GOVERNING AI: THE IMPORTANCE OF ENVIRONMENTALLY SUSTAINABLE AND EQUITABLE INNOVATION

by Henry Gunther and Julietta Rose

Julietta Rose graduated from Berkeley Law in May 2020 and plans to work in finance, cleantech, and renewable energy. Henry Gunther is a third-year law student at Washington University in St. Louis and plans to work in environmental law and policy.

Artificial intelligence (AI) and complex machine learning algorithms have come to play a profound role in many of our day-to-day activities. A quick Google search on the role of AI in everyday life (itself a complex algorithm that decides on the most relevant results)[1] will yield somewhere around 150 million results in about 0.88 seconds. AI can unlock your smartphone,[2] curate your social media feeds,[3] and predict your online purchases to improve shipping and delivery speeds.[4] And all of this happens without your awareness or engagement.

We have finally landed squarely in the age of ubiquitous computing—a stage of computer-society integration first predicted in 1988 by Mark Weiser at the Xerox Palo Alto Research Center, in which computer systems would "vanish into the background," weaving "themselves into the fabric of everyday life until they are indistinguishable from it."[5] Only a small fraction of total computer processing power is actually in the computers we use. Most is now networked with billions of sensors that surround us—from refrigerators to hair dryers, scales to garage door openers, bikes to watches, and of course the smart home systems that link all of your smart appliances together[6]—providing increasingly complete pictures of our everyday lives, activities, and propensities.[7]

This fusion of digital technologies and blurring of the human and digital boundaries is a new form of industrialism. Similar to all other industrial revolutions, these advancements and rapid shifts in productivity are outpacing our understanding of the potential costs and benefits. As AI becomes more prevalent in all areas of life, we need to turn our attention to the interactions between AI and our physical environment, to harness the potential of this technology while avoiding environmental and societal harms. Technological revolutions may fail to materialize, but when they do, they may have unforeseen consequences that leave us little time to prepare.

I. Definitions—Waves of AI

When we think about AI in relation to environmental impacts, it is easy to recall the Volkswagen (VW) emissions scandal and the "defeat devices" programmed to allow their cars to cheat emissions tests.[8] It is hard to forget such an intentionally designed violation of the law. What makes the VW emissions scandal so captivating is that human programmers are easy to blame. We can almost picture the engineers working in secret to design the device with malintent.

However, the VW devices are far from the sophisticated algorithms and machine learning systems that have become common in the field of AI. These devices would fall under the "first wave" of AI. The truth is that the consequences and potential harms from AI are far more complicated than the VW scandal, and culpability for harm is much more difficult to assign.[9]

An example of the widespread *unintended* environmental impacts of algorithmic failure dates back to the study of ozone depletion that led to the Montreal Protocol. Scientists had designed algorithms to monitor and record ozone

1. Google, How Search Algorithms Work, https://www.google.com/search/howsearchworks/algorithms/ (last visited Sept. 15, 2020).
2. Apple, *About Face ID Advanced Technology*, https://support.apple.com/en-us/HT208108 (last visited Sept. 15, 2020).
3. Paige Cooper, *How the Facebook Algorithm Works in 2020 and How to Make It Work for You*, Hootsuite, Jan. 27, 2020, https://blog.hootsuite.com/facebook-algorithm/.
4. Alina Selyukh, *Optimized Prime: How AI and Anticipation Power Amazon's 1-Hour Deliveries*, NPR, Nov. 21, 2018, https://www.npr.org/2018/11/21/660168325/optimized-prime-how-ai-and-anticipation-power-amazons-1-hour-deliveries.
5. Stanford University, *Ubiquitous Computing*, https://web.stanford.edu/dept/SUL/library/extra4/weiser/ubiq.html (last updated Sept. 15, 2020).
6. John R. Delaney & Angela Moscaritolo, *What Is a Smart Home Hub (And Do You Need One)?*, PCMag, July 13, 2020, https://www.pcmag.com/news/what-is-a-smart-home-hub-and-do-you-need-one.
7. Matt Burgess, *What Is the Internet of Things? WIRED Explains*, Wired, Feb. 16, 2018, https://www.wired.co.uk/article/internet-of-things-what-is-explained-iot.
8. Guilbert Gates et al., *How Volkswagen's "Defeat Devices" Worked*, N.Y. Times, Mar. 16, 2017, https://www.nytimes.com/interactive/2015/business/international/vw-diesel-emissions-scandal-explained.html.
9. Andrew Tutt, *An FDA for Algorithms*, 69 Admin. L. Rev. 83 (2017), *available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2747994.

The term **"artificial intelligence"** (AI) can be used to refer to many different types and complexities of computing systems. The Defense Advanced Research Projects Agency (DARPA) defines "AI" as the "programmed ability to process information," and has identified three waves, or stages, of AI system complexity:

First wave. Computer engineers "create sets of rules to represent knowledge in well-defined domains." In these types of systems, "the structure of the knowledge," and the inferential processing steps, are "defined by humans," while "the specifics are explored by the machine." These systems are not capable of learning new methods of information processing or integrating information that developers do not explicitly include. Early autonomous vehicles are an example of the first wave of AI.

Second wave. This level of AI system is defined by statistical learning abilities, where "engineers create statistical models for specific problem domains and train them on data." In other words, these systems can be trained to learn new ways to process information that were not explicitly programmed by developers. While these systems have "nuanced classification and prediction capabilities," they lack "contextual capability and [have] minimal reasoning ability." Handwriting recognition, facial detection, and language processing systems are all examples of the second wave of AI.

Third wave. Developers are working on a new, more sophisticated generation of AI systems that will be able to create "contextual explanatory models for classes of real-world phenomena." This type of system will be able to not only perceive and learn, but abstract and reason from prior inputs to generate new conceptual categories, and explain why it developed those categories.

Note: All citations in this section are to John Launchbury, Director I2O, DARPA, Presentation: A DARPA Perspective on Artificial Intelligence, https://www.darpa.mil/attachments/AIFull.pdf.

levels detected by (occasionally unreliable) satellites, and to disregard data reported by the satellites that were "obviously" inaccurate because it diverged too strongly from long-standing predictions. When ozone levels dropped far below what they were generally predicted to be, the algorithm wrote off the results as faulty data and did not flag the levels for researchers.[10] The algorithm designed to monitor the data was not inclusive enough.

Of course, we know now that those readings were in fact accurate, and it took a team of researchers on the ground to record and highlight what a computer had disregarded. A mistake made by an algorithm, based on a good-faith assumption by programmers, delayed the response to a massive global environmental problem and nearly under-

mined the success of one of the most effective international environmental law treaties ever signed. But like the VW scandal, these algorithms would be considered "first wave." They show that if we are not careful, the acceleration of AI without an accompanying environmental governance structure in some form could cause unintentional and unforeseen consequences.

This Comment is written not with the intent to stall research and innovation within the field of AI, but to encourage it. The promise of AI to address environmental issues such as climate change is a real source of hope, given the increasing need to act.[11] As Bill McKibben has noted, "We're running out of options, and we're running out of decades."[12] Researchers have identified dozens of problems in which AI offers promising solutions to mitigate the effects of climate change, from accelerating material science in order to develop stronger energy storage to precision agriculture that reduces the use of pesticides and other chemicals.[13]

Scientists are also utilizing AI in climate change adaptation, such as making climate predictions and modeling climate change effects.[14] The information and conclusions drawn by AI would enable policymakers to help communities adapt to climate change in vulnerable places like Alaska,[15] Dominica, Puerto Rico, Sri Lanka, and more.[16]

II. Double-Edged Sword

However, AI systems themselves—the hardware they run on and the software they run—can pose risks to the environment. Training a single AI system can produce the same greenhouse gas (GHG) emissions as five cars will in their lifetime.[17] Massachusetts Institute of Technology (MIT) researchers have developed ways to reduce this impact, but more efforts should be made to further improve the efficiency of these machines, like moving to all-renewable

10. *What Did NASA Know? And When Did They Know It?*, Real Climate, Dec. 24, 2017, http://www.realclimate.org/index.php/archives/2017/12/what-did-nasa-know-and-when-did-they-know-it/.

11. *Summary for Policymakers, in* Global Warming of 1.5°C. An IPCC Special Report on the Impacts of Global Warming of 1.5°C Above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty (Valérie Masson-Delmotte et al. eds., Intergovernmental Panel on Climate Change 2018), https://www.ipcc.ch/sr15/chapter/spm/.

12. Bill McKibben, *A Very Grim Forecast*, N.Y. Rev. Books, Nov. 22, 2018.

13. *Climate Change AI, Interactive Summary*, https://www.climatechange.ai/summaries (last visited Sept. 15, 2020).

14. Jackie Snow, *How Artificial Intelligence Can Tackle Climate Change*, Nat'l Geographic, July 18, 2019, https://www.nationalgeographic.com/environment/2019/07/artificial-intelligence-climate-change/; Nicola Jones, *How Machine Learning Could Help to Improve Climate Forecasts*, Nature Mag., Aug. 23, 2017, https://www.scientificamerican.com/article/how-machine-learning-could-help-to-improve-climate-forecasts/.

15. Alex DeMarban, *Young Alaskans Suing State Press for Reductions of Emissions Causing Climate Change*, Anchorage Daily News, Apr. 30, 2018, https://www.adn.com/alaska-news/environment/2018/04/30/young-alaskans-suing-state-press-for-reductions-of-emissions-causing-climate-change/.

16. David Eckstein et al., Germanwatch, Global Climate Risk Index 2019 (2018), https://germanwatch.org/files/Global%20Climate%20Risk%20Index%202019_2.pdf.

17. Karen Hao, *Training a Single AI Model Can Emit as Much Carbon as Five Cars in Their Lifetimes*, MIT Tech. Rev., June 6, 2019, https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/.

energy sources.[18] We must keep in mind that AI is a sophisticated tool that requires its own infrastructure and support, and comes with its own inherent costs beyond the outcomes it is able to achieve.

Additionally, AI is not exclusively a tool for environmentalists. The fossil fuel industry is investing heavily in AI research in order to improve the efficiency of their extraction and production. Analysts predict those investments may reach as high as $3 billion by 2022.[19] AI-enabled increases in fossil fuel production could negate the environmental gains from applying AI to areas such as grid management, energy efficiency, or carbon capture.

Fossil fuel companies could certainly compete for talented specialists in the AI field who are already in high demand.[20] While the rise in demand for machine learning engineers has led to a significant increase in certification programs, there are fewer than 20,000 certified machine learning engineers worldwide.[21] Though automated machine learning could eventually reduce the demand for trained data scientists, there will remain a subset of critical tasks that will resist automation.[22] Additionally, the powerful fossil fuel industry could lobby to weaken regulations surrounding the use of AI in their industry, and, if successfully utilizing AI as a resource, reduce the cost of fossil fuels and undermine the efforts of making alternative energy more affordable and available.

Beyond just the fossil fuel industry, AI's ability to increase efficiency can be harnessed by those in the business of resource extraction more broadly. Mining,[23] fishing,[24] and timber harvesting[25] are all examples of areas where AI is being implemented to make the industry more profitable and efficient. While AI has the ability to increase resource extraction and processing rates, it also has the ability to balance the interests of environmentalists and those within the industry to ensure that they remain profitable while protecting natural resources.

III. Application to Shipping

One industry becoming even more essential with the onset of the COVID-19 health pandemic is transportation and shipping. With typical in-person shopping limited or unavailable, consumers are turning to online ordering and shipping to meet their needs. While the rate of online shopping is beginning to fall again with some brick-and-mortar establishments reopening, online sales in May were still up 66% year over year.[26]

AI can play a critical role in the shipping of these items. After teaming up with Microsoft AI, a Hong Kong shipping company was able to generate $10 million in savings per year by tapping into the power of AI to efficiently plan and route shipping.[27] However, if the AI systems used to produce these savings are only considering the financial impact, there remains a missed opportunity to harness the potential of this technology to create next-generation logistics systems that are fast, efficient, and environmentally sound.

Suppose the system discovers a shortcut in a shipping route and begins to route cargo through a new channel in the ocean. Did the AI system consider the impact this change will have on the ecosystem of that area? What if the new route travels through a national marine sanctuary? Cargo shipping can have significant impacts on marine life and coastal habitats.[28] Additionally, do these new routes result in increased GHG emissions or other pollution?

What if, instead of just saving the company millions, the AI system could vary the routes of cargo ships to better protect the ocean's ecosystems? Better yet, suppose the company teamed up with scientists who use AI to track sensitive marine life, and used this data to determine new shipping routes based on the real-time movement and location of sensitive marine species. There is tremendous opportunity for AI to balance the interests of companies and environmentalists to create win-win situations in areas that are difficult to foster collaboration. However, the incentive to not consider the environment in these decisions must be overcome if the full potential of AI environmental protection is to be realized.

With all of AI's promise, we must keep in mind that it is merely a tool. Used intentionally and aimed at the right problems, AI can provide a part of the solution to some of the world's "wicked" environmental problems. However, AI can also be wielded in ways that are detrimental to the environment. Without appropriate governance in some format, AI could have serious consequences, both intentionally and unintentionally.

18. Rob Matheson, *Reducing the Carbon Footprint of Artificial Intelligence*, MIT News, Apr. 23, 2020, https://news.mit.edu/2020/artificial-intelligence-ai-carbon-footprint-0423.

19. Shane Randolph & Jim McBride, *AI & Machine Learning: The Next Transformation for Oil & Gas*, Opportune, Jan. 2019, https://opportune.com/Energy-Sector-Insights-Events/Insights/AI-Machine-Learning-The-Next-Transformation-for-Oil-Gas/.

20. Apteo, *Data Science Career Trends in 2020*, https://app.apteo.co/workspaces/2300510291097552329 (last visited Sept. 15, 2020).

21. Frederik Bussler, *Why You Should Get Google's New Machine Learning Certificate*, Towards Data Sci., July 20, 2020, https://towardsdatascience.com/why-you-should-get-googles-new-machine-learning-certificate-56af4204744f.

22. Holger Hürtgen et al., *Rethinking AI Talent Strategy as Automated Machine Learning Comes of Age*, McKinsey & Co., Aug. 14, 2020.

23. Zeshan Hyder et al., *Artificial Intelligence, Machine Learning, and Autonomous Technologies in Mining Industry*, 30 J. Database Mgmt. 67 (2019), *available at* https://www.researchgate.net/profile/Keng_Siau/publication/335542028_Artificial_Intelligence_Machine_Learning_and_Autonomous_Technologies_in_Mining_Industry/links/5e90eedca6fdcca7890a3c56/Artificial-Intelligence-Machine-Learning-and-Autonomous-Technologies-in-Mining-Industry.pdf.

24. Melina Kourantidou, *Artificial Intelligence Makes Fishing More Sustainable by Tracking Illegal Activity*, Conversation, July 12, 2019, https://phys.org/news/2019-07-artificial-intelligence-fishing-sustainable-tracking.html.

25. Scot McQueen, *How Artificial Intelligence, Robots Enhance Forest Sustainability in Finland*, ESRI Blog, December 16, 2019, https://www.esri.com/about/newsroom/blog/finland-enhances-forest-data-accuracy-for-automation/.

26. April Berthene, *Online Merchants Gain an Extra $107 Billion in 2020 Thanks to Pandemic*, Digital Com. 360, Sept. 14, 2020, https://www.digitalcommerce360.com/article/coronavirus-impact-online-retail/.

27. PacificGreen Technologies Group, *How Will AI Change the Shipping Industry*, https://www.pacificgreentechnologies.com/articles/how-will-ai-change-shipping-industry/ (last visited Sept. 15, 2020).

28. Laura Poppick, *Animals in Marine Sanctuaries Not Immune to Human Impact*, Live Sci., Nov. 1, 2013, https://www.livescience.com/40887-national-marine-sanctuaries-stress-top-predators.html.

IV. A Problem in the Commons

An immense amount of scholarship has emerged around Garret Hardin's famous 1968 *The Tragedy of the Commons*. It is beyond the scope of this Comment to dive into this literature to any significant degree, but interesting to point out that Hardin begins his famous article by stating that the commons, like nuclear war, is a "no technical solution problem," and goes on to argue that top-down governance or privatizing the resource is the only solution.[29] AI is proving this conclusion now to be somewhat, if not entirely, incorrect.

The examples of AI for "commons management" are becoming more numerous and a source of hope in addressing common resources. Mine operators use AI sensors to monitor the potential for environmental harms like the spilling of waste material or leaking of hazardous gases.[30] Human observers are being replaced by AI monitoring that can fuel cost-effective fisheries management and tackle illegal fishing more efficiently.[31] Microsoft's AI for Earth has been utilized by researchers to develop ways to more precisely and efficiently assess the status of forests and provide data-driven management that is crucial to carbon sequestration.[32] Google's DeepMind AI, using machine learning to train neural networks to predict the most efficient energy consumption in their data centers, was able to reduce the energy used to cool their centers by 40%.[33]

V. The State of Play

With all of these promising avenues for AI to take, it is tempting to incentivize this work. And indeed, those involved in deep decarbonization efforts are suggesting doing that by, for example, opening up expedited regulatory pathways for these systems to begin addressing environmental concerns.[34] The government is responding to calls for supporting AI development but not to calls for addressing environmental concerns.

A bill before the U.S. Congress called the Endless Frontiers Act will, if passed, grant an additional $100 billion over five years to the National Science Foundation (NSF).[35] The bill would create a technology directorate to ensure that the United States remains a leader in tech. The bill's sponsors believe the country that makes key advances in advanced technology like AI will be the "superpower of the future."[36]

This is a timely bill, as the NSF is set to be the organization coordinating the American Artificial Intelligence Initiative. This initiative is designed to increase research and development of AI so that the United States is a leader in the field.[37] Recently, NSF and the U.S. Department of Agriculture (USDA) announced a $140 million investment to establish a nationwide network of seven new AI research institutes, "to accelerate research, expand America's workforce and transform society for the future."[38]

However, the actions so far taken by the government have not included an intentional effort to consider environmental impacts or address environmental concerns. In fact, the initiatives do not pay much attention to AI governance in general and focus more on removing obstacles to AI innovation and implementation, favoring approaches that enable rather than potentially constrain development and commercialization.[39] Some of the areas identified by the Office of Management and Budget as being open to advancement through AI innovation are transportation, healthcare, manufacturing, financial services, agriculture, weather forecasting, and national security and defense.[40] Some of these areas are large producers of GHGs, and if environmental concerns could be addressed alongside the promotion of AI development, it could push the United States closer toward a more sustainable future.

Within the American Artificial Intelligence Initiative, the Donald Trump Administration instructed federal agencies to include AI research and development in their budgets. The U.S. Environmental Protection Agency (EPA) especially should take this directive as an opportunity to bolster AI environmental governance. Additionally, the initiative does make a general recognition of "challenges [that] lack near term market drivers" as well as other areas in AI development that "address agency-specific requirements or societal needs that are unlikely to lead to commercial profits."[41] Environmental issues certainly fall within this category, and it will be interesting to see if they receive the kind of attention and funding promised by the initiative.

29. Garrett Hardin, *The Tragedy of the Commons*, 162 Science 1243 (1968), *available at* https://pages.mtu.edu/~asmayer/rural_sustain/governance/Hardin%201968.pdf.

30. Ron Schmelzer, *AI Helping Extract Value in the Mining Industry*, Forbes, Aug. 9, 2019, https://www.forbes.com/sites/cognitiveworld/2019/08/09/ai-helping-extract-value-in-the-mining-industry/#3cb23d907006.

31. Kourantidou, *supra* note 24.

32. Lucas Joppa, *AI Can Be a Game-Changer for the World's Forests. Here's How*, World Econ. F., Sept. 11, 2018, https://www.weforum.org/agenda/2018/09/ai-can-help-us-map-forests-down-to-each-tree/.

33. Richard Evans & Jim Gao, *DeepMind AI Reduces Google Data Centre Cooling Bill by 40%*, DeepMind, July 20, 2016, https://deepmind.com/blog/article/deepmind-ai-reduces-google-data-centre-cooling-bill-40.

34. Legal Pathways to Deep Decarbonization in the United States: Summary and Key Recommendations (Michael B. Gerrard & John C. Dernbach eds., Envtl. L. Inst. 2018), *available at* https://www.eli.org/eli-press-books/legal-pathways-deep-decarbonization-united-states-summary-and-key-recommendations.

35. Jeffrey Mervis, *U.S. Lawmakers Unveil Bold $100 Billion Plan to Remake NSF*, Science, May 26, 2020, https://www.sciencemag.org/news/2020/05/us-lawmakers-unveil-bold-100-billion-plan-remake-nsf.

36. *Id.*

37. Office of Science and Technology Policy, The White House, American Artificial Intelligence Initiative: Year One Annual Report (2020), https://www.whitehouse.gov/wp-content/uploads/2020/02/American-AI-Initiative-One-Year-Annual-Report.pdf.

38. Press Release, USDA National Institute of Food and Agriculture, USDA-NIFA and NSF Establish Nationwide Network of Artificial Intelligence Research Institutes (Aug. 26, 2020), https://www.nifa.usda.gov/press-release/artificial-intelligence-research.

39. *See, e.g.*, Exec. Order No. 13859, Maintaining American Leadership in Artificial Intelligence, 84 Fed. Reg. 3967 (Feb. 14, 2019), *available at* https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/.

40. Identifying Priority Access or Quality Improvements for Federal Data and Models for Artificial Intelligence Research and Development (R&D), and Testing; Request for Information, 84 Fed. Reg. 32962 (July 10, 2019), *available at* https://www.govinfo.gov/content/pkg/FR-2019-07-10/pdf/2019-14618.pdf.

41. Office of Science and Technology Policy, *supra* note 37.

However, we must be careful about the trust we place in AI to solve problems for us and the speed with which we move toward these AI-driven solutions. Unlike the VW emissions scandal, AI can have consequences that are unintended and hard to predict, even when trained and programmed by individuals with the best intentions. The example with ozone monitoring was a mistake, made by the programmer, that resulted in an underinclusive algorithm. As AI technology has grown more complex, so have the ways in which AI can "make mistakes." To understand how this can happen, we need to understand the normative nature of AI systems.

VI.   AI Systems Are Normative Systems

Algorithms and AI, contrary to first appearances, make judgments, evaluations, and value comparisons. In other words, they are normative systems. Even though computers and the programs they run are "just" machines chugging through logical operations on 1s and 0s, when humans design and use algorithms, we imbue their symbols with meaning, and their operations with priorities and values. When an AI system ultimately makes a decision and implements it—for instance, about which route a ride-share driver should take to a destination—if developers have done their jobs correctly, the decision is the one that "should" be made in that situation, given all of the relevant data and the goals of the developers and users.

The values and priorities of AI systems are often perfectly consistent with the intentions, expectations, and background assumptions of developers and users, and can therefore be invisible. However, cases of algorithmic bias—when an AI produces a startling, unexpected, and discriminatory result that embodies values that we disavow as a society—throw these normative features of AI into relief, putting the values and priorities of AI systems on display.

For example, in 2010, a Nikon digital camera equipped with a new "blink detector" algorithm kept asking "did someone blink?" when an Asian person was in the photo.[42] In 2015, a user discovered that Google Photos was cataloguing photos of Black people as gorillas.[43] As recently as 2019, investigators discovered that an algorithm used by hospitals to help determine how to allocate medical resources and care was diverting medical care from Black patients in favor of white patients.[44] Earlier this year, a Black man was arrested on the basis of a false positive identification from facial recognition software used by police; it was later discovered that the software gives more false positives for Black faces than white faces.[45]

So, how can an algorithm so dramatically diverge from the intentions of its developers and users? Consider a hypothetical version of the algorithm used by hospitals to allocate medical care, mentioned above. Although race is not included in the types of data this hypothetical algorithm draws on, information on whether pain medication had been previously prescribed is included. This may seem like a non-normative and uncontroversial data set—after all, significant pain can be an important indicator of the severity of an injury or illness, and a prior prescription for pain medication is an intuitive indicator of whether pain is severe. Inclusion of the pain medication data set will cause the algorithm to allocate medical resources by prioritizing those patients with prior pain medication over those without (all else being equal).

However, this banal-seeming prioritization rationale will be foregrounded when the algorithm starts to prioritize white patients over Black patients in distribution of medical resources, because the pain medication data set is itself racially biased: a meta-analysis over 20 years of studies "covering many sources of pain in numerous settings found that Black/African American patients were 22% less likely than white patients to receive any pain medication."[46] This is not due to actual pain felt by patients, but to implicit and explicit biases held by medical professionals. When a patient's history of receiving pain medication in the past tracks the patient's race, and then that same data is used to allocate medical resources, the unintended result will be that medical resources will be at least partially allocated based on race. This is one example of how the normativity of the system and its data sets will be highlighted only when the algorithm's judgments diverge from the ones we would normally endorse.

Every decision that an AI developer makes about which data should be considered, and how, contains an implicit normative assumption about the relevance and reliability of that data to the decision.[47] For instance, the algorithms that operate Uber and Lyft consider data and variables relevant to the economic efficiency of routes and the relative availability of drivers near a rider's location; it is unlikely, however, that they consider the GHG emissions of possible routes in determining how to direct their drivers, and this failure to consider GHG emissions means that levels of emissions are effectively irrelevant.[48] The result is that

42. Adam Rose, *Are Face-Detection Cameras Racist?*, Time, Jan. 22, 2010, http://content.time.com/time/business/article/0,8599,1954643,00.html.

43. Jessica Guynn, *Google Photos Labeled Black People "Gorillas,"* USA Today, July 1, 2015, https://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/.

44. Carolyn Y. Johnson, *Racial Bias in a Medical Algorithm Favors White Patients Over Sicker Black Patients*, Wash. Post, Oct. 24, 2019, https://www.washingtonpost.com/health/2019/10/24/racial-bias-medical-algorithm-favors-white-patients-over-sicker-black-patients/.

45. Kashmir Hill, *Wrongfully Accused by an Algorithm*, N.Y. Times, Aug. 3, 2020, https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html.

46. Janice A. Sabin, *How We Fail Black Patients in Pain*, Ass'n Am. Med. Colleges, Jan. 6, 2020, https://www.aamc.org/news-insights/how-we-fail-black-patients-pain.

47. Joshua Adams, *Objective Algorithms Are a Myth*, OneZero, June 30, 2020, https://onezero.medium.com/objective-algorithms-are-a-myth-22b2c3e3d702.

48. Tony Barboza, *Taking an Uber or Lyft Pollutes More Than Driving, California Finds. Next Stop: Regulations*, L.A. Times, Mar. 7, 2020, https://www.latimes.com/environment/story/2020-03-07/uber-lyft-ride-hailing-air-pollution-greenhouse-gas-emissions; Andrew J. Hawkins, *Uber and Lyft Generate 70 Percent More Pollution Than Trips They Displace: Study*, Verge, Feb. 25, 2020, https://www.theverge.com/2020/2/25/21152512/uber-lyft-climate-change-emissions-pollution-ucs-study; Elijah Chiland, *Report: Uber, Lyft Trips Produce More Emissions Than Regular Car Trips*, Curbed L.A., Jan. 16, 2020, https://la.curbed.com/2020/1/16/21069214/uber-lyft-emissions-california-goals; Jose B. Tengco & Kathryn Phillips, *The Hidden Environmental Threat From Uber and Lyft*, CalMatters, Sept. 4, 2019, https://calmatters.org/uncategorized/2019/09/uber-greenhouse-gas/; Ben

widespread use of Uber and Lyft has foreseeably led to increases in GHG emissions.[49]

Mistaken evaluations of a data set's or variable's relevance or reliability can result in AI systems that produce discriminatory results; they can also produce results that are more broadly inconsistent with the values of developers, users, stakeholders, or members of the community—for instance, with the value of environmental conservation. Mistakes can arise because of a

- lack of diversity on development teams;
- failure to train the AI system on broad, representative data sets;
- failure to consult users and stakeholders throughout the development process;
- limited testing; and/or
- failure on the part of development teams to consider, at the outset of the development process, whether the data, variables, decision parameters, and so on, that they have identified are implicitly discriminatory, biased, or incomplete in some way.[50]

Given the scale and complexity of the problem, it is overwhelmingly likely that climate modeling and prediction and climate adaptation planning will rely on AI systems; as one researcher explained, "climate is now a data problem."[51] Consider a second hypothetical AI system, a climate adaptation decision support system. This system aggregates reams of data on climatic conditions (both present and predicted); geological features; county borders and water rights; the location, use, and economic value of various forms of infrastructure; socioeconomic data including property boundaries and market values of land, buildings, and other improvements; the location and long-term economic productivity of commercial and industrial centers ranging from town centers to shopping malls to industrial parks; and more. This algorithm will be used to help government decisionmakers determine which locations should be hardened to climate change threats through building seawalls, thermal shielding, and the like, and which locations should be abandoned. In other words: given limited resources, which communities are worth saving, and which are not?

Property values, both current and projected, are a significant chunk of the data the algorithm works on. However, residential property values can be racially biased; a recent study by the Brookings Institution concluded that after controlling for neighborhood quality, home quality, and local amenities, homes in majority Black neighborhoods

are worth "23 percent less ($48,000 per home on average, amounting to $156 billion in cumulative losses) . . . compared to those with very few or no black residents."[52] If the algorithm uses a cost-benefit model based on property values, current or projected, to allocate climate adaptation resources, the result will be that the algorithm's relocation recommendations will track the racial and class makeup of neighborhoods, with majority-Black neighborhoods more likely to be relocated and majority-white neighborhoods more likely to be protected.

Because Black Americans have much lower levels of family and intergenerational wealth than white Americans, the burdens of relocation will not only be racially distributed, but will also fall on those less able to afford them.[53] Thus, a superficially race-neutral climate adaptation decision support system could easily fuel racially discriminatory policy, with huge environmental and racial justice consequences—all because developers and stakeholders did not adequately interrogate the relevance and reliability of data sets, variables, and decision rationales.

Public awareness of discriminatory bias by AI systems is growing and the legal community is taking note, especially in criminal contexts.[54] However, there is still little discussion of the role that *environmental values* should play in the decisionmaking and decision-supportive AI systems that will become ubiquitous in our everyday life.[55]

VII. Governing AI

The normative nature of AI systems must not be forgotten in the rush to support AI innovation and implementation. The biases that get built into these systems can have a significant impact on many areas of life, and with the ubiquitous and invisible nature of AI, we might not realize these effects until after its integration. Governance of AI is crucial to mitigate the direct and indirect effects on environmental and societal systems.

Just as Congress saw the need in the 1960s to incorporate environmental concerns into federal decisionmaking by passing the National Environmental Policy Act (NEPA),[56] so too must we recognize the need to consider the environmental impacts of AI innovation now. A governance framework that intentionally addresses environmen-

Geman, *Uber and Lyft's Rise May Be Fueling Climate Change*, Axios, Feb. 27, 2020, https://www.axios.com/uber-lyft-ride-sharing-climate-change-172d9da3-784a-4257-9f45-e173c257971e.html.

49. *Id.*

50. *See* Brent Daniel Mittelstadt et al., *The Ethics of Algorithms: Mapping the Debate*, Big Data & Soc'y 1-21 (2016), https://journals.sagepub.com/doi/full/10.1177/2053951716679679; Kate Crawford & Ryan Calo, *There Is a Blind Spot in AI Research*, Nature, Oct. 13, 2016; Fairness, Accountability, and Transparency in Machine Learning, *Scholarship*, https://www.fatml.org/resources/relevant-scholarship (last visited Sept. 15, 2020) (offering a bibliography of work on algorithmic bias).

51. Jones, *supra* note 14.

52. Andre Perry et al., Metropolitan Policy Program at Brookings & Gallup, The Devaluation of Assets in Black Neighborhoods: The Case of Residential Property (2018), https://www.brookings.edu/wp-content/uploads/2018/11/2018.11_Brookings-Metro_Devaluation-Assets-Black-Neighborhoods_final.pdf.

53. Kriston McIntosh et al., *Examining the Black-White Wealth Gap*, Brookings, Feb. 27, 2020, https://www.brookings.edu/blog/up-front/2020/02/27/examining-the-black-white-wealth-gap/.

54. Andrea Roth, *Machine Testimony*, 126 Yale L.J. 1972, 1976 (2017); Andrea Roth, *Trial by Machine*, 104 Geo. L.J. 1245 (2016).

55. Some notable counterexamples include Dave Rejeski et al., Environmental Law Institute, When Software Rules: Rule of Law in the Age of Artificial Intelligence (2018); Fairness, Accountability, and Transparency in Machine Learning, *supra* note 50 (offering a bibliography of work on algorithmic bias).

56. Linda Luther, Congressional Research Service, The National Environmental Policy Act: Background and Implementation (2008), https://fas.org/sgp/crs/misc/RL33152.pdf; 42 U.S.C. §§4321-4370h, ELR Stat. NEPA §§2-209.

tal concerns surrounding the widespread use of AI is vital in utilizing AI for good while avoiding detrimental outcomes. This governance framework could, and probably ought to, take many forms.

One thing that AI governance can provide is a valuable perspective on the scope of problems that AI can solve. While this Comment details many of AI's promises for the environment, it is important too to keep in mind the limits of the technology. As the famous "law of the instrument" explains, when you have a hammer, every problem looks like a nail; so too may it be tempting to wield AI in every difficult environmental problem. We hope here to highlight both the potential of AI and the risks involved with no governance structures.

VIII. Private Governance

Calls for regulating AI have been made both by lawyers and AI developers themselves.[57] However, much of the "regulatory" activity with respect to the environment is currently governance within the private sphere.[58] It has been decades since a major environmental law has been passed, and while effective AI environmental governance legislation would be ideal, it seems unlikely that a sort of NEPA for AI will be passed in the near future—though there have been calls for a "[Food and Drug Administration] for algorithms."[59] Understanding how governance can be enhanced within the private sphere is an important piece of AI environmental governance.

The examples of industry AI use provided above outline many avenues for private AI development to address large environmental problems and highlight actions that are already underway. However, environmental protection is not currently a large consideration in AI development in general, and is certainly not obviously connected to profitability. It is rare for tech startups to focus on climate science, because the return on investment is so much higher in other fields.[60]

Environmental concerns fall within what the American Artificial Intelligence Initiative would likely deem "challenges [that] lack near term market drivers."[61] Coupled with the potential profit gain in using AI to extract more efficiency from legacy production processes or to exploit common resources, there is a strong incentive for private industry to underutilize AI's promise for addressing environmental concerns and underplay its environmental dangers—or even to "fake" private governance while discouraging government involvement, as may be the case in

antitrust enforcement.[62] Efforts should be made to increase the incentive for the sustainable use of AI such that the gains of short-term exploitation do not outweigh the benefits of long-term sustainability and preservation.

One interesting example aimed at addressing data and knowledge gaps, the Rockefeller Foundation's Lacuna Fund launched in July 2020, is designed to fill in gaps in important data sets that are "outdated, missing key information, or not representative of underserved populations—if they exist at all—resulting in biases and lower accuracy."[63] Google.org, the Rockefeller Foundation, and Canada's International Development Research Centre are the first funders; the first three initiatives will focus on natural language processing, agriculture, and health data related to COVID-19 and respiratory illness, with plans to expand the program.[64]

Such incentives might additionally include awards, similar to ENERGY STAR, that recognize the sustainable use of AI or recognition of industry leaders in managing the environmental risks of AI. If these incentives are properly balanced, AI could be the technical solution to manage the interests that once made Hardin deem common resource management a "no technical solution" problem.

IX. An Office of Algorithmic Governance

One possible route to achieve uniform AI governance *and* utilize AI's potential for environmental protection is for EPA to create an office to monitor and govern AI. With the increase in funding from the American Artificial Intelligence Initiative and a dedicated team, many of the concerns outlined here, as well as others that might be realized, could be addressed.

EPA has already recognized the advantages of AI in the environmental sphere. AI has been deployed to address a multitude of environmental concerns. Agency researchers have used AI to mimic human behavior in order to assess Americans' exposure to chemicals.[65] The AirNow program's sensors and algorithms report daily air quality in order to warn residents in areas where air quality is becoming hazardous.[66] Additionally, the Waste Reduction Algorithm (WAR) was designed to factor in environmental considerations in the production of chemicals. The algorithm can weigh different potential impacts like global warming potential or risk of human exposure.[67]

57. Tutt, *supra* note 9; Allison Matyus, *Elon Musk Warns That All A.I. Must Be Regulated, Even at Tesla*, Digital Trends, Feb. 18, 2020, https://www.digitaltrends.com/news/elon-musk-warns-artificial-intelligence-musk-be-regulated/.

58. Michael P. Vandenbergh, *Private Environmental Governance*, 99 Cornell L. Rev. 129 (2013), *available at* https://scholarship.law.cornell.edu/clr/vol99/iss1/3.

59. Tutt, *supra* note 9.

60. Christian Schroeder De Witt, *Climate Change: The Case for Artificial Intelligence*, Oxford Student, Feb. 17, 2019, https://www.oxfordstudent.com/2019/02/17/climate-change-the-case-for-artificial-intelligence/.

61. Office of Science and Technology Policy, *supra* note 37.

62. Jared Soares, *Big Tech Funds a Think Tank Pushing for Fewer Rules. For Big Tech.*, N.Y. Times, July 24, 2020, https://www.nytimes.com/2020/07/24/technology/global-antitrust-institute-google-amazon-qualcomm.html.

63. *Rockefeller Launches Funder Collaborative to Fill Global Data Gaps*, Philanthropy News Dig., July 14, 2020, https://philanthropynewsdigest.org/news/rockefeller-launches-funder-collaborative-to-fill-global-data-gaps.

64. *Id.*

65. U.S. EPA, *EPA Researchers Use AI to Mimic Human Behaviors That Could Affect Our Exposure to Chemicals*, https://www.epa.gov/sciencematters/epa-researchers-use-ai-mimic-human-behaviors-could-affect-our-exposure-chemicals (last updated Mar. 14, 2019).

66. AirNow, *How Is the NowCast Algorithm Used to Report Current Air Quality?*, https://www.airnow.gov/faqs/how-nowcast-algorithm-used-report/ (last visited Sept. 15, 2020).

67. U.S. EPA, *Waste Reduction Algorithm: Chemical Process Simulation for Waste Reduction*, https://www.epa.gov/chemical-research/waste-reduction-algorithm-chemical-process-simulation-waste-reduction (last updated Aug. 3, 2016).

These are just a few examples of the value of regulating AI and the way it can be used for environmental protection when good governance is employed. EPA should begin exploring its ability to govern algorithms outside of just the Agency's own use of AI, both to mitigate environmental harms and incentivize the use of AI to benefit the environment.

What would an initial governance framework for AI that considers environmental impacts look like? Given the sources of bias, examples of how AI has gone wrong in the past, and the ways in which AI is being utilized in various industries, here are a few key principles that ought to be addressed:

1. AI must be transparent.[68]

It is incredibly hard to fix and understand a problem with AI when the cause is hidden within an opaque black box. Therefore, algorithmic decisionmaking should be explainable and accessible to those who question the results or are adversely affected by the outcomes. Additionally, records should be kept on how the AI was trained, programmed, and so forth. For example, it may be worth implementing something similar to the European Union's "right to explanation" legislation promoted by the Alan Turing Institute.[69]

2. AI must be accountable.[70]

Those involved in the creation of AI must be aware of the problems that bias and blind spots can pose for the environment. More work needs to be done to develop a framework for AI accountability to identify those accountable if/when there are adverse consequences associated with AI—whether it is the institutions that developed the AI, those that reviewed it/certified its use, or others.

3. AI development must keep in mind potential environmental impacts.

A process-focused requirement like environmental impact statements under NEPA would ensure that AI developers keep their environmental impacts in mind.

Many of the consequences that AI could have on the environment might simply be a result of not considering the environment when implementing the system. Environmentalists and relevant stakeholders, including representatives of potentially affected communities, should be included in the conversations and planning phase of AI development.

4. There should be incentives for the use of AI in addressing climate change issues and efforts to disincentivize use of AI in adding to these problems.

Providing incentives, in whatever form, can give a huge boost to an industry that is constantly developing. If these incentives are designed with the environment in mind, then AI governance could provide the space for developers to become "disruptors" in the realm of energy and the environment. Currently, it is difficult for low-carbon innovators to excel when systems and policies are designed to support the high-carbon energy sector that has been around much longer.[71] With the American Artificial Intelligence Initiative promising to further support AI development, this concern of inequitable support for environmental AI use becomes even more important.

5. AI should capitalize on its ability to make common resource use more efficient but balance these efforts with sustainability.

There is a lot of room for private industries to collaborate on this work, and given the growing success of green marketing, these industries could have a lot to gain by making their work more sustainable.[72]

By no means is this an exhaustive list of areas to address when governing AI. Truthfully, we may not yet know the areas in which action will be needed to guide AI away from unintended (or intended) consequences. However, at the very least, these five primarily process-oriented elements can ensure that the development of AI is more intentional with respect to the environment and unforeseen consequences.

68. Association for Computing Machinery U.S. Public Policy Council, Statement on Algorithmic Transparency and Accountability (2017), https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf.

69. Alan Turing Institute, *A Right to Explanation*, https://www.turing.ac.uk/research/impact-stories/a-right-to-explanation (last visited Sept. 15, 2020).

70. There is significant controversy about how accountability works for AI; *see, e.g.*, Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 Harv. J.L. & Tech. 353 (2016), *available at* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2609777; Mittelstadt et al., *supra* note 50.

71. Rebecca Willis, National Endowment for Science, Technology, and the Arts, The Disrupters: Lessons From Low-Carbon Innovation From the New Wave of Environmental Pioneers (2007), http://sro.sussex.ac.uk/id/eprint/47867/1/NESTA_the_disrupters.pdf; Environmental and Energy Study Institute, Fact Sheet—Fossil Fuel Subsidies: A Closer Look at Tax Breaks and Societal Costs (2019), https://www.eesi.org/files/FactSheet_Fossil_Fuel_Subsidies_0719.pdf.

72. Jessica Lyons Hardcastle, *Will Multinationals' Sustainable Brand Buying Spree Continue?*, Env't & Energy Leader, Sept. 28, 2016, https://www.environmentalleader.com/2016/09/will-multinationals-sustainable-brand-buying-spree-continue/.