



Center for Strategic and Budgetary Assessments

IN THE NEWS

The Legal and Moral Problems of Autonomous Strike Aircraft

August 21, 2014 | *USNI News*

By: Dave Majumdar

Related Experts: Mark Gunzinger, Eric Edelman, Jim Thomas, Andrew F. Krepinevich, Jan van Tol

The U.S. Navy's move toward developing a carrier-based unmanned combat aircraft might eventually afford the service the ability to strike targets at long-range, but there are ethical and legal questions that linger should the Pentagon develop a fully autonomous system.

As currently envisioned, the Navy's Unmanned Carrier Launched Airborne Surveillance and Strike (UCLASS) will be autonomous, but it will have a "man-on-the-loop" according to Rear Adm. Mat Winter, the Naval Air Systems Command's Program Executive Officer for Unmanned Aviation and Strike Weapons. But the UCLASS is not going to be a penetrating strike aircraft such as many senior defense officials, academics and analysts had hoped. In fact, serious legal and ethical dilemmas might arise if the Pentagon were to pursue an unmanned penetrating strike aircraft.

There are many senior Pentagon officials—including Deputy Defense Secretary Bob Work—who argue for a deep penetrating unmanned strike aircraft that would launch from a carrier if the United States were serious about its "pivot" to the Pacific. But potential adversaries, such as Russia and China, are not stupid, and are certain to attack the vulnerable data-links that control such an unmanned bomber via electronic and cyber attacks.

One recently retired Navy official acknowledged that giving such a warplane full autonomy—to include launching weapons *without prior human consent*—might be the only effective way for a long-range unmanned strike aircraft to operate in a theater where the United States faces off against a near-peer potential adversary, but the prospect of such a system raises legal and moral questions.

Anti-Access Area Denial

The People's Liberation Army's DF-21D medium range ballistic missile, the so-called 'carrier killer.'

In the Western Pacific, China is building up its anti-access/area-denial (A2/AD) capabilities—including communications jamming, cyber-warfare and anti-satellite weapons. In the event of a conflict, Chinese forces are likely to attack those vital communications links than enable U.S. forces to operate cohesively. In those communications degraded/communications denied environments, unless a system is manned, autonomy might be the only way to go.

For the Navy there is an added dimension, as was postulated by Jan van Tol, Mark Gunzinger, Andrew Krepinevich and Jim Thomas at the Center for Strategic and Budgetary Assessments (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2250126): the service's aircraft carriers no longer have a haven in coastal waters 200 nautical miles offshore. With the rising threats to the aircraft carrier in the form of antiship cruise and ballistic missiles, those ships may be forced to stand off a significant distance—more than 1,000 nautical miles—from the enemy shoreline.

Additionally, with the proliferation of advanced integrated air-defense networks and low-frequency radars that can detect and track low-observable targets, existing stealth aircraft may not have the range or the survivability to operate in those theaters.

In that case, the best option for the Navy might be to develop a long-range unmanned strike aircraft with wide-band all-aspect stealth technology that would be able to persist inside even the densest of enemy air defenses. By necessity, given that such an advanced adversary would be able to deny or degrade communications significantly, such an aircraft would have to be fully autonomous. In other words, the unmanned aircraft would have to be able to operate independently of prolonged communications with its human masters and it would also need to be able to make the decision to release weapons without phoning home for a human operator's consent.

Moreover, the U.S. Air Force also faces basing challenges in the Western Pacific, as existing air bases such as Kadena and Misawa in Japan and Andersen Air Force Base in Guam are vulnerable to concerted air and missile attacks. A very stealthy long-range autonomous unmanned strike aircraft could be used to complement the service's prospective Long Range Strike Bomber—going into places that are far too dangerous for a manned aircraft or to perform missions like stand-in jamming from inside hostile territory.

The Autonomous Cost Equation

Contractors with Northrop Grumman responsible for operating the X-47B on the deck of USS Theodore Roosevelt (CVN-71) on Aug. 17, 2014. US Naval Institute Photo

While the initial cost of developing such an autonomous unmanned aircraft might be high, there might be significant savings longer-term. An autonomous unmanned aircraft only needs to be flown occasionally during peacetime to keep-up the proficiency of maintainers. Further, an autonomous aircraft has no need to fly training sorties or to practice—a computer can simply be programmed to do what needs to be done.

Additionally, such an autonomous unmanned aircraft would not need downtime between deployments—just the occasional depot-level maintenance overhaul. That means that the Navy—or the Air Force, if it bought some—would need only as many aircraft as required to fill the number of deployed carriers and account for attrition reserves and planes laid up in depot maintenance. There could also be significant personnel cost savings because a fully autonomous aircraft would not require pilots and the smaller fleet would require fewer maintainers.

The technology to develop and build such an aircraft mostly already exists. Most current unmanned aircraft like the General Atomics Aeronautical Systems MQ-1 Predator and MQ-9 Reaper are remotely controlled by a human operator. Others—like the Northrop Grumman MQ-4C Triton or RQ-4B Global Hawk—have far more autonomy but are not armed. Nonetheless, there are already a number of autonomous weapon systems that are either in service or that have reached the prototype stage that can engage hostile targets without human intervention.

Perhaps the two most obvious examples are cruise missiles and intercontinental ballistic missiles. Once those weapons are launched, they proceed autonomously to their preprogrammed targets without any human intervention.

If one were to imagine a U.S. Navy destroyer launching a Tomahawk cruise missile at a fixed target somewhere in the Western Pacific, there is a sequence of events that would be followed. The crew of the destroyer would receive orders to attack a particular target. The crew would then program that information into the missile. Once launched, the Tomahawk navigates its way to the target in a manner similar to a manned aircraft, but completely without human intervention.

Against a fixed target, for example a bunker or factory, a fully autonomous unmanned air vehicle would be very similar to a cruise missile. Like a Tomahawk cruise missile, the unmanned aerial vehicle (UAV) would receive a particular target location and instructions for how to engage that target with the correct weapons. Like the Tomahawk, the UAV would be able to navigate to that target completely autonomously. If the UAV were then to engage that fixed-target with a Joint Direct Attack Munition (JDAM) or some other weapon, in practical terms, there is no real difference between an unmanned aircraft and a cruise missile. The effect would be identical. The only change would be that the UAV could make a second pass, fly onto another target, or fly home to be rearmed. And it could be argued with its jet engine and wings, a Tomahawk is really just a small UAV on a one-way trip.

Expecting the Unexpected

Chinese anti-air missile system.

The more challenging scenario comes when there is an unexpected “pop-up” threat such as an S-400 surface-to-air missile battery that might be encountered by an autonomous unmanned combat air vehicle (UCAV) during a wartime sortie. Human pilots are assumed to inherently have the judgment to decide whether or not to engage such a threat. But those human pilots are making their decisions based on sensor information that is being processed by the aircraft’s computer. In fact, the pilot is often entirely dependent upon the aircraft’s sensors and the avionics to perform a combat identification of a contact.

The Lockheed Martin F-22 Raptor and F-35 Joint Strike Fighter epitomize this concept—especially in the realm of beyond visual range air-to-air combat. Both the Raptor and the F-35 fuse data correlated from the aircraft’s radar, electronic support measures and other sensors into a track file that the computer identifies as hostile, friendly or an unknown. The pilot is entirely reliant upon the computer to determine a proper combat identification; it would be a very small technological step for the system to engage targets autonomously without human intervention.

The air-to-ground arena is somewhat more challenging due to target location errors that are inherent in sensors and navigation systems (and also environmental effects and enemy camouflage). But with a combination of electro-optical/infrared cameras, synthetic aperture radar, ground moving target indication radar or even hyperspectral sensors, a computer can ascertain a positive combat identification of ground targets—assuming that the data being gathered is geo-registered. Once the computer can determine a positive identification—either a manned or unmanned aircraft—can engage a target. But at the end of the day, the computer is still making the determination that a contact is hostile.

In fact autonomous systems capable of identifying and attacking targets at their own discretion have existed in the past. One example is the Northrop AGM-136 Tacit Rainbow anti-radiation cruise missile, which was canceled in 1991. It was designed to be pre-programmed for a designated target area, over which it would loiter. It would remain in that designated box until it detected emissions from hostile radar. Once the Tacit Rainbow detected and identified an enemy emitter, the missile would zero in for the kill—all without any human intervention.

A later example is the Lockheed Martin Low-Cost Autonomous Attack System. The now-defunct miniature loitering cruise missile demonstrator was guided by GPS/INS to a target box. It would then use laser radar to illuminate targets and match them with pre-loaded signatures. The weapon would then go after the highest priority target while at the same time selecting the appropriate mode for the warhead to best engage the target autonomously without human intervention.

Other prominent examples include the Aegis combat system, which in its full automatic mode can engage multiple aircraft or missiles simultaneously without human intervention. Similarly, the shipboard Close-in Weapons System or Phalanx has an autonomous engagement capability.

Moral And Legal Questions

X-47B conducts its first night flight April 10 over Naval Air Station Patuxent River, Md. on April, 10 2014. US Navy Photo

What all of that means is that fully autonomous combat identification and engagement is technically feasible for unmanned aircraft—given sophisticated avionics and smart precision guided weapons. But while technically fully autonomous unmanned combat aircraft are feasible, what of the moral and legal implications?

The Pentagon has already preemptively issued policy guidance on the development and operational use of autonomous and semi-autonomous weapons in November 2012. DOD directive 3000.09 states: (<http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf>) “Autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.” But the policy does not expressly forbid the development of a fully autonomous lethal weapon systems, it merely states that senior Pentagon leadership would closely supervise any such development.

In order to prevent what the Defense Department calls an “unintended engagement”, those who authorize or direct the operation of autonomous and semi-autonomous weapon systems are required use “appropriate care and in accordance with the law of war, applicable treaties, weapon system safety rules and applicable rules of engagement,” the policy states.

Thus it would seem that the U.S. government views the use of autonomous weapon systems as legal under the laws of war—provided certain conditions are met. Indeed, a number of lawyers specializing in national security law have suggested that fully autonomous weapons are lawful. The responsibility for the use of such weapon would ultimately fall to the person who authorized its employment—which is similar to any other manned weapon.

But there are those who are adamantly opposed to any fully autonomous weapon systems—organizations such as Human Rights Watch (HRW). In a November 2012 report titled “Losing Humanity: The Case against Killer Robots,” (<http://www.hrw.org/reports/2012/11/19/losing-humanity-0>) HRW called for an international treaty that would preemptively ban all autonomous weapons. In fact it is likely that the DOD policy guidance on the development of autonomous weapons stems from the conclusions of the HRW report.

The HRW report makes three recommendations. The first: "Prohibit the development, production and use of fully autonomous weapons through an international legally binding instrument. The second: "Adopt national laws and policies to prohibit the development, production and use of fully autonomous weapons." The third: "Commence reviews of technologies and components that could lead to fully autonomous weapons. These reviews should take place at the very beginning of the development process and continue throughout the development and testing phases."

HRW asserts that autonomous systems are unable to meet the standards set forth under international humanitarian law. "The rules of distinction, proportionality and military necessity are especially important tools for protecting civilians from the effects of war, and fully autonomous weapons would not be able to abide by those rules," the report states.

But critics, such as legal scholar Benjamin Wittes at the Brookings Institution, have challenged such statements. Wittes has written (<http://www.lawfareblog.com/2012/12/does-human-rights-watch-prefer-disproportionate-and-indiscriminate-humans-to-discriminating-and-proportionate-robots/#.UunS7GRdXB8>) that there are situations where machines can "distinguish military targets far better and more accurately than humans can." Indeed, those familiar with unmanned technology, sensor hardware and software can attest that is indeed the case.

If a computer is given a certain set of parameters—for example a series of rules of engagement—it will follow those instructions precisely. If the autonomous weapon is designed and built to operate within the laws of war, then there should be no objection to their use. Under Article 36 of the 1977 Additional Protocol to the Geneva Conventions, weapons cannot be inherently indiscriminate and are prohibited from causing unnecessary suffering or superfluous injury. "The fact that an autonomous weapon system selects the target or undertakes the attack does not violate the rule," Hoover Institution legal scholars Kenneth Anderson and Matthew Waxman wrote in their paper "Law and Ethics For Autonomous Weapon Systems." (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2250126)

Technology is continuously moving forward and while autonomous systems may not be able to operate under all circumstances, it may only be a matter of time before engineers find a technical solution. While under many circumstances—with the right sensors and algorithms—autonomous systems would have the ability to distinguish lawful targets from unlawful targets, but that is not currently the case under all circumstances. Thus there are some limitations inherent to autonomous weapon systems for the time being.

However, those limitations will not always be there as technology continues its march forward and engineers continue to make progress. As Wittes correctly points out, “To call for a per se ban on autonomous weapons is to insist as a matter of IHL [international humanitarian law] on preserving a minimum level of human error in targeting.” Machines are generally far more precise than human beings.

Proportional Actions

B-52 bomber. US Air Force Photo

Along with being able to distinguish between targets, the law requires that combatants weigh the proportionality of their actions. “Any use of a weapon must also involve evaluation that sets the anticipated military advantage to be gained against the anticipated civilian harm,” Anderson and Waxman write. “The harm to civilians must not be excessive relative to the expected military gain.”

Though technically challenging, a completely autonomous weapon system would have to be required to address proportionality as well as distinction. But the difficulty is entirely dependent upon the specific operational scenario. For example, while an unmanned aircraft could identify and attack a hostile surface-to-air missile system deep behind enemy lines or an enemy warship at sea—where there is little chance of encountering civilians—targets inside a highly populated area are more difficult to prosecute.

Some of the most difficult scenarios—which would not necessarily be a factor in a high-end campaign against an A2/AD threats—would be challenging for a human pilot, let alone a machine. For example during counter-insurgency campaign, if there were two school buses driving side-by-side in a built-up area, but one of the vehicles was carrying nuns and the other carrying heavily-armed terrorists, it would be very difficult for a human pilot to determine which bus is the proper target until one of them commits a hostile act. The same would be true for an autonomous system—but in the near-term, it could be a technological challenge.

The human pilot would also have to determine what kind of weapon to use—judging the proportionality. Does he or she select a 2,000-pound JDAM or a smaller 250-pound small diameter bomb, or 20mm cannon, or do nothing since the risk of civilian casualties is too high? Likewise, an autonomous weapon system would need to be programmed to select an appropriate low collateral damage munition or to disengage if the danger of civilian casualties were too great once the target has been positively identified. But it would take time and investment before such an autonomous system could become a reality.

Baby Steps

Thus, for the near future, autonomous weapons would have to be developed incrementally starting with systems that could engage fixed targets and “obviously” military targets like surface-to-air missile sites or tank columns on the open battlefield during a conventional war. Likewise, in the maritime environment, where there are few civilians to speak of, autonomous systems could offer huge advantages with little in the way of any drawbacks.

Additionally, for the time being, autonomous weapons should not be utilized in complex scenarios—such as counter-insurgency—where there is significant possibility that it could cause inadvertent civilian casualties or unintended collateral damage. It may also be unwise to use a fully autonomous UCAV for missions like close air support—particularly during “danger close” type situations where friendly troops are in close contact with the enemy—until the technology has been proven operationally in other roles. Human pilots have a hard enough time with those types of missions.

While at present there are some technological limitations that do exist, those are not likely to remain to roadblocks forever. Autonomous technology is advancing rapidly and could one day be precise and reliable enough to not only distinguish correct targets but could also make proportionality judgments in complex scenarios based on parameters programmed into the machine. Those parameters would not be unlike rules of engagement given to human operators. Already, cameras and software exist that can identify individual human faces for example. Once a target is precisely identified, it would not be a huge leap then for an autonomous system to use a low-collateral damage weapons to eliminate hostile targets while minimizing any harm to civilians.

Much of the objection to fully autonomous weapons seems to stem from a sense of “fair play” rather than any real legal issues—most of which are likely to be overcome. But any time new weapons technology emerges, there is opposition from those who believe that the technology fundamentally unbalances war. Objections have been raised throughout history to new technologies—ranging from crossbows and longbows, to machine-guns and submarines—because the use of such weapons was considered to be “unfair” or “unsporting.” But ultimately, the use of such weapons became a fact of life. War is not a game, and as U.S. Air Force Col. Lawrence Spinetta, commander of the 69th Reconnaissance Group said: “Isn’t there a moral imperative on the part of a nation to minimize danger for its soldiers and airmen?”

Indeed there is no legal requirement for war to be fair—in fact throughout history war has been anything but. “The law, to be sure, makes no requirement that sides limit themselves to the weapons available to the other side; weapons superiority is perfectly lawful and indeed assumed as part of military necessity,” write Anderson and Waxman.