



Algorithmic predictions and pre-emptive violence: artificial intelligence and the future of unmanned aerial systems

Anthony Downey¹

Accepted: 6 November 2023 / Published online: 5 December 2023
© The Author(s) 2023

Abstract

The military rationale of a pre-emptive strike is predicated upon the calculation and anticipation of threat. The underlying principle of anticipation, or prediction, is foundational to the operative logic of AI. The deployment of predictive, algorithmically driven systems in unmanned aerial systems (UAS) would therefore appear to be all but inevitable. However, the fatal interlocking of martial paradigms of pre-emption and models of predictive analysis needs to be questioned insofar as the irreparable decisiveness of a pre-emptive military strike is often at odds with the probabilistic predictions of AI. The pursuit of a human right to protect communities from aerial threats needs to therefore consider the degree to which algorithmic auguries—often erroneous but nevertheless evident in the prophetic mechanisms that power autonomous aerial apparatuses—essentially authorise and further galvanise the long-standing martial strategy of pre-emption. In the context of unmanned aerial systems, this essay will outline how AI actualises and summons forth “threats” through (i) the propositional logic of algorithms (their inclination to yield actionable directives); (ii) the systematic training of neural networks (through habitually biased methods of data-labelling); and (iii) a systemic reliance on models of statistical analysis in the structural design of machine learning (which can and do produce so-called “hallucinations”). Through defining the deterministic intentionality, systematic biases and systemic dysfunction of algorithms, I will identify how individuals and communities—configured upon and erroneously flagged through the machinations of so-called “black box” instruments—are invariably exposed to the uncertainty (or brute certainty) of imminent death based on algorithmic projections of “threat”.

Keywords Artificial intelligence · Image processing · Drone warfare · Prediction · Pre-emption · Autonomous weapons systems

“The tragic reality of war became normalised not because we wanted it to, but because we had to. It was the only way we could keep living. It was the only way we could keep living, even though there was not much life, when the sky is not safe, no other place can be safe, even if you can escape or perhaps one day the crisis would end the trauma can follow you for the rest of your life.”

—Abrar Mechmechia, giving evidence to the Airspace Tribunal Toronto Hearing, The Power Plant Contemporary Art Gallery, 7 November 2020.

“I believe that modern developments in technology and telecommunication, instead of diminishing the realm of ghosts [ensures that they are] part of the future”

—Jacques Derrida, *Ghostdance* (Dir. Ken McMullen, 1983).

(1) At Baghdad International Airport on January 3, 2020, Major General Qasem Soleimani of the Iranian Islamic Revolutionary Guard Corps (IRGC) was killed by a missile launched from a MQ-9 Reaper drone.¹ Also known as a Predator B, the MQ-9 is referred to by the U. S. Air Force (USAF) as a Remotely Piloted Vehicle/Aircraft (RPV/RPA). It is but one component in a more expansive unmanned aerial system (UAS) that includes elements of human input—the so-called “human-in-the-loop” factor—in its standard operating procedures (SOP). The plan to assassinate Soleimani is believed to have been directed by the CIA from Creech Air Force Base in Nevada and was a “decapitation strike”; that is, a strike where the target was known beforehand (Read 2020). Although a “decapitation strike” differs to a

✉ Anthony Downey
anthony.downey@bcu.ac.uk

¹ Birmingham City University, Birmingham, United Kingdom



“signature strike” (insofar as the latter involves an unknown and as-yet-to-be-fully identified target), both operations serve a singular purpose: the eradication of an impending or *yet-to-be* fully substantiated instance of threat.²

In the aftermath of the fatal missile attack on Soleimani, the then president of the United States (US), Donald J. Trump, alluded to a confrontation that had already taken place. “They *attacked* us and we hit back”, he tweeted on January 5, 2020, without offering any contiguous evidence of the incident he was referring to nor, subsequently, any proof that a specific attack had indeed occurred in advance of Soleimani’s assassination.³ Prior to Trump’s tweet, the Department of Defense (DoD) announced that the missile strike was warranted due to the fact that Soleimani “was *actively developing* plans to attack American diplomats and service members in Iraq and throughout the region” (Statement by the Department of Defense 2020. Emphasis added). The DoD document, no doubt with one eye on international law, further volunteered the vague assertion that the missile strike on Baghdad Airport was “aimed at *detering future* Iranian attack plans” (Statement by the Department of Defense 2020. Emphasis added). The obfuscation at work here, with varying accounts of an attack that had happened, an attack that was about to happen (imminent), and an attack that might happen sometime in the future, should not distract us from a basic tenet of contemporary warfare: pre-emption.

Although it has been a historical mainstay of military tactics, the use of pre-emptive, or anticipatory, self-defense—the so-called “Bush doctrine”—is today seen as a core legacy of the attacks on the US on September 11, 2001.⁴ Despite no evidence of Iraqi involvement in the events of 9/11, the invasion of Iraq in 2003 was a *de facto* pre-emptive war waged by the US and its erstwhile allies in order to mitigate against such attacks happening again. Effectively institutionalising the logic of pre-emptive war, the Authorization of Use of Military Force Against Iraq, a law passed in 2002, declared that “the USA would reserve the right to attack any nation pre-emptively that it deemed to be a threat to its own national security and interests” (Gupta 2008, 181–182).⁵ It should therefore come as no surprise that the Trump administration, on February 14, 2020, proposed that it was legally authorised to assassinate Soleimani under the same law.⁶ This, alongside other justifications for the assassination of Soleimani, was roundly rebuked at the time, not least by Eliot Engel—the then sitting Chairman of the House Committee on Foreign Affairs—who argued that the “2002 authorisation was passed to deal with Saddam Hussein. This law had nothing to do with Iran or Iranian government officials in Iraq. To suggest that 18 years later [in 2020], this authorization could justify killing an Iranian official stretches the law far beyond anything Congress ever intended” (Engel, quoted in Stepansky 2020).

Although there remains considerable debate as to what constitutes the legality of a pre-emptive military strike under conditions of supposed, imminent or presumed threat (Bellinger 2020), the martial logic of pre-emption increasingly capitalises upon the predictive, presumptive logic of algorithmic extrapolation. The operative calculus of AI is preoccupied with an over-arching goal: calculated prognostication. The central role of prediction in US military strategy was apparent in a statement, made nine months before the invasion of Iraq, by George W. Bush when he announced the following: “[i]f we wait for threats to fully materialise, we will have waited too long” (Bush 2002). Implied in Bush’s statement, whether he intended it or not, was the unspoken assumption that the pre-emptive doctrine underwriting counter-terrorism planning—which would replace the Cold War doctrine of deterrence and containment—would be aided by semi, if not fully autonomous, weapons systems capable of maintaining and supporting the strategy of anticipatory and preventative self-defence.⁷ To predict threat you have to see further than the human eye and act quicker than the human brain; you have to both determine and eradicate, in sum, the “unknown unknowns”.⁸

The fact that the “Bush doctrine” was being articulated at a time when the ambitions and affordances of AI-powered surveillance systems were becoming more and more apparent, not least in their avowed prophetic capabilities, points to an aggregation of events and interests that raises a series of questions: To what extent, for one, can we understand pre-emptive acts of military violence as both the product of a martial logic of preventive self-defence, which is legally dubious, *and* the outcome of algorithmic predictions of presumed threats? Can we, thereafter, define the degree to which algorithmic prediction, in technical and practical terms, not only supports but accelerates the justifications and procedures involved in pre-emptive strikes? If violence is qualified and quantified through the affordances of mathematical predictions, based as they are upon statistical estimations of *potential* futures, then we likewise need to consider the degree to which the calculated abstractions—or “divinatory rationality” (Esposito 2015, 123)—of AI offer, *ex post facto*, a convenient excuse for the continued prosecution of pre-emptive models of autonomous warfare.⁹ Is there, furthermore, a correlation to be observed here between automation and abnegation: does the deferral of decision-making processes to autonomous apparatuses ensure, that is to ask, that we defer our moral responsiveness to the impact of such technologies alongside our legal, political, and individual responsibility for them?

(2) Throughout the algorithmic computation of risk, the apparently oracle-like apparatus of AI seeks to guarantee that the aporetic—that which is characterised by the irresolvable, undetermined and the yet-to-be fully identified—is



rendered not only knowable but, crucially, computationally detectable and, for that reason, theoretically and practically destroyable *in the future* (Downey 2023). Schematically inclined to summon forth events (threats) through prognostication, algorithmic rationalisations can and do result in injury and fatality for those corralled into the purview of AI-induced projections. Bearing in mind the collusive potential to be had in the martial logistics of pre-emption and the predictive models in use in unmanned aerial systems, how do we legally account for, and in turn render accountable, the concatenation of military and technological imperatives under current International Human Rights Law (IHRL)? Putting to one side the effectiveness of pre-emptive strikes and their questionable legality under international law, this is a crucial consideration in any attempt to ensure, through human rights legislation, the freedom to live without physical or psychological threat from above (Grief 2022). To this, we should note that the proven and demonstrable psychological and physiological impact of aerial surveillance, whilst less commented upon in discussions of the proportionality and traumatic aftermath of an aerial strike, remains crucial to any discussion of AI and the future of autonomous weapons systems (Loveday 2022).

In an era that sees an opportunistic affinity in the relationship being forged between models of AI-driven predictive analysis and the military deployment of pre-emptive strikes, we must likewise question the future deployment of AI in autonomous targeting systems.¹⁰ The pursuit of a human right to protect communities from aerial threats needs to contextualise the degree to which algorithmic auguries—evident in the predictive mechanisms of the machine learning systems that power autonomous aerial apparatuses—essentially authorise and further galvanise the long-standing martial strategy of pre-emption. This would then become not only a question of accountability—who, or more likely what, decides on the use of a pre-emptive strike—but also a question of proportionality: What is deemed a commensurate response in contemporary models of asymmetric warfare?

Although routinely presented as an objective “view from nowhere” and given the enthusiastic emphasis on extrapolation and prediction, AI-powered systems of unmanned aerial surveillance and autonomous weapons produce heuristic structures to justify the event of actual violence. They occlude paradigms of accountability whilst, based on the seemingly compelling “force of computation” (Bellanova et al 2021, 123) involved in algorithmic estimations, justify pre-emption as a fact of contemporary warfare. This is largely achieved, through the propositional logic of algorithms (their inclination to yield actionable directives), the systematic training of neural networks (through data-labelling), and the systemic reliance on statistical analysis in the structural design of machine learning models. The schematic intentionality, systematic bias, and systemic

(dys)function of algorithms contributes to a deterministic operative model that, when deployed in unmanned aerial systems, can lead to fatal results. This fatal determinism is often presented as the outcome of so-called “black box” systems. The latter metaphor, based on the metaphor of opacity, has become an all too expedient, if not defeatist, allusion to the purported “difficulties” involved in deconstructing the innermost, invariably proprietorial, workings of algorithmic apparatuses (Rudin and Joanna Radin 2019; Radin 2019). However, I want to foreground here an over-arching reality: working from the statistical prevalence of past features, patterns, and occurrences, the structural design of machine learning algorithms strives to generalise outputs from input data. These generalisations inevitably enable a degree of predictive analysis that, in kinetic and non-kinetic warfare, can and do result in death. Regardless of its so-called “black box” architectures, the *modus operandi* of machine learning and AI are, to be clear, directed towards probabilities not certainties: we can only ever “predict outcomes but not what is to come” (Vilém Flusser 2011, 159). To this end, no amount of deconstructing the inner workings of AI, or the “ethics” of its applications, is going to change the simple fact that algorithms are structurally and epistemically designed to generate conceivable and persuasive, rather than certified, prognostications of what *might* or *could* be.

The concern we therefore need to address when we consider developing a proposed new human right to address aerial threats—from autonomous weapons systems, in particular—is whether the codification and substantiation of future threats in algorithmic models of target validation technically and ideologically collude with the military objective of pre-emption. If this is the case, is AI being used to not only “generate” credible targets but also expand the geopolitical objectives of pre-emption? Through its apparently compelling models of computation, AI reifies threat as a way of projecting an apparently all-encompassing command over entire regions and the communities who inhabit them. Given the monolithic organisational structure of the Pentagon and the labyrinthine processes involved in commissioning new technologies, the predictive function of AI conceivably offers a convenient means by which to pursue, for now and in the foreseeable future, an algorithmically aided model of virtual occupation. Bearing these points in mind, are such systems concerned with occupying the future, rather than *just* monitoring the present? What does this mean for human rights more generally and, crucially, how can these processes be rendered accountable through a proposed new human right that will protect people and communities from aerial threats, be they from unmanned aerial surveillance systems or autonomous weapons?

(3) In a move that consolidated the use of advanced machine learning systems and computer vision, it was announced in



2020 that the US military would augment the MQ-9 Reaper Drone with AI.¹¹ From what we now know of the procedures leading up to the missile strike on Baghdad Airport on January 3, 2020, the MQ-9 Reaper drone—equipped with high-resolution digital imaging equipment—would have been instrumental in gathering data (rendered in the form of digitised images of the “target” and the wider environment), marshalling metadata (detailing, that is, the time and date when video footage was captured), and relaying intelligence to on-the-ground crews.¹² This information would be cross-referenced with the GPS (Global Positioning System) locations and known whereabouts of “targets” and human intelligence (HUMINT) obtained from on-the-ground operators. Crucially, this collation of data would be routinely organised and categorised through the use of onboard embedded image processing algorithms, a feature of unmanned aerial systems deployed by the US since at least 2013 (BAE 2012). Following the algorithmic rationalisation of collated information, a potential “target” is flagged to drone operators on the ground *before* a final decision is made about a missile strike. From the outset of authorising a missile strike, then, there is an abiding perception that the algorithmic quantification and qualification of collated data—the epitome of computational sanctioning—not only augments but potentially activates certain decisions.

In her analysis of how algorithms operate in relation to the “crowded data environment of drone images”, Louise Amoore (2020, 16. *Emphasis added*) argues that the “defining ethical problem of the algorithm concerns not primarily the power to see, to collect, or to survey a vast data landscape, but the *power to perceive and distill something for action*.” Insofar as this crucial insight returns us to the over-arching operative logic of an AI apparatus (its inclination to yield actionable directives), it is critical to consider whether artificially augmented auguries of future events not only lead towards but justify the prosecution of action, military and otherwise. Amoore (2020, 17) continues: “As an aperture instrument, the algorithm’s orientation to action has discarded much of the material to which it has been exposed. At the point of the aperture, the vast multiplicity of video data is narrowed to produce a single output on the object. Within this data material resides the capacity for the algorithm to recognize, or to fail to recognize, something or someone as a target of interest.” The innately machinic process of narrowing down an input (data) to a finite point of action (output), or prediction, implies that the affordances of AI have the potential to provoke the event of a missile strike. This is not, thereafter, merely a concern about the presence or non-presence of threat; rather, it is about the inexorable, if not profoundly deterministic, convention of an algorithmic “aperture” that, trained to produce an outcome (prediction), is procedurally focussed on summoning forth a “target”.¹³

The systematic models of data extraction and data labelling used in the training of AI systems likewise evince a singular purpose insofar as the data extruded from digital images provides the input upon which a prediction (output) is based. If intelligence is based on preconceived notions of threat, the human-defined labelling and inputting of data—that is, images extracted from zones of conflict—can only ever generate a paradigm of bias that is related to the *presumed* presence of an actual threat. Digital images of prior conflict (threat) will, theoretically and practicably, presuppose future threat when extruded through a deterministic model of prediction that is actively looking to isolate and highlight risks. In this *mise-en-abyme*, certain classes of images are significantly overrepresented or underrepresented compared to others, thus ensuring that any bias in the data-labelling or input stage will be subjected to “algorithmic amplification” (DiResta 2018) in the output stage of prediction. Any prediction based on input data—images extracted from conflict zones—recursively stimulates, in sum, computational exemplars of paranoid projection in the pursuit of extra-terrestrial dominion and terrestrial dominance.

These processes, largely understood as systematic (involving as they do the inputting of training data into neural networks), have given rise to a “data-driven killing apparatus” (Weber 2016, 28) based on extracted material that is *always already* encoded through the prism of imminent threat, be it contained in images of insurgency, collated metadata relating to apparently insurgent movements and communications, signal intelligence (SIGINT), or other more general forms of electronic intelligence (ELINT).¹⁴ The amassing of metadata—data about data—similarly concentrates the activity of lived life into sequences of anonymous information, so much so, as Pugliese observes (2016, 6. *Emphasis added*) that the “convergence of metadata systems and digitised identification systems exemplifies the rendering of life into an *orderable system* of information through the application of algorithmic formulae.” The focus on the nominally normative and non-normative behaviour systems of entire communities needs to be foregrounded here insofar as the algorithmic extraction and rationalisation of data allows for an “association matrix” to be formalised into a “social network analysis” (Commander’s Handbook for Attack the Network, quoted in Holland Michel 2019, 23–24). Suggesting a broader ambit of societal observation, rather than just a singular focus on insurgency, these association matrices and methods of social network analysis accommodate the identification of a capacious spectrum of individuals who can be targeted, captured, or killed. When we consider how counter-terror operations—alongside the computational summoning forth of events that have yet to occur—disclose how the logic of US military policy in the Middle East and elsewhere is concretised through models of algorithmic violence, these concerns become all the more palpable for the



people and communities who are subject to models of hyper-surveillance, autonomous targeting, and pre-emptive missile strikes. This is not so much about prototypes of threat prediction based on pattern recognition as it is a concern about pattern *precognition* or, as suggested, the summoning forth of *yet-to-be-realised*, threat-infused realities that corroborate and justify the logic of a pre-emptive strike. Such considerations are obviously more expansive and critically urgent than whether a pre-emptive missile strike on a known target—namely, Qasem Soleimani—was legal or not, inasmuch as this logic can be all-too-readily expanded to cover specific communities and, indeed, entire populations.

(4) Military action can be mobilised through the seemingly inevitable progressions involved in the operative logic of an algorithmic “aperture”—its intractable calculative logic—and the systematic labelling and categorising of input data. It can also, significantly, be summoned forth through the *systemic* logic of statistical distortion. Writing in *Resisting AI: An Anti-fascist Approach to Artificial Intelligence*, McQuillan emphasises that the training of the artificial neural networks (ANNs) used in machine learning algorithms and computer vision does not necessarily produce an efficient model of viable prediction. Rather, such activities reveal a system prone to generating probabilistic simulations (forecasts) based on the statistical transformation (extrusion) of data: “Let’s say we are dealing with a video: each pixel in a frame is represented by a value for red, green and blue, and the video is really a stack of these frames. So, when representing the video as numbers, the input into the algorithm is a huge, multidimensional block of data. As the input is passed through a deep learning network, the successive layers enact *statistically driven distortions and transformations of the data*, as the model tries to distil the latent information into output predictions” (McQuillan 2022, 19. Emphasis added). When we consider the use of ANNs to power machine learning and systems of advanced computer vision on UAS and autonomous weapons systems (AWS), it is this sense of systemic distortion and deviation that needs further assessment and qualification.

Operationally vulnerable to innate processes of mechanistic prediction—the deterministic validations of an “aperture”—and thereafter systematically laden (labelled) with political and martial considerations of apparently imminent threat and risk, the systemic calibration of input data can and does generate erroneous outputs (or so-called false positives). The existence of a “false positive” in the targeting apparatus of an autonomous weapons system—the statistically calculated conjecture that an object is a gun rather than, say, a camera—presents a drone operator, often far removed, with a decision that could define the difference between life and death for those coerced into an algorithmic radius of presumed threat. We could note here, to take but one

example, an investigation undertaken by AlgorithmWatch. In this experiment AlgorithmWatch demonstrated that Google Vision Cloud, a computer vision service, “labeled an image of a dark-skinned individual holding a thermometer ‘gun’ whilst a similar image with a light-skinned individual was labeled ‘electronic device’” (Kayser-Bril 2020). Although Google, once alerted to the bias, fixed it, the article concluded that “the problem is likely much broader” (Kayser-Bril 2020). Whilst this is arguably an example of systematic bias, based on the under- or over-representation of certain images in a given training set, the issue is also systemic inasmuch as algorithms can and do hallucinate. We could note here a particularly germane study involving an InceptionV3 image classifier that consistently classified an image of a turtle as a “rifle” (Athalye et al 2018, 19). The authors of the paper noted that as “an example of an adversarial object constructed using our approach”, a 3D-printed turtle was “consistently classified as rifle (a target class that was selected at random) by an ImageNet classifier” (Athalye et al 2018, *ibid*).¹⁵ This occasionally dry technical analysis reveals a profound reality that remains intrinsic to the neural networks and deep learning systems involved in training machines to “see”. They are not only systematically prone to producing specious and biased outputs based on the data sets used in training, they are also systemically susceptible to inventing (or hallucinating) objects that do not exist (Downey 2024).

The fact that the apparatuses that power AI in UAS and AWS prototypes can be subject to machinic determinism (the logic of the “aperture”), systematic bias (in data labelling), and systemic failings (that is, statistical distortions of data) should give cause for concern when we consider how international humanitarian law, alongside any future legislation that is concerned with the ratiocinations of AI, can effectively counter the often fatal hermeneutics involved in machine learning systems. The direct link between autonomous AI-augmented systems of identification—*calculus*—and the eradication of threat—*violence*—finds an all too readily convenient degree of purchase in the presumptive judgements that underwrite the martial stratagem of pre-emption. To question such biases and failings is to enquire into how the alibi of imminent and predicted, but not predictable, threat effectively generates the apparent, albeit legally contested, *right* to act upon it. The projection of threat, as we saw with the missile strike on Qasem Soleimani, is closely aligned with the doctrine of a pre-emptive strike in military operations: to wait for a threat to materialise, it would seem, is to have waited too long. When we extrapolate this to communities and populations living under the threat of drone strikes, the key anxiety is whether or not the justification of deadly pre-emptive action is based on a presumed threat that reveals not so much the presence of an actual risk but, rather, the enduring existence of an algorithmic predisposition to render “real”—through a series of deterministic,



systematic and systemic procedures or formulae—the prospect of menace.

(5) In July, 2023, Alex Karp, the CEO of Palantir, wrote an opinion piece for *The New York Times*. Conceding that the use of AI in contemporary warfare needs to be carefully monitored and regulated, Karp proposed that those involved in overseeing such checks and balances, including his company, the US government, and the US military, face a choice similar to the one the world confronted following the invention of nuclear weapons in the 1940s. “The choice we face is whether to rein in or even halt the development of the most advanced forms of artificial intelligence, which some argue may threaten or someday supersede humanity, or to allow more unfettered experimentation with a technology that has the potential to shape the international politics of this century in the way nuclear arms shaped the last one” (Karp 2022). Admitting that the most recent versions of AI, including the so-called Large Language Models (LLMs) that have become increasingly popular in machine learning, are impossible to understand, for user and programmer alike, Karp further accepted that what “has emerged from that trillion-dimensional space is opaque and mysterious”.¹⁶ However, and despite this impenetrability, the dilemma in halting the use of AI in future wars is judged to be akin to the moral conundrum faced by the inventors of the nuclear bomb during World War II: if one side out-performs the other then the fine balance of the “zero-sum” game involved in paradigms of nuclear deterrence no longer holds and mutually assured destruction is guaranteed. Perceiving this quandary and asserting, somewhat vaguely, that although it will be essential to “allow more seamless collaboration between human operators and their algorithmic counterparts, to ensure that the machine remains subordinate to its creator”, Karp’s overall argument is that we must not “shy away from building sharp tools *for fear they may be turned against us*” (Karp, 2022. Emphasis added).

Karp’s summary of the dilemmas faced in the use of AI systems in warfare, including the peril of machines that turn on us, needs to be taken seriously insofar as he is one of the few people who can talk, in his capacity of being the CEO of Palantir, with an insider’s insight into their future deployment. Widely seen as the leading proponent of predictive analytics in warfare, Palantir does not shy away from advocating the expansion of AI technologies in contemporary theatres of war, predictive policing, information management, and data analytics more broadly. In tune with its avowed ambition to see AI deployed in theatres of war, its website is forthright on this matter. We learn, for example, that “[n]ew aviation modernisation efforts extend the reach of Army intelligence, manpower, and equipment to dynamically deter the threat at *extended range*. At Palantir, we deploy AI/ML-enabled solutions onto airborne platforms so

that users can *see farther*, generate insights faster and react at the speed of relevance” (Palantir 2023). As to what reacting “at the speed of relevance” means we can only surmise this has to do with the pre-emptive martial logic of autonomously anticipating and eradicating threat *before* it becomes manifest (Downey 2023).¹⁷ Palantir’s stated objective to produce projective models and AI solutions that enable military planners to “see farther”, autonomously or otherwise, is not only evidence of its reliance on the inferential, or predictive, qualities of AI but, given its ascendant position in relation to the US government and the Pentagon, a clear indication of how such technologies will determine the prosecution and outcomes of future wars.¹⁸

In September 2023, no doubt with such presentiments factored into their reckoning, the US Department of Defense (DoD) announced that Palantir had been awarded a contract worth \$250 million to research and experiment with artificial intelligence and machine learning (Demarest 2023). This was in addition to Palantir’s investment in the now infamous Project Maven, which it had originally adopted from Google and, in an allusion to the eponymous 1982 sci-fi film, re-named “TRON” (Peterson 2019).¹⁹ Given that the computational and predictive power of AI networks have grown exponentially over the last two decades, with a marked increase in the use of such systems since the 2003 invasion of Iraq, the inducements and advantages of their deployment in theatres of kinetic and non-kinetic warfare appear to be both irresistible and inevitable to companies such as Palantir and other, including Google, Amazon, and Microsoft (Fang 2018, 2019; Weinberger 2019; Capaccio, 2023).²⁰ If the future of global conflict, in the simulations of present-day autonomous aerial surveillance and targeting, is being predetermined and calibrated by algorithms, as Karp observes and public companies readily condone, we will need to deliberate over what forms of violence will be visited on communities and individuals by virtue of being computationally incorporated into this sphere of AI predictive analysis. Algorithms can, as argued, actualise *presumed* threat; they can offer plausible projections that concretise projected risk and, simultaneously, capitalise upon the paranoid of a “possibilistic, pre-emptive culture of technosecurity” (Weber 2016, 108). The epistemologically sanctioned domains of algorithmic prediction—the regimen of epistemic violence—can, that is to argue, both precipitate and warrant actual violence.

The datafication of targets (people) in this realm declares them both invisible to the ocular-centric, fleshy eye and yet profoundly visible within the ambit of the algorithmic gaze. The historical devolution of deliberative, ocular-centric models of seeing and thinking to the recursive realm of algorithms reveals the calculated ordering of subjects in terms of their future disposability, or replaceability—the latter paradigm being an abiding feature of colonialism and its



calculative approach to life and death (Downey 2021). The process of devolving decision-making processes—relating to questions of life and death—to autonomous, algorithmically augmented systems further divulges a causal, if not fatal, link between colonial technologies of representation and the opaque realm of unaccountable neo-colonial apparatuses (Downey 2023). What, we need to ask, will be the future of death in an algorithmic age and who—or, more precisely, what—will get to decide its biopolitical and legal definitions in an era when AI models of predictive analysis actively collude with the martial logic of pre-emption?

(6) Introducing as it did the phantasmagorical vision of unending and perpetual violence, the so-called “war on terror” established a dualism of contending forces that, in its apparently all-encompassing urgency and implied dangers, foreshadowed an entire region in terms of both atavistic and pending threat. To counter such hazards, the evolution of AI and autonomous systems of aerial surveillance is increasingly qualified through the spectre of a purportedly unending phantasm of violence—a haunting that is not so much of the present but of the future.²¹ The martial codification of pre-emption has ensured that it does not have to be contiguous with, nor evidenced through, a corroborated threat, spectral or otherwise. Thereafter, the *modus operandi* of pre-emption, as Brian Massumi (2015, 235) has suggested, is concerned with eradicating any and all threats that are “not-yet-taking-place”. Although we enter here into a speculative domain, in which events are *not-yet-taking-place*, the virtual manifestation of perceived threat—through the algorithmic prediction of threat—can, and often does, justify the summary sanctioning of a pre-emptive drone strike against individuals, communities, and populations more broadly.

For all the apparent viability, if not validity, of the AI apparatuses deployed in wide-area persistent surveillance systems (WAPSS) and autonomous weapons systems, we need to consistently foreground the degree to which “algorithms are political in the sense that they help to make the world appear in certain ways rather than others. Speaking of algorithmic politics in this sense, then, refers to the idea that realities are never given but *brought into being* and *actualized* in and through algorithmic systems” (Bucher 2018, 3. Emphasis added). Following such insights, any consideration of a human right drafted to protect communities from UAS and autonomous weapons needs to acknowledge that algorithms are implicitly involved in the “ranking, classifying, sorting, predicting, and processing” of data and to those ends remain explicitly “political in the sense that they help to make the world appear in certain ways rather than others” (Bucher 2018, 3. Emphasis added). The world, in terms of weather systems, stock markets, the movement of people, and so forth, is nonetheless stochastic: despite recurring patterns, it is ultimately random and unpredictable.²²

The claims for the viability of algorithmic projections would appear to not only elide this fact but deny that other futures can emerge. If political debates and military objectives are defined by a nightmarish, but ultimately self-serving, vision of eternal threat—historically engendered in the so-called “war on terror”—and the spectre of unending hostility, then the algorithms employed therein, being the by-product of a social and political order, can only ever become agents of that paranoia. As we embrace the dominion of politically expedient, ethically devolved, and algorithmically defined inferences of what constitutes *being* in such an abysmal realm, it is all the more urgent that we enquire into what rights protect subjects deemed to be “targets”. What legal structures protect them from becoming expendable in the algorithmic computation of impending risk and threat? We could rephrase this question and simply ask what ultimately guards against their *not* being in the world?

To the degree that the deployment of algorithms in aerial systems of surveillance and targeting summon forth targets, there has been, to date, little by way of legal consideration of this fact. The recognition of a proposed new human right—which would recognize the physical *and* psychological harm of pre-emptive targeting and the widespread use of predictive models in autonomous weapons systems—would address such lacunae. In the domain of phantasmagorical, inexhaustible threat, recognized ontological frames of reference—those that define, albeit ineffectually, current human rights legislation and international laws regarding a subject’s right to life—are challenged by algorithmic predictions that are understood to be the confirmation of an actual as opposed to a phantasmal threat. These apprehensions, in our era of supposedly unending emergency (a convenient clarion call for unremitting forms of aerial surveillance), are far from regional. The subject of contemporary and future wars will be formulated in the shadow of these systems, and we will need to identify and distinguish, as a matter of due legal process, the impact of algorithms on the material realities of everyday life and, indeed, death.

Notes

- 1 Soleimani, who was also the commander of the Quds Force (one of the five branches of the IRGC), was killed alongside the Iraqi politician and military commander Abu Mahdi al-Muhandis and seven others (Scruton et al 2020).
- 2 For a discussion of the effectiveness or otherwise of such strikes, see Jordan (2020) and Khan (2021). For a debate on the perceived necessity of pre-emption, see Sofaer (2003)
- 3 On January 3, Trump addressed reporters from Mar-a-Lago in Palm Beach, Florida, asserting that Soleimani



was “plotting imminent and sinister attacks on American diplomats and military personnel, but we caught him in the act and terminated him.” Donald J. Trump, quoted in Joseph Stepansky (2020).

- 4 In the aftermath of 9/11, it has been argued that contemporary media outlets sought to anticipate or extrapolate their coverage onto future events rather than focusing on the realities of the present or, indeed, the legacies of the past (Grusin 2010). Networked media, Grusin proposed, were increasingly invested in the mediation of *future* events, often mobilizing collective anxieties through the projection of threat-based scenarios. In the run up to the Iraq war, it “was not that any particular military or political scenario was put forward but rather that so many different possible scenarios were premediated that war with Iraq came to seem an *inevitable event*, indeed seemed in many senses to have already been a televisually mediated news event.” Grusin (2010, 43. Emphasis added).
- 5 The 2002 Authorization of Use of Military Force Against Iraq was a joint resolution passed by the United States Congress in October 2002. It effectively authorized the use of pre-emptive force against Saddam Hussein’s government and has since become a mainstay of American foreign policy. See Public Law 107–243—October 16, 2002.
- 6 While there may be, as Eliav Liebllich (associate professor of law at Tel Aviv University) suggests, some debate about the actual status and methods of anticipatory self-defence, “[p]reventive self-defence is quite clearly unlawful.” (Liebllich, quoted in Mia Swart 2020).
- 7 For a discerning account of how the aftermath of the invasion of Iraq in 2003 gave impetus to US investment in the analytical and targeting capabilities of wide-area persistent surveillance systems (WAPSS), see Holland Michel (2019).
- 8 The legal case for pre-emption was progressed by a number of individuals, not least the United States Secretary of Defense Donald Rumsfeld when he uttered his gnomic warning as to why the presence of “unknown unknowns” should give pause to anyone citing the lack of any evidence linking the Iraqi government with weapons of mass destruction or nay intention to distribute them to terrorist groups. For a copy of the transcription of Rumsfeld now infamous remarks, see <https://archive.ph/20180320091111/http://archive.defense.gov/Transcripts/Transcript.aspx?TranscriptID=2636>
- 9 Esposito (2015: 123; emphasis added) proposes that the ‘divinatory’ logic of the web “leads to the emergence of problems relating to the autonomy of subjects and the *openness of the future*.”
- 10 Debates about whether drone-based targeted killing programmes are indicative of an incipient moral indifference to death have become widespread within military, ethical, and legal fields. See, for example, Michael J. Boyle (2015) and Henriksen Ringsmose (2015). For a counter-argument, see Strawser (2010).
- 11 Hambling (2020) notes that the aim was to enable the MQ-9 Reaper to “carry out autonomous flight, decide where to direct its battery of sensors, and to recognize objects on the ground.”
- 12 The baseline settings and configurations for a MQ-9 Reaper drone, as of March 2021 and according to the US Airforce (2021), include “Multi-Spectral Targeting System, which has a robust suite of visual sensors for targeting. The MTS-B integrates an infrared sensor, colour, monochrome daylight TV camera, shortwave infrared camera, laser designator, and laser illuminator. The full-motion video from each of the imaging sensors can be viewed as separate video streams or fused.”
- 13 There is a separate discussion to be had here about the actual training of the artificial neural networks that power machine learning and computer vision prototypes in unmanned aerial surveillance and targeting systems. The training of a system to “see” objects in the world-out-there, that is to note, presupposes a direct relationship between images (in the form of digital or other models of representation) and reality. Such assumptions, needless to say, are not only suspect but prone to levels of ambiguity and ambivalence that often render the deterministic logic of AI improbable and unconvincing (Downey, 2024).
- 14 For an in-depth discussion of how datafication increasingly ensures that war is rendered largely inaccessible to human perception and intelligibility, see Hoskins and Illingworth (2020).
- 15 The video accompanying this study can be viewed at <https://www.youtube.com/watch?v=YXy6oX1iNoA>.
- 16 Elsewhere, Karp has observed that “[w]e understand that all technology, including ours, is dangerous, and that software can be used as a weapon. Lives have been saved and taken as a result of the software products we have built.” (Karp 2022).
- 17 It is notable that in June 2022, Karp travelled to meet with President Zelensky of the Ukraine, while, at a more recent event in the Netherlands in February, 2023, he is reported to have stated the following: “If you go into battle with old school technology and you have an adversary that knows how to install and implement *digitalized targeting* in A.I., you obviously are at a massive disadvantage” (Karp, quoted in Lipton 2023. Emphasis added).
- 18 Palantir was originally funded from In-Q-Tel, the Central Intelligence Agency’s (CIA) not-for-profit venture capital arm. Apart from the CIA, the company has long-standing partnerships with, amongst other, the US



National Security Agency (NSA), the Federal Bureau of Investigation (FBI), the Department of Homeland Security (DHA), the Joint Improvised-Threat Defeat Organisation (JIDO), and the U.S. Immigration and Customs Enforcement.

- 19 Google, under pressure from its employees, pulled out of the contract for Project Maven in 2018. “Palantir is working with the Defense Department to build artificial intelligence that can analyze video feeds from aerial drones... Internally at Palantir, where names of clients are kept close to the vest, the project is referred to as “Tron,” after the 1982 Steven Lisberger film ...” (Peterson 2019).
- 20 In the case of drone footage used to train the AI systems originally in use in Project Maven, before Palantir acquired and started working on it, video from conflict zones was uploaded into an artificial neural network in the form of training data (input) for the purpose of producing efficient models of object identification and prediction (output). Project Maven used Google’s TensorFlow Application Programming Interface (API), a popular open-source machine learning framework that provide tools and libraries for building and training various types of neural networks. For a fuller discussion of these AI systems and their broader impact in the development of autonomous weapons and aerial surveillance, see Downey (2023).
- 21 For a perspicacious reading of how technologically-induced spectres haunt our futures, see Derrida and Stiegler (2022, 113–134).
- 22 Used to define a random sequencing of events, a stochastic process defines the unpredictable evolution of a given event over time. Weather patterns, or stock markets, as noted, are both subject to fluctuations based on numerous, if not innumerable, factors, thus rendering their evolution probabilistic rather than predetermined or definitive. They are, in sum, stochastic rather than deterministic. In AI, a Stochastic Gradient Descent (SGD) is a model used in machine learning to find the best solution by performing steps in random directions—over multiple “epochs” or iterations—so as to attain the best possible answer, classification, or prediction. Generative Adversarial Networks (GANs) often use a stochastic gradient descent (SGD) model, alongside other stochastic variants, as part of their training process and are used widely in models of computer vision including the prototypes deployed in unmanned aerial systems (Downey 2024).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing,

adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Amoore, Louise. 2020. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others* Durham NC: Duke University Press.
- Athalye, Anish, Engstrom, Logan, Ilyas, Andrew and Kwok, Kevin. 2018. Synthesizing Robust Adversarial Examples. arXiv, 7 June. Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, PMLR 80. <https://arxiv.org/pdf/1707.07397.pdf>
- BAE Systems. 2012. Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System—Argus-Is. <https://www.baesystems.com/en/product/autonomous-realtime-ground-ubiquitous-surveillance-imaging-system-argusis>
- Bellanova, Rocco, Kristina Irion, Lindskov Jacobsen, Ragazzi Katja, Saugman Francesco, and Rune Suchman, Lucy. 2021. Toward a Critique of Algorithmic Violence. *International Political Sociology* 15 (1): 121–150.
- Bellinger III, John B. 2020. Does the U.S. Strike on Soleimani Break Legal Norms? Centre for Foreign Relations (CFR), January 6. <https://www.cfr.org/in-brief/does-us-strike-soleimani-break-legal-norms>
- Bjering, Jens, Engberg-Pedersen, Anders, Gade, Solveig and Strandmose Toft, Christine. eds. *War and Aesthetics: Art, Technology, and the Futures of Warfare*. Cambridge, MA: MIT Press.
- Boyle, Michael, and J. 2015. The legal and ethical implications of drone warfare. *The International Journal of Human Rights* 19 (2): 105–126.
- Bucher, Taina. 2018. *If...Then: Algorithmic Power and Politics*. Oxford: Oxford University Press.
- Bush, George W. 2002. President Bush Delivers Graduation Speech at West Point. June 1, <https://georgewbush-whitehouse.archives.gov/news/releases/2002/06/20020601-3.html>
- Derrida, Jacques, and Bernard Stiegler. 2002. *Echographies of Television*. Cambridge: Polity.
- Demarest, Colin. 2023. Palantir wins \$250 million US Army AI research contract. C4ISRNET, 27 September. <https://www.c4isrnet.com/artificial-intelligence/2023/09/27/palantir-wins-250-million-us-army-ai-research-contract/>
- DiResta, Renee. 2018. Computational Propaganda: Public Relations in a High-Tech Age. *Wired*, 30 August. <https://yalereview.org/article/computational-propaganda>
- Downey, Anthony. 2021. The Algorithmic Apparatus of Neo-Colonialism: Or, Can We Hold “Operational Images” To Account. *Nordic Journal of Aesthetics* 30 (61–62): 78–82.
- Downey, Anthony. ed. 2022. *Shona Illingworth: Topologies of Air*. Berlin: Sternberg Press & Toronto: The Power Plant.
- Downey, Anthony. 2023. *Neocolonial Visions: Algorithmic Violence and Unmanned Aerial Systems*. Ljubljana: PostScript^{UM} #47 (Aksioma). https://aksioma.org/pdf/aksioma_PostScriptUM_47_ENG_downey.pdf
- Downey, Anthony. Forthcoming, 2024. *The Future of Death: Algorithmic Design, Predictive Analysis, and Drone Warfare*. In



- Jens Bjerling et al. *War and Aesthetics: Art, Technology, and the Futures of Warfare*. Cambridge, MA: MIT Press.
- Esposito, Elena. 2015. Digital Prophecies and Web Intelligence. In Hildebrandt, Mireille and de Vries, Katja. eds. *Privacy, Due Process and the Computational Turn: The Philosophy of Law Meets the Philosophy of Technology*. Abingdon, Oxford: Routledge, pp.121–142.
- Fang, Lee. 2018. Google Is Quietly Providing AI Technology for Drone Strike Targeting Project. *The Intercept*, 6 March. <https://theintercept.com/2018/03/06/google-is-quietly-providing-ai-technology-for-drone-strike-targeting-project/#:~:text=Google%20has%20quietly%20secured%20a,intelligence%20solutions%20to%20drone%20targeting>
- Fang, Lee. 2019. Google Hired Gig Economy Workers to Improve Artificial Intelligence In Controversial Drone-Targeting Project. *The Intercept*, 4 February 4. <https://theintercept.com/2019/02/04/google-ai-project-maven-figure-eight/>
- Flusser, Vilém. 2011; orig. 1985. *Into the Universe of Technical Images*. Minneapolis: University of Minnesota Press.
- Gupta, Sanjay. The Doctrine of Pre-Emptive Strike: Application and Implications during the Administration of President George W. Bush. *International Political Science*, 29 (2): 181–196.
- Grief, Nick. 2022. The Airspace Tribunal: Developing the Human Rights Dimension of Airspace and Outer Space. In Downey, A. ed. Shona Illingworth: *Topologies of Air*, Berlin: Sternberg Press & Toronto: The Power Plant, pp.233–238.
- Grusin, Richard. 2010. *Premeditation: Affect and Mediality After 9/11*. London: Palgrave Macmillan.
- Hambling, David. 2020. MQ-9 Reaper Drones With Artificial Intelligence. *Forbes*, 11 December. <https://www.forbes.com/sites/davidhambling/2020/12/11/new-project-will-give-us-mq-9-reaper-drones-artificial-intelligence/?sh=275445447a8e>.
- Henriksen, Anders, and Jens Ringsmose. 2015. Drone warfare and morality in riskless war. *Global Affairs* 1 (3): 285–291.
- Hildebrandt, de Vries., ed. 2015. *Privacy, Due Process and the Computational Turn: The Philosophy of Law Meets the Philosophy of Technology*. Abingdon, Oxford: Routledge.
- Michel, Holland, and Arthur. 2019. *Eyes in the Sky: The Secret Rise of the Gorgon Stare and How It Will Watch Over Us All*. Boston: Houghton Mifflin Harcourt.
- Hoskins, Andrew, and Shona Illingworth. 2020. Inaccessible War: Media, Memory, Trauma, and the Blueprint. *Digital War* 1 (1): 74–82.
- Jordan, Jenna. 2014. Attacking the Leader, Missing the Mark: Why Terrorist Groups Survive Decapitation Strikes/. *International Security* 38 (Spring): 7–38.
- Kayser-Bril, Nicolas. 2020. Google apologizes after its Vision AI produced racist results. *AlgorithmWatch*, 7 April. <https://algorithmwatch.org/en/google-vision-racism/>
- Karp, Alex. 2022. Annual Letter from the Chief Executive Officer. Palantir, 14 February. <https://www.palantir.com/newsroom/letters/2022-annual-letter/en/>
- Khan, Azmat. 2021. “The Human Toll of America’s Air Wars”, *The New York Times*, December 19. <https://www.nytimes.com/2021/12/19/magazine/victims-airstrikes-middle-east-civilians.html>
- Lipton, Alex. 2023. Start-Ups Bring Silicon Valley Ethos to a Lumbering Military-Industrial Complex. *The New York Times*, 21 May. <https://www.nytimes.com/2023/05/21/us/politics/start-ups-weapons-pentagon-procurement.html>
- Loveday, Catherine. 2022. From the Air: How Trauma and Threat Hijack our Memory System. In Downey, A. ed. Shona Illingworth: *Topologies of Air*. Berlin: Sternberg Press & Toronto: The Power Plant, pp. 255–257
- McQuillan, Dan. 2022. *Resisting AI: An Anti-fascist Approach to Artificial Intelligence*. Bristol: Bristol University Press.
- Massumi, Brian. *Ontopower: War, Powers, and the State of Perception*. Durham, NC: Duke University Press.
- Palantir. 2023. Powering Software Advantage Across Environments. <https://www.palantir.com/offers/defense/army/#airborne>
- Peterson, Becky. 2019. Palantir grabbed Project Maven defense contract after Google left the program. *Business Insider*, 10 December. <https://www.businessinsider.com/palantir-took-over-from-google-on-project-maven-2019-12?r=US&IR=T>
- Public Law 107–243—October 16, 2002. Authorization For Use of Military Force Against Iraq Resolution of 2002. <https://www.govinfo.gov/content/pkg/PLAW-107publ243/pdf/PLAW-107publ243.pdf>
- Pugliese, Joseph. 2016. Death by Metadata: The Bioinformationalisation of Life and the Transliteration of Algorithms to Flesh. In *Security, Race, Biopower: Essays on Technology and Corporeality*, ed. Holly Randell-Moon and Ryan Tippet, 3–20. Basingstoke: Palgrave MacMillan.
- Randell-Moon, Holly, and Tippet, Ryan. eds. *Security, Race, Biopower: Essays on Technology and Corporeality*. Basingstoke: Palgrave MacMillan.
- Read, Russ. 2020. World’s Most Feared Drone: CIA’s MQ-9 Reaper Killed Soleimani. *Washington Examiner*, 3 January. <https://www.washingtonexaminer.com/policy/defense-national-security/worlds-most-feared-drone-cias-mq-9-reaperkilled-soleimani>
- Rudin, Cynthia and Radin, Joanna. 2019. Why Are We Using Black Box Models in AI When We Don’t Need To? A Lesson From an Explainable AI Competition, *Harvard Data Science Review*, 1(2). <https://hdsr.mitpress.mit.edu/pub/f9kuryi8/release/8>
- Rudin, Cynthia. 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* 1: 206–215.
- Scruton, Paul., Clarke, S., Malone, T., and Watson. C. 2020. A visual guide to the US airstrike that killed Qassem Suleimani. *Guardian*, 3 January. <https://www.theguardian.com/world/2020/jan/03/visual-guide-airstrike-that-killed-qassem-suleimani-us-iran>
- Sofaer, Abraham D. 2003. On the Necessity of Pre-emption. *European Journal of International Law* 14 (2): 209–226.
- Statement by the Department of Defense, 2020, Jan. 2. <https://www.defense.gov/News/Releases/Release/Article/2049534/statement-by-the-department-of-defense/>
- Stepansky, Joseph. 2020. Timeline of Trump’s shifting justifications for Soleimani killing. *Al Jazeera*, February 19. <https://www.aljazeera.com/news/2020/2/19/timeline-of-trumps-shifting-justifications-for-soleimani-killing>
- Strawser, Bradley Jay. 2010. Moral Predators: The Duty to Employ Uninhabited Aerial Vehicles. *Journal of Military Ethics* 9 (4): 342–368.
- Swart, Mia. 2020. ‘To stop a war’: Did Soleimani killing violate international law? *Al Jazeera*, 5 January. <https://www.aljazeera.com/news/2020/1/5/to-stop-a-war-did-soleimani-killing-violate-international-law>
- US Airforce. 2021. MQ-9 Reaper. <https://www.af.mil/About-Us/Fact-Sheets/Display/Article/104470/mq-9-reaper/>
- Weber, Jutta. 2016. Keep Adding. Kill Lists, Drone Warfare and the Politics of Databases. *Environment and Planning D. Society and Space* 34 (1): 107–125.
- Weinberger, Sharon. 2019. The Everything War. *MIT Technology Review* 122 (6): 26–29.

Anthony Downey is Professor of Visual Culture in the Middle East and North Africa (Birmingham City University). He sits on the editorial boards of *Third Text*, *Digital War and Memory*, *Mind & Media*, respectively, and is the series editor for *Research/Practice* at Sternberg Press. Downey is the Cultural and Commissioning Lead on a four-year multi-disciplinary AHRC Network Plus award, where his research



focuses on cultural practices, digital methods, and educational provision for children with disabilities in Lebanon, the occupied Palestinian territories, and Jordan (2020–2024). His forthcoming book is titled

Decolonising Vision: Algorithmic Anxieties and the Future of Warfare (2025).

