Vantage: Journal of Thematic Analysis
DOI: https://doi.org/10.68903/vjta.2026.v07i02.02
Vol. 7(2), October, 2026, 121-132
ISSN(E): 2582-7391

Original Research Article

Who Controls, Who Suffers and Who Counts? Rethinking Accountability in AI-Driven One Health

Anshu Kumar,
Department of Philosophy, University of Delhi, Faculty of Arts, University Enclave, Delhi, 110007, India.
Corresponding Author: anshu.philosophy@gmail.com

Received: 06/06/2026 | Accepted: 20/09/2026 | Published: 30/10/2026

ABSTRACT

One Health AI systems have already been implemented rather than being conceptualised. These systems predict disease outbreak, recognise patterns of antibiotic resistance and monitor wild animal populations. How these systems will be governed needs to be evaluated. Current governance manifests a failure where control over design is in the hands of developers removed from frontline work and oversight is in the hands of policymakers without the capacity to determine the adequacy of the design. Those scientists who know the system's failings the best have only an advisory position. In fact, with the current governance frameworks, the absence of control is in effect harmful. This second failure is even less addressed. One Health's primary assertion is that human, animal and environmental well-being are connected. The EU AI Act (2024) and the UNESCO Recommendation (2021) extend moral standing only to humans. The frameworks do not recognise harm to African elephants, nor effects on the wildlife whose pathogen dynamics are being surveyed. The Layered Relational Accountability model offered here gives obligations a more appropriate distribution based on actual ability, expands moral concern to the three areas (human, animal, environment) and provides a framework for active accountability.

Keywords: One Health; AI accountability; Control-impact mismatch; Cross-species ethics; Moral responsibility

1. INTRODUCTION

The scientific and ethical foundations of One Health have been formally recognised by major global institutions. In essence, One Health asserts that human, animal and ecosystem health are inextricably linked. This creates a moral imperative to improve the health of all three and none can be neglected. The application of One Health relies heavily on the deployment of Artificial Intelligence (AI) systems. One Health practitioners use AI to predict when and where human populations are at risk of infection by transboundary animal diseases, to identify antimicrobial resistant (AMR) genes in wastewater, to monitor the movements of diverse wildlife populations globally and to determine where best to allocate scarce resources.

There has been rapid growth in scholarship in the field of AI ethics and such work has begun to address some of the challenges presented by these systems. The issue of accountability, especially as it relates to complex, multi-actor systems, has been extensively studied (Thompson, 1980, pp. 905-916; French, 1984, pp. 30-45; Bovens, 1998, pp. 45-60). The need for explainable AI and consequently, the need for transparency and traceability in AI systems, has received significant attention (Arrieta et al., 2020, pp. 82-115). More

© The Author(s) 2026. This work is licensed under a Creative Commons Attribution 4.0 International License which permits its use, distribution and reproduction in any medium, provided the original work is cited. https://vantagejournal.com

recently, the environmental aspects of AI ethics have begun to receive scrutiny (van Wynsberghe, 2021). However, two challenges specific to AI and One Health have been largely neglected and their confluence has created an unprecedented structural governance deficit.

The first concern is the disparity between the individuals and organisations that control AI systems and those who bear the consequences of their deployment. In One Health AI deployments, these groups are often entirely different. This disparity results from the global development of AI. This gap means that stakeholders who assume the most responsibility for developing AI systems are often shielded from the consequences, while those who experience the consequences have the least control over the design. Several accountability frameworks have been developed to address the relationships between control and consequence and are at least nominally aligned. However, these frameworks cannot address the disparity caused by the structures described above.

The second concern is a direct contradiction in the governance of One Health. One Health practitioners are committed to animal and environmental welfare for their own sake, rather than instrumentally for human health. However, most AI governance frameworks applicable to One Health, including the EU AI Act (2024) and the UNESCO Recommendation on the Ethics of AI (2021), consider AI systems primarily through the lens of their impact on humans. Animals and ecosystems are mentioned only as background. There is a lack of recognition that the welfare of animals and ecosystems imposes a responsibility on the governance of AI. This inconsistency is far from peripheral. It undermines the core values of the systems these paradigms are intended to serve.

Generally, AI ethics literature has not clearly defined either of these problems. Although the many hands problem (shared responsibility across multiple actors dissolving into no one's responsibility) has been well described, it is distinct from the control-consequence mismatch and the problems and solutions should not be conflated. In the limited scholarship on environmental AI, the anthropocentric focus of AI ethics has been briefly mentioned. However, it has not been framed as a contradiction of One Health's normative commitments, nor has a governance model been proposed to address this gap.

This paper seeks to address these gaps. It argues that both the redistribution of governance responsibilities and the genuine expansion of moral consideration regarding who and what matters in AI governance, are necessary to resolve these issues. The proposed Layered Relational Accountability (LRA) model is based on these dual requirements. It is grounded in well-established philosophical theory, which has yet to be applied to this governance context and is intended for implementation within existing international governance frameworks, rather than as a replacement.

2. THE CONTROL-CONSEQUENCE GAP: A STRUCTURAL PROBLEM DISTINCT FROM MANY HANDS

The many hands problem is the most well-known issue in assigning responsibility for the outcomes of complex AI systems. When many people contribute small actions that collectively cause harm, it is difficult to identify a single contribution significant enough to warrant full moral responsibility and responsibility for the outcome tends to disappear (Thompson, 1980, pp. 905-916). This is a real problem in AI-driven One Health. A single model could be built by one organisation, validated by a second, authorised by a third, modified by a fourth and run by a fifth.

Nonetheless, the many hands problem does not explain the full extent of the governance issue in this context. A specific and more differentiated structural issue exists in parallel that

has not received significant attention: a systematic structural absence of correspondence between those who control AI systems with real power and those affected by their real-world effects.

The absence of such correspondence is not itself a new observation. Beck's account of organised irresponsibility describes the condition in industrial modernity, and the environmental justice literature has documented the concentration of harm among populations with least influence over its production. Within AI ethics, the responsibility gap has been examined since Matthias (2004), and Santoni de Sio and Mecacci (2021) have distinguished four such gaps. What has not been examined is the configuration described here, in which those who control a system and those who bear its consequences are systematically different parties, and in which some of those affected have no standing within the governance framework at all. The many hands problem concerns the many contributions through which responsibility disappears. The problem identified here concerns the separation of control from consequence.

There are three facets to this disparity, each involving a different type of actor. At the first level are the developers, who possess the greatest degree of technical control. This control arises from their decisions on model architecture, selection of training sets and definition of performance targets. They set behavioural boundaries for deployed systems. However, developers are unlikely to spend time in the settings where their systems are ultimately used. They are often unaware of the constraints and challenges faced by users, the conditions of intended deployment and the differences between the populations where their systems are applied and those where the systems were originally developed. Their control is extensive, but their exposure to the impact is negligible.

Policymakers possess institutional control. They are responsible for authorising or prohibiting deployment, defining the terms and conditions of system use and determining the range of system applications. This power is central to any governance framework because large-scale deployment of a system is likely to cause significant and widespread harm to many people and to the integrity of many ecosystems. However, in most cases, policymakers are in the weakest position to evaluate whether a system has been designed to be fit for purpose in its intended context. To do this, they must rely on technical inputs from system developers and scientists. The quality of these inputs depends on the honest communication of concerns, which are likely to relate to system failure modes. However, this is unlikely to be a high priority given most existing incentive structures.

Reviewers and validators of AI models are familiar with the gaps, failures and challenges presented by the models. AI models are usually unable to represent and include all populations and contexts of the natural world. Therefore, working with AI carries significant risk, especially when models are deployed. Those who validate AI models often know the risks and these risks are often significant. However, within most institutions, the validation of AI models is advisory. This means validation may inform a decision but will not control or dictate it. Most scientists who validate AI models and do the most significant work to identify the models' challenges and failures have the least control and power.

This creates a governance framework in which power flows in one direction, while risk and harm occur in another. Those who make the most significant and foundational decisions are unlikely ever to experience the negative consequences of those decisions. Those who do experience the negative outcomes have no significant influence over the decisions that produced them. The negative impacts of AI models result from the current AI model development system, not from malicious or evil intent.

Nissenbaum (1994, pp. 72-80) examined an

early version of this issue in the context of computing and moral responsibility. When computer systems are positioned between the agent and the outcome, responsibilities such as foresight and control become difficult to maintain. Control of this kind is a standing condition of moral responsibility (Fischer & Ravizza, 1998, pp. 1-25), which is precisely what the structure of these systems erodes. The challenge presented in One Health AI, as identified in the literature, is even more complex, as the mediation is not only technological but also institutional, cross-organisational, cross-national and cross-species. Recent research into the attribution of causal responsibility in human-AI systems (Qi et al., 2024, pp. 1-14) offers guidance for the more limited domain of clinical decision support systems, but the challenges posed by systems that cause harm across multiple organisations and species and that affect animals and ecosystems absent from governance systems are more complex and cannot be addressed by that research.

The so-called accountability gap is not simply a problem of distributed responsibility. It is a specific kind of responsibility distribution problem in which those with the most power is remote from the impacts of that power and those most vulnerable to its effects (including non-human and ecologically vulnerable populations) are excluded from decision-making. To address this, the governance and responsibility framework must be expanded to include non-human and vulnerable populations.

3. THE ANTHROPOCENTRIC CONTRADICTION IN ONE HEALTH AI GOVERNANCE

One Health is founded on the principle that human, animal and environmental health are practically and integrally connected. This paradigm seeks to dismantle the hierarchy of concern and to promote the view that human health is not the sole priority within the system. This perspective fails to grasp the science of pathogen dynamics, resistance and the ecology of disease. It also fails to appreciate the importance of animal and environmental welfare (World Health Organization [WHO], Food and Agriculture Organization [FAO], World Organisation for Animal Health [WOAH], & United Nations Environment Programme [UNEP], 2021). Such a claim is not an afterthought in One Health; it is the paradigm's locus classicus and the source of its distinctive and intrinsic normative character.

It is therefore concerning that most major AI governance mechanisms applicable to One Health extend moral standing only to humans. According to the EU AI Act (2024), health AI systems are classified as high-risk and are therefore subject to certain standards regarding technical documentation and post-market monitoring. However, the assessment criteria are entirely human-centred. They evaluate systems for accuracy, equity and fairness across all human demographic groups, human user transparency and human rights. In this context, animals and other parts of the biosphere are not and cannot be recognised as moral subjects. Therefore, they do not and cannot have claims. They may exist as contextual variables, but only in relation to human health. In this respect, the UNESCO Recommendation on the Ethics of AI (2021) is more advanced than most, as it incorporates consideration of environmental sustainability. However, ecosystem health is presented as a matter of potential impact and thus falls short of recognising non-humans as moral subjects with a right to exist. What One Health requires of AI governance and what is currently provided are vastly different and largely unaddressed in the literature.

There are real-world impacts of conservation AIs that extend beyond the philosophical realm. These negative impacts result from conservation AI systems generating false reports, forcing game rangers to continually relocate and alter the social

structures of elephants. Such impacts can increase the mortality rate of elephant calves and disrupt their movement across the landscape. The harm pathway from repeated human intrusion is well documented: disturbance elevates glucocorticoid levels and depresses reproductive output (Gobush et al., 2008; Tang et al., 2020), and sustained disturbance degrades the social network structure on which calf survival depends (Wiśniewska et al., 2022). What has not been examined is the activation of that pathway by algorithmic false positives. Current governance and management systems completely ignore this disruption. These systems may record the conservation AI systems as performing adequately, since adequacy is based on human-centred measures, meaning the elephants' suffering from inadequate system performance is entirely overlooked.

Physiological stress in reservoir populations is a recognised risk factor for infection and shedding (Plowright et al., 2008), and the condition of the reservoir host is among the barriers that govern spillover (Plowright et al., 2017). These effects arise from changes in the behaviour and immune systems of the animal populations. A governance system that neglects animal welfare will also disregard the resulting negative impacts and feedback from these systems. Such systems will persist without corrective measures, even when severe negative impacts are documented.

A reconstructed account of the moral consideration of animals is unnecessary, as it has been effectively provided elsewhere. Singer (1975, pp. 1-20) argued that, from a utilitarian perspective, the capacity for suffering determines moral consideration, not species membership. If suffering is morally significant (and no major ethical position would deny this), then the suffering of sentient beings is significant. From a rights-based perspective, Regan (1983, pp. 243-256) maintained that individual animals have subjective experiences and possess inherent worth that cannot be justified solely by the greater good. Developing the capabilities approach first advanced by Sen (1999), Nussbaum (2006, pp. 325-340) argued that justice requires a long-term commitment to preserving the capabilities necessary for each species to flourish, including freedom from chronic pain, social contact and movement. Aldo Leopold (1949, pp. 201-226) asserted that living systems (organisms, soils, waters and the ecological processes that sustain life) form a community that must be of the greatest concern to any moral practice and that it is right to sustain the community and wrong to violate it.

These four positions do not entirely overlap, resulting in some constructive tensions. According to Singer's utilitarian framework, in certain circumstances, individual cases of animal suffering may be traded off for a greater overall good. Regan's rights-based position would not permit such trade-offs. Similarly, under Leopold's holism, individual animal suffering may be justified for the greater good of the entire system, such as maintaining ecosystem balance. Regan (1983, pp. 243-256) strongly criticised this approach and considered such trade-offs morally unacceptable. These tensions are genuine and should not be resolved; the more productive course is to make them explicit and govern them (Whittlestone et al., 2019, pp. 195-200). Even so, there is significant convergence among these positions, relevant to the present context, indicating that sentient animals and ecosystems are legitimate moral entities and that their welfare creates obligations for those whose actions may affect them. This convergence strongly suggests that the current exclusion of animals and ecosystems from One Health AI governance is not philosophically defensible. It is a choice embedded in governance policies and frameworks developed without considering whether the normative assertions of One Health should provide a fundamental structure for these systems. It is a choice based on moral and practical errors.

The Cross-Domain Moral Assessment (CDMA) discussed in this paper offers a relatively practical solution to the problem presented by this convergence. CDMA incorporates consideration for human welfare, animal welfare and environmental integrity within the decision-making framework of AI-enhanced One Health systems. Human flourishing, animal flourishing and biotic community flourishing yield different outcomes for the three groups and as a result, the three moral domains come with different sets of criteria. Human welfare is assessed through the study of human populations, outcomes and inequities. For animal welfare, the analysis is aided and informed by behavioural, physiological and neurological indicators (Sneddon et al., 2014, pp. 201-212). When evaluating environmental integrity, the assessment uses ecological indicators, the Biodiversity Intactness Assessment (Scholes & Biggs, 2005, pp. 45-49) and the monitoring indicators of the Kunming-Montreal Global Biodiversity Framework (Convention on Biological Diversity [CBD], 2022). CDMA requires that, in assessing the moral domains, the appropriate criteria and assessment frameworks are applied, findings are documented and made available in their unrefined form and any decision to allow harm to one moral domain in pursuit of benefit to another is made with full awareness and explicit justification, not merely through the model's objectives. In situations requiring reconciliation of animal and environmental harm, where the aims of the two domains diverge, CDMA does not supply a rule of priority. What it requires is that the divergence be identified as such, that the reasons for the resolution adopted be recorded and made available, and that the burden of justification rest with the party authorising the harm rather than with the domain sustaining it. The tensions set out above are not thereby resolved. They are made to operate as governance constraints, rather than being settled tacitly by the model's objective function or by institutional default.

4. THE LAYERED RELATIONAL ACCOUNTABILITY MODEL

This model addresses the two structural problems outlined in Sections 2 and 3. It is described in three layers. The first layer adjusts the obligations of the chain of actors according to their actual capacity, rather than theoretical capacity. The second layer incorporates Cross-Domain Moral Assessment (CDMA) into the evaluation of outcomes and governance decisions. The third layer assigns residual accountability at the institutional level, ensuring responsibility does not disappear within the many hands problem. The three layers work together: Layer 1 prevents the control–consequence gap from enabling unaccountable harm; Layer 2 ensures that non-human harms are recognised and have consequences; Layer 3 ensures that when such harm occurs, there is an identifiable subject responsible.

Layer 1 is based on the principle that responsibility follows capacity. Iris Marion Young's (2011) social connection model of responsibility states that those structurally positioned to prevent or reduce harm are responsible for acting, even if they are not the cause of the harm. In AI governance, the question is not who should be held responsible, but who has the capacity to prevent harm and what is required of them.

Developers who maintain technical control bear greater responsibility. Developers must account for likely harms in other domains before deploying the system. They must consider not only the system's performance in its primary task or under ideal conditions, but also the likely harms to animal populations, ecosystems and human communities under varying resource conditions. These assessments must identify gaps in training data that may be harmful to animals or ecosystems. They must also describe the positive and negative errors the system will produce and the likely impacts of those errors in the three moral domains. Additionally, they must outline

the risks of ecological feedback, where system control may alter the dynamics of the system itself. Non-specialist governance actors must be able to understand and directly contribute to the findings beyond documentation. The practical effects of documentation should not be undermined by requirements that are superficially easy to satisfy.

Policymakers must ensure that cross-domain impact assessments meet relevant metrics and that post-deployment monitoring of all three moral domains is initiated for considered uses before granting authorisation. Scientists who validate models are obligated to express concerns about model performance honestly, including concerns for non-human domains and to do so in governance statements of uncertainty rather than certainty. Field implementers are required to monitor and report using formally protected reporting channels. For all four groups of actors, the obligation to act is proportional to their capacity. An actor unable to foresee the consequence is not responsible for it. However, an actor with that capacity has a duty to enable prevention of consequences.

Through Moral Outcome Mapping (MOM), Layer 2 gives significant governance effects to CDMA. Table 1 shows the main differences between MOM and standard performance evaluations.

MOM does not eliminate the need to evaluate technical performance. Rather, it situates the assessment of technical performance within a framework that makes additional (non-human) harm visible and accountable to governance. In cases of cross-domain harms that are likely to occur but cannot be precisely measured, a precautionary principle is implemented and the burden of proof is placed on the proposer to demonstrate that the deployment is justified.

Layer 3 aims to address the gap of residual accountability. In situations characterised by the many hands problem, where no individual can be considered responsible for harm resulting from numerous inconsequential actions (in the sense of a single contribution) performed by multiple organisations, accountability appears to disappear. LRA addresses this by distributing residual accountability to the collective of various institutions: the developers, deployers, regulators, funders and research governance consortia of a given AI system. This collective, rather than any individual, is liable to make redress for the harm and is required to modify its practices to prevent further harm.

Whether collectives can bear responsibility is a contested question, dividing what Bovens (1998, pp. 45-60) describes as corporatists from individualists. The position taken here follows French (1984, pp. 30-45): what qualifies a collective as a bearer of responsibility is not its legal personality but its possession of a Corporate Internal Decision Structure, through which recognition rules confer an intentional character on corporate acts. The entities involved (health ministries, universities, international health agencies, conservation agencies, bilateral aid agencies) satisfy that condition. The LRA model therefore does not create a new moral subject; it allocates residual accountability to entities whose decision structures already render them capable of bearing it. In low-income settings (especially in sub-Saharan Africa, where numerous One Health AI deployments occur), the governance infrastructure this layer presupposes may be absent. In such cases, it is the responsibility of the international actors and donors launching the project to create the necessary governance infrastructure. This is a requirement of international research ethics. International research funders must demonstrate the existence of adequate governance infrastructure as a prerequisite for funding. Governance should not be considered a concern that can be postponed until after the project has been initiated.

An objection arises at this point, since accountability may simply disappear again at the collective level. Bovens (1998, pp. 45-60)

Table 1: Standard performance assessment compared with Moral Outcome Mapping

Dimension	Standard Approach	Moral Outcome Mapping
Outcomes assessed	Human health metrics only	Human welfare, animal welfare and environmental integrity
Harm assessment	Demographic fairness measures	Cross-domain harm assessment using CDMA
Foreseeability	Technical error bounds	Ecological feedback modelling; animal capability impact mapping
Trade-off handling	Implicit, by the model's objective function	Explicit deliberation, recorded and publicly available
Uncertainty	Confidence intervals on primary metric	Honest uncertainty profiles for each domain

states the difficulty plainly, observing that broad models of collective accountability invite free-rider problems and risk an outcome in which everyone is accountable after the event and therefore no one is. Three features of Layer 3 are intended to prevent this. The first is that the layer is residual and does not replace the obligations established at Layer 1, which hold independently of it. Collective accountability is what remains once those obligations have been discharged, not an alternative to discharging them. The second is that what the collective owes is specified rather than open-ended, since under French's principle of responsive adjustment it owes not an expression of regret but the alteration of the practice that produced the harm. Disclosure of what was altered, and on what evidence, forms part of that obligation. The third is that the obligation must be assigned rather than left general, and where a consortium is constituted its governance agreement should designate the locus within it that carries the responsive-adjustment duty. The obligation then attaches to a specific body rather than to the consortium at large. Without that designation, Layer 3 would reproduce the disappearance of responsibility it is intended to prevent.

Two instances demonstrate how the model operates. The first relies on the recorded use of AI in the conservation of elephants in Southern and Eastern Africa. AI has been used to forecast likely instances of poaching and aid in the deployment of rangers (Brickson et al., 2023, pp. 1-13). When the AI makes a false positive, rangers are prompted to move without reason. The repeated intrusion by rangers disturbs the protected area. It is unethical to stress the herd of elephants. It also puts a strain on the group of elephants, raises their stress levels, increases calf mortality and disrupts their movements (Gobush et al., 2008; Wiśniewska et al., 2022). These damages are not considered when the governance system measures detection accuracy. According to the LRA model, the developers must provide measures of the false alarm rates and their predicted ecological costs in the pre-deployment assessment. For authorisation to be granted, policymakers must provide a cross-domain cost of false positives not previously considered. This makes governance possible.

The second example concerns zoonotic surveillance AI in sub-Saharan Africa. Spillover risk detection relies on the sampling of large populations of wildlife. This intensive sampling is likely to cause changes in behaviour and immune response in the wildlife, resulting in stress and displacement. This may in turn alter pathogen dynamics and increase the risk of spillover – the very outcome the system is intended to prevent. The evidence that sampling alone is sufficient to shift these dynamics remains incomplete, which

is itself the point: a governance system that never measures the effect cannot establish whether it occurs. Structurally, human-centred governance cannot detect this phenomenon. Under CDMA, the phenomenon becomes a governance problem. Developers are required to incorporate the health of the wildlife into the models, which will improve the identification of cross-domain harm and make the models more precise. Policymakers are required to conduct forward-looking ecological assessments prior to approving the sampling. In the case of inadequate governance architecture for the consortium, international funders must fulfil the requirements to establish governance before the research begins.

4.1 Additional Empirical Evidence

Beyond elephant and zoonotic conservation and surveillance, other empirically grounded examples in AI governance illustrate the fundamental challenges facing One Health. In precision agriculture, AI-enabled pesticide application systems are designed to optimize crop yields; however, they disregard the non-target effects on the health and viability of the soil microbiome and pollinator populations (Tscharntke et al., 2012, pp. 53-59). Marine conservation AI used to identify illegal fishing focus on economic losses with less regard to the deaths of bycatch of non-commercial species, even though the loss of keystone predators will have an impact on all other species in the ecosystem (Sala et al., 2021). Wildlife corridor planning algorithms favor metrics of connectivity for charismatic megafauna, and disregard impacts on invertebrate dispersal and plant gene flow and would treat these impacts as externalities and not as integral aspects of ecosystem health (Tucker et al., 2021). In each of these cases, the embedded objective functions of AI automatically place a valuation on ecological impacts, though governance structures place control far away from those that experience the impacts.

4.2 Theoretical Implications

Several significant issues arise when considering the implications of the presumed flattening of humans and non-humans in One Health AI governance. First, removing humans as the only moral subject does not grant moral equivalence to non-humans. Instead, it requires a means of resolving the inescapable conflicts of interest among the welfare of sentient beings, the integrity of ecosystems, and human interests when all cannot be reconciled. Second, the opposition to ecocentric governance from some institutions may be attributable to the concerns of trade-offs not from a lack of bureaucratic zeal. If, within the rights framework, we cannot balance the suffering of animals with human gain, and within a holistic framework, we cannot balance ecosystems, then many forms of human technological progress become ethically untenable. This creates a major dilemma in determining how far we can take non-anthropocentric AI governance within contemporary technological progress. Finally, while all moral subjects, including non-humans, together with AI may have the capacity to challenge human interests, this assumes a conflict of interests, an assumption challenged by the One Health policy in prioritizing reconciliation over acceptance of zero-sum outcomes.

5. CONCLUSION

The aim of this article is to delineate the categories of governance failure displayed by the application of AI to One Health. Unlike governance failures associated with opacity and individual responsibility, research on these types of failures is scarce.

One of these governance failures is the structural imbalance of control and consequences. In One Health, the governance of AI fails when individuals with structural control over the design and implementation of systems end up being, by structural design, insulated from the consequences of that control. In contrast, control over the

Systems is exercised by elements of the Field (communities, animals, ecosystems) and they do bear the consequences. This cannot be characterised as the "many-hands problem," and is not likely to be solved by any increase in the quality of deliberation among the various constituents of the Systems. In fact, it will call for a re-distribution of responsibilities in accordance with what each constituent of the Systems is likely to make a substantial contribution to and will be governed by the principle, "where there is capacity, there is responsibility."

Another example of these governance failures is the clash between the normative frameworks of One Health and the 'anthropocentric' moral frameworks of AI governance, which are imposed on many of today's AI governance frameworks. One Health model requires the integration of humans, animals and the environment. AI governance models within One Health are absent. In the existing literature, this absence is primarily articulated as a contradiction, while in the real world, this creates negative impacts on animals and ecosystems, which is a problem that the existing governance frameworks do not address and regulate.

The Layered Relational Accountability (LRA) Model has three governance layers which enable the calibration of the constraining obligations cited above and the problems associated with ensemble accountability. The first layer consists of role-adjusted future obligations, the second layer focuses on cross-domain moral assessment through Moral Outcome Mapping and the third layer is ensemble accountability. The EU AI Act (2024) incorporates obligations for deployers aligning with the institutional accountability layer of the LRA Model. Within the framework of the EU Act, cross-domain impact provisions, when incorporated into high-risk One Health AI EU regulatory compliance AI assessments, would enable the Act's obligations on environmental sustainability. The UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), with its cross-border normative framework and concern for ecosystems, situate the CDMA layer. Within the World Organisation for Animal Health (WOAH) Terrestrial Animal Health Code of 2024, the welfare assessment of harm to animals of surveillance and sampling, along with the Kunming-Montreal Global Biodiversity Framework (CBD, 2022), provide comparable frameworks to evaluate ecosystem health. An example of how governance obligations can be structured within financially constrained contexts is the African Union Data Policy Framework (2022).

The model has certain defined limitations. Currently, there are no widely accepted tools for assessing cross-domain harm that can be integrated into a governance model for various contexts of One Health. Although there are certain animal welfare standards (Sneddon et al., 2014, pp. 201-212) and frameworks for the assessment of ecology (Scholes & Biggs, 2005, pp. 45-49), governance-sensitive assessments for both domains that can be utilised by non-experts are lacking. Rather, this portrays a significant gap in research and should not be a reason to suspend governance discussions. This model requires empirical validation. The elephant conservation case study utilised evidence from earlier uses, the model was illustrated in the surveillance study, but the model's application was neither described, nor was its potential for altering institutional practices and reducing cross-domain harm. This constitutes a gap that empirical work would best resolve. More research is also needed on the cross-border institutional liability, speciesism and its effect on conservation and public health Artificial Intelligence (AI) in the context of governance and the ability to operate (constrain) in lower resource environments.

There is no valid reason for delaying action on the issues that remain unaddressed. As has been outlined here with some brevity, effects are already being experienced in governance systems

that remain unaware of them. In the context of One Health, governance of AI needs to incorporate animals and ecosystems and may be as fragmented as governance of humans, if not more so. The Layered Relational Accountability model is a governance model that may be able to address what has been articulated in terms of structure, operationalism and realism.

Conflict of Interest: There are no competing interests in relation to the financial or personal spheres that may have impacted the research reported in this article.

AI Use Disclosure: The author acknowledges the use of AI-assisted tools for grammar checking and citation formatting. All conceptual development, argumentation, analysis, and writing remain the author's original work. AI tools were used solely for editorial assistance in accordance with journal guidelines.

REFERENCES

African Union. (2022). *Data policy framework*. African Union Commission. https://au.int/en/documents.

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82–115. https://doi.org/10.1016/j.inffus.2019.12.012.

Beck, U. (1992). *Risk society: Towards a new modernity*. Sage.

Bovens, M. (1998). *The quest for responsibility: Accountability and citizenship in complex organisations*. Cambridge University Press.

Brickson, L., Zhang, L., Vollrath, F., Douglas-Hamilton, I., & Titus, A. J. (2023). Elephants and algorithms: A review of the current and future role of AI in elephant monitoring. *Journal of the Royal Society Interface, 20*(208), Article 20230367. https://doi.org/10.1098/rsif.2023.0367.

Convention on Biological Diversity. (2022). *Kunming-Montreal Global Biodiversity Framework* (CBD/COP/DEC/15/4). United Nations. https://www.cbd.int/gbf/.

European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union*. https://eur-lex.europa.eu.

Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge University Press.

French, P. A. (1984). *Collective and corporate responsibility*. Columbia University Press.

Gobush, K. S., Mutayoba, B. M., & Wasser, S. K. (2008). Long-term impacts of poaching on relatedness, stress physiology, and reproductive output of adult female African elephants. *Conservation Biology, 22*(6), 1590-1599. https://doi.org/10.1111/j.1523-1739.2008.01035.x.

Leopold, A. (1949). *A Sand County almanac, and sketches here and there*. Oxford University Press.

Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology, 6*(3), 175-183. https://doi.org/10.1007/s10676-004-3422-1.

Nissenbaum, H. (1994). Computing and accountability. *Communications of the ACM, 37*(1), 72-80. https://doi.org/10.1145/175222.175228.

Nussbaum, M. C. (2006). *Frontiers of justice: Disability, nationality, species membership*. Harvard University Press.

Plowright, R. K., Field, H. E., Smith, C., Divljan, A., Palmer, C., Tabor, G., Daszak, P., & Foley, J. E. (2008). Reproduction and nutritional stress are risk factors for Hendra virus infection in little red flying foxes (Pteropus scapulatus). *Proceedings of the Royal Society B: Biological Sciences, 275*(1636), 861-869. https://doi.org/10.1098/rspb.2007.1260.

Plowright, R. K., Parrish, C. R., McCallum, H., Hudson, P. J., Ko, A. I., Graham, A. L., & Lloyd-Smith, J. O. (2017). Pathways to zoonotic spillover. *Nature Reviews Microbiology, 15*(8), 502-510. https://doi.org/10.1038/nrmicro.2017.45.

Qi, Y., Schölkopf, B., & Jin, Z. (2024). *Causal responsibility attribution for human-AI collaboration* (arXiv:2411.03275). arXiv. https://doi.org/10.48550/arXiv.2411.03275.

Regan, T. (1983). *The case for animal rights*. University of California Press.

Sala, E., Mayorga, J., Bradley, D., Cabral, R. B., Atwood, T. B., Auber, A., ... & Lubchenco, J. (2021). Protecting the global ocean for biodiversity, food and climate. *Nature*, 592(7854), 397-402.

Santoni de Sio, F., & Mecacci, G. (2021). Four responsibility gaps with artificial intelligence: Why they matter and how to address them. *Philosophy & Technology, 34*(4), 1057-1084. https://doi.org/10.1007/s13347-021-00450-x.

Scholes, R. J., & Biggs, R. (2005). A biodiversity intactness index. *Nature, 434*(7029), 45-49. https://doi.org/10.1038/nature03289

Sen, A. (1999). *Development as freedom*. Oxford University Press.

Singer, P. (1975). *Animal liberation: A new ethics for our treatment of animals.* Random House.

Sneddon, L. U., Elwood, R. W., Adamo, S. A., & Leach, M. C. (2014). Defining and assessing animal pain. *Animal Behaviour, 97*, 201-212. https://doi.org/10.1016/j.anbehav.2014.09.007.

Tang, R., Li, W., Zhu, D., Shang, X., Guo, X., & Zhang, L. (2020). Raging elephants: Effects of human disturbance on physiological stress and reproductive potential in wild Asian elephants. *Conservation Physiology, 8*(1), coz106. https://doi.org/10.1093/conphys/coz106.

Thompson, D. F. (1980). Moral responsibility of public officials: The problem of many hands. *American Political Science Review, 74*(4), 905-916. https://doi.org/10.2307/1954312

Tscharntke, T., Clough, Y., Wanger, T. C., Jackson, L., Motzke, I., Perfecto, I., ... & Whitbread, A. (2012). Global food security, biodiversity conservation and the future of agricultural intensification. *Biological Conservation*, 151(1), 53-59.

Tucker, M. A., Böhning-Gaese, K., Fagan, W. F., Fryxell, J. M., Van Moorter, B., Alberts, S. C., ... & Mueller, T. (2021). Moving in the Anthropocene: Global reductions in terrestrial mammalian movements. *Science*, 359(6374), 466-469.

UNESCO. (2021). *Recommendation on the ethics of artificial intelligence.* United Nations Educational, Scientific and Cultural Organization. https://unesdoc.unesco.org.

van Wynsberghe, A. (2021). Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics, 1*(3), 213-218. https://doi.org/10.1007/s43681-021-00043-6.

Whittlestone, J., Nyrup, R., Alexandrova, A., & Cave, S. (2019). The role and limits of principles in AI ethics: Towards a focus on tensions. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 195-200). ACM. https://doi.org/10.1145/3306618.3314289.

Wiśniewska, M., Puga-Gonzalez, I., Lee, P., Moss, C., Russell, G., Garnier, S., & Sueur, C. (2022). Simulated poaching affects global connectivity and efficiency in social networks of African savanna elephants – An exemplar of how human disturbance impacts group-living species. *PLOS Computational Biology, 18*(1), Article e1009792. https://doi.org/10.1371/journal.pcbi.1009792

WHO, FAO, World Organisation for Animal Health (WOAH), & UNEP. (2021). *One Health High-Level Expert Panel annual report.*

Whittlestone, J., Nyrup, R., Alexandrova, A., & Cave, S. (2019). The role and limits of principles in AI ethics. In Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society.

World Health Organization, Food and Agriculture Organization, World Organisation for Animal Health, & United Nations Environment Programme. (2021). *Tripartite and UNEP support OHHLEP's definition of One Health.* World Health Organization. https://www.who.int.

World Organisation for Animal Health. (2024). *Terrestrial animal health code* (32nd ed.). WOAH. https://www.woah.org.

Young, I. M. (2011). *Responsibility for justice.* Oxford University Press.

COPYRIGHT AND LICENSE

By submitting and publishing the article in **Vantage: Journal of Thematic Analysis**, author(s) remains the copyright owner. This is a peer-reviewed multidisciplinary Platinum OA publication of the Centre for Research, Maitreyi College, University of Delhi (https://vantagejournal.com). The author(s) grant the publisher a non-exclusive, worldwide, perpetual license to publish, reproduce, distribute, archive, index, and communicate the work in all media and formats.

This is an Open Access article distributed under the terms of the **Creative Commons Attribution 4.0 International (CC BY 4.0) License**. This license permits unrestricted use, sharing, adaptation, distribution, and reproduction in any medium or format, provided appropriate credit is given to the original author(s) and the source.

HOW TO CITE?

Anshu Kumar (2026). Who Controls, Who Suffers and Who Counts? Rethinking Accountability in AI-Driven One Health. *Vantage: Journal of Thematic Analysis*, 07(02), 121-132. https://doi.org/10.68903/vjta.2026.v07i02.02