
KANTIAN ETHICS AND ARTIFICIAL INTELLIGENCE

Ms. Shreya Goswami, School of Law, GD Goenka University

Ms. Suhani Kumar, School of Law, GD Goenka University

ABSTRACT

As artificial intelligence (AI) systems increasingly permeate sectors like healthcare, transportation, and law enforcement, the ethical implications of their design demand urgent philosophical scrutiny. This paper explores how Immanuel Kant's ethical framework, especially the Categorical Imperative, can guide the responsible development and deployment of AI. Kantian principles such as universality, autonomy, and respect for persons offer a robust foundation for evaluating the moral dimensions of AI behavior. Through case studies in autonomous vehicles, healthcare, and law enforcement, this paper illustrates practical ethical challenges and proposes a comprehensive framework for embedding Kantian values into AI design. The paper argues for integrating these principles into AI systems through ethical guidelines, transparency mechanisms, and continuous evaluation to ensure the protection of human dignity and autonomy.

Keywords: Kantian Ethics, AI.

Introduction

Artificial Intelligence (AI) represents one of the most transformative technologies of the 21st century, radically reshaping societal functions, industrial processes, and even fundamental human interactions. As AI systems are increasingly integrated into sectors such as healthcare, transportation, law enforcement, and national security, they are being entrusted with tasks that involve complex, value-laden decisions. These decisions have significant implications for individuals' rights, autonomy, and dignity. Consequently, the ethical considerations surrounding AI design and deployment have moved to the forefront of philosophical, legal, and technological discourse. In this context, it becomes essential to ground AI ethics in robust philosophical frameworks that can navigate the moral complexity of these emerging technologies. Among the diverse traditions in moral philosophy, Kantian ethics offers a compelling and principled foundation. Developed by Immanuel Kant in the 18th century, this deontological framework is grounded in the notions of moral law, rational agency, and respect for persons. The central pillar of Kantian ethics is the Categorical Imperative, which mandates that one should act only according to maxims that could be willed as universal laws¹. Equally fundamental are Kant's principles of autonomy—the capacity of rational beings to legislate moral laws for themselves—and human dignity, which demands that individuals never be treated merely as means to an end. Unlike consequentialist theories, such as utilitarianism, which evaluate actions by their outcomes, Kantian ethics insists on the intrinsic morality of actions and the inviolability of human dignity, making it uniquely suited to address the ethical challenges posed by AI². This paper is motivated by a set of core objectives. It seeks to analyze the potential of Kantian ethics to serve as a moral compass for AI development and deployment, focusing on the application of its principles in contexts where machines are entrusted with autonomous decision-making. It explores how AI systems can be aligned with the Categorical Imperative, particularly in situations requiring respect for universalizable moral norms. Additionally, the study examines how Kant's concept of autonomy can be used to interrogate the growing discourse around machine autonomy or “machine will,” raising crucial questions about moral agency and accountability. A significant portion of the analysis also involves practical applications, with a focus on AI in healthcare, autonomous vehicles, and law enforcement, which are environments where ethical tensions are frequently acute. By doing so, this paper aims to develop ethical design guidelines for AI systems that prioritize transparency, informed consent, and the preservation of user dignity.

In pursuing these objectives, the research seeks to answer several interrelated questions. What is the relevance of Kantian principles like the Categorical Imperative, autonomy, and dignity in the development of AI systems? Can AI systems be said to meet the criteria for moral agency as articulated by Kant? If not, what are the implications for responsibility and accountability? How can Kantian ethics be practically embedded in the architecture of AI decision-making systems, particularly in ethically sensitive sectors? Further, what are the ethical implications of using AI in civilian versus military contexts? Finally, how does Kantian ethics compare with other prominent ethical frameworks such as utilitarianism and virtue ethics in addressing these challenges? A growing body of literature has begun to explore these questions, providing a foundational basis for this research. The edited volume "Kant and Artificial Intelligence" by Hyeongjoo Kim and Dieter Schönecker presents an interdisciplinary examination of how Kantian philosophy intersects with contemporary AI debates, highlighting the challenges of applying rational moral principles to algorithmic systems. Similarly, Ozlem Ulgen's work "Kantian Ethics in the Age of Artificial Intelligence and Robotics" explores the practical tensions between Kant's moral imperatives and real-world AI applications, especially in contexts that involve significant human vulnerability. These texts serve as critical interlocutors for this paper, framing its central inquiry while also pointing to gaps in the application of Kantian theory to the technical architecture of AI systems.

The methodology adopted in this research is doctrinal. This approach involves a close analysis of primary philosophical texts—chiefly the works of Immanuel Kant—as well as key legal and ethical frameworks relevant to AI. Doctrinal research facilitates a detailed, normative exploration of the core concepts of Kantian ethics, enabling a robust theoretical grounding for evaluating the ethical implications of AI. It also provides a means of interpreting how existing legal and ethical standards can be expanded or reformed in light of Kantian moral principles. This method allows the study to contribute meaningfully to both philosophical scholarship and the evolving discourse on AI ethics by articulating a principled, coherent, and actionable ethical framework grounded in one of moral philosophy's most enduring traditions.

Through a Kantian lens, this paper seeks to illuminate a path forward for the ethical development of AI—one that does not merely aim for functionality or efficiency but aspires to

¹ Immanuel Kant, *Groundwork for the Metaphysics of Morals*, trans. Mary Gregor (Cambridge: Cambridge University Press, 1998), 421.

² Luciano Floridi and Josh Cowls, "A Unified Framework of Five Principles for AI in Society," *Harvard Data Science Review* 1, no. 1 (2019),

uphold the fundamental moral values that define our shared humanity.

Kant's Categorical Imperative

Core Principles

Kant's Categorical Imperative serves as the cornerstone of his moral philosophy, positing that moral obligations are universal and applicable to all rational beings. The Categorical Imperative is often summarized in two main formulations: universality and respect for persons. The universality principle dictates that one should act only according to that maxim by which one can at the same time will that it should become a universal law³. In other words, before taking action, individuals must consider whether the underlying principle of their action could be universally applicable without contradiction. This formulation emphasizes consistency in moral reasoning and the necessity of universalizability in ethical decision-making. The respect for persons principle asserts that individuals must treat humanity, whether in oneself or others, always as an end and never merely as a means to an end⁴. This principle emphasizes the intrinsic worth of every individual and the necessity of recognizing each person's autonomy and dignity in moral deliberation.

These core principles of Kantian ethics challenge us to consider the moral implications of our actions and the responsibilities we hold toward others. In the context of AI, these principles become critical in evaluating how these systems interact with humans and the ethical considerations necessary for their design and implementation.

Autonomy and Moral Agency

Kantian ethics places a significant emphasis on the concept of autonomy, defining it as the capacity for rational self-governance and moral agency. According to Kant, rational beings are capable of recognizing moral laws and acting in accordance with them. This notion of autonomy is essential for understanding the ethical implications of AI systems, particularly those designed to make decisions independently.

³ Immanuel Kant, *Critique of Practical Reason* (Mary Gregor trans., Cambridge Univ. Press 1997).

⁴ Christine Korsgaard, Kant's Formula of Humanity, in *Kant's Groundwork for the Metaphysics of Morals: A Critical Guide* 119–43 (Paul Guyer ed., Cambridge Univ. Press 2010).

For Kant, autonomy is foundational to moral agency. He argues that moral agents must be free to act according to rational principles⁵, which requires the ability to deliberate and choose based on moral considerations. Autonomy implies not only the capacity for rational thought but also the responsibility that accompanies moral decision-making. As AI systems become more autonomous, the question arises as to whether they can fulfill Kantian criteria for moral agency. Can AI systems, which operate based on algorithms and data processing, truly be considered autonomous in the same way humans are? This raises critical questions about the nature of decision-making in AI, particularly regarding the extent to which these systems can engage in moral reasoning.

Dignity and Respect for Persons

Kant's concept of dignity is intricately linked to his views on autonomy and moral agency. He posits that all rational beings possess inherent dignity due to their capacity for rational thought and moral action⁶. This concept has profound implications for the design of AI systems that interact with humans. Kantian dignity is not merely a social construct; it is rooted in the rational capacity of individuals. This means that any actions or technologies that undermine human dignity—by treating individuals merely as means to an end—are inherently unethical. Respecting human dignity requires that AI systems recognize and uphold the intrinsic worth of every individual.

In designing AI systems, developers must prioritize features that respect and enhance human dignity. This includes ensuring that AI technologies do not manipulate, deceive, or exploit users. Instead, AI should empower individuals, providing them with tools that enhance their autonomy and decision-making capabilities.

Applications of Kantian Ethics in Artificial Intelligence

The integration of artificial intelligence into morally complex human environments compels us to reconsider how ethical principles are embedded in machine behavior. As these systems increasingly make decisions in domains involving life, death, privacy, and autonomy, the demand for a principled ethical framework becomes urgent. Kantian ethics—with its emphasis on universality, rationality, autonomy, and respect for persons—offers a rigorous and morally

⁵ Allen W. Wood, *Kant's Ethical Thought* (Cambridge Univ. Press 1999).

⁶ Kant defines dignity as the "absolute inner worth" of rational beings, which is beyond all price.

consistent foundation to assess and guide the development and deployment of AI⁷. This section explores how Kant's philosophy applies to AI systems, particularly focusing on questions of autonomy, moral agency, and ethical decision-making, illustrated through real-world contexts such as transportation, healthcare, and law enforcement.

The development of autonomous AI systems poses especially significant challenges when viewed through the Kantian lens. Kantian autonomy is not merely the ability to act independently, but the capacity to will moral laws for oneself—acting not according to instinct or preference but according to reasoned moral maxims that could be universalized. When designing AI capable of autonomous decision-making, developers must grapple with whether such machines can meaningfully approximate this kind of autonomy. A Kantian analysis demands that an autonomous system not only act without human intervention but that it be capable of evaluating its actions according to moral principles that respect human dignity and could be affirmed by all rational agents. This introduces the philosophical and technical question of whether AI can be programmed to act in accordance with universal moral law. Ethical algorithms would have to encode reasoning structures that simulate moral deliberation—not just rule-following, but the evaluation of maxims against standards of universality and moral worth⁸.

This leads directly to the question of moral agency, a status Kant reserves for beings capable of rational reflection and moral responsibility. Moral agents, according to Kant, must be capable of recognizing duty and acting for the sake of duty. They must also be held accountable for their choices. At present, AI systems lack key components of moral agency: they do not possess consciousness, cannot comprehend moral reasons, and cannot be held accountable in a meaningful sense⁹. Although they can simulate human-like decision-making, their operations are still determined by coded instructions and probabilistic modeling rather than rational volition. Therefore, from a Kantian perspective, current AI cannot be considered moral agents. The burden of moral responsibility thus remains with the human designers, programmers, and

⁷ Thomas Metzinger, "Ethics Washing Made in Europe," *Der Tagesspiegel*, May 2019. Metzinger critiques the superficial use of ethical language in AI discussions and emphasizes the need for substantive frameworks like Kantian ethics.

⁸ Thomas Arnold et al., "Ethical Considerations in Artificial Intelligence Courses," *AI & Society* 35, no. 1 (2020): 1–11.

⁹ John Danaher, "The Ethics of Artificial Intelligence and Robotics," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Fall 2021 Edition)

institutions that create and deploy these systems.

A related philosophical concern involves the moral status of AI—whether these entities should be treated merely as tools or granted some form of ethical recognition. Kant draws a sharp distinction between persons, who are ends in themselves, and things, which are means to ends. Persons have intrinsic worth by virtue of their capacity for rational agency. As AI does not yet possess this capacity, it cannot be treated as a moral patient or holder of rights under Kant's framework. However, the way humans design and interact with AI systems reflects back on our own moral character. If AI is used to manipulate, deceive, or exert control in ways that undermine others' autonomy, we fail not the AI, but the people affected by it. The true Kantian test lies in whether AI systems are used in a way that respects the autonomy and dignity of all human individuals.

These philosophical principles take on sharp practical significance when applied to real-world case studies. Autonomous vehicles offer one of the clearest examples. These systems must often make split-second decisions that could involve life and death trade-offs. A Kantian approach insists that vehicles must not be programmed to make purely outcome-driven decisions, such as sacrificing one person to save many, if doing so violates the principle of treating individuals as ends in themselves¹⁰. Instead, their ethical programming must reflect a commitment to each person's inviolable dignity, avoiding any calculation that treats human lives as interchangeable units in a moral cost-benefit analysis.

Similarly, the application of AI in healthcare must be approached with deep sensitivity to patient autonomy and dignity. Healthcare AI systems are increasingly involved in diagnostics, treatment recommendations, and even patient interaction. From a Kantian standpoint, these systems must be transparent and supportive of informed consent, ensuring that patients retain agency over their medical decisions. Moreover, systems must be attentive not just to physical outcomes, but to the emotional and ethical needs of individuals, recognizing them as whole persons rather than data points or algorithmic subjects.

Finally, the use of AI in law enforcement and surveillance raises critical concerns regarding privacy, discrimination, and accountability. Kantian ethics demands that individuals not be

¹⁰ Sven Nyholm and Jilles Smids, "The Ethics of Accident-Algorithms for Self-Driving Cars: An Applied Trolley Problem?," *Ethical Theory and Moral Practice* 19, no. 5 (2016): 1275–1289.

treated merely as means to state ends, such as efficiency or public order. Surveillance technologies that monitor, predict, or control human behavior without transparency or consent undermine personal autonomy and reduce individuals to objects of control. AI systems deployed in this domain must be designed and governed in a way that upholds the moral worth of every citizen, with mechanisms in place to ensure fairness, due process, and recourse for wrongful actions.

In all these areas, the application of Kantian ethics serves as a moral compass, compelling designers, regulators, and societies to prioritize principles of respect, fairness, and universal moral law. Rather than allowing expediency or utility to dictate how AI functions, Kant's philosophy demands that we uphold the same moral standards in our machines that we expect in our human interactions. The path toward ethical AI thus involves not only technological sophistication but a deeper commitment to moral clarity and human dignity.

Case Studies

Autonomous Vehicles

The ethical implications of autonomous vehicles provide a compelling case study for applying Kantian ethics in AI¹¹. Autonomous vehicles must navigate complex moral dilemmas, particularly in collision scenarios where decisions could result in harm to human life. Applying Kantian ethics requires these vehicles to evaluate their actions against the principles of universality and respect for persons. In the event of an unavoidable accident, an autonomous vehicle must consider how its actions will affect the dignity and lives of all individuals involved. This may involve programming the vehicle to prioritize the safety of passengers while also respecting the lives of pedestrians or other road users. The ethical guidelines governing these decisions must reflect a commitment to human dignity and the principles outlined in Kantian ethics.

Healthcare AI

The use of AI in healthcare presents another critical area for examining Kantian ethics¹². AI systems utilized in patient diagnostics and treatment must respect the autonomy of patients.

¹¹ Patrick Lin, The Ethics of Autonomous Cars, *The New Atlantis*, no. 27, 10–23 (2010).

¹² Jason J. Fodeman et al., Ethics of AI in Health Care: A Systematic Review, *Artificial Intelligence in Medicine* 101: 101764 (2020).

This includes providing clear and transparent information about medical procedures and ensuring that patients are informed participants in their care decisions. Healthcare AI must also recognize and uphold the dignity of patients. This involves designing systems that prioritize patient welfare, privacy, and the emotional aspects of healthcare interactions. For example, AI systems should not only consider medical data but also the individual experiences and preferences of patients in their decision-making processes.

AI in Law Enforcement and Surveillance

The ethical concerns surrounding AI in law enforcement and surveillance highlight the necessity of applying Kantian ethics¹³ to prevent potential abuses of power. AI systems used for surveillance must respect the privacy and autonomy of individuals. This involves ensuring that data collection and monitoring practices do not infringe upon the rights of citizens or treat them merely as means to an end. The deployment of AI in policing raises questions about accountability and the ethical implications of using AI to make decisions that affect human lives. Kantian principles necessitate that law enforcement AI systems adhere to moral standards that prioritize human dignity, prevent discrimination, and maintain accountability for actions taken.

Challenges in Implementing Kantian Ethics in AI

Programming Kantian Principles

The technical and philosophical challenges in programming AI to adhere to Kant's ethical principles are substantial¹⁴. Incorporating Kantian principles into AI decision-making algorithms involves navigating the complexities of moral reasoning. AI systems must be able to assess various scenarios, weigh competing ethical considerations, and arrive at decisions that reflect Kantian values. The diverse and often ambiguous nature of human experiences complicates the task of programming AI to respect Kantian ethics. AI systems must be equipped to interpret context and adapt their decision-making processes accordingly, which poses significant technical challenges.

¹³ Elizabeth Adams, The Ethical Implications of AI in Law Enforcement, *Journal of AI & Ethics* 1(2): 125–39 (2021).

¹⁴ Ryan Calo, Artificial Intelligence Policy: A Primer and Roadmap, *Stanford Technology Law Review* 22: 1–48 (2019).

AI's Lack of Consciousness and Intent

A fundamental obstacle to aligning AI with Kantian ethics is its inherent lack of consciousness and intent. Kantian ethics requires that moral agents possess the capacity for rational thought and the ability to make decisions based on ethical principles. AI systems, which operate based on algorithms and data inputs, do not possess this level of consciousness or understanding, raising questions about their ability to fulfill Kantian standards. Given that AI lacks consciousness, the question of responsibility for actions taken by these systems becomes contentious. Who is accountable when an AI system acts in a manner that contradicts Kantian principles? Addressing these issues is crucial for developing ethical frameworks that govern AI behavior.

Ethical Dilemmas and Inherent Limitations

The application of Kantian ethics in AI is not without challenges, particularly in scenarios where ethical frameworks may conflict. Situations may arise where Kantian ethics conflicts with other ethical considerations, such as utilitarianism or consequentialism. For example, a utilitarian approach may justify actions that maximize overall good, even if they compromise individual rights, while Kantian ethics prioritizes respect for persons and intrinsic moral worth. The complexity of real-world situations can make it difficult to apply Kantian principles consistently. For example, in emergency scenarios, the need for rapid decision-making may necessitate trade-offs that challenge the rigid application of ethical norms. Finding a balance between adherence to Kantian principles and the practical realities of AI deployment is essential.

Proposed Framework for Kantian Ethical AI

Integrating the Categorical Imperative

To ensure that AI systems operate in accordance with Kantian ethics, practical methods for embedding the Categorical

Imperative into AI decision-making processes must be established. One approach is the development of ethical decision trees that allow for systematic evaluation of potential actions against Kantian principles. These trees would prompt AI systems to assess whether the proposed action aligns with universal moral laws and respects the dignity of affected

individuals. Additionally, regular ethical audits of AI systems can help identify potential areas of non-compliance with Kantian principles. These audits should involve multi-disciplinary teams that include ethicists, technologists, and representatives from affected communities to ensure diverse perspectives are considered.

Guidelines for Respecting Human Autonomy and Dignity

Establishing clear guidelines for AI systems is essential to ensure that they respect human autonomy and dignity in their operations¹⁵. AI systems should be equipped with protocols that prioritize informed consent, ensuring that users fully understand the implications of their interactions with the technology. This includes providing accessible information regarding data usage, decision-making processes, and potential risks associated with AI deployment. Furthermore, AI technologies should be designed with user empowerment in mind, providing individuals with the tools necessary to exercise control over their data and interactions with AI systems. This may involve features that allow users to customize settings, opt out of certain functions, or easily access their data.

Evaluation Metrics

To assess AI's ethical compliance from a Kantian perspective, specific metrics must be established. AI systems should be evaluated based on their adherence to the principle of universalizability. This involves assessing whether the underlying maxims guiding AI decisions could be universally applied without contradiction, ensuring that ethical consistency is maintained. Metrics that gauge how well AI systems respect user autonomy in their interactions can provide valuable insights into the ethical implications of these technologies. This may include user surveys, feedback mechanisms, and assessments of how effectively AI promotes informed decision-making. Developing criteria to evaluate how well AI systems uphold human dignity is essential, which could involve examining the impact of AI interactions on users' well-being, satisfaction, and trust in the technology¹⁶.

¹⁵ *Autonomy*, in the Kantian context, refers to the capacity of rational agents to legislate moral law for themselves. See: Allen W. Wood, *Kantian Ethics* (Cambridge: Cambridge University Press, 2008), 66–70.

¹⁶ Shannon Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (Oxford: Oxford University Press, 2016), 173–181. Vallor highlights the importance of dignity and trust in AI-human interactions

Implications for AI Policy and Regulation

The integration of Kantian ethics into AI development has profound implications for policy and regulation. Policymakers must develop national and international ethical standards that reflect Kantian principles, guiding the responsible development and deployment of AI technologies. These standards should emphasize accountability, transparency, and respect for human dignity. Regulatory bodies should be established to oversee the implementation of ethical guidelines and evaluate AI compliance with Kantian principles. This oversight could involve regular audits, assessments, and updates to ensure that AI technologies remain ethically sound. Engaging the public in discussions about AI ethics is crucial for fostering a collective understanding of the implications of AI on society. Educational initiatives can help demystify AI technologies and encourage informed participation in ethical debates.

Comparison with Other Ethical Frameworks

The comparison of Kantian ethics with other ethical frameworks—such as utilitarianism and virtue ethics—provides valuable insights into the strengths and limitations of each approach in addressing the complexities of AI ethics. Utilitarianism emphasizes outcomes and maximizing overall happiness, which can lead to moral justifications for actions that may violate individual rights. While this approach can offer practical solutions to dilemmas, it may conflict with Kantian principles that prioritize respect for persons and intrinsic moral worth. Virtue ethics focuses on the character and moral integrity of individuals, emphasizing the importance of moral virtues in ethical decision-making. While this framework aligns with the cultivation of ethical AI design, it may lack the prescriptive clarity of Kantian ethics regarding universal moral laws. Ultimately, a nuanced approach that incorporates elements from multiple ethical frameworks may be necessary to navigate the complex landscape of AI ethics effectively. This hybrid approach can enhance the ethical robustness of AI technologies while ensuring that fundamental Kantian principles are upheld.

Future Directions

As AI continues to evolve, future research should explore Kantian approaches to emerging challenges¹⁷, such as the ethical implications of advanced AI systems that may surpass human

¹⁷ Wendell Wallach & Colin Allen, *Moral Machines: Teaching Robots Right From Wrong* (Oxford Univ. Press 2008).

intelligence and decision-making capabilities. Understanding how Kantian ethics applies in scenarios where AI systems operate beyond human control is critical for safeguarding human dignity and autonomy. Additionally, examining the ethical considerations of collaborative systems that integrate human and AI decision-making is essential. Research could focus on how to design AI that complements human agency, ensuring that ethical principles remain central to collaborative processes. Exploring how Kantian ethics interacts with diverse cultural values and ethical frameworks across the globe will contribute to a more comprehensive understanding of moral philosophy in the context of AI. Understanding the intersectionality of ethics can help refine Kantian principles to accommodate a broader spectrum of human experiences.

Conclusion

The intersection of Kantian ethics and artificial intelligence offers a profound framework for understanding the moral implications of emerging technologies. By grounding AI development in the principles of universality, respect for persons, and dignity, we can ensure that AI systems are designed to enhance rather than undermine human agency. This paper has explored the theoretical foundations of Kantian ethics, its application in various AI contexts, and the challenges inherent in integrating these ethical principles into AI systems. By proposing a framework that incorporates Kantian principles into AI design and decision-making processes, we pave the way for a future where technology aligns with our moral values. The successful implementation of Kantian ethics in AI will require interdisciplinary collaboration, robust policy frameworks, and ongoing dialogue among stakeholders. As we navigate the complexities of AI development, a commitment to respecting human dignity and autonomy must remain at the forefront of our ethical considerations, guiding us toward a future where technology serves humanity in a manner consistent with our highest moral aspirations. As AI continues to shape the future, the necessity for ethical frameworks that prioritize Kantian principles will become increasingly urgent, ensuring that technology reflects our deepest values and respects the dignity of all individuals.

REFERENCES

1. Allen W. Wood, *Kant's Ethical Thought* (Cambridge Univ. Press 1999).
2. Jason J. Fodeman et al., Ethics of AI in Health Care: A Systematic Review, *Artificial Intelligence in Medicine* 101: 101764 (2020).
3. John Danaher, "The Ethics of Artificial Intelligence and Robotics," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Fall 2021 Edition)
4. Kant, Immanuel. *Groundwork for the Metaphysics of Morals*. Cambridge University Press, 1997.
5. Kant, Immanuel. *Critique of Practical Reason*. Cambridge University Press, 1997.
6. Kim, H. & Schönecker, D. (2022). *Kant and Artificial Intelligence*. Berlin, Boston: De Gruyter.
7. Shannon Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting* (Oxford: Oxford University Press, 2016).
8. Ryan Calo, Artificial Intelligence Policy: A Primer and Roadmap, *Stanford Technology Law Review* 22: 1–48 (2019).
9. Ulgen, O. (2017). Kantian ethics in the age of artificial intelligence and robotics