LAWFARE

HARD NATIONAL SECURITY CHOICES

The Lawfare *Research Paper Series*

RESEARCH PAPER NO. 1 – 2012

Understanding "The Loop": Autonomy, System Decision-Making, and the Next Generation of War Machines

- by -

William C. Marra & Sonia K. McNeil
J.D. Candidates, 2012
Harvard Law School

May 2012

Working Draft – Subject to Revision

Suggested Citation: 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), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2043131.

ABSTRACT

This paper responds to persistent confusion in public discussion about "autonomous" technologies. Unmanned aerial systems or "drones" are the focus of an especially intense debate that urgently requires greater clarity. The capabilities of these drones, and their relationship to their human operators, are too often described using imprecise language that unhelpfully complicates important questions about the technology. Should humans stay "in the loop" or "on the loop"? What is "the loop," and what does it mean for humans to be "out" of it? Without a basic understanding of these technologies and a clear vocabulary with which to discuss them, lawyers and policymakers risk grounding foundational decisions about drones and other rapidly advancing, increasingly autonomous systems on inaccurate assumptions. Precision matters, for whether and to what degree technologies should be autonomous will be one of the most important public policy debates of the next generation.

This paper offers a lexicon that lawyers and policymakers can use to analyze sophisticated systems and evaluate potential regulatory frameworks. Part I makes the case for clearer terminology and introduces the concepts of automation and autonomy. Part II first describes the common ground shared by automated and autonomous systems. In almost any machine system, the decision-making sequence is built on four basic stages: "observe," "orient," "decide," and "act." These four stages form the "loop" that is often invoked but rarely unpacked in debates about drones, especially the use of drones as a means of delivering kinetic force. After explaining the "loop," Part II explores distinctions between automation and autonomy in complex systems through the three key attributes of machine autonomy: the frequency of the human operator's interaction with the machine; the machine's ability to function successfully despite environmental uncertainty; and the machine's level of assertiveness, or freedom to select among different courses of action. This investigation suggests that autonomy is better conceptualized as a spectrum subject to calibration by humans than as a technological end state. Part III turns to drones and traces their past development, present state, and future evolution. It concludes that while today's drones are merely automated, future versions may approach true autonomy. Part IV transitions from technology to policy by sketching an answer to the question of what, exactly, is so new about advanced drone technology if today's systems are not markedly more autonomous than older models. Part V explains how this paper's discussion of machine systems and autonomy can be used today by lawyers and policymakers to structure a regulatory regime for tomorrow's advanced drones. Advanced drones represent only one type of complex system, however, and so the vocabulary developed here also has broader applications. Part VI concludes.

ACKNOWLEDGEMENTS

Grateful thanks to Kenneth Anderson, Gabriella Blum, Jack Goldsmith, Troy Jones, Jana Schwartz, Benjamin Wittes, and Juan Zarate for helpful comments and conversations. All views and errors are the authors' own.

AUTHOR CONTACT

The authors welcome comments at will.marra@gmail.com and sonia.k.mcneil@gmail.com.

TABLE OF CONTENTS

I. INTRODUCTION: AUTOMATION AND AUTONOMY

There has been a surge of public and academic commentary in recent years on advanced aerial systems technologies, colloquially dubbed "drones."[1] Contemporary conversations include now-familiar debates over the legality of the technology under international law and on the U.S. military's use of drones to conduct targeted killings.[2] One strand of discussion focuses on whether and to what extent drones leave humans "in the loop," "on the loop," or "out of the loop."[3] In some ways, this discussion is a relatively unstructured effort to identify the differences between advanced drones and the technologies that preceded them. The differences are important, for the extent to which humans remain in, on, or out of "the loop" in the operation of drones may inform our willingness to apply legal and political constraints designed for less advanced machines to increasingly sophisticated systems. However, debates about humans and "the loop" tend to rely heavily on language that is too imprecise to successfully draw out and analyze the differences between drones and predecessor technologies.

For some, drones simply feel different.[4] A sense that something has changed goes only so far, though.[5] Unexplored intuitions cannot fully and should not alone shape policy, particularly with technology. The stakeholders in debates about the appropriate uses of

---

[1] "Drones" are also variously characterized as "unmanned aerial vehicles," "remotely piloted aircraft," and "unmanned aircraft systems," among other names.

[2] *See, e.g.*, Ryan J. Vogel, *Drone Warfare and the Law of Armed Conflict*, 39 DENV. J. INT'L L. & POL'Y 138 (2010); Chris Jenks, *Law from Above: Unmanned Aerial Systems, Use of Force, and the Law of Armed Conflict*, 85 N.D. L. REV. 649, 671 (2009); Richard Murphy & Afsheen John Radsan, *Due Process and Targeted Killing of Terrorists*, 31 CARDOZO L. REV. 405 (2009); Kenneth Anderson, *Targeted Killing in US Counterterrorism Strategy and Law,* SERIES ON COUNTERTERRORISM AND AMERICAN STATUTORY LAW, at 34 (May 2009); Mary Ellen O'Connell, *To Kill or Capture Suspects in the Global War on Terror*, 35 CASE W. RES. J. INT'L L. 325 (2003).

[3] Bernd Debusmann, *More drones, more robots, more wars?*, REUTERS, Jan. 31, 2012. *See also*, Shane Harris, *Out of the Loop: The Human-free Future of Unmanned Aerial Vehicles* (Hoover Institute 2012), *available at* http://media.hoover.org/sites/default/files/documents/EmergingThreats_Harris.pdf; Markus Wagner, *Taking Humans Out of the Loop: Implications for International Humanitarian Law*, 21 J.L. INFO. & SCI. (forthcoming 2011), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1874039.

[4] *See*, *e.g.*, Mary Ellen O'Connell, *Seductive Drones: Learning from a Decade of Lethal Operations*, 19 J.L. INFO. & SCI. EAP 2 (2011), *available at* , http://ssrn.com/abstract=1912635 (asking, "Is it true that it true that existing law regulating manned systems is adequate to regulate killing with UCVs ["unmanned combat vehicles"]? . . . [E]vidence indicates that the availability of UCVs lowers political and psychological barriers to killing").

[5] *Cf.* KENNETH LIEBERTHAL & PETER W. SINGER, CYBERSECURITY AND U.S.-CHINA RELATIONS, THE BROOKINGS INSTITUTION 7-10 (Feb. 2012), http://www.brookings.edu/~/media/Files/rc/papers/2012/0223_cybersecurity_china_us_lieberthal_singer/0223_cybersecurity_china_us_lieberthal_singer_pdf_english.pdf

drones are varied,[6] the technology itself is developing rapidly,[7] and the resolution of questions like these will reverberate on a global scale.[8] The stakes are high. Without a basic understanding of the technology driving key questions and a common vocabulary with which to engage them, domestic and international policymakers risk speaking past one another and causing frustration, if not hostility.

Regulating advanced complex systems, of which aerial drone technologies are just one variant, requires understanding how the human-machine relationship has changed and may yet still shift. A thoughtful grasp of this relationship requires, in turn, an understanding of how the machines have evolved as a technical matter. This paper contributes to discussions about humans and "the loop" by identifying and explaining criteria that lawyers and policymakers can use for drone technology, as well as automated and autonomous systems generally. This paper chooses to focus on technology at the systems level because the drones at the center of many of today's discussions[9] — technological platforms equipped with sometimes-lethal payloads predominantly used in combat theater or conflict areas — are not the only advanced systems meriting policymakers' attention. Autonomous drones could be used domestically for purposes ranging from recreation[10] to law enforcement.[11] It also bears remembering that an autonomous system need not be shaped or operated like a drone. Autonomous systems could provide a variety of services, including healthcare and entertainment.[12] As a result, the vocabulary offered here will be useful beyond its immediate, relatively narrow context.

---

[6] *See*, *e.g.*, Andrew Stobo Sniderman & Mark Hanis, *Drones for Human Rights*, N.Y. Times, Jan. 30, 2012, at A25, *available at* http://www.nytimes.com/2012/01/31/opinion/drones-for-human-rights.html.

[7] *See, e.g.,* PETER W. SINGER, WIRED FOR WAR: THE ROBOTS REVOLUTION AND CONFLICT IN THE 21ST CENTURY 109–22 (2009) [hereafter SINGER, WIRED FOR WAR].

[8] *See*, *e.g.*, W.J. Hennigan, *New Drone Has No Pilot Anywhere, So Who's Accountable?*, L.A. TIMES, Jan. 26, 2012, *available at* http://articles.latimes.com/2012/jan/26/business/la-fi-auto-drone-20120126.

[9] *See*, *e.g.*, John O. Brennan, Assistant to the President for Homeland Security and Counterterrorism, "The Ethics and Efficacy of the President's Counterterrorism Strategy" (Apr. 30, 2012), http://www.wilsoncenter.org/event/the-efficacy-and-ethics-us-counterterrorism-strategy.

[10] For more information on an especially active community of drone enthusiasts, *see generally* DIY DRONES, http://www.diydrones.com/ (last visited May 5, 2012).

[11] *See*, *e.g.*, Brian Bennett, *Police Departments Wait for FAA Clearance to Fly Drones*, L.A. TIMES (Apr. 29, 2012), http://articles.latimes.com/2012/apr/29/nation/la-na-drone-faa-20120430. *See also* Greg McNeal, *A Primer on Domestic Drones: Legal, Policy, and Privacy Implications*, FORBES (Apr. 10, 2012, 8:12PM), http://www.forbes.com/sites/gregorymcneal/2012/04/10/a-primer-on-domestic-drones-and-privacy-implications/; Benjamin Wittes & John Villasenor, *Regulating Domestic Drones on a Deadline*, WASH. POST (Apr. 19, 2012), http://www.washingtonpost.com/opinions/faa-regulation-of-drones-will-challenge-our-privacy-expectations/2012/04/19/gIQA9IH8TT_story.html.

[12] *See* Ronald C. Arkin et al., *Moral Decision Making in Autonomous Systems: Enforcement, Moral Emotions, Dignity, Trust, and Deception*, Proceedings of the IEEE Vol. 100, No. 3., at 1 (Mar. 2012) [hereafter Arkin et al., *Moral Decision Making in Autonomous Systems*], *available at* http://smartech.gatech.edu/jspui/bitstream/1853/40769/1/IEEE-ethicsv17.pdf (observing, "needs [for

Bringing greater terminological rigor to these discussions is challenging in two ways. Not only might the technology itself be difficult to understand, but the words "automation" and "autonomy" — and especially "autonomy" — are plagued by ambiguity. The term "autonomy" is "an attempt to make sense of a tangled net of intuitions, conceptual and empirical issues, and normative claims."[13] It is difficult to agree on a definition even when the word is applied to humans.[14]

In moral and political philosophy, for instance, autonomy has been variously described as "a combination of freedom and responsibility . . . a submission of laws that one has made for oneself,"[15] as the state achieved when one "is acting from principles that we would consent to as free and equal rational beings,"[16] and as a "second-order capacity of persons to reflect critically upon their first-order preferences, desires, wishes, and so forth."[17] Other descriptions of autonomy emphasize perceived or actual authority, rather than a tincture of responsibility, consent, or reflection. "To regard himself as autonomous in the sense I have in mind," observed one philosopher, "a person must see himself as sovereign in deciding what to believe and in weighing competing reasons for action."[18] To another, the concept of autonomy seemed more basic yet. "I am autonomous if I rule me," he wrote, "and no one else rules I."[19] Autonomy is, in short, a slippery concept. "About the only features held constant from one author to another," one thinker despaired, "are that autonomy is a feature of persons and that it is a desirable quality to have."[20]

Notwithstanding the difficulty of pinning down a coherent theory of autonomy, the term plays an important role in the law, including in criminal law and constitutional law. In criminal law, for instance, a defendant's degree of autonomy, in the sense of his or her volition, is in some instances explicitly tied to his or her responsibility for a criminal act.[21] These themes are explored, among other places, in a growing body of literature on the invocation of defenses grounded in diminished autonomy by defendants suffering from Post-

---

advanced robotics and autonomous systems] transcend the warfighting domain and are pervasive, extending to eldercare, robot nannies, and other forms of service and entertainment platforms").

[13] GERALD DWORKIN, THE THEORY AND PRACTICE OF AUTONOMY 7 (1988).

[14] *Id.* at 7.

[15] ROBERT PAUL WOLFF, IN DEFENSE OF ANARCHISM 14 (1970) (describing the theories of Immanuel Kant).

[16] JOHN RAWLS, A THEORY OF JUSTICE 516 (1971).

[17] DWORKIN, *supra* note 13, at 20.

[18] Thomas Scanlon, *A Theory of Freedom of Expression*, 1 PHILOSOPHY AND PUB. AFFAIRS 204, 215 (1972).

[19] Joel Feinberg, *The Idea of a Free Man*, *in* EDUCATION AND THE DEVELOPMENT OF REASON 161 (1972) (R.F. Dearden, ed.).

[20] DWORKIN, *supra* note 13, at 6.

[21] *See, e.g.*, People v. Pettibone, No. 9632C (Sonoma Super. Ct. Cal., Feb. 29, 1980), wherein the defendant successfully asserted an "automation defense." *See also* People v. Lisnow, 151 Cal. Rptr. 621 (1978).

Traumatic Stress Disorder.[22] Autonomy also wends through constitutional law. Various conceptualizations of autonomy appear in and influence the Supreme Court's decisions on controversial and deeply divisive issues, including sex,[23] contraception and abortion,[24] and the existence and scope of a right to die.[25]

Small wonder that confusion pervades discussions about whether an advanced technological system is "autonomous," and what the policy implications of that autonomy might be. The dialogue uses terms imported from other contexts where the application and meaning of the language is equally variable and contested. If these debates could be rebooted and begun anew, perhaps a different and less loaded term of art would be preferable — although, given the tendency to anthropomorphize technology, the choice of words with human baggage to describe non-human systems is not entirely surprising. In any case, the conversation is already well underway, and so this paper takes it on its own terms. Autonomy is no longer solely a feature of humans, but also one of machines. Whether and to what degree it is a desirable quality for systems to have will be one of the most important public policy debates of the next generation.

II. MACHINE FUNCTIONING AND THE DIFFERENCE BETWEEN AUTONOMY AND AUTOMATION

A machine system's degree of autonomy is a function of the system's ability to operate and make decisions independent of a human operator. Accordingly, to understand autonomy and how it differs from automation, one must first explore how machine decision-making processes operate. The "OODA Loop"[26] is a particularly effective tool for

---

[22] *See*, *e.g.*, Thomas L. Hafemeister & Nicole A. Stockey, *Last Stand? The Criminal Responsibility of War Veterans Returning from Iraq and Afghanistan with Posttraumatic Stress Disorder*, 85 IND. L.J. 87, 132–34 (2010); Christopher Hawthorne, *Bringing Baghdad into the Courtroom: Should Combat Trauma in Veterans Be Part of the Criminal Justice Equation?*, 24 CRIM. JUST. 4 (2009); Comment, *PTSD: Effective Representation of a Vietnam Veteran in the Criminal Justice System*, 68 MARQUETTE L. REV. 647 (1985).

[23] *See*, *e.g.*, Lawrence v. Texas, 539 U.S. 558, 562 (2003), where the Court declared, "Liberty presumes an autonomy of self that includes freedom of thought, belief, expression, and certain intimate conduct. The instant case involves liberty of the person both in its spatial and in its more transcendent dimensions."

[24] *See*, *e.g.*, Gonzales v. Carhart, 550 U.S. 124, 170 (2007) (Ginsburg, J., dissenting) ("In reaffirming *Roe*, the *Casey* Court described the centrality of 'the decision whether to bear . . . a child,' to a woman's 'dignity and autonomy,' her 'personhood' and 'destiny,' her 'conception of . . . her place in society.'") (citing Eisenstadt v. Baird, 405 U.S. 438, 453 (1972)) (internal citations omitted)).

[25] *See*, *e.g.*, Washington v. Glucksberg, 521 U.S. 702, 726 (1997) (reversing the Court of Appeals, which had concluded, "Like the decision of whether or not to have an abortion, the decision how and when to die is one of the most intimate and personal choices a person may make in a lifetime, a choice central to personal dignity and autonomy." Compassion in Dying v. Washington, 79 F. 3d 790, 813-14 (9th Cir. 1996)).

[26] For detailed discussions of John Boyd's development of the OODA Loop, see generally FRANS P.B. OSINGA, SCIENCE, STRATEGY AND WAR: THE STRATEGIC THEORY OF JOHN BOYD (2006); GRANT T.

understanding complex systems, including aerial drones carrying lethal payloads, for it offers a language shared by engineers and the military. For example, an Air Force strategy document on the future uses of unmanned aerial vehicles (UAVs) states that "[o]ne of the most important elements" driving increased reliance on UAVs "is the potential for [UAVs] to rapidly compress the observe, orient, decide, and act (OODA) loop."[27] The terminology has entered public discussion as well. When commentators debate whether drone warfare leaves humans "in the loop," "out of the loop," or simply "on the loop," it is the OODA loop they are talking about.[28]

A. *Understanding "the Loop": OODA and How Machines Work*

Why did American F-18 fighter planes get the better of Soviet MiG-5 jets during the Korean War? Air Force pilot and military strategist John Boyd's answer to this question transformed how the U.S. military defines the key to winning in battle.[29] Boyd's OODA loop modeled human decision-making using a four-step process: *O*bserve, *O*rient, *D*ecide, *A*ct.[30] Under this four-step process, an individual first *observes* the world around her, gathering refined information about her environment through the array of nuanced senses with which human are equipped.[31] Second, she *orients*, or interprets the information she has gathered and converts it into knowledge upon which she can act.[32] Third, she weighs the potential courses of action based on the knowledge she has accumulated and *decides* how to act.[33] Fourth and finally, she *acts*, or executes the decision she has made.[34]

Boyd's insight was that in a dogfight, the advantage lay with the fighter pilot whose OODA loop was faster and more accurate than his opponent's, and who was able to throw his opponent's loop out of sync.[35] Boyd's elegant theory is still used by the military today.[36] It

---

HAMMOND, THE MIND OF WAR: JOHN BOYD AND AMERICAN SECURITY (2004); ROBERT CORAM, BOYD: THE FIGHTER PILOT WHO CHANGED THE ART OF WAR (2002).

[27] United States Air Force, *Unmanned Aircraft Systems Flight Plan, 2009–2047*, at 16 (2009), *available at* http://www.fas.org/irp/program/collect/uas_2009.pdf [hereafter Air Force Flight Plan].

[28] See Harris, *supra* note 3; Wagner, *supra* note 3. The Air Force Flight Plan projects that "[i]ncreasingly humans will not longer be 'in the loop' but rather 'on the loop'—monitoring the execution of certain decisions." Air Force Flight Plan, *supra* note 27, at 41.

[29] 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), *available at* http://www.dodccrp.org/events/10th_ICCRTS/CD/papers/365.pdf; Scott E. McIntosh, *The Wingman-Philosopher of MiG Alley: John Boyd and the OODA Loop*, 58 AIR POWER HIST. 24, 26–28 (2011).

[30] CORAM, *supra* note 26, at 334.

[31] Brehmer, *supra* note 29, at 2; McIntosh, *supra* note 29, at 27.

[32] Brehmer, *supra* note 29, at 2; McIntosh, *supra* note 29, at 27.

[33] Brehmer, *supra* note 29, at 2; McIntosh, *supra* note 29, at 27.

[34] Brehmer, *supra* note 29, at 2; McIntosh, *supra* note 29, at 27.

[35] McIntosh, *supra* note 29, at 26–27.

[36] For example, the Marine Corps' *Warfighting* manual discusses the OODA Loop and endorses Boyd's view that "the party who consistently completes the cycle faster gains an advantage that

has gained purchase in other fields too, including in business and sports — essentially, “anywhere a competitor seeks an edge.”[37] Even engineers have borrowed the concept to illustrate how machine systems operate, make decisions, and interact with the world.[38] For example, Thomas Sheridan, an engineer and a leading scholar of autonomy and robotics, suggests a four-stage information-processing model that is functionally identical to the OODA loop. Sheridan’s model, which is also designed to mirror human information processing, tracks OODA: (1) Information Acquisition; (2) Information Analysis; (3) Decision Selection; and (4) Action Implementation.[39]

The OODA loop is not without flaws. Even its proponents admit that the four-stage model is a “gross oversimplication” of both human and robot information processing, in part because the four stages often overlap temporally.[40] Indeed, the “loop” is not a clean linear process, for it includes constant feedback and integration among the different stages. When illustrated with all its complexity, the OODA loop appears as follows:

FIGURE 1 — Boyd’s OODA Loop[41]

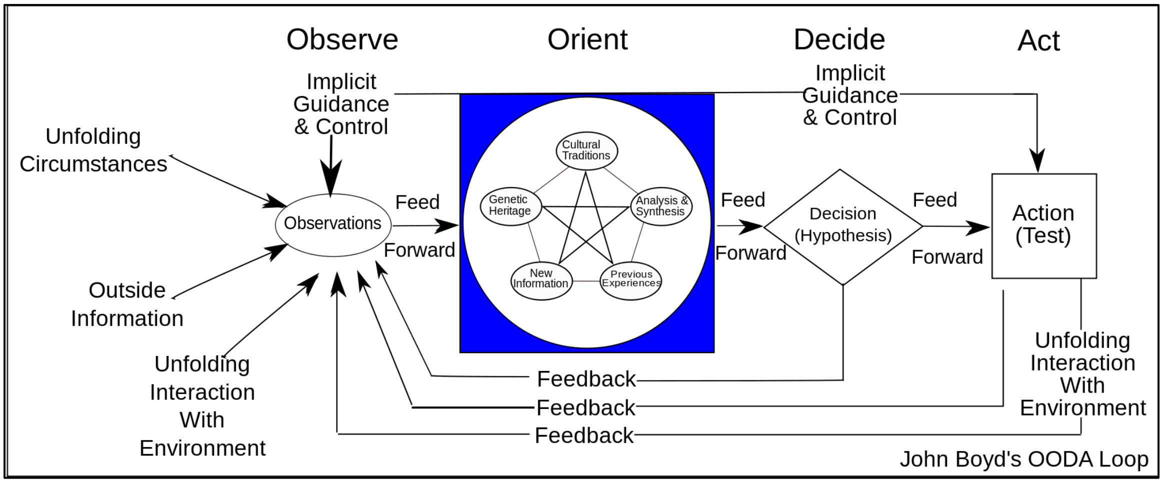

---

increases with each cycle.” See U.S. Marine Corps, *Warfighting*, at 40 n.18 (1997), *available at* http://www.dtic.mil/doctrine/jel/service_pubs/mcdp1.pdf.

[37] McIntosh, *supra* note 29, at 26.

[38] See, e.g., Eric Sholes, *Evolution of a UAV Autonomy Classification Taxonomy*, at 1 (remarks at the 2007 IEEE Aerospace Conference), *available at* http://ieeexplore.ieee.org/stamp/stamp.jsp?tp =&arnumber=4161585.

[39] Gilles Coppin & François Legras, *Autonomy Spectrum and Performance Perception Issues in Swarm Supervisory Control*, 100 PROCEEDINGS OF THE IEEE 590, 592 (2012), *available at* http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=06108329; Raja Parasuraman, Thomas B. Sheridan, & Christopher D. Wickens, *A Model for Types and Levels of Human Interaction with Automation*, 30 IEEE Transactions on Systems, Man, and Cybernetics Part A Systems and Humans 286, 288 (2000), *available at* http://ieeexplore.ieee.org/stamp/stamp.jsp?tp= &arnumber=844354.

[40] Parasuraman et al. *supra* note 39, at 288.

[41] John Boyd’s OODA Loop, WIKIPEDIA, *available at* http://upload.wikimedia.org/wikipedia/ commons/thumb/3/3a/OODA.Boyd.svg/2000px-OODA.Boyd.svg.png.

Despite its limitations, the OODA loop provides a useful lens for understanding system design. The loop is a powerful conceptual device for engineers at the design stage.[42] It also permits a relatively straightforward comparison of systems based upon their technological capabilities.[43] Thus, before one can understand the distinction between autonomy and automation, and before it is possible to know whether a human is inside, outside, or on "the loop," much less whether and how the "loop" matters at all, it is necessary to understand the OODA loop.

To see how the OODA loop works in practice, begin with a human being — call him Dave[44] — who is walking down a road and encounters a large boulder blocking his path. Dave wants to make his way past the boulder and continue on his walk. Described using the OODA loop framework, Dave first *observes* his surroundings, using all his senses to take in the world around him. He sees the boulder with his eyes and ascertains its height (can he step or jump over the boulder?), gauges its density (can he push it aside?), and observes the path on either side of the boulder. He also uses his senses of hearing, smell, touch, balance, temperature, and perhaps even taste, to absorb as much information about his environment as possible. Dave then *orients* himself by synthesizing the information he has observed and begins to convert it into knowledge upon which he can act. He may determine, for instance, that the boulder is too heavy to push aside and too tall to scale, but note that there is open space to the left of the boulder. Next, Dave *decides*. He weighs his options: he could retreat and go back, or he could attempt to scale the boulder, albeit not without risk of physical injury. After reflecting, Dave decides to take the open path to the left of the boulder. Finally, Dave *acts.* By moving his legs and walking to the left, past the boulder, Dave is able to continue down the path.

This description of Dave's encounter with the boulder may seem labored and didactic, but that is precisely the point. These tasks are complex, and require many mental and physical processes. Under normal conditions, humans may be able to perform the entire OODA loop subconsciously and, depending on the act and its context, almost instantaneously. If we now consider how a machine might interact with the same boulder, we can begin to understand the many technological processes the machine must be able to execute in order to get past the rock in the road.

Hal, the machine, must first *observe* its surroundings. This simple task is apt to be comparatively far more difficult for Hal than for Dave. Hal must have some technological capability to sense the world around it, including the ability to identify and classify objects according to their mass and density, their height, width and length, for instance. Hal needs more than the raw capability to sense what is before it: it also needs some scope of vision, so that it can identify not only the boulder immediately in front of it, but also the land on either side of the rock. Without the capacity to collect this information and to observe its

---

[42] See Coppin & Legras, *supra* note 39, at 593.

[43] *Id.* at 593.

[44] *See* 2001: A SPACE ODYSSEY (MGM 1968) (Dave Bowman: "Open the pod bay doors, HAL." Hal: "I'm sorry, Dave. I'm afraid I can't do that.").

environment in detail, Hal may not be able accurately to understand possible paths around the boulder. Observational deficiencies at the first stage in the OODA loop will curtail the options available to Hal at later stages in the process. Depending on the characteristics of the environment and how quickly it changes, the process of sensing and developing complete observations may be more or less complex. Incomplete or inaccurate observations increase the risk that Hal will err in the course of achieving its objective and potentially doom its mission. As the saying goes, "garbage in, garbage out" — or, to use the OODA loop, faulty information at the *observe* stage affects the rest of the process.

Hal must then attempt to *orient* itself, or transform its observations into actionable knowledge. To do this, Hal is required to weave together many independent data points into a coherent portrait of the world around him. Here, processing speed is also a factor. While Dave may be able to orient himself and draw conclusions based on his observations without breaking his gait, Hal may not be able to absorb and incorporate all the information flooding its processors from its sensors at once. Hal may have to slow down or even come to a complete stop while it processes the information and orients itself; if he is too slow, he may simply crash into the boulder.

Next, Hal must *decide* which action to take. This is perhaps the most difficult stage of the OODA loop for a machine to execute. The complexity of Hal's decision-making ability will depend on its technological sophistication. Hal's decision-making ability could be very basic and rudimentary. For instance, Hal could be programmed such that if it identifies a roadblock ahead, it simply stops in its tracks and calls its human operator for help. One can also easily imagine a more complex and varied menu of options, however. For example, upon identifying a roadblock, Hal could be programmed to turn left 90 degrees, advance five feet, and then turn right. It is also possible to structure decisions in sequence. If, after executing the pre-programmed maneuver prompted by the identification of an obstacle, Hal determines that there are no more roadblocks, its instruction might be to move forward. In contrast, if the roadblock persists, Hal's programming might dictate that it move five more feet to the left, turn right 90 degrees, and so on, until the machine's sensors show that the path ahead is clear. Of course, identifying possible actions is only half the battle. The machine must still make the *best* decision in light of the information available, probably taking into account a range of factors that include safety, speed, and efficiency. This decision does not necessarily have to occur without any input from a human operator, of course. The more rudimentary the technology, or the more complex the decision, the more appropriate it might be for Hal to request operator assistance.

Having made its decision, Hal must then *act*. As with the three prior stages of the OODA loop, Hal's range of possible actions is limited by its technological design, just as Dave's abilities are limited by his biological design. Dave is not likely to be able to fly over the boulder, for instance — though perhaps Hal, depending on its design, could. If Hal lacks arm-like appendages, moreover, it cannot move the boulder. If Hal is a land-bound rover, it cannot jump over the boulder. Even Hal's ability to go around the boulder will be limited if it is not equipped to traverse the terrain on either side of the boulder successfully.

These encounters with Hal and Dave show that a machine's capacities may vary significantly at each stage of the OODA loop. As the discussion suggests, a machine's ability to perform all four stages of the OODA loop — and to go through those stages with speed and accuracy — depends entirely upon the technology built into it. Depending on the state of the art, Hal may be better at observing than deciding, or better at acting than orienting.

The more Hal can achieve by itself through automated processes, the less the machine may need to communicate with an operator for instructions or recommendations in order to augment the capabilities it lacks with human senses, organic thought processes, analytical evaluation, or action. As the distance between the machine and its operator increases and the amount of interaction between the human and robot decreases, however, even action that is ultimately only the result of layers of complex and refined processes operating in tandem with one another — highly advanced automation — may begin to look qualitatively different. It is most often at this stage that the question of whether the machine is "autonomous" is raised.

B. Towards a Distinction Between "Autonomy" and "Automation"

Whatever might be said about humans, all machines are not created equal. Compare, for instance, the 1983 Apple Lisa to the contemporary MacBook; or contrast a standard vacuum cleaner with Rosie, the Jetsons' zippy robotic maid; or weigh the capabilities of a Cold War era U-2 spy plane against those of the surveillance and combat drones of today and tomorrow.[45] In each example, it is evident that the latter differs meaningfully from the former, including in speed and versatility. It is less clear, however, whether the distinction represents a quantum leap rather than a linear advance. Debates about advanced aerial systems reflect an intuition that these machines are qualitatively different by ascribing to them, largely without careful evaluation, the characteristic of "autonomy." The best way to actually understand autonomy, however, and thus to begin to evaluate whether and how these machines actually are different, is to study the distinction between "automation" and "autonomy."

Engineers evaluate a machine's level of autonomy by measuring its level of dependence on humans while executing the OODA loop — in short, the greater the machine's ability to observe, orient, decide, and act on its own, the greater its autonomy.[46] Yet while the terms "automation" and "autonomy" are alike in that both "refer to processes that may be executed independently from start to finish without any human intervention,"[47] the processes' differences are more revealing than their similarities. Automated systems are not self-directed. They also lack decision-making capability. Instead, these systems simply have "the

---

[45] See *infra* Part III.

[46] See SINGER, WIRED FOR WAR, *supra* note 7, at 74 ("That a machine can make a decision on its own, with or without a human in the loop, does not define whether or not it is a robot. The relative independence of a robot is merely a feature called 'autonomy.'").

[47] WALT TRUSZKOWSKI ET AL., AUTONOMOUS AND AUTONOMIC SYSTEMS: WITH APPLICATIONS TO NASA INTELLIGENT SPACECRAFT OPERATIONS AND EXPLORATION SYSTEMS at 10 (2009), *proof available at* http://eprints.ulster.ac.uk/2581/1/book-last-preproof.pdf.

capacity to operate without [human intervention]."[48] By contrast, autonomous entities are capable of being "independent in the establishment and pursuit of their own goals."[49] In sum, "[a]utomated processes simply replace routine manual processes with software/hardware ones that follow a step-by-step sequence that may still include human participation. Autonomous processes, on the other hand, have the more ambitious goal of emulating human processes rather than simply replacing them."[50]

Automation can be understood as both a precursor to and a crucial component of autonomy. Engineers have defined autonomous systems as having "independence of comportment," which is possible when it "can be described as possessing, to some degree, several defining characteristics. The first is *automation*: the capacity to operate without outside intervention. Although necessary, this alone is insufficient for significant autonomy."[51] Autonomy also requires some decision-making agency, captured by the two additional characteristics of *volition*, or "choice in action or thought," and *intent*, or deliberate "pursuit of goals."[52] Truly autonomous machines may also be able to actually learn, meaning they can draw conclusions based on past experience and incorporate these lessons into the machines' future conduct, enabling the choice of a greater variety of both means and ends than the machines were originally programmed to consider.[53]

This baseline distinction between automation and autonomy provides a useful starting point, but it is still incomplete.[54] Autonomy is a complex concept with many components that cannot be captured simply through a distinction based on decision-making independence. A capsule definition is too simple because there is no bright-line distinction between automated and autonomous technologies, and one may properly speak of technologies that exhibit different degrees of autonomy. The better approach, therefore, is to define autonomy in terms

---

[48] O. Grant Clark et. al., *Mind and Autonomy in Engineered Biosystems*, 12 ENGINEERING APPLICATIONS OF ARTIFICIAL INTELLIGENCE at 10 (1999), *available at* http://www.sciencedirect.com/science/article/pii/S095219769900010X.

[49] *Id.* at 2.

[50] TRUSZKOWSKI ET AL., *supra* note 47, at 10.

[51] Clark et al., *supra* note 48, at 10.

[52] *Id.* at 10. The technical definitions provide here track dictionary definitions of autonomy. "Automation" is defined as "[t]he technique of making an apparatus, a process, or a system operate automatically," or "[a]utomatically controlled operation of an apparatus, process, or system by mechanical or electronic devices that take the place of human labor." MERRIAM-WEBSTER ONLINE DICTIONARY, http://www.merriam-webster.com/dictionary/automation. "Autonomy" is defined as "[t]he quality or state of being self-governing; especially: the right of self-government." *Id.* at http://www.merriam-webster.com/dictionary/autonomy.

[53] See SINGER, WIRED FOR WAR, *supra* note 7, at 75–77.

[54] See TROY JONES & MITCH LEAMMUKDA, REQUIREMENTS-DRIVEN AUTONOMOUS SYSTEM TEST DESIGN: BUILDING TRUSTING RELATIONSHIPS 1 (The Charles Stark Draper Laboratory, Inc. 2011), *available at* http://www.itea-wsmr.org/ITEA%20Papers%20%20Presentations/2010%20ITEA%20Papers%20and%20Presentations/itea_lvcc_2010_uast_track2_draper_jones_leammukda_paper.pdf (noting that "[t]here are as many definitions of 'autonomous systems' as there are papers that define it," but arguing that all of these definitions are "incomplete").

of "common sets of traits that can be specified for any automated/autonomous/ robotic/intelligent system."[55]

To understand autonomy, we must identify and unpack those traits. A machine's capabilities across three dimensions define the degree to which a machine is merely automated or truly autonomous: (1) the *frequency of operator interaction* that the machine requires in order to function; (2) the machine's ability to function successfully notwithstanding *environmental uncertainty*; and (3) the machine's *level of assertiveness* as to each one of the various operational decisions that structure and enable the machine to complete its mission, which may be dictated by its operator or generated by the machine itself.[56] If we can understand these three attributes of autonomy, we can begin to understand what it means for a machine to be autonomous rather than just highly automated.

1. The first attribute of autonomy: Frequency of operator interaction

The first attribute of autonomy is frequency of operator interaction. An anthropomorphic shorthand for this attribute of autonomy is "independence." To oversimplify, autonomous machines require a lower frequency of operator interaction than automated machines. In an extreme case, interaction could be limited merely to the operator pressing "go" and leaving the machine to execute the entire mission without further human supervision or intervention. As engineers have explained, "an entirely autonomous system would require a single mission statement and it would execute that mission without further assistance."[57] The more frequently the operator must interact with the machine and give it commands, the less autonomous and more merely automated the system becomes, until it finally "degenerate[s] into an entirely remote-controlled system" where "[t]he human operator is constantly in contact with the system providing direct commands to accomplish the task."[58]

To better understand this attribute, revisit Hal and the boulder. Assume that Hal's goal (as established by its human operator) is to travel one mile down the road, and that he encounters the boulder mid-way through his trek. If Hal has the capability to autonomously complete the mission, its operator or supervisor does not need to take any additional action to help Hal get around the boulder. Hal would have the capability to independently observe the environment, orient itself, decide to go around the boulder's left flank, and execute the movement. At the other extreme, imagine Hal is a remote controlled robot, like a remote-controlled car or hobby plane. In this scenario, Hal requires continuous interaction with a human to move down the path. If the human stops giving commands to Hal, the machine will be unable to proceed.

Naturally, there is a broad middle ground between true autonomy and complete automation. Perhaps Hal is able to independently proceed straight down the road, but if it

---

[55] *Id.* at 2.
[56] *Id.* at 2–10.
[57] *Id.* at 5.
[58] *Id.*

encounters any potential obstacles, it is programmed to stop in its tracks and request operator permission to change course and turn left to get around the boulder. Or perhaps the human operator has veto power over Hal's decisions, such that the human is *permitted* but not *required* to interact with Hal. Moreover, Hal's dependence upon a human operator might vary across each of the four stages of the OODA loop. For example, Hal might have the capacity to observe its environment and orient itself based on the information it gathers without the need for any human interaction, but it might still lack the technical capacity to form and execute a decision. All of these configurations turn on the frequency of operator interaction. Each one represents an intermediate position between complete automation and true autonomy.

2. The second attribute of autonomy: Environmental uncertainty

The second attribute of autonomy is a machine's ability to deal successfully with varying degrees of environmental uncertainty. The critical issues here include the system's capability to detect and avoid collisions with objects in its environment, as well as the range of environmental obstacles the system is capable of confronting.[59] An anthropomorphic shorthand for the concept of environmental uncertainty is "perception": the more the machine is able to perceive, the better it can navigate its environment, and the more it can operate independently of a human controller.[60] A machine with high adaptability to environmental uncertainty is able to accommodate and navigate within a wide range of environments, including environments not previously encountered in a laboratory setting or otherwise anticipated by its programming. Flexibility tends to make the machine more autonomous. In contrast, a machine with low tolerance for environmental uncertainty may be able to operate optimally only in the narrow band of environments it was programmed to anticipate, or equipped to sense. It will have substantially slower reaction times in novel settings, assuming it can respond at all.

Return again to Hal and the boulder. Hal's ability to function successfully in its environment may vary substantially based on its capacity to adapt to environmental uncertainty. At one extreme, a robot with a very low tolerance for environmental uncertainty might have no "eyes" or other sensory perceptions whatsoever. Configured like so, Hal would move down the road essentially blind, unable even to tell that there is a boulder before it, much less to decide to avoid it. In this example, Hal will crash into the boulder without human interaction. At the other extreme, Hal could be equipped with sophisticated perceptive sensors and advanced, highly flexible processing capabilities that enable it to navigate through significant environmental uncertainty. Together, these abilities might permit Hal to observe and map out the terrain around the boulder without regard to whether Hal's programmer had anticipated a specific need for Hal to navigate inert obstacles of the boulder's specific size and mass. The ability to understand the terrain and independently to adapt to environmental uncertainty is critical to Hal's capability to act autonomously. If Hal cannot tell whether there is a boulder ahead but Hal is used in an area where boulders may be present, a human will be

---

[59] *Id.* at 2.
[60] *Id.*

required to steer it.

As with frequency of operator interaction, Hal's capacity to navigate in uncertain environments may vary along the different stages of the OODA loop. Hal's ability to manage different and potentially unfamiliar environments is clearly tied to its ability to *observe* its environment, or to perceive different objects in the world. However, environmental uncertainty also matters at other stages in the OODA loop, too. For example, if Hal has a low tolerance for environmental uncertainty, it may be unable to recognize that previously known objects have moved since its last visit, degrading or destroying its ability to orient itself to a room. Similarly, if Hal cannot identify the objects around him with specificity, it might be unable to identify the full range of possible decisions available to it. Depending on its adaptability to environmental uncertainty, a machine may be more or less truly autonomous.

3. The third attribute of autonomy: Level of assertiveness

The third attribute of autonomy is a machine's level of assertiveness. "Level of assertiveness" refers to the extent to which a system is able to change its operating plan in order to complete its assigned mission without guidance from a human operator.[61] An anthropomorphic shorthand for assertiveness could be "discretion." The gravamen of assertiveness is the system's ability independently to alter the *means* it uses to achieve the human-designed *ends*. A system approaching true autonomy may even have the authority to alter the *ends* that it will pursue without human real-time authorization or intervention.

When Hal encounters the boulder, there are many different ways it could react. In a scenario where the machine has a high level of assertiveness, Hal could be instructed simply to reach a destination a specified distance away, and be provided with both the capability and the authority to navigate around or eliminate any obstacles that might come into its path. On the other hand, Hal could have a very low level of assertiveness. In this scenario, Hal's programmers may instruct it to travel forward on the road but stipulate that if the robot encounters any obstacles, it must stop immediately in its tracks and await further instruction before either proceeding or taking any other action. As these two extremes show, the higher Hal's level of assertiveness, the more vigorous and potentially creative its attempts to complete its mission in the face of obstacles or hindrances will be.

Hal's level of assertiveness determines its ability to progress from one stage of the OODA loop to another independently. A machine's level of assertiveness affects its ability and authority to *decide* how to act. The breadth of Hal's freedom to choose among potential courses of action affects the degree to which Hal may alter the mission plan, as defined by its programmer or operator, or even the mission objective itself. Notice here the link between assertiveness and risk. A machine's assertiveness is most implicated when the machine becomes "stuck" — that is to say, for instance, when its "decide" process reports that no movements are possible. An assertive machine may be more successful in finding ways to get

---

[61] *Id.* at 9.

"unstuck," perhaps by taking a risk and executing an action that might not have as high probability of success. For example, imagine Hal concludes that the only way past the rock is to jump over it, but that Hal's leap has only an 80% chance of success and a 20% likelihood of crashing and causing damage to itself or its surroundings. Imagine further that under normal circumstances, Hal will only evaluate actions that have a 90% likelihood of success at the *decide* stage in the OODA loop. If Hal is unassertive, it will simply stop in its tracks, retreat, or call for operator assistance. But a more assertive machine, determining that the only way to complete the mission is to jump over the rock, might lower its threshold of success for this particular decision and attempt to leap over the boulder. The higher level of assertiveness, the closer the machine draws to true autonomy.

C. The Autonomy Spectrum

By now, it should be clear that there is no bright line between "automated" and "autonomous" machines. Instead, autonomy is a function of the three variables described above. A system is autonomous when it acts with sufficiently infrequent operator interaction, is able to successfully function in uncertain or unfamiliar environments, and achieves mission objectives with a high level of assertiveness. Still, though, the qualifiers that necessarily are attached in the previous sentence to each one of the three attributes of autonomy make it clear that there is no line in the sand between autonomy and automation. As one group of robotics engineers argue, "[l]ike intelligence and consciousness, autonomy should be measured on a continuous scale,"[62] or on a spectrum.

Some systems lie clearly on the "automated" end of the spectrum, like a robot welding doors in a Ford Motors plant. Other systems might be closer to autonomous — think of a futuristic drone able to seek, identify, and kill an enemy without any human interaction. Most systems, however, will fall somewhere between these two extremes, just as with the many different versions of Hal discussed above, each of which could be placed on a spectrum according to its level of autonomy. In some versions, Hal is operated like a tool. In others, the robot comes closer to resembling a teammate.[63] In other versions, it is somewhere in between.

Thomas Sheridan, the engineer and scholar introduced above, has developed a 10-level spectrum of autonomy usefully illustrating both the absence of a clear line dividing automated from autonomous systems, as well as the difficulty of defining precisely when a system moves from automated to autonomous.[64] At Level 1, the system is not autonomous,

---

[62] Clark et al., *supra* note 48, at 10.

[63] *Cf.* Johnson et al., *The Fundamental Principle of Coactive Design: Interdependence Must Shape Autonomy* 172, *in* COORDINATION, ORGANIZATIONS, INSTITUTIONS, AND NORMS IN AGENT SYSTEMS VI, Vol. 6541 (2011), *available at* http://www.springerlink.com/content/a54371r947t2n453/; *see also* Bradshaw et al., *From Tools to Teammates: Joint Activity in Human-Agent-Robot Teams*, *in* HUMAN CENTERED DESIGN, Vol. 5619 (2009), *available at* http://www.springerlink.com/content/l12348134205684h/.

[64] Parasuraman et al., *supra* note 39, at 287.

while at Level 10, it is fully autonomous. At any other level, the system is between automation and autonomy. As is shown below, Levels 2 through 4 concern the allocation of decision-making authority between the human and machine. By contrast, Levels 5 through 9 grant the initial decision-making authority to the machine, and give the human operator varying levels of approval or veto authority over the machine's initial decision.[65]

TABLE 1 — Sheridan's 10 Levels of Autonomy[66]

Level	Description
1	The computer offers no assistance, human must do it all.
2	The computer offers a complete set of action alternatives, and
3	Narrows the selection down to a few, or
4	Suggests one, and
5	Executes that suggestion if the human approves, or
6	Allows the human a restricted time to veto before automatic execution, or
7	Executes automatically, then necessarily informs the human,
8	Informs the human after execution only if the human asks, or
9	Informs the human after execution if it, the computer, decides to do so.
10	The computer decides everything and acts autonomously, ignoring the human completely.

Applying Sheridan's 10-level spectrum requires attention to one important wrinkle: a machine's autonomy can vary along each stage of the OODA loop. Thus, a machine might exhibit greater autonomy in observing its environment and orienting itself, but require greater dependence upon humans at the decision and action stages.[67] This phenomenon was illustrated by Hal's variable performance in each of the three attributes of autonomy in Part II.B.

When the Air Force Research Lab (AFRL) released its own 11-level autonomy spectrum, it took this important fact into account.[68] The Air Force's spectrum is particularly helpful because it describes a system's ability to observe, orient, decide, and act at each level of autonomy. Notably, the Air Force spectrum was developed in part to further an understanding of autonomy in situations where a single human operator controls multiple unmanned air or ground vehicles.[69] It is therefore addressed in part to the autonomy of multiple-machine systems.[70]

---

[65] Coppin & Legras, *supra* note 39, at 592.
[66] See Parasuraman et al., *supra* note 39, at 287.
[67] Parasuraman et al., *supra* note 39, at 289; Sholes, *supra* note 38, at 1.
[68] Sholes, *supra* note 38, at 2–3.
[69] *Id.* at 2.
[70] *Id.*

TABLE 2 — AFRL's 11 Levels of Autonomy[71]

Level	Level Description
0	Remotely piloted vehicle
1	Execute pre-planned mission remotely
2	Changeable mission
3	Robust response to real time faults/events
4	Fault/event adaptive vehicle
5	Real time multi-vehicle coordination
6	Real time multi-vehicle cooperation
7	Battlespace knowledge
8	Battlespace single cognizance
9	Battlespace swarm cognizance
10	Fully autonomous

The AFRL spectrum describes the levels of autonomy at each stage of the OODA loop. A small sampling of the chart is sufficient to understand its value:

TABLE 3: Autonomy Spectrum and the OODA Loop[72]

Level	Observe	Orient	Decide	Act
0	Flight control sensing and on-board camera.	Telemetered data; remote pilot commands.	None. Off-board pilot.	Control by remote pilot.
5	Local sensors to detect external targets, fused with off-board data.	Group action diagnosis and resource management.	On-board trajectory planning; optimize for current & predicted conditions; collision avoidance.	Group accomplishment of tactical plan as externally assigned; air collision avoidance.
10	Cognizant of all within the battlespace.	Coordinates as necessary.	Capable of total independence.	Requires little guidance of any sort.

---

[71] *Id.* at 3.
[72] *Id.*

As this table shows, a system can effectively mix and match autonomy and automation, achieving different autonomy levels depending upon the OODA loop stage in which the system is operating. For example, a machine might operate at Level 10 at the *observe* stage and thus have the ability to be cognizant of all objects in its environment. However, it might only be at a level 5 at the *decide* stage, able to avoid collisions with objects in its environment but still requiring human assistance to realize more substantial objectives.[73]

D. Humans Control the Autonomy Spectrum

Technology constrains where a machine falls along the autonomy spectrum. Today, technology is not sufficiently advanced to create complex machines capable of operating at higher levels of autonomy along the entire OODA loop. But even if the technology exists, the engineer can still control a machine’s level of autonomy through the initial coding. Just because a machine *can* be programmed to act autonomously does not mean it must be programmed that way. Thus, a system’s autonomy across the entire OODA loop, and across all three attributes of autonomy — frequency of operator interaction, environmental uncertainty, and level of assertiveness — is “entirely controllable by the customer and the development team.”[74] Cultural norms, politics, and civilian and military demand, in additional to the state of the art, play an important role in the design of our machine systems.[75]

Engineers can control a machine’s autonomy in several ways. First, the engineer can control *which functions* are automated as opposed to autonomous, giving the machine the capability to perform only certain tasks autonomously. For example, an unmanned aerial drone may have autonomous control over its flight path — where and when it should fly. But, it might only be automated when it comes to firing a missile at an enemy — meaning that its human operator would retain absolute control of when and at whom to fire.

Second, the engineer can control *when* a machine is automated rather than autonomous. In other words, the engineer can allow the human operator to toggle manually or automatically between different levels of autonomy at different times, depending on the mission.[76] This allows “systems to adapt to users and contexts, especially in terms of dynamically tuning the allocation of tasks — quantitatively and qualitatively — between operators and machines.”[77]

America’s war drones today exhibit characteristics both of automation and autonomy,

---

[73] *Id.* at 3–5.

[74] JONES & LEAMMUKDA, *supra* note 54, at 9.

[75] See *id.* at 287; *see also* BILL YENNE, BIRDS OF PREY: PREDATORS, REAPERS AND AMERICA’S NEWEST UAVS IN COMBAT 27 (2010). *See infra* Part V. Of course, particularly when it comes to militarized weapons systems, there may be great pressure to make a system as technologically advanced as possible. Just as it is dangerous to bring a knife to a gunfight, it would probably be ill-advised to send a Level 1 system into combat against an enemy fielding a Level 10 system.

[76] Parasuraman et al., *supra* note 39, at 289.

[77] Coppin & Legras, *supra* note 39, at 593.

though they are primarily only automated systems. For example, the Global Hawk reconnaissance drone has the ability to take off and land unassisted.[78] The human operator need only press a button and direct it to take off.[79] However, the Global Hawk lacks the ability to autonomously direct its camera at areas it independently deems to be of interest.[80] For other functions, the human operator can choose among different levels of autonomy.[81] For example, the Predator and Reaper drones each have three flight modes, which are selected by the human operator: manual flying (remote-control piloting), semi-autonomous monitored flight, and pre-programmed flight.[82] The point is that even in these seemingly novel and advanced drone systems, humans control and can actively calibrate the level of autonomy. As is discussed below, this fact is important to understanding the evolution of drone technology, and drone warfare.

III. DRONE WARFARE: YESTERDAY, TODAY, AND TOMORROW

Recent legal and policy debates about the use of drones in warfare focus on whether drones take humans "out of the loop." This discussion is fraught with terms that are both loaded and vague. What does it mean for a human to be "inside," "outside," or "on" the loop? And why does it matter? Further confusing the issues, the debate over "drones" covers a tremendous range of technologies, including those deployed today and their even more sophisticated progeny of tomorrow. It is difficult to structure a thoughtful debate and design an effective regulatory regime without first understanding the technology, including how it functions and how it differs from yesterday's weapons of war. Part II of this paper has begun to address that gap by providing a lens through which legal and policy analysts can understand how machines function and by offering bases upon which to compare different types of machine systems. Part III builds upon that understanding by describing, with reference to the framework laid out in Part II, the drones of yesterday, today, and tomorrow.

A. Yesterday's Drones

2001: A Space Odyssey, Stanley Kubrick's dystopian masterpiece, addressed not just the future of technology, but also the advent of warfare. In the movie's famous opening sequence, early Man discovers that his tools may be used as weapons — and he promptly takes over control of a competing group's watering hole by clubbing its leader to death. Warfare has been evolving ever since. Scholars even have a term for significant advances in warfare technology: they call these paradigm shifts "revolutions in military affairs," or RMAs for short.[83] One RMA happened when the English introduced longbow archers during the Middle Ages, ending the reign of horse-mounted knights. Another RMA came in the mid-

---

[78] SINGER, WIRED FOR WAR, *supra* note 7, at 36.
[79] *Id.*
[80] *Id.*
[81] Air Force Flight Plan, *supra* note 27, at 26–27.
[82] *Id.*
[83] SINGER, WIRED FOR WAR, *supra* note 7, at 181.

fifteenth century, with the dawn of the gunpowder age. In all, historians estimate there have been at least ten RMAs since 1300.[84]

To understand today's warfare technology, one naturally compares it with the technologies that came before. It is necessary, though, to identify the *relevant* points of comparison. Comparing today's Predator drones with the clubs, arrows, and muskets of yesteryear is interesting but of limited use. Contrast today's technologies with yesterday's weapons that exhibited some level either of automation (the use of machines to replace human functions) or autonomy (conceptualized as the use of machines to replace human decision-making) is far more helpful.

Although even basic autonomy is a relatively recent technological phenomenon, automated machines have been used in warfare at least since the beginning of the twentieth century. Nikola Tesla mastered wireless communication in 1893.[85] Five years later, before a crowd at Madison Square Garden, Tesla demonstrated the ability to remotely control the movements of a motorboat.[86] With the combination of wireless radio guidance with advances in engine technology, remote-controlled warfare was underway.[87]

Some of the earliest drones were used as target practice for American fighter pilots during World War II. The most famous example was the "Dennymite," a radio controlled plane invented by Reginald Denny, a British World War I pilot turned Hollywood star, that was used by the U.S. military.[88] The Germans also managed to weaponize remote-controlled machines during the War. For example, they deployed the land torpedo called "Goliath," which was essentially a remote-controlled car loaded with explosives that would be driven into enemy tanks and bunkers.[89] The Germans also developed a device called the "Fritz," which was a drone-like bomb with wings. The Fritz would be dropped from a plane, and a pilot would guide the bomb into a target via remote control.[90] The Dennymite, the Goliath, and the Fritz were primitive technologies that required constant interaction with their human operator. Each lacked any ability to navigate the OODA loop independently, and thus each falls squarely at Level 0 — "remotely piloted vehicle" — on the Air Force's 11-level autonomy spectrum.

Another important early use of drones was for surveillance and reconnaissance missions. A leading example is the Ryan Firebee II, which was used during the Vietnam War and had a range of nearly 900 miles and a flight endurance of more than one hour.[91] Another reconnaissance drone was the "Lightning Bug," a remote controlled plane used extensively in

---

[84] *Id.* at 182.
[85] *Id.* at 46.
[86] *Id.* at 46.
[87] YENNE, *supra* note 75, at 9.
[88] SINGER, WIRED FOR WAR, *supra* note 7, at 49; YENNE, *supra* note 75, at 10.
[89] SINGER, WIRED FOR WAR, *supra* note 7, at 47.
[90] *Id.* at 48.
[91] YENNE, *supra* note 75, at 14.

Southeast Asia between 1964 and 1975.[92] Both the Firebee and the Lightning Bug, like other drones of their generation, were "electro-optically guided, meaning that the controller at a remote control station could watch their progress on television as though he was aboard."[93] Here again, these technologies belong at Level 0 on the Air Force's 11-level autonomy spectrum.

The rudimentary state of guidance technology was among the single largest constraints on the development of early remote-operated weapons systems. Guidance capability — the ability to make an object follow a moving target or deviate from its course as necessary — is what separates an unmanned weapon from a bullet.[94] A significant step toward precision bombing came in 1920 when Carl Norden took advantage of the latest computer technology to invent the Norden bombsight.[95] The Norden's computer improved an aerial bomber's accuracy by automatically launching a missile at the right time to hit, or at least come near to hitting, the desired target.[96] When Paul Tibbets dropped "Little Boy" from the *Enola Gay* onto the city of Hiroshima on August 6, 1945, a Norden bombsight was used ensure the uranium bomb was dropped at the appropriate moment.[97]

War drives demand for warfare technologies, and the Cold War was no exception. But the rise of computer technology shifted attention and money away from robotics, dampening both demand for and development of automated warfare technologies.[98] Those programs that did receive funding fared poorly. One such noteworthy failure was the U.S. Army's ill-fated Aquila program, which was launched in 1979 with the goal of developing a remotely operated drone plane that would fly over enemy territory and conduct surveillance tasks.[99] In 1987, eight years and $1 billion dollars later, the program was cancelled with only a few primitive prototypes to show for the military's efforts.[100]

Although advances happened only gradually, progress did occur. During the Gulf War, the U.S. military deployed the Pioneer drone, a remote-controlled reconnaissance drone that conducted both surveillance and battlefield damage assessments.[101] In a symbolic step forward, if not a technological one, a Pioneer drone became responsible for the first-ever

---

[92] *Id.* at 15.
[93] *Id.*
[94] *Id.* at 9.
[95] SINGER, WIRED FOR WAR, *supra* note 7, at 50.
[96] Peter W. Singer, *In the Loop? Armed Robots and the Future of War*, BROOKINGS INSTITUTION (Jan. 28, 2009), *available at* http://www.brookings.edu/articles/2009/0128_robots_singer.aspx [hereafter Singer, *In the Loop?*]
[97] SINGER, WIRED FOR WAR, *supra* note 7, at 51.
[98] *Id.* at 53.
[99] *Id.* at 55.
[100] *Id.*
[101] *Id.* at 57; YENNE, *supra* note 75, at 27.

surrender of human soldiers to an unmanned system when Iraqi soldiers surrendered to the Pioneer by waving white bedsheets and undershirts.[102]

Perhaps the most significant development in unmanned war technology in the twentieth century was the invention of laser-guided bombs and cruise missiles. Laser-guided bombs are not actively steered and driven by human operators.[103] Rather, the human operator illuminates the target using a laser, and the bomb uses the laser guidance to steer into the target.[104] With early models, the human had to continue to keep the laser on the target. Later technology, however, allowed the human to exit the scene while the bomb completed the mission on its own.[105] This was something new: a weapon that could follow a target and did not need to be constantly guided and steered by a human operator. Nonetheless, these laser-guided bombs were still quite primitive, especially if evaluated on the autonomy scale. The systems lacked decision-making ability and exhibited only the limited capacity to execute a pre-programmed command to follow and crash into a specific target, placing them at a Level 1, at most, on the Air Force's 11-level autonomy spectrum. Moreover, laser-guided bombs had the critical weakness of poor functioning in bad weather.[106] Dust, haze, or smoke could render the bomb's laser guidance capabilities useless.[107]

Cruise missiles were more sophisticated weapons than laser guided bombs that flew themselves using preset coordinate or recognition software.[108] The first cruise missiles, including the pioneering German Fiesler Fi.103, were introduced during World War II.[109] Early cruise missiles were limited because they lacked precision guidance, meaning their ability to hit a specified enemy target was quite limited.[110] Technological advances eventually lead to the Tomahawk missile, which wreaked havoc in the Middle East during the Gulf War.[111] Yet when we study the Tomahawk missile through the lens of the OODA loop and autonomy spectrum, its limitations become apparent. Operating a Tomahawk required setting its target *before* take off.[112] The Tomahawk could not be programmed to attack a particular target midflight and "[i]t could not react to change" in its environment.[113] Furthermore, while the ability to instruct a Tomahawk to embed itself at a particular site diminishes the need for constant human guidance while in flight, the missile must travel over terrain that has already been mapped out and preprogrammed into its computer memory.[114]

---

[102] SINGER, WIRED FOR WAR, *supra* note 7, at 57; YENNE, *supra* note 75, at 27.
[103] SINGER, WIRED FOR WAR, *supra* note 7, at 57.
[104] *Id.*
[105] *Id.*
[106] *Id.*
[107] *Id.*
[108] *Id.*
[109] YENNE, *supra* note 75, at 12.
[110] *Id.*
[111] SINGER, WIRED FOR WAR, *supra* note 7, at 57.
[112] *Id.* at 58.
[113] *Id.*
[114] *Id.* at 57.

While the Tomahawk's reliance on its human operator was minimal once it was fired, its ability to dynamically manage environmental uncertainty was limited.[115]

The history presented here is necessarily brief, but these vignettes offer a few important lessons that apply broadly to early forays into mechanized warfare. The most important point is that while these technologies do exhibit some characteristics of automation, they have almost none of the characteristics of autonomy discussed in Part II. They are all at Level 0 or Level 1 on the Air Force and Sheridan autonomy spectrums. Most of the technologies were remote controlled, including the German Goliaths and Fritzes of World War II and the Pioneer drones of the Gulf War. This means they required nonstop interaction with the human operator, so one of three attributes of autonomy was wholly missing. Moreover, if the communications link failed, the systems were useless. Indeed, a number of reconnaissance drones crashed over North Vietnam during the Vietnam War when the data link to the human operators cut out.[116]

Those early machines that did not require frequent interaction with their human operators generally lacked the ability to navigate uncertain environments and had virtually no mission assertiveness. The Tomahawk missile is a good example, for it could only travel over known terrain and was unable to change targets or to stray into unknown lands. These technological constraints limited the machines' functionality. Most early unmanned weapons systems were essentially either target drones created to be destroyed in practice, or missiles designed to destroy enemy facilities. These systems were automated weapons, to be sure, but they had virtually no ability to navigate the OODA loop with anything resembling autonomy.

B. *Today's Drones*

The military's use of unmanned remote-operated vehicles predates the September 11 terrorist attacks, as the discussion above shows. Still, the advent of the post-9/11 global War on Terror led to vastly expanded use of military drones.[117] Several factors drove this trend. By this time, technology had improved to the point where the widespread use of drones became both feasible and attractive. While the early drone technology described above was hampered by technical challenges, including inadequate communications equipment and a limited ability to integrate an airplane's flight and targeting systems,[118] new technologies have permitted drones to become more effective. GPS technology in particular has been critical, for knowing precisely where a drone is positioned makes them far easier to maneuver and fly remotely.[119]

---

[115] *Id.*

[116] YENNE, *supra* note 75, at 23.

[117] SINGER, WIRED FOR WAR, *supra* note 7, at 61 (quoting Larry Dickerson, who wrote that "[p]rior to 9/11, the size of the unmanned vehicle market had been growing, but at an almost glacial place").

[118] YENNE, *supra* note 75, at 23.

[119] SINGER, WIRED FOR WAR, *supra* note 7, at 58.

Political and cultural factors have also driven the push towards drone technology. Terrorists are unlike the conventional enemies of World War I and II in important respects. Terrorists are not confined to a particular battle space, driving demand for 24/7 global surveillance that can be more effectively accomplished with a drone than with a human pilot. Targeted drone killings have also proved to be a relatively efficient way to confront terrorists that hide among civilian populations.[120] Diminished political tolerance for military casualties has also made drones more attractive, since they keeps soldiers farther from the battlefield, preserving the lives of American servicemen.[121] Increased judicial scrutiny of detention practices may have also increased the appeal of targeted drone killings.[122] Over time, the military has lowered its cultural opposition to drones, particularly its resistance to the notion that a plane should be piloted by a computer rather than a human pilot.[123] And of course, there are also cost concerns: in some cases, drones may be significantly cheaper to operate than manned vehicles.[124]

1. Aerial Drones

Unmanned aerial vehicles are the most well-known and widely deployed segment of the U.S. military's unmanned weapons portfolio. The queen of America's aerial drone program is the Predator, a 27-foot, 1,130-pound airplane that can spend 24 hours in the air at a time.[125] The Predator boasts incredibly sophisticated surveillance capabilities. It has the ability to loiter over a target for long periods,[126] read a license plate from two miles away,[127] and monitor terrain below during day or night, through fog, cloud cover, or gloom of night.[128]

---

[120] Jonathan Ulrich, *The Gloves Were Never on: Defining the President's Authority to Order Targeted Killing in the War Against Terrorism*, 45 VA. J. INT'L L. 1029, 1053 (2005). Drones have much greater endurance ability than humans, so they can patrol a given area for much longer without needing rest and without mental fatigue. See SINGER, WIRED FOR WAR, *supra* note 7, at 36.

[121] Report of the Special Rapporteur on extrajudicial, summary or arbitrary executions, Philip Alston, Addendum, Study on Targeted Killings, U.N. Doc. A/HRC/14/24/Add.6, at 9 (May 28, 2010) ("The appeal of armed drones is clear: especially in hostile terrain, they permit targeted killings at little to no risk to the State personnel carrying them out, and they can be operated remotely from the home State.").

[122] GABRIELLA BLUM & PHILIP B. HEYMANN, LAWS, OUTLAWS, AND TERRORISTS: LESSONS FROM THE WAR ON TERRORISM 88 (2010) ("[T]he killing of a terrorist often proves a simpler operation than protracted legal battles over detention, trial, extradition, and release. The more complicated detention has become, the more attractive the option of targeted killing seems to be.").

[123] YENNE, *supra* note 75, at 27.

[124] *See* SINGER, WIRED FOR WAR, *supra* note 7, at 114.

[125] Robert Valdes, *How the Predator UAV Works*, HOW STUFF WORKS (Apr. 1, 2004), *available at* www.science.howstuffworks.com/predator.htm (last visited Apr. 20, 2012).

[126] YENNE, *supra* note 75, at 48.

[127] *Id.*

[128] *Id.* at 32–33.

The Predator was originally designed to conduct surveillance and was first deployed on reconnaissance missions over the Balkans in 1995 and 1996.[129] After 9/11, however, the drone was outfitted with a payload of laser-guided Hellfire missiles, turning the Predator into an actual predator.[130] The Predator's surveillance capabilities are still critical, however, and this drone model has been used to provide continuous, useful real-time surveillance to soldiers on the ground.[131] Each Predator mission requires significant dedication of human manpower, though. A Predator's crew consists of one pilot and two sensor operators, but it takes a total of 82 people, including technical support staff, to fly it.[132]

In many ways, the Reaper drone is a souped-up version of the Predator. The Reaper can fly twice as fast and at double the altitude of the Predator while carrying a 3,750-pound payload, a nearly ten-fold increase over the Predator's 450-pound payload weight limit.[133] Instead of the Hellfire, the Reaper often carries the Maverick missile, which is capable of more "heavy duty" tasks like busting tanks.[134] The Reaper lacks the Predator's loiter ability, however, and with a flight endurance period of approximately 18 hours, it falls short when compared to the Predator's 22-hour limit.[135] If two 1,000-pound external fuel tanks are added to the Reaper, however, its endurance jumps to 42 hours.[136] The Reaper's extra capabilities come at a higher cost of $53.5 million per unit, compared to the Predator's $20 million per-unit price tag.[137]

While the Predator and Reaper drones have more advanced surveillance and attack capabilities than their predecessors, in many ways these drones are still remote-operated planes that operate at the lowest end of the autonomy spectrum. In sum, "[a]s an aircraft, the Predator UAV is little more than a super-fancy remote-controlled plane."[138] Both the Predator and Reaper can be flown on one of three modes, as was mentioned above, ranging from remote-control flying, semi-autonomous monitored flight, and pre-programmed flight.[139] Each of these modes, however, requires frequent human interaction. The birds are unable to suggest actions to human operators or decide independently how to act.

---

[129] *Id.* at 39.
[130] *Id.* at 43–45.
[131] *Id.* at 69.
[132] Valdes, *supra* note 125.
[133] Air Force Flight Plan, *supra* note 27, at 26–27; United States Air Force, *MQ-1B Predator*, Jan. 5, 2012, *available at* http://www.af.mil/information/factsheets/factsheet.asp?id=122; United States Air Force, *MQ-9 Reaper*, Jan. 5, 2012, *available at* http://www.af.mil/information/factsheets/factsheet.asp?id=6405.
[134] YENNE, *supra* note 75, at 79.
[135] Air Force Flight Plan, *supra* note 27, at 26–27.
[136] YENNE, *supra* note 75, at 79.
[137] United States Air Force, *MQ-1B Predator*, *supra* note 133; United States Air Force, *MQ-9 Reaper*, *supra* note 133.
[138] Valdes, *supra* note 125. For a detailed account of one Predator pilot's experience during the Iraq and Afghanistan conflicts, see MATT J. MARTIN & CHARLES W. SASSER, PREDATOR: THE REMOTE-CONTROL AIR WAR OVER IRAQ & AFGHANISTAN: A PILOT'S STORY (2010).
[139] Air Force Flight Plan, *supra* note 27, at 26–27.

The military also deploys surveillance drones that, unlike the Predator and Reaper drones, do not carry a weaponized payload. At the larger end is the Global Hawk drone, a high-altitude, long-range reconnaissance drone with a 28-hour endurance limit and the capability to conduct observations and collect information from 65,000 feet up.[140] The Global Hawk has the ability to watch what is below it day and night, and through most types of weather.[141] And the Global Hawk may take off and land almost entirely unassisted — "the operator just clicks to tell it to taxi and take off, and the drone flies off on its own."[142] The Global Hawk, of course, still requires a human to decide *when* the drone should take off, and *where* it should fly. In most ways, the Global Hawk is no higher than a Level 1 or Level 2 on the Air Force's autonomy spectrum.

The military also uses a number of smaller hand-launched surveillance drones. The Raven, for example, is a hand-launched surveillance drone with a range of seven to ten miles that can either be manually flown or programmed via GPS to follow a pre-determined flight path.[143] To launch the three-foot long Raven, a soldier simply picks up the bird and heaves it in the air — there is little more takeoff technology than a paper airplane. But while takeoff is mechanical, the Raven is remote-controlled while in flight, and the human operator determines where the Raven goes and where its camera looks.[144] The Wasp III, the Raven's mechanical cousin, is a hand-launched surveillance drone with a range of up to three miles that can either be manually flown or programmed via GPS to follow a predetermined flight path.[145] Like the Global Hawk, these drones operate at the lower ends of the Air Force's autonomy spectrum.

2. Land-bound drones

Not all unmanned drones are aerial vehicles. The PackBot and the SWORDS drone are two leading examples of land-based drones. The PackBot is a 42-pound robot the size of a lawnmower that is deployed by the military in Iraq and Afghanistan to help detect IEDs.[146] It is normally remote-controlled by a nearby American service member, though the drone does have a limited capability to drive independent of a human operator.[147] The PackBot is capable of performing many different functions. For example, it carries with it a mine detector, chemical and biological weapons sensor, and extra power packs. But the machine depends on its user to switch between and deploy the different tools.[148] While PackBot is

---

[140] Air Force Flight Plan, *supra* note 27, at 27.
[141] Gayle Putrich, *Next generation of Global Hawks ready to roll*, FLIGHT INT'L (Aug. 16, 2010), *available at* http://www.flightglobal.com/news/articles/next-generation-of-global-hawks-ready-to-roll-346116.
[142] SINGER, WIRED FOR WAR, *supra* note 7, at 36.
[143] Air Force Flight Plan, *supra* note 27, at 26.
[144] SINGER, WIRED FOR WAR, *supra* note 7, at 37.
[145] Air Force Flight Plan, *supra* note 27, at 25.
[146] SINGER, WIRED FOR WAR, *supra* note 7, at 22.
[147] *Id.*
[148] *Id.*

equipped with a camera in order to peer at possible IEDs, it simply transmits that information back to its controlling soldier and does not actually interpret the visual information itself.[149] The PackBot lacks the capability to independently observe the objects in its environment and decide for itself which sensor to deploy at a given time. Instead, the bot depends on regular interaction with its operator to fulfill its tasks. If the communications link cuts out, the bot is unable to continue its rounds and search the area for IEDs. The PackBot, like the other technologies discussed before, ranks at the low end of the Air Force's autonomy spectrum, probably no higher than Level 1.

The SWORDS drone lives up to its name: it is a robot equipped with a mounting that can carry almost any weapon under 300 pounds, including a machine gun, grenade launcher, and antitank rocket launcher.[150] On the autonomy spectrum, however, SWORDS is in many ways even more primitive than PackBot. The SWORDS drone must always be remote controlled and the human operator is responsible for loading its gun, as well as for aiming, firing, and reloading. SWORDS is an archetypical automated but non-autonomous robot. It replaces human functions (lock, load, and fire) but does not supplant human decision-making.[151] SWORDS may be placed at Level 0 on the Air Force's autonomy spectrum.

3. An automated, non-autonomous fleet

Technological problems still retard the operation of America's drones. The communications links between drones and their remote pilots are particularly unreliable and "regularly cut out, forcing the robotic aircraft into automatic holding patterns."[152] Even today's drones have low levels of assertiveness, one of the critical attributes of autonomy, for they have little ability to complete a mission without close human direction. These drones depend upon very frequent operator interaction to complete and close the OODA loop. The primary advancement and greatest advantage of the Predator and Reaper drones is *not* autonomy or independence from humans, but rather simply that these drones are able to execute missions "without ever exposing the pilot to the hostile environment."[153] Thus, these drones are useful because they *replace human actions* — in other words, because they are highly *automated* machines. Predators, Reapers, and their mechanical kin differ from earlier versions of drones not in terms of increased autonomy, but rather because they are better at automation.

Few of the military's weapons systems outsource any decision-making authority over when and at whom to shoot. There are some notable exceptions, such as Counter Rocket Artillery Mortar (CRAM) technology, an air defense technology that "automatically tracks

---

[149] *Id.*
[150] *Id.* at 30.
[151] *Id.* at 30.
[152] Noah Schachtman, *Report: U.S. Drone Goes Down Over Pakistan. Again.*, WIRED DANGER ROOM BLOG, Jan. 25, 2010, 11:02AM, http://www.wired.com/dangerroom/2010/01/us-drone-goes-down-over-pakistan-again/.
[153] Valdes, *supra* note 125.

and shoots down any missiles that have gotten past all other defenses and are to quick for humans to react to."[154] "R2-D2," as the technology is also known, is not without its risks. The system once locked in upon and prepared to shoot down an American helicopter over Baghdad, nearly blowing up the plane and potentially killing American soldiers.[155]

The Aegis sea defense system is another example of the benefits, and potentially dire costs, of machines programmed to kill. The Aegis has four modes, ranging from semiautomatic, in which the human controls the system and decides when and what to shoot, all the way to "casualty," where the system is authorized to shoot at its own will in order to keep the ship safe.[156] One risk associated with this system is that even in the less-autonomous modes, where a human retains a veto over the system's decisions, the humans will trust the machines better than their own judgment. This resulted in tragedy in 1988, when an Aegis system aboard the U.S.S. Vincennes shot down a civilian airliner it mistook for an Iranian fighter jet in the Persian Gulf.[157] Even though the Aegis system's human operators had a veto over the decision to shoot, they trusted the Aegis system more completely than they trusted themselves, and allowed the machine to shoot.[158] The story illustrates the limitations of engineering in a different and more painful way, for it shows that the availability of operator interaction may at times be academic.

Today, America's drone fleet makes the military more potent than it has ever been before. The drones permit more continuous surveillance of enemies and more accurate targeting at less financial cost, and at less cost to soldiers' lives. These drones are more effective than their predecessors, and they are deployed on a previously unprecedented scale. But despite these incredible advances, the drones are still primarily automated, non-autonomous systems. Virtually all presently operative drones are at Levels 0, 1, or 2 on the autonomy spectrum.

C. Tomorrow's Drones

Tomorrow's drones will exhibit greater autonomy along all four stages of the OODA loop and will be more sophisticated in each of the three attributes of autonomy described above. Future drones will require less human interaction, navigate greater levels of environmental uncertainty, and enjoy higher levels of mission assertiveness. The military predicts that these increasingly autonomous machines will be eventually be fully integrated into the fighting forces.[159] By 2025, the American armed forces are expected to be "largely robotic."[160] "The robots you are seeing here today, I like to think of as the Model T," a

---

[154] SINGER, WIRED FOR WAR, *supra* note 7, at 38.
[155] *Id.*
[156] *Id.* at 124.
[157] *Id.* at 124 –25.
[158] Harris, *supra* note 3, at 10*;* SINGER, WIRED FOR WAR, *supra* note 7, at 125.
[159] Air Force Flight Plan, *supra* note 27, at 48.
[160] SINGER, WIRED FOR WAR, *supra* note 7, at 133.

robotics executive said at a 2007 demonstration of robot prototypes.[161] "We are seeing the very first stages of this technology."[162]

The discussion above noted several reasons for the military's increasing reliance upon drone technology, including cost-effectiveness and diminished risk to human service members. Assuming these incentives persist, reliance upon drones will continue to grow. Other influences are also driving an increased appetite specifically for drones able to operate autonomously, or independent of human control. Machine systems dependent on human commands are vulnerable to enemy sabotage of communications, and so systems able to continue the missions even without contact and direction from an operator are especially valuable.[163] Systems deployed in combat or defensive positions must have the capability to act in circumstances that do not permit time for a human to react and give a command.[164] Finally, as the technology improves, drones will simply become better at executing a greater number of combat and surveillance tasks than humans.[165] As one commentator noted dryly: "They don't get hungry. They're not afraid. They don't forget their orders. They don't care if the guy next to them has just been shot. Will they do a better job than humans? Yes."[166]

The United States military has discussed the evolution of these machines in terms of the OODA loop. "One of the most important elements to consider with this battlefield is the potential for UAS [unmanned aerial systems] to rapidly compress the observe, orient, decide, and act (OODA) loop. Future UAS able to perceive the situation and act independently with limited or little human input will greatly shorten decision time."[167] The military is following Boyd's advice, albeit not necessarily in a way that he would have foreseen: through the development of autonomous systems that can go through the OODA loop better and faster than our adversaries. Increasingly, according to the U.S. military, "humans will not longer be 'in the loop' but rather 'on the loop' — monitoring the execution of certain decisions."[168]

The drones of tomorrow will generally exhibit increased autonomy. The next generation of drones will be able to take off and land without any human interaction[169] and even to perform one of the most difficult flight tasks — landing on an aircraft carrier — without any human control.[170] Drones will also have the capability to stay aloft for much longer periods of time. For example, the Defense Advanced Research Projects Agency (DARPA), a Department of Defense agency that funds research of new technologies, has

---

[161] *Id.* at 110.
[162] *Id.*
[163] Singer, *In the Loop?*, *supra* note 96.
[164] *Id.*
[165] SINGER, WIRED FOR WAR, *supra* note 7, at 63 (quoting Gordon Johnson of the Pentagon's Joint Forces Command).
[166] *Id.*
[167] Air Force Flight Plan, *supra* note 27, at 16.
[168] *Id.* at 41.
[169] *Id.* at 39.
[170] Hennigan, *supra* note 7.

invested in the development of drones that can stay aloft for up to five years.[171] Drones will also be able to conduct surveillance without the need for a human operator actively steering a camera to areas of interest. The future Reaper drones, for instance, are expected to have the capability to not only "recognize and categorize humans and human-made objects" it encounters, but also to "make sense of changes it is watching, such as being able to interpret and retrace footprints or even lawn mower tracks."[172]

The military plans that drone technology will eventually advance to the point where systems are able "to make combat decisions and act within legal and policy constraints without necessarily requiring human input."[173] Machine systems will be fully integrated into the armed forces, with commanders able to refine a system's level of autonomy by mission, just as they can now vary a human soldier's rules of engagement.[174] These drones will increasingly take to the sea and ground, both far more complex and difficult environments to navigate than the air. One such ground robot is the Gladiator, a well-armed combat robot that will be able to operate on semi-autonomous and fully autonomous modes.[175] In the water, there will be the Spartan Scout, now a prototype that, under its design, can be launched into the sea and will travel on its own for up to two days, carrying out surveillance missions, protectively patrolling harbors, and inspecting any suspicious ships it encounters.[176]

One of the most fascinating and novel trends is towards micro-drones and nano-drones, some bio-mechanical and as tiny as the width of a human hair.[177] These drones are expected to work together in large swarms, able to communicate with each other and thereby coordinate their activities. Publicly available military sources suggest these nano-drones portend a paradigm shift, with the "[d]evelopment of the nano/micro class" creating "capabilities never before realized."[178] According to one observer, nano-drones add to the push beyond mere automation because their tiny designs and numerosity "actually mandate that the systems will have to have high autonomy, carrying out their missions without human controllers."[179] This is so for several reasons. The drones are too numerous to each have an individual human controller, and flying the micro machines will actually make most human operators suffer from nausea.[180] Furthermore, the effectiveness of these drones comes from their ability to communicate with each other and coordinate their actions more effectively and rapidly than a group of human operators could possibly do.[181] One risk associated with these nano-drones is that because by definition they are not fully controllable by humans, they may

---

[171] SINGER, *supra* note 41, at 117.
[172] *Id.* at 116.
[173] Air Force Flight Plan, *supra* note 27, at 41.
[174] *Id.*
[175] SINGER, WIRED FOR WAR, *supra* note 7, at 111.
[176] *Id.* at 115.
[177] *Id.* at 118.
[178] Air Force Flight Plan, *supra* note 27, at 36.
[179] SINGER, WIRED FOR WAR, *supra* note 7, at 119.
[180] *Id.*
[181] *Id.*

not go exactly where the operator wants them to go, or do exactly what he or she wants them to do.[182] The human controller is responsible for directing the swarm and will have the ability to monitor the swarming unit, but his ability to control their individual actions may be quite limited.[183]

IV. THE PARADIGM SHIFT YET TO COME

The foregoing discussion of drone technology allows one to begin to see what is "new" about today's military technology, and how tomorrow's technology will be different as well. From an engineering perspective, today's drones are actually not so different from the automated systems developed earlier in the Twentieth Century. Despite the hype, today's drones remain primarily at Levels 0 or 1 of the autonomy spectrums discussed in Part II. Even the advanced Predator drone is normally operated by remote control.[184] The human pilot remains, albeit not in the cockpit. Today's technology has yet to march further down the autonomy spectrum into areas where the machine can navigate the OODA loop with little human interaction. Contemporary systems require frequent operator guidance, have only limited ability to navigate uncertain environments, and possess little assertiveness. Although drones might give the remote operator access to a greater range of capabilities than, say, a Tomahawk missile, both are fundamentally remote-operated war machines with little independent decision-making authority. Today's machines are far more effective automated robots than yesterday's automated robots, to be sure. But they are still automated, not autonomous.

Of course, lawyers and policymakers need not necessarily view the world through the lens of the engineer. And in many ways, today's drones do represent a revolution in warfare. First, drones keep their human operators farther from the battlefield than ever before. During World War II, for instance, the German "Fritz" had to be dropped from a human-operated plane above the battlefield.[185] Today, American pilots fly planes around the world from a relatively safe perch, often in Nevada, away from the battlefield.[186] Second, drones have much greater endurance than their human counterparts. A surveillance drone can fly for hours, and soon for days, at a time; a combat drone can stalk prey for long periods, awaiting the right moment to strike.[187] Third, drones today are deployed on a wholly new scale. The cost of technology is low enough, and the cost of human forces high enough, that drones may soon outnumber human soldiers.

Although engineers and policymakers may disagree about whether today's drones truly are different, the systems of tomorrow may unite the two camps. The race is on to develop autonomous systems that can complete the OODA loop better and faster than our

---

[182] *Id.* at 235.
[183] See Air Force Flight Plan, *supra* note 27, at 34.
[184] SINGER, WIRED FOR WAR, *supra* note 7, at 36; YENNE, *supra* note 75, at 40, 75.
[185] SINGER, WIRED FOR WAR, *supra* note 7, at 48.
[186] MARTIN & SASSER, *supra* note 138, at 2.
[187] *See supra* Part III.B.1.

rivals. The systems currently being developed on the floors of research-and-development labs and in the minds of inventors may prove to be truly *new* technologies. These machines will climb higher on the autonomy scale — through Levels 4, 5, 6, and beyond — and eventually be able to complete the OODA loop with complete independence from their human operators. Truly autonomous drones may have capacities not before seen. If policymakers are struggling with the implications of today's technology, the conversations about tomorrow's will be still more fraught.

With tomorrow's drones, humans will remain in the loop, but in a rather different way. Human intervention will still impact these systems, albeit at an earlier stage in what could be loosely dubbed a machine's "lifecycle." Even if operators' influence is reduced, engineers and systems designers will retain power over how the machines traverse the loop.[188] In other words, engineers will be able to program how systems process the information around them and, still more importantly, control the factors that machines consider during the decision-making processes. As pressure intensifies to deploy systems able to complete the loop faster than human synapses permit, hotly debated law and policy questions, including whether drones are capable of complying with international law (particularly *jus in bello* and *jus ad bello* requirements), must increasingly be addressed at the design stage. While a detailed discussion of these laws of war and their application to machine warfare is beyond the scope of this paper,[189] the underlying point remains. As the machines evolve, so too will the meaning of having a human "in the loop."

V. REGULATORY FRAMEWORKS FOR AUTONOMOUS SYSTEMS

Today, technology constrains our ability to make systems truly autonomous. This will not always be so. As technology continues to advance, legal, cultural, and political considerations will increasingly function as the primary limitations upon our weapons systems. Legal and ethical problems posed by tomorrow's drones are best addressed today, while the technology is still in a relatively nascent stage and the opportunity remains to control and constrain the development of these advanced systems.[190] The embryonic state of this technology presents a unique opportunity to think creatively and proactively about the systems' future development.[191] And the evidence indicates that legal, political, and cultural considerations do impose real boundaries.[192] For example, plans to construct a nuclear-powered drone able to stay aloft for years at a time were recently shut down over concerns

---

[188] While a discussion of the effect of increasingly autonomous technologies on supply chain security needs and the heightened potential for havoc to be wrought by insider threats is beyond the scope of this paper, these issues are grave, urgent, and clearly implicated by the shift towards greater autonomy.

[189] For thoughtful commentary, see Kenneth Anderson, *Efficiency in Bello and ad Bellum: Targeted Killing Through Drone Warfare*, SSRN (Sept. 23, 2011), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1812124; *see also* Jenks, *supra* note 2, at 671.

[190] See *supra* Parts III-IV.

[191] *Cf.* Arkin et al., *Moral Decision Making in Autonomous Systems*, *supra* note 12, at 1.

[192] Nick Fielding, *U.S. Draws Up Plans for Nuclear Drones*, THE GUARDIAN (Apr. 2, 2012), *available at* http://www.guardian.co.uk/world/2012/apr/02/us-plans-nuclear-drones.

about nuclear fallout from a drone crash and fear that terrorist groups might obtain the technology.[193] Lawyers and policymakers should reject technological determinism, which denies the ability of law and culture to impact system development.

This paper's discussion of autonomy and the OODA loop suggests ways to structure regulatory regimes, for it allows us to identify a number of points around which one could focus when designing legal controls that target specific concerns about technology. The current debate over whether humans are "in the loop" or "out of the loop" has an all-or-nothing feel to it: humans are either in or out. This frame is unhelpful because it is too simplistic. As this paper has shown, in truth machine systems do not work through a single loop; rather, a system will work through the OODA loop dozens, hundreds, or thousands of times during each mission. Each loop might be addressed to a different task — whether to move left around a boulder, shoot a missile, or take some other action.[194] Furthermore, this paper has also demonstrated that humans can calibrate how "wide" the loop may be — how autonomously a machine functions — and that the level of autonomy can vary across each stage of the OODA loop.[195] If we do want to regulate the technologies, adopting these more nuanced views of system operations and tailoring legal intervention accordingly will permit us to control their development and use without being overbroad, curtailing future innovation, and sacrificing some of the benefits the technology offers today.

To establish a successful regulatory framework, lawyers and policymakers must first consider the values they wish to promote. They can then constrain a machine's autonomy, and its independence inside the OODA loop, accordingly. We have some flexibility to keep humans "in the loop" when we want them there, and "out of the loop" when we do not. A machine's decision-making authority and independence from humans is a function of its engineering, and thus can be a function of deliberate choices made by policymakers. For example, when it comes to shooting a missile at an enemy, we may want to keep humans in the OODA loop, at least at the "decide" stage. But, we may feel more comfortable granting an aerial system complete autonomy over its flight path.

One way to structure a regulatory regime would be to restrict only certain stages of the loop, or to regulate different stages differently. At the *observe* stage, for instance, we could establish rules about the locations or types of activities the system is permitted to observe, or the duration of the observation. Depending on what we value, these rules might not necessarily decrease the amount of observation or restrict the places in which observation takes place. The amount and type of information that a machine gathers at the observe stage in the loop affects the machine's capacity to *orient* itself, the number and type of actions weighed when the machine *decides*, and the eventual *act* chosen and carried out. So, for instance, regulations could require the machine on a "kill" mission to remain at the observational stage of the loop for a longer period of time, or mandate that it collect a larger amount of data, or collect it using a wider variety of sensors, if an initial scan of the scene

---

[193] *Id.*

[194] *See supra* Part II.A.

[195] See *supra* Part II.B.

indicates that a proportionality analysis would be particularly complex (perhaps because children may be present — as might be suggested by an observed person's height relative to the ground or to others around him, or by his gait, the objects near him, the pitch of his voice, or the pattern of his movements). In some circumstances, a more robust data set, created by lengthier and more detailed observation, may enable the act eventually taken based upon the observations gathered to be more discriminating.

A regulatory regime could also distinguish between those decision-making loops the machine is permitted to carry out without human supervision or intervention and those requiring some level of human authorization. Constraints across loops can exist either instead of or in addition to the constraints imposed within individual decision loops. Laws could stipulate that a machine is permitted to autonomously cycle through loops where the act that results from the completed loop is some form of locomotion, or movement through space. In contrast, we could structure the human-machine authority differently with loops where the act at the loop's end may be fatal, such as releasing a missile or bullet. Using Sheridan's Levels of Autonomy, introduced above, one could bar any system from operating at a level of autonomy higher than Level 5 when shooting, regardless of whether the system would be technologically capable of operating at a level beyond Level 5 if not artificially constrained. That is, if the *act* that is at the end of the OODA loop is "shoot," the system could offer a complete set of action alternatives, suggest one (here, to shoot), and execute that suggestion only if the human approves — even if the system would be technologically capable of, for instance, shooting without human permission and informing the operator of the act only if queried about it, consistent with Level 8 on Sheridan's scale.[196]

Regulators could develop a list of acts for which human intervention is necessary either on an absolute basis — i.e., *any* decision to shoot — or tie the requirement that a human be given the option to intervene or to veto the action to other factors, such as the number of action alternatives the system has evaluated and discarded before settling on the act (whatever the act may be) proposed to be carried out. Although the necessity or propriety of most acts depends on context to at least some extent, policymakers could make it a priority first to establish the level of autonomy at which a system is allowed to operate for those acts that are highly likely to have lethal consequences, then consider others whose effects may be either less grave or potentially reversible.

Again, the regulatory approach may and should vary depending on which values we want to promote. For example, a regime that is primarily concerned about *accuracy* in systems decision-making could impose rigorous decision-making criteria on aerial systems, perhaps by requiring it to spend increased time at the *observe* and *orient* stages, or by allowing the machine to *decide* to take action only if it has reached a sufficiently high confidence level. An "accuracy regime" might also provide for human veto power, or even affirmative authorization, of significant machine decisions. A regime that is primarily concerned with *accountability* might look different. This regime might impose accountability

---

[196] *Id.*

by requiring affirmative human sign-off of each machine-executed kinetic action, or only of particular decisions. But the regime could also permit substantial decision-making autonomy by the machines, even over decisions to kill, so long as a human (probably the military commander, but perhaps the engineer) is responsible for any machine errors. In reality, a comprehensive regulatory regime will incorporate a number of values including accuracy, accountability, machine efficiency, and so forth. We can calibrate a machine's independence within and across loops to reflect our preferences for these different values.

This brief sketch elaborates only a few of the many ways in which better technological understanding allows us to craft more thoughtful regulation. On this view, the question of whether a human is or must remain "in the loop" can be parsed into more focused questions such as *which* loop and which stages of that loop matter, and how *wide* the machine's discretion to complete the loop independently should be. This approach allows us to begin thinking proactively about complex technologies now, even while some debates continue.

No regulations will necessarily be foolproof. Handing over decision-making authority to a machine necessarily entails some measure of risk that the machine will act unpredictably. Moreover, as discussed earlier, the OODA loop is itself a simplification of the decision-making process. As a result, it may not be possible to perfectly segment regulation by stage of the loop, or across different loops. Policymakers should be informed of these risks, and their tolerance for these risks should be incorporated into the policy calculus and regulatory design. We must also remain cognizant of potential limits on a machine's ability to replicate human decision-making in all ways.

Scholars and technical experts in the emerging field of robot ethics, for instance, are actively exploring how future systems can integrate morals at the level of a machine's programming.[197] Scientists and philosophers are debating in tandem the theoretical principles — whether consequentialist, utilitarian, deontological, or some admixture — in which that governing code should be grounded.[198] Through these conversations, the competing views of

---

[197] *See*, *e.g.*, Ronald C. Arkin, *Governing Legal Behavior: Embedding Ethics in a Hybrid Deliberative/Reactive Robot Architecture*, TECH. REP. GIT-GVU-07-11 (2009) [hereafter Arkin, *Governing Legal Behavior*], *available at* http://www.cc.gatech.edu/ai/robot-lab/online-publications/formalizationv35.pdf; Russell W. Robbins & William A. Wallace, *Decision Support for Ethical Problem Solving: A Multi-Agent Approach*, 43 DECISION SUPPORT SYS. 1571 (2007), *available at* http://www.sciencedirect.com/science/article/pii/S0167923606000418; Bruce M. McLaren, *Computational Models of Ethical Reasoning: Challenges, Initial Steps, and Future Directions*, IEEE INTELLIGENT SYS. 2-10 (July/Aug. 2006), *available at* http://tinyurl.com/88jbkae; Thomas M. Powers, *Deontological Machine Ethics*, AAAI TECH. REP. FS-05-06 79-86 (2005), *available at* http://www.aaai.org/Papers/Symposia/Fall/2005/FS-05-06/FS05-06-012.pdf; Bruce M. McLaren, *Extensionally Defining Principles and Case in Ethics: An AI Model*, 150 ARTIFICIAL INTELLIGENCE J. 145 (2003), *available at* http://tinyurl.com/6w8wwhr. *See also* SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS (2011).

[198] *See*, *e.g.*, Carl Shulman et al., *Which Consequentialism? Machine Ethics and Moral Divergence* at 2 (2009), *available at* http://singinst.org/upload/machine-ethics-moral-divergence.pdf (commenting that "the picture that emerges" when one considers establishing a canonical theory for machine ethics

the content and meaning of autonomy that lurk in discussions about the use of advanced technologies will come home to roost in machine systems architecture. This will force the explication, quantification, and systemization of law and values with more exquisite precision than much of what has come before.

The decision sequence of the United States Army Soldier's Guide provides a useful glimpse of a core lingering difficulty of translating carbon ethics into silicon programming.[199] "The Ethical Reasoning Process" offered in the Army Soldier's Guide is designed to help soldiers respond thoughtfully and correctly to ethical dilemmas, or situations in which soldiers "cannot simultaneously honor two or more values and follow given rules while accomplishing the mission.[200] One such hypothetical scenario is the dilemma of "The Checkpoint":

> Two days after a suicide car-bombing killed four soldiers at a checkpoint, another unit is operating a similar checkpoint some distance away. The unit was recently involved in offensive operations but was beginning the transition to stability operations. Unit training has emphasized the importance of helping the citizens return to a "normal" lifestyle. Nonetheless, the events of the previous day demonstrate that the enemy is still active, and will use civilian vehicles loaded with explosives to kill themselves in an attempt to also kill U.S. soldiers. At this time, soldiers at the checkpoint notice a large civilian passenger vehicle approaching at a high rate of speed.[201]

The Guide sets out a four-step process for soldiers to think through this dilemma.[202] Step 1 directs the service member to define the problem.[203] Step 2 instructs the soldier to "[k]now the relevant rules and values at stake," including the law, administrative rules, rules of engagement, command policies, and Army values.[204] At Step 3, the individual must "[d]evelop possible courses of action (COA)" and "evaluate them" using criteria excerpted here:

> a. Rules — Does the course of action violate rules, laws, or regulations?
> b. Effects — After visualizing the effects of the course of action, do you foresee bad effects that outweigh good effects?
> c. Circumstances — Do the circumstances of the situation favor one of the values or rules in conflict?

"is one of deep confusion: we have no idea what utility function to endow an explicit moral agent with," and concluding, "current moral theories are inadequate for machine ethics").

[199] *See* UNITED STATES ARMY FIELD MANUAL NO. 7-21.13, THE SOLDIER'S GUIDE (Feb. 2004), http://www.milsci.ucsb.edu/UCSB_Army_ROTC/Links_files/FM%20721.13,%20The%20Soldier's%20Guide.pdf.

[200] *Id.* at 1-29–1-30.

[201] *Id.* at 1-30.

[202] *Id.*

[203] *Id.*

[204] *Id.*

d. "Gut check" — Does the course of action "feel" like it is the right thing to do? Does it uphold Army values and develop your character or virtue?[205]

At Step 4, the final stage in the process, the Guide concludes, "Now you should have at least one COA that has passed Step 3. If there is more than one COA [that has passed Step 3], choose the course of action that is best aligned with the criteria in Step 3."[206]

Although it may be possible to code the processes at Steps 1, 2, 3(a)-(c), and 4 into automated systems, for now, the Guide's Step 3(d) remains, "clearly outside the scope of autonomous systems"[207] even according to those dedicated to and assiduously seeking ways of embedding ethics into machines. Scientists have successfully created a mechanical gut,[208] but the check proposed by the Army Guide needs a stomach of a different kind. Some doubt strongly that any artificial process could make a machine humane, much less imbue it with the ineffable human qualities the Army Soldier's Guide's Step 3(d) is intended to muster.[209] "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," scholars of this view argue, "it could not be ethical in the same way as is an ethical human. . . . In most real-world situations, these [decisions] are a matter of interpretation."[210]

This aspect of the debate over ethical robotics refracts through the prism of technology a long and profound disagreement over the nature of morality and the extent to which an act sounding in moral judgment can be disaggregated into component parts in the first instance[211] — much less coded, however painstakingly, into an algorithmic rational choice process. The gap between those who believe the human decision-making process to be replicable and those who believe it is indissoluble and irreplaceable might not be bridged. More to the point, the disagreement will not likely be settled before the practical question of whether and how to proceed with the development and use of increasingly autonomous technologies presses too urgently to avoid. This paper's discussion of machine systems and autonomy provides a language that can be used by lawyers and policymakers to structure a regulatory regime for tomorrow's technologies today. Duly equipped, policymakers should act.

---

[205] *Id.*

[206] *Id.*

[207] Arkin, *Governing Legal* Behavior, *supra* note 197, at 51.

[208] *Scientists Create Artificial Gut*, BBC (Nov. 10, 2006, 3:17PM), http://news.bbc.co.uk/2/hi/health/6136546.stm.

[209] *See* PATRICK LIN, KEITH ABNEY & GEORGE A. BEKEY, ROBOT ETHICS: THE ETHICAL AND SOCIAL IMPLICATIONS OF ROBOTICS 121 (2012) ("To be humane is, by definition, to be characterized by kindness, mercy, and sympathy, and 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.").

[210] *Id.*

[211] For explorations, *see, e.g.*, PHILIPPA FOOT, NATURAL GOODNESS (2003); Elizabeth Anderson, *Practical Reason and Incommensurable Goods*, *in* INCOMMENSURABILITY, INCOMPARABILITY, AND PRACTICAL REASON 90-109 (1998) (Ruth Chang, ed.).

VI. CONCLUSION

With the comprehensive understanding of autonomy and the OODA loop this paper has provided, one can see that both truly autonomous systems and automated machines work through the OODA loop thousands or millions of times in the course of each mission. Each stage in the loop may be addressed to a different decision — whether to move left or right around a boulder, whether to spy on a target, or even whether to drop a bomb — and produce a different outcome. The present debate over whether humans are “in the loop” or “out of the loop,” however, has an all-or-nothing feel. It does not adequately account for the complexity of the technology it considers.

Technological complexity actually can be helpful to policymakers, however, so long as it is explored clearly and thoughtfully. It means that we can keep the human “in the loop” when we want her there, leave her “out of the loop” when we determine that her attention is not needed or is more usefully directed elsewhere, and determine how “wide” the loop should be, or how free of human intervention, based on the capabilities of the system or its mission. Our detailed discussion of drones and machine decision-making shows that these and a host of other permutations are possible, for there are many degrees of autonomy and many ways in which machines may or may not exhibit it. This insight will help us to begin thinking about smart, targeted regulatory regimes that serve both national security and our values. We should do so now, while these systems remain relatively rudimentary.