ETHICS OF TECHNOLOGY

Satyagraha and Ubuntu in the Age of AI: Gandhian and Mandela's Principles of Non-Violence as a Foundation for Ethical AI and Responsible Digital Governance

Xoliswa Mazibuko

Department of Political Science, Bangalore University, Bengaluru, Karnataka, India

Published November 2026

Abstract

The rapid proliferation of Artificial Intelligence (AI) and digital technologies has raised urgent questions about ethics, accountability, and human dignity in governance frameworks worldwide. While technical and regulatory approaches to AI ethics dominate current discourse, this paper argues that enduring philosophical traditions rooted in non-violence offer an underexplored yet valuable foundation for responsible digital governance. This study undertakes a comparative analysis of Mahatma Gandhi's principle of *Satyagraha* (truth-force) and Nelson Mandela's philosophy of *Ubuntu* (shared humanity and interconnectedness), examining how these frameworks of non-violent resistance, reconciliation, and moral restraint can inform contemporary debates on ethical AI and data governance. Gandhi's emphasis on truth, transparency, and non-harm (*Ahimsa*) resonates with growing demands for algorithmic accountability, explainability, and the prevention of AI-driven harm to individuals and communities. Similarly, Mandela's Ubuntu philosophy—"I am because we are"—offers a relational, community-centered ethic that challenges purely individualistic or profit-driven approaches to data privacy and technological deployment, advocating instead for inclusive, human-centered governance structures. Both leaders' commitment to reconciliation over domination further provides a moral compass for addressing power asymmetries inherent in AI systems, including concerns around surveillance, data colonialism, and unequal access to digital resources between the Global North and Global South. Employing a qualitative, comparative-analytical methodology, this paper draws on primary writings, speeches, and secondary literature on Gandhi and Mandela, mapping their core principles onto contemporary frameworks of ethical AI and responsible digital governance. The study contributes to interdisciplinary scholarship by bridging political philosophy, ethics, and technology governance, proposing a values-based model that complements existing regulatory and technical approaches. It argues that revisiting non-violent leadership traditions can offer policymakers, technologists, and institutions a more humane, accountable, and globally inclusive foundation for navigating the ethical challenges of the AI era.

Keywords: Satyagraha; Ubuntu; non-violence; ethical AI; digital governance; Gandhi; Mandela

1 Introduction

The governance of artificial intelligence has become one of the defining regulatory and ethical challenges of the twenty-first century. From the European Union's Artificial Intelligence Act to the OECD AI Principles and the UNESCO Recommendation on the Ethics of Artificial Intelligence, a global architecture of soft law and hard law has emerged in a remarkably short span of time. Yet a recurring critique of this architecture, voiced by scholars working from non-Western traditions, is that it is

normatively thin: it recycles a narrow set of Kantian, utilitarian, and liberal-rights concepts — autonomy, fairness, transparency, accountability — without asking what conception of the human person, of community, and of truth underwrites those concepts in the first place. A systematic review of AI ethics guidelines found extraordinary convergence around a handful of abstract principles but almost no engagement with how those principles are to be weighed against one another, or with alternative moral traditions that might ground them differently.

This paper takes up that critique by turning to two moral-political traditions that were forged not in seminar rooms but in decades of organised struggle against colonial and racial domination: Mohandas Gandhi's philosophy and practice of Satyagraha, developed in South Africa and India between 1906 and 1948, and the Southern African philosophy of Ubuntu, most influentially embodied in Nelson Mandela's leadership of South Africa's transition from apartheid and in the work of the Truth and Reconciliation Commission (TRC) that he set in motion. Both traditions share a preoccupation with three themes that are directly relevant to AI governance: the relationship between means and ends, the constitution of personhood through relationship rather than isolated autonomy, and the possibility of holding power accountable without resorting to domination or erasure of the adversary. These are not merely historical curiosities. They speak directly to contemporary anxieties about algorithmic violence, data extraction, opaque accountability structures, and the exclusion of affected communities from the design of the systems that govern them.

The paper proceeds in five further parts. The first reconstructs the philosophical core of Satyagraha, emphasising its non-violence (ahimsa), its insistence on truth (satya) as a lived and provisional practice, and its unity of means and ends. The second reconstructs Ubuntu as articulated by Mandela and by Southern African philosophers, with particular attention to its relational conception of personhood and its embodiment in the restorative justice model of the TRC. The third surveys the dominant paradigms of contemporary AI ethics and identifies the gaps that a Gandhian-Ubuntu synthesis can address. The fourth develops this synthesis, which the paper calls relational non-violence, and applies it across four domains of digital governance. The fifth considers objections and limitations before the paper concludes.

2 Gandhian Satyagraha: Truth-force, non-violence, and the unity of means and ends

Gandhi coined the term Satyagraha — from satya (truth) and agraha (firmness or force) — to distinguish his method of resistance from what he regarded as the passive and essentially weak connotations of the English phrase “passive resistance.” Satyagraha was, in Gandhi's own account, an active and demanding discipline: the deliberate willingness to accept suffering upon oneself rather than inflict it upon an opponent, in pursuit of a truth that both sides might ultimately come to share. Three features of this philosophy are especially pertinent to AI governance.

2.1 Ahimsa as a constraint on instrumental reason

Ahimsa, usually translated as non-violence, is frequently misunderstood as mere restraint from physical harm. For Gandhi it was a far more exacting standard: a refusal to treat any human being, including one's political adversary, as a mere instrument for one's own ends. This is significant for AI governance because so much of the current discourse on “AI harm” is instrumentally framed — harm is what a system produces as an externality of an otherwise legitimate optimisation objective (engagement, efficiency, profit, security). A Gandhian frame would instead ask a prior question: does the design of the system treat the people whose data it consumes, and whose lives it affects, as ends

in themselves, or as means to the system's objective function? Applied to recommender systems that maximise engagement through outrage, to predictive policing tools that treat entire neighbourhoods as risk populations, or to labour-platform algorithms that treat gig workers as fungible inputs, ahimsa reframes the ethical question from harm-minimisation at the margins to a more fundamental interrogation of the objective itself.

2.2 Satya: Truth as provisional and dialogic, not possessed

Gandhi repeatedly insisted that no one possesses the whole truth, and that Satyagraha therefore requires humility: the satyagraha acts on their best present understanding of truth while remaining open to correction through dialogue and even through the suffering inflicted upon them by the struggle itself. This has a direct analogue in debates about AI transparency and explain ability. Much of the current discourse treats explain ability as a technical property to be engineered into a system and then delivered, complete, to a user. A Gandhian conception of satya suggests instead that truth-telling about an AI system is an ongoing, dialogic relationship between developers, regulators, and affected communities — one in which the system's own limitations, uncertainties, and failure modes must be disclosed honestly and revisited as new evidence emerges, rather than asserted once in a static model card or terms-of-service document.

2.3 The unity of means and ends

Perhaps Gandhi's most quoted formulation is that “means and ends are convertible terms”: an unjust means can never reliably produce a just end, because the means shape the character of the outcome and of the actors who employ them. This principle bears directly on the common practice, in both AI development and platform governance, of justifying ethically questionable means — covert data collection, manipulative design patterns, adversarial testing on unconsenting populations — by appeal to beneficial downstream ends such as improved public health models or safer autonomous vehicles. The Gandhian critique is not that such ends are illegitimate, but that the legitimacy of the end cannot retroactively launder an illegitimate means; the means must themselves satisfy the standard of non-violence and truthfulness.

Finally, Gandhi's practice of Satyagraha was never purely individual; it depended on organised, collective, and disciplined civil resistance — the Salt March, the Champaran satyagraha, the mill workers' strikes in Ahmedabad — in which ordinary people, not only elites, exercised agency over decisions that affected their lives. This anticipates, in a striking way, contemporary calls for participatory AI governance, in which the people subject to algorithmic decision-making are not merely consulted after the fact but are structurally empowered to contest and shape the systems that govern them.

3 Ubuntu and Mandela's praxis: Relational personhood and restorative accountability

Ubuntu is a cluster concept found across Southern African Bantu languages, frequently rendered through the aphorism (umuntu ngumuntu ngabantu) — “a person is a person through other persons.” Where much of Western liberal moral and political theory begins from the autonomous individual as the basic unit of ethical and political analysis, Ubuntu holds that personhood itself is constituted relationally: one becomes a full moral person through one's relationships, obligations, and participation in community, and one's dignity is bound up with the dignity of others. Scholars have identified social harmony, social progress, and human dignity as three interdependent pillars of

the Ubuntu ethical framework, in contrast to frameworks that centre individual rationality and autonomy as the ground of moral status.

3.1 Ubuntu in Mandela's political practice

Nelson Mandela did not merely invoke Ubuntu rhetorically; he institutionalised it. The most consequential expression of Ubuntu in his presidency was his sponsorship of the Truth and Reconciliation Commission, chaired by Archbishop Desmond Tutu, which offered conditional amnesty to perpetrators of apartheid-era human rights violations in exchange for full public disclosure of the truth. Tutu himself explained the TRC's model explicitly in Ubuntu terms: retributive justice, focused on punishment, was subordinated to restorative justice, focused on the restoration of relationships, the dignity of victims, and the reintegration of offenders into a shared moral community. Mandela's own conduct — his willingness to negotiate with, and later govern alongside, individuals and institutions responsible for his twenty-seven years of imprisonment — modelled this same principle: accountability without dehumanisation of the adversary.

3.2 Ubuntu as a critique of extractive, individualist data regimes

Contemporary AI systems are built on data infrastructures that treat individuals as discrete, severable sources of information — a paradigm sometimes described by critics as data colonialist or surveillance capitalism, in which behavioural data is extracted from populations, often without meaningful consent, and converted into predictive assets for commercial or state actors. Ubuntu offers a structural critique of this paradigm rather than a merely procedural one. If personhood is relationally constituted, then data about a person is never purely individual property to be traded in a bilateral contract between a platform and a user; it is also, in part, data about the communities, families, and networks in which that person is embedded, and its use carries communal as well as individual stakes. This suggests that consent frameworks built solely on individual notice-and-consent are conceptually inadequate to the relational reality of data, and that governance mechanisms should incorporate collective and community-level consent and benefit-sharing, an idea increasingly echoed in African Union and African national AI strategy documents that explicitly invoke Ubuntu as a governance touchstone.

3.3 Restorative, not merely punitive, accountability

Much of current AI accountability policy is modelled on liability and punishment: fines under data protection law, sanctions for non-compliance, and litigation for algorithmic discrimination. An Ubuntu-inflected model would not dispense with sanction, but it would insist that accountability mechanisms also aim at restoration — repairing harmed relationships between technology companies and affected communities, requiring public acknowledgement and disclosure of algorithmic harms (analogous to TRC testimony), and reintegrating remediation into ongoing governance relationships rather than treating a fine as the end of the moral transaction. This has practical implications, discussed further in Section 5, for how algorithmic audit and redress mechanisms might be designed.

4 The contemporary AI ethics landscape and its gaps

A well-known survey of the global proliferation of AI ethics guidelines found that documents issued by governments, companies, and civil-society bodies converge overwhelmingly on a small set of principles: transparency, justice and fairness, non-maleficence, responsibility, and privacy.

Complementary work, such as the AI4People synthesis of AI ethics frameworks, has attempted to map these onto the four classical principles of bioethics — beneficence, non-maleficence, autonomy, and justice — with an additional principle of explicability. The UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted in 2021, extended this convergence to a genuinely global, intergovernmental instrument, and explicitly acknowledged the need to reflect cultural and regional diversity in AI ethics, including African, Asian, and Indigenous ethical traditions, rather than treating Euro-American liberalism as a universal default.

Nonetheless, critical scholarship has identified at least three persistent gaps in this convergent, principle-based approach. First, it tends to be procedurally and technically oriented — specifying what a system should disclose or how it should be audited — without addressing the underlying question of whose conception of the good, and whose conception of personhood, the system is designed to serve. Second, it is frequently critiqued as a form of "ethics washing," in which voluntary principles substitute for enforceable regulation and allow developers to claim moral legitimacy without altering underlying business models. Third, and most relevant here, it has been criticised for epistemic parochialism: a 2020 Harvard-based analysis argued that mainstream AI ethics imports, often uncritically, a liberal individualist anthropology in which the autonomous, rights-bearing individual is the basic unit of moral concern, and that this anthropology is poorly equipped to address harms that are fundamentally relational or communal in character — such as the erosion of a community's collective epistemic environment through algorithmically amplified misinformation, or the disproportionate surveillance of racial communities as aggregates rather than as individuals.

These gaps are precisely where Satyagraha and Ubuntu have most to offer: not as a wholesale replacement for existing principles such as fairness and transparency, but as an alternative philosophical grounding that changes how those principles are interpreted, prioritised, and operationalized, and that supplies concepts — non-violence as a constraint on instrumental optimisation, relational personhood, restorative accountability, truth as dialogic practice — that the dominant frameworks lack.

5 Relational non-violence: A synthesised framework for ethical AI

This paper proposes relational non-violence as a term for the synthesis of Satyagraha and Ubuntu applied to AI governance. The framework rests on four propositions, each of which is then applied to a domain of digital governance practice.

5.1 Proposition one: Non-violence by design

Following Gandhian ahimsa, an AI system should be evaluated not only for the harms it causes as an externality but for whether its core objective function instrumentalists the people subject to it. In practice, this suggests governance requirements that go beyond harm audits after deployment to require an ex ante justification of a system's optimisation target: does a content-recommendation algorithm optimise for sustained human attention and wellbeing, or for engagement metrics that are agnostic to, or actively exploit, psychological vulnerability? Regulatory instruments such as the EU AI Act's prohibition on AI systems that deploy subliminal or manipulative techniques causing significant harm gesture toward this idea, but a Satyagraha-informed standard would go further, requiring developers to demonstrate affirmatively that a system's objective treats affected persons as ends, not merely as sources of predictive signal.

5.2 Proposition two: Data justice as collective truth-telling

Combining Gandhian satya with Ubuntu's relational conception of data, governance regimes should treat transparency and truth-telling about AI systems as an ongoing, dialogic, and collective obligation rather than a one-time disclosure. This has concrete institutional implications: model cards and system disclosures should be living documents updated as failure modes are discovered; affected communities, not only individual data subjects, should have standing to request disclosure and to contest data use, echoing both the Gandhian emphasis on collective civil resistance and the Ubuntu-informed calls in African AI governance scholarship for community-level data stewardship rather than purely individualistic consent.

5.3 Proposition three: Restorative algorithmic accountability

Where an AI system causes harm — wrongful denial of a loan or parole recommendation, discriminatory hiring screen, or biased medical triage tool — accountability mechanisms modelled on the TRC's restorative logic would supplement financial penalties with mandated public acknowledgement of the harm, direct engagement with affected individuals and communities, and structural remediation of the system rather than mere compensation. Some emerging algorithmic impact assessment regimes, including provisions in Canada's Directive on Automated Decision-Making and elements of the EU AI Act's post-market monitoring requirements, already gesture toward iterative, disclosure-based accountability; a restorative framing would make explicit what is currently implicit, treating disclosure and repair of relationship as co-equal with financial sanction.

5.4 Proposition four: Participatory governance as collective Satyagraha

Gandhi's mass civil-disobedience campaigns and Ubuntu's communal decision-making practices both model forms of governance in which those subject to power are structurally enabled to contest and shape it, rather than being limited to post hoc consultation. Applied to AI governance, this suggests institutional mechanisms such as binding community review boards for high-risk public-sector AI deployments (e.g., predictive policing, welfare-fraud detection, immigration risk-scoring), participatory data trusts in which affected populations hold real governance rights over how their aggregated data is used, and formal channels through which affected communities can exercise something analogous to civil disobedience — for example, rights to collective opt-out or algorithmic contestation — without being penalised for doing so.

Table 1 summarises how each proposition maps onto a corresponding gap in mainstream AI ethics discourse: non-violence by design responds to the gap in objective-function scrutiny; collective truth-telling responds to the gap in relational data ethics; restorative accountability responds to the gap in purely punitive liability models; and participatory governance as collective Satyagraha responds to the gap in meaningful, structurally empowered public participation.

Table 1. Propositions of relational non-violence and the gaps they address in mainstream AI ethics.

Proposition	Gap in mainstream AI ethics discourse
Non-violence by design	Objective-function scrutiny
Data justice as collective truth-telling	Relational data ethics
Restorative algorithmic accountability	Purely punitive liability models
Participatory governance as collective Satyagraha	Meaningful, structurally empowered public participation

5.5 Illustrative application: Facial recognition and predictive policing

The four propositions can be illustrated through a single recurring governance controversy: the deployment of facial recognition and predictive policing tools by law-enforcement agencies, which have been repeatedly documented to misidentify and over-surveil racialised and low-income communities. A non-violence-by-design review would ask, prior to procurement, whether the system's objective function — typically, maximising predictive "accuracy" against historical arrest data — encodes and amplifies the very patterns of over-policing that produced the training data, thereby instrumentalising the communities most already burdened by state violence. A collective-truth-telling requirement would obligate the deploying agency to disclose, in an ongoing and public manner, the system's error rates disaggregated by race, neighbourhood, and other salient community categories, rather than relying on a vendor's proprietary accuracy claims. A restorative-accountability mechanism would require the agency, upon a documented wrongful identification, to acknowledge the harm publicly, to engage directly with the affected individual and community rather than settling the matter through a sealed settlement, and to revise the underlying system rather than merely compensating the victim. Finally, participatory governance as collective Satyagraha would give the affected community itself — not only elected officials or agency technocrats — a binding role in deciding whether the tool is deployed at all, echoing the way satyagrahis and Ubuntu-based community councils have historically claimed the right to consent to, or refuse, interventions imposed upon them from above. Several municipalities that have banned or paused police use of facial recognition after sustained community organising, and jurisdictions that have instituted binding community oversight boards for surveillance technology procurement, illustrate that elements of this model are already emerging in practice, even where they are not explicitly framed in Gandhian or Ubuntu terms.

A second brief illustration concerns generative AI and misinformation. A Gandhian conception of satya as dialogic and provisional truth-telling suggests that platforms hosting generative AI content should treat truthfulness not as a filter applied once at the point of publication but as a continuing conversation with users — surfacing uncertainty, provenance, and correction over time — while an Ubuntu-informed lens reminds designers that the harm of synthetic misinformation is rarely borne by an isolated individual; it degrades the shared epistemic commons on which communities depend to reason and act together, and governance responses should therefore include collective remedies, such as community-level media literacy investment and platform-funded local journalism, alongside individual content-moderation tools.

6 Objections, tensions, and limitations

Three objections deserve direct engagement. First, both Gandhi and the post-apartheid Ubuntu-based settlement have been criticised, including by contemporaries such as B. R. Ambedkar in Gandhi's case, for tolerating or insufficiently confronting structural hierarchy — caste in Gandhi's India, and unresolved economic inequality in South Africa's negotiated transition. Any application of these traditions to AI governance must avoid romanticising them and must explicitly ask whether a non-violent, restorative approach to algorithmic harm risks under-correcting for entrenched structural injustice, particularly where corporate power asymmetries are vast and the incentive to treat restorative gestures as reputational cover is strong. The framework proposed here therefore treats restorative accountability as a supplement to, not a replacement for, binding regulation and financial liability.

Second, there is a risk of philosophical appropriation: extracting Satyagraha and Ubuntu from their specific historical, linguistic, and political contexts to serve a global AI-governance agenda authored predominantly by Western or elite institutions would itself replicate the epistemic extraction that Ubuntu-centred critiques of AI ethics warn against. Scholarship on Ubuntu and AI governance has explicitly cautioned that Ubuntu is sometimes invoked as a decorative gloss on frameworks that remain substantively unchanged, without engaging African philosophers, policymakers, and communities as co-authors of the resulting governance models. This paper's synthesis should therefore be read as an invitation to further, contextually grounded research and participatory policy development rather than as a completed governance blueprint.

Third, non-violence and restoration are demanding ethical standards that presuppose good-faith actors willing to be held accountable through dialogue; they may be poorly suited to governing adversarial actors — authoritarian states deploying AI for mass surveillance, or malicious actors using generative AI for disinformation and fraud — who have no interest in truth-telling or restoration. For these cases, the framework developed here should be understood as governing the design and accountability obligations of legitimate developers and public institutions, operating alongside, not instead of, robust security, law-enforcement, and sanction regimes for bad-faith actors.

7 Conclusion

This paper has argued that Gandhian Satyagraha and Mandela's Ubuntu-informed praxis, though developed in anti-colonial and anti-apartheid struggles far removed from contemporary AI systems, supply conceptual resources that mainstream AI ethics frameworks largely lack: a constraint on instrumental reason grounded in non-violence, a dialogic and provisional conception of truth-telling, a relational rather than atomistic conception of personhood and data, and a restorative rather than purely punitive model of accountability. The resulting synthesis, relational non-violence, does not ask policymakers to discard existing principles of fairness, transparency, and accountability, but to reground them in an anthropology that takes seriously the communal and dialogic nature of human dignity. As AI systems increasingly shape who is seen, heard, believed, employed, policed, and cared for, the question of which philosophical traditions get to define “responsible” governance is not incidental but constitutive of justice itself. Bringing Satyagraha and Ubuntu into the centre of that conversation — as living traditions with their own internal debates and limitations, rather than as decorative citations — is one step toward a genuinely plural foundation for ethical AI and responsible digital governance.

References

[1] African Union (2020). Digital Transformation Strategy for Africa (2020–2030). *Addis Ababa: African Union Commission*.

[2] Ambedkar BR (1945). What Congress and Gandhi Have Done to the Untouchables. *Bombay: Thacker*.

[3] Bostrom N, Yudkowsky E (2014). The ethics of artificial intelligence. In: *K. Frankish & W. M. Ramsey (Eds.), The Cambridge Handbook of Artificial Intelligence*. Cambridge University Press. p. 316–334.

[4] European Commission (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.

[5] Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, et al. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.

[6] Gandhi MK (1927). An Autobiography, or The Story of My Experiments with Truth. *Ahmedabad: Navajivan Publishing House*.

[7] Gandhi MK (1938). Hind Swaraj, or Indian Home Rule. *Ahmedabad: Navajivan Publishing House*.

[8] Gandhi MK (1951). Satyagraha in South Africa (V. *G. Desai*, Trans.). Ahmedabad: Navajivan Publishing House.

[9] Goffi ER (2022). Respecting cultural diversity in ethics applied to AI: A new approach for multicultural governance. *Revista Misión Jurídica*, 15(23).

[10] Jobin A, Ienca M, Vayena E (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.

[11] Mandela N (1994). Long Walk to Freedom: The Autobiography of Nelson Mandela. *Boston: Little*, Brown and Company.

[12] Mantena K (2012). Another realism: The politics of Gandhian nonviolence. *American Political Science Review*, 106(2), 455–470.

[13] Mhlambi S (2020). From Rationality to Relationality: Ubuntu as an Ethical and Human Rights Framework for Artificial Intelligence Governance (Carr Center Discussion Paper Series 2020-009). *Cambridge*, MA: Harvard Kennedy School, Carr Center for Human Rights Policy.

[14] Metz T (2011). Ubuntu as a moral theory and human rights in South Africa. *African Human Rights Law Journal*, 11(2), 532–559.

[15] Metz T (2021). A Relational Moral Theory: African Ethics in and beyond the Continent. *Oxford: Oxford University Press*.

[16] O'Neil C (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. *New York: Crown*.

[17] Presidential Advisory Committee (2020). Ethical AI in Africa: Ubuntu and context-aware governance. In *O. Emeka & A*. Ogude (Eds.), Ubuntu and Contemporary Global Issues. Pretoria: African Institute Press.

[18] Truth and Reconciliation Commission of South Africa (1998). Truth and Reconciliation Commission of South Africa Report (Vols. *1–5)*. Cape Town: TRC.

[19] Tutu D (1999). No Future Without Forgiveness. *New York: Doubleday*.

[20] UNESCO (2021). Recommendation on the Ethics of Artificial Intelligence. *Paris: United Nations Educational*, Scientific and Cultural Organization.

[21] Van Norren DE (2021). The ethics of artificial intelligence, UNESCO and the African Ubuntu perspective. *Journal of Information, Communication and Ethics in Society*, 21(1), 112–128.

[22] Winner L (1980). Do artifacts have politics? *Daedalus*, 109(1), 121–136.

[23] Zuboff S (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. *New York: PublicAffairs*.