Issues in GOVERNANCE STUDIES

Number 50 August 2012

Understanding "The Loop": Hum-ans and the Next Drone Generations

William C. Marra and Sonia K. McNeil


William C. Marra and **Sonia K. McNeil** are recent graduates of Harvard Law School.

Drones have revolutionized warfare. They may soon transform civilian life, too. Three machines have already been patrolling the Mexican border[1] and assisting with other law enforcement efforts.[2] And Congress has voted to further expand the use of drones at home, directing the Federal Aviation Administration to unshackle restrictions on domestic drones by 2015.[3]

U.S. Air Force personnel observe a test flight of a drone in the microaviary lab at Wright Patterson Air Force Base in Dayton, Ohio. 12/07/2011. Reuters.

As amazing as today's aerial drones may seem, they are merely the "Model T" of robot technology.[4] Most are souped-up remote-controlled airplanes, still with a human pilot, though he or she now sits at a military base, rather than in the cockpit. Today's drones do not think, decide, and act on their own. In engineering speak, they are merely "automated."

Tomorrow's drones are expected to leap from automation to "autonomy." These highly sophisticated machines will have the ability to undertake missions

---

[1] William Booth, *More Predator drones fly U.S.-Mexico border*, WASH. POST, Dec. 21, 2011.
[2] Nick Paumgarten, *Here's Looking at You: Should we worry about the rise of drones?*, THE NEW YORKER, May 14, 2012, at 46.
[3] Nick Wingfield & Somini Sengupta, *Drones Set Sights on U.S. Skies*, N.Y. TIMES, Feb. 17, 2012.
[4] PETER W. SINGER, WIRED FOR WAR: THE ROBOTS REVOLUTION AND CONFLICT IN THE 21ST CENTURY 110 (2009).

with little or no guidance from a human operator. The difficult policy questions raised by today's automated drones will seem pedestrian compared to the ones created by tomorrow's technologies.

Today, humans are still very much "in the loop."[5] Humans generally decide when to launch a drone, where it should fly, and whether it should take action against a suspect. But as drones develop greater autonomy, humans will increasingly be "out of the loop." Human operators will not be necessary to decide when a drone (or perhaps a swarm of microscopic drones) takes off, where it goes, and how it acts.

Autonomy is no longer solely a feature of humans. Whether it is a desirable quality for machines will be among the most important policy questions of the coming years.

Regulations for today's airborne drones should be crafted with an eye toward tomorrow's technologies. Policymakers must better understand how the next generation of autonomous systems willchange, compared to today's merely automated machines. As we discuss, language useful to the policymaking process has already been developed in the same places as drones themselves — research and engineering laboratories across the country and around the globe. We introduce this vocabulary here to explain how tomorrow's drones will differ and suggest possible approaches to regulation.

Autonomy is no longer solely a feature of humans. Whether it is a desirable quality for machines will be among the most important policy questions of the coming years.

Demystifying Machine Decision-Making: Introducing the "OODA Loop"

To understand the technology behind the unmanned drones of the future, begin in the past — with the dogfights between manned fighter jets during the Korean War. Air Force pilot and military strategist John Boyd wondered why American F-80 fighter planes got the better of Soviet MiG-15 jets in battle. Boyd realized that the advantage lay with the pilot whose decisions were faster and more accurate than his opponent's.[6] He distilled this decision-making process into a four-step loop: Observe, Orient, Decide, Act.[7] This "OODA Loop" transformed the way military leaders think about combat. It also influenced other fields, including business, sports, and engineering[8] — in short, "anywhere

[5] *See* Bernd Debusmann, *More drones, more robots, more wars?*, REUTERS, Jan. 31, 2012; Shane Harris, *Out of the Loop: The Human-free Future of Unmanned Aerial Vehicles* (Hoover Institute 2012).

[6] Scott E. McIntosh, *The Wingman-Philosopher of MiG Alley: John Boyd and the OODA Loop*, 58 AIR POWER HIST. 24, 26–27 (2011).

[7] ROBERT CORAM, BOYD: THE FIGHTER PILOT WHO CHANGED THE ART OF WAR 334 (2002); Berndt Brehmer, *The Dynamic OODA Loop: Amalgamating Boyd's OODA Loop & the Cybernetic Approach to Command and Control*, at 2 (remarks at the 10th Annual International Command & Control Research and Technology Symposium 2005); McIntosh, *supra* note 6, at 27.

[8] *See, e.g.*, Eric Sholes, *Evolution of a UAV Autonomy Classification Taxonomy* (remarks at the 2007 IEEE Aerospace Conference); Raja Parasuraman et al., *A Model for Types and Levels of Human Interaction with Automation*, 30 IEEE TRANSACTIONS ON SYSTEMS, MAN, & CYBERNETICS PART A SYS. & HUMANS 286, 288 (2000).

a competitor seeks an edge."[9]

In Boyd's model, an individual first *observes* her environment, gathering raw information about her surroundings through the array of human senses. Second, she *orients* herself, or interprets the information she has gathered and converts it through analysis into conclusions. Third, she weighs potential courses of action based on the knowledge she has accumulated and *decides* what to do. Fourth and finally, she *acts*, or executes the decision she has made.

This description simplifies a complex, non-linear process.[10] Fully articulated, the decision-making process includes constant feedback and refinement. In all its complexity, the OODA Loop appears as follows:

Figure One: Boyd's OODA Loop[11]

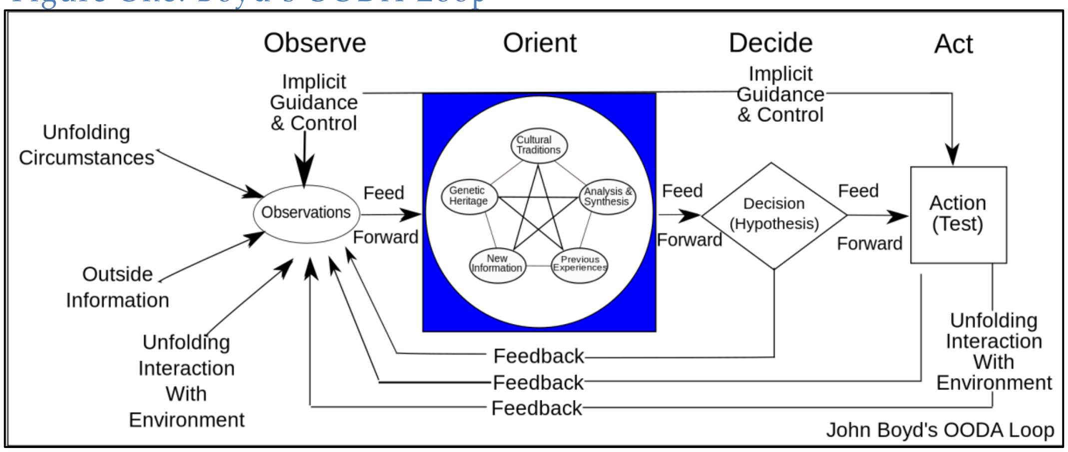

The OODA Loop offers a useful way to understand system design.[12] Boyd's theory of decision-making shows how machine systems like drones operate, make decisions, and interact with the world[13] — all points at which regulations could permit or constrain drone activity.

Consider a drone hovering over a busy highway during rush hour. A state-of-the-art mounted camera would first *observe* the scene below, collecting raw data about individual cars and drivers. The machine must then *orient* itself by processing the information that the camera or other sensors collected. Some data points will be more relevant than others, depending on the drone's programming. For a drone on the lookout for speeding, a car's velocity will be

---

[9] McIntosh, *supra* note 6, at 26.
[10] Parasuraman et al. *supra* note 8, at 288.
[11] Wikipedia, OODA loop, http://upload.wikimedia.org/wikipedia/commons/thumb/3/3a/OODA.Boyd.svg/2000px-OODA.Boyd.svg.png (as of Aug. 23, 2012, 18:00 EST).
[12] *See* Gilles Coppin & François Legras, *Autonomy Spectrum and Performance Perception Issues in Swarm Supervisory Control*, 100 PROCEEDINGS OF THE IEEE 590, 593 (2012).
[13] Sholes, *supra* note 8.

key. Drone-led drunk driving prevention, in contrast, might focus on gauging a car's lateral movements across or within traffic lanes.

Next comes a crucial moment: the drone must *decide* how to act. Should it descend to the road and interact with a car? Or wait and gather more information? Here again, the mission matters. Steps that are sensible in one context may not be suitable in another. Having decided, the drone *acts*: it follows or even pulls over the suspect, or perhaps not.

Today's drones require help from humans in order to complete the OODA Loop. For instance, contemporary drones often lack the sophisticated software necessary for on-board processing of the video they capture. Instead, humans are primarily responsible for interpreting data, considering variables, gauging risks, and deciding how to act. If we were to use drones as traffic cops today, we would need a human in the (virtual) driver's seat.

In the coming years, a machine may not need to communicate with an operator in order to augment machine capabilities with human senses and organic thought processes.

But as the state of the art improves, drones will be able to cycle through the OODA Loop with less and less human assistance. In the coming years, a machine may not need to communicate with an operator in order to augment its capabilities with human senses and organic thought processes. The distance between the machine and its operator will increase. And the technology will begin to feel qualitatively different.

The difference lies in the distinction between "automation" and "autonomy." Automated machines can operate without humans by stringing together rote, pre-programmed operations in sequence.[14] Automated machines "simply replace routine manual processes."[15] Autonomous machines, in contrast, "have the more ambitious goal of emulating human[s]."[16]

Defining Autonomy: Independence, Adaptability, Discretion

There is no bright-line distinction between automation and autonomy. Instead, the shift occurs in degrees.[17] For engineers, the degree to which a machine is autonomous turns on three capabilities: the *frequency of operator interaction* needed in order for the machine to function; the ability of the machine to successfully navigate *environmental uncertainty*; and the machine's *level of assertiveness* as to each one of the operational decisions required in the course of a mission.[18]

The first attribute of autonomy is the frequency with which an operator

[14] O. Grant Clark et. al., *Mind and Autonomy in Engineered Biosystems*, 12 ENG'G APPLICATIONS OF ARTIFICIAL INTELLIGENCE at 10 (1999).
[15] WALT TRUSZKOWSKI ET AL., AUTONOMOUS & AUTONOMIC SYSTEMS: WITH APPLICATIONS TO NASA INTELLIGENT SPACECRAFT OPERATIONS AND EXPLORATION SYSTEMS 10 (2009).
[16] *Id.*
[17] *See* Troy Jones & Mitch Leammukda, *Requirements-Driven Autonomous System Test Design: Building Trusting Relationships*, at 1 (The Charles Stark Draper Laboratory, Inc. 2011).
[18] *Id.* at 2–10.

must interact with the machine – in shorthand, its "independence."[19] Autonomous machines require a human operator's guidance less often than automated ones.[20] A human may still *supervise* the machine, but the human may not necessarily direct every action. A completely autonomous system "would require a single mission statement and it would execute that mission without further assistance."[21]

The second attribute of autonomy is a machine's tolerance of environmental uncertainty – in shorthand, its "adaptability."[22] A machine with high adaptability is able to navigate a wide range of environments, including ones it was not programmed to anticipate. A machine with low tolerance for environmental uncertainty, in contrast, may not be able to operate optimally in novel settings, if it even functions at all.

The current state of the art does not allow drones, such as our hypothetical robotic traffic cop, to achieve complete autonomy. This will change.

The third attribute of autonomy is a machine's level of assertiveness – for short, its "discretion."[23] A machine with a high level of assertiveness may have the ability to change its operating plan in order to complete the assigned mission, with minimal or no outside guidance. Put another way, the gravamen of assertiveness is the system's capacity to independently alter the *means* used to achieve the human-designed *ends*. A system approaching true autonomy may even have the capacity to modify the *ends* it pursues without human intervention.

If we think of automation and autonomy as endpoints on a spectrum,[24] some systems, like welding robots in a Ford Motors plant, clearly are nearer to automation. Other systems might be closer to autonomy — think of a drone able to seek, identify, and target a suspect, all without human commands. Most systems will fall somewhere in between. And indeed, machines might exhibit different levels of autonomy at different stages of the OODA Loop. A drone might be highly autonomous in *observing* its environment and *orienting* itself, but less autonomous at the *decision* and *action* stages.[25]

Technology still constrains where a machine falls on the autonomy spectrum. The current state of the art does not allow drones, such as our hypothetical robotic traffic cop, to achieve complete autonomy.[26] This will change. Policymakers have important opportunities now to think proactively and creatively about machines' future development.[27] A machine's autonomy is

[19] *See id.* at 5.
[20] *Id.*
[21] *Id.*
[22] *See id.* at 2.
[23] *See id.* at 9.
[24] Clark et al., *supra* note 14, at 10.
[25] Parasuraman et al., *supra* note 8, at 289.
[26] *See* Jones & Leammukda, *supra* note 17, at 9.
[27] *Cf.* Ronald C. Arkin, *Governing Legal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture*, TECH. REP. GIT-GVU-07-11 (2009).

within the control of human engineers and operators.[28]

Regulating "The Loop": Structuring Policy for Autonomous Systems

A machine's autonomy is within the control of human engineers and operators. To say that humans are "in the loop" or "out of the loop," full stop, would be too simplistic.

As we can now see, to say that humans are "in the loop" or "out of the loop," full stop, would be too simplistic. Machine systems do not operate in a single loop; rather, they cycle through the loop hundreds or thousands of times in a single mission. Each loop might be addressed to a different task — for instance, the *act* at the loop's culmination might be to follow a speeding car or even to disable it. If we want to regulate the design and operation of increasingly autonomous technologies, we should adopt this more nuanced view of system operations. Tailoring legal interventions accordingly will permit us to control the development and use of technology without crudely and unnecessarily curtailing innovation.

Regulations could, for instance, restrict only certain stages of the OODA to establish rules about the locations or types of activities that a drone is permitted to monitor, or limit the duration of that observation. Depending on our goals, such rules might not necessarily decrease surveillance. The amount and type of information that a machine gathers at the *observe* stage of the OODA Loop affects the machine's capacity to *orient* itself, the number and variety of actions weighed when the machine *decides*, and of course, the eventual *act* chosen and carried out. If the *act* that is contemplated is high risk and our paramount concern is its accuracy, it may be sensible for a longer fact-gathering period to be required.

Regulations could also distinguish between decision-making loops that the machine may carry out without human supervision and those requiring some level of authorization. For example, laws could specify that a machine is permitted to autonomously complete loops where the *act* that results is simply traveling through empty airspace. In contrast, we could differently distribute the authority to act with loops where the *act* requires the drone to engage with humans. Regulators could develop a list of acts for which human intervention is necessary either on an absolute basis — i.e., *any* decision in a given category — or tie the requirement of human permission or veto to other factors, such as the number of alternatives the system evaluated and discarded before settling on a particular act. The difference between pre-approval and veto is more than semantic. Scholarship on default rules suggests that the one we choose may affect the outcome.[29]

Our approach might also vary depending upon the values we seek to promote. For example, a regime primarily concerned with *accountability* might require an affirmative human sign-off of each machine-executed kinetic action, or

[28] Jones & Leammukda, *supra* note 17, at 9.
[29] *Cf.* Cass R. Sunstein, *Switching the Default Rule*, 77 N.Y.U. L. REV. 106 (2002).

perhaps only of particular decisions. An accountability-oriented framework could still permit highly autonomous machine decision-making, so long as a human was held responsible for any machine errors. By contrast, a regime concerned principally with *privacy* might limit the amount of time that a drone may linger or restrict the areas it can observe, much as the Fourth Amendment limits when and where police officers may permissibly search people, things, and places. A comprehensive regulatory framework would incorporate any number of values, including accuracy, accountability, privacy, and more.

The brief sketch we present here shows the importance of better understanding drone technology. It can allow us to craft more thoughtful regulations. And it lets us take initiative today through regulations that also fit the machines of tomorrow. Our choice is not starkly between a society without drones and one with uncontrolled swarms buzzing overhead. Policymakers can limit how and when drones are used. They can shape the human relationship to them. And engineers have already created a vocabulary we can use to navigate these decisions.

As machines evolve, so too will the meaning of having a human "in the loop."

Conclusion

As machines evolve, so too will the meaning of having a human "in the loop." We can keep humans "in the loop" when we want them there, stay "out of the loop" when our attention is more usefully directed elsewhere, and determine how "wide" or free from human intervention the loop should be.

As technology advances, legal, cultural, and political considerations will increasingly act as the primary limits on the capabilities of machine systems. These concerns do impose real boundaries.[30] Many regulatory permutations are possible, for there are many degrees of autonomy. With this insight, we can begin developing smart laws that serve both our security and our values.

Grateful thanks to Kenneth Anderson, Gabriella Blum, Jack Goldsmith, Troy Jones, Jana Schwartz, Benjamin Wittes, and Juan Zarate for helpful comments and conversations.

This essay is adapted from William C. Marra & Sonia K. McNeil, Understanding "The Loop": Autonomy, System Decision-Making, and the Next Generation of War Machines, Lawfare Research Paper Series No. 1-2012 *(May 2012).*

Governance Studies
The Brookings Institution
1775 Massachusetts Ave., NW
Washington, DC 20036
Tel: 202.797.6090
Fax: 202.797.6144
www.brookings.edu/governance.aspx

Editors
Christine Jacobs
Stephanie C. Dahle

Production & Layout
Stefani Jones

Email your comments to gscomments@brookings.edu. This paper is distributed in the expectation that it may elicit useful comments and is subject to subsequent revision. The views expressed in this piece are those of the authors and should not be attributed to the staff, officers or trustees of the Brookings Institution.

[30] Nick Fielding, *U.S. Draws Up Plans for Nuclear Drones*, THE GUARDIAN (Apr. 2, 2012) (describing how political opposition halted plans to develop nuclear-powered drones).