7 Killing Made Easy: From Joysticks to Politics

Noel Sharkey

To fight from a distance is instinctive in man. From the first day he has worked to this end, and he continues to do so.
—Ardant du Picq[1]

Robots will change the way that wars are fought by providing distant "stand-ins" for combatants. Military robots are the fruit of a long chain of weapons development designed to separate fighters from their foes. Throughout the history of war, weapon technology has evolved to enable killing from ever- increasing distances. From stones to pole weapons to bows and arrows to cannon to aerial bombing to jet propelled missiles, killing has become ever easier.

Not only have distance weapons led to a more effective killing technology, but attacking from a distance also gets around two of the fundamental obstacles that warfighters must face: fear of being killed and resistance to killing. Fear is one of the greatest obstacles to a soldier's effectiveness in battle (Daddis 2004). It is obvious that the greater the distance from the enemy, the less fear will play in the action. Many battles throughout history have been lost by men bolting in panic as fear swept through the ranks—often from a misunderstanding of the action (Holmes 2003).

Army historian Brigadier General Marshall ([1947] 2000), following after-action interviews with soldiers in the Pacific and European theaters of operation during World War II, claimed that only about 15 to 20 percent of riflemen were either able to or willing to fire. This means that around 80 percent of the U.S. infantry in World War II either were not firing their weapons when they could see their enemy, or were firing over enemy soldiers' heads. There have been some very sharp criticisms of Marshall's research methods, and the exact percentages may not be correct, but the nature of his findings—that many soldiers are unwilling to kill—has received general support from other analyses of historical battles.

In his book *Acts of War*, Holmes (2003) argues that the *hit rates* in a number of historical battles show that many soldiers were not prepared to fire directly at the enemy when they were in sight. A group of British soldiers entirely surrounded by

Zulu warriors fired at point-blank range, but had a hit rate of only one to every thirteen rounds fired. At the battle of Wissembourg in 1870, the French fired 48,000 rounds at the Germans advancing across open fields, but only managed to hit 404 of them. In the Vietnam War, it was estimated that over 50,000 bullets were fired for every soldier killed. Holmes also tells the World War I story of Lieutenant George Roupell, who, to stop his men firing in the air, patrolled the trenches, hitting them on the backsides with his sword, telling them to fire low.

The killologist, Lieutenant Colonel David Grossman, argues that "not firing" is not cowardice, but really a compulsion of ordinary men not to kill (Grossman 1995). He gives several examples in his book, *On Killing*, from the U.S. Civil War of low killing rates from close-distance musket fire. In one instance, the Battle of Gettysburg, of 27,574 muskets retrieved from the battlefield, 90 percent were still loaded or multiply loaded—one musket had even been loaded twenty-three times without being fired.

Grossman also points out that killing distance can be psychological as well as physical. He cites Clausewitz and du Picq for expounding at length on how the vast majority of deaths in battle occurred when the victors chased the losing side in retreat. Du Picq suggests that Alexander the Great lost fewer than seven hundred men over all his battles because there never was a victorious enemy to pursue his army—and so his soldiers never retreated. Grossman argues that across the battlefields of Europe and in the U.S. Civil War, the majority of casualties and deaths were inflicted by artillery. In his view, the greater the distance the artillery is from its targets, the greater its effectiveness will be. We see the same phenomena with increasingly high-altitude aerial bombing and the use of long-range missiles.

Now we are embarking on new territory, where the new battlefield robots should not be considered as distance weapons in the traditional sense. Yes, a cruise missile can be considered to be a robot, for after it is launched it can alter its course with built-in GPS. But it has a single purpose—to strike and destroy a target. The new battlefield robots are different. They can stand in directly for soldiers or pilots at greater and greater distances. These robots are coming into their own as a new form of automated killing machine that may forever alter how war is waged. Unlike missiles or other projectiles, robots can carry multiweapon systems into the theater of operations. How they are to be deployed in the theater need not be decided in advance, as they can act flexibly once in place. Eventually, they may be able to take the place of human combatants without risk to the lives of their operators. Killing will become so much easier—but not without moral risk.

7.1 The Ultimate Distance Weapon Systems

Nowadays, so many robots are being deployed in the Middle East conflict zones that it is difficult to get an accurate estimate of their numbers. The figures for ground robots range from 6,000 to 12,000. Even the lower figure shows the dramatic increase in the use of robots since 2004, when there were only 150, and it testifies to their military usefulness. The robots have mainly been deployed for dull, dirty, and dangerous tasks, such as disrupting or detonating improvised explosive devices and for surveillance of dangerous environments, such as caves and buildings that may be housing insurgents. Roadside bombs are the most common killer of allied soldiers, and robots are used to drive ahead and search cars or prod suspected packages. Robots have saved many soldiers' lives.

The first blood drawn by a ground robot was actually by the small and relatively cheap four-wheeled MARCbot, which looks like a toy truck with a camera stalk (Singer 2009). Its main purpose was to inspect underneath cars and trucks for explosives. But one U.S. unit had a clever idea. Its soldiers started loading MARCbotswith Claymore antipersonnel mines and went looking for insurgents hiding in alleyways to ambush them. When they found any, they killed them by exploding the mine. But this was an unofficial use of the robot and it took time to surmount some of the legal and physical difficulties of using special-purpose armed ground robots. Nonetheless, if there is an opportunity to use armed robots to separate soldiers from danger, commanders are likely to use them.

In June 2007, the first three armed Talon SWORDS (Special Weapons Observation Reconnaissance Detection System) were sent to Iraq at a reported cost of $200,000 each. These can be equipped with M240 or M249 machine guns, Barrett 0.50 caliber rifles, 40mm grenade launchers, or antitank rocket launchers. As far as it is possible to tell, they were not deployed in action. One explanation given by Kevin Fahey (the U.S. Army's executive officer for ground forces) was that when the SWORDS was first switched on, the gun had begun to move when it should not have moved (Sofge 2008). Another explanation, given to the *Defense Review* journal by U.S. Special Forces, is that SWORDS is jokingly referred to as the TVR, or Taliban Re-supply Vehicle, because "Taliban fighters will hide and wait for the weaponized Talon robot/SWORDS to roll by, sneak up on it, tip it over, remove the machine gun (or any other weapon) and ammo from it, and then use it/them against U.S. forces" (Crane 2008).

The SWORDS was essentially a test of concept to try the robots with soldiers on the battlefield. It has influenced the development of the next generation of armed ground robots, which is well under way. More powerfully armed robots, such as the tank-like MAARS (Modular Advanced Armed Robotic System) from Foster-Miller, are to replace the SWORDS.

But it is the robot planes and drones that are currently the ultimate in distance weapons systems. Missions are flown by "pilots" of the 432nd Air Expeditionary Wing at the Creech Air Force base in the Nevada desert, thousands of miles away from the operations. The operators sit at game consoles, making decisions about when to apply lethal force. Sometimes, all the operator has to do is to decide (in a very short space of time) whether or not to veto the application of force. The planes can be flown around the clock, as it is easy for pilots to take a break from "battle" at any time, or even go home to have dinner with their children. According to some, the sharp contrast between home life and the battlefield within the same twenty-four-hour period is apparently causing a new kind of battle stress that has not been witnessed before.

The Unmanned Combat Air Vehicle (UCAV), the MQ-1Predator, which carries a payload of two Hellfire missiles, flew 250,000 hours up until June 2007. As a mark of its military usefulness, it clocked an additional 150,000 hours in the Afghan and Iraqi conflicts in the subsequent fourteen months, and passed the one-million flight hours mark in 2010.

In October 2007, the Predator was joined by the much larger and more powerful MQ-9 Reaper. The MQ-9 Reaper carries a payload of up to fourteen Hellfire missiles, or a mixture of missiles and bombs. These "hunter-killer" unmanned aerial vehicles (UAVs) have conducted many decapitation strikes[2] since they were first deployed in Afghanistan in October 2007. There is a demand to get many more operational as soon as possible. The number of Reapers flying over the conflict zones has doubled to twenty during their first year of operation (2007–2008)—a year ahead of schedule—and there has been a push from the U.S. Air Force (USAF) for General Atomics to increase production levels above the current four per month. In late 2008, $412 million was added to the USAF budget for training more nonaerial pilots.

There was no change of direction under the Obama administration. Although there were cutbacks to conventional weapons, the robot programs received more cash than predicted. In 2010, the Air Force aimed to spend $2.13 billion on unmanned technology, with $489.24 million to procure twenty-four new heavily armed Reapers. The U.S. Army planned to spend $2.13 billion on unmanned vehicle technology. This includes the purchase of thirty-six more unmanned Predators. The U.S. Navy and Marine Corp targeted $1.05 billion for unmanned vehicles, including armed MQ-8B helicopters.

Outside of these conventional forces, there is a considerable Central Intelligence Agency (CIA) use of the drones for decapitation strikes. Indeed, it was the CIA that carried out the first missile strike from an armed Predator in Yemen in 2002. The CIA has now effectively got its own air force flying over Somalia, Yemen, Afghanistan, and Pakistan. The legality of such attacks was questioned at the UN General Assembly meeting in October 2009 by Philip Alston, UN special reporter on extrajudicial killings. He made a request for U.S. legal justification for how the CIA is accountable for the targets that they are killing. The United States turned down the request, stating that these are covert operations.

A rebuttal by Harold Koh, legal adviser, U.S. Department of State, insisted, "US targeting practices, including lethal operations conducted by UAVs, comply with all applicable law, including the laws of war" (Koh 2010). However, there are no independent means of determining how the targeting decisions are being made. A commander of a force belonging to a state acting against the United States would be a legitimate target. Intelligence errors made in the Vietnam War and its aftermath about the standard of evidence used for assassinations led to Presidential Order 12333, prohibiting the assassination of civilians. And it is now unclear what type and level of evidence is being used to sentence nonstate actors to death by Hellfire attack without right to appeal or right to surrender. It sits behind the cloak of national secrecy. A subsequent report by Alston (2010) to the UN General Assembly[3] discusses drone strikes as violating international and human rights laws because both require transparency about the procedures and safeguards in place to ensure that killings are lawful and justified: "a lack of disclosure gives states a virtual and impermissible license to kill." The debate continues.

All of the armed drones are currently "man in the loop" combat systems. This makes very little difference to the collaterally damaged villagers in Waziristan, where there have been repeated Predator strikes since 2006. No one knows the true figures for civilian casualties, but according to reports coming from the Pakistan press, drone attacks have killed fourteen al-Qaeda leaders, and this may have been at the cost of over six hundred civilians (Sharkey 2009b).

7.2 In, On, or Out of the Loop

There is now massive spending going on, and plans are well under way to take the human "out of the loop," so that robots can operate autonomously to locate their own targets and destroy them without human intervention (Sharkey 2008a). This is high on the military agenda of all the U.S. forces: "the Navy and Marine Corps should aggressively exploit the considerable war-fighting benefits offered by autonomous vehicles (AVs) by acquiring operational experience with current systems, and using lessons learned from that experience to develop future AV technologies, operational requirements, and systems concepts" (Committee on Autonomous Vehicles in Support of Naval Operations National Research Council 2005). There are now a number of autonomous ground vehicles, such as DARPA's "Unmanned Ground Combat Vehicle and Perceptor Integration System," otherwise known as the Crusher (Fox News 2008). BAE systems recently reported in an industry briefing to United Press International (2008) that they have "completed a flying trial which, for the first time, demonstrated the coordinated control of multiple UAVs autonomously completing a series of tasks."

The move to autonomy is clearly required to fulfill the current U.S. military plans. Teleoperated systems are more expensive to manufacture and require many support personnel to run them. One of the main goals is to use robots as force multipliers, so that one soldier on the battlefield can be a nexus for initiating a large-scale robot attack from the ground and the air. Clearly, one soldier cannot remotely operate several robots alone.

In the U.S. Air Force's *Unmanned Aircraft Systems Flight Plan 2009–2047*, autonomy was also discussed for swarm technologies: "SWARM technology will allow multiple MQ-Mb aircraft to cooperate in a variety of lethal and nonlethal missions at the command of a single pilot" (United States Air Force 2009, 39). Such a move will require decisions being made by the swarm—human decision making will be too slow and not able to react to the control of several aircraft at once.

There is also a considerable push to shrink the role of "the man in the loop." To begin with, autonomous operation will be mainly for tasks such as take-off, landing, and refueling. As unmanned drones react in micro- or nano-seconds, the "humans will no longer be 'in the loop' but rather 'on the loop,' monitoring the execution of certain decisions. Simultaneously, advances in AI will enable systems to make combat decisions and act within legal and policy constraints, without necessarily requiring human input" (United States Air Force 2009, 41).

The main ethical problems arise because no autonomous robots or artificial intelligence systems have the necessary sensing properties to allow for discrimination between combatants and innocents. This is also understood clearly by some within the military. Major Daniel Davis, a combat veteran of Iraq 1991 and Afghanistan 2005, writes: "Suggesting that within the next 12-plus years technology could exist that would permit life-and-death decisions to be made by algorithms is delusional. A machine cannot sense something is wrong and take action when no orders have been given. It doesn't have intuition. It cannot operate within the commander's intent and use initiative outside its programming. It doesn't have compassion and cannot extend mercy" (2007).

Davis quotes Colonel Lee Fetterman, training and doctrine capabilities manager for Future Combat Systems FCS, who has a high regard for the unmanned PackBot that he used in Afghanistan to search caves and buildings. However, he has strong opinions about robots making decisions about killing. "The function that robots cannot perform for us—that is, the function we should not allow them to perform for us—is the decide function. Men should decide to kill other men, not machines," he said (Davis 2007). "This is a moral imperative that we ignore at great peril to our humanity. We would be morally bereft if we abrogate our responsibility to make the life-and-death decisions required on a battlefield as leaders and soldiers with human compassion and understanding. This is not something we would do. It is not in concert with the American spirit" (Davis 2007).

Allowing robots to make decisions about who to kill could fall foul of the fundamental ethical precepts of a just war under *jus in bello*, as enshrined in the Geneva and Hague conventions and the various protocols set up to protect the innocent: only combatants/warriors are legitimate targets of attack—all others, including children, civilians, service workers, and retirees, should be immune from attack. In fact, the laws of protection even extend to combatants that are wounded, have surrendered, or are mentally ill (but see also Ford 1944).

These protections have been in place for many centuries. Thomas Aquinas, in the thirteenth century, developed the "doctrine of double effect." Essentially, there is no moral penalty for killing innocents during a conflict provided that (1) you did not intend to do so, or (2) killing the innocents was not a means to winning, or (3) the importance to the defense of your nation is proportionally greater than the number of civilian deaths.

There are many circumstances in a modern war where it is extremely difficult, if not impossible, to fully protect noncombatants. For example, in attacking a warship, some noncombatants, such as chaplains and medical staff, may be unavoidably killed. Similarly, but less ethically justifiable, it is difficult to protect the innocent when large explosives are used near civilian populations, or when missiles get misdirected. In modern warfare, the equivalent of the doctrine of double effect is the principle of proportionality, which "requires that the anticipated loss of life and damage to property incidental to attacks must not be excessive in relation to the concrete and direct military advantage expected to be gained" (Petraeus and Amos 2006).

In the heat of battle, both the principles of discrimination and proportionality can be problematic, although their violation requires accountability and can lead to war crimes tribunals. But the new robot weapons, which could violate both of these principles, cannot be held accountable for their decisions (Sharkey 2008b). You cannot punish an inanimate object. It would be very difficult to allocate responsibility in the chain of command or to manufacturers, programmers, or designers—and being able to allocate responsibility is essential to the laws of war.

The problem is exacerbated further by not having a specification of "civilianness" (see Roberts, forthcoming, for the difficulties in trying to find a definition of a civilian). A computer can compute any given procedure that can be written down in a programming language. We could, for example, give the computer on a robot an instruction such as, "if civilian, do not shoot." This would be fine, if and only if there was some way to give the computer a precise definition of "civilian." We certainly cannot get one from the laws of war that could provide a machine with the necessary information. The 1949 Geneva Convention requires the use of common sense, while the 1977 Protocol 1 essentially defines a "civilian," in the negative sense, as someone who is not a combatant:

1. A civilian is any person who does not belong to one of the categories of persons referred to in Article 4 A (1), (2), (3), and (6) of the Third Convention and in Article 43 of this Protocol. In case of doubt whether a person is a civilian, that person shall be considered to be a civilian.
2. The civilian population comprises all persons who are civilians.

3. The presence within the civilian population of individuals who do not come within the definition of civilians does not deprive the population of its civilian character. (Protocol 1 Additional to the Geneva Conventions, 1977 [Article 50])

And even if there were a clear computational definition of civilian, we would still need all of the relevant information to be made available from the sensing apparatus. All that is available to robots are sensors, such as cameras, infrared sensors, sonar, lasers, temperature sensors, ladars, and so on. These may be able to tell us whether something is a human or at least an animal, but not much else. In the labs there are systems that can identify someone's facial expression or that can recognize faces, but they do not work well on real-time moving people. And even if they did, how useful could they be in the fog of war? British teenagers beat the surveillance cameras just by wearing hooded jackets.

In a conventional war where all of the enemy combatants wear clearly marked uniforms (or better yet, radio frequency tags), the problems might not be much different from those faced in conventional methods of bombardment. But, asymmetrical warfare is increasingly making battle with insurgents the norm, and, in these cases, sensors would not help in discrimination. Knowing whom to kill would have to be based on situational awareness and having a theory of mind, that is, understanding someone else's intentions and predicting their likely behavior in a particular situation. Humans understand one another in a way that machines cannot. Cues can be very subtle, and there are an infinite number of circumstances where lethal force is inappropriate. Just think of children being forced to carry empty rifles, or insurgents burying their dead.

7.3 An Ethical Code for Robots?

The military does consider the ethical implications of civilian deaths from autonomous robots, although this is not their primary concern. Their role is to protect their country in whatever way is required. In the United States, all weapons and weapons systems are subjected to a legal review to ensure compliance with the Law of Armed Conflict (LOAC). There are three main questions to be asked before a weapon is authorized:

1. Does the weapon cause suffering that is needless, superfluous, or disproportionate to the military advantage reasonably expected from the use of the weapon? It cannot be declared unlawful merely because it may cause severe suffering or injury.
2. Is the weapon capable of being controlled, so as to be directed against a lawful target?
3. Is there a specific treaty provision or domestic law prohibiting the weapon's acquisition or use?

Regardless of these rules, we have already seen a considerable number of collateral casualties resulting from the use of semi-autonomous weapon systems. The argument then is one of proportionality, as stated in the first question, but there is no quantitative measure that can objectively determine military costs against civilian deaths. It is just a matter of political argument, as we have seen, time and time again.

Another concern is the question of what constitutes a new weapon. Take the case of the Predator UCAV. It was first passed for surveillance missions. Then, when it was armed with Hellfire missiles, the Judge Advocate General's office said that because both Predators and Hellfires had previously been passed, their combination did not need to be (Canning et al. 2004). Thus, if we have a previously used autonomous robot and a previously used weapon, it may be possible to combine them without further permission.

Armed autonomous robots could also be treated in a legally similar way to submunitions, such as the BLU-108 developed by Textron Defense Systems.[3] The BLU-108 parachutes to near the ground, where an altitude sensor triggers a rocket that spins it upward. It then releases four Skeet warheads at right angles to one another. Each has a dual-mode (active and passive) sensor system: the passive infrared sensor detects hot targets, such as vehicles, while the active laser sensor provides target profiling. They can hit hard targets with penetrators, or destroy soft targets by fragmentation.

The BLU-108 is not like other bombs because it has a method of target discrimination. If it had been developed in the 1940s or 1950s, there is no doubt that it would have been classified as a robot, and even now it is debatably a form of robot. The Skeet warheads have autonomous operation and use sensors to target their weapons. The sensors provide discrimination between hot and cold bodies of a certain height, but like autonomous robots, they cannot discriminate between legitimate targets and civilians. If BLU-108s were dropped on a civilian area, they would destroy buses, cars, and trolleys. Like conventional bombs, discrimination between innocents and combatants requires accurate human targeting judgments. A key feature of the BLU-108 is that it has built-in redundant self-destruct logic modes that largely leave battlefields clean of unexploded warheads, and it is this that keeps it out of the 2008 international treaty banning cluster munitions (Convention of Cluster Munitions).

To use robot technology over the next twenty-five years in warfare would, at best, be like using the BLU-108 submunition, in other words, it can sense a target, but cannot discriminate innocent from combatant (Sharkey 2008c). The big difference with the types of autonomous robots currently being planned and developed for aerial and ground warfare is that they are not perimeter-limited. The BLU-108 has a footprint of 820 feet all around. By way of contrast, mobile autonomous robots are limited only by the amount of fuel or battery power they can carry. They can potentially travel long distances and move out of line of sight communication.

In a recent sign of these future weapons, the U.S. Air Force sent out a call for proposals for "Guided, Smart Sub-munitions": "This concept requires a CBU (Cluster Bomb Unit) munition, or UAV capable of deploying guided smart sub-munitions, that has the ability to engage and neutralize any targets of interest. The goal of the sub-munitions is very challenging when considering the mission of addressing mobile and fixed targets of interest. The sub-munition has to be able to reacquire the target of interest it is intended to engage" (United States Air Force 2008). This could be very much like an extended version of the BLU-108 that could pursue hot-bodied targets. Most worrying are the words "reacquire the target of interest." If a targeted truck were, for example, to overtake a school bus, the weapons might acquire the bus as the target rather than the truck.

A naval presentation by Chief Engineer J. S. Canning subtitled "The Difference between 'Winning the War' and 'Winning the Peace'" discusses a number of the ethical issues involved in the deployment of autonomous weapons. The critical issue for Canning is that armed autonomous systems should have the ability to identify the legality of a target. His answer to the ethical problems is unnervingly simple: "let men target men" and "let machines target other machines" (Canning 2006). This restricts the target set, and, Canning believes, may overcome the political objections and legal ramifications of using autonomous weapons.

While machines targeting machines sounds like a great ethical solution on the drawing table, the reality is that it belongs to mythical artificial intelligence, not real-world AI. In most circumstances, it would not be possible to pinpoint the weapon without also pinpointing the person using it, or even to discriminate between weapons and nonweapons. I have the mental image of a little girl being blown away because she points her ice cream at a robot to see if it would like some. And what if the enemy tricks the robot into killing innocent civilians by, for example, placing weapons on a school or hospital roof? Who will take the responsibility?

A different approach, suggested by Ronald Arkin from the Georgia Institute of Technology, is to equip the robotic soldier with an *artificial conscience* (Arkin and Moshkina 2007). Arkin had funding from the U.S. Army to work on a method for designing an ethical autonomous robot, which he refers to as a humane-oid.[4]At first glance, this sounds like a move in the right direction. At the very least, it gets the army to consider the ethical problems raised both by the deployment of autonomous machines and even those of the soldier on the ground. Another of Arkin's concerns that he addresses in a public survey, and it is a good one, is "to establish what is acceptable to the public and other groups, regarding the use of lethal autonomous systems" (Arkin and Moshkina 2007).

Despite the good intentions, I have grave doubts about the outcome of this project. No idea is presented about how this could be made to work reliably, and reliability is a key issue when it comes to human lives. It is not just about having incredibly good sensors and camera inputs, or being able to make appropriate discriminations. A robot could actually have to make decisions in very complex circumstances that are entirely unpredictable.

It turns out that the plan for this conscience is to create a mathematical decision space consisting of constraints, represented as prohibitions and obligations derived from the laws of war and rules of engagement (Arkin 2009). Essentially, this consists of a bunch of complex conditionals (if–then statements). Reporting on Arkin's work, *The Economist* (2007) gives the example of a Predator UAV on its way to kill a car full of terrorists. If it sees the car overtaking a bus full of school children, it will wait until it has overtaken them before blasting the car into oblivion. But how will the robot discriminate between a bus full of school children and a bus full of guards? Admittedly, this is not one of the tasks that Arkin cites, but it is still the kind of ethical decision that an autonomous robot would have to make. The shadow of mythical AI looms large in the background.

Arkin believes that a robot could be more ethical than a human because its ethics are strictly programmed into it, and it has no emotional involvement with the action. The justification for this comes from a worrying survey, published by the Office of the Surgeon General (Mental Health Advisory Team 2006) that tells of the aberrant ethical behavior and attitudes of many U.S. soldiers and marines serving in Iraq. Arkin holds that a robot cannot feel anger or a desire for revenge, but neither can it feel sympathy, empathy, or remorse. Surely, a better way to spend the money would be on more thorough ethical training and monitoring of the troops.

Even if a robot was fully equipped with all of the rules from the Laws of War, and had, by some mysterious means, a way of making the same discriminations as humans make, it could not be ethical in the same way as is an ethical human. Ask any judge what they think about blindly following rules and laws. In most real-world situations, these are a matter of interpretation.

Arkin's anthropomorphism in saying, for example, that robots would be more humane than humans does not serve his cause well. To be humane is, by definition, to be characterized by kindness, mercy, and sympathy, or to be marked by an emphasis on humanistic values and concerns. These are all human attributes that are not appropriate in a discussion of software for controlling mechanical devices. More recently, Arkin has taken to talking about adding sympathy and guilt to robots. However, the real value of the work would be to add safety constraints to autonomous weaponized robots to help to cut down the number of civilian casualties. This is easy to understand, and may help the work to progress in a clearer way. The anthropomorphic terms create a more interesting narrative, but they only confuse the important safety issues and create false expectations.

The number of possible moral and ethical problems in a military operations theater full of civilians could be infinite, or at least run into extremely large numbers. Many different circumstances can happen simultaneously and give rise to unpredictable or chaotic robot behavior. From a perhaps cynical perspective, the "robot soldier with a conscience" could at some point be used by military public relations to allay political opposition, amounting to lots of talk while innocent civilians keep on dying: "Don't worry, we'll figure out how to use the technology discriminately eventually."

As Davis says about other defense experts talking up robot warfare, "such statements are dangerous, because men disconnected from the realities of warfare may sway decision-makers regarding future force decisions and composition" (Davis 2008). On the same basis, the "artificial conscience" idea could perhaps also be employed as an argument to shift the burden of responsibility for collateral fatalities from the chain of command onto inanimate weapons.

No civilized person wishes to see their country's young soldiers die in foreign wars. The robot is certainly a great defensive weapon, especially when it comes to roadside bombs. It is the moral responsibility of military commanders to protect their soldiers, but there are a number of far-reaching consequences of "risk-free" war that we need to consider.

- Having more robots to reduce the "body bag count" could mean fewer disincentives to start wars. In the United States, since the Vietnam War, body-bag politics has been a major inhibitor of military action. Without bodies coming home, citizens will care a lot less about action abroad, except in terms of the expense to the taxpayer. It could mean, for example, that with greatly reduced public and political opposition (passing the so-called Dover[5]), it is a lot easier for the military to start and run more "defensive" wars. This is an ethical and moral dilemma that should be engaging international thinking.
- Armstrong warns about the use of robots in "the last three feet" and asks if the United States really wants to have a robot represent the nation as a strategic corporal. You can't hope to win hearts and minds by sticking armed robots in the face of an occupied population (Armstrong 2007).
- It has been suggested that a country engaged in risk-free war will put its civilian population more at risk from terrorist attacks at home and abroad (Kahn 2002).
- It is more like policing—a term used for the Kosovo war—but policing requires a different set of rules than war; for example collateral civilian deaths are unacceptable for policing. Those suffering from policing need to be demonstrably morally guilty (Kahn 2002).
- There will clearly be proliferation (the indications are already there), and so the risk-free state could be short lived. As Chief Engineer Canning has pointed out: "What happens when another country sees what we've been doing, realizes it's not that hard, and begins to pursue it, too, but doesn't have the same moral structure we do? You will see a number of countries around the world begin to develop this technology on

their own, but possibly without the same level of safeguards that we might build-in. We soon could be facing our own distorted image on the battlefield" (Canning 2005).

A related concern is that when we say robot weapons save lives, we implicitly mean only the lives of *our* soldiers and their allies. Of course, in the middle of a vicious war, that is what we want. But let us not forget that such sentiments allow us to hide from ourselves the fact that the robot weapons could take a disproportionate toll of lives on the other side, including many innocent civilians. Autonomy could greatly increase fatal errors.

7.4 The Problem of Proportionality

According to the laws of war, a robot could potentially be allowed to make lethal errors, providing that the noncombatant casualties were proportional to the military advantage gained. But how is a robot supposed to calculate what is a proportionate response? There is no sensing or computational capability that would allow a robot such a determination. As mentioned for the discrimination problem described earlier, computer systems need clear specifications in order to operate effectively. There is no known metric to objectively measure needless, superfluous, or disproportionate suffering.[6] It requires human judgment.

No clear objective means are given in any of the laws of war for how to calculate what is proportionate (Sharkey 2009a). The phrase "excessive in relation to the concrete and direct military advantage expected to be gained" is not a specification. How can such values be assigned, and how can such calculations be made? What could the metric be for assigning value to killing an insurgent, relative to the value of noncombatants, particularly children, who could not be accused of willingly contributing to insurgency activity? The military says that it is one of the most difficult decisions that a commander has to make, but that acknowledgment does not answer the question of what metrics should be applied. It is left to a military force to argue as to whether or not it has made a proportionate response, as has been evidenced in the recent Israeli–Gaza conflict (Human Rights Watch 2009).

Uncertainty needs to be a factor in any proportionality calculus. Is the intelligence correct, and is there really a genuine target in the kill zone? The target value must be weighted by a probability of presence/absence. This is an impossible calculation unless the target is visually identified at the onset of the attack. Even then, errors can be made. The investigative journalist Seymour Hersh gives the example of a man in Afghanistan being mistaken for bin Laden by CIA Predator operators. A Hellfire was launched, killing three people who were later reported to be local men scavenging in the woods for scrap metal (Hersh 2002, 66). This error was made using a robot plane with a human in the loop. There is also the problem of relying on informants. The reliability of the informant needs to be taken into account, and so does the reliability of each link in the chain of information reaching the informant before being passed onto the commander/operator/pilot. There can be deliberate deception anywhere along the information chain, as was revealed in investigations of Operation Phoenix—the U.S. assassination program—after the Vietnam War. As Hersh pointed out, many of the thousands on the assassination list had been put there by South Vietnamese officials for personal reasons, such as erasing gambling debts or resolving family quarrels.

It is also often practically impossible to calculate a value for actual military advantage. This is not necessarily the same as the political advantage of creating a sense of military success by putting a face to the enemy to rally public support at home and to boost the morale of the troops. Obviously there are gross calculations that work in the extreme, such as a military force carrying weapons sufficient to kill the population of a large city. Then, it could be possible to balance the number of civilians killed against the number saved. Military advantage, at best, results in *deterrence* of the enemy from acting in a particular way, *disruption* of the social, political, economic, or military functions (or a combination of these), and *destruction* of the social, political, economic, or military functions (or a combination) (Hyder 2004, 5). Proportionality calculations should be based on the likely differences in military outcome if the military action killing innocents had not been taken (Chakwin, Voelkel, and Scott 2002).

Despite the impossibility of proportionality calculations, military commanders at war have a political mandate to make such decisions on an almost daily basis. Commanders have to weigh the circumstances before making a decision, but ultimately it will be a subjective metric. Clearly the extremes of wiping out a whole city to eliminate even the highest-value target, say Osama bin Laden, is out of the question. So there must be some subjective estimates about just how many innocent people killed equal the military value of the successful completion of a given mission.

Yes, humans do make errors and can behave unethically, but they can also be held accountable. Who is to be held responsible for the lethal mishaps of a robot? Robert Sparrow argues that it certainly cannot be the machine itself, and thus it is not legitimate to use automated killing machines (Sparrow 2007). There is no way to punish a robot. We could just switch it off, but it would not care any more about that than my washing machine would care. Imagine telling your washing machine that if it does not remove stains properly, you will break its door off. Would you expect that to have any impact on its behavior? There is a long causal chain associated with robots: the manufacturer, the programmer, the designer, the Department of Defense, the generals or admirals in charge of the operation, and the operator. It is thus difficult to allocate responsibility for deliberate war crimes, or even mishaps.

7.5 Conclusion

We discussed at the outset how killing is made easier for combatants when the distance between them and their enemies is increased. Soldiers throughout history have found it difficult to kill at close range when they can clearly see whom they are killing. Distance, whether physical or psychological, helps to overcome the twin problems of fear of being killed and resistance to killing that particularly dog the infantry.

Robots are set to change the way that wars are fought by providing flexible "stand-ins" for combatants. They provide the ultimate distance targeting that allows warriors to do their killing from the comfort of an armchair in their home country—even thousands of miles away from the action. Robots are developing as a new kind of fighting method different from what has come before. Unlike missile or other projectiles, robots can carry multiweapon systems into the theater of operations, and act flexibly once in place. Eventually, they may be able to operate as flexibly as human combatants, without risk to the lives of their operators that control them. However, as we discussed, there is no such thing as risk-free warfare. Apart from the moral risks discussed, asymmetrical warfare can also lead to more insurgency and terrorist activity, threatening the citizens of the stronger power.

The biggest changes in warfare will come with the further development of autonomous military robots that can decide who, where, and when to kill, without human involvement. There are no current international guidelines or even discussions about the uses of autonomous robots in warfare. These are needed urgently, since robots simply cannot discriminate between innocents and combatants.

If there was a strong political will to use autonomous robot weapons, or even a serious threat to the state that has them, then legal arguments could be constructed that leave no room for complaints.[7] This is especially the case if they could be released somewhere where there is a fairly high probability that they will kill a considerably greater number of enemy combatants (uniformed and nonuniformed) than innocents (i.e., the civilian death toll was not disproportionate to the military advantage).

At the very least, it should be discussed how to limit the range and action of autonomous robot weapons before their inevitable proliferation (forty-three countries now have military robot programs). Even if all of the elements discussed here could be accommodated within the existing laws of war, their application needs to be thought through properly, and specific new laws should be implemented to not just accommodate their use, but to constrain it as well. We don't know how autonomous robots will affect military strategy of the future, or if they will lead to more subjugation of weak nation-states and less public pressure to prevent wars.

Notes

1. See du Picq 1946. The book was compiled from notes left by Colonel Ardant du Picq of France after he was killed in battle by a Prussian projectile in 1870.

2. Decapitation is a euphemism for assassination of suspected insurgent leaders. The word *decapitation* was used to indicate cutting off the head (leader) from the body of the insurgents.

3. Thanks to Richard Moyes of Landmine Action for pointing me to the BLU-108 and to Marian Westerberg and Robert Buckley from Textron Defense Systems for their careful reading and comments on my description.

4. Contract #W911NF-06–1-0252 from the U.S. Army Research Office.

5. Dover, Delaware, is the U.S. Air Force base where the bodies of soldiers are returned from the front line in flag-draped coffins. The Dover test concerns how much the electoral chances of the national political administration are affected by the numbers of dead.

6. Bugsplat software and its successors have been used to help calculate the correct bomb to use to destroy a target and calculate the impact. It is only used to help in the human decision-making process and it is unclear how successful this approach has been in limiting civilian casualties.

7. Regardless of treaties and agreements, any weapon that has been developed may be used if the survival of a state is in question. The International Court of Justice *Nuclear Weapons Advisory Opinion* (1996) decided that it could not definitively conclude that in every circumstance the threat or use of nuclear weapons was axiomatically contrary to international law; see Stephens and Lewis 2005.

References

Alston, Philip. 2010. *Report of the Special Reporter on Extrajudicial, Summary, or Arbitrary Executions*. The UN Human Rights Council, fourteenth session, A/HRC/14/24/Add.6, May 28.

Arkin, Ronald. 2009. *Governing Lethal Behavior in Autonomous Robots*. Boca Raton, FL: Chapman and Hall/CRC Press.

Arkin, Ronald, and Lilia Moshkina. 2007. Lethality and autonomous robots: An ethical stance. Paper presented at the IEEE International Symposium on Technology and Society, June 1–2, Las Vegas.

Armstrong, Matthew. 2007. Unintended consequences of unmanned warfare. Presentation to Proteus Management Group Futures Workshop at the U.S. Army War College, Carlisle Barracks, Pennsylvania, August 15.

Canning, John. 2005. A definitive work on factors impacting the arming of unmanned vehicles. Dahlgren Division Naval Surface Warfare Center report NSWCDD/TR-05/36.

Canning, John. 2006. A concept of operations for armed autonomous systems. Presentation for the Naval Surface Warfare Center, Dahlgren Division.

Canning, John, G. W. Riggs, O. T. Holland, and C. J. Blakelock, 2004. A concept for the operation of armed autonomous systems on the battlefield. Paper presented at the Association for Unmanned Vehicle Systems International conference, Anaheim, CA, August 17.

Chakwin, Mark, Dieter Voelkel, and Enright Scott. 2002. Leaders as targets. Joint Forces Staff College, Norfolk, VA. Seminar #08.

Committee on Autonomous Vehicles in Support of Naval Operations National Research Council. 2005. *Autonomous Vehicles in Support of Naval Operations*. Washington, DC: The National Academies Press.

Crane, David, 2008. G-NIUS Guardium UGV: World's first operational security robot. *Defense Review* (August): 23. <http://www.defensereview.com/g-nius-guardium-ugv-worlds-first-operational-autonomous-security-robot/> (accessed April 3, 2011).

Daddis, Gregory. 2004. Understanding fear's effect on unit effectiveness. *Military Review* (July–August): 22–27.

Davis, Daniel. 2007. Who decides: Man or machine? *Armed Forces Journal*. <http://www.armedforcesjournal.com/2007/11/3036753> (accessed April 2, 2011).

du Picq, Ardant. 1946. *Battle Studies*. Part 2, chapter 1. Harrisburg, PA: Military Service Publishing Co.

The Economist. 2007. Robot wars: An attempt to build an ethical robotic soldier. April 17.

Fox News. 2008. Pentagon's "Crusher" robot vehicle nearly ready to go. February 27. <http://www.foxnews.com/story/0,2933,332755,00.html> (accessed November 27, 2010).

Ford, John S. 1944. The morality of obliteration bombing. *Theological Studies* 23: 261–309.

Grossman, David. 1995. *On Killing: The Psychological Cost of Learning to Kill in War and Society*. New York: Little, Brown and Co.

Hersh, Seymour. 2002. Manhunt: The Bush administration's new strategy in the war against terrorism. *New Yorker* (December): 64–68.

Holmes, Richard. 2003. *Acts of War: The Behaviour of Men in Battle*. London: Cassell.

Human Rights Watch. 2009. *Precisely Wrong: Gaza Civilians Killed by Israeli Drone-Launched Missiles*. New York: Human Rights Watch.

Hyder, Victor. 2004. *Decapitation Operations: Criteria for Targeting Enemy Leadership*. Monograph/report. Fort Leavenworth, KS: School of Advanced Military Studies United Sates Army Command and General Staff College.

Kahn, Paul. 2002. The paradox of riskless war. *Philosophy and Public Policy Quarterly* 22 (3): 2–8.

Koh, Harold, 2010. Speech to the American Society of International Law, Washington, DC, March 25.

Marshall, S. L. A. [1947] 2000. *Men against Fire: The Problem of Battle Command*. (First published by William Morrow & Company.) Norman: University of Oklahoma Press.

Mental Health Advisory Team. 2006. *Operation Iraqi Freedom 05–07*. Final report, November 17.

Petraeus, David, and James Amos. 2006. *Counterinsurgency*. Field Manual FM 3–24 MCWP 3–33.5, Section 7–30. Washington, DC: Headquarters of the Army.

Roberts, Adam. Forthcoming. What is a civilian? In *The Changing Character of War*, ed. Hew Strachan and Sibylle Scheipers. Oxford, UK: Oxford University Press.

Sharkey, Noel. 2008a. Cassandra or the false prophet of doom: AI robots and war. *IEEE Intelligent Systems* 23 (4) (July/August): 14–17.

Sharkey, Noel. 2008b. The ethical frontiers of robotics. *Science* 322 (5909): 1800–1801.

Sharkey, Noel. 2008c. Grounds for discrimination: Autonomous robot weapons. *RUSI Defence Systems* 11 (2): 86–89.

Sharkey, Noel. 2009a. Death strikes from the sky: The calculus of proportionality. *IEEE Science and Society* (Spring): 16–19.

Sharkey, Noel. 2009b. March of the killer robots. *Daily Telegraph*, June 15.

Singer, Peter Warren. 2009. *Wired for War: The Robotics Revolution and Conflict in the 21st Century*. New York: The Penguin Press.

Sofge, Erik. 2008. Non-answer on armed robot pullout from Iraq reveals fragile bot industry. *Popular Mechanics*, April 8. <http://www.unsysinst.org/forum/viewtopic.php?t=388&sid=b9e24762d3e32d5b72d9a00f540a5640> (accessed November 28, 2010).

Sparrow, Robert. 2007. Killer robots. *Journal of Applied Philosophy* 24 (1): 62–77.

Stephens, Dale, and Michael W. Lewis. 2005. The law of armed conflict—A contemporary critique. *Melbourne Journal of International Law* 6 (1): 55–85.

United Press International. 2008. BAE Systems tech boosts robot UAVs IQ. Industry Briefing, February 26. <http://bae-systems-news.newslib.com/story/3951-3226462/> (accessed April 3, 2011).

United States Air Force. 2008. Guided smart munitions. Call for proposals, topic number AF083–093, August 25.

United States Air Force. 2009. *Unmanned Aircraft Systems Flight Plan 2009–2047*. Headquarters of the United States Air Force, Washington, DC, May 18.