

INTEGRITY IN BIG DATA AND UTILITARIANISM: THE RESPONSE OF BERNARD WILLIAMS

AGUNBIADE Saheed Anuoluwapo*

Abstract

This paper examines the problem of integrity in the decision of the big data-driven algorithms. Big data refers to datasets that are too large, complex, or fast-moving for traditional processing methods to handle effectively. Data-driven algorithms often rely on historical data to make predictions, but when this data is biased, the outcomes are similarly skewed. The reliance on a utilitarian approach in big data practices; prioritizing efficiency and statistical outcomes for the majority, risks compromising individual integrity by neglecting personal autonomy, dignity, and the complexity of human experience. It is on the basis of this, this paper examines the ethical implications of utilitarianism in big data and algorithmic decision-making, using Bernard Williams' concept of integrity as a theoretical framework. The paper argues that the utilitarian justification for maximizing utility through big data often leads to the erosion of personal integrity, especially when individuals are treated merely as means to an end. Through evaluative and conceptual analysis methods, this paper advocates for an ethical reevaluation of big data practices, proposing that Williams' insights into integrity should guide the development of more ethically responsible technologies that preserve personal and moral authenticity.

Keywords: Algorithmic Bias, Bernard Williams, Big Data, Integrity, Utilitarianism,

Introduction

The advent of big data and its integration into machine learning technologies has revolutionized decision-making processes across multiple sectors, from healthcare to criminal justice, finance, and beyond. Big data refers to large, complex datasets that require advanced computational tools, like machine learning algorithms, to analyze and extract meaningful insights (Mayer-Schönberger & Cukier, 2013). These technologies are often grounded in utilitarian principles, which focus on maximizing overall well-being for the greatest number of people (Bentham, 1789). Utilitarianism, when applied to big data, often leads to decisions aimed at optimizing efficiency, reducing costs, or increasing overall happiness. While these outcomes may seem appealing, they present significant ethical dilemmas—particularly in terms of individual rights and the integrity of those affected by the decisions made by algorithms. The ethical challenge of integrity arises when big data, driven by utilitarianism, overlooks the moral rights and personal autonomy of individuals. Integrity, in this context, refers to the coherence between one's personal values and actions, as well as the respect for individuals' moral autonomy (Williams, 1973). In utilitarianism, decisions are made by weighing the aggregate benefits and costs, often justifying actions that harm minorities or violate personal dignity in the pursuit of collective good. This tension becomes even more pronounced in the era of big data, where individuals are often reduced to mere data points, with their rights subordinated to statistical outcomes that prioritize societal benefits (O'Neil, 2016).

One of the philosophers who offer a critique of such utilitarian approaches is Bernard Williams.

His theory of integrity stresses the importance of aligning individual actions with personal values, rather than subjugating them to abstract moral calculations. Williams (1973) argues that true moral action requires respect for the individual's moral coherence, which cannot be compromised for the sake of maximizing overall well-being. This perspective stands in contrast to utilitarianism, which often demands that individuals sacrifice their personal ethical principles in the interest of broader societal benefits. (Zeng et al., 2017). By virtue of this, the paper examines the ethical implications of utilitarianism in big data and algorithmic decision-making, exploring how Bernard Williams' critique of utilitarianism offers an essential framework for addressing the ethical issues embedded in data-driven decision-making. By integrating Williams' concept of integrity, it will be shown that ethical frameworks for machine learning systems must balance societal good with respect for individual rights and moral coherence. Nonetheless, before we delve into the main argument, the paper shall start with the idea of utilitarianism.

Utilitarianism and its Greatest Good

Utilitarianism is a well-known ethical theory centered on consequences, with its key ideas being the principle of utility and the Greatest Happiness Principle (Omeregbe, 1993: 233). John Stuart Mill, one of the major advocates of this theory, defines the principle of utility as the belief that "utility" serves as the foundation of morality, while the Greatest Happiness Principle suggests that "actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the opposite of happiness" (Mill, 1984: 246). In utilitarian thought, the term "utility" refers to happiness or pleasure, meaning we should always act in ways that result in the greatest balance of positive outcomes for all individuals affected by our actions. For utilitarians, anything that contributes to happiness is considered good. Rational behavior, from this perspective, aims to maximize net utility by calculating the benefits and subtracting the costs. In essence, utilitarians assert that the means must lead to the desired ends, and what makes a means "good" is its capacity to achieve the intended goal. Thus, the maxim of utilitarianism becomes: "happiness is the only thing intrinsically valuable, and everything else is valuable only as a means to achieving it" (Fasora, 2005:23). This is also known as "The Greatest Happiness Principle."

In utilitarianism, the morality of an action is solely determined by its consequences, particularly how it affects the happiness of the majority. To assess the outcomes of actions, utilitarians rely on theories of intrinsic value, which claim that the value of an object or action comes from certain inherent qualities (Skorupski, 2003:209). To determine the best outcome, utilitarians believe it is possible to compare the intrinsic values of different actions and predict which one will lead to the better overall result. Jeremy Bentham, a leading figure in utilitarian thought, introduced the "hedonic calculus," a method for objectively evaluating the morality of actions by comparing their potential pleasure and pain based on seven criteria: intensity, duration, certainty, propensity, fecundity, purity, and extent (Omeregbe, 1993: 234-5). Two main forms of utilitarianism exist: act utilitarianism and rule utilitarianism. Act utilitarianism, associated with Jeremy Bentham, suggests that we must evaluate the consequences of a particular action in a given situation for everyone involved. If the action leads to more overall 'good' than any alternative, it is the right choice (Shaw, 2002:49). Bentham asserts that actions are morally right if they contribute to the greatest human welfare and wrong if they do not, maintaining that moral decisions should directly appeal to the principle of utility and seek the course of action that produces the greatest net good

(Frankena, 1973: 35).

In contrast, rule utilitarianism, which is linked to John Stuart Mill, emphasizes following rules that, if universally adopted, would lead to the greatest benefit for the most people. This approach argues that no rule should be defended if it generates better outcomes for the majority while causing harm to the minority. Rule utilitarianism focuses on the consequences of rules and holds that these rules, rather than individual actions, determine what is morally right or wrong. Mill suggests that when comparing two pleasures, the one that is preferred by those who have experienced both is the more desirable, indicating that those with the greatest experience are the best judges of which pleasures are superior (Mill, 1984: 225). These competent judges tend to favor the higher pleasures over the lower ones (Ethel et al., 1984: 225). Having considered the idea of utilitarianism, the next thing is to consider the concept of big data and its ethical nature. This will be discussed in the next section.

Big Data in Machine Learning

Big Data refers to datasets so large, complex, or fast-moving that traditional data-processing software is inadequate to handle them. The nature of Big Data is multi-dimensional, not only defined by its size but also its variety, velocity, and veracity. Big Data, therefore, involves more than just big numbers. It is characterized by the intricate relationships between structured and unstructured data, dynamic systems, and the need for powerful computational tools and algorithms for processing. The idea Big Data can be defined in various ways depending on the context in which it is used. In general, Big Data refers to data that exceeds the capabilities of traditional data-processing systems in terms of volume, complexity, and speed. The following are some commonly accepted definitions:

1. **Big Data as Volume:** Traditionally, Big Data was defined in terms of the sheer volume of data being produced. It was initially considered a challenge to store such large datasets, but with advances in storage technology, it became clear that the real challenge lies in processing and analyzing this data (Chen, Chiang, & Storey, 2012).
2. **Big Data as Complexity:** Big Data also involves complex data types, including unstructured data such as text, images, video, and social media posts, which cannot easily be stored in traditional databases. The complexity of Big Data lies in processing and deriving meaning from these unstructured forms (Mayer-Schönberger & Cukier, 2013).
3. **Big Data as Speed:** The velocity of data generation is another crucial aspect of Big Data. With the rise of real-time data collection through IoT devices, social media, and other sensors, Big Data systems must be able to handle and process information at high speeds (Manyika et al., 2011).
4. **Big Data as Value:** A more recent shift in the conceptualization of Big Data emphasizes not just the quantity but also the value that can be derived from it. This perspective stresses that Big Data is valuable only if it can be analyzed in ways that lead to actionable insights (Chen, Chiang, & Storey, 2012).

Therefore, Big Data cannot simply be defined by its size; it incorporates different dimensions that make it valuable for decision-making, analysis, and prediction. This includes: volume, variety, Velocity, and value. As identified by Laney (2001), the 3Vs; volume, variety, and velocity, are the

traditional defining features of Big Data. These features often extend into the domain of business intelligence, enabling organizations to derive insights that drive strategic decisions (Chen, Chiang, & Storey, 2012). Nevertheless, the origin of Big Data can be traced back to the rapid development of digital technologies over the last few decades. The growth of the internet, social media, mobile devices, and sensor-based systems has led to an explosion in data production. The concept of Big Data first gained traction in the early 2000s when companies began recognizing the challenges and opportunities associated with managing vast quantities of digital information.

One of the earliest recorded mentions of Big Data was by Roger Mougalias in a 2001 paper where he outlined the challenges of processing large datasets (Mougalias, 2001). However, the term was popularized in 2005 when Doug Laney, an analyst at META Group (now Gartner), described the 3Vs model. Since then, the data landscape has continued to evolve, with modern-day advancements such as cloud computing, machine learning, and artificial intelligence (AI) further expanding the potential of Big Data. As businesses began to embrace the idea that data could be used as a resource for competitive advantage, the need for specialized infrastructure and tools to handle Big Data became critical. Advances in data storage technologies, particularly distributed databases like Hadoop and cloud services, made it possible to store and analyze vast amounts of data in real-time (Gartner, 2011).

However, as Big Data continues to shape decision-making across various sectors, several philosophical issues have emerged regarding its ethical implications, particularly its biases and discrimination in artificial intelligence. The data-driven algorithms often rely on historical data to make predictions or decisions, but if this data is biased, the outcomes will be too. Nevertheless, this problem will be examined in the next section.

Algorithms Bias in Big Data and Utilitarianism

The concept of an “algorithm” originates from the name of the Persian mathematician Al-Khwarizmi, who lived during the 9th century and made important contributions to mathematics and astronomy. The word is derived from his name and refers to a systematic method for solving problems through a series of defined steps or rules. Algorithms have been in use for centuries and continue to play a crucial role in shaping the modern world. Essentially, algorithms are step-by-step instructions that computers follow to address problems. They enable automation, data analysis, and faster decision-making. For instance, algorithms are applied in search engines, recommendation systems, and social media feeds to display relevant content (George & O’Neill, 2000). In essence, an algorithm functions like a set of instructions that tells a computer how to perform a task. Similar to a recipe for cooking, algorithms guide computers in performing a wide range of activities we engage in daily, such as searching the internet, suggesting movies or music, and providing directions. They are essential for helping computers solve problems and make decisions efficiently and accurately.

Machine learning algorithms are embedded in many aspects of daily life, influencing decisions with little transparency. These algorithms mimic human learning, becoming more accurate over time. They support decision-making in areas like business, healthcare, and e-commerce by identifying patterns in data. For instance, supervised classification helps in tasks like facial

recognition or product recommendations (Yapo, 2018). The machine learning process involves collecting, preparing, training, testing, and deploying models. By analyzing vast amounts of data, machine learning helps organizations respond quickly to trends (Suresh & Guttag, 2021). Its applications include chatbots, email filtering, virtual assistants, and fraud detection (Suresh & Guttag, 2021). Machine learning algorithms influence many decisions, such as search engine results and Facebook's news feed, with little visibility. However, it is important to recognize that AI systems are built using data and parameters defined by humans, and these systems can inherit biases from historical data, which may include discrimination against marginalized groups. When AI models are trained on biased data, they can perpetuate existing prejudices (Marr, 2021). AI tools can be biased not only in terms of technology but also due to the environments and individuals who design them. Most large AI systems are currently developed by a small group of elite institutions, which historically lack diversity. These institutions are often predominantly white, affluent, male, and technically focused, contributing to biases in AI development (Yapo, 2018).

Bias can enter machine learning at various stages, from data collection to algorithm development. For example, when developing a machine learning algorithm, if a company's focus is profit-driven, it could unintentionally lead to discriminatory outcomes. Bias can also emerge in the data collection phase, especially when the data fails to represent reality, leading to skewed results. Moreover, the data preparation phase can introduce bias depending on the selected attributes that shape the model (Hao, 2020). Despite the benefits machine learning offers to various sectors, research has uncovered numerous instances of bias in algorithms that result in negative societal impacts (Yapo, 2018).

Nonetheless, Utilitarianism influences algorithm design in two main ways: first, by optimizing for efficiency and utility, often favoring solutions that benefit the majority; and second, by justifying algorithmic outcomes based on overall societal benefit. Utilitarian approaches often prioritize maximizing positive outcomes, but at times, this leads to situations where certain groups; often the minority, are disproportionately harmed. For instance, if an algorithm used in a criminal justice system is designed to predict recidivism rates and is trained primarily on data from a specific demographic, it may favor decisions that benefit the majority group but unfairly disadvantage minorities. While this might be justified under a utilitarian perspective if it benefits the majority of society by improving the accuracy of predictions, the ethical implications for minority groups are significant.

A key example of utilitarianism in algorithmic bias is the use of risk assessment tools in the U.S. criminal justice system. These tools, such as the COMPAS algorithm, assess the likelihood that an offender will reoffend. However, several studies have shown that these tools exhibit racial bias, with black defendants being more likely to be flagged as high-risk, despite similar criminal histories compared to white defendants (Angwin, Larson, Mattu, & Kirchner, 2016). The aim of these risk assessment algorithms is to reduce human error and increase fairness in sentencing. However, the use of biased data—such as the overrepresentation of black defendants in historical criminal justice data—means that the algorithm is, in practice, reinforcing existing racial disparities. From a utilitarian viewpoint, these algorithms may increase overall efficiency and

reduce recidivism, but they perpetuate harmful biases for minority groups, which undermines their ethical justification. Also, a study conducted by Carnegie Mellon University in 2015 revealed gender bias in Google's advertising algorithm. The algorithm showed ads for high-paying executive jobs predominantly to male users, despite both male and female users searching for similar job-related queries. Research indicated that while Google showed ads for high-paying jobs 1,852 times for men, it only showed the same ads 318 times for women. This form of algorithmic bias can be justified by a utilitarian perspective if it is seen as optimizing for maximum ad engagement. However, such bias disproportionately disadvantages women, particularly when it comes to career advancement opportunities, perpetuating existing gender disparities.

One of the primary ethical criticisms of utilitarianism in algorithmic design is its tendency to prioritize the majority at the expense of the minority. Algorithms designed with the goal of maximizing societal benefit may inadvertently perpetuate inequality, especially if the data used to train the algorithm is already skewed in favor of certain groups. For example, an algorithm that optimizes for consumer behavior might benefit the majority, but it could overlook the needs of niche communities or marginalized consumers. In this way, the algorithm may perpetuate economic inequalities by favoring products or services that cater to a majority demographic, while ignoring those that could benefit minority groups.

Another central issue in utilitarian approaches to algorithmic bias is the lack of transparency in decision-making processes. Many machine learning models, particularly deep learning models, operate as "black boxes," where the decision-making process is not clearly understood by the public, developers, or even the organizations using them. This opacity can lead to situations where biased outcomes are justified without fully understanding or acknowledging the harm caused to affected groups. In the case of Facebook's news feed algorithm, for instance, the system's lack of transparency about how it amplifies certain types of content; such as fake news or polarizing political material—has led to significant societal consequences. While the algorithm may be optimizing for engagement (clicks, shares), it has led to an increase in misinformation, social division, and electoral interference, all of which have profound ethical implications.

Utilitarianism, while providing a clear framework for maximizing well-being, fails to address issues of equity and justice. A utilitarian approach may argue that it is acceptable to harm a minority if the overall benefit to the majority is large enough. This principle can be seen in several algorithmic systems, such as predictive policing algorithms that disproportionately target certain communities, or healthcare algorithms that may optimize for cost-saving at the expense of providing adequate care to vulnerable populations. Therefore, the ethical problem here is that utilitarianism often does not consider the distribution of benefits and harms based on the integrity of the subject or individuals. It is insufficient to only maximize the overall good if certain individuals or groups are being consistently harmed by those decisions without considering their personal integrity. Nevertheless, the way forward will be considered in the next section on the algorithmic bias in big data.

The Greatest Happiness or Personal Integrity: An ethical Reconstruction of Big data

Bernard Williams was a staunch critic of utilitarianism, particularly the classical form that emphasizes maximizing happiness as the central moral criterion. Williams' primary concern was that utilitarianism, with its focus on outcomes, often leads to morally questionable decisions by demanding individuals to disregard their personal commitments and integrity in favour of achieving the greatest good. He argued that this utilitarian calculus reduces moral actions to mere calculations of happiness, disregarding the complexities of human experience, personal responsibility, and ethical constraints. He argued that:

The moral integrity of the individual consists in the alignment of their actions with their deeper values, rather than the submission of those values to a utilitarian calculus that reduces them to mere instruments of maximizing happiness (Williams, 1985: 137). In his influential works, such as *Ethics and the Limits of Philosophy* (1985) and *Moral Luck* (1981), Williams explored how utilitarianism fails to account for the importance of moral integrity—the alignment of an individual's actions with their deeper values and commitments. He contended that utilitarianism often forces individuals to prioritize impersonal, abstract values like overall happiness over more personal, context-specific values that shape one's life and identity. In this sense, Williams' critique focuses on the alienating nature of utilitarianism, which strips individuals of their moral agency and forces them to treat their actions as mere tools for maximizing collective welfare. For Williams, moral philosophy should not be reduced to a simple cost-benefit analysis, as utilitarianism often demands. Instead, he emphasized the importance of moral integrity; the coherence of one's actions with their deeper sense of self and commitment to others. Integrity, for Williams, involves respecting individual autonomy, recognizing the significance of personal relationships, and confronting the complexities of moral choices without resorting to the simplistic calculus of maximizing happiness. This focus on authentic moral experience and the protection of individual integrity is particularly pertinent when addressing the ethical concerns surrounding big data and algorithmic decision-making.

One of the central ethical concerns in the context of big data is algorithmic bias—the phenomenon where algorithms produce discriminatory or unjust outcomes due to biased data, flawed assumptions, or systematic inequalities in society. Algorithmic bias can manifest in a variety of ways, such as biased hiring algorithms that disadvantage women or minority candidates, or predictive policing algorithms that disproportionately target black and brown communities. From a utilitarian perspective, algorithmic bias may seem acceptable if it leads to positive outcomes for the majority. For example, a predictive policing algorithm that targets high-crime neighborhoods might reduce crime rates overall, even if it disproportionately affects certain racial groups. In this case, the benefits of reducing crime are seen to outweigh the harm caused by discrimination, thus aligning with the utilitarian maxim of maximizing happiness for the greatest number.

However, Bernard Williams would strongly critique this approach. His concern with utilitarianism's application to big data and algorithmic decision-making is that it overlooks individual rights, moral agency, and the moral significance of fairness. From a Williamsian perspective, algorithmic bias is morally indefensible because it violates the integrity of individuals, particularly those who are disproportionately harmed by biased systems. Williams'

critique would argue that big data systems should not simply be evaluated by their outcomes (e.g., the total reduction in crime or improvement in efficiency), but by how they respect the moral rights and dignity of individuals. He argued that :

Moral integrity, for the individual, is not merely a matter of consistency or the following of rules; it involves a deep alignment between one's actions and the self-understanding of one's commitments, which cannot simply be reduced to utilitarian calculations (Williams, 1981: 83).

For Williams, moral integrity is compromised when individuals are reduced to statistical probabilities, as happens in the case of algorithmic decision-making. In his view, big data systems that perpetuate discrimination or inequality violate the integrity of those affected by them because they treat individuals as mere instruments to achieve a statistical goal, rather than as autonomous moral agents with inherent worth. The consequence of such systems is the alienation of individuals from their own moral agency and the erosion of their ability to make meaningful choices that align with their personal values.

Nevertheless, Williams' concept of moral integrity offers a human-centered critique of big data and its use in decision-making processes. Integrity, in Williams' view, is not merely about consistency or adherence to rules; it is about maintaining coherence between an individual's actions and their deeper values and commitments. In the context of big data, this means that the deployment of data-driven algorithms should consider not only their utilitarian benefits but also the moral costs; particularly in terms of how they affect the rights and dignity of individuals. Williams' critique of the greatest happiness principle would argue that the use of big data to optimize happiness is morally problematic when it undermines individual integrity. For example, consider a predictive policing algorithm that disproportionately targets minority communities based on historical crime data. Even if the algorithm reduces overall crime rates and is deemed to maximize social welfare, it does so at the expense of the moral integrity of the communities it targets. Individuals in these communities are treated as statistical variables rather than autonomous moral agents, and their rights to equal treatment and fair justice are compromised. From a Williamsian perspective, the moral cost of such systems; the erosion of individual dignity, the violation of autonomy, and the perpetuation of social inequality, must be weighed against the supposed utilitarian benefits. For Williams, this would represent a moral failure, as the pursuit of happiness for the majority cannot justify the undue harm done to marginalized individuals who are subjected to biased, data-driven decisions.

However, in response to algorithmic bias, Bernard Williams would advocate for a more nuanced and ethically grounded approach to the use of big data. Rather than relying on the simplistic utilitarian calculus of maximizing happiness, Williams would argue for an ethical framework that respects individual integrity, autonomy, and social justice. This would involve: Acknowledging the moral agency of individuals and recognizing their inherent worth, rather than reducing them to statistical probabilities or mere tools to achieve broader social goals; ensuring fairness and justice in the design and deployment of data-driven algorithms, with a focus on eliminating bias and protecting vulnerable populations from harm. Promoting transparency in algorithmic decision-making processes, ensuring that individuals have a clear understanding of how their data is being used and are able to hold institutions accountable for biased or discriminatory outcomes.

For Williams, the application of big data should be guided by a commitment to moral integrity, ensuring that technological advancements do not undermine the dignity and autonomy of individuals. This requires moral reflection, ongoing self-examination, and a willingness to challenge unjust systems; even if doing so comes at the cost of efficiency or the pursuit of aggregate happiness.

Conclusion

The discourse on Bernard Williams' critique of utilitarianism and its application to big data underscores the tension between the pursuit of collective welfare and the preservation of individual integrity. Williams' philosophical stance challenges the utilitarian notion that maximizing happiness for the greatest number justifies the sacrifice of personal commitments and values. His emphasis on moral integrity, as the alignment of one's actions with deeply held personal values and relationships, offers a compelling alternative to the utilitarian framework that dominates many contemporary ethical discussions, especially in the context of big data and algorithmic decision-making. In the era of big data, where algorithms increasingly shape decisions that affect people's lives, Williams' critique becomes particularly salient. The dangers of algorithmic bias, which often reflect societal inequalities, are exacerbated by a utilitarian approach that prioritizes overall efficiency or aggregate happiness at the expense of marginalized individuals. From Williams' perspective, this utilitarian calculus is morally inadequate because it undermines the integrity of those who are most vulnerable to discriminatory or unjust outcomes. By reducing individuals to mere statistical data points, big data systems can erode personal autonomy and alienate individuals from their moral agency. In response, Williams advocates for a more ethically grounded approach to technology, one that respects individual dignity and moral integrity. This perspective calls for fairness in algorithmic design, transparency in decision-making processes, and a commitment to social justice; elements that are crucial in safeguarding the moral rights of individuals. The protection of personal integrity, in Williams' view, is not an afterthought but a central ethical consideration that must guide the development and application of big data technologies.

Hence, Bernard Williams' critique offers a powerful moral framework for navigating the complexities of the modern world, reminding us that the true value of human life cannot be reduced to mere calculations of happiness or efficiency. It is through the preservation of individual integrity and the recognition of each person's moral agency that we can create systems; whether social, technological, or political, that reflects the full dignity and worth of human beings. This requires a shift in how we approach ethics in the age of big data, ensuring that technological advancements are aligned with values that protect and enhance the moral integrity of every individual.

***AGUNBIADE Saheed Anuoluwapo**

Lagos State University

Department of Philosophy

Saheed.agunbiade@Lasu.edu.ng

References

- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165-1188. <https://doi.org/10.2307/41703503>
- Ethel, A., Theodore, D., & Sheldon, P. (Eds.). (1984). *Great traditions in ethics* (5th ed.). Wadsworth.
- Fasoro, J. O. (2005). *Theories of ethics: An introductory reading in moral philosophy*. Majab Publishers.
- Frankena, W. (1997). *Ethics* (2nd edition). New Delhi: Prentice Hall of India.
- Gartner. (2011). *The definition of big data* [White paper]. Gartner. Retrieved from <https://www.gartner.com>
- George, W. R., & O'Neill, H. (2000). How algorithms work in decision-making systems. In *Handbook of decision support systems* (pp. 87-103). Springer. https://doi.org/10.1007/978-1-4471-0177-0_5
- Hao, K. (2020). The problem with bias in machine learning—and how to fix it. *MIT Technology Review*. <https://www.technologyreview.com/2020/10/01/1040078/the-problem-with-bias-in-machine-learning-and-how-to-fix-it/>
- Laney, D. (2001). 3D data management: Controlling data volume, velocity, and variety. *META Group Research Note*. Retrieved from <https://www.gartner.com>
- Manyika, J., Chui, M., Brown, B., Bughin, J., Dobbs, R., Roxburgh, C., & Huyett, J. (2011). Big data: The next frontier for innovation, competition, and productivity. *McKinsey & Company*. Retrieved from <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/big-data-the-next-frontier-for-innovation-competition-and-productivity>
- Marr, B. (2021). The big tech dilemma: AI and data bias. *Forbes*. <https://www.forbes.com/sites/bernardmarr/2021/07/07/the-big-tech-dilemma-ai-and-data-bias>
- Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.
- Mougalas, R. (2001). *Big Data: Challenges and opportunities for the digital economy*. Meta Group Report.
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown Publishing Group.
- Omeregbe, J. (1993). *Ethics: A systematic and historical study*. Joja Publishers.
- Ozumba, G. O. (2001). *A course text on ethics*. O.O.P. Limited.
- Shaw, W. (2002). *Business ethics* (4th ed.). Eve Howard.
- Skorupski, J. (2003). Ethics (2nd ed.). In N. Bunnim & E. P. Tusi-James (Eds.), *The Blackwell companion to philosophy* (pp. 945-964). John Wiley & Sons.
- Suresh, H., & Guttag, J. V. (2021). *Machine learning for healthcare: Predictive analytics and models*. MIT Press.

- Williams, B. (1981). *Moral luck: Philosophical papers 1973-1980*. Cambridge University Press.
- Williams, B. (1985). *Ethics and the limits of philosophy*. Harvard University Press.
- Yapo, D. (2018). *The ethics of machine learning and algorithmic decision making*. Springer.
- Yapo, E. (2018). Bias in machine learning algorithms: A case study of facial recognition and predictive algorithms. *Journal of AI Ethics*, 5(3), 5359–5366.