# Responsibility Practices in Robotic Warfare

**Deborah G. Johnson, Ph.D., and Merel E. Noorman, Ph.D.**

UNMANNED AERIAL VEHICLES (UAVs), also known as drones, are commonplace in U.S. military operations. Many predict increased military use of more sophisticated and more autonomous robots.[1] Increased use of robots has the potential to transform how those directly involved in warfare, as well as the public, perceive and experience war. Military robots allow operators and commanders to be miles away from the battle, engaging in conflicts virtually through computer screens and controls. Video cameras and sensors operated by robots provide technologically mediated renderings of what is happening on the ground, affecting the actions and attitudes of all involved.

Central to the ethical concerns raised by robotic warfare, especially the use of autonomous military robots, are issues of responsibility and accountability. Who will be responsible when robots decide for themselves and behave in unpredictable ways or in ways that their human partners do not understand? For example, who will be responsible if an autonomously operating unmanned aircraft crosses a border without authorization or erroneously identifies a friendly aircraft as a target and shoots it down?[2] Will a day come when robots themselves are considered responsible for their actions?[3]

*Deborah G. Johnson is the Anne Shirley Carter Olsson Professor of Applied Ethics in the Science, Technology, and Society Program at the University of Virginia. She holds a Ph.D., M.A., and M.Phil. from the University of Kansas and a B.Ph. from Wayne State University. The author/editor of seven books, she writes and teaches about the ethical issues in science and engineering, especially those involving computers and information technology.*

*Merel E. Noorman is a postdoctoral researcher at the eHumanities group of the Royal Netherlands Academy for Arts and Sciences. Before joining the eHumanities group, she was a postdoctoral research associate in the Science, Technology, and Society Program at the University of Virginia, where she worked with Deborah Johnson on the research this paper is based on. She received her Ph.D. from Maastricht University and studied artificial intelligence and science and technology studies at the University of Amsterdam and Edinburgh University.*

In principle, humans retain control of—and responsibility for—the behavior of autonomous machines. However, establishing precisely who is responsible for a machine's behavior is challenging. Autonomous machines, no matter how they are defined, developed, or used, operate as part of broad sociotechnical systems involving numerous individuals and organizations.

We advocate the concurrent development of new responsibility practices *with* the development of new technologies rather than before or after those technologies are developed and adopted for use. This is necessary because literature in the field of science and technology studies shows that the trajectory of a technology's development is unpredictable; how a technology takes shape depends on complex negotiations among relevant social groups.[4] The technologies eventually adopted and used are not predetermined by nature or any other factor. No one can predict with certainty how a developing technology will turn out or what new technologies will emerge. In the course of development, a new technology may change in response to many factors, including changes in funding, historical events such as wars, changes in the regulatory environment, and market indicators. The technologies that succeed (i.e., that are adopted and used) are the outcome of complex negotiations among many actors, including engineers and scientists, users, manufacturers, the public, policymakers, politicians, and others.

Negotiations among the actors involved with a new technology are part of the overall discourse around that technology from its earliest stages of development. The discourse about responsibility and autonomous military robots is a case in point; current discourse provides an opportunity to observe issues of responsibility being worked out early in the technology's development. The negotiations between researchers, developers, engineers, philosophers, policymakers, military authorities, lawyers, journalists, and human rights activists are taking place in the media and academic journals, at conferences and trade shows, through drafting new policies and regulations, in negotiating international treaties, and also through designing and developing the technologies. This process contrasts starkly with the all-too-common idea that issues of responsibility are decided after a technology is developed or separately from technological design.

Framing robots as autonomous challenges ordinary notions of responsibility. Autonomy in daily life and moral philosophy implies acting on one's own, controlling one's self, and being responsible for one's actions. On the other hand, being responsible generally means that individuals have some kind of influence or control over their actions and the outcomes of those actions. The idea of the autonomy of robots suggests that humans are not in control of the robots. Hence, at first glance, it may seem that humans should not be held responsible for autonomous robot behavior. However, this narrative of future autonomous robots operating on their own, without human control, is somewhat misleading, and it draws attention away from important choices about responsibility—choices made at the level of design and implementation.

Our analysis of the discourse on autonomous artificial agents and responsibility shows that delegating tasks to autonomous technologies *is* compatible with holding humans responsible for the behavior of those technologies. This is so for at least two reasons. First, the definition of machine autonomy has numerous interpretations, but all involve various kinds and degrees of human control. Second, humans decide who is responsible for the actions of a machine. Their decisions are affected by, but not entirely determined by, the nature of technology. Responsibility for the behavior of autonomous machines is and must continue to be determined by ongoing negotiations between relevant interest groups during the development of new technologies.

## Negotiating Autonomy

Popular accounts of future military robots often portray these technologies as entities with capabilities that rival or surpass those of humans. We are told that robots of the future will have the ability to think, perceive, and even make moral decisions. In *Discover Magazine*, for instance, Mark Anderson writes, "As surely as every modern jetliner runs primarily on autopilot, tomorrow's military robots will increasingly operate on their own initiative. Before the decade is out, some fighting force may well succeed in fielding a military robot that can kill without a joystick operator behind a curtain elsewhere in the world."[5] Such narratives raise concerns about the lack of human control, and as a result, they confound the determination of human responsibility.

However, in robotics and computer science, autonomy has many different meanings. It tends to be used metaphorically to emphasize certain features of a computational system that set it apart from other systems. Three conceptions of machine autonomy—as high-end automation, as something other than automation, or as collaborative autonomy—illustrate that humans do not necessarily lose control when tasks are delegated to autonomous systems. Rather, the delegation of tasks to these systems transforms the character of human control.

**Autonomy as high-end automation.** In its report, "The Role of Autonomy in Department of Defense Systems," the Defense Science Board Task Force characterizes autonomy as "a capability (or a set of capabilities) that enables a particular action of a system to be automatic or, within programmed boundaries, 'self-governing.'"[6] Capability, here, refers to a particular process (or processes) consisting of one or more tasks, such as navigation or flight control. This definition echoes a more traditional way of conceptualizing machine autonomy as at the high end of a continuous scale of increasing automation. In this way of thinking, automation involves the mechanization of tasks, where routine actions are translated into some formalized and discrete steps such that a machine can perform them.[7] At the high end of the automation scale are systems in which the automated machine performs most or all of the steps in a process. At the low end of the scale are systems in which decision making and control of the process are left largely to human operators. Autonomy is attributed to those systems with higher levels of automation. Such systems close the control loop over a process, i.e., most of the tasks in the process are automated while human operators make few, if any, decisions.

Machine autonomy, in this way of thinking is bounded; it extends only as far as the automated process. Therefore, in this kind of machine autonomy, humans are in control of what the machine does, even if they do not directly intervene or are not 'in the loop,' because they fully specify the process and the routine tasks the machine performs.

**Autonomy as something other than automation.** However, some participants in the discourse sharply distinguish machine autonomy from automation. They argue, for example, that autonomous systems (of the future) will be different from automated systems because their behavior will not be preprogrammed. Autonomous systems will only have to be instructed what to do, not how to do it.[8] Human operators and designers will not have to specify in advance all the behavior sequences that should follow a particular input.

In its *Unmanned Systems Integrated Roadmap FY2011–2036*, the Department of Defense (DOD) provides an illustration of this second take on machine autonomy. It argues that autonomous systems are "self-directed toward a goal in that they do not require outside control, but rather are governed by laws and strategies that direct their behavior."[9] Their behavior in response to certain events is not fully specified or preprogrammed. According to this

Lt. Gen. Jeffrey Talley (then Brig. Gen.), commander of 926th Engineer Brigade, Multi-National Division, watches a demonstration of robotic route-clearing equipment at 6th Iraqi Army Division headquarters motor pool, Iraq, 5 January 2009.

U.S. Army Sgt. Benjamin D. Parker, an explosive ordnance disposal team leader, and Spc. Chase Donnelly, a robotics operator, prepare their robot to inspect a suspected improvised explosive device in eastern Afghanistan's Nangarhar Province. (U.S. Army, Sgt. Tracy J. Smith)

2011 update of the *Roadmap*, "an autonomous system is able to make a decision based on a set of rules and/or limitations. It is able to determine what information is important in making a decision."[10] By contrast, the DOD argues, automatic systems are fully preprogrammed. They can "act repeatedly and independently of external influence or control," but they "follow a predefined path," and their behavior has to be fully specified in advance.[11]

Machine autonomy, from this perspective, refers to robotic systems that would somehow be more flexible and unpredictable, compared to automated systems, in deciding how to operate—given predefined goals, rules, or norms. Those that make this distinction about autonomy tend to point to artificial intelligence technologies—such as machine learning or probabilistic reasoning methods—as technologies that would enable these kinds of robotic systems.[12] Robots equipped with these kinds of technologies would be able to learn from experience and adapt to changing circumstances as well as deal with uncertain or missing data. Such descriptions of autonomy seem to suggest that human operators as well as developers would have less control over the behavior of the system. The machine would not only operate independently of the human operator, but also, to a certain extent, independently of its human creators.

Nevertheless, even here, autonomy does not mean that machines are free in the decisions they make; the conditions for making a decision are carefully set by humans. As the DOD's 2011 conception of autonomy shows, laws and strategies provided by humans will still govern the behavior of autonomous systems. The envisioned systems could vary their behavior as long as they stayed within these predefined constraints. Note that this would be a remarkable feat, as it would mean these robots could interpret laws and strategies, applying them appropriately in ever-changing sociotechnical contexts.

Regardless of whether this is possible, developers and designers would delimit the problem any robotic system is intended to solve. If the envisioned robotic technologies were based on artificial intelligence methods now in development, then those artificial intelligence methods would limit any robotic system's abilities to act independently. Although programmers and developers would not have to specify all the possible situations with which the software has to contend, designers would have to generate a model that approximates the behavior of particular aspects of the world and their uncertainties. Learning and probabilistic algorithms would be able to operate more flexibly than a preprogrammed deterministic algorithm because they would allow for variations and could respond to certain unanticipated contingencies. Nevertheless, this flexibility is a function of the problem definitions and the world models that the developers or programmers of the algorithm have formulated. Therefore, even where machine autonomy is considered more than high-level automation, the autonomy of the machine does not mean there is no human control because humans design, choose, and plan for the strategies employed by the machine.

***Humans exert their influence by defining the conditions for machine behavior.***

**Collaborative autonomy.** Both conceptions of machine autonomy described above (autonomy as high-level automation and autonomy as something other than automation) focus on what machines can do *without* direct human control. However, machine autonomy does not necessarily mean that humans will be taken out of the loop. Human operators may still be involved in the decision-making processes that autonomous robots execute. As explained in an earlier edition (published in 2009) of the *Roadmap (Unmanned Systems Integrated Roadmap FY 2009–2034)*: "First and foremost, the level of autonomy should continue to progress from today's fairly high level of human control/intervention to a high level of autonomous tactical behavior that enables more timely and informed human oversight."[13] The text implies an expectation that robots will operate in support of and in close communication with human actors, such that human oversight remains possible.

The ability of robots to engage in joint activities with humans has received more attention in human–computer interaction research, where researchers use terms such as collaborative control, situated autonomy, or adaptive autonomy.[14] Robin R. Murphy and David D. Woods, for example, have argued for what they call "situated autonomy," a notion that stresses the responsiveness of robots to humans.[15] They contend that robots should have the capability to respond to humans, as appropriate to the humans' roles. That is, a robot's behavior should be attuned to the relationships and social roles of the humans with which it interacts. Thus, a robot may request a confirmation from a superior when it receives a command from a human operator that exceeds the operator's level of authority, or it may decide to transfer control back to an operator when appropriate for the situation. The requirement of responsiveness, Murphy and Woods argue, captures a new form of autonomy, "not as isolated action but the more difficult behavior of engaging appropriately with others."[16] This type of autonomy places the emphasis on the interaction between humans and robots. It implies that robots should be designed so that control can be transferred smoothly from the human operator to the robot and back.

The Defense Science Board Task Force, cochaired by Murphy, uses a similar collaborative conception of autonomy. In their document *Task Force Report: The Role of Autonomy in DOD Systems,* the Task Force argues that many of the DOD studies of autonomy focus too much on machines and not enough on the human–machine system. They argue instead for the adoption of an "autonomous systems reference framework" that focuses on the explicit allocation of functions and responsibilities between human and computer, recognizing that these allocations may vary depending on the context. The framework should also make choices explicit about trade-offs inherent in technological design, such as optimization versus resilience or centralized information systems versus distributed systems. Human decisions about the allocation of control are thus an explicit part of the design process—a process that places overall control firmly in the hands of humans.

## Human Influence over Military Robots

None of the three approaches to the autonomy of robots described above implies that humans are not in control of the technology they create and deploy. The Defense Science Board Task Force even argues that "it should be made clear that all autonomous systems are supervised by human operators at some level, and autonomous systems' software embodies the designed limits on the actions and decisions delegated to the computer."[17] Instead of no human control, robot (or machine) autonomy appears to mean that humans have different kinds of control. Humans exert their influence by defining the conditions for machine behavior. They choose the mathematical and probabilistic models that will guide the behavior of the robotic system and determine the margins of error on what the robot can and cannot do. Designers, developers, managers, and operators set constraints on the behavior that robotic systems are allowed to exhibit.

As military robots become more autonomous, it would seem that they should only be allowed to operate autonomously if they exhibit predictable and reliable behavior. For example, an unmanned helicopter would be allowed to fly into an unknown environment only if the software controlling the helicopter would adhere to certain expectations and norms. The helicopter should not fly into trees, it should execute given instructions, and it should fly between waypoints in a limited amount of time. If the helicopter would not perform as expected, it would be regarded as malfunctioning.

It should not be surprising, then, that the idea of more autonomous robotic systems comes with an increased emphasis on reliability of and trust in technology, along with the need to develop better methods for verification and validation. In the *Report on Technology Horizons: A Vision for Air Force Science & Technology 2010–2030,* the U.S. Air Force chief scientist argues that although it is possible to develop systems with relatively high levels of autonomy, the lack of suitable verification and validation methods stands in the way of certifying these technologies for use.[18] The report claims that in the near- to mid-term future, developing methods for "certifiable trust in autonomous systems is the single greatest technical barrier that must be overcome to obtain the capability advantages that are achievable by increasing use of autonomous

U.S. Army soldiers operate a pack robot at Forward Operating Base Hawk, Iraq, 18 September 2008. (U.S. Air Force, Staff Sgt. Manuel J. Martinez) systems."[19] These observations reflect the need for reliability and predictability that one would expect from an organization in which command responsibility is a guiding principle. At the same time, they show that control of autonomous systems is partly in the hands of those who develop verification and validation methods or other methods of ensuring trust and confidence in these systems.

The three different conceptions of autonomy illustrate that autonomy does not mean that robots are or will be out of the control of humans. The different approaches to machine autonomy may, nevertheless, have effects on how issues of responsibility are understood and managed since they shape the activities and attitudes of humans and relations between them.[20] Robots that can operate with more flexibility in unknown environments, for instance, may affect the way humans perceive and experience what it means to be in control of outcomes and thus what it means to be responsible. Further automation of decision making brings in developers, testers, and others, which may influence how responsibility is distributed.

Nevertheless, although technologies shape how participants in the system perceive, experience, and behave with autonomous systems, they do not determine these things. Nor do they determine responsibility. Responsibility for the behavior of an autonomous system is a matter that humans involved in the development and use of autonomous military robots negotiate and will likely continue to negotiate.

## Responsibility Practices for Military Robots

As our discussion of different concepts of autonomy hints at, functioning military robots are not simply machines; they are not even simply intelligent and autonomous machines. They are sociotechnical systems. That is, the robot machine is a component in a system consisting of human actors and artifacts. The behaviors of both combine to produce a system that achieves (or attempts to achieve) human goals. When a robot is used in a military operation—say, an unmanned aircraft is sent to eliminate a target—the operation consists of human and nonhuman behavior. Humans decide if there is sufficient evidence of an appropriate target; a few of them decide whether and when to deploy a UAV; others sign off on the decision; yet others monitor and communicate with the UAV; and so on. And, of course, many humans have been involved in the design and manufacturing of the UAV and its delivery to a particular location.

Although sociotechnical systems function by delegating specific tasks to each component of the system, delegation of tasks is not the same as delegation of responsibility. Artifacts—such as machines, software, and mechanical parts—might be considered responsible for the performance of particular tasks, but this use of "responsible" is limited to performance and possibly the effects of tasks. Because other senses of responsibility—such as moral, legal, and professional—require conscious deliberation and voluntary actions, they apply, conventionally at least, only to human beings or groups of human beings. In philosophical and moral traditions, ascribing responsibility for a particular outcome to a person requires that the person acted freely and was able to consider the consequences. Machines typically lack such capabilities.

Ascribing responsibility can nevertheless be a challenge. Conventional moral notions can be difficult to apply in practice because individuals rarely have full control over outcomes, and they seldom know exactly what the consequences of their actions will be.[21] In sociotechnical systems, it can be difficult to figure out what happened and who was responsible for which actions or consequences following an untoward event.[22] Numerous individuals and institutions act with and in sociotechnical systems, and human and technological components affect each other in contingent ways. However, although ascribing responsibility can be challenging, this is not to say that no one is responsible.

Human responsibility can best be understood as constituted through a set of *responsibility practices.* Responsibility practices are the established ways that people within a particular environment or community understand, assign, and ascribe responsibility based on shared values and ideas about fairness and utility. These practices involve accepted ways of evaluating actions, holding others to account, blaming or praising, and conveying expectations about obligations and duties. They are also about prevailing norms and moral principles that guide the behavior of the members of a community.

Responsibility practices are both forward- and backward-looking. Forward-looking responsibility involves specifying which tasks and duties are assigned to which individuals and to which nonhuman components. Such practices might be promulgated through job descriptions, instruction manuals, ethical codes, observation of past practices, training before taking on a role, and so on. Backward-looking responsibility involves practices of tracing back what happened and identifying what went wrong. When a failure occurs, humans will seek out the cause of the failure, and humans operating in the system will be asked to account for their behavior. Backward-looking responsibility generally relies on, or at least presumes something about, forward-looking responsibility. That is, to understand what went wrong or what is to blame, we have to understand how tasks and responsibilities were assigned.

The extent to which individuals operating in the system are perceived to have responsibility or feel themselves in a position of responsibility is not simply a matter of tasks being delegated. It also depends on how responsibility practices convey expectations about the duties and obligations of the humans and hold actors accountable for performing or failing to perform as expected. Whether someone is considered responsible depends, as well, on evolving notions of what it means to be in control and able to think about the consequences of certain actions.

In a given system, adoption of a new technology may lead to negotiations about changes to existing responsibility practices, creation of entirely new practices, or both. Established practices may not accommodate the changes produced from introducing the new technology, e.g., changes in activities, attitudes, and relationships between people. The real-time stream of data that current UAVs produce is a good example here. The role of pilots has changed insofar as they now monitor video images and continuously communicate with others who have access to the same data and images (e.g., the sensor operator, the mission intelligence coordinator, and the data analysts miles away in an information fusion center). This has transformed the way targeting decisions are made, compared to manned operations. Decision making has become more shared and less compartmentalized. As a result, established ideas about what various human actors are supposed to do and what they have to account for have had to be adjusted.[23]

New norms and rules have to be established to govern the activities a new technology makes possible. Duties and obligations have to be reevaluated and redefined. Mechanisms for evaluating actions and holding others to account have to be adjusted or created. This will also be the case for future autonomous technologies. Regardless of how machine autonomy is interpreted, whether someone is responsible for the behavior of the system will not only depend on what the machine can and cannot do, it will also depend on the practices that prevail in the context.

***Human responsibility can best be understood as constituted through a set of responsibility practices.***

Shared values and principles may shape the establishment of new practices. In the case of UAVs, organizational values and national and international laws provide a moral framework for and set limits on the new activities these technologies enable. Take the principle of distinction, a key principle in international law that states civilians should be distinguished from combatants. This principle is intertwined with established responsibility practices within military organizations, as they are part of their routines, protocols, and procedures.

Yet, these shared values and principles are subject to negotiation. Achieving an interpretation of them may be challenging because of the introduction of new technologies and also because of social, political, and economic developments. The current debates about the use of drones provide a pertinent example. One contentious issue is that, according to anonymous government officials, the U.S. government regards all military-age males killed in a drone strike as combatants unless proven otherwise.[24] Such a controversial and broad interpretation of a key principle of the law of war affects responsibility practices significantly, at least in the sense that soldiers involved in deploying drones are held to a certain standard of responsibility for harm to noncombatants.

Responsibility practices are continuously negotiated and renegotiated. This can often be seen when something goes wrong with a new technology, and investigators trace back the cause of the failure.

U.S. airmen with the 62nd Expeditionary Reconnaissance Squadron speak to Afghan men and children about an MQ-9 Reaper unmanned aerial vehicle during the 2012 Kandahar Air Wing open house in Kandahar, Afghanistan, 1 January 2012. (U.S. Air Force, Staff Sgt. David Carbajal)

They may discover that something should have been done—and in the future should be done—differently. For instance, the precise role of UAV operators was not immediately clear when UAVs were first introduced. Various investigative reports following UAV mishaps and accidents have made recommendations to adjust and enhance training programs, procedures, communication protocols, and task assignments. The recommendations are targeted to delineate clearly who is responsible for what and to enhance the conditions under which individuals make decisions.[25] Such reports reveal evolving notions of the kind of skills and knowledge that operators need as well as changing norms that govern their behavior.

Negotiations about responsibility practices also may involve adjustments to the technology. In its report on autonomy in DOD systems, the Defense Science Board, for example, stressed the need for a more careful consideration of human factors.[26] Neglect of human–robot interaction in the early UAV development programs resulted in a relatively high number of mishaps. Operators made mistakes due to confusing interfaces and information overload. The Defense Science Board's report calls for changes to the existing interfaces.

Therefore, responsibility is best conceived of as a set of practices built on the foundation of a distribution of tasks. Responsibility practices are reinforced by activities that promulgate expectations about what individuals are supposed to do and what happens when failures occur. Among other things, organizations create expectations through policies and through their organizational culture. Responsibility practices develop expectations of how human and nonhuman components will behave (i.e., who is responsible for doing what) and specify what should or will happen when there is a failure to live up to expectations. These expectations and ideas about responsibility influence the design and eventual use of technologies. Increasingly autonomous technologies may necessitate changes to existing responsibility practices and creation of some entirely new practices in the future.

## Conclusion

The use of autonomous artificial agents raises significant issues of responsibility and accountability, and this is especially so when the artificial agents are part of military operations. Whether and in what ways humans are responsible for the behavior of artificial agents is not just a matter of delegating tasks to machines. Negotiations about responsibility for the behavior of these agents are ongoing with the development of the technologies. These negotiations involve a variety of actors, including the scientists and engineers designing the technologies and users such as the military or the public. Although it is difficult to predict where current negotiations will end up (and that is not our goal), our analysis shows that different notions of autonomy are being used, and each has distinctive implications for how we think about responsibility. At the same time, issues of responsibility are not determined by the design of the artificial agent. Decisions about responsibility, i.e., who is responsible for what behavior, are also made in the development and evolution of social practices that constitute the operation of artificial agents.

None of this is to say we should stop being concerned about the tasks assigned to the nonhuman components of military robotic systems. On the contrary, concerns about responsibility should be an important part of the negotiations. They should shape the delegation of tasks to the human and nonhuman components of these systems. The danger in concentrating on the technological side of autonomous robots is that the development of responsibility practices will be neglected. Instead of focusing on whether robots or humans can be held responsible for robots' behavior, we should focus on the best allocation of tasks and control among human and nonhuman components *and* how best to develop responsibility practices. ***MR***

*This paper is based on the research done for the National Science Foundation (Grant # 1058457) and draws on several papers that have been written as part of this project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.*

NOTES

1. Unmanned aerial vehicles (UAVs) are sometimes referred to as remotely piloted aircraft to emphasize the need for a pilot or unmanned systems to underline that the vehicle is inextricably linked to a larger system of human actors and technologies. For clarity, we use the terms drone and UAV. For different perspectives on the future of drones, see T. Adams, "Future Warfare and the Decline of Human Decision-Making," *Parameters* (Winter 2001-02): 55-71; Patrick Lin, George Bekey, and Keith Abney, *Autonomous Military Robots: Risk, Ethics, and Design* (California Polytechnic State University, 2008), <http://ethics.calpoly.edu/ONR_report.pdf>; Peter W. Singer, *Wired for War: The Robotics Revolution and Conflict in the 21st Century* (New York, New York: Penguin, 2009); U.S. Air Force Chief Scientist, *Report on Technology Horizons: A Vision for Air Force Science & Technology 2010–2030* (AF/ST-TR-10-01-PR, Maxwell Air Force Base, Alabama, September 2011). <http://www.defenseinnovationmarketplace.mil/resources/AF_TechnologyHorizons2010-2030.pdf>.

2. Peter M. Asaro, "How Just Could a Robot War Be," *Current Issues in Computing And Philosophy*, edited by Philip Brey, Adam Briggle and Katinka Waelbers, (Amsterdam, The Netherlands: IOS Press, 2008), 50-64.

3. Lin, Bekey, and Abney.

4. For an overview of science and technology studies, see Wiebe E. Bijker, Thomas P. Hughes, and Trevor Pinch, The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology (London, UK: The MIT Press, 1987).

5. Mark Anderson, "How Does a Terminator Know When to Not Terminate?" *Discover Magazine*, 31(4)(May 2010): 36.

6. DOD, Defense Science Board, *Task Force Report: The Role of Autonomy in DOD Systems* (Washington, DC: U.S. Government Printing Office [GPO], July 2012), <http://www.acq.osd.mil/dsb/reports/AutonomyReport.pdf>, 1.

7. See Thomas B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control* (Cambridge, MA: MIT Press, 1992).

8. See for example Bruce T. Clough, Metrics, Schmetrics: *How the Heck Do You Determine a UAV's Autonomy Anyway?* (Technical report, Air Force Research Lab., Wright-Patterson AFB, Ohio, 2002).

9. DOD, *Unmanned Systems Integrated Roadmap* FY2011-2036, (Washington, DC: GPO, 2011), <http://www.acq.osd.mil/sts/docs/Unmanned%20Systems%20Integrated%20Roadmap%20FY2011-2036.pdf>, 43.

10. Ibid., 43.

11. Ibid., 43.

12. U.S. Air Force Chief Scientist, *Report on Technology Horizons*.

13. DOD, *Unmanned Systems Integrated Roadmap FY2009-2034* (Washington, DC: GPO, 20 April 2009), <www.dtic.mil/dtic/tr/fulltext/u2/a522247.pdf>, 27.

14. See for a brief overview Matthew Johnson, et al., "The Fundamental Principle of Coactive Design Interdependence Must Shape Autonomy," in COIN@AAMAS'10, *Proceedings of the 6th International Conference on Coordination, Organizations, Institutions, and Norms in Agent Systems*, (2010), 172-191.

15. Robin R. Murphy and David D. Woods, "Beyond Asimov: The Three Laws of Responsible Robotics," *IEEE Intelligent Systems*, 24(4) (July-August, 2009):14-20.

16. Ibid., 18.

17. DOD, Defense Science Board, *Task Force Report*, 1.

18. U.S. Air Force Chief Scientist, *Report on Technology Horizons*, IX.

19. Ibid., 42.

20. Peter-Paul Verbeek, "*Materializing Morality*," *Science, Technology and Human Values*, 31(3) (2006): 361-380.

21. Hans Jonas, *The Imperative of Responsibility: In Search of an Ethics for the Technological Age* (Chicago: The Chicago University Press, 1984); Mark Coeckelbergh, "Moral Responsibility, Technology, and Experiences of the Tragic: From Kierkegaard to Offshore Engineering," *Science and Engineering Ethics*, 18 (2012): 35-48.

22. Deborah G. Johnson and Thomas M. Powers, "Computer Systems and Responsibility: a Normative Look at Technological Complexity," *Ethics and Information Technology*, 7(2) (2005): 99-107.

23. Peter M. Asaro, "The Labor of Surveillance and Bureaucratized Killing: New Subjectivities of Military Drone Operators," *Journal of Social Semiotics*, 23(2) (2013): 196-224.

24. Jo Becker and Scott Shane, "Secret 'Kill List' Proves a Test of Obama's Principles and Will," *New York Times*, (May 29, 2012), <http://www.nytimes.com/2012/05/29/world/obamas-leadership-in-war-on-al-qaeda.html?pagewanted=all>.

25. See for example David S. Cloud, "Anatomy of an Afghan War Tragedy," *Los Angeles Times* (April 10, 2011) <http://articles.latimes.com/2011/apr/10/world/la-fg-afghanistan-drone-20110410>; Sharon D. Manning, Clarence E. Rash, Patricia A. LeDuc, Robert K. Noback and Joseph McKeon, *The Role of Human Causal Factors in U.S. Army Unmanned Aerial Vehicle Accidents* (U.S. Army Aeromedical Research Laboratory Report # 2004-11, 2004); Kevin W. Williams, *A Summary of Unmanned Aircraft Accident/Incident Data: Human Factors Implications* (Washington, DC: U.S. Department of Transportation, Federal Aviation Administration, Office of Aerospace Medicine. Technical Report Publication No. DOT/FAA/AM-04/24, 2004).

26. DOD, Defense Science Board, *Task Force Report*.