



The X-45 Unmanned Combat Air Vehicle technology demonstrator that is eventually intended to fly high-risk air combat missions.

Future Unmanned Combat Air Vehicles (UCAV) and the Ethics of Responsibility

by Mark A. Sandner

Major Mark Sandner is an Air Combat Systems Officer currently posted to VX-1 Air Test and Evaluation Squadron at Naval Air Station Pax River, in Maryland, USA. He recently completed the RCAF Aerospace Studies Program at the RCAF William Barker V.C. College at 17 Wing, Winnipeg, and has worked extensively with remotely piloted aircraft (RPAs) in the past while posted to Australia from 2015–2017.

Introduction

In the article “*The Swarm, The Cloud, and the Importance of Getting There First*,” Major Blair Helms and Captain Nick Helms of the United States Air Force (USAF) push for a manned-unmanned synergy of operations that allows technology and automation to amplify what is currently possible in the world of remotely piloted air power. They argue that the limiting factor for achieving true operational fusion is *cultural*, not *technological*. Once the cultural acceptance of remotely-piloted aircraft (RPAs) and unmanned combat aerial vehicles (UCAVs) catches up with the technology, then true breakthroughs in capability can be realized. Semi-contrasting this belief is

Captain Michael Byrnes, also of the USAF, who in his article “*Nightfall: Machine Autonomy in Air-to-Air Combat*,” argues that a fully- autonomous unmanned aircraft will bring new unparalleled lethality to the air-to-air combat world.¹

Although the authors of both articles agree that the future of air power lies in autonomous UCAVs taking centre stage, Byrnes goes slightly further in stating that the technical and performance aspects of UCAVs will inevitably lead to a UCAV-dominant air environment. Helms and Helms, by contrast, argue that manned supervision will most likely always be required, and the degree of autonomy that should be given to UCAVs will benefit greatly from having manned supervision. In a purely air-to-air combat scenario, where the UCAV will primarily be used, I argue that the realistic view of the future of air combat is unmanned, due to the rapidly-increasing developments in technology, as well as the pure economics of fielding an air force of unmanned aircraft. And yet, current technology is still in its infancy in terms of machine learning, and there are still questions to be considered in terms of *responsibility* and *ethics* when a machine makes decisions to kill autonomously.

Background

R PAs and UCAVs have evolved considerably over the past one hundred years. The first unmanned aircraft was flown in 1916, less than fifteen years after the Wright brother's historic flight.² The Hewitt-Sperry Airplane, named after the two inventors, was a project funded by the United States Navy (USN). Evolution from the Hewitt-Sperry Airplane has spawned a diverse range of modern RPAs, leading to the first trans-Pacific UAS flight that occurred in 2001,³ performed by a USAF *Global Hawk* from Edwards Air Force Base, USA to Royal Australian Air Force Base Edinburgh, Australia. The flight demonstrated an RPA capability to fly for an extended period at high altitude without ground radar coverage, autonomously. The *Global Hawk* flight was a milestone for unmanned aircraft, and was a precursor of what the future held for the unmanned aerial systems (UAS) industry. Since that flight, leaps in technology have allowed for greater autonomy for RPAs and UCAVs. However, true breakthroughs are not yet present with respect to UCAVs replacing manned fighter aircraft.

Current technologies for UCAVs, although greatly advanced in the last ten years, are still in their infancy in terms of full automation and machine learning. To date, there does not exist

“Current technologies for UCAVs, although greatly advanced in the last ten years, are still in their infancy in terms of full automation and machine learning.”

feasible and credible UCAVs optimized for air combat. The reason for this technology being as-of-yet undeveloped lies in the sheer complexity of the machine logic required to make combat decisions as well as a pilot with years of training and experience. USAF fighter pilot training is, according to Major Kreuzer of the USAF, “largely an algorithmic function.”⁴ Junior pilots learn the basics of

air combat first: manoeuvres straight from the textbook designed to instill a form of muscle memory in pilots when certain circumstances occur in the air. This gives junior pilots the intuition in the air that is such an advantage for an experienced air force. As pilots gain more experience and training, the wisdom of being in the aviation world comes into play, and pilots develop an advanced knowledge of air-to-air combat that one cannot truly learn in a book; the knowledge that comes from hundreds of hours of mastering a craft, when the basics of flight have become second-nature, and the

mind can concentrate on higher-order demands. The tactics and experience that a fighter pilot develops over the many years of scenarios in which one may find themselves would be something that any form of artificial intelligence (AI) would have to master in order to truly be considered more worthwhile than having a manned aircraft in the air.

In their current state, UCAVs are only now being certified for flight outside military controlled airspace, meaning autonomous flight where the aircraft itself is responsible for safety-of-flight



DVIDS/U.S. Air Force photo/6517624

©ERIK HILDEBRANDT

An RQ-4 *Global Hawk* unmanned surveillance and reconnaissance aircraft flying over Patuxent River, Maryland.

duties. These duties refer to the requirements of an aircraft to maintain level altitude, airspeed, on a flight path, and to be able to avoid traffic that is *either* emitting a GPS position *or not* doing so. The technology required for a UCAV to monitor its own flight path and make deviations based upon other air traffic is called a sense-and-avoid (SAA) system. The goal of an SAA system is to enable "...current and future UCAVs to be able to replicate the human see-and-avoid capability at a comparable or superior level upon replacing the onboard pilot."⁵ Accomplishing this goal means that UCAVs will be permitted to operate around the world in any airspace, provided it has the necessary equipment onboard, something that is also under development. The SAA system becomes the eyes of the UCAV. It is therefore vital for matters of flight safety and future combat operations for a UCAV to have a robust sensor suite to detect possible air traffic and enemy intruders, which will allow it to take corrective action while still operating within international flight rules.

The next step for UCAVs once SAA technology becomes commonplace is to develop a machine algorithm that is capable of *fighting* and *winning* in an air combat environment against manned aircraft. Byrnes outlines a hypothetical UCAV in his article which, due to its lack of reliance upon a human pilot, is able to outperform a manned aircraft. The exploitation of a smaller cross-section, lighter weight, and the ability to pull larger positive and negative "Gs," permit this hypothetical aircraft to close in and win a dogfight with 5th generation fighters.⁶ The decision-making process for an aircraft envisioned by Byrnes would be significantly more complex than current RPAs and UCAVs. The ability for a UCAV to learn and adapt to situations based upon the tactics of enemy aircraft is a major challenge for the future of UCAVs.

Discussion on Responsibility

The ability for a UCAV to differentiate between a civilian and a military target is easier said than done. Considering the altitudes at which an aircraft flies, it is a difficult and a time-consuming task for even a manned aircraft to discern whether a person on the ground is friend or foe. To ask a machine to do the same would require an incredibly-complex level of machine learning that currently does not exist. If something like that were to become available, however, the political and legal framework for governing it would need to be in place and ready to accept the technology. Again, easier said than done. Tactical autonomy, as Byrnes calls it, is the ability for a UCAV to make a decision to fire weapons or to perform a set of manoeuvres in response to an enemy's own manoeuvres and weapons. The ability to do so at a level of skill higher than that of a human requires an incredible amount of information coming from many different sources outside the aircraft.

Critics of tactical autonomy state that the information a UCAV would gather to make tactical decisions on its own is subject to spoofing and jamming by the enemy, and thus, the information gleaned and the resultant decision can never be trusted.⁷ Although this is a legitimate consideration, the same can be said for the information being received by a manned aircraft, and the subsequent decisions made by the pilot, based upon that information. In this age of connectivity and cyber-warfare in which the world presently exists, it is increasingly important to protect systems against such jamming or spoofing, and there is no doubt that any future UCAV would include protection from



DND photo BN2011-0028-54



DVIDS/U.S. Navy photo by Mass Communication Specialist Seaman Keenan Daniels/5402645

A Phalanx Close-In Weapon System being fired aboard USS Nimitz.

such attacks, given the direction of the evolving technology. No machine is completely jam-resistant, which can be a problem when that machine is expected to make decisions regarding life and death. If an enemy were to jam a friendly UCAV, that UCAV could begin to make erroneous decisions with respect to who is *friendly*, who is an *enemy*, and who is a *civilian*. The link between the UCAV and its human operator is also at risk, and the severing of such a link would cause the UCAV to operate on pre-programmed settings. Errors such as this could have vast repercussions, not just at the *tactical* level, but also *strategically* and *politically*. Strategically, UCAV assets could not be seen as reliable, and might be removed from the battlefield entirely until either the electronic warfare (EW) threat is removed, or the UCAV can be proven to be making the correct decisions. Until that point, friendly units would face a serious detriment in terms of air support. Politically, a country with easily-jammable UCAVs would be more of a liability than an asset. This could have repercussions in terms of where allies will want the country to operate, and which operations the country would not be allowed to mount, due to national security considerations.

Giving tactical autonomy to a UCAV brings about another set of questions regarding the ethics of permitting a machine to make a decision to kill or to cause harm. Autonomous machines being given the ability to kill is nothing new. From the

cheap-but-effective landmine, to the complex technology of a *Phalanx* close-in-weapon-system, humans have constantly been looking for ways to increase security by offloading some duties to machines and equipment. An automated weapon means that it is capable of acting independently of immediate human control, something of a fire-and-forget system. Weapons such as this have been used for many years and do not raise any ethical questions beyond that of traditional long range weapons.⁸ In 2012, Human Rights Watch, a group that “regularly addresses the issue of robotics and warfare,”⁹ examined the difference between an *automatic weapon*, and an *autonomous machine*, such as a UCAV. The group found there was an acceptable distinction between autonomous weapons that were human-supervised, and automated weapons. Given this finding, UCAV find themselves in less questionable territory in terms of ethics, as long as there is a certain level of human control in place. The human element becomes important, not only in terms of the ethics of machine killing, but of who is responsible for a UCAV taking a life.

Responsibility for a weapon lies with the officer or official in charge, be they the aircraft captain in relation to an aircraft, or a base commander if it is a stationary weapon. This responsibility is fed back to the state to which the individual belongs, and is tied to the laws that govern warfare for allied states. If states were to use UCAVs in a killing role, the responsibility for those weapons

“Responsibility for a weapon lies with the officer or official in charge, be they the aircraft captain in relation to an aircraft, or a base commander if it is a stationary weapon.”



Wreckage of a car destroyed by a US drone air strike that targeted al Qaeda militants in August 2012 in al-Qatn, Yemen.

and the decisions of the autonomous systems still needs to fall upon the parent state.¹⁰ This is an important specific requirement for state users of UCAVs, because it will stop states from shirking responsibility when a UCAV fires a weapon at the wrong target. An example of this would be if a UCAV fired a weapon at a target and unintended civilians were hit. It would be a legally grey area regarding the state responsibility, and thus, specificity in this area will become extremely important. Another example would be if a UCAV fired a weapon when it was not supposed to do so, based upon an error in an algorithm, or upon its machine learning. This would present a difficult situation in deciding just who is responsible for the accidental deaths, since the officer in charge did not intend upon firing the weapon, and the UCAV did it autonomously.

Robert Sparrow is an Adjunct Professor in the Centre for Human Bioethics at Australia's Monash University, where he works on ethical issues raised by new technologies. A leading authority in the field, his book, "*Killer Robots*," provides some interesting discussion regarding who should be held responsible for possible war crimes in a situation that involves a UCAV making incorrect decisions with respect to taking lives.

The Programmer

Sparrow posits that it could be easy to blame the person who designed or programmed the UCAV's decision-making algorithm, since they are the ones that incorrectly *designed* the system. He then argues that this is not the case for two reasons: the possibility that the UCAV may attack wrong targets could be a known limitation of the UCAV (it was designed with these limitations, and they were not an oversight), or the possibility that the UCAV made a choice other than that programmed or predicted, due to its autonomous, machine learning nature.¹¹ The fact that in this case, the UCAV made a choice autonomously proves that the choice was not an original design, which would prove that it is truly autonomous. It would not be feasible that the programmer would be at fault for designing a system that makes its own decisions, even if sometimes those decisions are erroneous, for that was the stated requirement at the outset.

The Commanding Officer

Sparrow states that the argument for the commanding officer to have responsibility for UCAV decisions lies in the fact that traditionally, "...the officer who ordered the deployment

of the weapons system should instead be held responsible for the consequences of its use.”¹² This seems to be preferred for states utilizing UCAVs, since it fits in with the traditional rules that govern conflict of modern militaries, and it makes the most sense. The officers that use these weapons should be the ones held responsible for their misuse. However, this argument’s flaw is that it does not take into account the autonomous, ‘smart’ natures of *future* UCAVs. The prime advantage of future UCAVs will be that they can make their own decisions, sometimes with better information than the commander. If future UCAVs are treated in the same way as ‘dumb bombs,’ then there is no real difference between the two, and the advantage of using such weapons is gone. In order to have effective autonomous UCAVs, with the ability to make their own decisions with respect to where to drop weapons, the military must accept that sometimes it will make mistakes, just as with any manned weapon. The more autonomous a UCAV becomes, the more risk the military must accept regarding its independent decisions being right or wrong. However, the smarter these UCAVs become, the higher the probability that the decision they make will eventually be the correct one. And at some point, it will likely not be fair to hold the commanding officer responsible for the UCAV’s decisions.¹³

The Machine

Sparrow’s third discussion point regarding responsibility submits that the UCAV itself should be responsible for its own decisions. The idea that a machine could be held morally responsible for causing a death is an odd one, for machines do not (presently) understand the difference between right and wrong, good or bad. They merely understand what they are taught or programmed. The

human ideals of good and evil are too complex for a machine, and thus, we cannot hold a machine guilty for something of which it has no understanding. Sparrow raises the point that to hold a machine morally responsible for an action, there must exist a possibility that it can be rewarded or punished for good or bad behaviour. This is another difficult concept to consider, for punishment and reward stem from being satisfied with an outcome, or from feeling a sense of suffering. How to ensure a machine is punished or rewarded for right or wrong actions is an entirely different discussion that borders upon the futuristic, or the realm of science fiction, and it introduces further complexities to the discussion at hand. It may be possible at some point in the future that machine AI could be held responsible for its actions, based upon the moral beliefs of the machine, and the human operators/supervisors would then be absolved of responsibility. However, UCAV technology is nowhere close to being evolved to that point.

The three arguments that Sparrow advances all lead to issues regarding the technology used, and they do not offer a clear-cut solution to the issue of responsibility. The programmer is one that can be ruled out; just as a tradesman does not blame the manufacturer of his tools for the tradesman’s improper use of them, one cannot blame the programmer for designing a UCAV with an AI that makes its own decisions. An AI making mistakes proves that it is truly autonomous. Placing responsibility on the UCAV is also not a solution that would work, given present technology. The idea of a morally-responsible machine simply is not possible just yet. Machines are not at a point where the ideals of good and bad can be taught to them. Thus, the only practical, responsible solution at this point in time is for UCAVs to have a human being be responsible for its actions, and then to be held



An MQ-9 *Reaper* flies a training mission over the Nevada Test and Training Range, 15 July 2019.



Major Bishane, a 432nd Aircraft Maintenance Squadron MQ-9 Reaper pilot, controls an aircraft from Creech Air Force Base, Nevada, 6 May 2015.

accountable or rewarded as appropriate.¹⁴ This is a requirement, not just in accordance with the laws of armed conflict, but for the public to provide support for the use of such weapons. Admittedly, this is not the perfect solution. The fact that a human being could be responsible for the actions of a UCAV that he or she did not personally order is unjust, and it sparks other moral discussions. However, in comparison to the other options, it presents the most feasible alternative until technology provides more effective options.

Requiring human operators and officers to be responsible, and to approve all decisions made by UCAVs, will appease most critics of current UCAV operations. The problem in the future, however, is when technology reaches a point where machines are able to make decisions on life and death, due to the fact that they simply have more information than the human supervisor. The advantage of having a fully-autonomous UCAV will not be fully realized if that capability cannot be fully exercised. There would be no point to developing the technology if a military did not intend to use it to its full potential. If resources were heavily invested in developing fully-autonomous UCAVs, there would be immense political and military pressures to use them as intended. The advantages of having autonomous UCAVs would be clear upon first use; quicker decision making, possible savings in human life, less supervision for automated

“Requiring human operators and officers to be responsible, and to approve all decisions made by UCAVs, will appease most critics of current UCAV operations.”

systems leading to less human labour requirements, and the ability to field more assets simultaneously. If autonomous UCAVs develop past a certain technological milestone, manned operators or supervisors could very well be either the weak-point, or simply serve as a disadvantage in conflict.¹⁵ If enemy states develop the same technology, the ability to have autonomous machines make split-second decisions would be of even more importance, and

the ‘human behind the screen’ would be even more of a liability. Requiring a communication link with a manned operator would also continue to serve as a weakness or challenge in future autonomous UCAV operations. The ideal situation would be to leave the UCAV to make its own decisions, regardless of whether a human is supervising it or not. This would negate the need for a constant satellite link, and it would shore up a known weakness or limitations of UCAVs.

To keep a human in a future unmanned system is also a weakness for other, less-visible reasons. While the psychological stressors on

operators of UCAVs and other unmanned systems has become a talking point with various users of the technology throughout the world, the main research on this subject has been in the USA. Studies conducted by the United States Air Force School of Aerospace Medicine (USAFSAM) have produced interesting results regarding the mental health of prolonged operations

be diminished if the person understands that they did not give the order to kill.¹⁷

It is unknown at this time whether taking the human completely out of UCAV operations is the best course of action for future autonomous flight. Filtering war down to reports of confirmed kills and data being read back from a UCAV seems like a dark moral and ethical path for humans to take, and it may make states more likely to wage war if they know that they would not have to experience it first hand, or risk the lives of their own soldiers on the front lines. One of the items that receives the most public outcry when a state is waging war or conflict is when the public sees sol-

diers returning home in coffins. This is a stark reminder that war is real, and it can have repercussions affecting all the citizens of a nation, not just those personally involved in war. Although it cannot be denied that having less soldiers die in conflict is a good thing, the risk of de-sensitization to conflict through the use of fully autonomous UCAVs and other machines is something that may come to fruition in the future. At this time, it is difficult to predict the ramifications of such technology, and it would be an interesting point of research going into the future.

Conclusion

This article has briefly discussed the genesis of UCAV autonomous technology: from where it exists today in the form of sense-and-avoid technology where manned supervision of basic duties is still required, but is in the process of being phased out, to a future that would include fully-autonomous UCAVs being given the responsibility and the authority to make decisions with respect to the lives of humans. The technology that will enable UCAVs that can think for themselves remains a future prospect, but it can be said with certainty that this technology will be developed eventually. The issue, then, will not be *when* the technology is developed, but *how* it is used, given the current laws of armed conflict and the rules of a just war. Who or what will be given the responsibility for decisions regarding life and death is still a very grey area, and several options have been laid out this article. None of them have provided a concrete solution to the problem as it is perceived today. If a more viable solution cannot be argued with confidence, then most likely it would be best to leave the system in status quo for the time being and keep the responsibility for killing with commanding officers and operators of nation's militaries around the world. This is a system with which



A remotely piloted aircraft crew flies a simulated training mission on an MQ-9 Reaper, Creech Air Force Base, 8 May 2014.

of unmanned systems in conflict areas. The most recent study conducted in 2014 from USAFSAM reported that 10.72% of operators self-reported experiencing high levels of mental distress, partially from the shift work, long hours, and low unit manning that results from the high operational tempo of UCAV squadrons.¹⁶ Mental issues such as this cost the US Government millions of dollars a year in support expenditures, as well as lost man-hours. Increased support for these personnel at risk is an effective measure. However, this appears to be a 'band aid solution.' A more effective solution would probably be to address the operational tempo required of UCAVs.

With respect to this last point, autonomous UCAVs would be able to address the operational tempo as well as the mental distress of UCAV operators because of the minimal supervision and personnel required to operate them, compared to present day aircraft. UCAVs that are able to operate autonomously without human supervision, to make decisions for themselves, and then report back automatically regarding the results of remote operations would significantly reduce the workload of current unmanned aircraft personnel. Fully autonomous UCAVs would also require less analysts and intelligence personnel tasked with observing sometimes gruesome deaths or dismemberments of enemy forces. A smaller number of personnel exposed to such images would equate to less risk of mental disorders, such as post-traumatic-stress-disorder (PTSD). Taking the human out of the decision to kill reduces stress in response to such traumatic events. Since stressors related to being either a witness or a participant in traumatic events varies greatly, the preference would be to involve personnel where the risk of mental stress is less. Acting as a *witness* to killing, but not being a *participant* is suggested to evoke less stress from members. This is not to suggest that watching war-like images will not have some form of effect upon a person, but the effect would

states are familiar, and it is one that works. Until autonomous UCAVs reach a point of intelligence where they are able to make equal-or-better decisions than the human operators, the proven system in place should be the one that the world utilizes. The more autonomous a system becomes, the less that one can reasonably argue that the person who designed it or the human in charge could be responsible for its actions. Either militaries continue to hold officers responsible for deaths that were not their decision, or they accept the fact that there may be unwanted deaths on the battlefield where the only one to blame is a machine that processed information incorrectly.

Increasing the use of autonomous UCAVs will also bear witness to an equal increase in moral or ethical concerns. The desensitization of conflict to military and public alike could possibly make countries more likely to wage war, or be willing to enter into conflict. This could present completely unknown problems as the world forges into a future where machines and not humans become the decider of who lives and who dies. Militaries will need to carefully consider these trade-offs when the time comes for fleets of fully autonomous UCAVs to take flight.



DVIDS/U.S. Air Force/Staff Sergeant Ramon A. Adelan/4909286



1st Lieutenant Murphy, a 7th Reconnaissance Squadron RQ-4 *Global Hawk* pilot, reviews pre-flight checklists, 24 October 2018, at Naval Air Station Sigonella, Italy.

NOTES

- 1 Captain Michael W. Byrnes, "Machine Autonomy in Air-to-Air Combat," 2014, p. 49.
- 2 Kimon P. Valavanis, and George J. Vachtsevanos, (eds.), *Handbook of Unmanned Aerial Vehicles*. Dordrecht, Netherlands: Springer, 2015, at: <https://doi.org/10.1007/978-90-481-9707-1>, p. 60.
- 3 "Global Hawk Unmanned Reconnaissance System Sets Aviation Record with Deployment to Australia." Northrop Grumman Newsroom. Accessed 15 January 2018 at: <https://news.northropgrumman.com/news/releases/global-hawk-unmanned-reconnaissance-system-sets-aviation-record-with-deployment-to-australia>.
- 4 Major Michael Kreuzer, "Nightfall and the Cloud," n.d., p. 59.
- 5 S. Ramasamy, R. Sabatini, and A. Gardi. "Cooperative and Non-Cooperative Sense-and-Avoid in the CNS+A Context: A Unified Methodology," in *2016 International Conference on Unmanned Aircraft Systems (ICUAS)*, pp. 531–539, 2016, at: <https://doi.org/10.1109/ICUAS.2016.7502676>, p. 532.
- 6 Kreuzer, p. 51.
- 7 *Ibid.*, p. 57.
- 8 Robert Sparrow, "Building a Better WarBot: Ethical Issues in the Design of Unmanned Systems for Military Applications." in *Science and Engineering Ethics* 15, No. 2 (June 2009), pp. 169–187 at: <https://doi.org/10.1007/s11948-008-9107-0>, p. 65.
- 9 Kreuzer, p. 65.
- 10 *Ibid.*, p. 64.
- 11 "Killer Robots." *Journal of Applied Philosophy* 24, No. 1 (February 2007), pp. 62–77 at: <https://doi.org/10.1111/j.1468-5930.2007.00346>, p. 69.
- 12 Sparrow, pp. 169–187, at: <https://doi.org/10.1007/s11948-008-9107-0>, p. 70.
- 13 *Ibid.*, p. 71.
- 14 *Ibid.*, p. 74.
- 15 *Ibid.*, p. 68.
- 16 Wayne L. Chappelle, Kent D. McDonald, Lillian Prince, Tanya Goodman, Bobbie N. Ray-Sannerud, and William Thompson. "Symptoms of Psychological Distress and Post-Traumatic Stress Disorder in United States Air Force 'Drone' Operators." in *Military Medicine* 179, No. 8S (August 2014), pp. 63–70, at: <https://doi.org/10.7205/MILMED-D-13-00501>, p. 67.
- 17 *Ibid.*, p. 10.