April 2021

Ethics and Artificial Intelligence

A Policymaker's Introduction

CSET Policy Brief

AUTHOR
James E. Baker

Executive Summary

Policymakers contemplating the burgeoning field of artificial intelligence will find, if they have not already, that existing laws leave huge gaps in deciding how (and whether) AI will be developed and used in ethical ways. The law, of course, plays a vital role. While it does not guarantee wise choices, it can improve the odds of having a process that will lead to such choices. Law can reach across constituencies and *compel*, where policy *encourages* and ethics *guide*. The legislative process can also serve as an effective mechanism to adjudicate competing values as well as validate risks and opportunities.

But the law is not enough when it contains gaps due to lack of a federal nexus, interest, or the political will to legislate. And law may be too much if it imposes regulatory rigidity and burdens when flexibility and innovation are required. Sound ethical codes and principles can help fill legal gaps. To do so, policymakers have three main tools:

- Ethical Guidelines, Principles, and Professional Codes
- Institutional Review Boards (IRBs)
- Principles of Corporate Social Responsibility (CSR)

Below is a primer on the limits and promise of these three mechanisms to help shape a regulatory regime that maximizes the benefits of AI and minimizes its potential harms.

This paper addresses specific considerations for policymakers:

1. Where AI is concerned, ethics codes should include indicative actions illustrating compliance with the code's requirements. Otherwise, individual actors will independently define terms like "public safety," "appropriate human control," and "reasonable" subject to their own competing values. This will result in inconsistent and lowest-common-denominator ethics. If the principle is "equality," for example, an indicative action might require

training data for a facial recognition application to include a meaningful cross-section of gender and race-based data.

2. Most research and development in AI is academic and corporate. Therefore, Institutional Review Boards and Corporate Social Responsibility practices are critical in filling the gaps between law and professional ethics, and in identifying regulatory gaps. Indeed, corporations might consider the use of IRBs as well.

3. Policymakers should consider the Universal Guidelines for Artificial Intelligence (detailed below) as a legislative checklist. Even if they don't adopt the guidelines, the list will help them make purposeful choices about what to include or omit in an AI regulatory regime consisting of law, ethics, and CSR.

4. Academic leaders and government officials should actively consider whether to subject AI research and development to IRB review. They should further consider whether to apply a burden of proof, persuasion, or a precautionary principle to high-risk AI activities, such as those that link AI to kinetic or cyber weapons or warning systems, pose counterintelligence (CI) risks, or remove humans from an active control loop.

5. Corporations should create a governance process for deciding whether and how to adopt CSR national security policies answering the question: What does it mean to be an American corporation? They should consider adopting a stakeholder model of CSR that is, in essence, a public-private partnership that includes input from consumers and employees as well as shareholders and the C-Suite.

6. Policymakers, lawyers, and corporate leaders should communicate regularly about the four issues that may define the tone, tenor, and content of government-industry relations: uniformity in response, business with and in China and Russia, encryption, and privacy.

7. Where government agencies, corporations, and academic entities have adopted AI Principles, as many institutions now have, it is time to move from statements of generic principle to the more difficult task of applying those principles to specific applications.

Ethical Guidelines

Commentators are quick to observe that artificial intelligence poses ethical challenges, but not as quick to detail those challenges or identify their solutions. These ethical questions tend to derive from the use of AI in decision-making, data management, and bias.

Ethics, as the National Security Commission on Artificial Intelligence and Defense Innovation Board (DIB) have noted, are what will (or could) distinguish American or democratic use of AI from the authoritarian use of AI. For example, ethics embedded in corporate policies and law will help determine the extent to which U.S. companies partner with authoritarian regimes in the development and deployment of AI systems used to monitor and control domestic populations. The ethical use of AI will help to attract talent to, or keep it in, the United States, including in industry, academia, and government. Ethical use of AI will also encourage security and economic alliances from like-minded governments and entities. Conversely, the perception that the United States is using AI in unethical manners will deter AI talent from working in the United States, for the U.S. government, and hinder international cooperation. Finally, the transparent and ethical use of AI will more likely garner public trust and support, leading to the sustained commitment needed to maximize the security and economic advantages of AI and mitigate the risks.

Ethical choice will supplement legal requirements, or, where policymakers cannot agree on policy or law, fill the vacuum. One ethical approach, in the absence of a comprehensive law or policy framework, is to adopt general principles to apply to specific scenarios. This is done, for example, with the law of armed conflict, which does not seek to address every possible combat scenario with a rule, but rather requires the application of several binding principles: necessity, proportionality, minimization of suffering, and military objective. The DIB, for example, recommended that DOD not deploy AI until an application was demonstratively responsible, equitable, traceable, reliable, and governable. The Department subsequently adopted these five "Ethical Principles for Artificial Intelligence." The Office of the Director of National Intelligence has adopted a similar set of six "Principles of Artificial Intelligence

Ethics for the Intelligence Community," while also promulgating lists of related questions to address and espouse the importance of a stakeholder model of ethical governance. This model encourages, but does not compel, the inclusion of lawyers and civil liberties and privacy offices and officers in the design, testing, and use of AI.

However, principles only go so far if they do not indicate how they should apply to a specific application in context. Thus, to this list of DIB and IC principles, one might add the notion that policymakers should not approve the use of an AI application until they purposefully determine how the applicable principles apply in specific context. They should do the same with the following additional principles:

- **Agency:** The right to determine how one's image, voice, or data are used and by whom.
- **Privacy:** Conscious consideration and accountable choice regarding (1) how, by whom, and for what purpose data is aggregated; (2) how, by whom, with what purpose, for how long, and with what degree of notice data sets are collected and stored; (3) the application of security safeguards commensurate with the AI value of data, not necessarily the value at the time of which it was collected; and (4) whether data is allowed to be sold or transferred, including overseas.
- **Accountability:** With each AI application, policymakers should ask, who is authorized to make a request to use the application? According to what standards? According to what process of review? And with what record of use? They should ensure there is a method in place to determine algorithmic design and to record its accuracy.
- **Accuracy:** Designers and users of AI have a duty to ask who has validated the accuracy of the data or AI in use or accounted and mitigated for any bias embedded in the design or operation of the application. They also have a duty to test, validate, and as necessary, mitigate the use of AI on an ongoing basis.

- **Equality:** In cases where it is legal to collect data, one should ask if it is ethical to do so based on the principle of equality. Are disfavored or disadvantaged groups treated differently for reasons not related to empirical need or programmatic purpose, and if so, why? Is the data being used for AI development derived from unwitting or unwilling subjects or through contracts of adhesion, and if so, does it matter? Similarly, one should consider whether the use of the AI has any disproportionate, unintended, and negative effects on a particular group. For example, we know that facial recognition applications have been less accurate in the case of women and minorities, with algorithms trained on data principally drawn from male images.

- **Professional Codes as Models to Promote Ethical Conduct:** In certain professions, ethical codes apply to research and development of AI. For instance, the National Society of Professional Engineers (NSPE) promulgates a code of ethics with several relevant provisions. "Fundamental Canons" 1 and 6, for example, state: "Engineers, in the fulfillment of their professional duties, shall: (1) hold paramount the safety, health, and welfare of the public. . . . (6) conduct themselves honorably, responsibly, ethically, and lawfully so as to enhance the honor, reputation, and usefulness of the profession."

Note that such language is broad in scope and thus, even where binding, is too vague or general to guide or direct outcomes presenting complex or competing values, as opposed to conduct that is squarely off-code, such as financial self-dealing, which is expressly prohibited in most professional codes.

Specialized fields of engineering also promulgate ethical codes, such as the EC-Council in the case of Certified Ethical Hackers. The most important such code for AI is likely that of the Institute of Electrical and Electronics Engineers (IEEE). The IEEE is the primary professional association for industry and academic professionals in

computer science, electronics, and electrical engineering. The IEEE code of ethics states, in part:

"We, the members of the IEEE, . . . do hereby commit ourselves to the highest ethical and professional conduct and agree:

> 1. To hold paramount the safety, health, and welfare of the public, to strive to comply with ethical design and sustainable development practices, and to disclose promptly factors that might endanger the public or the environment; ...
>
> 5. To improve the understanding by individuals and society of conventional and emerging technologies, including intelligent systems."

To be sure, this is general stuff. "Strive to comply" in an ethics document is like a UN resolution that "calls on" states to take "appropriate" action. The IEEE, however, has done more. In 2016, the IEEE addressed eight issue areas with respect to AI in a document titled "Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems." It addresses numerous AI ethics issues, including transparency, accountability, and black-box components. However, many of the issues remain in question form rather than answer form and are not linked to an enforcement mechanism.

In addition to ethical codes, there are several statements of ethical principles issued by members of the scientific research and civil society groups addressing AI, including its security use. Scholars cite two documents in particular.

The 2017 Asilomar AI principles were developed in conjunction with a conference hosted by the Future of Life Institute at the Asilomar Conference Center in California. Roughly 100 practitioners in science, engineering, and ethics drafted the 23 principles. "Principle 6: Safety" states: "AI systems should be safe and secure throughout their operational lifetime, and verifiably so where applicable and feasible." "Principle 9: Responsibility" states: "Designers and builders of advanced AI systems are stakeholders in the moral implications of their use, misuse, and actions, with responsibility and opportunity to shape those implications."

"Principle 12: Personal Privacy" states: "People should have the right to access, manage, and control the data they generate, given AI systems' power to analyze."

In 2018, a broad coalition of civil society organizations under the umbrella of the Public Voice Coalition issued a statement of Universal Guidelines for Artificial Intelligence in conjunction with the International Conference of Data Protection and Privacy Commissioners. The preamble states:

> "We propose these universal guidelines to guide the design and use of AI. These guidelines should be incorporated into ethical standards, adopted in national law and international agreements, and built into the design of systems. We state clearly that the primary responsibility for the AI systems must reside with those institutions that develop and deploy them."

Twelve guidelines follow, including:

> **(1) Right to Transparency**. All individuals have the right to know the basis of an AI decision that concerns them. This includes access to the factors, the logic, and the techniques that produced the outcome.
>
> **(2) Right to Human Determination**. All individuals have the right to a final determination made by a person.
>
> **(4) Accountability Obligation**. Institutions must be responsible for decisions made by an AI system.
>
> **(9) Cybersecurity Obligation**. Institutions must secure AI systems against cybersecurity threats.

The guidelines have no legal or governmental standing and are not binding. However, like the Asilomar principles, they warrant the attention of policymakers because they:

- Highlight some of the strengths and weaknesses of existing law and ethical codes.
- Offer insight into views within the profession.

- Offer a framework for considering an AI legal regime.

In contrast, the American Bar Association (ABA) Model Rules of Professional Conduct, are generally binding on attorneys through state licensing mechanisms that adopt parallel rules, but these rules are not specifically addressed to AI. Thus, attorneys are required to be diligent and competent and take reasonable measures to keep client information confidential; however, what do "diligent," "competent," and "reasonable" mean when it comes to advising clients on the uses (e.g., discovery) and risks (cyberattack) of AI? It is time now to provide more guidance.

Sound ethical codes and principles can help to identify professional concerns before they become legislative concerns. A prudent legislator or regulator would be wise to watch the ethical debates within the IEEE and ABA to forecast the sorts of issues policymakers and legislators should address with AI.

While complicated, the processes for amending ethical codes, whether at the IEEE or ABA, are faster and more certain than the legislative process. And in any event, policymakers and legislators need to be familiar with the ethical codes, so that law and ethics work in a parallel and complementary manner, rather than at cross-purposes. Most professional codes also include requirements to report codal violations, including safety concerns. One follow-up question for government, industry, and academic policymakers is: Do their institutions provide timely and effective mechanisms to raise design and deployment concerns associated with AI, which means mechanisms people will actually use?

But as we have noted, ethics alone are inadequate to regulate AI.

First, many of the relevant ethical codes do not bind. At best, they guide. "Strive to comply" is not a rigorous standard. Moreover, voluntary codes do not hold up well against prevailing professional incentives. They do not compete well, for example, with the financial incentives of industry or the Everest complex of academia, the desire to be the first to break new ground. Just ask an engineering professor at MIT, Caltech, or Carnegie Mellon: which is the stronger incentive—the IEEE ethical code or the desire to be

first to the patent or publication finish line? Moreover, accountability does not always align with responsibility. With software engineering, for example, coders may be responsible for cyber-vulnerabilities, but it is security officers who are held to account.

Second, ethical codes are usually too general to effectively guide. For example, Asilomar principle 5, "Race Avoidance," states, "Teams developing AI systems should actively cooperate to avoid corner-cutting on safety standards." This reads more like a warning than a principle of conduct. Specific examples are not provided.

Third, where ethical codes are binding on a profession, they tend to reflect the lowest common denominator of agreement, the basement of permitted conduct, not its ceiling, which is noteworthy because it is guidelines and principles where one finds the greatest concern about bias, privacy, and data transparency, not the binding professional codes.

Finally, professional ethical codes accent the views and interests of single disciplines—for example, engineers or lawyers. And they may not represent the views of the profession at large. Rather, they accent the views of the members of the profession who have opted into the professional associations involved and the views of members who are part of the constitutive process of drafting and approving rules.

Institutional Review Boards

The primary procedural mechanism for reviewing the ethical conduct of research at universities is known as an Institutional Review Board (IRB). In the case of federally funded research, grant recipient institutions are required to have an IRB review for any research involving human subjects. Known as the Common Rule, it is binding in the case of federally funded research and applies to research and experimentation involving humans.

The Common Rule is not the only federal rule or mechanism addressing research ethics, nor the only academic rule. Specific, grants and contracts may impose additional ethical requirements.

Government restrictions and regulations also apply to federally funded research involving embryonic stem cells. The “dual use research of concern” (DURC) policy covers life-science (pertaining to living organisms and their products) research that could be used for both benevolent and harmful purposes and thus characterized by the United States government as “dual use research.” Dual use research of concern is defined as “Life sciences research that, based on current understanding, can be reasonably anticipated to provide knowledge, information, products, or technologies that could be directly misapplied to pose a significant threat with broad potential consequences to public health and safety, agricultural crops and other plants, animals, the environment, material, or national security.”

The DURC policy is intended to encourage “a culture of responsibility” and “to ensure that dual use research of concern is identified at the institutional level and risk mitigation measures are implemented as necessary.” These measures include steps to provide for biosafety, physical security, and personnel reliability. The DURC covers research involving 15 high-consequence agents and toxins, such as anthrax. In addition, it covers seven categories of experiments, such as “enhances the harmful consequences of the agent or toxin.” Specific compliance responsibilities are assigned to funding agencies, recipient institutions, and principal investigators (the academic term of art for responsible grant officials).

In addition, the government regulates certain biotechnology applications pursuant to the coordinated framework for regulation of biotechnology. Established in 1986 under the auspices of the Office of Science and Technology Policy (OSTP), within the executive office of the president, this framework addresses biotechnology regulation in academia and industry. The framework links policy to a patchwork of enabling laws assigning responsibilities and authorities to different agencies, such as the FDA, USDA, and EPA, over certain food, drugs, plants, and animals. Although not immediately applicable to AI, the framework is an example of a federal response to emerging technologies and warrants review for lessons applicable to AI. On the positive side, it illustrates what a comprehensive and lasting policy approach might

look like in terms of a government process. On the negative side, the regulation is static. Only listed items are covered. An emerging technology such as AI requires a more fluid response.

As a matter of policy, many universities subject other forms of research to IRB review as well. Some universities, for example, have institutional animal care and use committees, or their equivalent, to address the ethical treatment of animals in research. Some universities subject human embryonic cell research to limitations and oversight beyond that subject to federal funding restrictions. Academic institutions also review the research use of DNA or RNA, biological agents and toxins, radiation, and hazardous materials for compliance with safety, licensing, and training requirements. However, in the absence of government reporting requirements or funding, the system generally depends on self-initiation and reporting. Academia may impose laboratory safety and training requirements as well. Princeton University has a small, unmanned aircraft systems policy (SUAS). The policy: (1) imposes restrictions on the operation of SUAS on university property; (2) prohibits certain persons and certain types of UAS from being flown on university property; and (3) includes an enforcement mechanism referencing federal, state, and local law. The point is universities have discretion to do more when it comes to AI if they choose.

Finally, many universities, or components of universities, require online training through the Collaborative Institutional Training Initiative (CITI). This is a consortium of schools that contribute to the production of online education in "research ethics, compliance, and professional development." The CITI website notes that its training is "used worldwide by over 2,200 organizations and more than 1 million users." The program could serve as an AI platform as well.

The policy questions for academics and government actors are: what type of AI research should be subject to IRB review and is it time for "a common rule for AI?" Such a rule might consider requirements for: (a) a counterintelligence plan; (b) a data management and integrity plan; (c) proof of algorithmic design prior to deployment; (d) a bias mitigation plan; (e) research

parameters and limitations; and (f) declarations of responsibility as to which specific humans would be in, on, or out of the loop.

Corporate Social Responsibility

The private sector drives AI research, development, and deployment. This is true of national security as well as commercial applications. That makes purposeful choices about the regulation of private industry a national security necessity. It also makes corporate social responsibility, especially as it relates to national security, a compelling subject for policy consideration.

Where the law is silent or inadequate, or government policy uncertain, CSR may be the primary source of policy influence to guide corporate behavior. With AI, this might be known as corporate ethics and also corporate *security* responsibility. CSR may derive from patriotism, as was the case with AT&T and electronic surveillance before FISA established a system of court review and orders requiring carrier compliance. It is also true of the traditional defense companies associated with the defense industrial base, where business and patriotism often align. CSR can also derive from a sense of market self-interest, client pressure, altruism, employee pressure, or all four at once.

That was the apparent case with Google's participation in Project Maven. In an open letter to Google CEO Sundar Pichai in August 2018,[1] over 3,000 employees implored the company to cancel the project which they described as "a customized AI surveillance engine that uses 'Wide Area Motion Imagery' data captured by U.S. government drones to detect vehicles and other objects, track their motions, and provide results to the Department of Defense." The letter invoked the company's then-motto — "Don't Be Evil"— and stated, among other things, "This plan will irreparably damage Google's brand and its ability to compete for talent," and, "We cannot outsource the moral responsibility of our technologies to third parties." Google subsequently canceled its participation in the project.

In 2020, Google adopted its own AI Principles, including four categories of "AI applications we will not pursue," including

"technologies that cause or are likely to cause overall harm," a description that like so much else with AI may not take on real meaning without an understanding of the application and its implementation detail. In 2021, some Google employees went a step further by forming the Alphabet Workers Union, not to address employment conditions, but to address the company's societal role and company culture, according to interviews conducted by NPR.[2] Indeed, unions and consumer/citizens groups can and may play an increasing role with AI ethics and CSR.

There is a school of thought that the paramount duty of a corporation is to its shareholders. Directors and officers have a fiduciary duty to the interests of the corporation and must act with a duty of care and in good faith. However, corporations have multiple stakeholders, beyond shareholders, including employees, customers, and the communities in which they work. Corporations, and those who lead corporations, have more discretion in how they define their social and security responsibilities with respect to AI than observers may think.

There is no obligation for corporations to apply CSR principles; however, to the extent they do not do so, they may encounter increased government pressure in the form of unwanted publicity, litigation, and the prospect of legislative or regulatory compulsion.

Perhaps the most visible and serial debate about CSR in national security involves encryption and the going-dark debate. One example is the debate over whether Apple should be compelled to create a means for government to access iPhone data when ordered to do so by a court. What CSR means when it comes to AI will also depend on whether and when U.S. corporations provide AI expertise to authoritarian regimes like China's.

Levers for Influencing Corporate Behavior

Law and government policy can influence corporate conduct in multiple ways. Where the government is a consumer, as with military hardware, it can set standards and impose contractual requirements. However, this type of influence is limited to those companies competing in this space. The government can also

shape corporate practice using the bully pulpit, as illustrated by the 2019 campaign to discredit Huawei as a 5G security risk.

Federal or state law can impose or compel behavior through direct regulation, as in the case of environmental standards, like the California fuel efficiency standards or the federal prohibition on the use of DDT. The government can also condition market access through the licensing and permit process. The CFIUS process is a case in point, where the government can require foreign-owned or foreign-directed corporations to mitigate national security concerns or bar their market entry.

Federal and state law can also incentivize behavior, a form of indirect regulation. This can be done with tax credits, favorable grant and loan terms, or antitrust exemptions, all of which are authorized in the Defense Production Act. The government can also influence behavior through the threat of criminal sanctions, such as those found in export control laws and those regulating financial transactions with governments, entities, and individuals designated by the president under the International Emergency Economic Powers Act and specific sanctions legislation.

Most notably for AI, section 230 of the Communications Decency Act of 1996, with limited exception, exempts social media companies from liability for what is posted on their platforms. One policy question now is whether the law has worked too well, allowing ISPs and social media companies a free pass when it comes to third party conduct on their platforms. A second question is whether the role and responsibility of platforms is best addressed through law, regulation, or self-regulation in the form of CSR policy.

Leverage can also be asserted through regulation of the insurance industry, insurance costs, and limitations on coverage. Litigation may have the same effect, because it may influence the cost of insurance or affect a corporation's reputation and thus market share.

Finally, the government can influence corporate behavior by establishing best practice standards and certifying those practices.

Corporations may have a market and marketing incentive to adopt those standards to validate their products and service quality. Insurance rates, for example, might be linked to government AI certifications.

Going Forward

Ideally, law and ethics work hand in hand; ethics filling the void where law does not, or should not, reach. But there are important differences. Law, in theory at least, is binding and thus removes the element of choice. Ethics, by definition, are non-binding, voluntary, and subject to our moral choices.

Those developing policy to regulate AI must ensure not only that the legal regime is sound but that the ethical framework is as well, providing specific, meaningful, and contextual guidance. Contextual means guidance that comes with effective processes to reflect and assert that guidance, and do so in a transparent manner that can be evaluated and adjusted if need be—for example, how government agencies assert "appropriate human control" over applications or when corporations or academic entities share AI knowledge overseas. Contextual guidance also provides indicative actions, which illustrate how a principle should be applied, not just state what it is. We hope this paper serves as a good starting point.

Author

The Honorable James E. Baker, a CSET Distinguished Fellow, is a professor at Syracuse University College of Law with a courtesy appointment in the Maxwell School. Judge Baker also serves as Director of the Institute for Security Policy and Law. He is the author of *The Centaur's Dilemma: National Security Law for the Coming AI Revolution* (Brookings: 2021).

Acknowledgments

For feedback and assistance, the author would like to thank Chuck Babington, Danny Hague, Mark Hanin, Laurie Hobart, Matt Mittelsteadt, Mark Rosen, Adrienne Thompson and Lynne Weil.

Document Identifier: doi: 10.51593/ 20190022

Relevant Reading

"Artificial Intelligence at Google: Our Principles," Google AI, https://ai.google/principles/.

"Asilomar AI Principles," Future of Life Institute, 2017, https://futureoflife.org/ai-principles/.

"CITI Program: Research, Ethics, and Compliance Training," Collaborative Institutional Training Initiative, https://about.citiprogram.org/en/homepage/.

"Code of Ethics," EC-Council (International Council of Electronic Commerce Consultants), https://www.eccouncil.org/code-of-ethics/.

Communications Decency Act of 1996, 47 U.S.C. § 230 (1934) https://www.govinfo.gov/content/pkg/USCODE-2011-title47/pdf/USCODE-2011-title47-chap5-subchapII-partI-sec230.pdf.

Defense Innovation Board, *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense* (Washington, DC: October 2019), https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF.

"IEEE Code of Ethics," Institute of Electrical and Electronics Engineers, https://www.ieee.org/about/corporate/governance/p7-8.html.

IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, "Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems" (IEEE, 2019), https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf.

Joint Artificial Intelligence Center, *Ethical Principles for Artificial Intelligence* (Washington, DC: Department of Defense, 2020), https://www.ai.mil/docs/Ethical_Principles_for_Artificial_Intelligence.pdf.

"Model Rules of Professional Conduct," American Bar Association, https://www.americanbar.org/groups/professional_responsibility/publications/model_rules_of_professional_conduct/.

National Security Commission on Artificial Intelligence, *Final Report*, (Washington, DC: March 2021), https://www.nscai.gov/wp-content/uploads/2021/03/Full-Report-Digital-1.pdf.

"NSPE Code of Ethics for Engineers," National Society of Professional Engineers, https://www.nspe.org/resources/ethics/code-ethics.

Office for Human Research Protections, *Federal Policy for the Protection of Human Subjects ('Common Rule')* (Washington, DC: U.S. Department of Health and Human Services), https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html.

Office of Science Policy, National Institutes of Health, *United States Government Policy for Institutional Oversight of Life Sciences Dual Use Research of Concern* (Washington, DC: Department of Health and Human Services, 2014), http://www.phe.gov/s3/dualuse/Documents/durc-policy.pdf.

Office of the Director of National Intelligence, *Principles of Artificial Intelligence Ethics for the Intelligence Community* (Washington, DC: 2020), https://www.dni.gov/files/ODNI/documents/Principles_of_AI_Ethics_for_the_Intelligence_Community.pdf.

"Princeton Use of Small Unmanned Aircraft Systems (sUAS) Policy," Princeton University, January 16, 2017, https://drones.princeton.edu/learn-more/policies-and-procedures/princeton-suas-policy.

The Unified Website for Biotechnology Regulation, "About the Coordinated Framework," U.S. Department of Agriculture, U.S. Food and Drug Administration, and U.S. Environmental Protection Agency, https://usbiotechnologyregulation.mrp.usda.gov/biotechnologygov/about.

"Universal Guidelines for Artificial Intelligence," The Public Voice, October 23, 2018, https://thepublicvoice.org/ai-universal-guidelines/.

Endnotes

[1] "Letter to Sundar Pichai Against Project Maven," Google Employees, April 4, 2018, https://static01.nyt.com/files/2018/technology/googleletter.pdf.

[2] Bobby Allyn, "Google Workers Speak Out About Why They Formed A Union: 'To Protect Ourselves,'" *NPR*, January 8, 2021, https://www.npr.org/2021/01/08/954710407/at-google-hundreds-of-workers-formed-a-labor-union-why-to-protect-ourselves; Alina Selyukh, "After Years of Activism, More Than 200 Google Employees Form a Union," *NPR*, January 4, 2021, https://www.npr.org/2021/01/04/953314490/after-years-of-activism-more-than-200-google-employees-form-a-union; Alina Selyukh, "Google Workers Launch Union to Press Grievances With Executives," *NPR*, January 4, 2021, https://www.npr.org/2021/01/04/953198140/google-workers-launch-union-to-press-grievances-with-executives.