ISSN- 0301-1216
Indian J. Prev. Soc. Med. Vol. 57, No.3, 2026

Artificial Intelligence in Public Health: A Boon or Bane

Abhishek Jaiswal[1]

ABSTRACT

We outsource our memory to silicon. Soon, we will outsource our clinical intuition. Integrating machine learning into epidemiology forces a brutal re-evaluation of how medical professionals actually think. Neural networks mimic executive function. They do not merely archive patient demographics. At the primary care level, delegating clerical friction to algorithms feels like oxygen. It liberates a medical officer's working memory, preserving biological bandwidth for authentic causal inference. We democratize triage through diagnostic algorithms. Yet outsourcing cognitive struggle invites a profound Socratic atrophy. Algorithms remain spectacular correlation engines entirely blind to causation. When biostatisticians passively accept automated risk stratifications without dissecting the underlying societal mathematics, epidemiological intuition rots. Relying on flawed proxy variables hardwires historical prejudices directly into local surveillance infrastructure. It mathematically disenfranchises the exact vulnerable cohorts who already avoid formal health systems. Epistemic stewardship is the only viable countermeasure. Predictive models require aggressive scepticism. We must demand explainability. We must abandon lazy list wise deletion in messy community surveys, forcing advanced statistical imputation instead. The machine advises. Human intellect decides. Artificial intelligence will ruthlessly amplify our existing structural intent, but it cannot replace the productive, messy struggle of genuine public health inquiry.

Keywords: Cognitive Offloading, Algorithmic Bias, Digital Epidemiology, Community Medicine, Predictive Surveillance, Health Equity.

Authour(s) Details:

1. Assistant Professor, Department of Community Medicine, Institute of Medical Sciences, Banaras Hindu University, Varanasi- 221 005, U.P., India

Corresponding Address: Dr. Abhishek Jaiswal, Assistant Professor, Department of Community Medicine, Institute of Medical Sciences, Banaras Hindu University, Varanasi, U.P., India; Phone: 7982298020, 9205624317

Citation: Jaiswal A. Artificial Intelligence in Public Health: A Boon or Bane. Indian J Prev Soc Med, 2026; 57 (3): **208-212.**
DOI: https://doi.org/10.67212/ijpsm.v57i3.265

Sequence of Article: **Submission 08.08. 2026** **Accepted: 10.08.2026** **Published: 30.09.2026**

Prior Publication: Nil; Source of Funding: Nil; Conflicts of Interest: None, Article #

Introduction

Socrates despised the alphabet. He viewed papyrus as a cognitive threat. The Greek philosopher argued violently that externalizing human memory onto written scrolls would rot innate intellect. Subsequent generations, he warned, would inherit the mere illusion of wisdom rather than its agonising, internally generated reality. We mock his paranoia today. We write everything down. Yet his foundational anxiety remains violently relevant inside a modern community health centre.

Machine learning does not just store district health registry data. It mimics human executive function. It predicts. It categorizes. It triages. This forces an existential clinical question upon the medical community. Are predictive algorithms elevating our epidemiological intellect, or are they blindly automating our historical prejudices? The answer is not a binary switch. It depends entirely on whether we treat neural networks as cognitive prosthetics or as wholesale intellectual replacements.

Thesis

Epidemiological fieldwork is exhaustive. The human brain operates on a bottlenecked working memory. When researchers burn their mental stamina finding associations across chaotic district-level registries, they lose the capacity to think deeply. Harmonizing incomplete morbidity data from auxiliary healthcare workers destroys biological bandwidth. There is nothing left for causal inference.

The machine becomes indispensable here. Cognitive psychologists term this transfer of computational effort "cognitive offloading" [1]. By dumping low-level processing tasks onto external systems, we free up neural real estate. Eric Topol conceptualises this specific technological pivot as high-performance medicine[2]. Delegating relentless data drudgery to silicon theoretically liberates human cognitive reserves. We reclaim the capacity to synthesise complex social determinants. We look patients in the eye again. We rebuild fractured community trust.

Deep learning models transform the muddled sludge of electronic health records into clinical foresight[3]. At the grassroots level, this offloading is not a luxury. It is survival. The mismatch between community morbidity and available specialist hours routinely collapses into mathematical absurdity. A single medical officer in an under-resourced outpatient department might consult two hundred patients before noon. The clinic bleeds time. Diagnosis becomes a reflex rather than an intellectual process.

Convolutional neural networks alter this grim calculus. Look at dermatology. Algorithms can now classify malignant melanoma with a diagnostic competence rivalling board-certified specialists[4]. They detect microvascular damage in retinal fundus photographs with terrifying speed and accuracy[5]. This diagnostic frontier even extends to breast cancer screening. Algorithmic systems consistently outperform human radiologists in identifying malignant tissue while drastically reducing false-positive biopsies[6].

Picture a fatigued physician in a primary health clinic. They point a smartphone at an abnormal chest radiograph. A localised algorithmic triage system flags a high-probability tuberculosis lesion in seconds. The geographical divide vanishes instantly. We democratise diagnostic expertise.

The burnout epidemic currently gutting the global medical workforce demands this exact intervention. Clinical fatigue is not a moral failing of the physician. It is a systemic mathematical certainty. Survey data consistently reveals a desperate clinician hunger to delegate administrative burdens to machines just to survive the week[7]. Liberated from staring at health management information portals, physicians reclaim the relational empathy that algorithms cannot simulate.

Predictive digital epidemiology pushes this boundary even further. Traditional infectious disease surveillance is done retrospectively. We only realize an outbreak is accelerating once local fever wards overflow. Predictive modelling acts as the bedrock of modern learning healthcare systems[8]. By ingesting localized wastewater pathogen metrics, global transit vectors, and shifting climatic data simultaneously, administrators stop reacting. They can now anticipate. They gain the foresight to deploy micro-containment zones before transmission peaks. They target vector-control resources precisely without shuttering entire local economies. Intellect is not replaced here; it is mathematically amplified.

Antithesis

But the psychological trap remains severe. The boundary separating cognitive offloading from cognitive atrophy is microscopic. Outsource the manual data entry, and you thrive. “Outsource the productive struggle of causal inference. The machine may deliver, but the human intellect will wither.” Public health intelligence relies on understanding why a disease propagates through a marginalised block. Recognising the statistical pattern of transmission is never enough. Deep learning models are correlation engines. They remain completely blind to causation[9]. If a biostatistician passively accepts a machine-generated regression model without aggressively interrogating the underlying societal variables, their intuition rusts. *We risk treating statistical symptoms. We remain entirely ignorant of the sociological disease.*

A decade ago, tech giants attempted to predict influenza spikes purely by tracking search engine queries. It failed spectacularly. The algorithm vastly overestimated flu prevalence because it lacked human epidemiological context (10). Researchers now term this phenomenon ***big data hubris***. Relying on automated models without understanding the underlying mathematics induces a digital Socratic atrophy. We project the illusion of epidemiological wisdom. We possess zero actual comprehension.

This uncritical embrace of algorithmic systems actively threatens global health equity. We must aggressively dismantle the myth of algorithmic neutrality. Algorithms traffic in cold probabilities. They wear a deceptive cloak of mathematical

objectivity. Yet foundational healthcare algorithms are predominantly trained on cohorts drawn from affluent, digitized, high-resource environments. Transposed onto marginalised populations, their predictive efficacy collapses.

Diagnostic models systematically underdiagnose underserved patient populations[11]. They effectively label sick patients as healthy. The algorithms deny them critical secondary care. Algorithms act as mathematical mirrors. They reflect every historical inequity embedded in their training sets. ***Garbage in, bias out.***

Implementing machine learning in healthcare introduces severe ethical dilemmas that software engineers cannot solve in isolation[12]. We do not have to guess what algorithmic disenfranchisement looks like in practice. A landmark investigation examined a commercial algorithm managing healthcare resources for tens of millions of people. The team uncovered a staggering racial disparity[13]. Black patients were vastly sicker than White patients who shared the exact same algorithmic risk score.

The machine was not programmed with overt malice. It relied on a deeply flawed proxy variable. The designers used historical healthcare spending to predict future healthcare needs. Because systemic poverty and institutional biases historically suppressed medical spending on Black communities, the algorithm falsely deduced they were simply healthier. It locked them out of vital preventative care management programs. When we prioritise computational scale over rigorous human scrutiny, we hardwire historical discrimination directly into the architecture of future health systems. We weaponise our own data against the vulnerable.

Synthesis: Epistemic Stewardship

Resolving this tension requires a recalibration of how we govern health technology. We cannot treat artificial intelligence as proprietary software. It is foundational public health infrastructure. Mitigating existential risk demands an uncompromising roadmap for responsible machine learning that prioritises patient safety over rapid deployment[14]. We cannot patch ethical guardrails onto an algorithm after it disenfranchises a vulnerable community. Transparency must be engineered into the conceptualisation phase.

Data inclusivity requires strict regulatory mandates. Algorithmic models must prove statistical fairness across geographic and socioeconomic lines before field deployment. In routine rural surveillance, missing data is not a random statistical anomaly. It is a footprint of marginalisation. When dealing with messy community health surveys, advanced biostatistical techniques like multiple imputation must replace lazy list wise deletion. The latter practice notoriously silences the data of groups who lack consistent interactions with formal healthcare infrastructure. If we delete the incomplete rows, we delete the poorest patients from our predictive models.

Predictive accuracy can never justify clinical opacity. The medical community must demand explainable artificial intelligence. Current approaches to algorithmic explainability often provide a false sense of security, masking deeper structural flaws in the model[15]. Yet a strict human-in-the-loop framework remains non-negotiable. If an algorithm flags a rural district as a high-risk zone for an impending dengue outbreak, the medical officers deploying resources must know exactly which variables triggered the alert. Did the model detect shifting vector habitats, or did it misinterpret a flawed socioeconomic proxy? The machine advises. The human fundamentally decides.

Artificial intelligence must actively dismantle health disparities rather than passively observing them. Ethical integration requires deliberate policy frameworks designed to leverage technology specifically for the marginalized[16]. Instead of training models to predict hospital profitability or bed turnover rates, we should calibrate algorithms to identify populations facing the steepest barriers to clinical access. We map vulnerability. We direct proactive social support exactly where it is needed most.

Artificial intelligence is not an autonomous monster coming to destroy medical intellect. It is an incredibly powerful amplifier of our existing human intent. Deployed recklessly, it automates our worst biases. It serves as the exact cognitive crutch Socrates warned against. Yet, stewarded with epistemic humility and transparent methodology, it represents our greatest opportunity to rethink public health in a century. We must never stop questioning the machine. ***The true threat to the intellect of humanity is our own willingness to stop thinking.***

Conclusion

Artificial intelligence is neither an autonomous monster destined to dismantle the medical intellect nor a magical panacea that will effortlessly cure systemic public health crises. It is simply a tool- a powerful amplifier of our existing human intent.

If deployed recklessly, driven entirely by commercial convenience, computational haste, and unexamined datasets, machine learning will automate our historical biases and serve as the exact cognitive crutch Socrates warned against. It will strip the soul out of community medicine, leaving us with sterile, mathematically pristine models that remain completely blind to genuine human suffering.

Yet, if stewarded with epistemic humility, transparent methodology, and an uncompromising commitment to health equity, it represents our greatest opportunity to rethink public health delivery in a century. The survival of human intellect in this algorithmic age relies on one simple, non-negotiable imperative: we must never stop interrogating the machine. ***The true threat is not artificial intelligence; it is our own willingness to stop our own human intelligence.***

References

1. Risko EF, Gilbert SJ. Cognitive offloading. Trends Cogn Sci. 2016; 20 (9):676-688. doi: 10.1016/j.tics.2016.07.002. Web: https://doi.org/10.1016/j.tics.2016.07.002
2. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med. 2019;25 (1):44-56. doi: 10.1038/s41591-018-0300-7. Web: https://doi.org/10.1038/s41591-018-0300-7
3. Rajkomar A, Dean J, Kohane I. Machine learning in medicine. N Engl J Med. 2019; 380 (14):1347-1358. doi: 10.1056/NEJMra1814259. Web: https://doi.org/10.1056/NEJMra1814259
4. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017; 542 (7639):115-118. doi: 10.1038/nature21056. Web: https:// doi.org/ 10.1038/nature21056
5. Gulshan V, Peng L, Coram M, Stumpe MC, Wu D, Narayanaswamy A, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA. 2016;316 (22):2402-2410. doi: 10.1001/jama.2016.17216. Web: https://doi.org/10.1001/jama.2016.17216
6. McKinney SM, Sieniek M, Godbole V, Godwin J, Antropova N, Ashrafian H, et al. International evaluation of an AI system for breast cancer screening. Nature. 2020; 577 (7788):89-94. doi: 10.1038/s41586-019-1799-6. Web: https://doi.org/ 10.1038/s41586-019-1799-6
7. Pedro AR, Dias MB, Laranjo L, Cunha AS, Cordeiro JV. Artificial intelligence in medicine: A comprehensive survey of medical doctor's perspectives in Portugal. PLoS One. 2023;18(9):e0290613. doi: 10.1371/journal.pone.0290613. Web: https://doi.org/10.1371/journal.pone.0290613
8. Ankolekar A, Eppings L, Bottari F, Pinho IF, Howard K, Baker R, et al. Using artificial intelligence and predictive modelling to enable learning healthcare systems (LHS) for pandemic preparedness. Comput Struct Biotechnol J. 2024;24:412-419. doi: 10.1016/j.csbj.2024.05.014. Web: https://doi.org/10.1016/j.csbj.2024.05.014
9. Panch T, Mattie H, Celi LA. The "inconvenient truth" about AI in healthcare. NPJ Digit Med. 2019;2:77. doi: 10.1038/s41746-019-0155-4. Web: https://doi.org/10.1038/s41746-019-0155-4
10. Lazer D, Kennedy R, King G, Vespignani A. The parable of Google Flu: traps in big data analysis. Science. 2014; 343(6176):1203-1205. doi: 10.1126/science.1248506. Web: https://doi.org/10.1126/science.1248506

11. Seyyed-Kalantari L, Zhang H, McDermott MBA, Chen IY, Ghassemi M. Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations. Nat Med. 2021;27(12):2176-2182. doi: 10.1038/s41591-021-01595-0. Web: https://doi.org/10.1038/s41591-021-01595-0
12. Char DS, Shah NH, Magnus D. Implementing machine learning in health care addressing ethical challenges. N Engl J Med. 2018; 378 (11):981-983. doi: 10.1056/NEJMp1714229. Web: https://doi.org/10.1056/NEJMp1714229
13. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019; 366 (6464):447-453. doi: 10.1126/science.aax2342. Web: https://doi.org/ 10.1126/ science. aax2342
14. Wiens J, Saria S, Sendak M, Ghassemi M, Liu VX, Doshi-Velez F, et al. Do no harm: a roadmap for responsible machine learning for health care. Nat Med. 2019; 25 (9):1337-1340. doi: 10.1038/s41591-019-0548-6. Web: https://doi.org/ 10.1038/s41591-019-0548-6
15. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. Lancet Digit Health. 2021; 3 (11):e745-e750. doi: 10.1016/S2589-7500(21)00208-9. Web: https://doi.org/ 10.1016/S2589-7500(21)00208-9
16. Thomasian NM, Eickhoff C, Adashi EY. Advancing health equity with artificial intelligence. J Public Health Policy. 2021; 42 (4):602-611. doi: 10.1057/s41271-021-00319-5. Web: https://doi.org/10.1057/s41271-021-00319-5