

Submitted: 6 February 2025  
Accepted: 24 September 2025  
Published online: 7 April, TAPS 2026, 11(2), 131-133  
<https://doi.org/10.29060/TAPS.2026-11-2/II3665>

# Teaching ethics in the age of AI: Reflections from a medical educator

Pacifico Eric Eusebio Calderon<sup>1,2,3</sup>

<sup>1</sup>St. Luke's Medical Center, <sup>2</sup>National Children's Hospital, Quezon City, Philippines; <sup>3</sup>Faculty of Laws, University College London, United Kingdom

## I. INTRODUCTION

Artificial intelligence (AI) is now a familiar presence in healthcare. Frequently introduced as a means of augmenting clinical work, it also invites reflection on how the character of medical practice is evolving. AI may influence not only clinical decision-making (Byrne, 2023), but also the production of medical knowledge, the framing of ethical questions, and the assignment of responsibility when outcomes are uncertain, or contested (Aquino, 2023). As these technologies become embedded in the routines of care, they may begin to reshape prevailing conceptions of clinical judgement, moral attentiveness, and professional responsibility.

This article reflects on how the increasing integration of AI into clinical settings may be subtly reconfiguring the ethical landscape of medicine and considers how such shifts might be addressed in ethics education. It critiques three domains in which new tensions emerge: the erosion of space for moral discernment, epistemic injustice within data-driven systems, and the fragmentation of responsibility across increasingly distributed environments.

In place of technical prescriptions, the paper invites educators to reflect on the kinds of moral sensibilities we seek to cultivate in those learning to practise medicine—whether students, trainees, or professionals in continuing formation. How might ethical capacities be fostered in healthcare systems increasingly configured by technologies that clinicians do not design and cannot fully control? What dispositions might be required to

remain attentive, critical, and responsive within datafied systems of care?

These questions are pursued through a series of reflections on how AI is reshaping attentiveness, knowledge, and responsibility—and on how ethics education might engage with these shifts with nuance and care.

## II. PRESERVING ATTENTIVENESS IN ALGORITHMIC ENCOUNTERS

The clinical encounter between doctor and patient remains foundational to medical practice. Such moments are rarely straightforward. They require not only clinical reasoning but also the capacity to navigate uncertainty, emotional nuance, and what is often unspoken. Ethical significance in these situations is not always immediately visible; it may emerge in a hesitation, a glance, or an absence that nonetheless invites moral attention. Attending to these subtleties requires what might be called moral attentiveness: the ability to notice what might otherwise be missed, and to recognise that ethical meaning is not always legible within procedural norms.

This form of attentiveness finds philosophical resonance in Tronto's (1993) articulation of care as relational, situated, and responsive to particular needs. On this view, good care cannot be reduced to procedural fidelity or technical adequacy. It involves a willingness to remain present, to slow down, and to engage meaningfully with the lived experience of the person before us.

Yet this space for attentiveness may be increasingly constrained by the integration of AI systems into clinical work (Dalton-Brown, 2020). Many such systems are designed to promote speed, consistency, and institutional efficiency. They may generate clinical suggestions before a patient is even seen, structure how documentation is produced, and guide decisions in ways that encourage adherence to predefined pathways (Byrne, 2023). Whilst these tools may support workflow, their underlying logic can narrow the reflective space needed for ethical discernment. When clinical attention is structured in advance by algorithmic cues, the opportunity to pause, to wonder, or to respond to the unexpected may begin to contract (Dalton-Brown, 2020).

This shift presents a challenge not only for practice but also for pedagogy. If AI systems increasingly shape how care is delivered, then ethics education must consider how to support learners in sustaining forms of attentiveness that resist automation. What pedagogical approaches might preserve interpretive openness in contexts structured around procedural closure? This may call for renewed emphasis on cultivating presence, responsiveness, and moral imagination (Tronto, 1993)—qualities that remain vital to ethical practice but are difficult to codify, even harder to delegate to machines.

### III. RECOGNISING EXCLUSIONS IN DATA-DRIVEN KNOWLEDGE

AI systems are often introduced with the promise of improving efficiency, promoting consistency, and mitigating bias or human error in clinical practice (Byrne, 2023). Yet the data on which such systems rely on is rarely neutral. Most are developed in high-resource environments and trained on datasets that reflect the clinical norms, priorities, and assumptions of those contexts. As a result, some experiences of illness are amplified, whilst others are excluded, distorted, or remain unrecognised altogether (Aquino, 2023). These exclusions are not merely technical gaps but carry ethical implications, shaping whose suffering is acknowledged and whose is not.

This form of marginalisation has been theorised by Fricker (2007) as epistemic injustice: harm that arises when individuals or groups are excluded from contributing to shared knowledge, or when their insights are misrepresented, dismissed, or devalued. In healthcare, for instance, this may occur when symptoms presented by certain populations are not recognised by AI systems trained on different demographics, or when non-standard forms of expression—whether cultural context, embodied experience, or vernacular language—are treated as deviations rather than legitimate sources of insight.

For learners, the effects of these omissions may unfold incrementally. What is consistently absent from training tools may come to feel irrelevant; what is frequently represented may appear normative. Over time, these patterns can come to shape how clinicians perceive credibility, construct clinical knowledge, and attend to suffering. The narrowing of epistemic horizons is rarely intentional, but it has moral consequences (Fricker, 2007). Certain voices come to dominate, and some forms of distress remain invisible within algorithmic frames (Aquino, 2023).

Ethics education might respond by fostering what could be described as epistemic humility: an awareness that all systems of knowledge, however advanced, can be partial and situated. This involves not only recognising what is missing but also cultivating the capacity to dwell with uncertainty and remain attentive to the margins of representation. Especially in global or resource-constrained settings—where imported AI systems may misrepresent local realities—this disposition is not only prudent, but also pedagogically essential. The task is not to reject such tools outright, but to approach them with critical distance, sustained attentiveness, and moral care.

### IV. NAVIGATING RESPONSIBILITY IN DISTRIBUTED SYSTEMS

The deeper integration of AI into medical work is also reshaping how professional responsibility is perceived. AI is often viewed as a form of support—something that augments rather than replaces the clinician (Byrne, 2023). Yet in practice, the distinction between assistance and authority may be far from straightforward. When outputs appear confident and their reasoning opaque, clinicians may feel compelled to defer, even in the presence of doubt.

Efficiency is frequently presented as the primary feature of such tools. Yet efficiency is rarely neutral. It tends to reflect the priorities of institutions—throughput, documentation, predictability—rather than the relational demands of ethical care. The logic of efficiency that underpins many AI systems often aligns with these institutional imperatives. In doing so, it may shift the moral orientation of practice away from responsiveness to particular needs and toward standardised procedures. As Tronto (1993) reminds us, responsibility is not simply the performance of tasks; it involves attentiveness to needs that unfold slowly or resist resolution. When time saved is redirected toward institutional metrics, the more reflective dimensions of medical work may be compromised.

Within such systems, responsibility can become fragmented and elusive. Clinical decisions often arise

through a convergence of human reasoning, algorithmic suggestion, and organisational structure (Aquino, 2023). Yet when outcomes are contested, accountability frequently reverts to the individual clinician. For learners, this may create a disorienting professional ethical terrain. They are expected to exercise moral judgement in contexts that may increasingly constrain their agency.

In response, ethics education might offer more than abstract principles. It can support learners in reflecting on what it means to assume responsibility in conditions where control is partial and in navigating situations where the line between autonomous professional judgement and systemic compliance is blurred. Discernment—understood here as the capacity to act with care in the face of uncertainty, complexity, or constraint—becomes central to this pedagogical task. It is perhaps not a matter of identifying the right answer, but of cultivating the sensitivity to decide well when clarity is elusive.

## V. CONCLUDING REFLECTIONS

The discussion has traced how the integration of AI into clinical practice may be reshaping the moral contours of medicine—not through sudden rupture, but through subtler shifts in how clinicians attend, decide, and take responsibility. It explored three such developments: the narrowing of interpretive space in clinical encounters; the exclusions embedded in data infrastructures; and the dispersal of professional responsibility across distributed systems. These changes do not call for rejection, but for careful recalibration—one that sustains moral attentiveness, epistemic humility, and ethical discernment within systems increasingly structured around speed, efficiency, and procedural logic. Each domain also opens space for pedagogical reflection, prompting us to ask not only how we teach ethics, but what kinds of moral sensibilities we hope to preserve.

What forms of teaching might support the cultivation of these capacities? How might empirical inquiry illuminate the lived ethical consequences of AI integration across diverse institutional and cultural contexts? And how can educators, ethicists, clinicians, and curriculum designers engage in shared dialogue about the values we wish to uphold amid technological transformation? Much, however, remains unsettled. The task ahead may lie in cultivating—in learners, and in ourselves as educators—a disposition to remain with ethical demands that technological systems cannot resolve. Such a pedagogy would rest not on certainty but on reflective presence, epistemic humility, and a sustained attentiveness to the forms of care we still hope to practise in a world increasingly shaped by algorithmic reasoning.

## Notes on Contributors

The author solely conceptualised, drafted, and revised the manuscript.

## Ethical Approval

As this is a theoretical study, it does not involve human participants or data collection. Accordingly, ethical approval was not applicable.

## Funding

This study did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

## Declaration of Interest

The author declares no conflict of interest.

## References

- 
- Aquino, Y. S. J. (2023). Making decisions: Bias in artificial intelligence and data-driven diagnostic tools. *Australian Journal of General Practice*, 52(7), 439–444. <https://doi.org/10.31128/AJGP-12-22-6630>
- Byrne, M. F., Parsa, N., Greenhill, A. T., Chahal, D., Ahmad, O., & Bagci, U. (Eds.). (2023). *AI in clinical medicine: A practical guide for healthcare professionals*. John Wiley & Sons.
- Dalton-Brown, S. (2020). The ethics of medical AI and the physician–patient relationship. *Cambridge Quarterly of Healthcare Ethics*, 29(1), 115–121. <https://doi.org/10.1017/S0963180119000847>
- Fricke, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198237907.001.0001>
- Tronto, J. (2020). *Moral boundaries: A political argument for an ethic of care*. Routledge. <https://doi.org/10.4324/9781003070672>

---

\*Pacífico Eric Eusebio Calderon  
Faculty of Laws, University College London  
4-8 Endsleigh Gardens,  
London WC1H 0EG  
United Kingdom  
Email: [pacifico.calderon.24@ucl.ac.uk](mailto:pacifico.calderon.24@ucl.ac.uk)