Why We Should Not Build Autonomous Robotic Weapons

William M. Fleischman
Departments of Computing SciencesandMathematicalSciences
Villanova University
Villanova, Pennsylvania 19085, U. S. A.
william.fleischman@villanova.edu

Abstract. We discuss robotic weapons, their advantages and disadvantages, and their effect on the way humans wage war. We consider the factors favoring the development of lethal robotic weapons that can operate autonomously. We discuss the attempt to mitigate the dangers inherent in such weapons by means of an ethical controller implemented in software. We conclude that this is impossible to achieve and therefore that autonomous lethal robotic weapons should not be developed.

Keywords: Computer ethics, Robotic weapons, Autonomous moral agents.

1 Introduction

In this paper, we argue that fully autonomous robotic weapons that have the capacity to kill should not be developed or deployed. We begin with an overview of current robotic devices that are deployed or under development by the military. We discuss the advantages and disadvantages that these weapons confer in combat and in the larger context of decisions to wage war and attitudes toward the conduct of war. We consider the thorny problem of keeping humans "in the loop" in situations where these weapons are used with potentially lethal effects and discuss efforts to develop a so-called "ethical governor" to restrict the behavior of robotic weapons capable of autonomous operation. We present an instructive example from the history of the Cold War that underscores the importance of human deliberation in situations of belligerent confrontation. This example is followed by a discussion of the problem of responsibility and accountability as it applies to autonomous robotic weapons. We conclude with a section of general observations about the choices involved in opting to invest important material and human resources in the development of lethal autonomous weapons.

2 An Overview of Robotic Devices in Use and Under Development

We begin with a short discussion of the recent accelerated development of robotic weapons – unmanned ground and aerial vehicles (UGVs and UAVs) under the impetus of the wars in Iraq and Afghanistan. Numbers tell at least part of the story. There were few of either type of system deployed in the 2003 invasion of Iraq. By 2011, there were an estimated 12,000 UGVs and 7,000 UAVs in the inventory of the

U.S. military forces. Significantly, the U.S. Air Force currently trains more UAV operators than fighter and bomber pilots combined. [1]

Enemy deployment of IEDs in Iraq created an instant demand for Packbots – a ground-based, essentially defensive device developed by iRobot, the Boston-area company originally famous for manufacturing the Roomba robotic vacuum cleaner. The Packbot was used to detect and, if necessary, disarm IEDs without the risk of loss of human life. Initially, it was simply thought of as a "mobile pair of binoculars." With the addition of simple effector arms and grippers the Packbot acquired the capability to disarm and destroy improvised explosive devices concealed by the enemy. [2]

The initial problem addressed was that of locating and identifying non-human threats. A related problem, of course, is the location and identification of human threats – enemy snipers. In this application, however, once a threat is identified, the next job is to eliminate it by killing the sniper. Quite logically, a mobile device that carries a weapon in addition to its cameras provides the possibility of eliminating the threat by aiming and firing remotely under control of a soldier who does not have to appear in the sight of the sniper's weapon. Once again, the desire to shield one's soldiers from situations in which their lives are at risk provides the incentive for development of a robotic device with additional capabilities. As has often been observed [3], the desire to increase the killing effectiveness of one's soldiers while increasing the distance between them and the enemy is a constant in the history of warfare. So this was a natural application for Packbot, Warrior, its more heavily armed successor, and congeners such as the Talon and SWORDS devices manufactured by Foster-Wheeler, a second Boston-based robotics firm. In essence, arming the Packbot or similar robotic device is simply a next step in this historical process.

Not all UGVs have direct combat roles. The special dangers of the role of human medics serving in battlefield situations has led to the development of a version of the Packbot that can search for wounded soldiers and provide a video feed that allows a distant human controller to deploy medical equipment on the so-called "med-bot" in order to evaluate and treat the wounded individual. [2]

UAVs have undergone a similar rapid transformation. Perhaps the best-known UAV is the twenty-seven foot long Predator, capable of carrying out 24-hour reconnaissance and surveillance missions, returning high quality images day and night by means of normal and infrared cameras. Furthermore, the Predator's synthetic-aperture radar can provide valuable information even where the terrain is obscured by clouds, smoke, or dust. As Singer notes, [t]he exact capabilities of the system are classified, but soldiers say they can read a license plate from two miles up." [4]. The same logic that has driven changes in the design of UGVs has resulted in the arming of UAVs which now can carry out offensive missions under the direction of a human pilot or operator located thousands of miles away. In addition, UAV technology has spread in both directions along the size and mission-length continua with the Raven (thirty-eight inches in length and ninety minutes in the air), the Wasp (fifteen inches in length, forty-five minutes of endurance), with micro-UAVs the size of insects in the planning stage. At the other end of the spectrum, the newer Reaper [get some specs if you intend to include this UAV in the discussion] and the Global Hawk (nearly forty-eight feet long with an endurance of thirty-five

hours) provides both wide-area search and high-resolution single target identification and has the capability of autonomous operation between the signals to taxi, take off, and land provided by its human operator. [4]

In addition to deploying its own force of UAVs, the U.S. Navy is also developing various types of unmanned surface and underwater vessels (USVs and UUV's). [4]

3 Advantages and Disadvantages of Robotic Weaponry

It is not hard to see (and it is very hard to resist) the advantages of robotic weaponry. The first, and most compelling for an armed force possessing these weapons, is that they replace humans on the battlefield and therefore reduce the number of human casualties this force will sustain. Beyond this, they are markedly superior to humans in what military strategists describe as the "three D's" – situations that are dangerous, dirty, and dull.

Dirty environments include not only those, like desert battlefields affected by smog, smoke, sand and dust, but also those which have been contaminated by biological, chemical, or radioactive agents. Robots have a very clear advantage in these environments where humans would be encumbered by bulky protective suits and related gear.

Many military missions require concentration over long periods of time. In addition to the physical stress of the activity, there is the psychological stress of paying steady attention in otherwise boring circumstances. Humans can do this for limited periods of time and need downtime or pauses to recover the necessary level of acuity. By contrast, robots don't need to sleep, to eat, or to take a break for "rest and recreation."

The human body is limited in the speed and limits of reaction to threats and forces to which it is exposed in combat situations. From g-forces acting on human pilots of advanced aircraft to speed of recognition and reaction to battlefield dangers, robotic systems appear to have a clear advantage. As already noted, the first advantage of a robot in a dangerous environment is that its destruction involves the loss of a machine (although this may be more consequential if it falls into the hands of an enemy who can study and copy it) and not the loss of a human life.

Related to these factors is calculation regarding risk. Singer notes that, "The unmanning of [an] operation also means that the robot can take risks that a human wouldn't otherwise, risks that might mean fewer mistakes." [4] He cites friendly fire incidents during the Kosovo campaign in 1999 in which the imperative to avoid loss of NATO pilots resulted in orders that planes not be flown at altitudes below 15,000 feet. One of the most grievous errors of this nature occurred when NATO planes flying at these altitudes bombed a convoy of buses carrying Kosovar refugees mistakenly identifying them as a convoy of Serbian tanks. Singer also notes that the "removal of risk allows decisions to be made in a more deliberate manner than normally possible. Soldiers describe how one of the toughest aspects of fighting in cities is how you have to burst into a building and, in a matter of milliseconds, figure out who is an enemy and who is a civilian." In this situation, a robot that can enter a room and shoot only at someone who shoots first has a distinct advantage over the

human who must take fire and somehow instantly manage to determine the source, return fire, and avoid hitting any civilians. [4]

Another advantage that robots have in situations of combat is that they do not suffer from human emotions of rage against adversaries who have caused harm or death to a soldier's comrades. We know of many episodes where otherwise good individuals have given way to extreme emotion and committed atrocities after experiencing the loss of or grievous harm to someone with whom they have bonded and upon whom they have depended in situations of danger. Surely eliminating the danger of such episodes is an important advantage favoring robotic agents over humans.

With all these advantages noted, what could possibly be the downside of the use of robotic weapons? These may be more subtle and harder to see but, in a certain sense, the disadvantages of these weapons are identical with their advantages. One of these disadvantages, clearly recognized by those in command positions in the military, is that over a long time and haltingly we have negotiated barriers against barbaric behavior in war. The Geneva Conventions and treaties barring the use of chemical and biological weapons are among these barriers. When, however, one side in a conflict has such technological superiority, when there is marked asymmetry in the resources each brings to battle, there is an inescapable lessening of the respect that each side owes the other out of recognition of the parity of the risks the combatants share. The sense that the weaker forces can be eradicated like insects by the "magic" of advanced technology acts, in a mutually reinforcing manner, on both sides to undercut the restraints erected against barbarity. [5]

Perhaps the most serious disadvantage of robotic weapons has to do with another set of barriers. General Robert E. Lee, commander of the Confederate forces in the American Civil War of the 19th century once wrote, "It is good that we find war so horrible, or else we would become fond of it." [3] The act of declaring war is or should be a grave existential decision for any country. But we have seen, perhaps most notably in the case of the ill-considered invasion of Iraq by the United States, how consciousness of technological superiority lowers the barrier against waging war.

Paradoxically, to the extent that atrocities committed by otherwise decent soldiers of our military remind us of the horror of war, they serve as a factor that should give pause to anyone contemplating "loosing the dogs of war."

4 Keeping Humans 'In the Loop'

This section is the easiest to write and the most frightening. When it comes to giving robotic weapons lethal capabilities, official military policy seems to be very clear and emphatic: "Humans must be kept in the loop." The meaning of this is, or should be, that a human must give authorization before any robotic weapon can fire on a human target. In fact, however, whenever this matter is raised in serious discussion, the result is averted eyes and a change in topic. The reasons for this are also clear. Although the ideal is to keep humans in the [command] loop, there are so many factors militating against this that in practice it seems impractical. Why, if there is risk of loss of life on your side in the interval between identification of a lethal threat

and authorization to fire issued by a human controller, should the robotic weapon not be given the capability to fire immediately upon locating the threat? Since the authorization requires communication between controller and weapon, and this communication can be cut or disrupted by the enemy, why should there not be an emergency back-up capability for the weapon to operate autonomously in this situation?

Singer points out that the logic of human control of robotic weapons seems to demand a many-one correspondence between weapons and controllers. But humans are notoriously ill equipped and unreliable for the task of controlling multiple units at one time, even under relatively calm conditions. A Pentagon-funded report notes that, "Even if the tactical commander is aware of the location of all his units, the combat is so fluid and fast-paced that it is very difficult to control them." [3]

Further, as Singer points out, human control of automated weapons systems has already been seriously compromised by the human tendency to "believe what the computer says." The paradigmatic example of this is the case of the downing of Iran Air flight 665 over the Persian Gulf in July 1988 by an American naval vessel patrolling the gulf during the Iran-Iraq war. Iran Air Flight 665 was an Airbus passenger jet on a commercial flight from Tehran, Iran to Dubai via Bandar Abbas. On the morning of the flight, the U. S. Navy guided missile cruiser, the Vincennes, equipped with the Aegis combat system, an integrated weapon control system that uses powerful computers and radars to coordinate, track, and guide weapons to destroy enemy targets. Even though the passenger jet was climbing after takeoff from Bandar Abbas, flying a consistent course, and "squawking" the appropriate radio signal that proclaimed it to be a civilian airliner, the Aegis system radars on board the Vincennes seemed to identify the plane as an assumed enemy fighter jet on a descending attack profile. Even though most members of the crew of the Vincennes and almost everyone on board its sister ships on patrol that morning were reading data that accurately identified the nature of the flight, not one of the eighteen sailors and officers of the Vincennes were willing to question the Aegis system's apparent mistaken designation of an attacking enemy fighter aircraft. As a result, the captain of the Vincennes, an officer with a known penchant for aggressive action, gave the authorization to fire resulting in the destruction of Iran Air Flight 655, killing all 290 passengers and crew, among them sixty-six children. [3, 6]

The problem with Singer's analysis and the flaw in the conclusion drawn is that a software design or software-engineering error that should not have evaded the eye of even undergraduate software engineering students was one of the principal factors implicated in the mistaken characterization of flight 655. In fact, in the process of coordinating data on the radars of the three ships in the patrol, a tag used within the previous hour to label a (friendly) fighter jet making a landing (thus descending) was reassigned as the label for Iran Air Flight 655 on the radars of the Vincennes. [6] So while it is not entirely inaccurate to think of this as an illustration of the way in which humans defer to the "judgment" of computer-controlled systems, it is far more relevant to see this as a warning against placing too much trust in the reliability of even state of the art software engineering.

5 Compensating for the Human 'Out of the Loop'

If the superior capabilities of robotic weapons and the limitations of humans acting as controllers so far compromise the military principle of always keeping the human in the loop, then perhaps we can substitute an "ethical governor" implemented in software for the absent human controller. Properly programmed, weapons acting in autonomous mode could perhaps be constrained to "act ethically in war," observing all the articles of the Geneva Conventions, the laws of war, and, in the local context of the combat in which they are deployed, the relevant rules of engagement. And since they are not subject to the psychological and emotional stresses that affect human combatants, we might even expect that they would act more morally than the human soldiers whose combat roles they assume.

In fact, the NSF and U.S. government agencies associated with the Department of Defense have funded an initiative of precisely this nature. Singer quotes the assertion of Ronald Arkin, a professor of computer science at Georgia Tech who has received support various agencies of the government for just such a project, "Ultimately these systems could have more information to make wiser decisions than a human could make. Some robots are already stronger, faster, and smarter than humans. We want to do better than people, to ultimately save more lives." [5]

In a recent paper, Gerdes and Øhrstrom discuss the possibility of devising a Moral Turing Test, which, in their words, "might enable us to distinguish principles for evaluating morally correct *actions* rather than (as in the original Turing test) skills of articulation." [7] Such a test would constitute a necessary but not sufficient condition for the development of what is referred to as an Artificial Moral Agent. Their analysis, rooted in the work of the logician A. N. Prior [], leads to the conclusion that "Prior was right in claiming that the formulation of a formal system which correctly incorporates all aspects of moral reasoning would in principle require a complete description not only of all relevant moral rules and laws but also of all relevant aspects of the situation in question. However, having such descriptions is tantamount to having a God's eye view of all relevant aspects of reality." Although they conclude that it may still be "possible to formalize important aspects of ethical reasoning in a specific context and thereby contribute to a system which may pass a comparative Moral Turing Test," I take their paper as indicating that even this partial approach to creating an Artificial Moral Agent represents a software engineering project of considerable difficulty and complexity. Since the conditions of actual combat constitute a context of such fluidity and rapid change as to defy the simple description of "a specific [i.e., closed] context," it is not unreasonable to conclude that the project envisioned by Arkin has an impossible goal. The similarity of this case with that of the Strategic Defense Initiative (the so-called 'Star Wars' project) from which David L. Parnas withdrew in a well-known letter and series of critical papers [8] suggests that the appropriate response of computer scientists of good conscience toward Arkin's project or any other claiming to have the purpose of devising an "ethical governor" for autonomous robotic weapons should be to condemn it.

In this light, I think it is important to ask, "What is the purpose of the NSF in funding this "research?" Why should anyone want to do this? One possible motivation is as a salve to the consciences of those who are participating in and drawing public funds from the Department of Defense and the National Science

Foundation in research that they know to be, in the last analysis, destructive and anti-human. We are building these lethal autonomous robotic weapons but they are going to be "stronger, faster, and smarter than humans." We are going to do better than mere humans and we will save many lives. We believe (or convince ourselves that) we can achieve this chimera and therefore we must try (and, of course, inure ourselves to the burden of accepting the public's money in furtherance of this grotesque illusion.)

Again, I want to insist on the question, "Why should anyone **want** to do this?" In the words of Joseph Weizenbaum, "Technological inevitability can thus be seen to be a mere element of a much larger syndrome. Science promised man power. But, as so often happens when people are seduced by promises of power, the price exacted in advance and all along the path, and the price actually paid, is servitude and impotence. Power is nothing if it is not the power to choose. Instrumental reason can make decisions, **but there is all the difference between deciding and choosing.**"[9, emphasis added] What is it that we are choosing when we choose to develop the ability to make war in a way that is better than the way humans wage war?

6 An Instructive Story

As In mid-October of 1962, photographs taken during a U2 surveillance flight over Cuba revealed the presence of missile sites and Soviet missile components on the island. Assurances given both by Andrei Gromyko, the Soviet Foreign Minister, and Nikita Khrushchev, the leader of the Soviet Union, that no Soviet missiles would be installed in Cuba were thus revealed to be a deception. This precipitated what was in all probability the most dangerous episode of the Cold War, a period of fifteen days in which the two superpowers were on a path to war that would have involved attacks using nuclear weapons by each on the other. The consequences of this were and are unimaginable.

In a chapter of the excellent book, *Humanity: A Moral History of the 20th Century*, Jonathan Glover recounts the story of how Khrushchev and Kennedy managed to step back from the brink in spite of the strong forces – intense military competition, mutual suspicion and misjudgment, internal political pressures, the actions of military subordinates in the forces of both countries that exceeded their standing orders – that tended toward war and nuclear disaster. The conditions surrounding the Cuban Missile Crisis enumerated by Glover recapitulate, in an eerie correspondence, the set of misjudgments, miscalculations, and reckless actions that in 1914 led the European powers into a war that can only be considered a disaster for those who fought and for the generation that survived the conflict. How, then, did the leaders of the two superpowers in 1962 avoid the trap? It is a riveting and illuminating story worth the attention of anyone considering the role of autonomous weapons in war. [10]

The story is riveting because this was a very close call. There were pressures on both leaders – from both the political and military establishments as well as the Cuban leader Fidel Castro – to take actions (including on the U. S. side, an air attack and/or invasion of Cuba) which, with Soviet tactical nuclear weapons already deployed in

Cuba, would almost certainly have led to a catastrophic nuclear exchange. Among the factors that appear to have prevented this, there were two that are worthy of reflection in the context of this paper.

The first is that historian Barbara Tuchman had published earlier that year her study, *The Guns of August*, which carefully dissected European internal political pressures, misunderstandings in regard to treaty commitments, ambiguous signals, poor communication among allies and between potential belligerents, and the military preparations once begun that seemed impossible to roll back that led ineluctably to war and disaster for the continent. Both President Kennedy and his closest advisors (including his brother Robert) had read the book and referred to it during the meetings at which the possible responses to the Soviet threat were discussed. According to the memoirs of Robert Kennedy, quoted in Glover [10], JFK spoke with his brother about the European leaders in 1914 saying "they seemed to tumble into war through 'stupidity, individual idiosyncrasies, misunderstandings, and personal complexes of inferiority and grandeur.' He said, 'I am not going to follow a course which will allow anyone to write a comparable book about this time, The Missiles of October. If anybody is around to write after this, they are going to understand that we made every effort to give our adversary room to move. I am not going to push the Russians an inch beyond what is necessary.'"

Of equal weight, on the Russian side, Khrushchev, early in the crisis, sent a letter to President Kennedy in which he wrote: "Should war indeed break out, it would not be in our power to contain or stop it, for such is the logic of war. I have taken part in two wars, and I know that war ends only when it has rolled through cities and villages, sowing death and destruction everywhere … If people do not display wisdom, they will eventually reach the point where they will clash like blind moles, and then mutual annihilation will commence … You and I should not now pull on the ends of the rope in which you have tied a knot of war, because the harder you and I pull, the tighter this know will become. And a time may come when the knot is tied so tight that the person who tied it is no longer capable of untying it, and then the knot will have to be cut." [10]

Both the words of Nikita Khrushchev and the import of Barbara Tuchman's analysis that was present in the minds of John Fitzgerald Kennedy and his advisors resonated with the warning articulated by Robert E. Lee: "It is good that we find war so horrible, or else we would become fond of it." This was a crisis that both leaders understood would forever indelibly bear their signatures, however it unfolded. That personal sense of responsibility and the consciousness of the horrors of war were the factors that made it possible to pull back. Let us imagine the computer system, designed and implemented by individuals without names and without the wisdom of those who read and reflect and are conscious of the horror, let us indeed pause and imagine the system capable of the saving wisdom of Khrushchev and Kennedy.

7 The Question of Accountability and Responsibility

The "Problem of Many Hands," articulated by Helen Nissenbaum in her 1994 paper [11], has become a common-place, a cliché. We cite this problem by name, nod

knowingly in acquiescence of the certainty that any large software engineering project is bound to have some unanticipated failure modes with serious negative consequences. If, as is customary, the project is developed over a significant period of time by a team the membership of which is not fixed, it will be difficult, perhaps impossible, to determine who is responsible for the failure, to say who should be held accountable for harms ultimately engendered in the use of such a system. This is just a fact of life in our technologically sophisticated world. Get over it and move on.

Perhaps we can agree that there are areas of application where the expectation of future benefits resulting from the development of a new technology justifies accepting the risks of such negative consequences – without, however, relinquishing the understanding that, while we are waiting for the realization of such benefits, someone, some organization must be held accountable and accept responsibility for these harms. There are some areas of application where we can agree to take these risks. But there are assuredly areas where this attitude is unjustifiable. The development of autonomous robotic killing weapons is one of them.

Whose name will be on the disaster precipitated by the predictable malfunction of one of these weapons? Whose name will be attached, as Khrushchev and Kennedy were aware theirs would beto the nuclear disaster precipitated by a reckless gesture in the course of the Cuban Missile Crisis? Who will own the damage to what little of civilized culture we still imagine we possess? Certainly not foolish and opportunistic computer scientists like Arkin, whose names will have long been forgotten. In a sense, this is appropriate. However much their work contributes to this damage, the disaster will be ours as a society if we do not recognize the folly of the path we are taking.

8 Concluding Observations

Finally, it is important to recognize that, although many profound thinkers have contributed to our understanding of what it means to act ethically, our ideas about ethical behavior are as much a product of our experience and our emotional wisdom as of our analytical intelligence. Beware the scientist or engineer who claims that technique will substitute for human instinct and wisdom and enable us to program a machine to behave ethically. Even to approximate this would require the solution of a software engineering problem of forbidding complexity. A moment's reflection on our discouraging experience with such systems should give us pause. [12]

Long ago, Joseph Weizenbaum cautioned against the seduction of technique applied to problems for which its application is utterly inappropriate. "There are two kinds of computer applications that either ought not be undertaken at all, or if they are contemplated, should be approached with utmost caution. …The first kind I would call simply obscene. These are ones whose very contemplation ought to give rise to feelings of disgust in every civilized person. … I would put all projects that propose to substitute a computer system for human understanding for a human function that involves interpersonal respect, understanding, and love in [this] category. [9]

Beyond this, in choosing to invest in the chimerical pursuit of the ability to build machines that can "do better than people" at waging war, we are distorting the

priorities on which a civilized society should rest. We seem unable to make a commitment to educating or providing adequate health care for all the children who live among us, but we find it easy to lavish great sums in the pursuit of an obscenity, oblivious to the warning, "It is good that we find war so horrible, or else we would become fond of it."

References

1. Singer, P. W.,Military Robotics and Ethics: A World of Killer Apps, *Nature*, vol. 477, 22 September,2011, pp. 399-401
2. Singer, P. W., Wired for War: The Future of Military Robots, in *Wired (UK)*, August 2009, available at http://www.brookings.edu/opinions/2009/0828_robots_singer.aspx, last accessed 15 July 2013
3. Singer, P. W., *Wired for War: The Robotics Revolution and Conflict in the 21st Century*, Penguin Press, New York, (2009)
4. Singer, P. W., Military Robots and the Laws of War, *The New Atlantic*, Winter 2009 available at http://www.brookings.edu/research/articles/2009/02/winter-robots-singer, last accessed 15 July 2013
5. Singer, P. W., The Ethics of Killer Applications: Why Is It So Hard to Talk About Morality When It Comes to New Military Technology, *Journal of Military Ethics*, vol. 9 no. 4, pp. 299-312 (2010)
6. Iran Air Flight 655, in Wikipedia, at http://en.wikipedia.org/wiki/Iran_Air_Flight_655, , last accessed 15 July 2013
7. Gerdes, A. and Øhrstrom, P., Preliminary Reflections on a Moral Turing Test, in *Proceedings of ETHICOMP 2013, The Possibilities of Ethical ICT*, University of Southern Denmark, Kolding, Denmark, pp. 167-174 (2013)
8. Parnas, D. L, Letter to James H. Offutt, in *Introduction to Computer Ethics, Parts 1 and 2*, at *www.stanford.edu/class/cs181/materials/CS181-Parts1and2.pdf*, last accessed 15 July 2013
9. Weizenbaum, Joseph, *Computer Power and Human Reason: From Judgment to Calculation*, W. H. Freeman, New York, (1976)
10. Glover, Jonathan, *Humanity: A Moral History of the 20th Century*, *2nd edition*, Yale University Press, New Haven, (2012)
11. Nissenbaum, H., Computing and Accountability, Communications of the ACM, vol. 37, no. 1, pp. 73-80 (1994)
12. Fleischman, W., Electronic Voting Systems and the Therac-25: What Have We Learned?, in , *Proceedings of ETHICOMP 2010, The "Backwards, Forwards, and Sideways Changes" of ICT*, UniversitatRoviraiVirgili, Tarragona, Spain, pp. 170-179 (2010)