Critical Military Studies

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rcms20

Probabilities towards death: bugsplat, algorithmic assassinations, and ethical due care

John R. Emery

To cite this article: John R. Emery (2020): Probabilities towards death: bugsplat, algorithmic assassinations, and ethical due care, Critical Military Studies, DOI: 10.1080/23337486.2020.1809251

To link to this article: https://doi.org/10.1080/23337486.2020.1809251

Published online: 09 Oct 2020.

Submit your article to this journal

Article views: 17

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=rcms20

CRITICAL MILITARY STUDIES
https://doi.org/10.1080/23337486.2020.1809251

Probabilities towards death: bugsplat, algorithmic assassinations, and ethical due care

John R. Emery

Stanford University, Center for International Security and Cooperation, Stanford, CA, USA

ABSTRACT

This article explores the principle of due care in war and the myth that improved battlefield technology makes Western warfare inherently more ethical. The discursive construction – which I term *virtuous chaoplexic militarism* – of the US as ethical by virtue of its utilization of technologically advanced modes of killing, seeks to dissolve the ethico-political dilemmas of war into quantifiable problems to-be-solved. This article illustrates this dissolution by outlining the transformation within US military decision-making from an ethics of practical judgement to a computational techno-ethics. To do this, I evaluate two concrete cases of US algorithms of militarism. The first case traces the rise of collateral damage estimation algorithms, colloquially known as bugsplat. I examine how bugsplat is programmed, its fundamental design flaws, and its practical exploitation by commanders to erroneously tick the box of ethical due care. The second case explores the SKYNET machine-learning algorithm that was designed to construct 'legitimate targets' for US drone strikes via heterogeneous correlations of SIM card metadata. While drone strikes are widely praised for their capacity to individualize targeting, the algorithmic process of SKYNET ultimately erodes the individual subjectivity that is foundational for ethics of war through data constructions of 'terroristness.' As both cases demonstrate, the ultimate goal of this virtuous chaoplexic militarism is to render the ethico-political dilemmas of killing quantifiable, predictable, and solvable. There exists an urgent need to interrogate socio-technical interactions in the military setting; and specifically, the degree to which practical judgement has been outsourced to a morally problematic computational techno-ethics.

ARTICLE HISTORY
Received 31 July 2018
Accepted 10 August 2020

KEYWORDS
Drones; technology; ethics; algorithm; due care; militarism

Introduction

During the initial 'shock and awe' campaign of the 2003 Iraq War, the US military ran an algorithmic computer programme called the collateral damage estimation tool (CDET) – colloquially known as 'bugsplat.' It estimated the probable number of civilians that would be killed in a given kinetic strike. On opening day, the estimations presented to Gen. Tommy Franks 'indicated that 22 of the [30] projected bombing attacks on Iraq would produce what they defined as heavy bugsplat – that is, more than 30 civilian deaths per raid. Franks said, "Go ahead, we're doing all 30"' (quoted in Chamayou 2014, 216). Such

CONTACT John R. Emery ✉ jremery@stanford.edu Center for International Security and Cooperation, 616 Jane Stanford Way, Stanford University, Stanford, CA 94305-6055

© 2020 Informa UK Limited, trading as Taylor & Francis Group

algorithmic software for civilian casualty estimation claims to 'produce a large body of *scientifically valid* data, which enable weaponeers to *predict* the effectiveness of weapons against most selected targets' (Joint Targeting Publication 2013, II-15, emphasis added). However, bugsplat relies on wholly theoretical data of probabilities of collateral damage and has never taken into account the actual empirical numbers of civilian casualties in order to retrain the algorithms. What Franks demonstrates is not a one-off anecdote, but a systematic attempt to outsource ethical practical judgement to computation. More recently, CIA-operated drone strikes in Pakistan utilized a machine learning algorithm called SKYNET to gather and interpret SIM card metadata in order to construct 'legitimate targets' based on a risk assessment score of one's 'probability of terroristness.' From the 1990s to today, US militarism has employed algorithmic technologies to erroneously tick the box of ethical due care in war; a blind faith in the computational outputs, no matter how unscientifically generated. Ultimately, I argue that these probabilities towards death have further entrenched a discourse of scientism, objectivity, and techno-rationality that purports to make warfare inherently more ethical by virtue of its utilization of technologically advanced modes of killing.

Drones and the prospect of killer robots have shed light on the ethical dilemmas of technologies of US militarism.[1] Yet, what is often underexplored – whether in the case of the rise of precision-guided munitions (Zehfuss 2011), drones (Enemark 2019; Carvin 2015; Brunstetter and Braun 2011; Gregory 2017), or lethal autonomous weapons systems (Roff 2014; Morkevicius 2014; Horowitz 2016) – is the enabling algorithmic technologies that are constitutive of these diverse weapons systems. Such innovations in death and destruction claim to represent the culmination of the 'ethical war' – a Western ideal of a humanitarian framing of killing in warfare (Zehfuss 2018; Carvin and Williams 2015; Mabee 2016). However, such an overreliance on algorithmic logics ultimately enables what it seeks to constrain. Namely, constructing *us* as ethical because we target 'individuals' with the SKYNET algorithm, and kill with 'precision' munitions while running bugsplat. Nevertheless, technology does not inherently make war more ethical. Instead, these algorithms function to discursively replace due care with a techno-ethics of war that purports a fantasy of control over the inherent uncertainties of conflict. The politics of replacement in this instance allows decision-makers to tick the box of ethical due care by appealing to the advanced levels of technology used to kill, while displacing the moral responsibility for deaths by removing humans one causal step from the act of killing.

What then, is an algorithm? In his book *The Master Algorithm*, Pedro Domingos (2015) offers a simple definition: 'An algorithm is a sequence of instructions telling a computer what to do.' Algorithms are reducible to three logical operations: AND, OR, and NOT. While these operations can chain together in mind-bogglingly complex ways, at its core, algorithms are built out of simple rational associations (Brogan 2016). Nick Seaver takes a more holistic approach to algorithms as something more than a detached and neutral technology. Instead, he 'enacts them as part of culture, constituted not only by rational procedures, but by institutions, people, intersecting contexts, and the rough-and-ready sensemaking that obtains in ordinary cultural life' (2017, 10). Adopting this broader definition avoids the problems of treating the algorithms like bugsplat and SKYNET as if they were devoid of the very human contextual, cultural, linguistic, and organizational assumptions that go into writing, refining, and utilizing them. The methodological choice of exploring the two cases of bugsplat and SKYNET intends to fill an important gap in discussions of ethics of war and technology. Rather than theorizing about hypothetical futures of killer robots, I look empirically from the 1990s to today in order to study the complexities of socio-technical interactions as critical case studies in the context of US militarism. Ultimately, I investigate how meaning is constructed from these supposedly hermetically sealed, objective, and neutral algorithmic technologies that purport a more scientific and ethical means of killing.

For this study, I blend the work of James Der Derian (2000) and Antoine Bousquet (2009) to capture the complex nexus of techno-ethical militarism. First, Der Derian notes how virtuous and virtual both 'originated in the medieval notion of a power inherent in the supernatural, of a divine being endowed with natural virtue. And both carried a moral weight, from the Greek and Roman sense of virtue, of properties and qualities of right conduct' (2000, 772). In more modern usages, virtual has taken on a morally neutral tone in technical usage, while '"virtuous" lost its sense of exerting influence by means of inherent qualities' (*Ibid*). Thus, at 'the heart of virtuous war is the technical capability and ethical imperative to threaten and, if necessary, actualize violence from a distance – *with no or minimal casualties*' (*Ibid*., emphasis original). With virtuous war I add Bousquet's (2008, 2009) chaoplexic warfare – that uncertainty and disorder are simply temporary problems to be solved via technological innovation – as exemplary of contemporary US practices of militarism. Chaoplexic is a combination of chaos theory and complexity theory whereby the top-down military structure is deemed inefficient and therefore must be streamlined into a network-centric warfare for a decentralized and flexible force. Thus, *virtuous chaoplexic militarism* seeks to 'effect ethical change through technological and martial means' by melding technical capability and ethical imperative to actualize violence with asymmetric risk (Der Derian 2000, 772). War becomes more scientific and predictable in a virtuous chaoplexic global battlefield; ethical due care becomes an efficient keystroke.

The article proceeds in the following manner. First, I make the case for why due care is essential for just conduct in war and why it ought to remain solely a form of human practical judgement. Second, I explore the discursive politics of replacement by academics, policymakers, and military personnel that attempts to link *jus in bello* discrimination and proportionality to the level of technology being used to kill. This 'virtuous' linguistic trick shifts ethical decision-making from practical judgement to computation; the realm of ethical *questions* to that of scientific *answers*. In part three, I excavate the evolution of bugsplat, how it should work in theory, how it is utilized in practice, and its fundamental flaws. Part four brings to light how machine learning algorithms like SKYNET undermines foundational assumptions of subjectivity in ethics of war. By taking Louise Amoore's (2014) work on Big Data, I problematize the construction of a risk-assessment score of 'terrorist-ness' whereby individualized assassinations are justified based solely on aggregate population-based metadata. Finally, I conclude with a word of caution as new algorithmic technologies of 'chaoplexic' militarism, like killer robots, are further divorcing decision-makers from accountability for killing. The ultimate goals of virtuous chaoplexic militarism are to render all ethico-political dilemmas of killing into quantifiable, predictable, and solvable risk-assessment scores. This article ultimately calls on us to recognize how the complexities of socio-technical interactions and the moral weight of killing are dissolved into hermetically sealed computational answers in a techno-ethics of war.

Ethics of due care in war

What does it mean to exercise ethical due care in war? Michael Walzer in his classic book *Just and Unjust Wars*, illustrates what due care may look like in practice with the WWI memoir *Old Soldiers Never Die* by Private Frank Richards of the Royal Welsh Fusiliers:

> When bombing dug-outs or cellars [in France], it was always wise to throw the bombs into them first and have a look around them after. But we had to be very careful in this village as there were civilians in some of the cellars. We shouted down to them to make sure. Another man and I shouted down one cellar twice and receiving no reply were just about to pull the pins out of our bombs when we heard a woman's voice and a young lady came up the cellar steps … She and the members of her family … had not left [the cellar] for some days. They guessed an attack was being made and when we first shouted down had been too frightened to answer. If the young lady had not cried out when she did, we would have *innocently murdered* them all (2006, 154 emphasis added).

Walzer utilizes this case to illustrate that due care calls for the soldiers to put themselves at some risk in order to protect civilians. If there had been German soldiers in the cellar, they might have scrambled out firing and it would have been far more prudent to simply throw the grenades in the cellar without shouting down, which *military necessity* would have justified him doing so. However, Richards was 'surely doing the right thing when he shouted his warning. He was acting as a moral man ought to act; this is not an example of fighting heroically, but simply of fighting well. It is what we expect of soldiers' (Walzer 2006, 154). Due care is, in essence, a 'positive commitment to save civilian lives. Not merely to apply the rule of proportionality and kill no more civilians than is militarily necessary' (Walzer 2006, 155).

The work of Neta Crawford in *Accountability for Killing*, follows Walzer's understanding of due care. In exercising due care, commanders should attempt to reduce the risk to non-combatants or forego the attack altogether. Crawford views Walzer's formulation as a delicate balancing act where protecting civilians ought to be weighed more heavily than force protection, which is admittedly a difficult task (Crawford 2013a, 165). She concludes that 'the civilian deserves greater protection because although both soldiers and civilians are vulnerable in war, the greater relative vulnerability of civilians means that those who are relatively less vulnerable come second in our evaluation' (*Ibid*., 216). These categories are in a constant state of flux, and balancing military necessity, force protection, and non-combatant protection are always uncertain. Yet, Walzer, Crawford, and I err on the side of protecting civilians, while recognizing that commanders also have a lower moral duty to protect their soldiers. However, when the risk is asymmetric in airstrikes and drone strikes it undermines moral assumptions that underpin the right to kill in war (Renic 2018, 2020). Due care is a deliberative process, and issues arise when this human practical judgement is outsourced to algorithmic computation that outputs an answer to the complex ethico-political dilemmas that maintain degrees of uncertainty. This process is a meaningful and necessary *inefficiency* that is sought to be eliminated in virtuous chaoplexic war.

In line with Schwarz's (2016) question of how drones might shape our *capacity* to think ethically, this article examines how these algorithms of militarism and their technologics discursively function to *replace* ethico-political decision-making with 'objective, neutral, and quantifiable' risk assessments (Hagmann and Cavelty 2012). Yet, for ethical deliberation in asymmetrical war, the *jus in bello* criterion of discrimination should 'be more closely linked to a principle of due care than to considerations of proportionality' (Schwenkenbecher 2014, 95). Without mutual risk, due care means setting a high bar of positive commitment towards protecting civilians and avoiding foreseeable harm which the principle of proportionality and doctrine of double effect may otherwise permit (*Ibid.*, 100). Otherwise, absent mutual physical risk, there is no longer ethical war, '[f]or only a killing that is warlike is supposed therefore to be morally better than mere slaughter' (Enemark 2017, 10). Ultimately, as technology further insulates our troops from harm, bugsplat and SKYNET function to tick the box of ethical due care and exacerbate 'risk transfer militarism' from *our* soldiers to *their* civilians (Shaw 2002).

Techno-ethics: linking ethical due care to technical precision

Policymakers, practitioners, and academics alike appeal to technological innovation as making war inherently more ethical; the 'virtuous' element of chaoplexic militarism. In the early days of the Second Iraq War, Navy Captain Arthur Cebrowski, claimed that 'network-enabled armies kill more of the right people quicker. With fewer civilian casualties, warfare would be more ethical. And as a result, the US could use military might to create free societies without being accused of imperialist arrogance' (Shachtman 2007). Similarly, Donald Rumsfeld stated: 'Our military capabilities are so devastating and precise that we can destroy an Iraqi tank under a bridge without damaging the bridge. We do not need to kill thousands of innocent Iraqis to remove Saddam Hussein from power' (Kaag 2008). Here, technical capability is evoked as an end in and of itself, divorced from the empirical outcomes. This linguistic trick shifts the focus from the *jus ad bellum* considerations of whether the war itself is just, to the narrow *jus in bello* linkage of proportionality to mission-specific technical accuracy. Thus, virtuous militarism constructs technical precision as inherently ethical warfare. Although we 'need not kill thousands of innocent Iraqis', that is precisely what we did; the means are divorced from outcomes.

This techno-logic has become especially prevalent with the expansion of CIA drone strikes outside of declared warzones. Drones are touted as 'the Most Humane Form of Warfare Ever' (Lewis 2013), and 'morally obligatory' weapons (Strawser 2010). Kenneth Anderson (2012) believes that the drone:

> 'provides a *deus ex machina* and an escape from the *jus in bello* proportionality trap . . . The technology provides force protection to (one side's) combatants; it provides greater protection to civilians through precision targeting. What's not to like? No weighing up of perplexing values needs to take place, because everything is on the plus side, win-win.'

The fear with this techno-logic is that if 'precision weaponry is assumed to be inherently ethical, it may grant policymakers and strategists the chance to conflate the description of tactics with the prescription of normative judgements' (Kaag 2008). Such practices are depoliticized within life-affirming and humanitarian discourses, whereby the 'matrix of war invokes life as the ultimate purpose of its operations' (Jabri 2006, 60). Ultimately, the virtuous chaoplexic warfare of today is an ethics constituted solely by the technology used to kill, irrespective of outcomes.

The utilization of precision-guided munitions (PGMs) or ‘smart bombs’ rose in conjunction with the bugsplat algorithm throughout the 1990s. To better understand this historical unfolding, I build upon the work of Patricia Owens (2003) and Maja Zehfuss (2011) by expanding their understanding of language construction with these empirical cases. Owens discusses the rhetorical power of construction that non-combatant deaths caused by Western militaries are only ever ‘accidents’ because we could never *intentionally* target civilians. The question of intention is brought to light by an over-reliance on a techno-logic of algorithmic programming that not only rationalizes civilian deaths as *a priori* accidental, but also raises the deeper question that these acts may be ‘beyond intention’ (Owens 2003). Furthermore, I add to Zehfuss’s (2011) discussion of how ‘smart bombs’ are a production of ‘us’ as ethical because we ‘bomb precisely’, which depends on a ‘curious fusion of intent and outcome, *a fantasy of control*’ (561 emphasis added). Along these lines Beier (2017) makes the compelling case that smart bombs have been fundamental to the blurring of the agent and subjecthood. Constitutive of PGMs were collateral damage estimation algorithms like bugsplat, which claimed to predict probabilities of civilian casualties for any given kinetic strike. For Beier, there are three intertwined and mutually reinforcing rhetorical moves that surround the evolution of PGM discourse from indiscriminate to precision bombing that exemplify virtuous chaoplexic militarism:

> the denial of a viable oppositional subject position; the mystification of sites of subjecthood that is affected by discursive and semiotic construction of weapons averring varying degrees of autonomy; and the apparent predilection to impute agency to weapons themselves such that they may even be read to be occupying some measure of a subject position in the ethical practice of war (Beier 2017, 11).

This article expands these analyses by deepening the understanding of technologies of militarism where ‘faith in the ethical conduct of war has increasingly become coterminous with faith in the weapons’ via the algorithmic mechanisms that further enabled the virtuous chaoplexic discourse of techno-ethics (*Ibid*., 10). Ultimately the case studies of bugsplat and SKYNET fundamentally call us to reassess how ethico-political dilemmas in war are being discursively replaced by a computational techno-ethics with these probabilities towards death.

Bugsplat: epistemologies of algorithmic killing

Collateral damage estimation algorithms are today touted as a technical solution to the ethico-political dilemmas of killing civilians. Bugsplat claims to offer an objective, scientific, numerical probability of civilian casualties in any strike context, from which military commanders can exercise practical judgement. However, in practice, bugsplat ticks the box of ethical due care by constructing collateral damage as always already beyond intention. Thus, one can claim that we ran the algorithm and it predicted no casualties, but 100 civilians died; ethical due care becomes procedural rather than an ongoing evaluative practice. To understand problematic assumptions that undergird such algorithms, how they outsource judgement to computation and defer accountability for killing, it is important to trace the evolution of bugsplat in its various applications. To this end, in what follows I explore: the origins of bugsplat (the historical process by which the term bugsplat came to be used in virtuous chaoplexic militarism); what constitutes 'ideal' bugsplat (how algorithms *ought* to be utilized in conjunction with the practical judgement of commanders); the four fundamental flaws of bugsplat; and finally, by opening the black box of these particular algorithms, I assess how their design and utilization shifts ethical discourse of due care from judgement to computation.

Origins of bugsplat

The origins of bugsplat can be traced to the First Gulf War and the Joint Warfighting Analysis Center (JWAC) in 'developing precision targeting options' (Sewall 2016, 154). But after the immense and foreseeable humanitarian consequences of attacking Iraq's electrical grid and infrastructure, the US Air Force (USAF) began to ask: 'how could airpower more predictably and discretely disable infrastructure in the future, such that it would yield only desired military effects?' (*Ibid*.). During the Iraq no-fly-zone campaigns of the 1990s – Operation Northern Watch and Operation Southern Watch – these simplistic algorithms began to be trained on the effects of different munitions on various targets. However, during the 1990s Kosovo campaign immediate kinetic effects of bombs and other strategic applications began to take on a moral dimension by adding calculations of probably civilian casualties. Indeed, NATO began utilizing the algorithmic software – Conventional Casualty Estimation Tool and the Collateral Damage Estimation Tool (CDET) – which were run on over 400 targets in Kosovo. CDET was originally developed for 'preplanned attacks against fixed targets; it was not initially envisioned to have tactical applications. CDET modeling required extensive information on the materials in the target, population density, terrain, aim points, munitions, and so forth' (Sewall 2016, 158). The algorithm used three-dimensional modelling for a 'high fidelity assessment' of probabilities of collateral damage (Crawford 2013b, 242).

The CDET algorithm could, at its most advanced modelling levels, 'simultaneously replicate multiple parameters of a proposed strike and predict its impact on people, both directly through explosion and blast and also indirectly through the destruction of buildings or even the trajectory of shards of broken window glass' (Sewall 2016, 155). According to the Combined Joint Chiefs of Staff Instruction 3160.01 released in 2002, combatant commanders were required to estimate, evaluate, and mitigate potential collateral damage. Per this directive several algorithms already in use and in the developmental stages 'were to be utilized for evaluating potential targets and estimating both casualties and collateral damage' (cited in Crawford 2013a, 351). Given the level of complexity and information necessary, it took at least four hours to run CDET. In other words, the process was technically complex, lengthy, and too tedious for the new War on Terror.

The driving impetus for CDET was, at its core, to devise a way to gauge the predictability of military effects, yet the story told today *ex post facto*, is one of humanitarianism and minimizing civilian harm. The USAF modellers of CDET believed they had 'developed ways to estimate the direct physical impact on physical persons' computationally (Sewall 2016, 154). Here began the quest to quantify probabilities towards death, such that the 'processes of considering civilian casualties [were] replaced by computerized analysis and the algorithms' (Crawford 2013a, 349). Even at their best, these algorithms only take into account *immediate* kinetic effects, not lasting effects like the decimation of

Iraq's electrical grid during the First Gulf War. However, at its worst, bugsplat actually enables killing civilians by ticking the box of ethical due care irrespective of the outcome.

By the early combat days of Afghanistan in 2001 CDET was still utilized; but this was a different kind of fight that did not have the luxury of time for which the software was originally developed – counterinsurgency as opposed to strategic bombing campaigns. CDET was cumbersome for the accelerated War on Terror and 'led to overestimating likely collateral damage' (Lambeth 2010, 320). As Crawford (2013a, 350) notes, unhappy field commanders thought the 'target approval process was already too slow' and in Lambeth's view, the 'extra caution' of CDET was 'simply another factor in reducing military effectiveness.' Consequently, new software was developed by JWAC, known as the Fast Assessment Strike Tool – Collateral Damage (FAST-CD). This software drew a blob-like two-dimensional footprint of a proposed air strike's estimated blast radius, which gave it the colloquial name of 'Bugsplat'. According to Captain Mary Cohen 'one of Bugsplat's benefits is that it's far simpler to use' (Graham 2003). Speed, ease of use, along with the strategic and normative push to minimize civilian casualties were essential drivers in the development and evolution of these algorithmic software packages. On an accelerated global battlefield, speed and efficiency were the *modus operandi* of bugsplat; instead of four hours, it would take 'as few as five minutes and generally no more than 10′ to produce an output (Sewall 2016, 158).

This evolution in algorithmic CDET, along with 'smart bombs', and more capable intelligence gathering aircraft – i.e. drones – resulted in an increased assumption that civilian casualties could be reined in and the uncertainty of warfare could be tamed. Brigadier General Kelvin Coppock, director of intelligence for the Air Combat Command, stated that bugsplat was a 'significant advance' as 'it will allow us to target those facilities that we want to target with *confidence* that we're not going to cause collateral damage' (Graham 2003, emphasis added). These algorithmic programmes offer a *fantasy of control* to mitigate the unknown consequences, which paradoxically increase probabilities of civilian casualties as it decreases the liability and accountability of war-makers for foreseeable and preventable civilian casualties. Hence, the idea that precision munitions and collateral damage software make war less destructive or inherently more ethical gives practitioners a false sense that the killing of innocents is always already beyond intention; a virtuous war. The techno-logic is thus: the 'ethical war' is only a few software updates away, when FAST-CD fails to accurately predict the level of collateral damage in a timely manner, we develop Advanced CDET and FAST-CD 2.0. In virtuous chaoplexic militarism, ethical *questions* become technical problems to-be-solved with algorithmic *answers*.

'Ideal' bugsplat

Collateral damage methodology (CDM) and collateral damage estimation (CDE) tools demarcate a process of escalatory steps created for the purpose of 'assisting a commander in adhering to the Law of War' (CLAMO 2009). Ideally, bugsplat algorithms and commander judgement are a '[b]alance of science and art that produces the best judgment of potential damage to collateral concerns.' USAF JAG lawyers are keen to point out what CDE is not: First and foremost, it is 'not an exact science: Supporting technical data and processes of the methodology are derived from physics-based computer models, weapons test data, and operational combat observations.' With this, all technical data and gathering processes 'contain some degree of inherent error and uncertainty.' Consequently, the guidelines are not a decision itself, but '[m]erely informs a commander's decision. Its application relies on sound *judgement*' (CLAMO 2009, emphasis added). Crucially, the 'CDM does not predict the actual outcome of weapon employment. The operational environment, weapon's reliability, and fidelity of intelligence data are primary factors that account for a CDE output differing from actual combat employment' (CJCSI 3160.01, 2012). Though the CDM follows a rigid process and generates estimated values, 'neither analysts or commanders should be under the impression that these values in any way constitute ground truth, an exact science, or flawless data' (*Ibid*). But such nuance is lost in its practical application.

While the JAG methodology highlights the complex elements of judgement that rests ultimately in commander control, the Joint Warfare Targeting documents paint a much more optimistic picture of these collateral damage algorithms. The 2013 document discussed how a variety of military organizations 'have developed a number of quantitative techniques used to estimate weapon effectiveness and collateral damage risk' (Joint Publication 3-60 2013). These operational and analytical models are utilized to:

> measure and predict munitions effectiveness. These models produce a large body of *scientifically valid* data, which enable weaponeers to *predict* the effectiveness of weapons against most selected targets. Inputs to these calculations include target characteristics (e.g., size, shape, and hardness), desired damage criteria or probability of damage (PD) calculations, and delivery parameters (e.g., altitudes, speeds, dive angles). Model outputs include the predicted effectiveness of selected weapons and target pairings or the number of assets required to create desired effects using specified weapons and/or delivery systems (Joint Publication 3-60 2013, emphasis added).

While the empirical data on what particular bombs do to particular targets may be robust, the morphing of this data into a humanitarian discourse of saving civilians is not supported by how the algorithms were designed, or how they function in practice. The assumptions made in this portrayal of CDM point to the fundamental flaws of algorithms like bugsplat. Numerical 'objectivity' via techno-innovation is held up as an ethical end itself; it is to these fundamental flaws that I now turn.

Fundamental flaws of bugsplat

I identify four fundamental flaws of bugsplat for virtuous chaoplexic militarism: 1) an arbitrary ceiling of 30 civilian casualties 2) a lack of empirical data on civilian casualties 3) systematic overestimations 4) automation bias and black box algorithms. While each of these has major implications for collateral damage algorithms, I want to implore the reader to think about the broader implications of these findings for the future of a more technological battlefield. The socio-technical interactions between decision-makers and how these algorithms ultimately enable rather than constrain the acceptability of civilian casualties.

The first fundamental flaw is that if bugsplat predicts that 30 or more civilian casualties the strike is flagged and triggers further review. 'The ceiling of thirty potential civilian casualties [does] not mean that the strike would not occur, only that permission must be sought and given by a high-level commander or the president or Secretary of Defense' (Crawford 2013a, 355). The issue here is that this arbitrary ceiling of 30 civilian deaths became the threshold for *all targets* irrespective of its military necessity. As one JAG lawyer noted: 'Such approval thresholds were applied regardless of the value of the target – notably delinked from LOAC [law of armed conflict] standards of proportionality. The threshold has been enshrined in doctrine as the non-combatant casualty cut-off value [] and is now specified in operational ROE' (Sewall 2016, 155). This delinking of civilian casualties from military necessity undercuts the laws and ethics of war. Any strike with possibilities of collateral damage, proportionality should always be balanced against military necessity, and due care exercised; not applied universally to any target. In some cases 30 casualties may be acceptable depending on the target, other times it may be immensely disproportionate. Hence, making 30 predicted civilian casualties an arbitrary universal cut-off value risks skewing decision-making in two ways: If a 'legally questionable target can be eliminated with predicted zero civilian casualties, it is likely to move higher up on the preference list even if its military value is minimal. Similarly, political and military decision makers may feel constrained to attack a particular target set simply because it is predicted to cause few civilian casualties' (Sewall 2016, 157–158). Ultimately the technology itself is not an objective algorithmic output, but shapes commanders' capacity to think ethically in problematic ways.

In practice bugsplat appears to have significant effects on military decision-making. In her interviews with high level military commanders, Sarah Sewall of the Air Force Research Institute garnered some candid responses of how these algorithms are utilized in practice. Lt. Gen. David Deptula believes that: 'Modeling provides the ability to demonstrate to your political masters how you can achieve the results you want … You can lift these restraints if you can demonstrate the particular effects of specific weapons, and that will allow for greater application of force' (Sewall 2016, 154). Indeed, the concern is not with minimizing civilian harm, but practically for demonstrating the blast effects of particular weapons to tick the box of civilian oversight in order to gain more leeway in the application of force. Maj. Gen. Charles Dunlap described the algorithmic CDET system as 'a kabuki dance because you don't have the fidelity of systems; they can't tell you the reality.' Beyond the Orientalist implications of this quote, the point is that: 'Key assumptions could be adjusted so that estimates could stay below the specified level of political approval. For example, if a strike was predicted to have 31 civilian casualties, triggering White House review, it would be possible to adjust assumptions such that the recalculated numbers fell below 30 and thereby avoided presidential scrutiny' (Sewall 2016, 156). These interviews, along with the telling opening anecdote of Gen. Tommy Franks demonstrates the deference to the technology itself. They also reveal how the technology can be utilized instrumentally to avoid constraint, oversight, and accountability for civilian casualties. In sum, not only does bugsplat risk divorcing military necessity from proportionality, but killing civilians becomes a technical problem to-be-solved, deferring human judgement to algorithmic computation.

The second fundamental flaw lies with the fact that the bugsplat algorithms have *never* been programmed with the empirical data of civilian casualties; its collateral damage estimations remain completely theoretical. Instead of utilizing actual civilian casualty data to update the algorithms, the USAF does not 'do body counts', and thus, the 'scientific estimation' of bugsplat is a guessing game of outdated population density numbers and theoretical 'best times' to strike. With the focus on CDE itself, a consideration of civilian impact can be divorced from the actual empirical military effects. As one military lawyer complained, the process reflected a 'desire to make a decision based on some objective "number" – no matter how unscientifically reached or misunderstood – rather than a subjective "value"' (Sewall 2016, 157). Furthermore, the algorithmic process had the potential to blind operators by substituting a number for a *real* effect. One modeller complained that others 'want to maneuver the system so they get the results they want for airpower. They think "If I get approval for this, I need to know what to tell the tribal sheik the impact will be." They aren't thinking about how the locals will view the hole and the buildings. They're thinking how many, the numbers, to get approval' (Sewall 2016, 157). Most troublingly, Sewall notes: 'The USAF could have used the actual results of airpower on civilians to corroborate the CDET models or enable adjustments. Yet because the US military for decades dismissed the desirability and feasibility of conducting "civilian body counts," the USAF lacked data to validate or adjust its models' (2016, 156). Consequently, bugsplat was divorced from empirical effects on civilians as the USAF made the *political choice* 'to focus on prediction – theoretical modeling of what would happen – rather than on empirical data regarding civilian effects' (Sewall 2016, 153). Not only does this give a false sense of objectivity, it divorces decision-making from the outcomes; judgement becomes computation, no matter how unscientifically achieved.

The third fundamental flaw lies with systemic overestimations of blast patterns by the bugsplat algorithm. Questionable assumptions from data training during the 1990s Iraqi no-fly-zones were coded into subsequent updates of CDET algorithms. The 'weaponeering process was designed to *underestimate* the effects of a weapon (in order to ensure destruction), whereas civilian protection estimates should instead *overestimate* potential harm so as to understand the outer limits of effects' (Sewall 2016, 155–156, emphasis original). In other words, the algorithm was biased to need bigger bombs to ensure target destruction, not smaller bombs to protect civilian life. The *ex post facto* discourse of bugsplat's humanitarian impulse is not reflected in the coding itself. In Sewall's interview with one of the computer scientists that worked on CDET, the reason for the over-prediction model appeared to be largely a function of the legacy hardware-system requirements. Hence, the design itself has this tendency to underestimate casualties not for any reason other than the training data in the 1990s; i.e. it is much easier to build from what you had as opposed to creating an entire new algorithmic infrastructure. Moreover, if the algorithmic process is focused on blast fragmentation, no matter how sophisticated, in order to make a proportionality calculation it is reliant on accurate population density numbers. Once a war commences any semblance of an ability to accurately measure populations goes out the window. Indeed, the USAF *intentionally excluded* actual empirical data of civilian casualties to empower a virtuous chaoplexic militarism. Flaws are compounded doubly; first when they don't do body counts, second when a software update is assumed to solve the ethico-political dilemmas of the killing of innocents.

The fourth fundamental flaw is automation bias and treating the algorithms like a black box. Socio-technical interactions are complex and constantly evolving. Although there are no specific studies on military deferral to algorithmic technologies, there is a plethora of research across disciplines from the medical field, to aviation, finance, and driverless cars, that suggests humans frequently suffer from automation bias – the 'tendency of people to defer to automated technology when presented with conflicting information' (Wagner, Bornstein, and Howard 2018, 22). In autonomous driving, people might assume the technology 'has knowledge it does not possess' because they assume positive design intent in 'malfunctioning' systems; in fact even 'when presented with evidence of a system's bad behaviour or failure … users may still defer' to the technology (Wagner, Bornstein, and Howard 2018, 23). More troubling, '[a] utomation bias occurs in both naive and expert participants, [it] cannot be prevented by training or instructions, and can affect decision making in individuals as well as in teams' (Parasuraman and Manzey 2010). This automation bias deep implications for military applications of technology.

The assumptions that go into US military utilization of bugsplat algorithms and its variants are to treat the software as if it were a 'black box'; objective and neutral, forgetting that human judgement is always already buried within the code. As Frank Pasquale (2016, 107) argues in the financial sector, the attraction of the 'black box' algorithm is that it *promotes* an 'automation bias'. There is 'an assumption that a machine-driven, software-enabled system is going to offer better results than human judgement. And when the stakes are high enough, automation bias can degenerate into wishful thinking or worse: opportunistic misuse of models to validate' existing practices. Elke Schwarz (2018, 159) discusses the military context: 'techno-authority is implicit in most contemporary wars conducted by US and allied militaries … this authoritative relationship has an effect on our agency and ability to contest technological decision-making.' Furthermore, she highlights Kevin Miller's analysis of technologies of predictive policing: 'In decision-systems, study after study across numerous disciplines has confirmed the phenomenon of "automation bias [that] occurs in decision-making, because humans have a tendency to disregard or not search for contradictory information in light of a computer-generated solution that is accepted as correct"' (Miller 2014, 122). These observations can inform our understanding of the socio-technical interactions of the bugsplat algorithm in virtuous chaoplexic militarism.

The pervasiveness of automation bias and black box algorithms in this context, suggests that even in ideal circumstances, the technology has a powerful framing effect that alters the parameters of military decision-making and enables a discourse of virtuous war. Thus, the use of bugsplat serves an essential discursive purpose in setting the parameters of acceptability of deaths. However, these benchmarks are ultimately arbitrary and devoid of considerations of military necessity as required by the laws and ethics of war. Adjusting bomb size or fuse delay to get the number to 29 instead of 31 becomes the problem to-be-solved as opposed to thinking seriously through the military necessity or proportionality of any given strike. Citing Clausewitz, Olivia Garard (2016) argues: 'Targeting should not be determined by the means [] we use. Rather, it ought to be characterized by the judgment we use to select the target against which we intend to strike.' Furthermore, we 'must not let our desire and bias for action – targeting just because we are capable – to overshadow the need to question, before, during, and after, whether the target remains appropriate.' Nevertheless, the consequences of bugsplat is that the iterative process of ethical due care in targeting is relegated to an algorithmic output.

From judgement to computation

Instead of due care as imagined by Walzer, Crawford, and myself discussed above; the ethical parameters were established by a computer programmer, who inscribed in code an algorithmic protocol to apply to all situations irrespective of context. Yet, bugsplat was born out of concrete contextual circumstances and contemporary outputs reflect original assumptions coded in the algorithms. Although commanders are supposed to continue to exercise practical judgement, judgement has been replaced by computation. The automation bias and the 'kabuki dances' of commanders to tick the box of ethical due care demonstrates that bugsplat enables rather than constrains the killing of civilians. As Schwarz (2018, 159) warns us: 'expertise and technology meet in a powerfully commanding merger. As technology and biology become further entwined, and the logos of the human is predominantly framed as techno-logos, the horizon for contestation of technological decision-making is diminished and techno-paternalism finds its foothold.'

Ethical due care moves from judgement to computation in these complex socio-technical interactions of probabilities towards death. Crawford sums it up best when she says:

> The technical analysis is used to help decision makers stay within the law, but it may also serve to excuse decisions that we might otherwise believe were wrong and to defuse the moral responsibility for actions. The moral tension between military necessity, and discrimination and proportionality are not eliminated, but they are smoothed by the use of technical analysis. In a sense, some amount of authority over *jus in bello* was ceded to the military, then to military lawyers, is then ceded to technical analysis and a form of computer-assisted expertise (Crawford 2013b, 233).

However, excavating the black box of bugsplat reveals that this 'computer-assisted' expertise is rife with flaws that are exponentially amplified with the abuse of these algorithmic tools. Getting the results you want, divorced from empirical outcomes, enables commanders to defer accountability for killing. While technology continues to evolve, especially with the advent of drones, this virtuous chaoplexic militarism has become embedded in the ways the US determines targeting in drone strikes across the globe via SKYNET algorithms and machine learning assassinations.

Machine-learning assassinations

Drones represent perhaps one of the most studied phenomena of militarism of the past decade, and yet, mostly absent is a discussion of the SKYNET algorithms that are constitutive of enabling the link between technological and ethical superiority. Drone assassinations with SKYNET brings to light the culmination of probabilities towards death, from calculating 'bugsplat' with 'smart bombs' targeting buildings, to 'individualizing' targeting in drone strikes. While the US has the technical capability to target individuals globally, they are no longer individual subjects, combatants, or criminals being targeted. Human beings have become shadows of subjectivity, constructed by metadata attached to their movements and behaviours that claim to predict a probability of 'terroristness' now or at some unknown point in the future. As Schwarz (2016) explored in her article on drones and biopolitics 'that which might pose a risk is identified and selected as a justified target merely on the basis of identifiable markers, patterns and algorithmic calculations, and in most cases the exact factors that contribute to the algorithmic determination of targets remain opaque.' Subsequent revelations about the SKYNET program via a leaked NSA PowerPoint allows us to gaze deeper into the chaoplexic methodology of US targeting.

SKYNET was the joint NSA and CIA operation over Yemen and Pakistan where the NSA swept up a dragnet of SIM card metadata to determine targeting for drone strikes.

SKYNET works like any typical modern Big Data business application. The program collects metadata and stores it on NSA cloud servers, extracts relevant information, and then applies machine learning to identify leads for a targeted campaign (Grothoff and Porup 2016). Except, instead of trying to sell the targets something like the business applications, this campaign executes the CIA's 'Find-Fix-Finish' strategy using Hellfire missiles to destroy their targets (Greenwald and Scahill 2015). In addition to processing logged cellular phone call data (so-called 'DNR' or Dialed Number Recognition data, such as time, duration, who called whom, etc.), SKYNET collects user location, allowing for the creation of detailed travel profiles. Turning off a mobile phone gets flagged as an attempt to evade mass surveillance. Users who swap SIM cards, naively believing this will prevent tracking, also get flagged (the ESN/MEID/IMEI burned into the handset makes the phone trackable across multiple SIM cards). Using these sets of metadata, SKYNET pieces together people's typical daily routines – who travels together, have shared contacts, stay overnight with friends, visit other countries, or move around permanently. Overall, the leaked slides indicate that the NSA machine-learning algorithm uses more than 80 different properties to rate people on their probability of 'terroristness' (Grothoff and Porup 2016).

The paradox, then, is that targeting claims to be individualized – by a loitering drone striking a car in the desert – when in most cases, the shadow of subjectivity is all that is targeted; the individualization of killing has eroded the subjectivity of the individual. While 'signature strikes' – or targeting based on one's visual 'pattern of life' – raised alarms in 2013, it was assumed that targeting a 'kill list' was somehow ethically superior. We were not privy to how the intelligence was gathered to 'know' *who* was targeted. New subjectivities emerge in this stage of the virtuous chaoplexic war. These techno-practices of war are a concrete illustration of how the *who* targeted is no longer an individual subject, but a *what* of statistical correlations of probabilities of 'terroristness.' Radical homogeneity is *constructed* from heterogeneous data upon which life and death decisions are based. Hence, the idea of *who* is liable to be killed in war and thus becomes a legitimate target, has been eroded.

The ethical implications are staggering as the subjectivity of the combatant has been replaced by the process of data construction. This process undercuts the moral foundations for why it is ethically permissible to kill a combatant in war. The assumptions that undergird the warrior ethos and just war are both built upon a reciprocal threat between belligerents, which Renic (2020) persuasively argues is undermined by asymmetric risk. However, SKYNET further exacerbates reciprocity challenges as drones are no longer targeting an individual, but a shadow of subjectivity; an algorithmic output of population-based metadata. Ultimately, SKYNET goes far beyond bugsplat in that it is not just deeming all deaths by the US as *de facto* beyond intention; it is a sinister step to make data – no matter how unscientifically gathered and analysed – the sole basis for killing in war.

N = All: statistical death sentencing

At this juncture of the article, algorithmic militarism has moved from ticking the ethical box of killing civilians, to metadata assassinations in attempts to quantify the uncertainties of war. The essence of chaoplexic warfare – that uncertainty and disorder are simply temporary problems to be solved via technological innovation – in this instance is algorithmic quantification. SKYNET presents correlations as a 'risk-analysis formula, which assigns a numerical value to a risk theme by multiplying the probability of occurrence by a figure for the potential impact' presenting a 'rationalization of the future based on engineering risk-assessment methodology', but is nothing more than 'a glorified form of guesstimates' (Hagmann and Cavelty 2012, 81). Opening the black box algorithms of bugsplat and SKYNET it is easy to see how these statistical correlations are 'already enfolded the intuitive and inferential in its very objectivity' (Amoore 2014, 425). Big Data of today, makes important breaks with statistics that presents deep epistemological issues, especially with life and death decision-making.

In his book *The Taming of Chance*, Ian Hacking traces the intellectual and historical processes that led to the birth of modern statistics and the institutionalization of the 'probabilization' of Western intellectual thought. His analysis captures the avalanche of data that early statisticians of the 19th century explored, everything from suicide rates and crime rates, to jury sizes and birth rates, or whatever was of statistical interest. However, Big Data of today, makes a crucial epistemic break with early statisticians whereby *everything* can be quantified and analysed with massive computing analytics (Amoore and Piotukh 2015). This epistemic break in Big Data statistics shifts us from subsets of populations to an n = all dataset. Statisticians consistently argue that even within subsets, the bell curve cannot tell you anything about a particular individual in that group; yet that error is amplified when everything and everyone becomes quantifiable in n = all social world. In the context of drones, Chamayou sums up the concern succinctly: 'But the whole problem–at once epistemological and political–lies in this claimed ability to be able to correctly convert an assembly of probable indices into a legitimate target' (2014, 49).

The first implication is an *epistemology* of population where *n = all*, where Pakistani or Yemeni residents are reduced to a numerical object of interest. They are detached from the population as such, and relegated to a 'chain of analysis' in which the 'person of interest' emerges from the links of 'activities funded;' 'members of;' 'listed;' 'acquainted with;' 'traveled to;' etc. (Amoore and Pouch 2015, 359). The rise of big analytics has rendered *all data tractable*, which 'carves out radical heterogeneity into flat difference of degree, such that it appears as though everything is calculable, everything about the *uncertain* future is nonetheless decidable' (*Ibid*., 361, emphasis added). These algorithmic technologies tend to reduce difference in *kind* to differences in *degree*. This leads to a reduction and flattening of life's daily chaos, so that 'patterns of life' emerge. It is these heterogenous patterns of life that produce a risk assessment in the shadow of subjectivity, as the basis for lethal intervention. Indeed, General Michael Hayden (former director of the NSA and CIA) bluntly states: 'We kill people based on metadata' (Cole 2014). Ultimately, the turn to Big Data and epistemologies of n = all functions to replace judgement with computation under the guise of a 'scientific' calculation of risk of a (un)knowable future yet to come. However, such probability-based computation cannot dictate *values*, such as ethical due care, but 'it now lies at the basis of all reasonable choice made by officials' (Hacking 1990, 4). The paradox of targeting individuals that are a shadow of subjectivity is that the aim is not to confront a concrete dangerous situation that those individuals pose, but rather 'to anticipate all the possible forms of irruption of danger. "Prevention" in effect promotes suspicion to the dignified scientific rank of a calculus of probabilities' (Castel 1991, 288).

Conclusion

The algorithmic logics of SKYNET and bugsplat both enable what they seek to constrain; namely making killing more palatable to the liberal conscience while deferring accountability for killing. These are the most recent iterations in a long history of *virtuous chaoplexic militarism*, where technology is constructed as inherently more ethical, while simultaneously taming chance in war through scientism. However, ethics of due care is a *meaningful inefficiency* and cannot be relegated to computation. No algorithm can account for the totality of circumstances and the irruption of the improbable that characterizes the uncertain nature of warfare. The goal of scientism of an idealized Cartesian rational computer making ethical decisions devoid of emotion, neglects the central role of emotion in our ethical decision-making. As Valerie Morkevicius (2014) argues, 'emotions can help us to act morally in four ways that are particularly relevant for the ethics of war. By informing our moral intuition, generating empathy and holding us accountable for our choices, our emotions – as expressions of our inner soul or conscience – actually guide us towards more ethical behavior.'Ultimately, the ethico-political dilemmas of virtuous chaoplexic militarism cannot be divorced from human judgement; bugsplat and SKYNET represent radical challenges to ethical due care that have gone largely unquestioned in the historical unfolding of probabilities towards death.

In the technological era, the allure of the 'ethical war' seems within reach, the culmination of a decades-long trajectory of virtuous chaoplexic militarism. First with smart bombs and collateral damage estimation algorithms, continuing with machine learning assassinations in drone strikes, and culminating in the killer robots of tomorrow. This article has sought to raise important questions about the necessity of practical judgement in war ethics, and the inability to provide an algorithmic answer these ethico-political dilemmas. The concrete cases of bugsplat and SKYNET demonstrate that attempts to quantify the global battlefield of the US War on Terror, raise significant ethical predicaments. The systematic outsourcing of human judgement to algorithmic computation, has the effect of absolving decision-makers of accountability for killing and justifying existing practices. These empirical probabilities towards death provide a cautionary tale for future military development in the field of AI. A techno-ethics that divorces us from the weight of taking lives in virtuous chaoplexic war is fraught with peril because it relinquishes due care to morally flawed coding.

These military applications are symptoms of wider issues of attempts in late modernity to quantify the unquantifiable and tame chance. Yet such a futile endeavour remains 'a grandiose technocratic rationalizing dream of absolute control of the accidental understood as the irruption of the unpredictable. In the name of this myth of absolute eradication of risk, they construct a mass of new risks which constitute so many new targets for preventive intervention' (Castel 1991, 289). What is at stake in these techno-practices of war is nothing less than the erosion of effective constraints on the use of lethal force because the techno-rationalization of risk assessment has supplanted genuine ethical deliberation in contemporary conflicts. Bugsplat makes due care impossible, and SKYNET shifts categories of legitimate targets from a *who* to a *what* of metadata constructions of individuals. These techno-practices of virtuous chaoplexic militarism attempt to divorce the means of killing from the empirical outcomes of war, and further enable the war machine. It reduces the ethico-political dilemmas of killing and maiming to a 'scientific process' rife with flaws, that puts the virtuous war only a software update away.

Note

1. I frequently utilize the term militarism, because as Mabee (2016) argues, a historical sociological approach to understanding US 'militarism' rather than war, 'broadens out the critical analysis of present-day military practices, by focusing on their long-term institutionalization.'

Disclosure statement

No potential conflict of interest was reported by the author.

References

Amoore, L. 2014. "Security and the Incalculable." *Security Dialogue* 45 (5): 423–439. doi:10.1177/0967010614539719.

Amoore, L., and V. Piotukh. 2015. "Life beyond Big Data: Governing with Little Analytics." *Economy and Society* 44 (3): 341–366. doi:10.1080/03085147.2015.1043793.

Anderson, K. 2012. "Efficiency in Bello and Ad Bellum: Making the Use of Force Too Easy?" In *Targeted Killings: Law and Morality in an Asymmetric World*, edited by J. D. O. Claire Finkelstein and A. Altman, 374–402. Oxford: Oxford University Press.

Beier, J. M. 2017. "Short Circuit: Retracing the Political for the Age of 'Autonomous' Weapons." *Critical Military Studies:* 1–18. doi:10.1080/23337486.2017.1384978.

Bousquet, A. 2008. "Chaoplexic Warfare or the Future of Military Organization." *International Affairs* 84 (5): 915–929. doi:10.1111/j.1468-2346.2008.00746.x.

Bousquet, A. 2009. *Scientific Way of Warfare: Order and Chaos on the Battlefields of Modernity*. Oxford: Oxford University Press.

Brogan, J. 2016. "What's the Deal with Algorithms?" *Slate*. February 2. https://slate.com/technology/2016/02/whats-the-deal-with-algorithms.html

Brunstetter, D., and M. Braun. 2011. "The Implications of Drones on the Just War Tradition." *Ethics & International Affairs* 25 (3): 337–356. doi:10.1017/S0892679411000281.

Carvin, S. 2015. "Getting Drones Wrong." *International Journal of Human Rights* 19 (2): 127–141. doi:10.1080/13642987.2014.991212.

Carvin, S., and M. J. Williams. 2015. *Law, Science, Liberalism and the American Way of Warfare: The Quest for Humanity in Conflict*. Cambridge: Cambridge University Press.

Castel, R. 1991. "From Dangerousness to Risk." In *The Foucault Effect*, edited by G. Burchell, C. Gordon, and P. Miller, 281–298. Hertfordshire, Great Britain: Harvester Wheatsheaf.

Chamayou, G. 2014. *A Theory of the Drone*. Lloyd, J. Trans. New York: New Press.

CJCSI 3160.01A. 2012. "No-Strike and the Collateral Damage Estimation Methodology" https://info.publicintelligence.net/CJCS-CollateralDamage.pdf

CLAMO PowerPoint. 2009. "Slides Used in the US Army Judge Advocate General's School, Center for Law and Military Operations (CLAMO) Class: An Introduction to the Collateral Damage Methodology (COM) and the Collateral Damage Estimate (CDE)." *Government Attic*. Released 30 August 2011. https://www.governmentattic.org/5docs/ArmyJAG-CLAMO-Slides_2009.pdf

Cole, D. 2014. "We Kill People Based on Metadata." *The New York Review of Books*, May 10.

Crawford, N. C. 2013a. *Accountability for Killing: Moral Responsibility for Collateral Damage in America's Post-9/11 Wars*. Oxford: Oxford University Press.

Crawford, N. C. 2013b. "Bugsplat: US Standing Rules of Engagement, International Humanitarian Law, Military Necessity, and Noncombatant Immunity." In *Just War: Authority, Tradition, and Practice*, edited by A. Lang, C. O'Driscoll, and J. Williams, 231–251. Washington D.C.: Georgetown University Press.

Domingos, P. 2015. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. New York: Basic Books.

Enemark, C. 2019. “Drones, Risk, and Moral Injury.” *Critical Military Studies* 5 (2): 150–167. doi:10.1080/23337486.2017.1384979.

Garard, O. 2016. “Targeting Clausewitzian Judgments: Fusing Precision and Accuracy to Strategy and Tactics.” *Strategy Bridge*, September 20.

Graham, B. 2003. “Military Turns to Software to Cut Civilian Casualties.” *Washington Post*, February 21. https://www.washingtonpost.com/archive/politics/2003/02/21/military-turns-to-software-to-cut-civilian-casualties/af3e06a3-e2b2-4258-b511-31a3425bde31/?utm_term=.e7b1effacaef

Greenwald, G., and J. Scahill. 2015. “The Assassination Complex.” *The Intercept*, October 15. https://theintercept.com/drone-papers/the-assassination-complex/

Gregory, T. 2017. “Targeted Killings: Drones, Noncombatant Immunity, and the Politics of Killing.” *Contemporary Security Policy* 38 (2): 212–236. doi:10.1080/13523260.2017.1336296.

Grothoff, C., and J. Porup. 2016. “The NSA’s SKYNET Program May Be Killing Thousands of Innocent People.” *ARS Technica*, February 16. https://arstechnica.com/information-technology/2016/02/the-nsas-skynet-program-may-be-killing-thousands-of-innocent-people/3/

Hacking, Ian. 1990. *The Taming of Chance*. Cambridge: Cambridge University Press.

Hagmann, J., and M. D. Cavelty. 2012. “National Risk Registers: Security Scientism and the Propagation of Permanent Insecurity.” *Security Dialogue* 43 (1): 79–96. doi:10.1177/0967010611430436.

Horowitz, M. 2016. “The Ethics & Morality of Robotic Warfare: Assessing the Debate over Autonomous Weapons.” *Daedalus* 145 (4): 25–36. doi:10.1162/DAED_a_00409.

Jabri, V. 2006. “War, Security and the Liberal State.” *Security Dialogue* 37 (1): 47–64. doi:10.1177/0967010606064136.

Joint Publication 3-60. 2013. “Joint Targeting” 31 January. Available via *Just Security* at: https://www.justsecurity.org/wp-content/uploads/2015/06/Joint_Chiefs-Joint_Targeting_20130131.pdf

Kaag, J. 2008. “Another Question Concerning Technology: The Ethical Implications of Homeland Defence and Security Technologies.” *Homeland Security Affairs* 4(2). https://www.hsaj.org/articles/125

Lambeth, B. S. 2010. *Air Power against Terror: America’s Conduct of Operation Enduring Freedom*. Santa Monica, CA: RAND Corporation.

Lewis, M. W. 2013. “Drones: Actually the Most Humane Form of Warfare Ever.” *The Atlantic*, August 21.

Mabee, B. 2016. “From ‘Liberal War’ to ‘Liberal Militarism’: United States Security Policy as the Promotion of Military Modernity.” *Critical Military Studies* 2 (3): 242–261. doi:10.1080/23337486.2016.1184418.

Miller, K. 2014. “Total Surveillance, Big Data and Predictive Crime Technology: Privacy’s Perfect Storm.” *Journal of Technology of Law and Policy* 19: 105–146.

Morkevicius, V. 2014. “Tin Men: Ethics, Cybernetics and the Importance of Soul.” *Journal of Military Ethics* 13 (1): 3–19. doi:10.1080/15027570.2014.908011.

Owens, P. 2003. “Accidents Don’t Just Happen: The Liberal Politics of High-Technology `humanitarian’ War.” *Millennium: Journal of International Studies* 32 (3): 595–616. doi:10.1177/03058298030320031101.

Parasuraman, R., and D. H. Manzey. 2010. “Complacency and Bias in Human Use of Automation: An Attentional Integration.” *Human Factors* 52 (3): 381–410. doi:10.1177/0018720810376055.

Pasquale, F. 2016. *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press.

Renic, N. C. 2018. “Justified Killing in an Age of Radically Asymmetric Warfare.” *European Journal of International Relations* 25 (2): 408–430. doi:10.1177/1354066118786776.

Renic, N. C. 2020. *Asymmetric Killing: Risk Avoidance, Just War, and the Warrior Ethos*. Oxford: Oxford University Press.

Roff, H. 2014. “The Strategic Robot Problem: Lethal Autonomous Weapons in War.” *Journal of Military Ethics* 13 (3): 211–227. doi:10.1080/15027570.2014.975010.

Schwarz, E. 2016. “Prescription Drones: On the Techno-Biopolitical Regimes of Contemporary ‘Ethical Killing.’.” *Security Dialogue* 47 (1): 59–75. doi:10.1177/0967010615601388.

Schwarz, E. 2018. *Death Machines: The Ethics of Violent Technologies*. Manchester: Manchester University Press.

Schwenkenbecher, A. 2014. "Collateral Damage and the Principle of Due Care." *Journal of Military Ethics* 13 (1): 94–105. doi:10.1080/15027570.2014.910015.

Seaver, N. 2017. "Algorithms as Culture: Some Tactics for the Ethnography of Algorithmic Systems." *Big Data & Society* 4 (2): 1–12. doi:10.1177/2053951717738104.

Sewall, S. B. 2016. *Chasing Success: Air Force Efforts to Reduce Civilian Harm*. Maxwell Air Force Base, Alabama: Air University Press.

Shachtman, N. 2007. "How Technology Almost Lost the War: In Iraq, the Critical Networks are Social – Not Electronic." *Wired*, November 27.

Shaw, M. 2002. "Risk Transfer Militarism, Small Massacres, and the Historic Legitimacy of War." *International Relations* 16 (3): 343–359. doi:10.1177/0047117802016003003.

Strawser, B. J. 2010. "Moral Predators: The Duty to Employ Uninhabited Aerial Vehicles." *Journal of Military Ethics* 9 (4): 342–368. doi:10.1080/15027570.2010.536403.

Wagner, A. R., J. Bornstein, and A. Howard. 2018. "Computing Ethics: Overtrust in the Robotics Age." *Communications of the ACM* 61 (9): 22–24. doi:10.1145/3241365.

Walzer, M. 2006. *Just and Unjust Wars: A Moral Argument with Historical Illustrations 4th Ed*. New York: Basic Books.

Washington Post. 1998. "Fog of War: The Battle for Hearts and Minds." http://www.washingtonpost.com/wp-srv/inatl/longterm/fogofwar/vignettes/v8.htm

Watch, H. R. 2003. "The Conduct of the Air War." *Human Rights Watch*. https://www.hrw.org/reports/2003/usa1203/4.htm#_ftnref29

Zehfuss, M. 2011. "Targeting: Precision and the Production of Ethics." *European Journal of International Relations* 17 (3): 543–566. doi:10.1177/1354066110373559.

Zehfuss, M. 2018. *War and the Politics of Ethics*. Oxford: Oxford University Press.