU-Ebene allein kann jedoch ein globales technologisches Problem nicht lösen.

2 Terroristische Organisationen stehen bei der Nutzung von KI vor mehreren Hürden: ideologische Einschränkungen, technische Kompetenzlücken, Glaubwürdigkeitsprobleme, sicherheitsorientiertes Plattformdesign sowie das Risiko, von Strafverfolgungs- oder Nachrichtendiensten entdeckt zu werden. In salafistisch-dschihadistischen Kreisen agieren offizielle Medienkanäle vorsichtiger, auch weil synthetische Inhalte ihren Anspruch auf Authentizität und Wahrheit untergraben könnten. Unabhängige Anhänger hingegen scheinen eher bereit zu sein, KI-Tools zu nutzen, um extremistische Inhalte zu erstellen, zu verbessern, zu remixen und zu verbreiten.

3 Der derzeitige Einsatz von KI durch terroristische Akteure ist noch uneinheitlich und oft weniger revolutionär, als die öffentliche Debatte vermuten lässt. Die wichtigsten Anwendungen sind nicht unbedingt vollständig synthetische Propaganda, sondern vielmehr die KI-gestützte Aufbereitung bestehender, von Menschen erstellter Inhalte, die Produktion sozialmedientauglicher Bilder und Videos, das Umgehen automatisierter Moderation sowie die individuelle Befähigung von Unterstützern, denen zuvor die Fähigkeiten oder Ressourcen für einen „Medien-Dschihad“ fehlten. KI stärkt somit die anhaltende Online-Präsenz extremistischer Bewegungen, ohne ältere Formen der Propaganda zu ersetzen.

4 Hybride Bedrohungen und Einflussoperationen bleiben im Kern netzwerkbasiert. Inhalte sind nur dann relevant, wenn sie menschliche Zielgruppen erreichen, sich innerhalb von Gemeinschaften verbreiten und Verhalten beeinflussen – unabhängig davon, ob sie durch KI erzeugt wurden oder nicht. Ali Fisher warnte vor einem „KI-Tunnelblick“: Die ausschließliche Fokussierung darauf, ob Inhalte KI-generiert sind, verfehlt die wichtigere Frage, wie sich Informationen durch Netzwerke bewegen, wie Gemeinschaften sie interpretieren und wie Akteure bestehende Gewohnheiten, Konfliktlinien und Online-Strukturen ausnutzen.

5

Wirksame Gegenmaßnahmen sollten technische Erkennung mit einer menschenzentrierten Strategie verbinden. Die Referenten betonten die Bedeutung von Netzwerkanalyse, „Security by Design“, plattformübergreifendem Monitoring und globaler Zusammenarbeit. Zudem argumentierten sie, dass Resilienz nicht allein durch Gegenbotschaften aufgebaut werden kann. Sie erfordert das Zuhören gegenüber den betroffenen Gemeinschaften, die Anerkennung ihrer Handlungsfähigkeit, den Abbau von Marginalisierung sowie die Entwicklung kooperativer Strategien, in denen das Publikum nicht nur als passives Ziel von Einflussnahme betrachtet wird.

INTRODUCTION

The rapid expansion of generative artificial intelligence since the public emergence of large language models at the end of 2022 has created new opportunities for legitimate users, but also new forms of misuse. The panel opened from a European perspective, noting that EU and national actors have increasingly framed AI misuse as a security concern. Examples mentioned included debates in the European Parliament on AI misuse in disinformation and sexual offences, concerns about malicious exploitation of AI by EU staff, and Europol's attention to AI in the context of violent extremism and terrorism.

In his initial statements, Miron Lakomy argued that the terrorist use of AI is already visible, but not yet uniform or fully strategic. Different extremist milieus have adopted AI at different speeds. Far-right actors appear to have embraced it more readily than far-left actors, while Salafi-jihadi groups have faced theological and credibility-related questions. Ali Fisher placed this issue in the broader context of hybrid threats and influence operations. His central argument was that AI matters, but only in relation to human networks. Influence operations do not succeed because content exists; they succeed when content reaches specific communities and affects behaviour.

The panel addressed three main questions: how terrorist actors currently use AI; how AI changes the logic of extremist propaganda and hybrid threats; and what kinds of responses are likely to be effective.

AI MISUSE HAS BECOME A EUROPEAN SECURITY CONCERN

The discussion began with Miron Lakomy began by situating AI misuse within the European regulatory and security debate. He noted that European actors have already identified AI-related risks in disinformation, sexual offences, and terrorism. He referred to the Digital Services Act as one early attempt to address disinformation on mainstream platforms and to the AI Act as an instrument that introduces obligations such as labelling synthetic content and prohibiting some unacceptable AI uses.

At the same time, he described these measures as only initial steps. The challenge is not solely that harmful content may be generated more easily. It is that AI can alter how extremist actors radicalise online communities, avoid moderation, and maintain visibility across platforms. Europol's recent reporting was presented as recognising that violent extremist organisations are already experimenting with AI and that classic approaches to online terrorism are becoming less effective.

Regarding Europe's role in curtailing the threat of AI, Lakomy's conclusion was cautious. European regulation is necessary, but insufficient on its own. Since AI is a global technology, strict EU rules cannot by themselves shape how the technology is used in the Middle East, Africa, or other regions where terrorist groups and supporters may operate. He therefore called for global solutions, especially around security-by-design in AI development.

TERRORIST ADOPTION OF AI IS LIMITED BY IDEOLOGY, SKILLS, CREDIBILITY, AND SECURITY RISKS

A central part of Lakomy's presentation focused on the determinants of terrorist AI use

identifying several factors that shape whether and how extremist actors adopt AI.

First, there are ideological thresholds. In Salafi-jihadi communities, the permissibility of generating images of human beings or using chatbots to process religious texts has been contested. This partly explains why official Salafi-jihadi media structures initially appeared hesitant. The example of an AI-generated avatar used in an Islamic State-linked newsflash after the Moscow attack illustrated this point: followers reportedly criticised the producer because the avatar's face was visible, after which the face was covered in later versions.

Second, there are skill-related thresholds. Basic chatbot use is easy, but more advanced AI applications require technical skills, computer-science knowledge, or operational capacity. Not all terrorist organisations possess those resources or capacities.

Third, synthetic content can create credibility problems. Many terrorist organisations claim to be providers of truth. If their propaganda is openly artificial or fabricated, it may undermine their claim to authenticity. This is one of the reasons why official media channels still tend to rely heavily on human-made content, while unaffiliated followers are more active in using AI-generated or AI-enhanced material.

Fourth, mainstream AI platforms create security risks for terrorist users. Subscriptions, payment systems, account requirements, and possible law-enforcement access to chatbot conversations may discourage official actors from relying on major commercial tools. This makes locally installed open-source models more attractive, especially where they lack strong safety features.

AI ENABLES PROPAGANDA, CONTENT ENHANCEMENT, KNOW-HOW ACQUISITION, AND MODERATION EVASION

The conversation then moved from constraints to actual uses of AI. Lakomy emphasised that the greatest attraction of AI for terrorist actors is the possibility of mass-producing propaganda. AI can reduce time, skill, and resource constraints. One person can generate text, images, video, and audio at a scale that previously required larger media structures.

Lakomy also addressed the use of AI to obtain sensitive know-how on technical information. He mentioned research showing that chatbots can be jailbroken to provide instructions related to explosives, toxins, or cybersecurity. However, he qualified this point: in Salafi-jihadi circles, older human-made manuals remain more important and often more useful than AI-generated instructions.

More significant, in his view, is content generation and especially content enhancement. AI-generated images, videos, and audio circulate widely in extremist social-media chatter, though less often through official channels. Examples included AI avatars, deepfake-style audio, images of lions and horsemen linked to Islamic State symbolism, and AI-generated Christmas-related violent imagery. Yet the speaker argued that the more important trend is the enhancement of older propaganda. Existing videos, images, and symbols are repackaged with AI effects, filters, or platform tools such as CapCut, an AI video editor, to make them more appealing to younger social-media audiences.

AI can also support evasion. Lakomy described techniques using optical illusions or "illusion

diffusion" to transform recognisable extremist images into landscapes or objects that may bypass automated moderation and even confuse human moderators. Another technique embeds extremist images or videos into synthetic room scenes, such as a television screen in the background, making it harder for hash-matching systems to detect known terrorist content.

UNOFFICIAL SUPPORTERS MAY BENEFIT MORE THAN OFFICIAL ORGANISATIONS

One of the strongest points emphasized by Lakomy was that AI currently empowers unaffiliated supporters more than central terrorist media offices. Official propaganda outlets still have organisational structures, skills, established visual styles, and credibility concerns. They therefore often continue to rely on human-made material.

Individual supporters are in a different position. They may have ideological motivation but lack time, design skills, technical ability, or resources. AI reduces these barriers, allowing actors to produce synthetic content, enhance existing propaganda, remix material, create coded signals of support, and participate in online propaganda ecosystems.

In the Q&A, Lakomy linked this to the broader individualisation of terrorism. AI makes it easier for isolated sympathisers to manifest support online, especially on mainstream platforms such as TikTok. He referred to research identifying large numbers of Islamic State-linked or Islamic State-adjacent profiles on TikTok, ranging from passive viewers to active promoters. AI does not create this ecosystem from nothing, but it makes participation easier.

HYBRID THREATS ARE ABOUT HUMANS, NETWORKS, AND BEHAVIOUR

The conversation then shifted the focus from terrorist use of AI to hybrid threats and influence operations more generally. Ali Fisher argued that definitions of hybrid threats often become overly broad, combining overt and covert action, military and non-military instruments, state and non-state actors, political influence, propaganda, and security disruption. What connects them, in his view, is the human being behind the screen.

His core argument was that AI must be understood in relation to human producers and human recipients. Technology has evolved faster than human cognitive capacity. People still rely on habits, trusted sources, familiar communities, repeated online paths, and social networks to process information. Influence operations exploit those habits.

This means that information does not simply broadcast into society. It flows through networks. Fisher used examples from Twitter around the killing of Osama bin Laden, the Arab Spring, large dark-web mapping projects, and state-linked Russian and Iranian influence operations to show how information clusters form and how adversaries target communities rather than abstract publics.

For present-day hybrid threats, the implication is clear: AI-generated content only matters when it enters a network, reaches a target audience, and produces a behavioural effect. What is important is not whether the video, post, or image was generated by AI. Rather how content travels, who engages with it, and whether it changes what people do.

AI AMPLIFIES EXISTING MULTI-PLATFORM STRATEGIES RATHER THAN REPLACING THEM

AI is a legitimate tool not to be underestimated with arguments of its novelty, a point Fisher warned against exaggerating. Extremist and hostile influence actors already used automation, multi-platform distribution, and community targeting before today's generative AI boom. He used Salafi-jihadi online activity as an example. During Ramadan, al-Qaeda reportedly produced around ten hours of new video and Islamic State around three and a half hours, with little or no AI generation. The point was that these groups still need to show humans doing human things: fighting, speaking, celebrating, displaying spoils, and performing identity. AI may make distribution and repackaging more effective, but it does not remove the need for human performance and human reception.

Fisher also stressed the multi-platform character of extremist dissemination. Salafi-jihadi actors have long used hundreds of platforms. One example involved content distributed almost simultaneously across Telegram, Rocket, Matrix, Signal, WhatsApp, websites, darknet mirrors, and file-sharing platforms. Since the same file appeared across platforms in cryptographically identical form, the speaker inferred that the operation was coordinated before upload rather than casually reposted from one platform to another.

This matters for disruption. Once content has reached the primary audience across several platforms, removal becomes cleanup rather than prevention. AI may make dissemination more efficient, but the underlying problem of fast, coordinated, cross-platform release already exists.

RESPONSES SHOULD AVOID AI TUNNEL VISION

Both Lakomy and Fisher converged on the need to avoid over-focusing on AI as such. Lakomy compared the current AI debate to earlier hype cycles around social media, the metaverse, and other technologies. While Fisher compared AI to 3D printing in that it will be highly important in some fields, but it will not transform everything equally.

This has policy consequences. A response that only detects AI-generated content may miss human-made extremist content, older propaganda, platform migration, and community dynamics. Conversely, a response that treats every AI-generated item as strategically important may overstate irrelevant material that never reaches a meaningful audience.

Rather than focusing on AI outputs alone, Fisher suggested that responses to hybrid threats should start from the desired behavioural outcome. The goal is not always to say the opposite of what the adversary says. Some communities may respond to direct correction, but others may respond better to being heard, included, or engaged around their own priorities.

HUMAN-CENTRED RESILIENCE AND COLLABORATIVE STRATEGIES

The conversation concluded with a focus on resilience. Fisher argued that target audiences should not be treated as passive objects fought over by "us" and the adversary. They have agency, goals, beliefs, and grievances. If they are marginalised or ignored, they are less likely to cooperate with institutions seeking to protect them from hostile

influence. He therefore proposed a collaborative strategy cooperating with the communities being targeted by adversarial influence. In practice, this may include listening, exchange, facilitating community goals, and building trust before attempting direct counter-messaging.

This approach relies on employing both technical tools and human resources. Technical systems are needed to identify information flows and recognise harmful activity, but human actors are needed to interpret the results, engage communities, and decide when to respond directly and when to build longer-term resilience. Summarily, AI will be important, but the tools to fight hybrid threats will remain mostly human for the foreseeable future.

IMPRESSUM:

Österreichisches Institut für Internationale Politik – oiip
Austrian Institute for International Affairs
A-1090 Vienna, Währinger Straße 3/12
www.oiip.ac.at, info@oiip.ac.at

www.oiip.ac.at

Österreichisches Institut
für Internationale Politik
Austrian Institute
for International Affairs