The Ethics of Precision

Autonomous Drones in Warfare

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Xavier Castillo-Vieira
Spring, 2021

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Signature ______________________________ Date __________
Xavier Castillo-Vieira

Approved ______________________________ Date 5/4/2021
Hannah Rogers, Department of Engineering and Society

Abstract

With the onset of artificial intelligence technology in virtually every avenue of technology, it is inevitable that autonomy will find its way onto the battlefield. AI on the battlefield has already had important implications, including increased accuracy, maneuverability and response time, but it has also opened a lot of legal and ethical debates. While its applications are increasingly impactful, the military, the citizens of the United States, and foreign powers must consider the risks of military applications of autonomy and artificial intelligence. It is ultimately important to balance this technological ambition with a proper assessment of risk and to consider the ethical implications of allowing artificial intelligence the ability to take a person's life. Throughout this paper, risk society will be used as a lens to analyze the implied risks and hazards that this new technology brings. While it is important for the United States to maintain its edge on the current state of global military technology, unnecessary risks jeopardizing civilians should not be taken. This research paper explores the undeniable advantages of autonomy in drones, from increased maneuverability to faster decision making and more efficient reconnaissance, while also balancing the argument with ethics and moral debates. Through a variety of research methods and ethical debates, the conclusion is that while beneficial, autonomous technology in drones should not be deployed until its accuracy and safety is beyond doubt. Currently, it is impossible to predict the trajectory of artificial intelligence, namely deep learning, but it is safe to say that this discussion will only become more important as time and technology progress.

Introduction

Throughout the past two decades, automation and computerization have revolutionized warfare, increasing the deadliness and effectiveness of armed forces across the globe. In the past decade, automation and other technological developments have proven more and more pervasive on the battlefield, so much so that the only line of defense is more computer automation. One example of this self-propelling development is hypersonic missiles. Guided by computers and travelling at speeds of roughly one mile per second, these missiles will be impossible to defend from without artificial intelligence driven defense systems (Fryer-Biggs, 2020). Going further, the United States Military in conjunction with defense contractors are actively working on automating existing drone aircrafts. An example of the autonomous drone aircrafts in development is Northrop Grumman's X-47A Pegasus drone, which currently operates under remote piloting. In an interview with *The Drive* contributor Tyler Rogoway (2020), Richard Sullivan (Vice President of Program Management at Northrop Grumman) discusses the advantages of autonomy in aircrafts, adding "… (autonomy) will optimize from an advanced mission management perspective, the mission effectiveness with all the assets that it has…".

Ulrich Beck's risk society framework will highlight how the risk involved in high-stakes military operations carried out by drones has steered the development of autonomous drones (Beck, 2020). The framework highlights the impact of society's perspective of autonomous drones on their development and use and showcases the importance of analyzing risks in the development of new technologies. Throughout this paper, risk society will be used as a lens to analyze the implied risks and hazards that this new technology brings. In order to properly understand the current development of autonomous drones and their progression onto the battlefield, the following question must be addressed: What is the risk analysis of offensive autonomous drones in warfare and will their benefits on the battlefield justify their use? This question will be addressed over the course of this research paper by considering both sides of the debate. First, research will be done into the benefits that autonomy brings, then the negative factors of this technology will be addressed.

Methodologies

In order to address this research question, Ulrich Beck's Risk society framework will be applied to characterize the risks involved with a battlefield of autonomous drones. Anecdotal and statistical data on the performance of currently operating remotely piloted vehicles will be used to predict the efficacy of autonomous drones. This data will be gathered from official government publications as well as opinion pieces arguing for or against the use of drones. Some key words used in the procurement of sources regarding currently operating drones will include "military drone accuracy", "military drones safety", and "military drone statistics'. Information and data from official government sources will be used as a primary source, as it will provide facts for interpretation. Furthermore, research will be done on the current developments of autonomous drones uses. Keywords used for this portion of the research will include: "autonomous drone development" and "autonomous drone use". While current research and development is ongoing, there are currently no autonomous drones in operation so statistics regarding their operation are not available. Instead, commercial applications of similar technologies will be glossed over to explore the capability of state-of-the-art and autonomous technology. The most reliable source for statistics of ongoing drone uses as well as predictions for autonomous drone deployment will be sources in the United States State Department. For objective opinions and ethical discussions on current drone use and future autonomous drone use, more emphasis will be given to civilian and academic opinions. Opinions will yield the best ideas in the discussion of the risks and ethical dilemmas in the use of autonomous drones. Key words used in the search of such opinion-pieces will include: “ethics of autonomous drones”, “autonomous drone ethics”, etc. For this section of the research, there is no obvious method of categorizing the information because it is not very easily categorized due to the variety of viewpoints and different analysis opinions. This portion of the research will offer excellent interpretations of risk society. Many sources considering the risks and possible failures of autonomy will be exploring risk society even if it is not directly mentioned.

Background

Though the use of modern day unmanned aerial vehicles, also known as drones, began as far back as the early 1930’s, it was until the Vietnam war that drone technology really made its first reconnaissance appearance (Vyas, 2020). Going further back, Austrian generals dreamed of striking their enemies in Venice from a distance. During a campaign in 1849, the Austrian military launched 200 incendiary bombs attached to balloons and attempted to have them drift over Venice. Though this campaign was a failure, it was quite possibly the first recorded instance of bombing remote targets (Vyas). In the 1980’s drones began making rare offensive appearances but only in the early 2000’s did they gain widespread attention and scrutiny with the unveiling of the United States’ famous *Predator* drone.

Once the underlying technology progressed to a high-performance level, drone use quickly became a fundamental aspect of the United States’ battleground tactics and activity. After only a brief spell as a surveillance only drone, the *Predator* was beefed up to carry and fire hellfire missions in combat, according to investigative journalist and author Mark Bowden (2013). The

Predator, primarily operated in Afghanistan at the beginning, proved to be an incredibly valuable tool. Rather than risk military personnel for reconnaissance or offensives, the United States could inflict serious damage on enemy positions while ensuring no military casualties.

Furthermore, drones were soon equipped with satellite-based communications, allowing them to operate by a pilot stationed in the United States. This ability to carry out offensive operations without ever setting boots on the ground “insulated” U.S. pilots and soldiers from the risks of combat. Nowadays, drone technology is so pervasive on the battlefield that the United States military maintains permanent observation over large areas. Targets and persons of interest can be tracked continuously and photographed. In conjunction with other forms of intelligence, such as the interception of communications, drone enabled militaries can gather as much intelligence as they need and enact lethal strikes on command, with virtually no risk. In fact, Lockheed Martin’s RQ-170 Sentinel drone is responsible for a large percentage of the vast intel leading to the operation resulting in Osama Bin Laden’s Death (Bowden). In the words of Richard Pildes, a professor of constitutional law at New York University’s School of Law, “Drones are the most discriminating use of force that has ever been developed,”. Despite their effectiveness and utility, drone use has not steered clear of public criticism.

This criticism stems from the unintentional consequences of drone strikes- civilian casualties. Though it is nearly impossible to avoid one hundred percent of civilian casualties, despite the method of attack, the importance of civilian deaths is almost exasperated when a drone is responsible for the attack (Kreps, 2014). Deborah Johnson (2014), philosopher at the University of Virginia claims that “autonomous systems (of the future) will be different from auto- mated systems because their behavior will not be preprogrammed.” This is the key distinguishing factor between commonplace autonomous systems commonplace nowadays in every household and a truly autonomous system. The largest source of risk and technological hazard is the uncertainty in decisions that an autonomous system can make. Rather than being able to attribute a decision to a particular piece of code or wiring, modern drone autonomy uses more abstract decision-making systems. Like *Alpha Go*, convolutional neural networks act like "black-boxes", where decisions cannot be easily traced and dissected. Every decision made could be adversely affected by any one of the countless parameters guiding a decision. For this reason, there is a large reluctance to give up full decision-making control.

The risk of failure or mistakes is an important driving force in the public perception of future events or technology. Developed by German sociologist, Ulrich Beck, the Risk Society framework is an excellent tool to use when considering the reaction of the general population to risks. As humanity progresses further into the modern industrial era, there are more and more examples of technological hazards (Beck, 2020). These technological hazards, as described by Beck, are unlike natural hazards in the sense that human development created them. Because these technological hazards are produced by society, they are widely seen as avoidable, and put under extensive scrutiny because of that. The use of autonomous drones for offensive purposes falls under this same category.

Risk society is incredibly thorough in its characterization of human thoughts and emotions in the perception of risk consist of three main propositions:

> "(1) the scale and potential for catastrophe is increasing; (2) there is a loss of faith in experts and science to predict and protect people from these technological hazards; and (3) there are increasingly competing knowledge claims (erosion of expert consensus) regarding the management of technological hazards" (Beck,2020)

The previously mentioned propositions relate to autonomous drones. The scale of disaster in modern warfare is always increasing with the development of newer and more powerful explosives. Additionally, with the advent of autonomy, more aircrafts will be under operation at any given time, increasing the probability of catastrophe. This, however, is also matched with increased precision and intelligence leading to attacks. Secondly, there is a general distrust in the experts, in this case the military and weapons developers to protect civilians here or abroad from the unintentional consequences of attacks. Civilian casualties have the largest effect on the support for the use of force out of any motivator in studies (Walsh, 2015). It is also worth pointing out that now, with the widespread use of cellphones and the internet, civilians have a much better perspective on the state of events in foreign countries. Unlike a couple decades ago, now society is consistently flooded with images and recordings of the aftermath of battles in war zones. This has an enormous effect on the support for the use of force, and the sentiment has grown in recent years (Walsh, 2015). The use of offensive autonomous aircrafts in these possible attacks complicates the situation because it widens the range of responsibility for the unintended consequences of offensive operations. The use of autonomy, civilians feel more of the weight of responsibility for catastrophe since they have supported the nation's use of such technology. Lastly, there are "increasingly competing knowledge claims (erosion of expert consensus) regarding the management of technological hazards" (Beck, 2020). This ties directly into the discussion above, where the widespread use of cellphones paired with global internet connectivity allows for more diverse perspectives to be shared.

Analysis

The controversy surrounding the use of offensive drones, autonomous or not, revolves around the accidental casualties that occur during offensive operations. Recently, in July of 2016, the Director of National Intelligence report revealed alarming statistics on the efficacy and precision of the drones currently in use. The data released showed that armed drones operating over Yemen, Somalia and Pakistan killed an average of 59 times the number of humans per strike than similar fighter-jet based offensives in Syria and Iraq (Barela and Plaw, 2016). This statistic is incredibly alarming and reinvigorated the harsh debate over the ethical implications of drones in warfare. At the time, government officials, including the President himself, assured the public that drone strikes were more precise than piloted aircraft attacks. This was constantly drawn into question by critics of airstrikes and humanitarian organizations, though. Letta Tayler (2016), a correspondent for the Huffington Post and Associate Director of Crisis and Conflict Division at Human Rights Watch brings up questions about whether all the information provided was truly transparent:

> The DNI range of 2,372 to 2,581 "combatant" deaths includes a swing number of 209. Because no explanation is provided, we don't know whether the lack of certainty is for a valid reason, or whether the government is including people who should be considered civilians in its "combatants" list.

This question has frequently been brought up, showing a sense of distrust in the government's reports.

During a press conference on April 8th, 2016, former President Barack Obama himself was quoted saying: "What I can say with great certainty is that the rate of civilian casualties in any drone operation is far lower than the rate of civilian casualties that occur in conventional war." While all the details of military operations in warzones will never be These conflicting sources of information seemed to directly contradict information given to the American people at the time, furthering the "loss of faith in experts" and "erosion of expert consensus" consistent with second and third main propositions of the risk society framework (Beck, 2020).

Barela and Plaw, however, argue that this information is not as straight-forward and damning as it seems. Though the results seem alarming at first glance, there are several inconsistencies between the comparison between the drone use in Yemen, Somalia and Pakistan versus the predominantly fighter jet use in Iraq and Syria. First off, the fact that the enemies in Yemen, Somalia and Pakistan embed themselves in populated regions, surrounded by civilians rather than the conditions in Iraq and Syria, where conditions are much more like a conventional war. This distinction is important and highlights the fact that the comparison ignored strongly favorable conditions for strikes in Iraq and Syria. It is safe to say that any attack, whether drone or manual, is more dangerous to the civilian population if the location is densely populated. Additionally, Barela and Plaw point out that there is a very loose metric on the characterizations of attacks between the two cases. The main distinction in the report published by the DNI that classifies attacks in Yemen, Somalia and Pakistan as "drone attacks" is that 95% of the operations carried out in the region were executed by drones. While this is a valuable statistic, it ignores any non-drone attacks that also resulted in civilian casualties, only attributing the drones to the higher rate of civilian casualty. In fact, this region is classified as a "drone warfare" region, which is a classification made by the U.S. military to characterize ongoing military operations in the region. A specific anecdote supporting this claim is that in 2015, 41 out of a total of 61 deaths in the Yemen, Somalia and Pakistan region attributed to drone warfare were actually carried out by a ship-based cruise missile attack (Barela and Plaw, 2016).

While there is still a significant discussion to be had on the efficacy of drones versus human-piloted strikes, it is important to understand the context of the information being reported.

A fall-short of this research paper is the unavailability of military statistics on the usage of drones. The Director of National Intelligence report in 2016 is the most recent cache of information released by the United States, leaving a lot of questions on the current state of drones unanswered. Consequently, there are no available statistics on modern state-of-the-art drone technologies. Perhaps most notably, the military is currently relying on the General Dynamics MQ-9 reaper for the majority of its drone strikes (Hambling, 2020).

Despite controversy and concerns, research and development on autonomous military weapons, including autonomous drones, continues to grow: "In developing and deploying these (autonomous) weapons systems, the United States and other countries appear to be motivated largely by the aspirations of their own military forces, which see various compelling reasons for acquiring robotic weapons…" (Klare, 2019). Klare, defense correspondent of The Nation magazine, suggests that the development of autonomous weapons has only just begun despite concerns about precision and ethics.

Recently, autonomous military technologies hit a technological break-through with the Defense Advanced Research Projects Agency's (DARPA) debut of *Alpha Dogfight*. This technological marvel uses artificial intelligence, more specifically deep reinforcement learning, a machine learning tactic where the machine is run through back-to-back simulations where good behaviors are "rewarded" and bad behaviors are "punished" (Roblin, 2020). *Alpha Dogfight* blew its expectations out of the water, easily defeating seasoned Air Force fighter jet pilots in simulated dog fights during every trial. The performance of *Alpha Dogfight* proves that the next frontier of combat may well be in artificial intelligence and autonomy. Roblin states that autonomous technologies will inevitably surpass human pilots due to three main factors: lack of physical limitations, decision making time, and the ability to process complex data quickly.

Physical limitations are a huge limitation for human-piloted fighter jets, as humans are limited in their ability to perform extreme maneuvers and turns due to excessive G-forces on the pilot. Decision making time and data processing are another incredibly important factor in this comparison between AI and humans. The computer is able to process all the sensor data provided by the aircraft and process it at rates that a human cannot compete with. *Alpha Dogfight* was trained with over four billion iterations and test runs, learning off of each mistake throughout and improving its algorithms dynamically. These reasons bring up the serious debate as to whether the military is limiting itself by relying entirely on physical and remote human piloting.

Klare continues to argue that precision and safety for U.S. soldiers are not the only driving reason behind this surge in military AI research:

> institutional considerations, however, are not the only drivers for developing autonomous weapons systems. Military planners around the world are fully aware of the robotic ambitions of their competitors and are determined to prevail in what might be called an "autonomy race."

These preemptive competitive preparations in order to maintain a leading edge on other countries' developments has led modern day militaries into a quasi-arms race. Klare points out the highly plausible repercussions of this: "…in their haste to match the presumed progress of likely adversaries, states might field robotic weapons with considerable autonomy well before their abilities and limitations have been fully determined, ". While grim, this perspective opens up the conversation into the risk analysis of autonomous drones and weapons.

There are clear advantages and risks of autonomous drones in warfare. The biggest risk factor in the use of this autonomous technology is the possibility of a mistake. Preparing an algorithm to adapt to a dynamic and dangerous environment while assuring there will be no erroneous decisions or software mistakes is an incredibly difficult task that new technologies like *Alpha Dogfight* may not yet be equipped to do. The problem associated with this is that even a minor mistake could lead to massive consequences and loss of life when the system in question is equipped with an arsenal of explosives. There are still clear shortcomings on image recognition technology that could lead to false assessments and mistakes that cannot be afforded (Klare, 2019). This risk is consistently accentuated as the "spectrum" of autonomy is expanded, leaving more and more decisions to the computer as time progresses.

Conclusion

Though their use would have undeniable benefits and would provide an advantage over more traditional methods of drone piloting and maneuvering, it is important to analyze the risks involved with equipping hostile aircrafts with such technology. Currently, autonomous drones in warfare would be characterized as a hazard rather than a risk. This distinction can be made clearer in Beck's *Risk Society* (2000):

> Hazards are often latent and immanent, that is, invisible and untrackable to everyday perceptions. This social invisibility means that, unlike many other political issues, risk must clearly be brought to consciousness, only then can it be said that they constitute an actual threat

This difference is important to mention because currently, there are no autonomous drones operating on the battlefield. In relation to this issue, society is in the very early stages of a risk society, where there is no direct danger imminent, but when autonomous drones are eventually deployed there will be a significant risk and direct threat to societies. The most important factor and most consequential risk involved in the operation of autonomous drones is the unintentional loss of civilian lives. For this reason, from a risk society perspective, the technological hazard of rushing the use of artificial intelligence for offensive purposes is not advisable. From an ethical perspective, autonomy in drones should be limited to non-offensive purposes, such as reconnaissance, intelligence gathering and target acquisition until the underlying technology is beyond-a-doubt capable of discerning targets from civilians and bystanders. It is important for the development of technology, as well as a safe and fair society, to not overreach engineering into areas where it is not prepared to perform well in. Autonomous technology in offensive applications will most likely be prevalent in the future, and it is important to research and develop this technology, but currently, its rollout would be more detrimental than beneficial.

References

Barela, S. J., & Plaw, A. (2016, October 03). Are Drones Less Accurate than Piloted Aircraft? Retrieved October 15, 2020, from https://www.justsecurity.org/33333/drones-accurate-piloted-aircraft-2/

Barela, S. J., & Plaw, A. (2016, August 26). The Precision of Drones: Problems with the New Data and New Claims. Retrieved October 15, 2020, from https://www.e-ir.info/2016./08/23/the-precision-of-drones-problems-with-the-new-data-and-new-claims/

Beck, Ulrich (2000). Risk society revisited: theory, politics and research programmes. In Barbara Adam, Ulrich Beck & Joost van Loon (eds.), *The Risk Society and Beyond: Critical Issues for Social Theory*. Sage Publications. pp. 211--29.

Bowden, M. (2013, November 1). *How the predator drone changed the character of war.* Retrieved March 16, 2021, from https://www.smithsonianmag.com/history/how-the-predator-drone-changed-the-character-of-war-3794671/

Clare, M. T. (n.d.). Arms Control Today. Retrieved October 16, 2020, from https://www.armscontrol.org/act/2019-03/features/autonomous-weapons-systems-laws-war

Col Dyndal, G. L., & LtCol Bernsten, T. A. (2017, July 28). *Autonomous military drones: No longer science fiction*. NATO Review. https://www.nato.int/docu/review/articles/2017./07/28/autonomous-military-drones-no-longer-science-fiction/index.html

Gibbs, S. (2018, January 24). GM sued by motorcyclist in first lawsuit to involve autonomous vehicle. Retrieved October 16, 2020, from https://www.theguardian.com/technology./2018/jan/24/general-motors-sued-motorcyclist-first-lawsuit-involve-autonomous-vehicle.

Hambling, D. (2020, June 10). Why the Air force needs a Cheaper reaper. Retrieved April 03, 2021, from https://www.forbes.com/sites/davidhambling/2020/06/10/why-the-air-force-needs-a-cheaper-reaper/?sh=53b74f98946f

Howard-Williams, R. (2017, October 25). Terrorism and World Risk Society: Resilience, Resentment, and Spectacle. Retrieved November 23, 2020, from https://www.21global.ucsb.edu/global-e/october-2017/terrorism-and-world-risk-society-resilience-resentment-and-spectacle

Jenks, C. (2017, April 1). THE GATHERING SWARM: THE PATH TO INCREASINGLY AUTONOMOUS WEAPONS SYSTEMS. Jurimetrics: *The Journal of Law, Science & Technology*, 57(3), 341 - 359.

Johnson, D. G., & Noorman, M. E. (2014). Responsibility Practices in Robotic Warfare. *MILITARY REVIEW*, 12–20.

Jones, H. (2020, January 1). The Social Ethics of Self-Driving Cars: Public Perceptions and Predictions of Autonomous Vehicle Safety Risks. Contemporary Readings In Law & Social Justice, 12(1), 37 - 43.

Klare, M. T. (2019, March). Arms control today. Retrieved April 03, 2021, from https://www.armscontrol.org/act/2019-03/features/autonomous-weapons-systems-laws-war

Kreps, S. (2014, April 01). *Flying under the RADAR: A study of public attitudes towards unmanned aerial vehicles* - Sarah KREPS, 2014. Retrieved March 16, 2021, from https://journals.sagepub.com/doi/full/10.1177/2053168014536533

Lala, J. H., Landwehr, C. E., & Meyer, J. F. (2020, September 1). Autonomous Vehicle Safety:

Lessons from Aviation: How more than 25 years of experience with aviation safety-critical systems can be applied to autonomous vehicle systems. Communications of the ACM, 63(9), 28 - 31.

Roblin, S. (2020, August 31). Military AI vanquishes human fighter pilot in F-16 simulation. How scared should we be? Retrieved October 15, 2020, from https://www.nbcnews.com/think/opinion/military-ai-vanquishes-human-.fighter-pilot-f-16-simulation-how-ncna1238773

Roff, H. M. (2014, October 1). The Strategic Robot Problem: Lethal Autonomous Weapons in War. Journal of Military Ethics, 13(3), 211 - 227.

Sarmento, A. (2015, February 01). Aiming from a distance: The implications of the use of drones for security. Retrieved March 16, 2021, from https://www.e-ir.info/2014/12/28/aiming-from-a-distance-the-implications-of-the-use-of-drones-for-security/

Shin, H., Lee, J., Kim, H., & Hyunchul Shim, D. (2018, January 1). An autonomous aerial combat framework for two-on-two engagements based on basic fighter maneuvers. Aerospace Science and Technology, 72, 305 - 315.

Tang, M., Alemzadeh, H (advisor), & Seabrook, Bryn (advisor) (2020). Wearable Cognitive Assistant Systems for Emergency Response; An Analysis of Morality in Autonomous Vehicles. University of Virginia, School of Engineering and Applied Science, BS (Bachelor of Science), 2020: Charlottesville, VA. Retrieved from https://doi.org/10.18130/v3-8mxd-3d83

Tayler, L. (2020, October 28). *Airstrike Transparency We Can't Believe In*. Human Rights Watch. https://www.hrw.org/news/2016/07/08/airstrike-transparency-we-cant-believe#.

United States, Congress, Strategic Studies Institute, and Keene, S. D. *LETHAL AND LEGAL? THE ETHICS OF DRONE STRIKES*, United States Army War College Press, 2015, pp. 11–38.

Vyas, K. (2020, July 16). A brief history of DRONES: *The remote controlled unmanned aerial Vehicles (UAVs).* Retrieved March 16, 2021, from https://interestingengineering.com/a-brief-history-of-drones-the-remote-controlled-unmanned-aerial-vehicles-uavs

Walsh, J. I. (2015, October 1). Precision Weapons, Civilian Casualties, and Support for the Use of Force. Political Psychology, 36(5), 507 - 523.