AI for Peace

Noa Garcia Leonardo Impett Yannis Kalantidis Matt Mahmoudi
Evangelos Kazakos Sonia Fereidooni

1 Workshop summary

Artificial intelligence (AI), including machine learning and computer vision, has many historical roots in research for milliary applications [5, 6] and much of the scientific community remains deeply connected to the surveillance and defense industries [12]. Yet, even today, the military uses of AI are often obscured, whether research is directly funded by defense agencies or developed for civilian purposes but with dual-use implications. As a result, many researchers and developers remain unaware of how their work might be deployed in conflicts [18], and the extent to which they might contribute to intentional harm, including potentially violations of international law [7].

Although AI in conflict and surveillance is a key topic of public and policy debate [20, 13, 15], there are currently no formal spaces for AI researchers themselves within the main machine learning conferences to articulate and discuss their positions towards the weaponization of their research. Given this gap, ICLR–with its key position in the research field–is an ideal host for a forum to consider both the harms associated with research dissemination and design decisions, and the opportunities for affirmatively building our research and development agenda from the starting position of non-violent harm prevention, research ethics, and respect for international law, including international humanitarian and human rights law.

In this workshop, we aim to address the critically under-discussed issue of AI's dual-use nature [4], focusing on how machine learning technologies are being adapted for military purposes [14, 16], potentially without the researchers' knowledge or consent. While attending to the heightened risks associated with particular areas and systems of research, we will also be collectively thinking through what it looks like to engage productively in research and development activities that considers ethics and international law at its core. Our objectives are to:

- Increase transparency about the pipelines through which AI research enters into military and surveillance applications [12].
- Develop collective strategies to address ethical and legal risks as a community of researchers [11].
- Highlight and support research efforts that contribute to peace-building applications [19, 8], including those helping to surface or elucidate harmful applications of AI [2, 9].

A key avenue of exploration will be to invite parallels between current conversations in AI and similar debates (with longer histories) in other scientific fields, such as genetic biology and nuclear physics; where researchers have grappled with similar ethical challenges and proposed concrete professional responses.

Call for contributions We invite submissions for presentations at our workshop under two tracks:

- **Tiny papers track** presenting novel work or work in progress relevant to the topics of interest listed below. This track welcomes concise contributions that may not fit the scope of full conference papers but still offer valuable insights to the community. Submissions may be up to 2 pages (excluding references and supplementary material). All accepted tiny papers will be shared via OpenReview and presented as posters during the workshop, with a select few invited for short oral presentations.

Preprint. Under review.

- **Previously published work** related to the workshop topics of interest listed below. These submissions will be evaluated solely for eligibility and relevance. Accepted works will be presented as posters, with exceptional submissions considered for spotlight presentations.

While AI assistance is permitted, AI-generated papers are not allowed. Papers should follow the Policies on Large Language Model Usage at ICLR 2026, including the disclosure of LLM usage.

The reviewing process will be single-blind and managed through OpenReview. During submission, authors will be required to disclose all sources of funding related to the research. We welcome submissions that fall under one of the following topics:

- Studies investigating how basic research transitions into military applications.
- Effects of AI militarization on marginalized communities and global inequalities.
- Research efforts that contribute directly to peace-building applications in a non-violent way.
- Collective strategies to address ethical risks as a community.
- Historical or comparative analyses of dual-use technology governance.
- Sociotechnical audits of existing military AI systems.
- Whistleblowing mechanisms and academic responsibility in high-risk research.

Presentation format We anticipate receiving approximately 20 submissions and plan to accept about 10-15 contributions. All accepted tiny papers will be presented at least as poster presentations and given the option to be featured on the workshop website. Given the sensitive nature of the topic, authors will also have the option to publish their work anonymously, if desired.

Managing conflicts of interest We will follow the standard conflict of interest (COI) protocol commonly used in academic workshops. Both authors and reviewers will be asked to declare any conflicting email domains, which will be used to guide reviewer assignments and ensure a fair evaluation process. Additionally, organizers will recuse themselves from any final discussions involving papers with which they have a conflict of interest, including submissions from someone within the same organization.

Timeline The timeline for paper submission, notification, and publication will be as follows:

Dec. 1, 2025	Workshop acceptance notification.
Dec. 15, 2025	Submission site opens on OpenReview.
Jan. 30, 2026	Workshop contributions submission deadline.
Feb. 2, 2026	Reviewers assignments.
Feb. 20, 2026	Reviews due.
Feb. 25, 2026	Emergency reviews due.
Mar. 1, 2026	Decision notification and OpenReview publication.
Mar. 30, 2026	Website update with the final program.
Apr. 26 or 27, 2026	Workshop date.

All dates are 11.59pm AOE.

2 Tentative schedule

We propose an in-person full-day event split into two morning and afternoon sessions. The workshop will feature:

- Three 35-minute keynote presentations (including QA) by three invited speakers.
- Four 20-minute invited talks (including QA) by four invited speakers.
- A 60-minute oral session for accepted paper presentations.
- A 20-minute interactive spotlight session with 2 to 5 minute short talks and a joint QA session.
- A 90-minute poster session.
- A 100-minute panel session with all the invited speakers.

This structure provides time for both formal knowledge-sharing through keynotes and oral paper presentations, as well as plenty of discussion time in the spotlight, poster, and panel sessions. While the final program including talk titles and the accepted contributions will be made publicly available about a month prior to the event, a tentative schedule is:

Morning

Start	Duration	Session
9:00	5 min	Welcome
		Session 1
9:05	35 min	Keynote talk #1
9:40	20 min	Invited talk #1
10:00	60 min	Oral presentations
11:00	15 min	Coffee break
		Session 2
11:15	35 min	Keynote talk #2
11:50	20 min	Invited talk #2
12:10	20 min	Interactive spotlight

Afternoon

Start	Duration	Session
12:30	55 min	Lunch break
		Session 3
13:25	35 min	Keynote talk #3
14:00	20 min	Invited talk #3
14:20	90 min	Poster presentations
15:30	20 min	Coffee break
		Session 4
15:50	20 min	Invited talk #4
16:10	100 min	Panel discussion
17:50	10 min	Closing remarks

3 Invited speakers

Our list of confirmed invited speakers include seven leading experts in machine learning, computer vision, human-computer interaction, cognitive science, anthropology, sociology, and political science. The list of speakers, listed alphabetically by last name is:

- **William Agnew [Confirmed]** (Carnegie Mellon University)

  Topic: Mapping the Computer Vision Surveillance and Weapons Pipeline

  William Agnew is a CBI postdoc fellow at CMU. William received his Ph.D. from University of Washington with Sidd Srinivasa, where he worked on AI ethics, critical AI, and robotics. William also helped found Queer in AI. William is interested in developing and sharing tools and ideas that go beyond participatory design and allow marginalized individuals and communities to own and meaningfully control their data and models derived from that data. Building on ideas from usable security/privacy, usage licenses, and indigenous data sovereignty, William wants to contribute to data and AI futures where individuals and communities know where their data is and can remove, add, or change their data in different datasets.

- **Abeba Birhane [Confirmed]** (Trinity College Dublin)

  Topic: From machine learning research to downstream societal impacts

  Abeba Birhane founded and leads the TCD AI Accountability Lab (AIAL). She received her PhD from UCD in 2022 and is currently a Research Fellow at the School of Computer Science and Statistics in Trinity College Dublin. Abeba's research focuses on AI accountability, with a particular focus on audits of AI models and training datasets – work for which she was featured in Wired UK and TIME on the TIME100 Most Influential People in AI list in 2023. She also served on the United Nations Secretary-General's AI Advisory Body and currently serves at the AI Advisory Council in Ireland.

- **Timnit Gebru [Confirmed]** (The DAIR Institute)

  Topic: Mitigating Harms and Fostering Community-Rooted Computer Vision Research

  Timnit Gebru is DAIR's founder and executive director. Prior to that she was fired by Google in December 2020 for raising issues of discrimination in the workplace, where she was serving as co-lead of the Ethical AI research team. Timnit also co-founded Black in AI, a nonprofit that works to increase the presence, inclusion, visibility and health of Black people in the field of AI, and is on the board of AddisCoder, a nonprofit dedicated to teaching algorithms and computer programming to Ethiopian and Jamaican highschool students. She has received a number of accolades including being named one of Nature's Ten people who helped shape science and one of TIME 100's most influential people. She is currently writing The View from Somewhere, a memoir + manifesto arguing for a technological future that serves our communities instead of one that is used for surveillance, warfare, and the centralization of power by Silicon Valley.

- **Sophia Goodfriend [Confirmed]** (University of Cambridge)

  Topic: Downstream Harms of Militarized AI

  Sophia Goodfriend is an anthropologist, a research fellow at the University of Cambridge's Pembroke College, and a non-resident fellow of the Harvard Kennedy School's Middle East Initiative. Sophia's academic research examines the impact of big data and machine learning on military conflict. Her writing on warfare and automation has appeared in The London Review of Books, Foreign Policy, The Baffler, +972 Magazine, The Boston Review, among other outlets. She holds a PhD in Cultural Anthropology from Duke University and received a Master's in Social Sciences from the University of Chicago and a BA in American Studies (summa cum laude) from Tufts University. Her research has been supported by the Harvard Kennedy School's Belfer Center, the Fulbright-Hays Program, and the National Science Foundation, among other institutions.

- **Gisela Luján Andrade [Confirmed]** (Perú por el Desarme)

  Topic: Autonomous Weapons in Latin America

  Gisela Luján Andrade is a social communicator and political scientist with over 15 years of expertise in political advocacy, humanitarian disarmament, human rights, and political communication. Her work spans analysis, advocacy, and research on autonomous weapons systems (AWS), and their humanitarian and social impact on civilians. She is the founder of 'Perú por el Desarme', a civil association promoting awareness of the humanitarian risks of AWS and fostering a culture of peace. Gisela has also contributed to international and regional processes related to landmines, cluster munitions, nuclear weapons, and AWS. She has published on topics related to emerging military technologies, disarmament activism, and AI governance. She is Peru's representative for the SEHLAC Network, a member of Stop Killer Robots, holds two master's degrees in political science (Pontifical Catholic University of Peru and Panthéon Sorbonne University), and has pursued doctoral studies in sociology at the École des Hautes Études en Sciences Sociales (EHESS).

- **Joseph Redmon [Confirmed]** (Allen Institute for AI)

  Topic: The role of researchers in military AI

  Joseph Redmon is a computer scientist best known for creating the real-time object detection system YOLO ("You Only Look Once") and the neural network framework Darknet. His work includes YOLOv1, YOLOv2 / YOLO9000, and YOLOv3, which balance speed and accuracy in detecting objects in images and video. Beyond technical contributions, Joseph has been vocal about the ethical implications of computer vision technologies, especially their use in military applications and surveillance, and in 2020 announced he would step back from computer vision research over concerns about potential negative societal impacts.

- **David Gray Widder [Confirmed]** (University of Texas at Austin)

  Topic: Military AI Research Funding

  David Widder studies how people creating Artificial Intelligence systems think about the downstream harms their systems make possible, and the wider cultural, political, and economic logics which shape these thoughts. Before his current position at UT Austin as an Assistant Professor, he was a postdoctoral Fellow at the Digital Life Initiative at Cornell Tech, and earned his PhD from the School of Computer Science at Carnegie Mellon University. He has previously conducted research at Intel Labs, Microsoft Research, and NASA's Jet Propulsion Laboratory. His recent research has been accepted to FAccT, Nature, CSCW, and Big Data & Society. His scholarly and activist work has appeared in Motherboard, Wired, MIT Technology Review, the Associated Press, and the New York Times.

4 Organizers and biographies

Our organizing team consists of six members with expertise spanning both computer science and the social sciences. In addition, we have secured a list of researchers who have already confirmed their willingness to serve on the program committee.

4.1 Organizing committee

Sonia Fereidooni (sf752@cam.ac.uk) is a current PhD student in Digital Humanities at the University of Cambridge funded by Gates Cambridge, studying the weaponization of representation through Generative AI and its militarization within the military-industrial-academic complex. She completed her BS in Computer Science & Data Science, BA in Sociology, and MS in Computer Science & Engineering at the University of Washington. Her research has spanned AI Bias, Commonsense Reasoning, improving unsupervised Computer Vision models, and desinging equitable Computer Science education. She previously worked as a Research Engineer at Google Brain and DeepMind, developing and open-sourcing the AI/ML development frameworks T5X and SeqIO. She was also a Responsible AI Research Fellow at Google's Impact Lab involved in analyzing systemic biases within Google's search algorithms, and a Research Scientist at the Allen Institute for AI (Ai2) working on the Mosaic Team under Yejin Choi researching commonsense reasoning in AI models.

Noa Garcia (noagarcia@ids.osaka-u.ac.jp) is an Associate Professor at The University of Osaka. Her research lies at the intersection of computer vision, natural language processing, and art, with a recent focus on the societal impact of machine learning, particularly algorithmic fairness. She is an active member of the academic community by regularly publishing at NeurIPS, CVPR, ICCV, etc. She has co-organized several workshops, including the Vision for Art (VISART @ ECCV 2022, ECCV 2024), the Instance Level Recognition (ILR @ ICCV 2021, ECCV 2022, ECCV 2024, ICCV 2025), and has served as the lead organizer for the Critical Evaluation of Generative Models and their Impact on Society workshop (CEGIS @ ECCV 2024, ICCV 2025).

Leonardo Impett (li222@cam.ac.uk) Leonardo Impett is an Assistant Professor in Digital Humanities at the University of Cambridge and a Research Group Leader at the Max Planck Society. He did his PhD at the Image and Visual Representation Laboratory, EPFL, where he worked on computer vision and the digital humanities. His research focuses on the application of machine learning to cultural and artistic data. He was previously Assistant Professor in the AI and Human Systems Durham University, and a research fellow at the Max Planck Society and Harvard University. His current work investigates bias in computer vision and the history of computer vision and AI research. Impett has also collaborated with cultural institutions such as the Liverpool Biennial, the Royal Opera House, and the Whitney Museum of American Art on projects relating to machine learning in artistic and curatorial practice.

Yannis Kalantidis (skamalas@gmail.com) Yannis Kalantidis is a Principal Scientist at NAVER LABS Europe. He received his PhD in Computer Science from the National Technical University of Athens in 2014. Prior to joining NAVER LABS Europe in 2020, he was a Research Scientist at Yahoo Research in San Francisco and Facebook AI in Menlo Park. His research focuses on visual representation and multi-modal learning, including self-supervised learning and adaptive systems. He is also passionate about connecting the computer vision community with socially impactful tasks, datasets, and applications. Yannis has co-led the organization of several workshops at top-tier AI venues, including Computer Vision for Global Challenges (CV4GC @ CVPR 2019), Computer Vision for Agriculture (CV4A @ ICLR 2020), and Wikipedia and Multi-Modal & Multi-Lingual Research (Wiki-M3L @ ICLR 2022).

Evangelos Kazakos (evangelos.kazakos@cvut.cz) is a postdoctoral researcher at the Czech Institute of Informatics, Robotics and Cybernetics at the Czech Technical University in Prague (CIIRC CTU). He received his PhD in Computer Science from the University of Bristol, where he worked on audio-visual egocentric action recognition. Prior to joining CIIRC CTU, he was a research scientist at Samsung AI Research in Cambridge. His research focuses on multi-modal learning, with an emphasis on image/video and language understanding, spatio-temporal reasoning, and large multi-modal models. More recently, he has developed an interest in applying these models to robotic manipulation.

Matt Mahmoudi (mm2134@cam.ac.uk is an Assistant Professor in Digital Humanities at the University of Cambridge, and an Advisor on Artificial Intelligence and Human Rights at Amnesty International. Matt's research focus is on red-lining and resistance in digital cities and the "smart" reproduction of racial capitalism. At Amnesty, he has led research and advocacy work on AI-driven surveillance from the NYPD's surveillance machine to Automated Apartheid in the occupied

Palestinian territory. He is a Research Associate with the Centre of Governance and Human Rights, and an Affiliate Researcher with the DALOSS project at the University of Copenhagen. Matt is the author of Migrants in the Digital Periphery: New Urban Frontiers of Controls (University of California Press, 2025), and is a co-editor on Resisting Borders & Technologies of Violence (Haymarket, 2024) together with Mizue Aizeki and Coline Schupfer.

4.2 Program Committee

Table 1 lists researchers who have already confirmed their intention to participate in the program committee for our workshop.

Table 1: Computer scientists and their institutional affiliations.

Name	**Institution**
Katerina Adam	NTUA, Greece / CIIRC CTU, Czech Republic
Pierre Alquier	ESSEC Business School - ASIA PACIFIC, Singapore
Adrián Arnaiz-Rodríguez	ELLIS Alicante / University of Alicante, Spain
Tianwei Chen	ZOZO Research, Japan
Amanda Duarte	Barcelona Supercomputing Center / UPC, Spain
Andreu Girbau-Xalabarder	Denso IT Laboratory, Tokyo, Japan
Xavier Giró-i-Nieto	Amazon, Barcelona, Spain
Yusuke Hirota	NVIDIA Research, Taiwan
Amirpasha Mozaffari	Barcelona Supercomputing Center, Spain
Michael A Osborne	University of Oxford, United Kingdom
Bill Psomas	Czech Technical University in Prague, Czech Republic
Patrick Ramos	The University of Osaka, Japan
Ryan Ramos	The University of Osaka, Japan
Piera Riccio	ELLIS Alicante / University of Alicante, Spain
Enrique Sánchez-Lozano	Samsung AI Center Cambridge, United Kingdom
Joseph Schafer	University of Washington, USA
Vladan Stojnić	Czech Technical University in Prague, Czech Republic
Giorgos Tolias	Czech Technical University in Prague, Czech Republic
Nanne van Noord	University of Amsterdam, Netherlands
Lu Wei	The University of Osaka, Japan
Yankun Wu	The University of Osaka, Japan
Tong Xiang	The University of Osaka, Japan
Cheng Xu	University College Dublin, Ireland

5 Anticipated audience size

Based on our experience in hosting workshops, the large audience at past ICLR conferences, and the interest of the community in critical AI issues, we anticipate 100-150 participants in the room at all times and about 300-400 audiences in total throughout the event.

6 Plan to get an audience for a workshop (advertising, reaching out, etc.)

We will promote the workshop through both online and in-person channels. This includes outreach on social media and at major conferences such as NeurIPS 2025, EurIPS 2025, and AAAI 2026. Additionally, the Call for Contributions will be shared through established mailing lists and affinity groups, including LatinX in AI, Women in Machine Learning, Women in Computer Vision, and others. Several of our invited speakers have strong professional networks, and we will collaborate with them to further amplify the workshop's visibility.

7 Diversity commitment

We intentionally worked to ensure diversity of the organizers and invited speakers, with respect to geography, gender, race, affiliations, expertise, and seniority.

Organizers The organizing team spans five countries (Czech Republic, Italy, Japan, Spain, United Kingdom) and is composed of two women and four men. Three members of the team hold faculty positions in academia, one works in industry, one is a postdoctoral researcher, and another is a PhD student. The expertise of the organizers is interdisciplinary and covers machine learning, digital humanities, and social sciences.

Invited speakers Four out of seven invited speakers identify as female, an intentional effort to help address gender imbalance in the machine learning community. The speakers represent a range of research domains and institutions, offering a rich variety of perspectives.

The workshop is designed to reflect a broad range of geopolitical perspectives on the deployment of AI for military purposes. The confirmed speakers are based in four countries (Ireland, Peru, United Kingdom, and United States) and bring lived and professional experiences from across Africa, Europe, North America, and South America. This diversity is intentional, ensuring that discussions are informed by varied national contexts, including those from both the Global North and South. The workshop aims to foster inclusive dialogue that accounts for the global nature of AI's impacts, particularly on marginalized communities and across different security and policy environments.

Promoting an inclusive environment We are committed to creating a welcoming and inclusive environment for all attendees. Participants with disabilities will be encouraged to contact us in advance with specific accessibility requirements so that we can provide appropriate support. In addition, we are exploring hybrid or remote attendance options, particularly for researchers with disabilities and those from institutions with limited resources.

Remote attendance For those unable to attend in person or remotely due to time zone differences or other constraints, workshop materials will be made available on the website. In exceptional cases (such as visa denial, illness, or pregnancy), speakers who are unable to attend in person will be permitted to present remotely. Similarly, if authors of accepted contributions are unable to attend for these reasons, we will arrange for their posters to be displayed on their behalf.

8 Virtual access to workshop materials and outcome

To support broad accessibility, we will request permission from authors and speakers to share their materials online following the workshop. Additionally, a summary of key discussions and outcomes will be published on the workshop website to promote continued engagement and knowledge exchange within the wider research community.

9 Previous related workshops

This workshop is a new initiative that aims to bring this critical discussion to a major AI conference for the first time. While this is new workshop, it builds on the momentum of recent successful workshops related to governance and safety, including "Harms and Risks of AI in the Military" [1] (Mila 2024), ML Safety [10] (NeurIPS 2022), Multi-Agent Security [17] (NeurIPS 2023), and Private ML [3] (ICLR 2024).

References

[1] Workshop on Harms and Risks of AI in the Military at MILA, 2024. https://www.harms-risks-ai-military.org/.

[2] Amnesty International. Israel and occupied palestinian territories: Automated apartheid: How facial recognition fragments, segregates and controls palestinians in the OPT, 2023. https://www.amnesty.org/en/documents/mde15/6701/2023/en/.

[3] S. Avestimehr, T. Li, N. F. Mireshghallah, S. Oh, F. Tramer, and Z. Xu. Workshop on Privacy regulation and protection in machine learning at ICLR, 2024. https://pml-workshop.github.io/iclr24/.

[4] A. Brenneis. Assessing dual use risks in ai research: necessity, challenges and mitigation strategies. *Research Ethics*, 21(2):302–330, 2025.

[5] K. Crawford. *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press, 2021.

[6] J. E. Dobson. *The Birth of Computer Vision*. U of Minnesota Press, 2023.

[7] S. Fereidooni and V. Heidt. The fallacy of precision: Deconstructing the narrative supporting ai-enhanced military weaponry. In *Harms and Risks of AI in the Military*, 2024.

[8] J. Filipi, V. Stojnić, M. Muštra, R. N. Gillanders, V. Jovanović, S. Gajić, G. A. Turnbull, Z. Babić, N. Kezić, and V. Risojević. Honeybee-based biohybrid system for landmine detection. *Science of the total environment*, 803:150041, 2022.

[9] S. Goodfriend. Algorithmic state violence: Automated surveillance and palestinian dispossession in hebron's old city. *International Journal of Middle East Studies*, 55(3):461–478, 2023.

[10] D. Hendrycks, V. Krakovna, D. Song, J. Steinhardt, and N. Carlini. Workshop on Machine Learning Safety at NeurIPS, 2022. https://neurips.cc/virtual/2022/workshop/4.

[11] L.-A. Kaffee, A. Arora, Z. Talat, and I. Augenstein. Thorny roses: Investigating the dual use dilemma in natural language processing. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13977–13998, 2023.

[12] P. R. Kalluri, W. Agnew, M. Cheng, K. Owens, L. Soldaini, and A. Birhane. The surveillance AI pipeline. *arXiv preprint arXiv:2309.15084*, 2023.

[13] H. Khlaaf, S. M. West, and M. Whittaker. Mind the gap: Foundation models and the covert proliferation of military intelligence, surveillance, and targeting. *arXiv preprint arXiv:2410.14831*, 2024.

[14] A. Loewenstein. *The Palestine laboratory: How Israel exports the technology of occupation around the world*. Verso Books, 2024.

[15] S. Romansky. Lessons from the EU on confidence-building measures around artificial intelligence in the military domain. *SIPRI Publications*, 2025. https://www.sipri.org/sites/default/files/2025-05/eunpdc_no_97.pdf.

[16] P. Scharre. *Four battlegrounds: Power in the age of artificial intelligence*. WW Norton & Company, 2023.

[17] C. Schroeder de Witt, H. Milani, K. Krawiecka, S. Mehta, C. Cremer, and M. Strohmeier. Workshop on Multi-agent security: Security as key to ai safety at NeurIPS, 2023. https://neurips.cc/virtual/2023/workshop/66520.

[18] S. Schwartz, L. G. Guntrum, and C. Reuter. Vision or threat—awareness for dual-use in the development of autonomous driving. *IEEE Transactions on Technology and Society*, 3(3):163–174, 2022.

[19] R. Sefala, T. Gebru, L. Mfupe, N. Moorosi, and R. Klein. Constructing a visual dataset to study the effects of spatial apartheid in south africa. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.

[20] The Future of Life Institute. Autonomous weapons open letter: AI & robotics researchers, 2016. https://futureoflife.org/open-letter/open-letter-autonomous-weapons-ai-robotics/.