

## **ETHICS IN THE AGE OF ARTIFICIAL INTELLIGENCE: RECONCEPTUALISING THE TRADITIONAL ETHICAL THEORIES**

**Anthony Ugochukwu Nwokoye**  
Department of Philosophy  
Nnamdi Azikiwe University, Awka  
au.nwokoye@unizik.edu.ng

### **Abstract**

The rapid evolution of artificial intelligence (AI) has fundamentally disrupted traditional ethical theories, necessitating a re-conceptualization of moral frameworks in the age of AI. As AI systems grow increasingly sophisticated, they challenge long-standing notions of moral agency, responsibility, and ethical decision-making. This research examines how the three waves of AI—Predictive AI, Generative AI, and Agentic AI—reshape ethical paradigms. Predictive AI, with its data-driven algorithms, exposes inherent biases and raises critical questions about justice, fairness, and accountability in automated decision-making. Generative AI, capable of creating synthetic content, disrupts traditional concepts of authorship, authenticity, and intellectual property, forcing a re-evaluation of ethical norms in creativity and ownership. Agentic AI, with its capacity for autonomous action, pushes the boundaries of moral responsibility, challenging humans to reconsider the ethical implications of delegating decision-making to machines. These developments demand a rethinking of traditional ethical theories such as utilitarianism, deontology, and virtue ethics, which were designed for human moral agents but fall short in addressing the unique moral dilemmas posed by AI. The research highlights the limitations of these classical theories in dealing with AI's opacity, autonomy, and responsibility. The study examines key ethical challenges, including the issue of moral agency and algorithmic bias, and proposes the need for a new ethical framework that accounts for the collaborative nature of human-AI interactions, emphasizing distributed moral responsibility and the importance of human oversight in ensuring ethical outcomes.

**Keywords:** Artificial Intelligence, Moral Agency, Predictive AI, Generative AI, Agentic AI, Ethical Theories, Moral Responsibility

### **Introduction**

The emergence of artificial intelligence (AI) has ushered in an era of unprecedented technological innovation, fundamentally altering the way humans interact with the world and with

each other. Yet, as AI systems grow increasingly sophisticated, they also pose ethical challenges that demand a re-examination of the moral theories that have guided human thought for centuries. Traditional ethical theories—rooted in the works of philosophers such as Immanuel Kant, John Stuart Mill, and Aristotle—have long provided the foundation for understanding moral agency, responsibility, and decision-making. However, the unique capabilities of AI, particularly in its predictive, generative, and agentic forms, expose the limitations of these classical theories, revealing their inadequacy in addressing the novel ethical dilemmas of the digital age.

At the heart of traditional ethics lies the concept of moral agency, the capacity of an entity to make decisions and bear responsibility for its actions. Deontological ethics, championed by Kant, emphasizes duty, rules, and the intrinsic morality of actions. For Kant, moral worth is derived from adherence to universal principles, such as the categorical imperative, which demands that one act only according to maxims that could be universally applied (Kant, 2008). However, this theory struggles to accommodate AI systems, which operate not from a sense of duty or moral reasoning but from programmed algorithms and data-driven processes (Wallach & Allen, 2008). Can an AI system, devoid of consciousness and intentionality, be held to the same moral standards as a human agent? If not, who bears the responsibility for its actions—the designer, the user, or the machine itself? These questions remain unresolved within the deontological theory.

In a similar vein, utilitarianism, articulated by John Stuart Mill, offers a different perspective, focusing on the consequences of actions and the maximization of overall happiness. According to Mill, the ethical value of an action is determined by its outcomes, with the goal of achieving the greatest good for the greatest number (Mill, 1863). While this approach may seem adaptable to AI, particularly in optimizing decision-making processes, it falters when confronted with the inherent biases and unintended consequences of algorithmic systems (Zimmer, 2020). Predictive AI, for instance, often relies on historical data that may perpetuate systemic inequalities, raising questions about fairness and justice (O’Neil, 2016). Can a utilitarian framework adequately address the ethical implications of AI-driven decisions that benefit the majority while marginalizing vulnerable populations?

On the other hand, Virtue ethics, rooted in Aristotle’s emphasis on character and moral excellence, shifts the focus from rules or outcomes to the cultivation of virtuous traits. This approach seeks to answer the question, "What kind of person should I be?" by promoting qualities

such as wisdom, courage, and compassion (Aristotle, 350 BCE). Yet, virtue ethics faces significant challenges when applied to AI. Can an AI system, which lacks consciousness and the capacity for moral growth, embody virtues? Moreover, how can virtue ethics guide the behavior of human designers and users of AI, particularly when the technology operates in ways that are opaque or beyond human comprehension?

AI has evolved from rule-based systems to sophisticated models capable of prediction, generation, and autonomous action (Russell & Norvig, 2016). Predictive AI, which relies on data-driven algorithms to forecast outcomes, has become permeating in fields such as healthcare, finance, and criminal justice. Generative AI, exemplified by systems like OpenAI's GPT and DALL-E, can produce text, images, and even music that are indistinguishable from human-created content (Totlani, 2023). Agentic AI, which exhibits goal-directed behavior, represents the next frontier, with systems capable of making decisions and taking actions with minimal human intervention (Ogbu, 2024). In light of these developments, it becomes clear that traditional ethical theories, while foundational, are ill-equipped to address the unique challenges posed by AI. The principles of deontology, utilitarianism, and virtue ethics, though invaluable in guiding human behavior, fail to account for the complexities of intelligent systems that operate beyond human intuition and control (Boddington, 2017).

### **The Limitations of Traditional Ethical Theories in the Age of Artificial Intelligence**

Traditional philosophy, while robust in its historical and theoretical foundations, was developed in a pre-digital era and thus struggles to account for the difficulties introduced by AI. Classical ethical theories, such as deontology and utilitarianism, were designed to address human moral agency and decision-making. However, they fall short in addressing the ethical implications of AI systems that operate autonomously. As Mittelstadt, Allo, Taddeo, Wachter, and Floridi (2016) argue, the opacity and complexity of AI algorithms make it difficult to apply traditional ethical principles, such as accountability and transparency, in a meaningful way.

Immanuel Kant's deontological ethics emphasizes the importance of duty, intention, and universal maxims (Kant, 1785). Deontology, rooted in the philosophy of Immanuel Kant, focuses on duty, rules, and the intrinsic morality of actions. Kant argued that moral decisions should be

guided by universal principles, like the "categorical imperative," which requires us to act in ways that could be universally applied (Kant, 1785). But how do we apply this to AI? Machines don't have intentions or a sense of duty. For example, consider an autonomous car faced with a no-win scenario: should it swerve to hit one pedestrian to save five others? Deontology can't easily answer this, because the car isn't making a "moral choice"—it's following programmed instructions. The lack of intentionality in AI systems makes it hard to hold them accountable under Kantian principles. Who is accountable when an autonomous vehicle causes harm? The opacity of AI algorithms, often referred to as the "black box" problem, further complicates the application of traditional ethical frameworks (Mittelstadt et al., 2016). Unlike human agents, AI systems cannot be held morally responsible for their actions, making concepts like accountability and transparency to raise further philosophical questions.

On the other hand, utilitarianism, as articulated by Jeremy Bentham and John Stuart Mill, focuses on maximizing overall happiness or utility (Mill, 1863). While this approach has been useful in policy-making and cost-benefit analyses, it struggles with AI's complexity. Utilitarianism's focus on the "greater good" can overlook the harm done to specific groups, raising questions about fairness and justice. Can we really call an AI system ethical if it optimizes for overall utility but perpetuates inequality?

Virtue ethics, inspired by Aristotle, emphasizes moral character and virtues like honesty, courage, and compassion (Aristotle, *Nicomachean Ethics*). This approach is deeply human-centered, focusing on how individuals can cultivate good character over time. But AI systems don't have character or moral agency. For example, how do we judge the "virtues" of an AI that generates deepfake videos or manipulates public opinion? Virtue ethics offers little guidance here, because it's not designed to evaluate non-human entities. The question of whether AI can—or should—embody virtues remains unanswered.

### **The Need for a Philosophical Reimagining**

Given the limitations of traditional philosophical frameworks, there is an urgent need to reconceptualise the branches of philosophy to address the unique challenges posed by AI. This reimagining must take into account the relational and interconnected nature of human and machine

agency, the ontological significance of AI-generated entities, the hybrid nature of human and machine reasoning, and the collaborative nature of knowledge production in the age of AI.

### **Ethical Issues in the Age of Artificial Intelligence**

The rapid rise of artificial intelligence has forced humanity to confront new ethical questions that traditional moral frameworks struggle to answer or did not consider. From healthcare and criminal justice to entertainment and warfare, AI is reshaping how humans live, work, and interact. But as these systems grow more powerful and autonomous, they also reveal gaps in the traditional ethical theories. Traditional theories like deontology, utilitarianism, and virtue ethics were designed for a world where humans were the primary moral agents. In the age of AI, these theories often fall short, leaving humans unprepared to address the unique challenges posed by machines that can think, learn, and act in ways that mimic or even surpass human capabilities. Two of the numerous ethical challenges are examined: the issue of moral agency and algorithm bias.

#### **Moral Agency:**

The concept of moral agency is central to ethical theory, referring to the capacity of an entity to make decisions and bear responsibility for its actions (de Pedro, 2024). Traditional ethical frameworks, such as deontology and utilitarianism, are built on the assumption that moral agents possess intentionality, consciousness, and the ability to reason about right and wrong. However, AI systems lack these qualities, raising questions about their status as moral agents. AI systems operate based on algorithms and data, without any understanding of the ethical implications of their actions. According to Moor (2006), there are two types of ethical agents; implicit ethical agents and explicit ethical agents. Moor made a distinction between implicit and explicit ethical agents as a foundational concept in the field of machine ethics. According to him, Implicit Ethical Agents are machines programmed to follow ethical rules or avoid unethical behavior without explicitly understanding or reasoning about ethics (Moor, 2006). Their ethical behavior is constrained by the principles embedded in their design. For example, a robot vacuum cleaner programmed to avoid damaging furniture is an implicit ethical agent because it follows rules without understanding the ethical implications of its actions. On the other hand, Explicit Ethical

Agents are machines capable of reasoning about ethical dilemmas and calculating the best course of action based on ethical principles (Moor, 2006). They can explicitly represent and apply ethical rules, making them more sophisticated than implicit agents. An instance to illustrate this is an autonomous vehicle that can weigh the consequences of different actions in a moral dilemma (e.g., the trolley problem) would qualify as an explicit ethical agent.

Moor's framework assumes that explicit ethical agents can calculate the "best action" in ethical dilemmas using ethical principles. However, this assumes a level of transparency and interpretability that is often lacking in modern AI systems. Many AI systems, particularly those based on deep learning, operate as "black boxes"—their decision-making processes are opaque, even to their designers. However, when an autonomous vehicles that must make split-second decisions in life-or-death situations is put into consideration, the outcome differs. While these systems may be designed to follow ethical principles e.g., minimizing harm, their decisions are often based on complex algorithms that are not easily interpretable. This raises questions about whether such systems can truly be considered explicit ethical agents, as their reasoning processes are not transparent or understandable to humans. The lack of transparency in AI systems undermines Moor's assumption that explicit ethical agents can reliably calculate the best action in ethical dilemmas. If we cannot understand how an AI system arrives at its decisions, how can we trust it to act ethically?

The lack of intentionality and consciousness in AI systems further complicates the issue of moral agency. Unlike humans, AI systems do not possess beliefs, desires, or the capacity for moral reasoning (Sullivan & Fosso Wamba, 2022). They operate based on statistical patterns and correlations, without any understanding of the ethical principles that guide human behaviour (Mueller, Hoffman, Clancey, Emrey, & Klein, 2019). This raises questions about whether AI systems can ever be held morally responsible for their actions.

The autonomous nature of AI systems creates accountability gaps, particularly in high-stakes applications such as autonomous vehicles and medical AI. When an AI system causes harm, it is often unclear who should be held responsible—the developer who designed the system, the user who deployed it, or the organization that owns it. This ambiguity undermines trust in AI systems and raises ethical concerns about justice and fairness. AI-powered diagnostic tools have been introduced in some hospitals to improve the accuracy and efficiency of medical diagnoses.

While these tools have shown promise, they have also been involved in cases where incorrect diagnoses led to inappropriate treatments. In a scenario where an AI system misdiagnosed a patient with a rare condition, resulting in unnecessary surgery and significant harm to the patient, the question of who is responsible for such errors—the hospital, the developers of the AI system, or the healthcare professionals who relied on it—remains unresolved, illustrating the difficulties of assigning accountability in AI-driven decision-making.

### **Algorithmic Bias**

Algorithmic bias refers to the tendency of AI systems to produce skewed or unfair outcomes, often as a result of flawed data or design (Baker & Hawn, 2022). Bias in AI can manifest in various ways, including racial, gender, socioeconomic, and cultural biases, and it can have far-reaching consequences for individuals and communities (Bose, 2025). The root causes of algorithmic bias are multi-layered, but they primarily stem from two sources: biased training data and flawed algorithms (Kordzadeh & Ghasemaghaei, 2022). Biased training data is one of the most common causes of algorithmic bias. AI systems learn from historical data, which often reflects existing societal prejudices and inequalities. For example, if a hiring algorithm is trained on data from a company that has historically favored male candidates for technical roles, the algorithm may learn to prioritize male applicants, perpetuating gender bias (O'neil, 2017). Similarly, facial recognition systems trained on datasets that are predominantly composed of lighter-skinned individuals may perform poorly on darker-skinned faces, leading to racial bias (Buolamwini & Gebru, 2018). These examples illustrate how historical inequities can be encoded into AI systems, perpetuating and even amplifying existing biases. On the other hand, flawed algorithms also contribute to algorithmic bias. The design and implementation of AI algorithms can introduce biases, particularly when developers fail to account for the diverse contexts in which these systems will be deployed (Klaassen, 2024). For instance, predictive policing algorithms have been criticized for disproportionately targeting minority communities, not because these communities have higher crime rates, but because the algorithms are trained on historical arrest data that reflects biased policing practices (Eubanks 2018). In such cases, the algorithm's design reinforces existing inequalities, creating a feedback loop that exacerbates social injustices.

Many companies now use AI-driven tools to screen job applicants, with the goal of streamlining the hiring process and reducing human bias. However, these systems often replicate and even amplify existing biases. Amazon developed an AI recruiting tool that was trained on resumes submitted to the company over a 10-year period. Because the tech industry is predominantly male, the algorithm learned to favor male candidates, penalizing resumes that included words like "women's" or referenced all-female colleges (Dastin, 2022). This example highlights how biased training data can lead to discriminatory outcomes, even when the intent is to create a fairer hiring process.

The ethical implications of algorithmic bias are deep, particularly in terms of its role in reinforcing systemic inequalities and discrimination. AI systems are often deployed in high-stakes contexts, such as hiring, lending, healthcare, and criminal justice, where biased outcomes can have life-altering consequences for individuals and communities. When AI systems replicate and amplify existing biases, they contribute to the marginalization of already disadvantaged groups, perpetuating cycles of inequality. The reinforcement of systemic inequalities through algorithmic bias raises serious ethical questions about the role of AI in society. As AI systems become increasingly integrated into decision-making processes, there is a risk that they will entrench existing power structures and exacerbate social injustices. This is particularly concerning given the opacity of many AI systems, which makes it difficult to identify and address biases.

Algorithmic bias also poses significant challenges to fairness and justice in decision-making processes. Fairness is a fundamental ethical principle that requires decisions to be made impartially and without prejudice. However, achieving fairness in AI systems is complicated by the fact that different definitions of fairness often conflict with one another. The ethical challenges posed by algorithmic bias are not limited to fairness; they also extend to justice. Justice requires that individuals be treated equitably and that decisions be made in a transparent and accountable manner. However, the opacity of many AI systems makes it difficult to understand how decisions are being made, let alone to hold anyone accountable for biased outcomes. This lack of transparency undermines trust in AI systems and raises questions about their legitimacy as decision-making tools.

## The Way Forward: Building a New Ethical Framework for AI

The rapid advancement of artificial intelligence (AI) has exposed the limitations of traditional ethical theories in addressing the unique challenges posed by intelligent systems. Classical frameworks such as deontology, utilitarianism, and virtue ethics were developed in a pre-digital era and were designed to address human moral agency and decision-making. However, they fall short in addressing the ethical implications of AI systems that operate autonomously, often in ways that are opaque, context-dependent, and beyond human comprehension. As Mittelstadt et al. (2016) argue, the complexity and unpredictability of AI algorithms make it difficult to apply traditional ethical principles, such as accountability and transparency, in a meaningful way. To address these challenges, there is an urgent need to develop a new ethical framework that is specifically tailored to the realities of AI. This framework must integrate insights from philosophy, computer science, law, and social sciences, while also accounting for the unique characteristics of AI systems, such as their autonomy, opacity, and ability to learn and adapt.

Luciano Floridi is a leading philosopher in the ethics of information and technology, and he introduces the concept of distributed moral responsibility (DMR) as a framework for addressing the ethical challenges posed by artificial intelligence (Floridi, 2020). Floridi's concept of distributed moral responsibility is grounded in the idea that moral responsibility in complex systems is not confined to individual agents but is distributed across multiple actors, including designers, users, policymakers, and even the systems themselves. According to Floridi, DMR arises in situations where:

**Multiple Agents Are Involved:** Responsibility is shared among various stakeholders, each contributing to the outcomes of the system.

**Outcomes Are Emergent:** The behavior of the system cannot be attributed to any single agent but emerges from the interactions of its components.

**Intentionality Is Distributed:** The intentions behind the system's actions are not localized but are spread across the network of agents (Floridi, 2020).

Floridi argues that distributed moral responsibility is particularly relevant in the context of AI, where the complexity and autonomy of systems make it difficult to assign responsibility to any single individual or entity. For example, in the case of an autonomous vehicle involved in an

accident, responsibility may be distributed among the car manufacturer, the software developer, the regulatory body, and even the user who chose to rely on the system (Floridi, 2016). However, this can result to a situation where no one takes responsibility for unethical outcomes, as each stakeholder points to others as the primary culprits. For example, in the case of a biased AI system, developers may blame the data, companies may blame the regulators, and regulators may blame the developers, resulting in a lack of accountability (Taddeo & Floridi, 2018). On the other hand, Floridi's framework focuses on the responsibility of human actors but does not fully address the role of AI systems themselves in ethical decision-making. As AI systems become more autonomous, there is a growing need to consider their moral status and the extent to which they can be held responsible for their actions. This also leads to one of the key limitations of traditional ethical theories which focus on individual moral agency, and fail to account for the collaborative nature of ethical decision-making in the age of AI.

Thus, rather than viewing AI systems as independent ethical agents, AI should be considered as tools that require human oversight and intervention to ensure ethical outcomes. It becomes necessary to adopt a collaborative ethical framework to address the issue of human interaction and AI. Collaborative ethics places emphasis on the importance of human-AI collaboration in ethical decision-making. Collaborative ethics requires the development of AI systems that are designed to work in partnership with humans, rather than replacing or overriding human judgment (Bryson 2018).

Cañas (2022) in the article "Collaborative Ethics: A Framework for Addressing Ethical Challenges in AI" presents a compelling argument for the adoption of collaborative ethics as a framework for addressing the ethical challenges posed by artificial intelligence. Cañas (2022) define collaborative ethics as a framework that emphasizes the importance of human-AI collaboration in ethical decision-making. The argument centers on the fact that AI systems should not be viewed as independent ethical agents but as tools that require human oversight and intervention to ensure ethical outcomes. Collaborative ethics recognizes that ethical decision-making in AI involves multiple stakeholders, each with a role to play in ensuring ethical outcomes. This includes developers, who design the systems; users, who interact with them; policymakers, who regulate their use; and even the AI systems themselves, which must be designed to align with human values (Cañas, 2022). Secondly, the framework emphasizes the importance of human

oversight in AI systems, particularly in high-stakes applications such as healthcare, criminal justice, and autonomous vehicles. Humans must remain actively involved in the decision-making process, interpreting and acting on AI recommendations to ensure ethical outcomes (Cañas, 2022). Thirdly, Collaborative ethics requires AI systems to be sensitive to the specific contexts in which they operate, adapting their behavior to reflect the values and preferences of the communities they serve. This includes accounting for cultural, social, and institutional differences, as well as the unique needs of individual users (Cañas, 2022).

The postulations of Cañas (2022) on the concept of collaborative ethics provides a more comprehensive approach than traditional frameworks, recognizing the shared responsibility of all actors involved in the design, deployment, and use of AI systems. As rightly noted by Floridi (2016) this is particularly important in high-stakes applications, where the consequences of unethical AI can be severe. Also with the emphasis on the importance of human oversight, collaborative ethics ensures that humans remain actively involved in the decision-making process. This is critical for maintaining accountability and ensuring that AI systems operate in ways that align with human values. As observed by Topol (2019) that in healthcare, human oversight can help ensure that AI recommendations are interpreted and acted on in ways that prioritize patient well-being .

Notwithstanding the merit of collaborative ethics, there are still ethical challenges that demands for answers especially in the area of the risk associated with over-reliance on human oversight, particularly in contexts where humans may lack the expertise or resources to effectively monitor AI systems. Mittelstadt et al. (2016) opine that even when humans are actively involved in overseeing AI systems, they may struggle to understand how these systems arrive at their decisions, making it difficult to identify and address ethical issues. At the extreme, this can result to ethical overconfidence, where humans would now depend on AI systems to make ethical decisions. Ethical overconfidence is quite problematic because it can lead to a false sense of security, where humans assume that AI systems are inherently ethical or unbiased.

## References

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International journal of artificial intelligence in education*, 1-41.

- Bose, M. (2025). Bias in AI: A Societal Threat: A Look Beyond the Tech. In *Open AI and Computational Intelligence for Society 5.0* (pp. 197-224): IGI Global Scientific Publishing.
- Buolamwini, J., & Gebru, T. (2018). *Gender shades: Intersectional accuracy disparities in commercial gender classification*. Paper presented at the Conference on fairness, accountability and transparency.
- Cañas, J. J. (2022). AI and ethics when human beings collaborate with AI agents. *Frontiers in psychology, 13*, 836650.
- Dastin, J. (2022). Amazon scraps secret AI recruiting tool that showed bias against women. In *Ethics of data and analytics* (pp. 296-299): Auerbach Publications.
- de Pedro, J. (2024). Social Systems as Moral Agents: A Systems Approach to Moral Agency in Business. *Journal of Business Ethics, 195*(4), 695-711.
- Floridi, L. (2016). Faultless responsibility: On the nature and allocation of moral responsibility for distributed moral actions. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374*(2083), 20160112.
- Floridi, L. (2020). Distributed morality in an information society. In *The Ethics of Information Technologies* (pp. 63-79): Routledge.
- Kant, I. (2008). *Groundwork of the Metaphysics of Morals: by Immanuel Kant*: MobileReference.
- Klaassen, G. (2024). Confronting Bias In Artificial Intelligence: Building Transparent, Diverse, and Ethical AI Systems.
- Kordzadeh, N., & Ghasemaghahi, M. (2022). Algorithmic bias: review, synthesis, and future research directions. *European Journal of Information Systems, 31*(3), 388-409.
- Mill, J. S. (1863). Utilitarianism. In: London: Parker, Son, and Bourn.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society, 3*(2), 2053951716679679.
- Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics. *IEEE intelligent systems, 21*(4), 18-21.
- Mueller, S. T., Hoffman, R. R., Clancey, W., Emrey, A., & Klein, G. (2019). Explanation in human-AI systems: A literature meta-review, synopsis of key ideas and publications, and bibliography for explainable AI. *arXiv preprint arXiv:1902.01876*.
- O'neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*: Crown.
- Ogbu, D. (2024). Agentic AI in Computer Vision Domain-Recent Advances and Prospects.
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*: pearson.
- Sullivan, Y. W., & Fosso Wamba, S. (2022). Moral judgments in the age of artificial intelligence. *Journal of Business Ethics, 178*(4), 917-943.
- Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science, 361*(6404), 751-752.
- Topol, E. (2019). *Deep medicine: how artificial intelligence can make healthcare human again*: Hachette UK.
- Totlani, K. (2023). The evolution of generative AI: Implications for the media and film industry. *International Journal for Research in Applied Science and Engineering Technology, 11*(10), 973-980.
- Wallach, W., & Allen, C. (2008). *Moral machines: Teaching robots right from wrong*: Oxford University Press.
- Zimmer, M. (2020). "But the data is already public": on the ethics of research in Facebook. In *The ethics of information technologies* (pp. 229-241): Routledge.