

WORKSHOP: POLITICS, ETHICS, AND EXPERTISE IN THE DEBATE OVER EMERGING WEAPONS TECHNOLOGIES



University of Canterbury, Christchurch, New Zealand

August 30 – September 1, 2023

Wednesday, 30 th August – Undercroft 101	
5:00 – 6:30pm	Jack Poulson Keynote Address
Thursday, 31 st August – Elsie Locke 611	
9:30 – 10:00am	Welcome and Introductions
10:00 – 11:00am	Jack Poulson
11:00 – 11:10am	Break
11:10am – 12:10pm	Online session I: Lucy Suchman
12:10 – 1:30pm	Lunch
1:30 – 3:00pm	Jeremy Moses , Geoff Ford , and Sian Troath
3:00 – 3:10pm	Break
3:10 – 4:40pm	Michael Richardson and Daniel Chia Matallana
4:40 – 4:50pm	Afternoon tea
4:50 – 6:20pm	Online session II: Ingvild Bode , Anna Nadibaidze , and Natasha Karner
7pm –	Workshop Dinner – venue TBC
Friday, 1 st September – Elsie Locke 611	
9:00 – 10:30am	Online session III: Berenike Prem and Fynn Horstmanshoff
10:30 – 10:45am	Break
10:45am – 12:15pm	Bianca Baggiarini and Sean Rupka
12:15 – 12:45pm	Wrap-up discussion
12:45 – 2:30pm	Lunch

Curating international public procurement for surveillance and weaponry

Jack Poulson, Tech Inquiry

Tech Inquiry maps international ecosystems of surveillance and weapons, and the influence industries surrounding them. It draws upon a curated database of government purchasing across more than thirty countries, alongside lobbying records, tax filings, reporting, and manually tagged associations to identify, explore, and analyse relationships between corporations, governments, and nonprofits. In this address, Tech Inquiry's founder and developer, Dr Jack Poulson, will discuss how Tech Inquiry was built, the public data sources and technologies it uses, and what this reveals about connections between surveillance and weapons industries and government agencies in the US. His talk will explore case studies of cellphone location-tracking, facial recognition, and social media surveillance and how these relate to the US military, law enforcement, and vigilantism.

Il/legal war: Expanding the frame of meaningful human control

Lucy Suchman, Professor Emerita of Anthropology of Science and Technology, Department of Sociology at Lancaster University

This talk considers the enduring tension between militarism and democracy. To develop the argument I work from a central trope in current debates over the regulation of autonomous weapon systems (AWS), the concept of 'meaningful human control.' I begin with my own attempt, as a participant in discussions at the United Nations Convention on Certain Conventional Weapons (CCW), to set out an argument for the impossibility of developing fully automated weapon systems that could be adherent to International Humanitarian Law (IHL). Reflecting on the force but also the limits of that argument brings me to the question of how efforts to regulate AWS (pragmatically and perhaps necessarily) frame the bounds of the discussion, and what is placed outside of the frame. I then turn to a key moment in recent history, the US invasion of Iraq in 2003, to consider how the frame of meaningful human control might be expanded beyond the operation of weapon systems in war, to address the question of war's il/legality in particular instances. I conclude with some thoughts about the broader frame of militarism, and efforts to shift the balance away from interests vested in the perpetuation of militarism toward greater investment in alternatives to war.

Money and Expertise in the US Military-Industrial Complex: The Case of CNAS and LAWS

Jeremy Moses, Associate Professor of International Relations, University of Canterbury

This paper examines the position and influence of think tanks in the debate over autonomous weapons systems (AWS). In particular, it examines the influential role of the Center for a New American Security (CNAS) in shaping the debate over new weapons technologies. The paper draws on Thomas Medvetz's understanding of the role of think tanks as 'members of an interstitial field' that lies amidst academic, political, business, and media domains. This can be clearly recognised in the operation of CNAS and in the careers and commentary on AWS by key figures at CNAS, including Robert Work and Paul Scharre. Drawing on a knowledge graph built from CNAS' online biographies of experts, board members and affiliates, we

examine the networks of state, corporate, and military actors connected to CNAS that also have financial and political interests in the AWS debate. Existing reports on the relationship between think tanks and the interests of their major donors are used to show how tightly embedded the beneficiaries of US Government military contracts are with the research bodies that are most influential in promoting the need for investment in the military technology sector. The paper concludes with the argument that these relationships between the state, the weapons and tech industries and think tanks need to be more openly and transparently acknowledged and taken into account when research and advocacy on AWS is being conducted.

Tech Boosting through Expert Criticism: The case of the Future of Life Institute and AI weapons

Geoff Ford, Lecturer, University of Canterbury

The Future of Life Institute (FLI) has been influential in the public understanding of the dangers of autonomous weapons systems, particularly through the publication of open letters and the production of the 'Slaughterbots' short films. In this presentation of work in progress, we look at the activism of FLI in order to understand how the critique of autonomous weapons they advance sits in relation to their broader boosting of AI technologies. We argue that the FLI's warnings about the dangers of autonomous weapons effectively functions to promote the continuation of military AI research, only ruling out the pursuit of fully autonomous weapons systems that would be unlikely to be of use to militaries. In deploying well-known technologists and entrepreneurs to validate their messaging and pointing toward the threat of rogue actors and robots run amok, FLI's contributions to the autonomous weapons debate shows how expertise can be deployed in the service of interests that are not always self-evident.

Expertise in trust and trust in expertise

Sian Troath, Postdoctoral Fellow in International Relations, University of Canterbury

Who gets to decide what constitutes a trustworthy autonomous system? What makes a trust expert trustworthy? Trust is a central part of the discourse about the development and adoption of robotics, autonomous systems, and artificial intelligence (RAS-AI) for defence purposes. The people who gets to decide what constitutes trust, then, are influential figures in shaping the future of war. They hold influence over whether RAS-AI align with legal requirements, and how they will be regulated. Decisions on what constitutes trust feeds into design decisions, and choices of particular technologies to pursue. The language of trust is being used by academia and industry to market their expertise and their products to defence. This paper will explore who is being considered a trust expert in RAS-AI in Australia, the UK, and the US, and the politics of trust, ethics, and expertise.

'Expertise, Military Virtue, and the Limits of 'Ethics' in AI Research'

Michael Richardson, Associate Professor of Media, UNSW Sydney

Militaries, including the Australian Defence Force, tend to see ethics as instrumental and principally related to conduct on and off the battlefield, rather than enmeshed with larger questions of justice or societal obligation. When this approach encounters AI, the limits of 'ethics' as a framework for reducing the harms produced by AI become clear. Militaries are well behind the private sector in AI and big data development and expertise, a reality which is accelerating collaborative, industry-led processes that mimic aspects of the Silicon Valley model of agile development. This move fast and break things approach could have serious repercussions given the life and death situations in which military technology is often applied. But militaries are not only valuable clients for big tech, but also increasingly important sources of funding for academic research. In the context of military AI, 'ethics' possesses an economic function that helps frame engagements with industry and academic expertise. Ethics are typically posed as both values to hold and problems to solve. To show how military ethics functions within economies of virtue (Phan, et al 2022), this paper argues that three critical dynamics around 'ethics' are shaping the emerging military technology industry in Australia and its relation to academia. First, ethics are often seen as codes themselves – and as codable. Second, ethics serves an important discursive purpose within the rapidly growing defence industries in Australia. And third, ethics facilitates engagement with academic expertise, with research funding as a central factor, materialised through centres, networks, symposiums, and workshops, not to mention the buzzcuts, insignia and corporate polos that populate them. Mapping and analysing this confluence of AI research, industry and application is a critical task because it operates according to a different logic and economy of funding than is the norm within tech support for academic research on AI and big data.

Jurisprudential Imaginaries in Science Fiction: AI and Legal Reasoning

Daniel Chia Matallana, PhD Candidate, Victoria University of Wellington

Legal reasoning lies at the legal profession's core, encompassing judicial, legislative, and lawyerly practices. It is often characterised as "a matter of induction, deduction and analogy", taking on authoritative, inquisitorial, dialectical, or hermeneutical forms (Samuel, 2009, p. 182). Despite the sort of legal reasoning scheme used, for example, to apply a norm, read a statute or provide advice in a military situation to comply with international humanitarian law (IHL), the fact is that legal reasoning is a legal paradigm exclusive to humans. In Condello's (2020, p. 1185) words, "interpretation in law cannot be replaced by an algorithm". If such claims are valid, integrating artificial intelligence (AI) in legal decision-making requires a more collaborative approach between decision-makers and AI agents. This study investigates the role of AI technologies as decision-support mechanisms and their interaction with the principle of proportionality in the context of IHL. It will explore the contention put forth by some scholars that the law of armed conflicts "does not require human decision for selecting and engaging targets to be legal" (Jensen, 2020, p. 28). Such examination raises essential questions about the potential of AI in making warfare more humane versus its ability to legitimise violence. This research, part of my PhD thesis, seeks in science fiction literary works marks that jurisprudentially anticipate and speculate about contemporary and future legal paradigms regarding AI in today's and forthcoming society. In this specific case, it will analyse the impact

of AI on legal reasoning. I will conduct an analysis of the jurisprudential imaginaries of "Plague Birds" a novel by Jason Sanford, which entails the story of Crista and her integrated AI, both forming a singular plague bird tasked with administering justice in a world striving to regain its lost humanity. Thus, it is also a story about the interdependence of humans and machines.

Emerging technologies: what counts as expertise in United Nations arms control forums?

Ingvilde Bode, Associate Professor of International Relations, University of Southern Denmark

Various United Nations disarmament fora have long drawn on different forms of expertise. This has taken several forms, for example appointing individual experts on specific mandates or organising groups of governmental experts to deliver recommendations on moving forward. This is therefore not a new phenomenon, but the potential integration of so-called emerging technologies in weapon systems often requires drawing on a much wider range of experts. But on what basis is the United Nations (UN) actually constructing the input provided by external experts as authoritative? The paper unpacks this question with respect to the example of and expertise in the debate around the potential regulation of military applications of artificial intelligence. At the UN, this chiefly held under the auspices of the UN Convention on Certain Conventional Weapons in Geneva. I argue that answering who is constructed as an expert and why matters for questions of trust in international organisations, such as the UN, because such constructions end up being decisive for what is suggested as "appropriate" governance. In particular, I examine different basis for constructing expertise including objectivity, positivist versions of scientific rigour, and a hierarchy of the sciences. I conclude that based on these dynamics certain types of expertise, often so-called "technical" types, count for more than others with significant consequences for a critically informed public debate on emerging technologies in arms control.

Expertise in Understanding Responsible Military AI: The Role of the REAIM Conference

Anna Nadibaidze, PhD Fellow, University of Southern Denmark

The first edition of the Responsible AI in the Military Domain (REAIM) conference was held in 2023 in The Hague, organised by the Dutch government in collaboration with the Republic of Korea. It was hailed as a forum that brought together different actors (industry, governments, think tanks, academia, civil society) in discussing military AI beyond lethal autonomous weapon systems (LAWS), which remain the core focus of the Group of Governmental Experts (GGE) at the United Nations. Although the REAIM conference does not have formal ties with the GGE process, it has become an important part of the regulatory debate on autonomy in weapon systems. At the 2023 GGE meetings, several delegations described REAIM as a key step towards forging a shared understanding of what constitutes 'responsible' applications of AI in the military sphere, and especially in weapon systems. Having attended the GGE meeting in March, I have also noticed in private conversations how the REAIM conference left an impression on state delegations and other actors. In this work in progress, I explore whether and/or how this conference has contributed to promoting certain understandings, narratives,

and practices of ‘responsible AI’ in the military, as well as how the conference involved different forms of expertise in this process of promoting certain understandings and practices versus others. I argue that these practices matter for the debate about international regulation or national policies that states adopt in the area of military AI. The guiding questions for this project include: who was invited to the conference, who was present/absent (particularly in terms of Global North/South divisions)? Who was considered an expert (especially in plenary sessions)? Which narratives of responsible military AI were promoted by which actors, and with the help of which expertise? To answer these questions, I will rely on my own observations as a participant at REAIM, but also interview other participants from different backgrounds. Moreover, I will look at the mentions about REAIM in the GGE debate by accessing the recordings of the meeting and using my personal notes from the March 2023 session. Based on the feedback received during the workshop, I could engage in more focused observations during the second edition of REAIM, planned to be held in the Republic of Korea in 2024.

Technoscientific Rationality and AWS

Natasha Karner, PhD Candidate, RMIT University

This presentation will look at how a scientific rationality underpins Autonomous Weapons Systems (AWS), including assumptions that being “mechanised,” “remote,” and “precise” are ideal in the conduct of war. It provides an epistemological discussion on the role of technology in society, before expanding on to the influence of technoscientific rationality on war. For its methodology, it utilises a discourse analysis of key philosophical texts, historical records, speeches, and policy documents. Overall, this presentation shines light on the technoscientific rationale that underpins a push for modern, “dehumanised” war so that we may better understand the development and rationalisation of AWS.

Contested Expertise: Norm Antipreneurship and the Regulation of Autonomous Weapons Systems

Berenike Prem, Postdoctoral Researcher, University of Bremen

For almost ten years, the issue of autonomous weapons systems (AWS) is on the agenda of the United Nations (UN) where state delegates discuss the implications of increasing autonomy in weapons systems and potential aspects of a “normative and operational framework” governing AWS within the UN Group of Governmental Experts on Lethal Autonomous Weapons Systems (GGE on LAWS). Since its inception, proponents of a legally binding instrument have criticized the GGE process for its slow progress and lack of ambition. In the latest round of talks in 2023, states could agree on little more than the tautological statement AWS must not be used if they are incapable of being used in compliance with international humanitarian law (and are therefore already de facto legally prohibited). Scholars have attributed the current stalemate of the GGE to issue-specific attributes and misguided strategic choices of the Campaign to Stop Killer Robots (Bode and Huelss 2022, Maas 2019, Rosert and Sauer 2019). Unlike previous arms control efforts, there is no on statistical evidence of widespread harm, photographs of victims, or testimonies of survivors to document the problem. Strategies that have traditionally worked, such as stigmatizing AWS based on their indiscriminateness, have

therefore proven elusive (Rosert and Sauer 2020). While these explanations hold merit, they overlook the degree of consensus that states have already reached in the GGE due to the persistent knowledge work of experts and campaigners which have gone a long way to educate state delegates about the nature and consequences of increasing autonomy in weapons systems. Today, about 90 countries have joined the call for the negotiation of a legally binding instrument on AWS. The question therefore arises: why has there been so little substantial progress in norm evolution despite the growing convergence on the humanitarian, legal, ethical, strategic, or operational grounds against (certain) AWS and the need for their regulation? To explain this outcome, the paper invites us to pay closer attention to the contested dynamics of knowledge production in the GGE and the role of “norm antipreneurial” states. What currently hampers progress towards a normative framework is not merely the absence of knowledge or the lack of a shared understanding of the problem at hand, but also the active resistance of norm-antipreneurial states. They have been able to fend off claims for ambitious arms control norms by strategically mobilizing and at times actively creating uncertainty. To substantiate this claim, the paper will analyze the role of the US, Russia, and India in the GGE as some of the most active opponents of legally binding rules on AWS.

Expert Influence on Germany's Position on Autonomous Weapons Systems

Fynn Horstmannshoff, Graduate Student, University of Bremen

The aim of this Master Thesis is to expand the literature on expert influence by looking at the case of Germany's position on Autonomous Weapons Systems from 2013-2023. Since 2013, so called “(Lethal) Autonomous Weapon Systems ((L)AWS)” are being discussed under the forum of the United Nations Convention on Certain Conventional Weapons (UNCCW), after intensive lobbying by civil NGOs to get the topic on the UN's agenda. During this period different national positions have formed on the matter of these by critics called “killer robots” or more technically called “(L)AWS”. At the beginning of the debate about AWS, states have been confronted with uncertainty about the topic, for example, what constitutes AWS, how far developed they are, when states can expect a use of them if nothing is done, expected impacts on International Humanitarian Law (IHL) and other issues. This has created a demand for expertise as well as the need for a more depoliticized forum than the CCW in which states can exchange their views on the matter, without the necessity to frame it as a discussion for an addition to the UNCCW protocols. Therefore, the UNCCW established a Group of Government Experts (GGE) to provide this more informal forum. The demand for expertise has been filled by experts that were invited by the UNCCW itself but also on the national levels by domestic experts that were consulted by the government institutions for expertise on the matter. Over time, more fleshed out positions of states on the matter of AWS have formed in the GGE and the UNCCW on LAWS. The goal of the study is to look in particular at the German position and how it evolved and was developed by the government institutions, primarily the Federal Foreign Office of Germany and connected ministries like the Federal Ministry of Defense. The work will trace the process of position forming and how the position changed over the nine years of discussion on the matter and will investigate which role experts played in this process of positioning of Germany in the GGE and the UNCCW. The focus on national expertise on the matter is not unjustified: Emanuel Adler in 1992 already said that there is a need to pay more attention to the international influence of domestic epistemic communities in various fields, including arms control. Recent literature also hypothesizes that

a possible explanation for developments in national positions in the EU (especially in Germany) on LAWS might be traceable back to the engagement with expert groups such as the International Panel on the Regulation of Autonomous Weapons (iPRAW), a funded expert committee by the German Federal Foreign Office, and other civil society organizations. The aim of the study is to analyze documents, mainly statements and working papers that have been issued by Germany to the UNCCW or the GGE, publicly available documents, statements and speeches of the Federal Foreign Minister and its Ministry, documents that are issued about LAWS by the Federal Ministry of Defense in Germany as well as publicly available records and documents of events that were hosted to bring together experts with politicians or administrative decision makers to discuss the matter of LAWS. The study attempts to trace the development of the national position of Germany and will compare it with documents by experts that have provided analysis, discussions and policy advice to the government. For that key think tanks and academic institutes were identified that play a key role in the German debate about LAWS. Finally, the document analysis will be supplemented by a handful of interviews with key experts from Germany and personal from the Federal Foreign Office or the Defence Ministry, which will provide further insights into the processes and the relationship between experts and politics.

The Failure of “Ethical military AI”

Bianca Baggiarini, Lecturer in Military Studies, Australian National University

This paper will reflect upon my time working in Australia as an “ethicist” on projects associated with military AI. Although I am not an ethicist, and have never thought of myself as such, during my post-doc at the Australian Defence Force Academy people started referring to me as an ethicist (it was ethicists that had something to say – not just your ordinary social science type)! Beginning in 2019, we can point to a boom in AI ethics as applied to emerging, AI-enabled military technologies, especially in Australia: more broadly, many militaries, states, NGOs, non-profits, companies, and supranational entities began producing ethics principles. In response, and with this (contrived) hat of ethicist on, I was able to gain access to important conversations with Defence, and thus was offered a seat at the table (sort of). But why? The question of who is included as important, with something meaningful to say, or who is considered to be a stakeholder with legitimate interests, and thus who is excluded in the process, is an important, yet often overlooked one in the debates on LAWS. However, in Defence and Defence-adjacent spaces, this is deliberately overlooked. In shaping the debate about the future of autonomous weapons, for example, ethicists are included to the extent that they remain depoliticized even whilst they ironically call for AI ethics to be “fair,” “just,” and “equitable.” The more “codifiable” their ideas could be, the better. And thus ethicists emerged to smooth over the cracks and act as middlemen between the technologists (engineers, etc.) and Defence. Put differently: you cannot be an ethicist in this space and also be anti-war, or anti-AI. The debate is not about a democratic multiplicity of voices that may contradict “the plan”, but one driven by authoritarian consensus and tunnel vision where the “solutions” to “ethical dilemmas” are already known in advance by those with power. In 2023, Australia announced two major security/strategic initiatives (the DSR and AUKUS), neither of which have cited ethics as relevant, despite tremendous financial investment in this space and years of proliferating ethics discourse. What happened to ethics, and what explains the discursive boom and bust? My paper will reflect on the relationship between ethics and politics as it relates to

expertise around military AI, analysing why military AI ethics has so far ignored politics. AI ethics for militaries has failed, at least in Australia, for now – but in many ways it was always destined to.

Dereliction by Design: The Diffusion of Responsibility through the Technologization of Military Labour

Sean Rupka, PhD Candidate, UNSW Canberra

This paper will reflect ongoing work regarding the political effects of the transformation of military labour through processes of technologization and automation. I argue that certain ‘problems’ or challenges emerging from the shift in the nature of contemporary military work and operations, be they ethical, political or more specifically around the effects on the soldiery, are in fact not symptomatic but indeed intrinsic to logic of this technological shift itself. Such a logic is buttressed by very particular notions of expertise. The now ubiquitous implementation of autonomous systems, not only in Australia but globally, seems present as given the necessity of such systems for the future of modern warfare. Further, while military applications of such technologies remain as the focal point of most discussion, autonomous systems are no longer merely the sole purview of advanced militaries but rather the logic that drives their proliferation in security arenas is expanding to commercial and industrial areas of the public sector as well. In the military context, while much attention is directed towards questions surrounding *jus in bello*, or how to most efficiently employ human-machine teams, the ethical questions around the application of autonomous systems often miss foundational questions such as ‘what is the problem to which autonomous systems are the answer in the first place’? What political solutions are generated by a system that relies first and foremost on such tools? From a political perspective, increasing technologization appears unavoidable, a singular inescapable *telos* coming from a momentum generated by industries of expertise and discussion that take such technologies as inevitable. This is particularly apparent through ethical discussions that have proliferated in lock step (though always a step behind) with military development. Questions around Kantian categorical imperatives versus utilitarian trolley problems abound, but the drive for expanding and broadening technological implementation relies on a governing reason or rationality that insists on its own apoliticism. To an extent, even ethical discussions in this area seem to have almost run their course, for while substantial investment has occurred to generate academic discussion in partnership with national security/defense forces, the so far largely circular debates are seemingly being deemed less important going forward, if reports such as the recent Australian Defence Strategic Review in 2023 can be seen as indicators or gauges of current political will and interest. This is not to say that autonomous systems are a Pandora’s box that can be sealed back up, nor to say that important conversations cannot or should not occur around their proper implementation. My concern however, is around how such conversations may be used to obscure or elide discussions of a different nature. I suggest focussing on the radical potential effects of such technologies not as ‘undesirable but perhaps unavoidable’ collateral damage, but instead examining them as outcomes that are in fact generated by design. As such, my paper hopes to probe at a more basic level the historical logic that led in this present position and to interrogate how the supposed symptoms of such progress are in fact inherent to the dominant logic driving the system.