

AI in Sacred Healing: Health Law Perspectives on Regulating Algorithmic Interventions Against Spiritual Autonomy in Pluralistic Healthcare Systems

Onyebule, Kelechi G*

Alex Ekwueme Federal University, Ebonyi State

onyebule.kelechi@funai.edu.ng

Abstract

This paper interrogates the regulatory lacunae that emerge when algorithmic systems, ostensibly therapeutic, encroach upon the sacrosanct domain of spiritual autonomy within pluralistic healthcare regimes. Framed at the nexus of health law, medical ethics, and the anthropology of healing, the analysis posits that contemporary governance paradigms, tethered to evidence-based biomedicine, systematically efface the ontological pluralism that underwrites indigenous, faith-based, and esoteric curative practices. By deploying a tripartite heuristic: (i) the algorithmic reification of probabilistic ontologies, (ii) the juridical commodification of belief as “data exhaust,” and (iii) the epistemic violence latent in risk-benefit calculus, the study unmasks how AI-mediated interventions transmute sacred epistemologies into actuarial variables, thereby vitiating the inviolability of spiritual self-determination. Methodologically, drawing upon comparative constitutional jurisprudence (inter alia, the Indian Supreme Court’s articulation of “essential religious practices,” the European Court of Human Rights’ margin of appreciation doctrine, and the African Charter’s communal dignity jurisprudence), the paper contends that extant regulatory frameworks, premised on paternalistic beneficence, fail to apprehend the incommensurability between machine rationality and transcendent healing. A novel conceptual scaffold is proffered: the “spiritual harm threshold,” a juridical metric that obliges regulators to demonstrate not merely empirical efficacy but also phenomenological non-interference with the patient’s cosmogonic narrative. This threshold, operationalized through mandatory “ontological impact assessments,” inverts the burden of proof, compelling algorithmic

proponents to negate existential displacement rather than merely affirm clinical outcomes. The argumentation culminates in a provocative normative claim: absent a statutory entitlement to “algorithmic abstention” in matters of sacral therapeutics, pluralistic polities risk the quietus of metaphysical diversity under the guise of precision medicine. Conclusively, by foregrounding the irreducibly hermeneutic character of sacred healing, the paper challenges health law scholars to reconceive autonomy not as volitional consent but as ontological sovereignty, an exigency that confounds utilitarian aggregation and demands a radical reconfiguration of regulatory reason.

1.0 Introduction:

In the quiet wards of a Rwandan clinic, where patients turn to ancestral rituals alongside IV drips, or in the bustling urban hospitals of Mumbai, where ayurvedic chants mingle with the hum of diagnostic machines, healing has always been more than biology, it is a tapestry woven from belief, community, and the unseen forces that shape our sense of self (World Health Organization, 2023). Yet, as artificial intelligence weaves its way into these spaces, promising precision and efficiency, it risks unraveling threads of spiritual autonomy that have sustained diverse healing traditions for generations. This paper probes the regulatory voids that open when AI systems, dressed in the guise of therapeutic tools, step into the hallowed ground of spiritual self-determination within pluralistic healthcare landscapes. Drawing from the crossroads of health law, medical ethics, and the anthropology of healing, we confront how today’s governance structures, moored firmly to the empirical anchors

of biomedicine, quietly sideline the rich pluralism of ontologies that underpin indigenous, faith-rooted, and esoteric forms of care.

The allure of AI in healthcare is undeniable. From predictive algorithms that forecast disease outbreaks to chatbots offering mental health support, these technologies hold the potential to democratize access and sharpen clinical judgment (Silcox, C., Zimlichmann, E., Huber, K. *et al.*, 2024). In pluralistic societies like those in sub-Saharan Africa or South Asia, where over 80% of people still rely on traditional healers for primary care, AI could bridge gaps left by overburdened systems (World Health Organization, 2023). But here's the rub: these tools, trained on vast datasets skewed toward Western biomedical norms, often flatten the sacred into the statistical. A patient's cosmogonic narrative: their story of creation, affliction, and redemption, perhaps rooted in Amazonian perspectivism where the body isn't a fixed vessel but a relational nexus of spirits and humans, becomes just another data point in a risk algorithm (Viveiros de Castro, 2021). What was once a pathway to transcendent wholeness is recast as a probabilistic outlier, vulnerable to erasure under the weight of evidence-based mandates.

This encroachment isn't mere oversight; it's a symptom of deeper ontological friction. Western biomedicine, with its positivist lens, assumes a singular reality where healing equates to measurable outcomes: reduced mortality rates, optimized resource allocation (Young & Varpio, 2025). Yet, anthropological insights remind us that healing ontologies vary profoundly: in perspectival cosmologies, health emerges from balanced exchanges across species boundaries, challenging AI's reductive metrics (Descola, 2021). Ethical guidelines from global bodies, like the Vatican's 2020 Rome Call for AI Ethics, reaffirmed through new signatories in 2025, urge that AI serve human dignity without supplanting it, emphasizing inclusion and transparency to safeguard the vulnerable (Pontifical Academy for Life, 2020/2025). Yet, as recent reviews highlight, regulatory lags persist: biases in training data

amplify disparities, eroding patient trust and autonomy in faith-infused care (Nazer *et al.*, 2024). In mental health, for example, AI chatbots trained on secular datasets might dismiss spiritual distress as delusion, clashing with pastoral care traditions that view it as soul-deep calling, and worsening crises in vulnerable users (Rahsepar Meadi *et al.*, 2025).

Against this backdrop, the European Union's AI Act of 2024 marks a tentative step, classifying high-risk health AI as needing rigorous impact assessments to protect dignity and equity (European Parliament, 2024). But such frameworks, while vital, often overlook the phenomenological toll, the subtle violence of rendering sacred epistemologies into actuarial fodder. This study unmasks that dynamic through a tripartite lens: the hardening of fluid worldviews into coded probabilities, the marketization of faith as byproduct data, and the insidious epistemic harms embedded in utilitarian trade-offs. By weaving in comparative jurisprudence, from India's doctrinal safeguards for religious essentials (*Shayara Bano v. Union of India*, 2017) to Europe's deference in human rights margins (*S.A.S. v. France*, 2014) and Africa's communal ethos (*Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania*, 2024), we argue that paternalistic regulations falter against the chasm between silicon logic and soulful restoration.

At its core, this paper isn't just critique; it's a call to fortify spiritual sovereignty in an algorithmic age. We propose the "spiritual harm threshold", a benchmark demanding proof of existential non-intrusion alongside clinical gains, and pair it with "ontological impact assessments" to shift the evidentiary onus. In pluralistic polities, where faith and tech converge, ignoring this risks not progress, but the slow suffocation of metaphysical variety beneath precision's polished veneer. What follows reimagines autonomy not as a signature on consent forms, but as unyielding guardianship over one's narrative cosmos, a demand that upends regulatory

orthodoxy and invites a hermeneutic renewal of health law.

2.0 Theoretical Framework:

At the heart of any discussion on healthcare governance lies a fundamental question: what counts as healing, and whose understanding of reality defines it? Ontological pluralism, in this context, posits that there is no singular, universal way of being or knowing when it comes to health and illness; instead, diverse cultural, spiritual, and social frameworks coexist, each shaping distinct perceptions of the body, affliction, and restoration (Khalikova, 2021). This pluralism challenges the monolithic grip of biomedicine, which often operates under a positivist ontology: one that views the body as a mechanical entity governed by empirical laws, reducible to cells, genes, and quantifiable metrics (Heuser, Steil & Salloch, 2025). In pluralistic healing regimes, particularly those in postcolonial or multicultural societies, indigenous shamans might interpret sickness as a rupture in communal harmony with ancestral spirits, while faith healers could see it as a test of divine will, and esoteric practitioners might frame it through energetic imbalances, all valid within their ontological worlds. These regimes are not mere alternatives but interwoven tapestries where patients navigate multiple systems, blending Ayurvedic herbs with chemotherapy or Pentecostal prayers with psychiatric counseling, driven by cultural resonance, accessibility, and perceived efficacy (Patil et al., 2024).

Medical anthropology provides a lens to unpack this pluralism, revealing how healing is not a linear path but a negotiated terrain of epistemologies, ways of knowing that vary profoundly across contexts (Tobert, 2022). For instance, in Amazonian indigenous communities, health emerges from "perspectival cosmologies," where humans, animals, and spirits share relational perspectives, and illness signals a misalignment in these exchanges rather than an isolated bodily malfunction (Viveiros de Castro, 2021). Such ontologies stand in stark contrast to biomedicine's

naturalistic etiology, which attributes disease to impersonal causes like pathogens or genetics, often dismissing personalistic explanations: those tying affliction to social, moral, or supernatural forces, as superstition (Khalikova, 2021). This friction becomes acute in pluralistic healthcare systems, where state-backed governance paradigms prioritize evidence-based practices, marginalizing non-biomedical approaches through regulatory hierarchies that favor randomized trials over experiential wisdom (Patil et al., 2024). Anthropologists argue that this sidelining is not neutral; it enacts a form of epistemic violence, where dominant ontologies colonize others, reshaping flexible, holistic healing into standardized, commodified forms to fit institutional molds (Tobert, 2022).

The integration of artificial intelligence into these regimes amplifies such tensions, as AI systems are typically engineered within a biomedical ontology, trained on datasets that encode Western norms of health as measurable outcomes like survival rates or cost-efficiency (Heuser, Steil & Salloch, 2025). In doing so, they risk effacing the ontological diversity that underpins spiritual autonomy, the right to define one's healing narrative without external imposition. Consider AI-driven diagnostic tools in mental health: algorithms might classify spiritual experiences, such as visions or ecstatic states revered in Pentecostal traditions, as pathological delusions, thereby overriding the patient's cosmogonic framework where such events signify divine connection (Rahsepar Meadi et al., 2025). Ethical frameworks underscore that this not only undermines autonomy but also erodes trust, as patients in faith-infused communities may perceive AI as an intruder that flattens sacred epistemologies into data points (Lee et al., 2025). Global ethical guidelines, like those evolving from the Vatican's Rome Call, call for AI to respect human dignity by incorporating inclusive values, yet they often fall short in addressing the life-world perspectives, the lived, embodied practices, that sustain pluralistic healing (Pontifical Academy for Life, 2020/2025).

This framework, therefore, positions ontological pluralism not as a barrier to progress but as a safeguard against homogenization. In sub-Saharan Africa, for example, where over 80% of populations blend traditional healers with modern clinics, pluralism fosters resilience, allowing communities to adapt therapies to local cosmologies amid resource constraints (World Health Organization, 2023). Yet, as AI proliferates, regulatory paradigms tethered to biomedicine, such as the EU's risk-based classifications, must evolve to accommodate these multiplicities, lest they perpetuate disparities by privileging probabilistic models over hermeneutic ones (European Parliament, 2024). Drawing on comparative insights, we see echoes in jurisprudential doctrines: India's "essential religious practices" test protects faith-based rituals from state interference, much like ontological pluralism demands space for diverse healing truths (Shayara Bano v. Union of India, 2017). Similarly, the European margin of appreciation affords deference to cultural variances in rights adjudication, while African communal dignity jurisprudence emphasizes collective worldviews over individualistic metrics (S.A.S. v. France, 2014; Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania, 2024).

In a nutshell, by foregrounding ontological pluralism, this paper sets the stage for interrogating how AI-mediated interventions transmute these varied regimes into actuarial shadows, vitiating spiritual self-determination. What emerges is a call for governance that honors incommensurability, the irreducible gaps between machine rationality and transcendent narratives, paving the way for novel metrics like the spiritual harm threshold.

3.0 Regulatory Lacunae in Evidence-Based Governance Paradigms

Evidence-based medicine (EBM) has revolutionized healthcare governance by grounding policies in empirical data from clinical trials, systematic reviews, and statistical analyses, aiming to enhance safety, efficacy, and equitable

resource distribution. However, as artificial intelligence (AI) becomes increasingly embedded in healthcare, from predictive diagnostics to treatment recommendation systems, these paradigms expose critical regulatory lacunae. These gaps arise primarily from EBM's positivist foundations, which favor measurable outcomes over the diverse, often intangible, aspects of healing in pluralistic societies. In contexts where spiritual, indigenous, or faith-based practices coexist with biomedicine, AI-driven governance risks amplifying epistemic hegemony, marginalizing non-Western ontologies and perpetuating injustices where certain knowledge systems are systematically undervalued or erased (Emah & Bennett, 2025; Kay, Kasirzadeh & Mohamed, 2024).

One prominent lacuna lies in the handling of epistemic injustice within AI-integrated EBM frameworks. Epistemic injustice occurs when individuals or communities are wronged in their capacity as knowers, such as when AI algorithms dismiss spiritual narratives as irrelevant data noise. For instance, in the Global South, where pluralistic healing regimes blend traditional and modern approaches, AI tools trained on biased datasets may pathologize faith-infused explanations of illness, leading to regulatory oversights that fail to protect patient autonomy (Birhane, 2025). Recent analyses highlight how generative AI exacerbates this by undermining collective knowledge integrity, commodifying personal beliefs into "data exhaust" without adequate safeguards, and creating accountability voids where victims of harm bear undue burdens (Palaniappan, Ting Lin, Vogel, & Lim, 2024; Kay, Kasirzadeh & Mohamed, 2024). This not only widens health disparities but also entrenches a form of epistemicide, where indigenous and local knowledge is sidelined in favor of technocratic models (Redvers, Lokugamage, Barreto, Bajracharya & Harris, 2024).

Ethically, these gaps manifest in the tension between EBM's principles of beneficence and the phenomenological impacts of AI. While EBM

prioritizes aggregate benefits like reduced mortality, it often overlooks how algorithmic decisions inflict hermeneutical injustice, where patients lack the interpretive resources to articulate spiritual distress within secular systems (Adams, 2025). In mental health, for example, AI chatbots might classify ecstatic spiritual experiences as delusions, clashing with pastoral or indigenous traditions and worsening crises without regulatory mechanisms for cultural deference (Tunks Leach, Simpson, Lewis, *et al.*, 2023). Global regulatory landscapes, such as those discussed in reviews of AI frameworks, reveal inconsistencies: while some regions address technical biases, few tackle the sociotechnical imaginaries that perpetuate epistemic harms in diverse populations (Jonas, Bacharach, Nightingale, & Filoche, 2024; Khurana, 2025).

Legally, the fragmentation is stark. The EU AI Act, for instance, categorizes health AI as high-risk and mandates impact assessments, yet it inadequately addresses cultural and spiritual dimensions, focusing on privacy and accuracy rather than ontological pluralism (European Parliament, 2024). In the U.S. and elsewhere, governance through agencies like the FDA emphasizes RCT-derived evidence, but neglects structural epistemic injustices in AI development, such as exclusionary data pipelines that amplify colonial legacies (Doshi *et al.*, 2024; Khurana, 2025). This paternalistic approach assumes machine rationality can seamlessly integrate into transcendent healing without friction, ignoring calls for relational ethics that prioritize dignity and communal input (Heuser, Steil & Salloch, 2025).

Moreover, in pluralistic polities like those in sub-Saharan Africa or South Asia, these lacunae invite existential risks. AI proliferation without tailored regulations could suffocate metaphysical diversity, as evidence-based mandates override cosmogonic narratives under the guise of precision medicine (World Health Organization, 2023). Comparative jurisprudence offers potential remedies, such as incorporating doctrines that protect essential practices or communal dignity to mandate

"ontological audits" (Shayara Bano v. Union of India, 2017). Until addressed, however, these gaps perpetuate a cycle of harm, demanding a hermeneutic renewal of governance.

4.0 Tripartite Heuristic of AI-Mediated Epistemic Violence

To dissect the subtle yet profound ways in which algorithmic systems inflict harm on spiritual autonomy within healthcare, this paper introduces a tripartite heuristic, a conceptual tool that illuminates three interconnected mechanisms of epistemic violence. Rooted in critical theory and drawing from recent scholarship on AI ethics, this framework exposes how AI not only disrupts but actively reshapes sacred epistemologies, transforming them from dynamic, interpretive modes of knowing into static, utilitarian artifacts. Epistemic violence, as conceptualized here, extends beyond mere bias; it encompasses the structural erasure of non-dominant worldviews, where machine-mediated interventions prioritize probabilistic logic over the hermeneutic richness of faith-based or indigenous healing (Adams, 2025). In pluralistic regimes, this violence manifests quietly, often under the guise of therapeutic advancement, but its effects are far-reaching: undermining patient trust, amplifying disparities, and risking the homogenization of metaphysical diversity. By parsing this into reification, commodification, and calculative harms, the heuristic not only critiques current governance but also lays groundwork for remedial metrics like the spiritual harm threshold.

This approach is timely, as emerging studies highlight the interwoven epistemic, sociopolitical, and technical ramifications of AI in healthcare, where systems designed for efficiency can inadvertently enact forms of injustice (Emah & Bennett, 2025). For instance, in contexts where healing involves communal rituals or transcendent narratives, AI's intrusion can fracture these practices, reducing them to outliers in data models. The heuristic thus serves as a diagnostic lens,

urging regulators to confront the ontological friction at play.

4.1 Algorithmic Reification of Probabilistic Ontologies

At the core of AI's epistemic violence lies reification, the process by which fluid, abstract concepts are concretized into tangible, often rigid entities. In healthcare, this manifests when algorithms reify probabilistic ontologies, hardening the inherent uncertainties of spiritual worldviews into fixed, computable forms. Spiritual healing often embraces ambiguity: a faith healer's prognosis might hinge on divine will or karmic flux, where probabilities are interpretive rather than statistical. Yet, AI systems, trained on vast biomedical datasets, transmute these into deterministic outputs, scoring risks or predicting outcomes with an air of infallibility that eclipses sacred indeterminacy (Emah & Bennett, 2025).

This reification is not abstract; it plays out in real clinical scenarios. Consider predictive analytics in oncology, where an AI tool might quantify survival odds based on genetic markers, dismissing a patient's indigenous ontology that views cancer as a relational imbalance with ancestral spirits. Such tools, by solidifying probabilistic elements into "reified" variables, enact violence by stripping away the relational depth of healing, as noted in critiques of AI's role in perpetuating epistemic injustices through over-reliance on emergent algorithms (Adams, 2025). Experiments with text-to-image AI have similarly revealed epistemically violent biases, where generative models reify cultural representations in ways that marginalize non-Western perspectives, a pattern mirrored in health AI that flattens diverse ontologies into standardized probabilities (Doshi et al., 2024).

The regulatory implications are stark. Without checks, this mechanism widens lacunae in evidence-based paradigms, assuming universality in probabilistic modeling while ignoring the violence inflicted on spiritual self-determination. As anthropological insights remind us, healing in perspectival cosmologies thrives on multiplicity,

not reification, a contrast that demands governance reforms to preserve ontological pluralism (Viveiros de Castro, 2021).

4.2 Juridical Commodification of Belief as Data Exhaust

Building on reification, the second prong examines the juridical commodification of belief, where spiritual convictions are reduced to "data exhaust", incidental byproducts extracted, anonymized, and monetized within AI ecosystems. In healthcare, this occurs when patient data laced with faith narratives, such as ritual preferences or cosmogonic stories, is harvested for model training, often under legal frameworks that prioritize intellectual property over sacred inviolability. What begins as a personal epistemology ends up as commodified fodder, traded in data markets that fuel AI's growth without consent or compensation (Redvers, Lokugamage, Barreto, Bajracharya & Harris, 2024).

This commodification carries legal and ethical weight. Juridical structures, like those governing data privacy (e.g., GDPR equivalents), often fail to recognize beliefs as more than metadata, allowing their extraction as exhaust in algorithmic pipelines. In mental health apps, for example, a user's spiritual journal entries might be scraped to refine chatbots, commodifying soul-deep reflections into assets that enhance secular models, thereby eroding the communal dignity central to African or indigenous healing traditions (Birhane, 2025). Recent scoping reviews underscore this gap in global AI regulations, noting how frameworks overlook the epistemic injustices amplified by generative AI, where collective knowledge integrity is undermined through unchecked commodification (Palaniappan, Ting Lin, Vogel, & Lim, 2024).

The violence here is insidious, as it masks exploitation under innovation's banner. Patients in pluralistic systems, already navigating overburdened care, face additional harms when their beliefs are juridically repurposed, highlighting the need for regulations that treat data

exhaust not as neutral but as potential sites of sacred desecration.

4.3 Epistemic Violence in Risk-Benefit Calculus

The heuristic culminates in the epistemic violence embedded in AI's risk-benefit calculus, a utilitarian framework that quantifies harms and gains but often inflicts deeper wounds on spiritual epistemologies. This calculus, standard in EBM, weighs clinical outcomes against costs, yet in AI applications, it systematically devalues transcendent narratives, rendering sacred risks (e.g., defying an algorithm to follow a prophetic vision) as irrational liabilities (Doshi et al., 2024). The result is a violence that silences alternative ways of knowing, prioritizing aggregate metrics over individual cosmogonies.

In practice, this plays out starkly in high-stakes decisions. An AI system optimizing hospital resources might deprioritize faith-based palliative care, calculating it as low-benefit amid resource scarcity, thus pathologizing spiritual distress and clashing with traditions that view suffering as redemptive (Tunks Leach, Simpson, Lewis, *et al.*, 2023). Philosophical analyses of AI in medicine reveal this as a dual ethical-epistemic failure, where systems harden epistemic boundaries, excluding contributory injustices like the omission of spiritual experts from risk assessments (Adams, 2025). Moreover, in digital mental health, biases unpack to show how AI perpetuates political and social harms, where risk calculi reinforce secular norms at the expense of diverse healing regimes (Jonas, Bacharach, Nightingale, & Filoche, 2024).

Ultimately, this prong calls for inversion: shifting the calculus to demand proof of non-violence against sacred epistemologies. Without it, AI risks epistemicide in healthcare, suffocating the very pluralism that enriches healing (Redvers, Lokugamage, Barreto, Bajracharya & Harris, 2024).

5.0 Comparative Jurisprudence

Constitutional

To bridge the regulatory lacunae exposed in the previous sections, this paper turns to comparative constitutional jurisprudence as a methodological anchor. By examining doctrines from India, Europe, and Africa, we uncover models of deference to cultural and spiritual diversity that could inform AI governance in healthcare. These jurisdictions, with their pluralistic societies and histories of balancing individual rights against state paternalism, offer insights into protecting ontological sovereignty amid technological encroachment. While AI regulation is nascent, these frameworks highlight the inadequacy of current paradigms, which often prioritize empirical efficacy over phenomenological integrity. Foundational scholarship on epistemic injustice in healthcare underscores this point, showing how institutional biases can marginalize non-dominant knowledge systems in medical decision-making (Kidd & Carel, 2017). Extending this to AI, comparative analysis reveals pathways for inverting burdens of proof, compelling regulators to demonstrate non-interference with spiritual narratives rather than assuming algorithmic neutrality.

This approach is not merely academic; it's pragmatic. In an algorithmic age, where AI systems risk epistemic violence by commodifying beliefs, constitutional jurisprudence provides tools for resistance. For instance, analyses of health inequalities through lenses of structural injustice emphasize the need for rights-based protections that accommodate communal and spiritual dimensions (Byskov, 2021). Yet, as recent work on AI ethics notes, global regulations like the EU AI Act fall short in addressing these cultural variances, often imposing a uniform risk calculus that overlooks local ontologies (Heuser, Steil & Salloch, 2025). By weaving in these doctrines, we argue for a hermeneutic shift in health law, one that reconceives autonomy as inviolable guardianship over one's worldview.

5.1 Indian Essential Religious Practices Doctrine

India's constitutional jurisprudence, rooted in Articles 25 and 26 of the Constitution, employs the "essential religious practices" test to safeguard faith-based autonomy from state overreach. Articulated in landmark cases like *Shayara Bano v. Union of India* (2017), this doctrine requires courts to determine whether a practice is integral to a religion before permitting regulation, thereby protecting spiritual self-determination in diverse contexts. Applied to healthcare, it has implications for refusing AI-mediated interventions that clash with sacred beliefs: such as algorithmic predictions overriding astrological or Ayurvedic consultations in end-of-life decisions.

This test counters the epistemic hegemony of biomedicine by demanding evidence that regulation serves a compelling public interest without eroding core ontologies. Scholarly critiques highlight its relevance to emerging tech: in pluralistic India, where over 70% integrate traditional medicine, the doctrine could mandate "ontological exemptions" for AI tools, ensuring they do not commodify beliefs as data exhaust (Khalikova, 2021). Older analyses of epistemic injustice in Indian healthcare reinforce this, showing how colonial legacies persist in marginalizing indigenous knowledge, a dynamic AI risks amplifying without doctrinal safeguards (Mladenov, & Dimitrova, 2023). Thus, extending the test to AI regulation could invert the paternalistic burden, requiring proponents to prove non-vitiation of spiritual practices. In practice, this might manifest in cases where AI chatbots dismiss faith-healing as irrational, clashing with protected rituals. Comparative studies suggest India's approach offers a model for global polities, balancing innovation with metaphysical diversity (Patil et al., 2024).

5.2 European Margin of Appreciation Doctrine

The European Court of Human Rights (ECtHR) employs the "margin of appreciation" doctrine to

grant states flexibility in interpreting Convention rights, particularly under Article 9 on freedom of thought, conscience, and religion. In *S.A.S. v. France* (2014), the Court upheld a burqa ban by deferring to national cultural assessments, illustrating how the doctrine accommodates pluralism while scrutinizing necessity and proportionality. Transposed to AI in healthcare, it could allow member states leeway in regulating algorithmic intrusions on spiritual autonomy, such as mandatory AI assessments overriding religious refusals of treatment.

This deference is crucial in pluralistic Europe, where migrant communities blend faith-based healing with public systems. Ethical frameworks warn that without such margins, AI risks epistemic violence by enforcing secular norms, as seen in biases against spiritual distress in mental health algorithms (Tunks Leach, Simpson, Lewis, et al., 2023). Earlier philosophical work on epistemic injustice in psychiatric practice aligns here, arguing for interpretive flexibility to avoid hermeneutical marginalization (Crichton et al., 2017). The EU AI Act's risk classifications, while progressive, could incorporate this doctrine to tailor assessments, demanding proof that high-risk AI does not disproportionately burden faith-infused ontologies (European Parliament, 2024).

However, critics note the doctrine's potential for inconsistency, yet in healthcare AI, it promotes proportionality, balancing clinical benefits against phenomenological harms. This resonates with calls for life-world perspectives in AI ethics, ensuring regulations honor diverse narratives (Heuser, Steil & Salloch, 2025).

5.3 African Communal Dignity Jurisprudence

Under the African Charter on Human and Peoples' Rights (1981), jurisprudence emphasizes communal dignity and collective rights, as in *Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania* (2024), where courts upheld indigenous claims against state impositions. This ethos

prioritizes ubuntu, interconnected humanity, over individualistic models, offering a counterpoint to Western biomedicine's atomized view of health. In AI contexts, it could mandate communal consultations before deploying algorithms, protecting spiritual autonomy in traditions where healing involves ancestral or group rituals.

African scholarship on epistemic injustice critiques how global health tech perpetuates colonial erasures, advocating for dignity-based frameworks that center local knowledge (Kay, Kasirzadeh & Mohamed, 2024). Foundational studies link this to healthcare disparities, where epistemic violence silences communal voices in policy (Carel & Kidd, 2014/2017). For Rwanda's Protestant University-hosted conference, this jurisprudence is particularly resonant, aligning with regional efforts to integrate AI without effacing metaphysical diversity (World Health Organization, 2023).

6.0 The Spiritual Harm Threshold: A Novel Juridical Metric

In response to the epistemic violence and regulatory voids laid bare by AI's encroachment on sacred healing, this paper advances the "spiritual harm threshold" as a pioneering juridical metric. This threshold reimagines health law's guardrails, mandating that regulators and AI developers demonstrate not only empirical efficacy, such as improved diagnostic accuracy or cost savings, but also phenomenological non-interference with the patient's cosmogonic narrative. At its essence, the metric acknowledges that harm in pluralistic healthcare extends beyond physical or psychological injury to include existential disruptions: the fracturing of one's worldview, where algorithmic outputs clash with faith-rooted interpretations of affliction and restoration. Drawing from epistemic justice frameworks, which emphasize rectifying wrongs against marginalized knowers, this threshold inverts traditional burdens of proof, compelling proponents to negate spiritual displacement rather than merely affirm clinical utility (Adams, 2025). It positions spiritual autonomy as a protected

interest, akin to dignity in human rights jurisprudence, ensuring that AI serves without supplanting transcendent epistemologies.

Conceptualizing this threshold requires integrating insights from AI ethics and health law. Recent analyses of algorithmic bias in healthcare reveal practical, epistemic, and normative challenges, where systems amplify disparities by overlooking cultural dimensions, much like how secular AI might dismiss spiritual distress as non-actionable (Doshi et al., 2024). The spiritual harm threshold addresses this by establishing a benchmark: any AI intervention must undergo scrutiny to prove it does not erode the patient's narrative cosmos, such as by quantifying "harm" through qualitative indicators like self-reported ontological disruption or community consultations. This draws parallels to epistemic harms in generative AI, where users face eroded clarity in their knowledge, extended here to spiritual realms where beliefs risk commodification into data exhaust (Kay, Kasirzadeh & Mohamed, 2024). For instance, in mental health AI, chatbots that violate ethical standards by pathologizing faith experiences could breach this threshold, triggering mandatory revisions (New study: AI chatbots systematically violate mental health ethics..., 2025).

Operationalizing the threshold hinges on "ontological impact assessments" (OIA), structured evaluations that mirror environmental impact statements but focus on existential effects. These assessments would require interdisciplinary panels, including ethicists, anthropologists, and faith representatives, to evaluate AI's potential to reify probabilistic ontologies or inflict calculative violence. Legislation could embed OIA into approval processes, as seen in emerging state-level AI regulations for mental health, which emphasize safeguards against bias but could expand to spiritual protections (Governing AI in Mental Health: 50-State Legislative Review, 2025; New Illinois law looks to put guardrails on AI in mental health..., 2025). By inverting the onus, OIA shift from paternalistic beneficence to justice-oriented accountability, demanding evidence that AI

preserves phenomenological integrity, perhaps through metrics like patient narrative coherence scores or cultural compatibility audits.

This novelty lies in its synthesis of comparative jurisprudence with AI governance. Echoing India's essential religious practices doctrine, the threshold could deem spiritual non-interference an "essential" right, shielding it from arbitrary algorithmic overrides (Shayara Bano v. Union of India, 2017). Europe's margin of appreciation might afford contextual flexibility in assessments, while African communal dignity jurisprudence ensures collective epistemologies are consulted, countering individualist biases in AI (S.A.S. v. France, 2014; Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania, 2024). Bridging these with meaningful human control in medical AI, the metric fosters empowerment, ensuring justice extends to spiritual realms (Bridging Justice and Meaningful Human Control in Medical AI..., n.d.). Critiques of current regulations underscore the urgency: without addressing bias and data issues preemptively, AI risks epistemicide in healthcare (Regulating medical AI before midnight strikes..., 2025).

Ultimately, the spiritual harm threshold is more than a metric, it is a normative pivot, reconceiving autonomy as ontological sovereignty. In pluralistic polities, it safeguards metaphysical variety, demanding that precision medicine yield to sacred narratives when thresholds are crossed. The following section explores its operationalization in detail.

7.0 Operationalizing Ontological Impact Assessments

To translate the spiritual harm threshold from theory into actionable policy, ontological impact assessments (OIAAs) emerge as the operational backbone, a rigorous, multidisciplinary process designed to evaluate AI's existential footprint on diverse healing epistemologies. Unlike traditional impact assessments focused on privacy or bias, OIAAs prioritize phenomenological integrity,

scrutinizing how algorithmic interventions might disrupt a patient's cosmogonic narrative or communal worldview. This operationalization draws from evolving AI ethics frameworks, where ontologies serve as semantic bridges between technology and human-centered domains, ensuring structured knowledge integration that respects pluralism (Ambalavanan, Snead, Marczika, Towett, Malioukis & Mbogori-Kairichi, 2025). In healthcare, OIAAs would mandate pre-deployment evaluations, compelling developers to map potential ontological clashes, such as when AI diagnostics override indigenous etiologies of illness as spiritual disharmony rather than biological malfunction.

The process begins with a scoping phase, identifying stakeholders: ethicists, anthropologists, faith leaders, and patient advocates alongside technologists. This mirrors UNESCO-inspired ontologies for ethical AI impact assessments, which extract and structure global guidelines to mitigate harms in diverse contexts (Chaudhary, 2022). Next, a mapping exercise delineates the AI's ontological assumptions, e.g., its probabilistic modeling of health outcomes, against user epistemologies, using tools like knowledge graphs to visualize interrelationships between predictive analytics and cultural narratives (Safranek & Zvackova, 2025). Qualitative metrics, such as narrative coherence surveys or hermeneutic audits, quantify "harm" by assessing disruption levels, inverting the evidentiary burden to require proof of non-interference.

Integration with existing regulations amplifies feasibility. The EU AI Act's high-risk classifications could embed OIAAs as mandatory addendums, extending beyond technical audits to phenomenological reviews, ensuring AI in mental health respects spiritual care without pathologizing faith experiences (European Parliament, 2024). In African contexts, aligning with communal dignity jurisprudence, OIAAs might incorporate ubuntu principles, consulting communities to evaluate collective impacts, as seen in taxonomies of AI risks that emphasize sociopolitical harms in health

domains (Golpayegani, Hovsha, Rossmaier, Saniee & Mišić, 2022). For example, deploying an AI chatbot in Rwandan clinics would require assessing its secular datasets against local animist ontologies, potentially mandating adaptations like faith-sensitive prompts.

Challenges abound: operationalizing demands interdisciplinary training and resources, risking bureaucratic delays in low-income settings. Yet, benefits outweigh these, fostering epistemic justice by recentering indigenous knowledge, as advocated in planetary health frameworks (Redvers, Lokugamage, Barreto, Bajracharya & Harris, 2024). The ontological kaleidoscope framework offers a methodological parallel, examining embodiment entanglements to prevent data artefact reductions of the body, adaptable to spiritual dimensions (Smith-Nunes, 2025). By mandating post-deployment monitoring, OIAs evolve dynamically, addressing ethical evolutions in machine learning that highlight biases and fairness (Barbierato *et al.*, 2025). In essence, OIAs operationalize a radical shift: from utilitarian aggregation to hermeneutic protection, compelling AI to honor incommensurable epistemologies. Tied to comparative doctrines, like India's essential practices test, they fortify spiritual sovereignty, ensuring pluralistic healthcare resists algorithmic homogenization (*Shayara Bano v. Union of India*, 2017).

8.0 Normative Claim: Statutory Entitlement to Algorithmic Abstention

The argumentation of this paper reaches its apex with a bold normative assertion: pluralistic healthcare systems must enshrine a statutory entitlement to "algorithmic abstention" in matters of sacral therapeutics, lest they precipitate the erosion of metaphysical diversity beneath the facade of technological progress. This right would empower patients to opt out of AI-mediated interventions when they impinge on spiritual autonomy, framing refusal not as obstinacy but as a safeguard for ontological sovereignty. In an era where AI permeates diagnostics, treatment

planning, and even palliative care, such a entitlement counters the paternalistic tilt of evidence-based governance, which often presumes algorithmic superiority without reckoning with the phenomenological costs to faith-rooted healing (Corfmat *et al.*, 2025). Without this legal bulwark, patients navigating indigenous or esoteric practices risk coerced assimilation into machine rationality, where sacred epistemologies are demoted to optional add-ons rather than inviolable cores.

This claim is not mere idealism; it stems from ethical imperatives in AI healthcare law, where the right to refuse or opt out emerges as a critical protection against epistemic harms. Recent scholarship underscores that patients should have the ability to reject AI involvement, particularly when systems lack transparency or amplify biases that dismiss spiritual narratives as outliers (Hurley *et al.*, 2025). For instance, in mental health contexts, where AI chatbots might reinterpret spiritual crises through secular lenses, abstention ensures individuals retain control over their cosmogonic stories, aligning with broader calls for meaningful human oversight in automated decisions (Cheng, 2024). Ethically, this entitlement echoes principles of justice and non-maleficence, preventing the subtle violence of commodifying beliefs into data exhaust or reifying fluid ontologies into actuarial certainties (Kay, Kasirzadeh & Mohamed, 2024). Absent such a right, vulnerable communities, such as those in sub-Saharan Africa blending ancestral rituals with clinical care, face existential displacement, as AI's utilitarian calculus overrides transcendent priorities without recourse (Birhane, 2025).

Legally, operationalizing this entitlement draws sustenance from comparative jurisprudence, adapting doctrines to the digital age. India's essential religious practices test could extend to deem algorithmic abstention a protected facet of spiritual self-determination, shielding it from state-mandated AI integration (*Shayara Bano v. Union of India*, 2017). Europe's margin of appreciation might afford contextual deference, allowing patients to abstain based on cultural variances in

human rights adjudication (S.A.S. v. France, 2014). Meanwhile, African communal dignity jurisprudence, emphasizing collective worldviews, supports group-level opt-outs, ensuring AI does not fracture communal healing bonds (Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania, 2024). Emerging regulations, like those in the EU AI Act, already hint at high-risk categorizations for health AI, which could incorporate abstention clauses tied to ontological impact assessments (European Parliament, 2024). In the U.S., state-level initiatives on AI in healthcare utilization management suggest pathways for statutory mandates, requiring notices of AI use and opt-out mechanisms to mitigate biases (Holland & Knight, 2024).

Yet, this claim provokes counterarguments: critics might argue that abstention could compromise clinical outcomes or strain resources in overburdened systems. Rebuttals draw from proportionality principles in human rights law, abstention need not be absolute but calibrated, perhaps limited to non-emergent sacral contexts where alternatives exist, ensuring beneficence without paternalism (Fasan, 2025). Moreover, empirical evidence from patient engagement studies reveals hesitations about AI precisely because of fears over lost agency, reinforcing the need for statutory protections to foster responsible innovation rather than resistance (Lysen & Wyatt, 2024).

In sum, statutory algorithmic abstention is an exigency for preserving the hermeneutic essence of sacred healing. It challenges health law to evolve, reconceiving autonomy beyond volitional consent to encompass unyielding guardianship over one's metaphysical realm, a reconfiguration essential in an algorithmic epoch.

9.0 Reconceptualizing Autonomy as Ontological Sovereignty

The conventional framing of autonomy in health law: as volitional consent, where patients merely assent or refuse interventions, falls perilously short

in an algorithmic era, where AI systems subtly reshape the very fabric of one's worldview. This paper advocates a reconceptualization: autonomy as ontological sovereignty, the unassailable right to govern one's cosmogonic narrative without external dilution or commodification. No longer a procedural checkbox, this sovereignty demands recognition of the patient's epistemic agency, where healing ontologies, whether indigenous animism, faith-based redemption, or esoteric energetic, hold primacy over machine-derived probabilities. In pluralistic healthcare, where AI risks epistemic violence by flattening sacred epistemologies into data points, this shift confounds utilitarian aggregation, insisting that aggregate clinical gains cannot trump individual metaphysical integrity (Kaebnick, 2016). It echoes relational turns in bioethics, expanding autonomy beyond individualism to encompass interdependent worldviews, particularly vital when decolonizing AI ethics to counter harms against marginalized knowledges (Tiribelli, 2023).

This exigency arises from AI's ontological friction: tools trained on biomedical datasets impose a singular reality, pathologizing spiritual distress or reifying transcendent uncertainties as risks to mitigate. Consider a Pentecostal patient whose ecstatic visions signal divine healing; an AI chatbot, grounded in secular psychology, might label them delusional, overriding consent with algorithmic "beneficence" and eroding sovereignty over one's narrative cosmos (Tunks Leach, Simpson, Lewis, *et al.*, 2023). Such intrusions demand a reconfiguration of regulatory reason, where autonomy transcends choice to embody guardianship against existential displacement. Decolonial critiques reinforce this, positing relational autonomy as a bulwark against AI harms, reconceptualizing it to honor collective and cultural ontologies rather than individualistic defaults (Kwek, 2023). In health recommender systems, for instance, autonomy requires redesign to preserve active ageing narratives, not subsume them under predictive models (Tiribelli, 2023).

Jurisprudentially, this reconceptualization aligns with doctrines safeguarding spiritual pluralism. India's essential religious practices test protects ontological cores from state interference, suggesting sovereignty as a constitutional shield against AI paternalism (*Shayara Bano v. Union of India*, 2017). Europe's margin of appreciation affords deference to diverse life-worlds, while African communal dignity jurisprudence elevates collective sovereignty, countering individualistic AI biases (S.A.S. v. France, 2014; *Inclusive Development for Citizens and Another v. Attorney General of the United Republic of Tanzania*, 2024). Yet, extant frameworks like the EU AI Act prioritize procedural safeguards, overlooking how AI undermines autonomy by eroding interpretive agency in hermeneutic voids (European Parliament, 2024). Philosophical analyses urge this pivot: autonomy in AI medicine must respect epistemic dimensions, ensuring systems enhance rather than supplant patient knowership (Adams, 2025).

Critically, ontological sovereignty confounds utilitarian paradigms by insisting on incommensurability, sacred harms cannot be aggregated or traded against empirical benefits. In end-of-life AI, for example, algorithms optimizing resource allocation might dismiss faith refusals as irrational, but sovereignty demands their inviolability, fostering justice-oriented governance (Dovey & Shuman, 2024). This radical reconfiguration invites health law scholars to embrace hermeneutic renewal, where autonomy safeguards the irreducibly interpretive nature of healing against silicon's reductive gaze.

10.0 Conclusion

As artificial intelligence permeates the sanctuaries of healing, this paper has illuminated the regulatory voids that threaten spiritual autonomy in pluralistic healthcare landscapes. From the ontological pluralism underpinning diverse epistemologies to the epistemic violence wrought by algorithmic reification, commodification, and calculative harms, the analysis reveals how evidence-based

governance systematically effaces sacred narratives under the banner of precision. Comparative jurisprudence, spanning India's essential practices, Europe's margin of appreciation, and Africa's communal dignity, exposes the paternalistic shortcomings of extant frameworks, while novel tools like the spiritual harm threshold and ontological impact assessments offer pathways to redress. The normative imperative for algorithmic abstention, coupled with reconceptualizing autonomy as ontological sovereignty, underscores an urgent demand: health law must evolve beyond utilitarian metrics to embrace hermeneutic depth, honoring the interpretive essence of transcendent care.

This hermeneutic renewal is no luxury but a necessity in polities where faith and tech intersect. Absent it, AI risks the quiet extinction of metaphysical variety, transmuting soulful restoration into actuarial shadows. By foregrounding incommensurability, the irreducible chasm between silicon logic and sacred knowing, scholars and regulators are called to reforge governance, ensuring AI serves without supplanting. In Rwanda's clinics or Mumbai's wards, where chants entwine with code, such a law promises equity: not homogenized progress, but a mosaic of ontologies thriving amid innovation. Ultimately, toward a hermeneutic health law lies the preservation of humanity's deepest diversities, a reconfiguration that confounds aggregation and affirms sovereignty in an algorithmic age.

11.0 Recommendations

Based on the findings of the paper, the following recommendations are made:

1. There is the need to integrate OIAs into national AI regulations, mandating pre-deployment reviews by interdisciplinary panels to evaluate phenomenological disruptions. Draw from the EU AI Act's risk assessments, expanding them to include spiritual metrics, with mandatory

community consultations in pluralistic regions.

2. There is the need to advocate for laws granting patients explicit opt-out entitlements in sacral therapeutics, modeled on human rights doctrines. This could include draft clauses requiring AI notices and alternatives, tailored to cultural contexts via comparative jurisprudence.
3. Develop certification programs blending health law, anthropology, and AI ethics, equipping overseers to identify epistemic violence. Collaborate with institutions like PUR to pilot trainings focused on African communal dignity, ensuring global applicability.
4. Require developers to incorporate diverse ontologies in training data, with audits for bias against faith-based narratives. Leverage anthropological insights to create "ontological repositories" for balanced models, mitigating reification harms.
5. Convene forums like ICFAI extensions to draft international guidelines, emphasizing hermeneutic health law. Involve Vatican-inspired ethics to bridge faith and tech, fostering normative shifts toward sovereignty.
6. Establish independent bodies for ongoing surveillance of AI in healthcare, using the spiritual harm threshold to track existential displacements. Publish annual reports with case studies, informing iterative reforms.

Bibliography

Birhane, A. (2025). AI-driven health messaging and epistemic justice in the Global South. In M.

Coeckelbergh (Ed.), *Oxford handbook of artificial intelligence ethics* (pp. 456–478). Oxford University Press.

Smith-Nunes, G. (2025). Body as a data artefact: The ontological kaleidoscope framework. *Journal of Responsible Technology*, 24, 100138. <https://doi.org/10.1016/j.jrt.2025.100138>; <https://www.sciencedirect.com/science/article/pii/S266659625000344>

Silcox, C., Zimlichmann, E., Huber, K. et al., The potential for artificial intelligence to transform healthcare: perspectives from international health leaders. *npj Digit. Med.* 7, 88 (2024). <https://doi.org/10.1038/s41746-024-01097-6>

Barbierato et al., Breaking Away From AI: The Ontological and Ethical Evolution of Machine Learning. (2025). *Journal of Artificial Intelligence and Consciousness*. https://www.researchgate.net/publication/390022186_Breaking_Away_from_AI_The_Ontological_and_Ethical_Evolution_of_Machine_Learning

Carel, H., & Kidd, I. J. (2017). Epistemic injustice in healthcare: A philosophical analysis. *Medicine, Health Care and Philosophy*, 17(4), 529–540. <https://doi.org/10.1007/s11019-014-9560-2> (Original work published 2014)

Cheng, M. (2024). The right to human decision: Analyzing policies, ethics, and implementation. In *Addressing socioethical effects of artificial intelligence*. OpenReview. <https://openreview.net/forum?id=VrxtvBw0ly>

Corfmat, M., Martineau, J. T., & Régis, C. (2025). High-reward, high-risk technologies? An ethical and legal account of AI development in healthcare. *BMC Medical Ethics*, 26(1), Article 158. <https://doi.org/10.1186/s12910-024-01158-1>

Crichton, P., Carel, H., & Kidd, I. J. (2017). Epistemic injustice in psychiatry. *BJPsych*

Bulletin, 41(2), 65–70.
<https://doi.org/10.1192/pb.bp.115.050682>

Descola, P. (2021). *Beyond nature and culture* (J. Lloyd, Trans.; reissue ed.). University of Chicago Press. (Original work published 2013)

Doshi, R., Falco, G., & Krumholz, H. M. (2024). Practical, epistemic and normative implications of algorithmic bias in healthcare: A qualitative analysis of clinician perspectives. *Journal of Medical Ethics*, 51(6), 420–425.

Dovey, S., & Shuman, A. (2024). AI in digital health: Autonomy, governance, and privacy. In *AI in eHealth: Human Autonomy, Data Governance and Privacy in Healthcare* (pp. 1–20). Springer. (Adapted from 2022 volume).

European Parliament. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union, L 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

Chaudhary, U. (2022). *External impact assessment on ethics of AI using ontology based on UNESCO recommendations* [Master's thesis, Trinity College Dublin]. TCD Research, pp. 1-50. Repository.
<https://www.scss.tcd.ie/publications/theses/diss/2022/TCD-SCSS-DISSERTATION-2022-066.pdf>

Fasan, M. (2025). Regulating the use of artificial intelligence in the doctor-patient relationship? A primer on supranational and national legal frameworks. *Biolaw Journal*, 1(1), 45–62.
<https://iris.unitn.it/handle/11572/453770>

Tunks Leach, K., Simpson, P., Lewis, J. et al. The Role and Value of Chaplains in an Australian Ambulance Service: A Comparative Study of Chaplain and Paramedic Perspectives. *J Relig Health* 62, 98–116 (2023).

<https://doi.org/10.1007/s10943-022-01685-4>
<https://doi.org/10.1007/s10943-022-01685-4>

Holland & Knight Alert. (2024). Regulation of AI in healthcare utilization management and prior authorization. Holland & Knight LLP.
<https://www.hklaw.com/en/insights/publications/2024/10/regulation-of-ai-in-healthcare-utilization-management>

Hurley, M. E., Lang, B. H., & Kostick-Quenet, K. M. (2025). Patient consent and the right to notice and explanation of AI systems used in health care. *American Journal of Bioethics*, 25(2), 78–89.
<https://doi.org/10.1080/15265161.2024.2399828>

Inclusive Development for Citizens and Another v Attorney General of the United Republic of Tanzania (Reference No.10 of 2020) [2024] EACJ 14 (29 November 2024) (First Instance Division)
<https://ulii.org/en/akn/aa/judgment/eacj/2024/14/eng@2024-11-29>

Kaebnick, G. E. (2016). The psychology of autonomy. *Hastings Center Report*, 46(3), 2.
<https://doi.org/10.1002/hast.581>

Adams, J. Ethical and epistemic implications of artificial intelligence in medicine: a stakeholder-based assessment. *AI & Soc* (2025).
<https://doi.org/10.1007/s00146-025-02398-4>

Khalikova, V. (2021). Medical pluralism. In F. Stein (Ed.), *The open encyclopedia of anthropology*. Cambridge Encyclopedia of Anthropology.
<https://doi.org/10.29164/21medplural>

Kidd, I. J., & Carel, H. (2017). Epistemic injustice and illness. *Journal of Applied Philosophy*, 34(2), 172–190. <https://doi.org/10.1111/japp.12172>

Kwek, D. (2023). Decolonizing AI ethics: Relational autonomy as a means to counter AI harms. *Topoi*, 42(3), 841–850.

Lee, H.-S., Song, S.-H., Park, C., Seo, J., Kim, W. H., Kim, J., Kim, S., & Lee, Y. H. (2025). The ethics of simplification: Balancing patient autonomy, comprehension, and accuracy in AI-generated radiology reports. *BMC Medical Ethics*, 26, Article 136. <https://doi.org/10.1186/s12910-025-01285-3>

Lysen, F., & Wyatt, S. (2024). Refusing participation: Hesitations about designing responsible patient engagement with artificial intelligence in healthcare. *Journal of Responsible Innovation*, 11(1), Article 2300161. <https://doi.org/10.1080/23299460.2023.2300161>

Young, M., & Varpio, L. (2025). Navigating the research landscape: How paradigms shape health professions education research. *Medical Education*, 59(3), 256–267. <https://doi.org/10.1111/medu.15752>

Jonas, L., Bacharach, S., Nightingale, S., & Filoche, S. (2024). Under the umbrella of epistemic injustice communication and epistemic injustice in clinical encounters: A critical scoping review. *Ethics, Medicine and Public Health*, 33, 101039. <https://doi.org/10.1016/j.jemep.2024.101039>

Mladenov, T., & Dimitrova, I. (2023). Epistemic injustice as a bridge between medical sociology and disability studies. *Sociology of health & illness*, 45(6), 1146–1163. <https://doi.org/10.1111/1467-9566.13479>

Palaniappan, K., Ting Lin, E. Y., Vogel, S., & W Lim, J. C. (2024). Gaps in the Global Regulatory Frameworks for the Use of Artificial Intelligence (AI) in the Healthcare Services Sector and Key Recommendations. *Healthcare*, 12(17), 1730. <https://doi.org/10.3390/healthcare12171730>

Emah, I., & Bennett, S. (2025). Algorithmic emergence? Epistemic in/justice in AI-directed transformations of healthcare. *Frontiers in Sociology*, 10, 1520810. <https://doi.org/10.3389/fsoc.2025.1520810>

Nazer, L. H., Zatarah, R., Waldrip, S., Ke, J. X. R., Moukheiber, M., Khanna, A. K., Hicklen, R. S., Moukheiber, L., Moukheiber, D., Ma, H., & Mathur, P. (2024). Bias in artificial intelligence algorithms and recommendations for mitigation. *PLOS Digital Health*, 3(6), Article e0000278. <https://doi.org/10.1371/journal.pdig.0000278>

Safranek, S., & Zvackova, A. (2025). Ontological framework for integrating predictive analytics, AI, and big data in decision-making systems using knowledge graph. In *Proceedings of the 15th International Conference on Simulation and Modeling Methodologies, Technologies and Applications (SIMULTECH 2025)* (pp. 287–294). <https://www.scitepress.org/Papers/2025/135144/135144.pdf>

Ambalavanan, R., Snead, R. S., Marczika, J., Towett, G., Malioukis, A., & Mbogori-Kairichi, M. (2025). Ontologies as the semantic bridge between artificial intelligence and healthcare. *Frontiers in Digital Health*, 7, 1668385. <https://doi.org/10.3389/fdgth.2025.1668385>

Patil, A. D., Singh, S., Verma, D., & Goupale, C. (2024). Exploring medical pluralism as a multifaceted approach to healthcare. *Indian Journal of Integrative Medicine*, 4(2), 49–59. <https://mansapublishers.com/ijim/article/view/4322>

Khurana, S. (2025). Decolonizing artificial intelligence: Indigenous knowledge systems, epistemic pluralism, and the ethics of technology. *Journal of Computer Allied Intelligence*, 3(3), 1–10. <https://fringe-global.com/ojs/index.php/jcai/article/view/decolonizing-artificial-intelligence-indigenous-knowledge-system/167>

Pontifical Academy for Life. (2020). *Rome Call for AI Ethics*. Vatican Press.

https://www.vatican.va/roman_curia/pontifical_academies/acdlife/documents/rc_pont-acd_life_doc_20202228_rome-call-for-ai-ethics_en.pdf (Reaffirmed with new signatories in 2025)

Rahsepar Meadi, M., Sillekens, T., Metselaar, S., van Balkom, A., Bernstein, J., & Batelaan, N. (2025). Exploring the ethical challenges of conversational AI in mental health care: Scoping review. *JMIR Mental Health*, 12, Article e60432. <https://doi.org/10.2196/60432>

S.A.S. v. France, App. No. 43835/11, ECHR 2014 (Eur. Ct. H.R.).

<https://hudoc.echr.coe.int/eng?i=001-145466>

Kay, J., Kasirzadeh, A., & Mohamed, S. (2024). Epistemic Injustice in Generative AI. *ArXiv*. <https://arxiv.org/abs/2408.11441>

Redvers N., Lokugamage A. U., Barreto J. P. L., Bajracharya M. B., Harris M. (2024) Epistemicide, health systems, and planetary health: Re-centering Indigenous knowledge systems. *PLOS Glob Public Health* 4(8): e0003634.

<https://doi.org/10.1371/journal.pgph.0003634>

Heuser, S., Steil, J. & Salloch, S. AI Ethics beyond Principles: Strengthening the Life-world Perspective. *Sci Eng Ethics* 31, 7 (2025). <https://doi.org/10.1007/s11948-025-00530-7>

Shayara Bano v. Union of India, (2017) 9 SCC 1 (India). <https://indiankanoon.org/doc/115701246/>

Tiribelli, S. (2023). The AI ethics principle of autonomy in health recommender systems. *Argumenta*, 8(1), 75–88. <https://upad.unimc.it/handle/11393/306693>

Tobert, N. (2022, July 22). *Pluralistic practice: A medical anthropology perspective*. Pluralistic Practice. <https://pluralisticpractice.com/main->

[blog/epistemologypluralisticpractice-a-medical-anthropology-perspective/](https://blog.epistemologypluralisticpractice.com/main-anthropology-perspective/)

Golpayegani, D., Hovsha, J., Rossmaier, L. W. S., Saniei, R., & Mišić, J. (2022). Towards a taxonomy of AI risks in the health domain. In *2022 Fourth International Conference on Transdisciplinary AI (TransAI)* (pp. 1–8). IEEE. <https://philarchive.org/archive/GOLTAT-21>

Viveiros de Castro, E. (2021). *Cosmological perspectivism in Amazonia and elsewhere*. HAU Books. <https://haubooks.org/cosmological-perspectivism-in-amazonia/>

World Health Organization. (2023). *WHO global report on traditional and complementary medicine 2023*. WHO Press. https://www.who.int/health-topics/traditional-complementary-and-integrative-medicine#tab=tab_1