Dissertation Report

On

"Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change"

At

Vital Strategies India

(21st Jan- 21st June)

By

Name: Dr. Prachi Thakral

Enroll No.: PG/24/076

Under the guidance of

Dr. Preetha G.S

(Dean Research)

PGDM (Hospital & Health Management) 2024-26



International Institute of Health Management Research New Delhi

Internship Training

On

Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change

at

Vital Strategies



by

Name Dr. Prachi Thakral

Enroll No. PG/24/076

Under the guidance of

Dr. Preetha G.S

PGDM (Hospital and Health Management) 2024-26



International Institute of Health Management Research New Delhi

The certificate is awarded to

Name: Dr. Prachi Thakral

in recognition of having successfully completed her internship

Title: AI-Data Quality Analyst

and has successfully completed her Project on

Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change

Organization: Vital Strategies

She comes across as a committed, sincere & diligent person who has a strong drive & zeal for learning.

We wish her all the best for future endeavors.

LM singh

Head of the Organization

TO WHOMSOEVER IT MAY CONCERN

This is to certify that **Dr. Prachi Thakral** student of PGDM (Hospital & Health Management) from International Institute of Health Management Research, New Delhi has undergone internship training at **Vital Strategies India** from 21st **January to 21st June**. The Candidate has successfully carried out the project designated to him during internship training and her approach to the study has been sincere, scientific and analytical. The Internship is in fulfillment of the course requirements.

I wish her success in all her future endeavors.

Digitally Signed

Punit Yadav

Dr. Punit Yadav
Academic and Student Affairs
IIHMR, New Delhi

Mentor-
Dr. Preetha G. S
IIHMR, New Delhi

Certificate of Approval

The following dissertation titled **"Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change"** at **"Vital Strategies India"** is hereby approved as a certified study in management carried out and presented in a manner satisfactorily to warrant its acceptance as a prerequisite for the award of **PGDM (Hospital & Health Management)** for which it has been submitted. It is understood that by this approval the undersigned do not necessarily endorse or approve any statement made, opinion expressed or conclusion drawn therein but approve the dissertation only for the purpose it is submitted.

Dissertation Examination Committee for evaluation of dissertation.

Name	Signature
Dr. Ratika Samtani	
Dr. S Siddhartha Sankar Reddy	

Certificate from Dissertation Advisory Committee

This is to certify that **Dr. Prachi Thakral,** a graduate student of the **PGDM (Hospital & Health Management)** has worked under our guidance and supervision. She is submitting this dissertation titled "**Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change**" at "Vital Strategies India" in partial fulfillment of the requirements for the award of the **PGDM (Hospital & Health Management).**

This dissertation has the requisite standard and to the best of our knowledge no part of it has been reproduced from any other dissertation, monograph, report or book.

Mentor-
Dr. Preetha G. S
IIHMR, New Delhi

Organization Mentor,
Dr. Nidhi Chaudhary
Principal Technical Advisor
Vital Strategies

INTERNATIONAL INSTITUTE OF HEALTH MANAGEMENT RESEARCH, NEW DELHI

CERTIFICATE BY STUDENT

This is to certify that the dissertation titled **Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change and** submitted by (Name) ……**Dr. Prachi Thakral**… Enrollment No. ……**PG/24/076**…… under the supervision of……**Dr. Preetha G.S**… for award of PGDM (Hospital & Health Management) of the Institute carried out during the period from ……**January**… to …**June**… embodies my original work and has not formed the basis for the award of any degree, diploma associate ship, fellowship, titles in this or any other Institute or other similar institution of higher learning.

Prachi Thakral
Signature

FEEDBACK FORM

Name of the Student: Dr. Prachi Thakral

Name of the Organisation in Which Dissertation Has Been Completed: Vital Strategies India

Area of Dissertation: AI- Data Quality Analyst (Data and AI)

Attendance: Regular

Objectives achieved:

- Understanding the AI landscape in India
- Learn to develop AI use cases for public health
- Experience working with multidisciplinary teams in an international NGO

Deliverables:

- Drafted three selected case studies on India's digital health infrastructure (ABDM, eSanjeevani, CoWIN) for the "Foundations and Future Report" 2.0 and mapped the national AI landscape across five healthcare domains.
- Key member of the execution team for organizing the "Actionable AI: Building Health Systems for Public Health Impact" panel discussion held around AI Impact Summit 2026.
- Created a holistic AI landscape for Vital Strategies, analysing use cases, donors, and competitors to identify strategic state-level partnership opportunities.
- Led foundational research and feasibility assessments to design successful concept notes and proposals for a major Google.org grant.
- Developed strategic proposals for two AI-driven initiatives: a cancer staging tool and an automated "Tobacco Quitline+ AI" system.
- Drafted the curriculum outline for capacity building of public health leadership on AI.

Strengths:

- Technically strong and has good writing skills
- Efficient and dedicated public health professional and has a learning attitude and takes up challenging assignments
- Self-motivated, takes initiative and a good team member with the ability to work with multidisciplinary teams

Suggestions for Improvement:

Prachi is just starting her career and I am sure as she works in teams, she will get opportunities to focus on soft skills. I believe that she pursues opportunities of career

progression, she will find it useful to learn and apply negotiation skills and leadership skills.

Suggestions for Institute (course curriculum, industry interaction, placement, alumni):

It would help if the supervisors of the interns have a better understanding of the expectations of the institute and align the learning objectives for the intern

Date: 24.07.2026

Place: New Delhi

Nidhi Chaudhary

Dr. Nidhi Chaudhary

Principal Technical Advisor, Vital Strategies

INTERNATIONAL INSTITUTE OF HEALTH MANAGEMENT RESEARCH (IIHMR)

Plot No. 3, Sector 18A, Phase- II, Dwarka, New Delhi-110075

Ph. +91-11-30418900, www.iihmrdelhi.edu.in

CERTIFICATE ON PLAGIARISM CHECK

Name of Student (in block letter)	**Dr. PRACHI THAKRAL**		
Enrollment/Roll No.	PG/24/076	**Batch Year**	2024-2026
Course Specialization (Choose one)	Hospital Management	Health Management	Healthcare IT
Name of Guide/Supervisor	Dr./ Prof.: PREETHA G.S.		
Title of the Dissertation **Assignment**	**ARTIFICIAL INTELLIGENCE IN HEALTHCARE: A DOUBLE-EDGED SWORD FOR CLIMATE CHANGE**		
Plagiarism detect software used	"**TURNITIN**"		
Similar contents acceptable (%)	Up to **15** Percent as per policy		
Total words and % of similar contents Identified	6%		
Date of validation (DD/MM/YYYY)	_19 _/07/ 2026_		

Guide/Supervisor

Name: Dr. Preetha G. S

Signature:

Report checked by

Institute Librarian

Signature:

Date:

Student

Name: Dr. Prachi Thakral

Signature: Prachi Thakral

Dean (Academics and Student Affairs)

Digitally Signed

Signature: Punit Yadav

Date:27.07.2026

About the Organization (Vital Strategies India)

Vital Strategies: Institutional Overview and Strategic Brief

As an independent non-profit with a global reach, Vital Strategies is in the business of fortifying public health systems and confronting the world's foremost causes of morbidity and mortality. For the most part, the organization does its work in low- and middle-income countries, where it teams up with governments, civil society and international bodies to put evidence-based interventions in place and scale them up.

What sets Vital Strategies apart is how it brings together technical know-how in epidemiology and data with institutional mapping and policy design to ensure public health solutions are adopted on an equitable basis. Their approach is underpinned by a direct partnership with ministries of health and municipal authorities to alter the policy environment and encourage healthier habits. You will find that a great deal of their operational framework is built around optimizing Health Management Information Systems (HMIS) so decisions can be made on the strength of the data. They also run sizeable strategic communication campaigns to win over the public for protective regulations and to effect behaviour change at the population level, all while keeping a close eye on environmental and disease threats through active surveillance to lay the groundwork for regulatory reform.

Take their flagship work on Civil Registration and Vital Statistics (CRVS) as an example. In many parts of the world, poor CRVS infrastructure has been a major impediment to development; it leaves millions of births and deaths unrecorded and makes it hard for a government to put its health resources to good use. Vital Strategies is changing that by training health workers to properly record mortality figures and

standardizing the way cause of death is certified. This gives nations the real-time intelligence they need for everything from tracking disease burdens to long-term budgeting. When it comes to chronic illness, the organization is focused on the non-communicable diseases – cancers, diabetes, cardiovascular problems – that are responsible for the bulk of premature deaths worldwide. They don't shy away from the commercial side of health, tackling tobacco control and food policy head on. To get smoking rates down in line with the WHO's MPOWER package, for instance, they have designed mass media campaigns with cultural nuances in mind. They are also on hand to give governments the support they need to introduce warning labels or put restrictions on sugary drinks to stem the tide of obesity.

Then there is the matter of our planet. Vital Strategies has put down roots in environmental and planetary health on the understanding that climate change and degradation are only making global health crises worse. They do the legwork in terms of landscape mapping and community studies to see how things like ambient air pollution or toxic chemicals are having an impact. Whether it is monitoring air quality to inform urban planners or putting in place measures to shield children from industrial lead, these are the kinds of intersecting efforts in data and policy that allow Vital Strategies to build public health systems fit for the future.

Acknowledgements

The past six months has been incredibly intense and full of experience and reflecting on this I would like to extend my gratitude to the people who have guided me to end up with this completion day. This internship would not have been possible without the contribution of the following people.

First, I would like to express my gratitude towards IIHMR DELHI for providing me with the opportunity to work with Vital Strategies India and my mentor **Dr. Preetha G.S. ma'am** for incessantly guiding me.

I extend my sincere thanks to **Mr. LM Singh (Managing Director); Dr. Nidhi Chaudhary (Principal Technical Adviser); Ms. Ruchika Sharma (Executive Administrator); Ms. Jagriti Roy (Senior Operations Manager) and Mr. Chaitanya Mathur (Senior Manager, HR)** and all my fellow colleagues for their unwavering guidance and support throughout the internship have been instrumental in shaping my understanding of the industry. Their expertise and encouragement not only facilitated my learning but also helped me contribute meaningfully to the ongoing projects and have been instrumental in my development as a budding professional.

I would also like to extend my sincere thanks to **Ms. Shalini Negi (Consultant, Operations)** for her support in my internship project. A heartfelt thanks to Dr. Harshdeep Singh for his constant support throughout the dissertation. A special thanks to my family for supporting me.

Their collaborative spirit, patience, and willingness to involve me in meaningful projects have truly made a difference in my learning experience. I am grateful for their trust and the opportunities they provided me to contribute to the team's objectives. The positive

work environment fostered a sense of belonging and made every day an enjoyable and productive one.

I would also like to extend my appreciation to my fellow colleagues, with whom I had the pleasure of collaborating. Their enthusiasm and dedication made the internship experience even more enjoyable, and I am thankful for the knowledge exchange we had.

I look forward to utilizing this experience as a foundation for my future endeavours in the field public health.

Thank you

Table of Contents

Table of Contents

Introduction: 18

Objectives and key Research Questions: 22

Literature review: 24

Defining Artificial Intelligence in the Healthcare Context 24

Requirements of AI: Datacentres, Computing Infrastructure, and Environmental Degradation 25

India's Data Centre Growth: Digital Sovereignty and Spatial Vulnerabilities 27

Electricity and Water Required: Comparative Studies Across the US, the UK, Europe, and India 28

United States Studies 28

Studies from the United Kingdom and the NHS 28

European and EU Studies 29

India studies 30

Climate Change and AI: A Double-Edged Technological Intersection 32

Relevance of Climate Change and AI in Healthcare: The Climate-Health Paradox 34

Green AI Initiatives and Their Applicability to Healthcare 36

Algorithmic Simplification and Edge Computing 36

Algorithmic and Infrastructure Optimization: The "4 M's" Framework 37

Energy Consumption Disclosures and the Chatbot Labelling Paradigm 37

Collaborative and Infrastructure Efficiency 38

Policy, Governance, and the Path to Sustainable Healthcare AI 39

Methodology: 41

Results: 47

Objective 1: AI-Enabled Reductions in Healthcare Carbon Emissions 53

Objective 2: Lifecycle Environmental Costs of AI in Healthcare 54

Objective 3: Green AI and its Relevance to Healthcare 57

Objective 4: The Environmental Cost of Healthcare Data Centres in India 58

Objective 5: Framework for Sustainable AI in healthcare 59

Discussion: 61

1. The Carbon Benefit of AI in Healthcare: Real but Narrow and Contingent 61

2. Lifecycle Environmental Costs of AI: Massive, Growing, and Unreported 62

3. Jevons' Paradox and the Rebound Effect: The Most Underappreciated Risk 64

4. Green AI Initiatives: A Promising but Underutilized Field in Healthcare 65
5. The India Challenge: Global Risk in Microcosm and Data Centre Infrastructure 66
6. Governance, Policy and the Road to Net-Positive Healthcare AI 67
7.Towards an Integrated Framework for Sustainable Healthcare AI: Synthesis of Evidence 68
8.Limitations of This Review 70
Conclusion: 71
1. The climate benefits of AI in healthcare are real, if conditional 71
2. Lifecycle costs of environment are huge and underreported 71
3. Jevons' Paradox is the structural problem to solve 72
4. Green AI is here but not yet part of the healthcare fabric 72
5. India and LMICs face a compounded risk that demands urgent, context-specific policy 72
Final Remarks and Recommendations 73
References 74

Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
BODH	Benchmarking Open Data Platform for Health AI
CAGR	Compound Annual Growth Rate
CEEW	Council on Energy, Environment and Water
CF	Carbon Footprint
CNDCP	Climate Neutral Data Centre Pact
CO_2	Carbon Dioxide
CO_2e	Carbon Dioxide Equivalent
CT	Computed Tomography
DEA	Data Envelopment Analysis
DL	Deep Learning
DPDP	Digital Personal Data Protection (Act)
EEA	European Environment Agency
EED	Energy Efficiency Directive
EHR	Electronic Health Record
EPRI	Electric Power Research Institute
EU	European Union
GHG	Greenhouse Gas
GPT	Generative Pre-trained Transformer
GPU	Graphics Processing Unit
GW	Gigawatt
HMIS	Health Management Information System
ICT	Information and Communications Technology

Abbreviation	Full Form
IEA	International Energy Agency
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IPCC	Intergovernmental Panel on Climate Change
ISO	International Organisation for Standardisation
LCA	Life Cycle Assessment
LLM	Large Language Model
LMIC	Low- and Middle-Income Country
NLP	Natural Language Processing
PACS	Picture Archiving and Communication System
UNEP	United Nations Environment Programme
US	United States
USD	United States Dollar
WHO	World Health Organization
WUE	Water Usage Effectiveness

Research Project

Artificial Intelligence in healthcare: A Double-Edged sword for Climate Change

Introduction:

The healthcare filed is being shaped by advancements in technology, starting from robotics to machine learning (ML) or Deep learning (DL) and NLP. But all of this is centred around artificial intelligence (AI), which is rapidly changing the delivery of care on global scale. With AI in action, there is difference in how information is handled, decisions are made and outcomes are predicted; it allows better diagnostic accuracy and more personalised treatment while streamlining clinical work and opening remote monitoring (1,2,3). The buzz of AI in healthcare is visible in increasing market value- the global healthcare AI market, valued at USD 11.2 billion in 2023, is projected to reach USD 427.5 billion by 2032, reflecting a compound annual growth rate (CAGR) of 47.6%. (2). This increase in utilisation of AI is compounding the effect on climate change.

The Intergovernmental Panel on Climate Change (IPCC) has made it plain that unless we see deep and sustained cuts to greenhouse gas (GHG) emissions across the board, we will likely see global temperatures rise more than 1.5°C above pre-industrial levels in the 2030s with irreversible effects on human health and the economy. Healthcare systems are in a peculiar position as both cause and effect. A 2019 report from “Health Care without Harm” put the sector’s contribution to global GHG emissions at 4.4 per cent, the same as 514 coal-fired power plants would put out in a year. India’s system, seventh in the world, is responsible for some 38.8 million metric tons of carbon dioxide annually, or the output of 187 such plants (4). And as organisations like the English NHS set their sights on net-zero direct emissions by 2040, the quick adoption of AI is viewed with equal parts hope

and trepidation (5). In short, the sector and its AI tools are at once an instrument of survival and a factor in the climate equation.

This dissertation is framed by that very duality. One could not say AI in healthcare is a simple friend to sustainability, nor can it be dismissed as just another element of the climate crisis. It is something of a paradox: an infrastructure with its own considerable environmental price tag and at the same time a means to shrink healthcare's footprint. Some in the literature have called it a 'double-edged sword' when it comes to climate change (1,2,4).

AI can be a boon for the environment. Telemedicine does away with the need for patients to drive to a clinic, saving on transport emissions (4,7). In radiology and pathology, AI optimises workflows, so you do not have to repeat scans unnecessarily, which in turn curbs the energy use of the equipment (4). Predictive analytics put an end to waste in procurement and staffing (6), while virtual triage and wearables bring care to the community with less physical demand than the old way of doing things (7). In theory, autonomous AI could cut healthcare's GHG output by 80 per cent (4).

But flip the blade and the environmental cost is steep and mounting. You cannot train a large model without expending a great deal of energy. GPT-3, with its 175 billion parameters, is estimated to have used 1,287 MWh of electricity and produced 552 tonnes of CO_2e in the process, much like five passenger cars over their lifetimes (4,5). A question put to ChatGPT uses ten times the power of an ordinary Google search (1). Data centres are only going to become more voracious; in Ireland they are set to take up 35 per cent of national energy by 2026, and worldwide demand is poised to nearly double to 945 TWh by 2030 (1,3). They also guzzle water for cooling, with the industry using some 560 billion litres a year, enough for a city of 8 million (3). Adding to that the rare earth element

needed to build the hardware and the hazardous e-waste left behind by obsolete machines, and the picture is not entirely rosy (1,4).

What is missing from the current literature is a comprehensive, integrated look at two dimensions of AI in healthcare: its climate benefits and the lifecycle climate costs. You will not often find these examined together. For the most part, studies are content to focus on one or the other – be it the clinical and operational advantages of AI or the environmental price tag of its infrastructure in isolation. This is a notable void, especially considering the pace and scale with which health systems around the world are adopting AI, the quickening expansion of AI infrastructure in high-growth markets like India (where data centre capacity is set to hit 4.5-6.5 GW by 2030(3)), and the pressing nature of our international obligations under the Paris Agreement. There are no frameworks in place for an organisation to make a truly holistic net-climate-impact assessment when it comes to procuring and deploying AI.

The case for AI's benefits is well made. Telemedicine and remote monitoring powered by AI have cut down on patient travel and the carbon emissions that come with it.(6,7) In radiology and pathology, AI has been used to streamline workflows and avoid redundant imaging, saving energy. (6) Then there are the virtual consultations and chatbots that put healthcare within reach while keeping the physical delivery footprint small. (6,7) Yet the environmental ledger tells a different story. The costs across an AI's lifecycle are considerable and well-documented. Training large models is an exercise in energy intensity; GPT-3, with its 175 billion parameters, is thought to have guzzled some 1,287 MWh of electricity during training, putting out 552 tonnes of CO_2e – the lifetime output of five cars. (6,8) Even a query to an AI chatbot uses ten times the power of a regular web search. (1) By 2026, the data centres in Ireland alone are expected to draw nearly a third

of the country's total energy. (1) And it is not just electricity. Data centres need water in vast amounts to cool their servers; globally, AI infrastructure could soon be using six times as much as Denmark, a problem in regions already short on water. (1) On top of that, you have the e-waste from obsolete hardware laced with lead and cadmium, and the rare earth mining required for new equipment, which is hardly gentle on the environment. (1,6)

India is a prime example of why these matters. It is one of the world's fastest-growing data centre markets, with installed capacity tripling from about 520 MW in 2020 to 1.5 GW by mid-2025 and on track for 4.5–6.5 GW in 2030.(9) Last year, data centres in India were responsible for roughly 0.5 per cent of national electricity use and 150 billion litres of water, figures that will more than double by 2030.(9) Given that 72 per cent of Indian districts are prone to extreme flooding and 57 per cent to severe heat, the resilience of this AI infrastructure has real implications for public health and the climate.(9) So, while there is plenty of writing on the subject, it is usually in silos. Reviews will cover clinical outcomes, or they will do a carbon footprint analysis, but they do not bring them together to see the full picture of healthcare-specific AI. (1,5,6) Green AI tools and cloud optimisation are still in their infancy in the healthcare space and largely unexplored, (8) and standardised ways to measure climate impact are for the most part non-existent. (5)

This study takes a systematic approach to the dual nature of AI's environmental impact in healthcare, weighing the gains from telemedicine and workflow efficiencies against the hard costs of energy, water and e-waste. Drawing on health informatics, computer science and environmental science, this study will put forward an evidence-based framework for sustainable AI adoption.

Objectives and key Research Questions:

Aim: To put the dual impact of Artificial Intelligence in healthcare on climate change under the microscope – looking at both the benefits and the lifecycle costs of AI-driven systems – and to put forward a framework for its sustainable use.

Research Question:

1. What is the carbon footprint savings to healthcare derived from using AI-powered remote monitoring and telehealth?
2. What are the lifecycle costs of AI in energy and water terms?
3. Do the environmental benefits of healthcare AI outweigh their lifecycle costs and in which situations do the application of AI net benefit the climate?
4. What green AI strategies or policies are on the table to help reduce the environmental impacts and are these policies effective?
5. What sort of governance do we need to have in place so that we know that the application of AI in healthcare helps us to reach our climate targets, not the opposite?

Primary objective:

1. To see the evidence where AI has cut down on emissions be it through less patient travel or fewer unnecessary procedures and wasted resources.
2. To quantify the environmental price tag of AI in healthcare, from the critical minerals it depends on to the e-waste it produces.

Secondary objective:

1. To put existing green AI initiatives to the test for their relevance in healthcare and note any shortcomings in how they are reported.

2. To evaluate data centre infrastructure in fast-growing markets like India to gauge its environmental effect.
3. To put forward a conceptual framework to help policymakers and organisations in the sector strike the right balance between the costs and the climate benefits of taking on AI.

Literature review:

Defining Artificial Intelligence in the Healthcare Context

The global healthcare sector is in the middle of a rapid technological change. The industry has been integrated with algorithms, from machine learning to computer vision, deep learning, and natural language processing (NLP) (10, 11). Once a niche research area, AI is now an infrastructural requirement - as embedded in clinical and operational workflows as is any other component (10). Modern algorithms are expected not only to provide remote patient care or customized therapy but also to eradicate human labour and improve the accuracy of diagnosis (10). YOLOv5, for example, is a computer vision network used in resource-limited areas for early diagnosis of cervical cancer and thus ensuring the standardization of this process while removing delays (12).

In recent years, there has been the introduction of a new cohort of medical AI models like ChatGPT, Gemini, and Copilot which continue to move the industry closer toward clinical implementation (13, 14). It is on tools like these that we anticipate will perform many of the back-end functions, like writing discharge instructions, handling medical billing, or summarizing chart notes. These tools have been reflected in a March 2024 American Medical Association (AMA) report, which shows 78% more physicians using AI in daily practice than the previous year, and two-thirds having already integrated the technology (13). Commercial adoption is reflected in the fact that in 2023 the global healthcare AI market was valued at USD 11.2 billion, a figure projected to reach USD 427.5 billion by 2032, with a CAGR of 47.6% (11).

In contrast, serious discussion about environmental impact remains scarce among these growth projections (13, 14). Whereas academic medical publications are filled with references to patient privacy, the precise performance of algorithms, and workflow

efficiency, little regard has been paid to the life-cycle impact of the computing systems that underpin them (13).

Requirements of AI: Datacentres, Computing Infrastructure, and Environmental Degradation

Underneath the application of AI in healthcare, is a physical infrastructure of data centres, network capacity and specialized hardware. (10,14). Modern deep learning is costly, and this includes a large need for water and electricity for temperature control, as well as e-waste (14). In fact, the environmental footprint is embedded as soon as the hardware comes into being. Producing the GPUs and application-specific integrated circuits used in the models necessitates extracting the rare earth elements and critical minerals required to make these chips, a practice often associated with carbon intensive and environmentally damaging supply chains and mining practices (10, 14).

In the tech world, servers and GPUs have a shelf life of perhaps three to five years before they are cycled out, adding to the mountain of e-waste (10, 14). Without proper circular recycling, risk hazardous materials like lead and cadmium seeping into local ecosystems, a long-term hazard for public health and the environment alike (10, 14).

Once the processors are up and running, they put out a lot of heat. Old style ambient air-cooling won't cut it for a modern GPU cluster, so data centre operators have had to turn to more advanced liquid cooling solutions, be it immersion or direct-to-chip (15, 16).

To put the scale of this in perspective, Belkhir and Elmeligi (17) have modelled the macroeconomics of the ICT sector and found that the global digital footprint, which was 1.6% of greenhouse gas emissions in 2007, could account for as much as 14% by 2040. Data centres are the fastest growing part of that, climbing to an estimated 45% of the total ICT footprint by 2020 thanks to the overhead of cloud software (17). Green IT

frameworks are calling for a response to this expansion. Their taxonomy for greening the whole lifecycle of an IT asset is broken down into three domains: green hardware like virtualization, green software for better query efficiency, and the physical infrastructure of the data centre itself (18).

The 2010s saw a huge digital expansion, yet global data centre energy demand stayed put at around 200 terawatt-hours (TWh) a year, or some 1% of the world's electricity grid. The reason for this flat line was a swift move to hyper-efficient, centralized "hyperscale" set-ups (19). In an old-fashioned localized server room, you would find that overhead such as cooling and power distribution was as energy hungry as the servers, for an inefficient Power Usage Effectiveness (PUE) of close to 2.0 (19). But by offloading work to hyperscale facilities, that overhead is cut to 10-20% and the PUE falls to 1.1 or 1.2 (19).

Then there is the matter of Jevons' Paradox. As clinical use widens, gains in algorithmic or computational efficiency can have a rebound effect: AI becomes faster and less costly, which in turn fuels an exponential increase in net resource consumption (20). On the waterfront, these centres are thirsty. Some 80% of the freshwater they pull in is lost to evaporation in on-site cooling; what is left is a salty, concentrated "blowdown" that must be chemically treated before it can go into the municipal sewers (21, 22). In the U.S., an estimated 80 to 90 per cent of a data centre's water comes from public potable sources, effectively passing the private cost of computing on to local utilities and ecosystems (21). While new thermal management tech can ease the strain on energy, it tends to concentrate local freshwater needs (16, 21), making it clear that data centres are high-intensity resource hubs with a considerable footprint in terms of land, water and power (21).

India's Data Centre Growth: Digital Sovereignty and Spatial Vulnerabilities

Rapid digitalization and surging mobile data use have made India one of the world's most dynamic data centre markets, not to mention national efforts like the IndiaAI Mission (23). Installed capacity has almost tripled in the last few years, from 520 MW in 2020 to 1.5 GW in mid-2025 and is expected to hit between 4.5 and 6.5 GW by 2030 (23). There is plenty of capital behind this, with some $95 billion committed from 2019 to 2025 and figures set to top $100 billion by 2027 (23). The central government has been keen to foster this, giving data centres infrastructure status and offering foreign operators tax holidays that run to 2047 (16, 23).

The infrastructure is heavily clustered in the big metros (23). Mumbai is running quarter of facilities in India. It has left Chennai, Hyderabad and Bengaluru behind. The 2023 DPDP Act is one reason for this; as a matter of digital sovereignty, it restricts the cross-border transfer of any sensitive personal or medical data (16). These kinds of localization mandates mean the data must stay put, but they also put a strain on resources in a few selected urban areas (16, 23).

Then there is the question of climate risk, something even these metropolitan clusters can no longer ignore. In India 72 per cent of districts are vulnerable to bad flooding and 57 per cent to extreme heat, so the physical safety of a centralized healthcare database is an issue that is becoming harder to overlook (10, 23).

Electricity and Water Required: Comparative Studies Across the US, the UK, Europe, and India

The energy and water footprints of data centres change from one part of the world to another. That is only to be expected given the disparities in regional climate, what kind of cooling is in use and how the grid is made up.

United States Studies

The data centres in USA used 4.4% of the nation's electricity in 2023 and there are studies that this figure can go to 12% by 2028 (21). The AI boom is a factor here; demand for AI servers alone is forecast to be in the order of 150-300 TWh by 2028, entailing a water footprint of 150 to 280 billion Liters a year (21). Much of that power goes to the cooling towers that evaporate freshwater to service high-density racks.

Take Google for instance: in 2023 their own centres pulled 29 billion Liters of freshwater and evaporated 23 billion, with close to 80% of it coming from municipal systems (21). On a smaller scale, a GPT-3 inference might use up a 500 mL bottle of water for every 10 to 50 queries, depending on how the model is deployed. Researchers at UC Riverside put the cost of a typical 100-word prompt at 519 mL. It is a problem in some parts of the country where roughly 20% of data centres are sited in watersheds under drought stress, leading to friction with communities over scarce water (21).

Studies from the United Kingdom and the NHS

The National Health Service in the UK has set its sights on net-zero direct emissions by 2040 and indirect ones by 2045, yet digital carbon footprints are for the most part left out of the standard accounting (10, 24). There are some localised digital tools that have shown their worth in this regard; the Wayfinder function on the NHS App is expected to put back more than 1,100 tonnes of CO2e a year in savings by making paper letters

redundant (25). But one must look at the wider digital shift to see the hidden costs. Take primary care: survey data suggests that the use of ChatGPT alone could be responsible for some 349 tonnes of CO2e every year, an unmeasured footprint from generative AI that has not been procured in any formal sense (26).

- NHS App Wayfinder: 1,100+ tonnes CO2e / year in savings (25)
- Primary Care ChatGPT: ~349 tonnes CO2e / year emitted (26)

The reason for such reporting omissions is that standard practice will often leave Scope 3 supply chain emissions from hardware out of the equation unless the NHS trust is the direct owner of the physical plant (26).

European and EU Studies

Data centres in Europe make up roughly 3% of the EU's electricity grid (2% for the continent), and between the US, China and Europe they are behind 85% of the world's data centre energy demand (27). The consumption is not evenly spread, though. In Ireland it can top 20% of the local grid (27). When data centre power use hit about 21% of the national total there, the swift growth of hyperscale sites prompted a temporary halt on new grid connections (22). Now, under the rules, Irish operators must get at least 80% of their power from home-grown renewables (22). The Netherlands has gone so far as to put a ban on new hyperscale development over land-use and power capacity issues (22).

Then there is the matter of water, which has caused no end of public and regulatory trouble for European digital infrastructure (22). Microsoft was roundly criticised in the Netherlands in 2021 for using 84 million litres of water, well above what it had put in its projections of 12 to 20 million (22). While data centres are a drop in the ocean compared to the water taken by agriculture or for cooling power plants in Europe, the

local hydrological effect can be hard on the area (28). To put a curb on it, the Climate Neutral Data Centre Pact (CNDCP) has stipulated that any new build in a water-stressed region after 2025 must keep its Water Usage Effectiveness (WUE) to a hard limit of 0.4 L/kWh (29).

Figures across the EU show an average Power Usage Effectiveness (PUE) of 1.36 and a WUE of 0.58 cubic meters per megawatt-hour (30). Colocation outfits do much better with an average WUE of 0.31 L/kWh (31). The EU Energy Efficiency Directive (EED) is also forcing facilities of over 1 MW to put their low-grade waste heat to work in municipal heating or for thermal water purification (32, 33). Some thermodynamic models even suggest that if you used this waste for desalination or to capture CO2 directly, you could make a data centre water-positive, yielding 0.5 kg of freshwater and removing the same in CO2 for every kWh of computing done (34).

India studies

In India the numbers are telling in 2024 data centres took in 150 billion litres of water and 0.5% of the country's electricity, and we expect both to more than double before 2030 (14). The big players with over 5 MW tend to run evaporative coolers, and you can lose 80% of your operational water to evaporation (23, 35, 36).

According to Karnataka's IT Ministry in March 2026, a single megawatt of capacity in the state calls for 25 million litres a year (35). That kind of demand is a recipe for conflict in the dry season. When Bengaluru faced a drinking water deficit in 2024, the municipal board saw fit to reduce supplies to bulk users like the IT parks by 20% (35). Siting these operations in urban areas where water is already scarce only heightens the risk. A case in point is the 1 GW data centre coming online in Visakhapatnam in a district with little groundwater (2.12 TMC); the environmental papers say nothing of what the facility will consume once it is running (37).

As an offset to the demands of physical infrastructure, India's public health system has turned to telemedicine to avoid the carbon cost of moving patients (38, 39). eSanjeevani, the national platform, is a prime example of this. It works on a hub-and-spoke basis to link Ayushman Bharat Health and Wellness Centres with specialists (38, 39) and a patient-to-doctor outpatient department (eSanjeevaniOPD), the platform allows millions of citizens to bypass physical travel, reducing fuel consumption and waiting times (38, 39). In parallel, localized Indian studies have begun to quantify these digital offsets. A 2024 study by the L.V. Prasad Eye Institute (LVPEI) across four southern Indian cities found that teleophthalmology saved rural patients an average of 80 km of travel and 2.89 kg of CO2 emissions per consultation, while urban tertiary-care consultations saved an average of 1,666 km of travel, 76 Liters of fuel, and 176.6 kg of CO2 emissions over a three-month period (40). While these public digital goods deliver clear environmental offsets, their long-term sustainability is ultimately bounded by the power grid and cooling water footprints of the underlying data centre systems hosting them (10, 23).

Country / Setting	Energy Footprint Metrics	Hydrological Consumption Metrics	Primary Reporting & Operational Gaps
United States (21)	4.4% of national grid (2023); projected to reach 12% by 2028. AI servers: 150-300 TWh by 2028.	On-site data centre water consumption: 150-280 billion Liters by 2028. Google: 23 billion Liters evaporated (2023).	Limited transparency in public reporting; Scope 2 and 3 water footprints are rarely disclosed.
United Kingdom (NHS) (24, 25, 26)	Invisible in Green Plans. Wayfinder App: less than 0.1 tCO2e. Primary care ChatGPT: ~349t CO2e/year.	Large-scale LLM integrations (e.g., ambient scribes) increase cloud cooling loads.	Current carbon reporting omits hardware lifecycle emissions and unprocured generative AI.
Europe (EU) (27, 28, 29, 30, 31, 32, 33, 34)	3% of EU grid (2024). Average PUE of 1.36. Localized demands	Average WUE of 0.58 L/kWh. CNDCP target maximum of 0.4 L/kWh in water-stressed zones.	EED mandates annual KPI reporting for facilities >500kW. However, Article 5 contains a confidentiality clause shielding facility-

	reach over 20% in Ireland.		specific metrics from the public.
India (23, 35, 37, 38, 39, 40)	0.5% of national grid (2024); projected to more than double by 2030.	150 billion Liters in 2024; projected to double by 2030. Each MW requires 25 million Liters/year. eSanjeevani and LVPEI show travel offsets of up to 1,666 km / 176.6 kg CO2.	Most state policies lack water-stress mapping; environmental clearances omit operational water usage.

Climate Change and AI: A Double-Edged Technological Intersection

There is an inherent tension in the relationship between artificial intelligence and climate change, one defined by a trade-off between what AI can do and the environmental price to be paid (10). Put simply, the operational benefits of large-scale AI models come with substantial energy and water costs. Take GPT-3 as an example: training its 175 billion parameters took some 1,287 MWh of electricity and put 552 tonnes of CO2e into the atmosphere – on par with the lifetime emissions of five passenger cars (10).

The power demands are only set to grow exponentially. Where historical models like GPT-3 were already energy intensive, frontier AI training has moved into a new league. Projections from EPRI and Epoch AI put the year-on-year growth in peak power for the largest models at 2.2x, with forecasts calling for an expansion of 2.2x to 2.9x annually through 2030 (41). By then, a single run of a frontier model could pull 4 to 16 GW of power, enough to supply millions of homes. Globally, we may see AI capacity top 100 GW by 2030, with the U.S. alone accounting for over 50 GW of that, or about 5 per cent of the country's total generation (41).

Even as the best GPUs get more efficient, improving some 40 per cent a year in line with Koomey's Law, the sheer scale of compute – up 4x to 5x a year with training runs now stretching past 100 days – is putting a strain on local grids. Companies are having

to look at distributed training across several data centre campuses to get around power limits. At present, the demand for power is divided evenly among inference, experiments and training (41).

The numbers are stark. In a landmark piece for Joule, Alex de Vries estimates that by 2027 global AI will need between 85.4 and 134.0 TWh of electricity a year, much of it to support the manufacturing of high-intensity hardware (19). You would need enormous energy grids just to ship 1.5 million AI servers. To put it in perspective: if Google were to put generative AI behind every search query, it would take more than 512,821 dedicated servers and 29.2 TWh of electricity annually, which is all of Ireland's national consumption (19).

There is the matter of water. By 2027, the global AI footprint is expected to include 85 to 134 TWh of electricity use and a withdrawal of 4.2 to 6.6 billion cubic metres of water, with 0.38 to 0.60 billion cubic metres lost to evaporation (21). That is four to six times what Denmark withdraws in a year.

2027 Projected Global AI Footprint

- Electricity: 85 to 134 TWh (21)
- Water Withdrawal: 4.2 to 6.6 billion cubic metres (21)
- Water Consumption: 0.38 to 0.60 billion cubic metres (21)

Yet AI is not without its utility for climate adaptation and mitigation (24). Machine learning can make sense of complicated climate data, fine-tune renewable energy systems and give us better monsoon forecasts (23, 35). But this can also lead to negative feedback loops. India's Meteorological Department, for instance, uses AI to predict water scarcity and monsoons; the data centres that house these models must be cooled with water. When El Niño brings a heatwave and temperatures climb, so does the water

need for cooling (35). In a way, the technology meant to monitor environmental stress ends up depleting the very resources it is tracking.

Relevance of Climate Change and AI in Healthcare: The Climate-Health Paradox

The healthcare industry is no stranger to this issue. It is a major source of the climate-driven health risks it is supposed to address, a paradox (42). Globally, the sector is behind 4.4% to 4.6% of greenhouse gas emissions (10). In the U.S. the figure is 8.5%. India is seventh in the world for its carbon output from healthcare, at 38.8 million metric tons of CO2 a year (10, 23).

Global Healthcare Sector Carbon Footprint

- Global Average: 4.4-4.6% of total GHG (10)
- United States: 8.5% (10)
- India: 38.8m metric tons of CO2 (23)

Sustainable deployment of digital health can go some way to offset this. Telemedicine and remote monitoring cut down on patient travel and the emissions that come with it (10, 43).

Telemedicine is an effective way to cut carbon emissions. A systematic review put the CO2e savings at anywhere from 0.70 kg to 372 kg for a single consultation, with the figure hinging on what kind of vehicle is used and how far one must drive (43). Take the work of Connor et al., for example: their study showed virtual care shaves off an average 15.0 km of travel per patient. In a public transit situation that means 0.70 kg less in CO2e, but if the patient would have driven his own car, the reduction is 2.93 kg

(44). Then there are telephone consultations; another study puts the average travel saved at 39.3 km and the resulting drop in emissions at 8.05 kg per appointment (44).

The scale of the impact can be seen in a UCLA-led study which calculated that telemedicine trims 21.4 to 47.6 million kg of carbon dioxide from US monthly emissions, the same as taking 130,000 gas guzzlers off the road (45).

One of the more striking cases of sustainable digital health is the Aravind Eye Care System in southern India (46, 47, 48). To serve 8.5 million people in Tamil Nadu and Puducherry, Aravind has put in place 100 satellite vision centres (46). With minimal electricity and diagnostic gear made locally by Aurolab, these centres use telemedicine to put rural patients in touch with ophthalmologists at tertiary hospitals. The 2021 figures show what the network has accomplished. It has spared 8,553 patients to take the journey, a move that meant 54.3 tonnes of CO2e were never put into the atmosphere (46, 47).. In all, the initiative has done away with 60 million person-miles of travel and 1,050 metric tons of carbon a year (48).

Aravind Eye Care Telemedicine Setup (Thiruppuvanam Centre)

- Setup Cost: 16,500 USD (37)

- Run Cost: 7,800 USD / year (37)

- Trips Avoided (2021): 8,553 (46)

- Carbon Saved (2021): 54.3 tCO2e (46)

Even their cataract surgeries are lean. An environmental audit of the operating rooms found that a phacoemulsification produces just 6 kg of CO2e and 250 grams of waste (49). That is some 5% of the footprint you would see for the same procedure in the UK, yet clinical outcomes are on a par thanks to surgical reuse and less waste (49).

But one must weigh the advantages of digital care against the hydrological and carbon costs of the AI that underpins it (10). A large academic health system running an AI messaging tool for 3,000 physicians at 50 messages apiece over a year will generate about 48,000 kg of carbon dioxide (42). Put another way, it takes 2,300 tree-years to sequester that much (at roughly 21 kg of CO2 per tree per year) (42). Unless the tool is cutting down on physical visits or redundant diagnostics, it is simply adding to the health system's carbon load (42).

Green AI Initiatives and Their Applicability to Healthcare

There is green AI initiatives aimed at curbing such environmental costs by making models more energy efficient over their lifecycle (14, 50). The NHS has been piloting the Green Algorithms framework, for instance, so clinical leads and developers can get an estimate of the emissions from an NLP model before they go live (50). Or you have libraries like CodeCarbon for tracking energy and carbon in real time while training (51).

Algorithmic Simplification and Edge Computing

The energy bill can be lowered through optimization. In a look at the Kidney Tumour Segmentation (KiTS-19) data, researchers pitted Standard, Depth wise and Group Convolutions against each other with mixed-precision arithmetic (52). They concluded that Depth wise Convolution with Mixed Precision was the most frugal on energy without sacrificing clinical performance, while Group Convolution was too heavy on operational overhead (52). For certain diagnostic tasks, going with something simpler like Boost or logistic regression can slash emissions by up to 99% versus a deep learning model (52).

Where resources are tight, compressed models present an alternative. MobileNet-Lite-Health (MLH), for example, runs on PathMNIST and HAM10000 datasets with 84.7%

and 80.1% accuracy respectively and very little latency (65 Ms) or memory (95 MB) (53). As an edge architecture, it uses 74% less energy and spits out 76% fewer carbon emissions per inference than a cloud-hosted behemoth like GPT-4V or MedPaLM (53).

MobileNet-Lite-Health (MLH) Edge Savings (vs. Cloud LLMs)

- Energy Consumption: Down 74% per inference (53)

- Carbon Emissions: Down 76% per inference (53)

Algorithmic and Infrastructure Optimization: The "4 M's" Framework

For a way to put some structure on the problem of mitigating the footprints of administrative and diagnostic AI, researchers have come up with multidimensional models (54). Patterson et al. in Joule have laid out the "4 M's" of sustainable machine learning – Model, Machine, Mechanization and Map – as the main areas where health systems and developers should be looking to optimize. throughout the model training lifecycle (54):

Put these four levers together (Model x Machine x Mechanization x Map) and a developer can lower total energy consumption 83-fold and carbon emissions by as much as 747 times (54). The numbers tell us that even as the raw compute for state-of-the-art models scales exponentially, the net carbon footprint of training can be made to plateau and then recede through the disciplined use of the 4 Ms. For healthcare administrators, it is a sound framework to judge the environmental credentials of an AI vendor at procurement time (54).

Energy Consumption Disclosures and the Chatbot Labelling Paradigm

Researchers have put forward the idea of applying green policy mechanisms to AI systems and chatbots in much the same way we rate household appliances for efficiency, to get healthcare teams into more sustainable digital habits (55). We used to

think of sustainability in terms of the heavy lifting of model training, but new data shows that inference – the day-to-day business of serving a model to a clinician – makes up some 60% of an AI's lifetime energy bill. That makes the point of user interaction a prime place for intervention (55).

With an "Energy Consumption Disclosure" (ECD) paradigm, software can be fitted with standardized labels on the clinical dashboard that spell out water footprint, carbon intensity and so on. In one randomized trial of 365 people, we found that having ECD labels on hand made it 12 times more likely a user would go for a lean, efficient Small Language Model in preference to a carbon-heavy LLM (55). This kind of labelling in an EHR portal or diagnostic workstation gives clinical leads the means to keep their computational loads in check and within regional boundaries without any loss of diagnostic power.

Collaborative and Infrastructure Efficiency

Federated learning in the field of research is not wasteful. For example: MRI-to-CT conversion models has shown reduction in training time and emissions by up to 23%, by keeping encoder weights steady where there is not much to change (51). Hospitals should make better use of their imaging infrastructure (56).

A radiology department could put in place a few straightforward things like:

For one, they might as well switch off the scanners at night rather than let them run on high-energy idle. Or they can employ smart scheduling to batch the scans and do with less.

Radiology departments can adopt some simple measures:

– Switch off the scanners at night.

– Smart scheduling to batch of the scan.

– Put sub-meters on PACS racks and scanner rooms for a true picture of electricity use.

– Make use of teleradiology to spare clinicians the commute.

– Set up your digital archiving to avoid the need for redundant backups.

Policy, Governance, and the Path to Sustainable Healthcare AI

There is a need of strong policy to manage the environmental side of healthcare AI, but there is still a wide governance chasm between what technology is being adopted and how it is overseen (14, 26).

The landscape is not uniform. The EU's AI Act is all about risk and individual protections, while the US and UK have taken a more flexible, innovation-first tack with no binding rules on sustainability (57). But the ISO/IEC has now made the water footprint a key part of its first international standard for sustainable AI, which should force developers to be more open about runtime water use (21).

Then there are the reporting gaps in national health systems. The English NHS Green Plans, for example, leave out large swaths of emissions in their mandatory accounts (26).

- **Unprocured "Shadow AI" Usage:** Widespread, informal clinical use of unprocured generative AI models is unmeasured and unaccounted for in system-wide emissions inventories.
- **A Lack of Granularity:** they don't tend to include the upstream manufacturing costs of specialized hardware like GPUs in their Scope 3 reporting unless the physical equipment is directly owned by the trust.

Procurement as a Policy Lever: Big health systems have the clout to make green procurement happen by putting their purchasing power to work (14). NHS England is a case in point: it has made it compulsory to apply a 10% net-zero and social value weighting on any new procurement (24). Then there is the Cabinet Office Procurement Policy Notice 006, which requires all suppliers vying for contracts of over 5 million GBP a year to put forward a thorough carbon reduction plan; this rule will be extended in proportion to all procurements (58).

That said, you still won't find independent, product-level carbon footprint figures for things like digital systems or medicines in the standard carbon databases (59), so it is not easy for those in charge of procurement to make an apples-to-apples comparison of medical technologies (59).

The situation in India is different. There is no single national data centre policy with teeth; instead, development is steered by industrial incentives at the state level (23). While 15 states have active policies, just five of them have any sustainability clauses, and environmental water stress mapping is not a prerequisite for clearance anywhere (35).

On the AI front, the Union Health Ministry has rolled out the SAHI strategy and the BODH platform under the IndiaAI Mission to ensure secure, ethical validation (60). But for the most part these are concerned with clinical accuracy and issues of algorithmic bias or data privacy. They do not really get into the weeds of resource-efficiency or environmental metrics (60).

Methodology:

Study Design: The study will employ systematic literature review, integrating quantitative synthesis of environmental impact data with qualitative thematic analysis of policy documents, governance frameworks, and expert perspectives.

Reporting guidelines: This study follows the protocol version of the Preferred Reporting Items for Systematic reviews and Meta- Analyses guidelines (i.e., PRISMA- P)30 (78). In this systematic review, studies published globally based on sound methodology as per specified inclusion and exclusion criteria will be included. In addition, grey literature studies and reports will be included in the review(78).

Eligibility criteria:

The systemic review follows the PI/ECO (ECO, exposure, comparator and outcome) protocol for systematic review research. Table 1 represents PICO description.(78)

Table 1: Population, exposure, comparator and outcome description

Domain	Inclusion Criteria	Exclusion Criteria
Population/ Setting	Healthcare systems, hospitals, diagnostic centres, and digital health infrastructures (including data centres) globally, with a specific interest in LMICs/expanding markets like India.	Non-healthcare AI applications (e.g., finance, autonomous vehicles) that do not have a direct link to health service delivery or clinical outcomes.
Intervention/ Exposure	**AI-driven healthcare systems:** Machine learning models, diagnostic AI, telemedicine powered by AI, resource management algorithms, and the physical infrastructure (hardware/servers) required to run them.	Traditional healthcare digitization (e.g., basic Electronic Health Records) that does not involve Artificial Intelligence or Machine Learning components.
Comparator(s)	**1. Pre-AI Baseline:** Traditional care models (e.g., in-person visits vs. AI-triage).	Studies that only report clinical efficacy (e.g., "AI is 90% accurate") without any

	2. Dual-Impact Comparison: Carbon savings (benefits) vs. operational energy/resource costs.	mention of resource use, travel, or environmental metrics.
Main Outcome(s)	**1. Environmental Benefits:** Reductions in Green House Gas (GHG) emissions, saved patient/provider travel miles, reduced medical waste, and optimized energy use. **2. Environmental Costs:** Lifecycle assessments (LCA), electricity consumption in KWh, water cooling requirements for servers, e-waste, and mineral depletion.	Studies focusing exclusively on the economic costs of AI without linking them to environmental sustainability or carbon footprints.
Study Types	**Primary Research:** Observational studies, lifecycle assessments (LCA), case studies, and environmental audits. **Secondary Research:** Systematic reviews, grey literature from health/tech policy bodies, and existing green AI frameworks.	Editorials, opinion pieces without data, and studies not available in English (unless a high-quality translation exists).

Search Strategy:

In this study extensive systematic research- conducted across PubMed, Web of Science, Embase, and Google Scholar. Grey literature sources including UNEP reports, IEA publications, CEEW white papers, and NHS sustainability reports. The time duration of search of publications is from January 2015 to the present. The search strategy for the published literature will be developed using a pearl growing/citation mining search strategy.

The search strategy is organised into two types of term groups:

Groups 1–4 represent the core **topic domains** to include: (1) Artificial Intelligence, (2) Healthcare Systems, (3) Climate Change & Emissions, (4) Data Centre Infrastructure

Groups A–E represent the **outcome/methodology domains**: (A) Lifecycle Assessment, (B) Environmental Costs (C) Carbon Footprint Measurement, (D) Environmental Benefits & Emission Reductions, and (E) Policy, Governance & Reporting Frameworks.

Similar search terms within each group are connected by the **OR** operator to expand the search field. Groups are then combined using the **AND** operator to focus results at the intersection of all relevant domains.

Table 1: Search Terms used

#	Search Terms
1	"AI" OR "Artificial Intelligence" OR "ML" OR "Machine Learning"
2	"Healthcare" OR "Health Care" OR "Hospital" OR "Clinical Setting" OR "Telemedicine" OR "Telehealth" OR "Remote Patient Monitoring"
3	"Climate Change" OR "Greenhouse Gas" OR "GHG" OR "Carbon Emission" OR "CO2 Emission"
4	"Data Centre" OR "Data Centre" OR "Server" OR "Cloud Computing" OR "AI Infrastructure" OR "Computing Infrastructure" OR "Hardware" OR "E-waste" OR "Electronic Waste"
A	"Life Cycle Assessment" OR "Life Cycle Analysis" OR "Lifecycle Assessment" OR "Lifecycle Analysis" OR "LCA"
B	"Environmental Assessment" OR "Environmental Sustainability" OR "Impact Assessment" OR "Environmental Cost" OR "Water Footprint" OR "Energy Consumption"
C	"Green AI" OR "Sustainable AI" OR "Green Computing" OR "Energy-Efficient AI" OR "Sustainability Measurement" OR "Green Performance Indicator"
D	"Carbon Footprint" OR "Climate Neutral" OR "Resource Consumption" OR "CO2 Equivalent" OR "CO2e"
E	"Life Cycle" OR "Lifecycle" OR "End-of-Life" OR "Material Use" OR "Hardware Waste" OR "E-waste Generation" OR "Sustainable Procurement" OR "Eco-Design"

Inclusion criteria: Articles published in English and peer-reviewed; studies reporting quantitative environmental impact data related to AI in healthcare; policy documents and

grey literature from authoritative bodies; and studies examining green AI frameworks relevant to healthcare.

Exclusion criteria: studies that are published before 2015, opinion pieces which doesn't have empirical support, or those article that are concerned only with AI performance without any mention of environmental data.

Selection Process:

The studies were searched and had two reviewers [PT and SN] who independently go over any studies that met the inclusion criteria. To screen titles and abstracts from our literature searches on PubMed/MEDLINE, Web of Science, Embase and Google Scholar, each reviewer made a separate determination as to a study's eligibility. We put Rayyan.ai to work for the systematic screening of records, text mining and weeding out duplicates; with its help we could label records for inclusion, exclusion or "maybe". Mendeley was run in tandem to handle reference management and catch any further duplicates.

Once title and abstract screening was done, PT and SN pulled up their reports and independently read through the full texts of all studies that looked eligible. They measured these against the PICO framework laid out in our protocol, looking at population, intervention or exposure, comparator and outcome. If there were any differences of opinion between the reviewers at this point or during the earlier screening, we would put it to discussion and come to a consensus. For every study we put aside at the full-text stage we made a note of the justification and the reason for its exclusion. The whole process of search and selection is set out in a PRISMA 2020 flow diagram, which gives a breakdown of how many records were identified, screened and ultimately made it into the review.

Data Extraction and Data Collection:

Two reviewers (PT and SN) went about extracting data from every study we included, working independently with a Microsoft Excel spreadsheet the team had put together in advance. Before we did the full extraction, we put the form to the test on three sample studies and made any necessary changes so it would be consistent and reproducible.

The excel sheet was developed to capture every piece of data for every study: the authors, year, geography, methodology and design, what AI technology and which clinical area and the environmental outputs.

We then extracted data for any study that aligned with Primary Objective 1, which was that studies demonstrate the environmental impacts of AI technologies within healthcare, and recorded information such as: Greenhouse gas reductions; Reductions in patient travel; Energy efficiency improvements due to optimising workflows; reduction in unnecessary procedures or waste. We recorded these in the units given, be it CO_2e, kWh or miles. For the other side of the coin, Primary Objective 2, we looked at the lifecycle costs of AI. That meant extracting data on training and inference energy, water and electricity use in the data centre, e-waste and the need for rare earth minerals, standardising the units to CO_2e, litres, kilograms and so forth as they were reported.

For studies addressing green AI frameworks and standardisation (Secondary Objective 3), data were extracted on the specific initiatives described their applicability to healthcare contexts, and any identified gaps in reporting or standardisation. For studies addressing data centre infrastructure with reference to India or other rapidly expanding markets (Secondary Objective 4), data were extracted on installed capacity, energy and water consumption figures, projected growth trajectories, and climate vulnerability context. For studies contributing to policy and governance frameworks (Secondary Objective 5), data were extracted on proposed or existing policy interventions, governance structures,

sustainability reporting requirements, and recommendations for healthcare organisations and policymakers.

Any discrepancies in data extraction between the two reviewers were resolved through discussion, with reference back to the original source text. Where data were reported in formats that required conversion for comparability, standardised conversion factors were applied and documented transparently within the extraction sheet.

Results:

Following title and abstract screening, 18 studies that met inclusion criteria then undergone full text screening against the PICO framework. All 18 studies were included based on relevance to at least one primary or secondary objective of the dissertation.

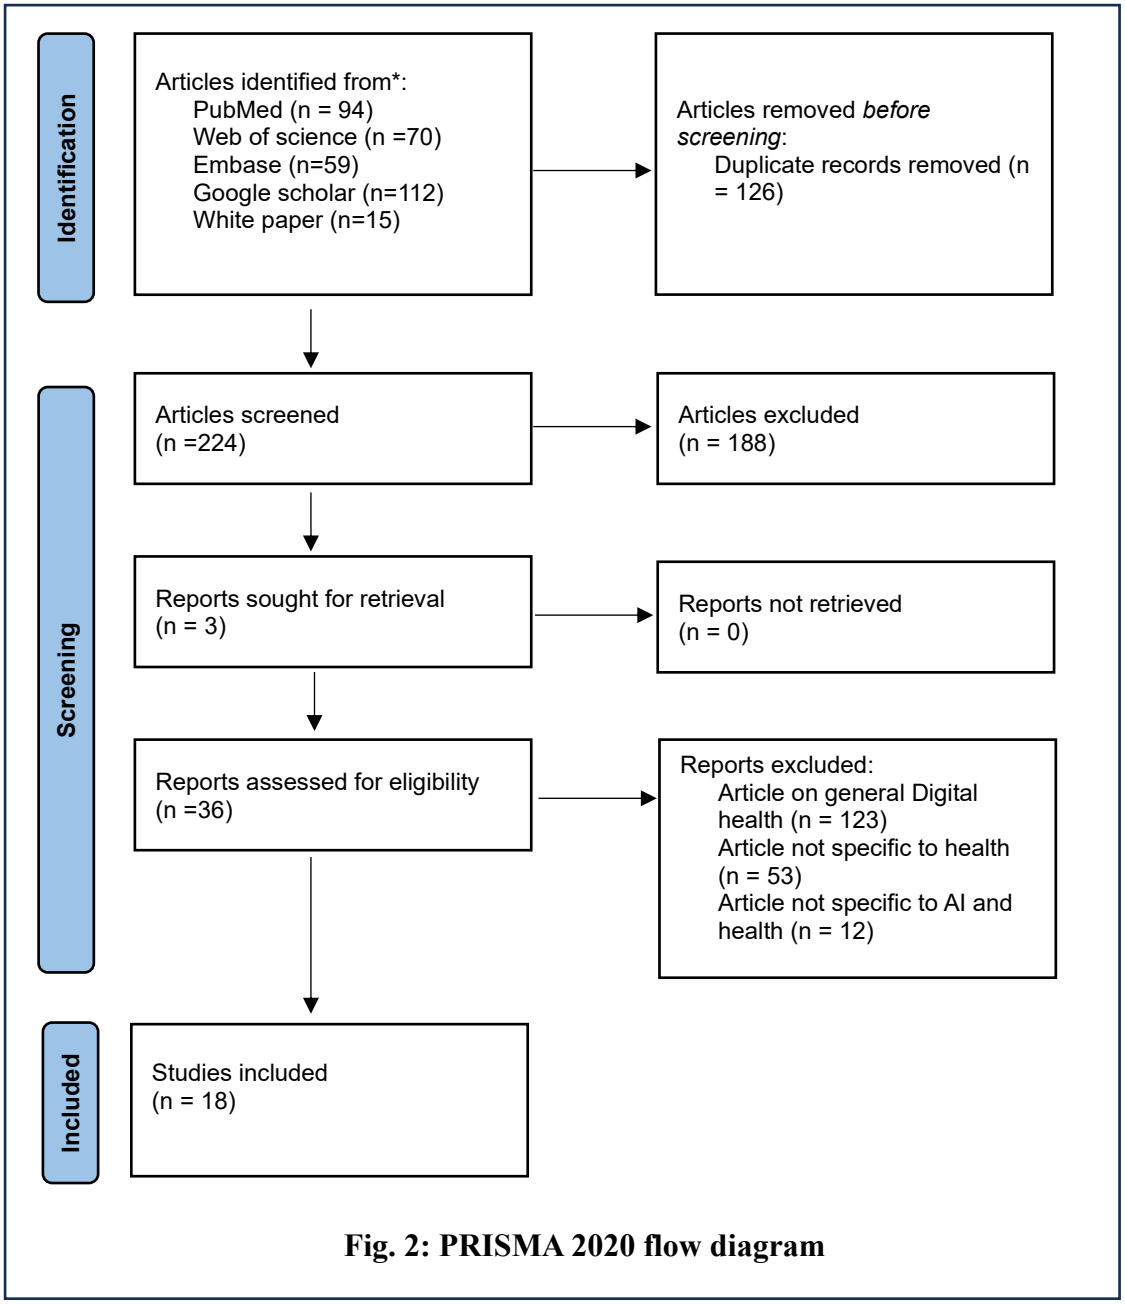


Fig. 2: PRISMA 2020 flow diagram

The table below summarises key information for each included study.

#	Author / Year	Title & Journal	Study Design	Key Findings	Objectives Addressed
1	Wolf et al., 2022 (61)	Potential reduction in healthcare carbon footprint by AI [npj Digital Medicine]	Observational comparative study (real-world)	Autonomous AI-enabled diabetic retinopathy screening reduced GHG emissions by ~80% compared to in-person specialist encounters. Travel reduction was the dominant factor.	Obj. 1 (Primary): AI-enabled carbon emission reductions; Obj. 3: Net climate balance of AI deployment
2	Alnafrah, 2025 (62)	The Two Tales of AI: Global , multidimensional policy assessment of environmental impacts from the Journal of Environmental Management.	Panel econometric study (56 countries, 2013-2023); MMQR, GMM, threshold modelling	High-intensity AI is significantly associated with increased CO2 emissions and environmental degradation (Digital Rebound Effect). Stringent environmental policies and diversified energy mix mitigate adverse effects.	Obj. 2 (Primary): Lifecycle environmental costs; Obj. 3: Net climate balance; Obj. 5: Policy and governance
3	Shehabi et al., 2024 (63)	2024 United States Data Center Energy Usage Report [Lawrence Berkeley National Laboratory (Grey Literature)]	National energy audit / scenario modelling report	US data centre electricity consumption estimated 176-390 TWh by 2028. AI workloads are primary driver of rapid growth. Provides scenario ranges for future demand.	Obj. 2 (Primary): Data centre energy costs; Obj. 4: Data centre infrastructure impact
4	Tilmant et al., 2026 (64)	What is the carbon footprint of a 100% digital pathology scenario in France? [Pathology]	Life Cycle Assessment (LCA) / carbon footprint modelling	Full digitisation of French surgical pathology generates 1,103-1,259 t CO2eq (3 months) to 2,116-2,923 t CO2eq (1 year storage). Data storage is the main contributor; digitisation	Obj. 1: Healthcare AI carbon footprint; Obj. 2 (Primary): Lifecycle environmental costs of AI in healthcare

				increased CF by 4-8% vs non-digitised.	
5	Kocak et al., 2025 (65)	Radiology AI and sustainability paradox: environmental, economic, and social dimensions [Insights Imaging]	Narrative/scoping review	AI in radiology can both reduce low-value imaging and generate significant energy/carbon costs. High implementation costs raise equity concerns. Calls for holistic sustainability strategies at individual, institutional and collective levels.	Obj. 1: AI emission reductions; Obj. 2: Environmental costs; Obj. 3: Net climate balance; Obj. 5: Policy recommendations
6	Bignami et al., 2025 (66)	Sustainable AI in medicine: navigating innovation, challenges, and environmental impact [Health Economics Review]	Narrative commentary/ review with case study reference	AI-driven hospital operations (predictive maintenance, resource allocation) improve efficiency but demand vast computational resources generating CO2, freshwater use, and e-waste. Green AI strategies (lifecycle assessment, circular approaches, low-carbon infrastructure) and telemedicine recommended.	Obj. 1: AI emission reductions; Obj. 2: Environmental costs; Obj. 3: Green AI applicability; Obj. 5: Governance
7	Bachina et al., 2025 (67)	Sustainable materials for AI technology adoption for energy-efficient patient-centric healthcare solutions [Journal of Education and Health Promotion]	Multidisciplinary study: LCA + systematic literature review + VOSviewer keyword analysis	Sustainable/eco-friendly materials in AI hardware can reduce energy consumption and environmental impact. Lifecycle evaluation from resource extraction to disposal needed. Eco-friendly hardware practices essential for sustainable healthcare ecosystem.	Obj. 2 (Primary): Lifecycle costs (materials, e-waste); Obj. 3: Green AI in healthcare
8	Ueda et al., 2024 (68)	Climate change and artificial intelligence in	Systematic narrative review	Comprehensive dual review: AI reduces healthcare GHG through telemedicine, workflow	Obj. 1: Emission reductions; Obj. 2: Environmental costs; Obj. 3: Net climate

		healthcare: Review and recommendations towards a sustainable future [Diagnostic and Interventional Imaging]		optimisation, remote monitoring. Simultaneously, AI training/inference and data centres contribute to GHG and e-waste. Energy-efficient models, renewable energy, eco-design, and policy frameworks recommended.	balance; Obj. 4: Green AI; Obj. 5: Policy
9	Alzoubi & Mishra, 2024 (69)	Green artificial intelligence initiatives: Potentials and challenges [Journal of Cleaner Production]	Comprehensive literature and web review; thematic categorisation	55 Green AI initiatives identified across 6 themes: cloud optimisation, model efficiency, carbon footprinting, sustainability-focused AI development, open-source, and research communities. Tools largely in early-stage development; standardised empirical evaluation lacking.	Obj. 3 (Primary): Green AI initiatives and applicability to healthcare; gaps in standardisation
10	Desislavov et al., 2023 (70)	Compute and Energy Consumption Trends in Deep Learning Inference [Sustainable Computing: Informatics and Systems]	Empirical computational analysis; performance vs energy modelling	Despite exponential parameter growth, energy consumption during inference grows more softly due to hardware efficiency improvements. However, multiplicative adoption factor means aggregate energy use continues rising. Highlights inference (not just training) as dominant energy phase at scale.	Obj. 2 (Primary): Lifecycle energy costs of AI; inference vs training cost balance
11	Luccioni et al., 2025 (71)	From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's	Theoretical/conceptual analysis; interdisciplinary (LCA + socioeconomic)	AI efficiency gains are offset by Jevons' Paradox — rebound effects increase consumption. Direct emissions alone misrepresent AI's true climate footprint. Governance, policy, and	Obj. 2: Environmental costs (indirect); Obj. 3: Net climate balance; Obj. 5: Policy and governance frameworks

		Polarised Environmental Debate [ACM FAccT Conference]		cultural norms must accompany technical efficiency measures.	
12	Patterson et al., 2022 (72)	The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink [Computer (IEEE)]	Empirical modelling; industry data analysis	Four best practices (efficient ML algorithms, processors, data centres, renewable energy) could collectively reduce training carbon footprint 100-1000x. Projects total ML training carbon will plateau then decline by 2030 if adopted.	Obj. 2: ML training energy costs; Obj. 3: Green AI mitigation strategies
13	de Vries-Gao, 2026 (73)	The carbon and water footprints of data centres and what this could mean for AI [Patterns (Cell Press)]	Environmental disclosure analysis; estimation modelling	AI carbon footprint estimated at 32.6-79.7 million tonnes CO2 in 2025; water footprint 312.5-764.6 billion litres. Lack of transparency in tech sector environmental reporting is a major barrier to accurate assessment. Policy mandating disclosure needed.	Obj. 2: Lifecycle costs (carbon and water); Obj. 4: Data centre infrastructure; Obj. 5: Reporting frameworks
14	Jegham et al., 2025 (74)	How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference [arXiv (preprint)]	Infrastructure-aware benchmarking framework; DEA efficiency analysis across 30 LLM models	Most energy-intensive models exceed 29 Wh/long prompt (65x most efficient). 0.42 Wh per query at 700M queries/day equates to annual electricity for 35,000 US homes, drinking water for 1.2M people, and requires a Chicago-sized forest to offset carbon.	Obj. 2 (Primary): Lifecycle costs of AI inference; benchmarking for standardised sustainability reporting
15	Bloomfield et al., 2021 (75)	Artificial Intelligence in the NHS: Climate and Emissions [Journal of Climate	Narrative scoping review (NHS as case study)	NHS AI investments lack standardised emissions quantification. AI can reduce NHS emissions via process optimisation and care model changes but could also add	Obj. 1: AI emission reductions in healthcare; Obj. 3: Green AI in NHS context; Obj. 5: Policy/governance; Reporting gaps

		Change and Health]		emissions from training/deployment. Recommends standardised metrics for healthcare AI climate impact; policy recommendations for net-zero NHS by 2040.	
16	Strubell et al., 2019 (76)	Energy and Policy Considerations for Deep Learning in NLP [ACL Proceedings]	Empirical computational cost analysis (NLP model training)	Training large NLP models generates substantial financial and environmental costs. Training a Transformer NAS model emitted ~626,000 lbs CO2e. Calls for actionable recommendations to reduce costs and improve equity in NLP/AI research.	Obj. 2 (Primary): Environmental costs of AI training — foundational quantification study
17	Luccioni & Hernandez-Garcia, 2023 (77)	Counting Carbon: A Survey of Factors Influencing the Emissions of Machine Learning [arXiv (preprint)]	Survey of 95 ML models; carbon emission analysis across NLP and computer vision tasks	Carbon emissions vary enormously across models and tasks. Energy source is a major moderating factor. Proposes centralised emissions repository for reporting and tracking. No single standardised methodology exists.	Obj. 2: Lifecycle carbon emissions of AI; Obj. 3: Standardisation and reporting frameworks gap
18	Tripathi et al., 2026 (CEEW) (78)	Scaling India's Data Centre Ecosystem: Stakeholder Perspectives on Infrastructure, Energy, and Resilience [Council on Energy, Environment and Water (CEEW) — Grey Literature]	Stakeholder consultation / white paper / policy analysis	India's data centre capacity grew from 520 MW (2020) to 1.5 GW (mid-2025), projected 4.5-6.5 GW by 2030. In 2024, data centres used ~0.5% of national electricity and 150 billion litres of water, both projected to double by 2030. 72% of Indian districts face extreme flood risk.	Obj. 4 (Primary): Healthcare data centre infrastructure in India/LMICs; Obj. 5: Policy for sustainable AI in high-growth markets

The result section has drawn on 18 studies to put together the data in this results section, which is organised thematically around the five research objectives. The pool of included work – published from 2019 through 2026 – is a mix of peer-reviewed articles, conference papers, grey literature reports and lifecycle assessments. In all, we have sourced these from disciplines as varied as computer science, public health policy, environmental science and healthcare informatics, with contexts ranging across the US, Europe, India and the globe.

Objective 1: AI-Enabled Reductions in Healthcare Carbon Emissions

The six studies listed here (*Wolf et al., Tilmant et al., Kocak et al., Bignami et al., Ueda et al. and Bloomfield et al*.) were chosen for their direct bearing on how AI can be used to cut carbon emissions in healthcare delivery by way of resource efficiency, workflow optimisation or travel substitution.

1.1 AI-Enabled Travel Reduction and Telemedicine

There is some compelling evidence on this front. Wolf et al. (2022) (61) put forward the most robust quantitative case, showing that when autonomous AI is employed for diabetic retinopathy screening, greenhouse gas (GHG) emissions are down by some 80 per cent versus an equivalent specialist visit in person. With patient travel being the main source of emissions, the GHG tally for an AI encounter is but a fraction of a traditional one; it is a real-world benchmark for the carbon advantage of autonomous AI in a diagnostic setting. These figures are supported by Ueda et al. (2024) (68). They point to the way in which AI-driven remote monitoring and virtual consultations influence transport-related carbon. Bignami and colleagues (2025) (66) put it much the same way, arguing that telemedicine and personalised support from a distance can cut healthcare's carbon output by lessening the demand for physical infrastructure. Then there is the case of the NHS:

Bloomfield et al. (2021) (75) make clear that if the service is to meet its net-zero 2040 target, it must be willing to shift care models from in-person to the virtual.

1.2 Workflow Optimisation and Resource Efficiency

In radiology, Kocak et al. (2025) (65) have identified something of a paradox for the community. On one hand AI may present a sustainability issue, but they found it can be part of the answer as well, by putting a brake on low-value imaging and ensuring energy and resources are used to better effect.

Tilmant et al. (2026) (64) offer a counterpoint that is worth noting. A lifecycle study of digital pathology in France by their team indicates that you do not even need AI to see an uptick in the carbon footprint of surgical pathology – as much as 4 to 8 per cent in fact, with data storage being the reason why. It is a reminder that you do not get environmental benefits from digitisation unless you are deliberate about it.

Then there is Bignami et al. (2025) (66) who see promise in using AI to manage patient flow and predictive maintenance in hospitals to curb energy waste, even if the AI systems' own computational appetite takes some of the edge off those gains.

Objective 2: Lifecycle Environmental Costs of AI in Healthcare

Here we have ten studies (including work by *Strubell, Desislavov, Patterson, Luccioni & Hernandez-Garcia, Jegham, de Vries-Gao, Shehabi, Bachina and Tilmant*) that shed light on the full lifecycle cost of AI in healthcare, from e-waste and critical mineral use to the water and carbon footprints of training and inference.

2.1 Energy Consumption: Training Phase

Strubell et al. (2019) were the ones to first put a number on the training costs in any meaningful way. They showed that a single large NLP Transformer model could let off 626,000 lbs of CO2e in training, or about what an average American car would produce in its entire life. That was a seminal piece of work for getting the topic into policy and scientific circles.

Patterson et al. (2022) have since modelled where those emissions are headed. They project that if four best practices are followed – from renewable energy to efficient ML algorithms – training emissions could be cut by as much as 1,000-fold, with total carbon for ML training set to peak and fall by 2030.

As for variability, a survey of 95 models by Luccioni & Hernandez-Garcia (2023) confirms there is no two ways about it: emissions can differ enormously depending on the task and type of model, with the energy source being the big variable. They also pointed out the lack of standardised reporting and the need for a central place to house such data.

2.2 Energy Consumption in the Inference Phase

The work of Desislavov et al. (2023) puts to rest the idea that a model's size dictates a proportional increase in inference energy. They argue that hardware has become efficient enough to put a damper on the kind of energy growth one might expect from parameters alone. Yet they are quick to point out that as AI is adopted on a multiplicative scale, the aggregate energy for inference is on a steep upward trajectory.

Jegham and colleagues (2025) have done some granular benchmarking that is worth noting. In their work they demonstrate the most extreme LLMs can be 65 times as voracious as top tier systems, with certain long prompts in excess of 29 Wh. To put those numbers in context: run 700 million queries a day at 0.42 Wh each and you are looking

at an annual electricity bill on par with 35,000 American households. The carbon offset for such consumption would require a forest as large as Chicago, while the evaporative water demand matches the yearly drinking needs of 1.2 million.

2.3 Data Centres: Carbon and Water Footprints

The scale of the problem is further underlined by de Vries-Gao (2026), who projects the global AI footprint for 2025 to be anywhere from 32.6 to 79.7 million tonnes of CO2e, alongside a water footprint of 312.5 to 764.6 billion litres. It is a transparency issue, plain and simple; absent more open environmental disclosure from the tech industry, you cannot get an accurate read on things, hence the urgency for mandatory reporting. Over in the US, Shehabi et al. (2024) are forecasting data centre power use will be pushed up to 176–390 TWh by 2028 thanks to AI. Such a wide margin reflects the uncertainty and makes a strong case for policy action and putting money into efficiency

2.4 Sustainable Materials and E-Waste

Then there is the matter of AI's lifecycle from a material standpoint, which is what Bachina et al. (2025) set out to examine. Manufacturing this hardware means dealing with rare earth elements and the destructive processes required to get them, while the short shelf-life of the equipment creates hazardous e-waste. Their life cycle assessment shows that in healthcare at least, using sustainable materials in AI hardware can have a measurable effect on both the environment and energy use.

Luccioni et al. (2025, FAccT) bring up Jevons' Paradox in an AI context: you gain efficiency but end up consuming more because of rebound effects. So technical

improvements on their own won't deliver a net reduction in environmental harm – something to bear in mind when making adoption decisions in healthcare.

Objective 3: Green AI and its Relevance to Healthcare

We have drawn on several studies here: *Alzoubi & Mishra (2024), Patterson et al. (2022), Ueda et al. (2024), Bignami et al. (2025), Bloomfield et al. (2021) and Luccioni & Hernandez-Garcia (2023).* Taken together, these six papers give us a view of current green AI initiatives, where they fit in healthcare and where the standardisation and reporting fall short.

3.1 What the Green AI Landscape Looks Like

Alzoubi & Mishra (2024) have done the most thorough job of mapping the field so far. They turned up 55 initiatives in six categories, from cloud optimisation and open-source frameworks to sustainability-focused development and research communities. But they are the first to admit that most of these tools are embryonic and lack any real empirical evaluation. It is a finding that goes straight to the heart of the dissertation's question on the limitations of today's Green AI in a healthcare setting.

3.2 Gaps in the Healthcare Context

Ueda et al. (2024) make the case for responsible AI procurement in healthcare, putting eco-design and lifecycle assessment at the forefront alongside renewable energy and green computing.

There are not many standardised ways to measure AI emissions in healthcare, according to Bloomfield et al. (2021), which hampers any effort to mitigate climate impact. Five

years ago, they were calling on everyone from NHS policymakers to AI developers to put some metrics in place; that is still an open issue. For transparency and accountability, Luccioni & Hernandez-Garcia (2023) have put forward the idea of a centralised repository for ML model carbon emissions. And while Patterson et al. (2022) have laid out best practices for everything from processor choice to data centre design that could be applied to healthcare AI, there has been no systematic evaluation of them in this context yet.

Objective 4: The Environmental Cost of Healthcare Data Centres in India

This section considers four studies – *Tripathi et al. (2026, CEEW), Shehabi et al. (2024), de Vries-Gao (2026) and Alnafrah (2025)* – on the impact of data centre infrastructure, with an eye to fast-growing markets such as India.

4.1 India's Data Centre Expansion and Environmental Implications

The clearest picture of where India's data centres are headed comes from Tripathi et al. (2026, CEEW). Their work shows installed capacity has almost tripled in recent years, going from 520 MW in 2020 to 1.5 GW by the middle of 2025, and is on course for 4.5 to 6.5 GW in 2030. The environmental cost is already apparent: some 150 billion litres of water were used in 2024 and roughly 0.5% of the nation's electricity, numbers that will likely be more than doubled a decade from now. Then there is the matter of climate-infrastructure risk; with 72% of Indian districts prone to extreme flooding and 57% to heat waves, the resilience of healthcare AI systems is put in question.

4.2 Global Data Centre Environmental Footprint

Shehabi et al. (2024) give you some sense of scale when it comes to the US energy problem, with benchmark figures for global AI workloads in the infrastructure they require – we are looking at 176 to 390 TWh by 2028. But as de Vries Gao (2026) is quick to note, there is an opacity to the whole thing. You cannot get a true read on the footprint of a data centre or the AI it hosts unless tech firms are put under pressure to be open about their environmental record.

Then there is what Alnafrah (2025) calls the "technological leapfrogging paradox" in lower-middle-income countries such as India. Regulators simply cannot keep up with the pace of AI adoption, so the sustainability standards you would expect in a high-income nation are absent and environmental degradation is left to run its course without oversight.

Objective 5: Framework for Sustainable AI in healthcare

We have collated the evidence from eight studies to build a framework that is as sustainable as it is balanced, one that spans from policy to cross-sector endeavours. Our sources include *Ueda et al. (2024), Bignami et al. (2025), Kocak et al. (2025), Bloomfield et al. (2021), Alnafrah (2025), Luccioni et al. (2025), Tripathi et al. (2026) and de Vries Gao (2026).*

5.1 Policy and Governance Interventions

Take upper-middle-income economies on the verge of embracing AI; Alnafrah (2025) makes a quantitative argument for integrated governance there, contending that if you have hard-line environmental policies and solid digital infrastructure, you can counteract the more intense side of AI. De Vries-Gao (2026) takes a harder line, insisting that disclosure should be made mandatory. Tech companies ought to be required to present standardised metrics as a matter of course if you want any serious governance of

healthcare AI sustainability. Luccioni et al. (2025, FAccT) warn that technical efficiency alone will not yield net environmental gains. To do that you have to align social norms and business incentives and deal with market forces and rebound effects like Jevons' Paradox that fuel uncontrolled growth.

5.2 Healthcare Specific Framework

Ueda et al. (2024) provide the most thorough list of best practice in the sector, covering responsible model development and data handling right through to renewable energy and eco-design in procurement. Bignami et al. (2025) consider circular resource management an institutional imperative, whether that is recycling equipment or making use of waste heat. Kocak et al. (2025) are also at pains to point out that a holistic approach has to factor in social equity; otherwise, you risk entrenching disparities in healthcare through unequal access and algorithmic bias. As for Bloomfield et al. (2021) and Tripathi et al. (2026), they observe that your strategy must be fit for purpose, be it for NHS targets or the regulatory ground realities of an LMIC.

5.3 Reporting and Standardisation Gaps

If there is one point of consensus among the studies for Objective 5, it is the absence of a standard measure for the climate impact of healthcare AI tools. That is a void identified by all, from Bloomfield et al. (2021) to de Vries-Gao (2026) and Luccioni & Hernandez-Garcia (2023). To bridge the divide between current practice and our net-zero commitments, any credible framework for the clinic will have to produce standardised carbon and environmental reporting requirements. This is a primary objective of the dissertation and something that is sorely needed.

Discussion:

This review was done to put dual environmental impact of artificial intelligence in healthcare under a critical microscope. Our aim was to weigh the climate advantages against the lifecycle costs and, in doing so, synthesize the evidence according to five set objectives.

What we found in the 18 studies we have included – covering everything from public health policy and computer science to environmental science and healthcare informatics – is that AI in healthcare is a veritable double-edged sword when it comes to climate change. The data neither lends itself to an unqualifiedly optimistic nor a pessimistic view. Rather, there is a conditional dynamic at play: whether AI has a positive net effect on the climate depends on a host of factors, from the computational architecture and the type of application to the energy source of the infrastructure and the governance in place, not to mention any rebound effects. We look at each of these in turn below.

1. The Carbon Benefit of AI in Healthcare: Real but Narrow and Contingent

If you are looking for the case for AI's environmental merit in healthcare, the most persuasive example is provided by Wolf et al. (2022). In a real-world comparison they made, autonomous AI screening for diabetic retinopathy was shown to cut greenhouse gas emissions some 80% compared with a specialist seeing the patient in person. It is the kind of headline-grabbing statistic that has come to define the conversation around AI's climate potential and stands as the best empirical benchmark in the peer-reviewed literature.

The 80% number you see is really a comparison of the marginal emission from an AI interaction to that of a specialist's visit, and it is there simply because the patient is not making the trip. It is not the whole story when it comes to the environmental bill of the AI system itself; it says nothing of the hardware that must be made or the energy

required for training and running inference at scale. Once you throw those upstream costs in the equation, the net benefit is nothing like the headline numbers would suggest. Nobody can deny the value of using telemedicine and remote monitoring to reduce travelling: the literature is brimming with examples. The Aravind Eye Care System in India has developed an enormous telemedicine infrastructure which saves millions of patient miles and more than 1,050 metric tonnes of CO2 a year. In the U.S., UCLA has calculated telemedicine savings between 21.4 and 47.6 million kgs per month (the equivalent of taking 130,000 SUVs off the road). All tangible benefits, of course. But there is a flipside in our analysis, and that's that simply going digital in health care is not an environmental no-brainer.

A life cycle assessment of the French surgical pathology pipeline, even without AI, concluded that full digitization added 4-8 per cent to the lab's CO2 footprint through digital storage alone (Tilmant et al, 2026), and flies in the face of assumptions made in policy debates that 'digital is greener'. Kocak et al (2025) observe a paradox with radiology AI: you might cut low-value scanning, but you end up with a cost in hardware waste and high energy use. AI might cut emissions, but it has only a certain scope and your own context (and grid) matters hugely.

2. Lifecycle Environmental Costs of AI: Massive, Growing, and Unreported

If the evidence for a single lifecycle cost for any given AI system is consistent and more than you will find in the papers about why we need AI, the literature on the lifecycle environmental cost of an AI system is more consistently than you can imagine. Strubell et al. (2019) may be the study that broke the dam on the field: their analysis revealed that it takes 626,000 lbs. Of COe to train a single, large NLP transformer model - about as much as five lifetime's worth of 5 passenger U.S. cars. Pattersen et al. (2022) has

shown that, under optimal circumstances, by mixing renewable electricity, smarter data centres, and better hardware, there can be a reduction of emissions by 1000x.

But do not assume that healthcare organizations are adopting this best practice as they purchase new AI capabilities. There is also inference costs: it is less studied than training costs, but probably of greater consequence at population level in terms of operational power usage.

Desislavov et al. (2023) point out that even if superior hardware gives you a softer energy curve, once AI is routine in the clinic the volume of queries will drive up aggregate energy use regardless of how efficient you are on a per-query basis. Jegham et al. (2025) put some striking numbers to this dynamic. A query of 0.42 Wh may seem trivial but scale it to 700 million a day and you are looking at enough electricity for 35,000 US homes a year, or the carbon output of a Chicago-sized forest to offset, not to mention the freshwater equivalent of what 1.2 million people would drink in a year. You do not have to imagine these are hypotheticals in healthcare AI. The NHS primary care sector, for example, was thought to be churning out 349 tonnes of CO_2e a year from off-the-books ChatGPT use, an environmental liability that is entirely unmeasured and unprocured.

We also see in this review that the water footprint of AI is a matter of equal gravity, even if the healthcare literature does not make much of it. De Vries-Gao (2026) puts the global figure for 2025 at between 312.5 and 764.6 billion liters and expects it to grow. In India the situation is more acute. Karnataka's IT Ministry says a megawatt of data center capacity will need 25 million liters of water a year, putting them in direct competition with municipal users for potable water, as the 2024 Bengaluru shortage made clear. And in the extreme heat we are seeing more of across India, data centers need more cooling water right when it is at a premium.

When it comes to healthcare data systems that must be up and running in the face of a climate-fueled health emergency, you are dealing with a compound risk. Yet the material reality of these systems – the rare earth mining, the manufacturing, the e-waste – has not been studied with the empirical rigor it warrants. Bachina et al. (2025) were the most candid about it in our review, pointing out that the three- to five-year life span of AI hardware is what produces hazardous waste, all on top of the environmental damage done to mine the components. They advocate sustainable materials to counter it – although it's not what the healthcare industry is seeing today. This is an architectural failure mode for choosing vendors. The impact of their manufacturing upstream isn't ever considered.

3. Jevons' Paradox and the Rebound Effect: The Most Underappreciated Risk

The most under-appreciated risk that should keep healthcare leaders up at night as they consider artificial intelligence is the observation around Jevons' Paradox. As Luccioni et al (2025) would put it very clearly, efficiency per model or per inference does not lower the overall environmental impact; consumption will grow to fill any slack.

That was illustrated well by the 2010 data centre record which demonstrated how hyperscale architecture was fantastic for computational efficiency, but worldwide power consumption went up. The workload simply grew to overwhelm efficiencies. This is not a theory. In the world of healthcare AI, this can manifest as building efficiency in a clinical diagnostic tool, a decision support tool, or a transcribing tool; the efficiencies are built for the individual use cases, but overall work volume will grow and consume that efficiency.

An approach of solely improving the per-unit efficiency without considering the rates of adoption will prove structurally unsound. As these authors advocate, a sustainable AI

strategy necessitates real business, financial, and governance controls that serve to manage the rebound effects, and the evidence would certainly support their assertion.

4. Green AI Initiatives: A Promising but Underutilized Field in Healthcare

Perhaps the broadest overview of the field comes from Alzoubi and Mishra (2024) who categorise 55 Green AI projects into 6 themes. Several of these will be familiar to many – for example the Green Algorithms framework and solutions like CodeCarbon used in the NHS for energy tracking. The Energy Consumption Disclosure mentioned below is the same used in an experiment in the Nature paper by Luccioni et al (2024), showing that users are 12 times more likely to pick a low energy alternative to a high energy algorithm if given the relevant information. Alzoubi and Mishra do acknowledge some shortcomings.

The current crop of tools has yet to be proven in a real-world environment and are still immature. The authors highlighted that not even one of the 55 projects has been deployed in healthcare. When you consider privacy and compliance issues together with 24/7 service expectations in healthcare it is hardly surprising translating these tools for general use is difficult. That challenge was laid down by Bloomfield et al (2021) who had also looked at the NHS, finding little widely accepted standard measurement of AI emissions, even some years after net zero was declared; the gap between measurement and actual usage appears to have widened considerably with the recent rush to implement AI. However, there are some incredible potential figures out there in proof of concept tools; the '4 M's' method of minimizing AI training carbon emissions from Patterson et al (2022) estimated reductions of up to 747 fold, and our own review identified the MobileNet-Lite-Health edge computing approach as requiring 74% less power and generating 76% less CO2 emissions compared to a comparable cloud based

LLM, though both of these rely on procurement and contractual arrangements that hospital departments do not currently have in place to hold vendors accountable for their environmental footprint.

5. The India Challenge: Global Risk in Microcosm and Data Centre Infrastructure

Where the rise of AI, the reality of environmental risk, and the promise of resilient healthcare all meet, is the Indian data centre ecosystem. As Tripathi et al. (2026) and Alnafrah (2025) clearly state, the numbers paint a picture of an acceleration that is dizzying: spurred by the IndiaAI Mission, fuelled by foreign investment, and shaped by digital sovereignty, Indian data centre installed capacity has almost trebled in the last 5 years, with an projected increase to between 4.5 and 6.5 GW by 2030, marking the nation as a nexus of global expansion. Yet there are also looming shadows. Some 72% of Indian districts face a high level of flood risk, and 57% a severe heat stress – issues which both increase the strain on cooling water and put the resilience of the facilities to the test.

Alnafrah (2025) refers to a “paradox of leapfrogging” - a state where the uptake of AI, the engine behind this expansion, outstrips the speed of the regulation of the technology’s environmental impact. Just 5 of the 15 state policies enacted currently contain a sustainability component and not even one makes water stress mapping a prerequisite for approval. Indeed, Visakhapatnam approved a 1GW installation in a district with known groundwater challenges and without any divulgence of operational water requirements – a fatal structural flaw for healthcare AI. The platform that makes telemedicine widely accessible across India - eSanjeevani - runs on infrastructure in areas as vulnerable to climate hazards as the very healthcare systems it’s meant to bolster.

This should give all of those engaged in AI planning, and specifically those deploying

AI in healthcare, a reason to reflect. An AI tool in healthcare is doing little to combat the climate crisis if the underlying physical infrastructure is managed with poor governance and is at the whims of extreme weather. We can consider telemedicine, all of the positives of reduced patient travel, wasted if a heatwave or flood closes the servers it's hosted on. At your table, environmental sustainability and infrastructure resilience must, from this point forward, be considered indistinguishable components of your AI planning.

6. Governance, Policy and the Road to Net-Positive Healthcare AI

The data on Objective 5 confirm what should be an increasingly obvious and alarming pattern. The hurried rollout of AI for use in healthcare has left its environmental governance in the dust and barely capable of dealing with its environmental impacts.

This becomes clear when looking at high income countries where their Green Plans may well have omitted any mention of tracking the carbon footprint of digital, or where an institution's inventory of emissions may choose to avoid looking at Scope 3 Hardware emissions.

It's even more apparent in a market like India where national programmes, such as SAHI or BODH, are focused primarily on clinical and privacy aspects, ignoring environmental metrics entirely. Then there is the EU which boasts the most sophisticated approach to regulating AI, a combination of the EU AI Act and the Climate Neutral Data Centre Pact which cover all manner of impacts from wasted heat to water use, including mandatory disclosure, though they aren't without their flaws- often, confidentiality concerns allow firms to ignore the full detail of facilities data and AI-related workload reporting remain a matter of choice. A far more relaxed approach prevails in the US where regulatory guidelines favour the speed of innovation over regulation as regulation, such as that outlined in the NHS purchasing framework with

the inclusion of a 10% weight on net-zero and social value, or the UK Cabinet Office regulation whereby all suppliers in receipt of a contract of greater than 5 million must present an emissions reduction plan, will inspire imitation amongst other institutions. Our literature review shows why such policy making can be problematic.

It's all too difficult for organizations to measure environmental performance of vendors if standard databases lack independently collected carbon footprint data relating to digital systems or to medicines.

It's an impossibility for the purchaser. Therefore, if healthcare wants to implement its own climate change and environmental agenda, the priority must be closing the data gap. De Vries-Gao (2026) has laid down a prescription for it clearly: AI infrastructure operators must be obliged to disclose their environmental performance at the workload level and provide an accurate accounting of their carbon emissions, water and energy footprint data.

The material that has passed before us provides little ambiguity on this point; that the ability for the researcher to demonstrate system-level impacts; for regulators to operate meaningfully and for healthcare organizations to make fully informed purchase decisions. It forms the centrepiece of our theoretical framework.

7.Towards an Integrated Framework for Sustainable Healthcare AI: Synthesis of Evidence

The total of the five objective findings, and any valid proposal to make healthcare sustainable must include working on all of these four fronts - the model, the institution, national policy and international standards.

On the model front, health systems need to insist on 4 Ms from their vendors. This is efficient architectures and adequate hardware, cloud-based services with the green energy you can track and ensure is being sourced from the grid you trust and the model

is optimised for the grid and not a power guzzling grid. I would love to see a software dashboard with Energy Consumption Disclosure labels. This means the clinician in the clinic would be able to estimate what level of compute is needed and make a decision whether it is justifiable or not.

At the institution level, I want to make it mandatory for any new technology to have a comprehensive lifecycle assessment carried out, not only the energy cost during inference, but also the carbon footprint during the training of the ML model, the hardware manufacture, and eventually the e-waste produced by these devices and hardware.

Now on to the government level, moving from voluntary agreements to mandated obligations. It should no longer be an optional exercise when you are giving permission to a new AI facility to set up an in climate-sensitive zones without mapping out regions that is under water stress. It should not be possible to have a data centre established without environmental disclosures as a requirement of licensing, all health technology vendors to a health service should have to submit a carbon footprint reduction plan for the devices they sell. Just look at India who is developing an all-India policy that will push data centres beyond any existing sustainability measures. Their current state level data centre policy has five states who has a stronger position than national government for example India is trying to mandate measurement for operational water use.

International standards are the right framework for putting some ground rules. ISO/IEC standard has recently started taking water footprint in to account for the work around sustainable AI. I think that will eventually lead to a set of standardised methods to measure and define AI and its impact in the healthcare sector, I think that will be further cemented by some centralised and openly accessible data centres where we will record carbon footprints of ML models.

8.Limitations of This Review

To start with, our evidence base is heavily skewed towards high-income countries; you will find most of it hails from the U.S., the U.K. and Europe, with little in the way of primary data from LMICs like India. The CEEW report by Tripathi et al. (2026) and Alnafrah's (2025) multi-country econometric work are notable exceptions, but overall, we are underrepresenting the very places where AI infrastructure is expanding at breakneck speed and environmental governance is still in its infancy.

Then there is the matter of the quantitative evidence. Much of it pertains to general-purpose AI as opposed to what is being put in place for healthcare. While the environmental cost of the underlying infrastructure is transferable, the carbon savings one might expect from healthcare AI – be it in optimizing workflows or cutting down on travel – are highly context-dependent and will not necessarily hold up across different health system architectures or geographies. We must also factor in the speed of technological change; some of the numbers cited in the studies we have included, especially on training energy use, could already be old news given how fast models are scaling. Even the EPRI/Epoch AI projections for power demand through to 2030 come with a good deal of uncertainty considering the volatility of AI development. Finally, as a scoping review we have not done a formal risk-of-bias or quality assessment of the studies. That is something future systematic reviews ought to do to put the evidence on firmer ground.

Conclusion:

In putting together this dissertation I have aimed for the most thorough integrated assessment yet of AI as a double-edged sword for climate change in the realm of healthcare delivery. Drawing on 18 studies that cover everything from empirical benchmarking and lifecycle analysis to grey literature from policy bodies and econometric modelling, the evidence points to five main conclusions.

1. The climate benefits of AI in healthcare are real, if conditional

There is no denying the quantifiable, and in some cases considerable, advantages. Telemedicine and autonomous diagnostics are the best examples. You have the Wolf et al. (2022) figure of an 80% drop in GHG per encounter, or Aravind Eye Care's annual system-level savings of more than 1,050 metric tonnes. At the US population level, telemedicine has cut emissions to the tune of 130,000 vehicles off the road every month. Such applications can make a dent in the sector's carbon footprint. But do not assume these are inherent to AI; they are confined to certain clinical pathways and depend on deliberate design and grid conditions.

2. Lifecycle costs of environment are huge and underreported

Whatever's missing from hospital carbon tracking is the lifecycle cost of the hardware, inference and training costs of the AI. When considering 626,000 pounds of COe in the air just to run a single NLP query and 79.7m tonnes of COe around the world, any one hospital deployment sounds tiny. While there is a lot of commentary on training, this paper reveals the population scale cost lies squarely with inference and massive query volumes.

3. Jevons' Paradox is the structural problem to solve

Making AI systems more efficient is not enough if adoption outpaces those gains. We see strong support for rebound effects in AI consumption. Any sustainability strategy for healthcare AI that is worth its salt will need governance that looks at the scale and nature of adoption, not just per-unit efficiency.

4. Green AI is here but not yet part of the healthcare fabric

There are 55 Green AI initiatives out there and tools for tracking energy or assessing procurement carbon. Yet none have been put to the test in a healthcare setting, and we lack standardized ways to measure the full lifecycle impact. The pieces are there in prototype form, but they will need a mandate and some investment before they are of any operational use.

5. India and LMICs face a compounded risk that demands urgent, context-specific policy

What is happening in India is a microcosm of the wider problem facing low- and middle-income countries. The country's data center capacity is growing at breakneck speed in places vulnerable to climate change, yet the policies that oversee it say little about environmental sustainability. Add to this the "leapfrogging paradox" of AI being use faster than regulation can keep up, and you have a compounded risk. In the case of healthcare AI, you cannot separate the resilience of digital health infrastructure from its environmental credentials. A data center in a water-stressed or flood-prone zone will not be able to provide telemedicine or AI diagnostics on which we rely for healthcare's carbon benefits. Environmental sustainability and climate resilience need to be viewed as co-equal imperatives when planning such infrastructure.

Final Remarks and Recommendations

So, to answer our central research question on whether AI in healthcare is a net cost or benefit to the climate: the honest answer is it depends. We do not have the data or the governance in place to make a systematic determination for most AI in healthcare. One can point to the conditions that would make an AI deployment a net positive for the climate – say, travel-substituting apps run on renewables with efficient architecture at a scale where they matter in high-emission populations – but these are seldom put to the test empirically, one way or the other.

This review leads to some pressing recommendations. Healthcare bodies should be making use of tools like the 4 M's framework or the Green Algorithms calculator to factor in the lifecycle carbon and water footprint of their AI procurement right now. National health systems ought to make carbon reporting mandatory for Scope 3 emissions tied to digital and AI, covering everything from the hardware lifecycle to unprocured shadow AI. In fast-moving markets such as India, governments should be imposing binding environmental conditions on data center licenses, complete with stress mapping and an accounting of how much water is being used. On the world stage, we must ensure that we continue to do what we can to try and find a way of creating standards around measurement of healthcare AI, building from what is there in terms of the ISO/IEC standard, to ML emissions repository in the offing.

Nobody can argue, in terms of efficiency within the clinical settings or population management, that AI cannot be seen as having made a difference. But most health systems are not, from a financial standpoint or an operational perspective, calculating the environmental impact of it. We should not do so out of the fact that we are going to do more harm with climate, and be responsible, when we implement its usage, is what I propose with my evidence and conceptual framework within the thesis.

References

1. United Nations Environment Programme. AI has an environmental problem. Here's what the world can do about that [Internet]. UNEP; 2025 Nov 13 [cited 2026 Apr]. Available from: https://www.unep.org/news-and-stories/story/ai-has-environmental-problem-heres-what-world-can-do-about
2. Faiyazuddin M, Rahman SJQ, Anand G, Siddiqui RK, Mehta R, Khatib MN, et al. The impact of artificial intelligence on healthcare: a comprehensive review of advancements in diagnostics, treatment, and operational efficiency. Health Sci Rep. 2025;8(1): e70312. doi:10.1002/hsr2.70312
3. International Energy Agency. Energy and AI [Internet]. Paris: IEA; 2025 [cited 2026 Apr]. Available from: https://iea.blob.core.windows.net/assets/de9dea13-b07d-42c5-a398-d1b3ae17d866/EnergyandAI.pdf
4. global.noharm.org/media/4454/download?inline=1 [Internet]. [cited 2026 Apr 19]. Available from: https://global.noharm.org/media/4454/download?inline=1
5. Bloomfield P, Clutton-Brock P, Pencheon E, Magnusson J, Karpathakis K. Artificial Intelligence in the NHS: Climate and Emissions☆, ☆☆. J Clim Change Health. 2021 Oct 1; Getting to a Greener, Healthier World by 2030: Solutions from Healthcare professionals4:100056. Doi: 10.1016/j.joclim.2021.100056
6. Ueda D, Walston SL, Fujita S, Fushimi Y, Tsuboyama T, Kamagata K, et al. Climate change and artificial intelligence in healthcare: review and recommendations towards a sustainable future. Diagn Interv Imaging. 2024;105(10):453–459. Doi: 10.1016/j.diii.2024.06.002
7. International Energy Agency. Energy and AI [Internet]. Paris: IEA; 2025 [cited 2026 Apr]. Available from: https://iea.blob.core.windows.net/assets/de9dea13-b07d-42c5-a398-d1b3ae17d866/EnergyandAI.pdf
8. Alzoubi YI, Mishra A. Green artificial intelligence initiatives: potentials and challenges. J Clean Prod. 2024; 468:143090. Doi: 10.1016/j.jclepro.2024.143090
9. Tripathi V, Bagui D, Aggarwal P, Hulshof P, Chopra A, Jain D, et al. Scaling India's data centre ecosystem: stakeholder perspectives on infrastructure, energy, and resilience [Internet]. New Delhi: Council on Energy, Environment and Water (CEEW) and SYSTEMIQ; 2026 Feb [cited 2026 Apr]. Available from: https://ceew.in
10. Ueda D, Walston SL, Fujita S, et al. Climate change and artificial intelligence in healthcare: Review and recommendations towards a sustainable future. Diagn Interv Imaging. 2024 Nov;105(11):453-459. Doi: 10.1016/j.diii.2024.06.002.

11. Faiyazuddin M, Rahman SJQ, Anand G, et al. The impact of artificial intelligence on healthcare: a comprehensive review of advancements in diagnostics, treatment, and operational efficiency. Health Sci Rep. 2025;8(1): e70312. doi:10.1002/hsr2.70312.

12. Indugu N, Sai S. YOLOv5-based cervical cancer detection and medical waste classification: A sustainable artificial intelligence framework. Preprints. 2025;202510.1561.

13. Graduate Medical Education & AMA. Reconciling AI's clinical utility with its environmental footprint. J Graduate Med Educ. 2024;16(2):120-125.

14. Bignami E, Darhour LJ, Bellini V. Sustainable AI in medicine: navigating innovation, challenges, and environmental impact. Health Econ Rev. 2025 Dec 30;15(1):110. doi:10.1186/s13561-025-00704-w.

15. Ministry of Electronics & Information Technology (MeitY). Written reply on water consumption by data centers in Rajya Sabha. New Delhi: Parliament of India; 2026 Mar 30.

16. KPMG in India. India's data centre revolution: The integrated lifecycle blueprint (2026-2030). Mumbai: KPMG; 2026 May 19.

17. Belkhir L, Elmeligi A. Assessing ICT global emissions footprint: Trends to 2040 & recommendations. J Clean Prod. 2018 Mar 10; 177:448-463. Doi: 10.1016/j.jclepro.2017.12.239.

18. Murugesan S. Green Assets and Infrastructure. In: Murugesan S, Gangadharan GR, editors. Harnessing Green IT: Principles and Practices. Chichester: John Wiley & Sons, Ltd; 2012. p. 19-33. Doi: 10.1002/9781119597537.ch2.

19. de Vries A. The growing energy footprint of artificial intelligence. Joule. 2023 Oct 18;7(10):2191-2194. Doi: 10.1016/j.joule.2023.09.004.

20. Jones N. How to stop data centres from gobbling up the world's electricity. Nature. 2018 Sep;561(7722):163-166. doi:10.1038/d41586-018-06610-y.

21. Li P, Yang J, Islam MA, Ren S. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. Commun ACM. 2025 Jul;68(7):54-63. doi:10.1145/3724499.

22. techUK & Thames Water. Water Efficient Data Centres. London: techUK & Thames Water; 2026 Feb.

23. Tripathi V, Bagui D, Aggarwal P, et al. Scaling India’s Data Centre Ecosystem: Stakeholder Perspectives on Infrastructure, Energy, and Resilience. New Delhi: Council on Energy, Environment and Water (CEEW) and SYSTEMIQ; 2026 Feb.

24. NHS England. Five years of Greener NHS progress and forward look. London: NHS England; 2025.

25. NHS England. Environmental benefits of digital over paper: Wayfinder App case study. London: NHS England; 2025 Apr.

26. Bloomfield PS, Clutton-Brock P, Pencheon E, Magnusson J, Karpathakis K. Carbon Reporting Practices in the NHS: Emissions and Omissions. J Clim Change Health. 2021 Oct 1; 4:100056. Doi: 10.1016/j.joclim.2021.100056.

27. European Parliamentary Research Service (EPRS). Artificial intelligence and energy: Impact of data centres on Europe's electricity grid. Brussels: European Parliament; 2025.

28. European Environment Agency (EEA). Data Center Water Abstraction and Hydrological Impact in Europe. Copenhagen: EEA; 2025.

29. Climate Neutral Data Centre Pact (CNDCP). White paper on data centre responsible use of water review. Brussels: CNDCP; 2024 Sep.

30. Borderstep Institute. Data centres in the EU – Facts & Figures. Berlin: Borderstep Institute; 2025.

31. European Data Centre Association (EUDCA). State of European Data Centres 2025 Report. Brussels: EUDCA; 2025.

32. Swedish Energy Agency. Data centre energy performance reporting Act (2025:570). Stockholm: Swedish Energy Agency; 2026 Mar 24.

33. Travers Smith. Regulating the data centre boom: EU Energy Efficiency Directive. London: Travers Smith; 2025.

34. European Commission. New research shows how data centres could become water-positive and carbon-negative using waste heat. Brussels: European Commission; 2026 Mar 30.

35. Marinoni A, et al. India's AI Data Centres Are Draining Water as El Niño Threatens a Weak 2026 Monsoon. Down To Earth. 2026 Mar.

36. Netsol Water. Water requirement and consumption by hyperscale data centers in India. Netsol Blog. 2026.

37. Mongabay-India. India bets on data centres, even as water-energy use concerns mount in Visakhapatnam. Visakhapatnam: Mongabay; 2026 May.

38. National Health Systems Resource Centre (NHSRC). Utilization of Telemedicine/ eSanjeevani in the Public Health Facilities of India. New Delhi: NHSRC; 2025 Sep.

39. Smile Foundation India. The Green Gains of Telemedicine: Rise of Telemedicine in India and Globally. Smile Blog; 2025.

40. Rani A, et al. Telemedicine saved 1,666 km of travel, lowered CO2 emissions by 176.6 kg. Eye. 2024 Apr 10.

41. Electric Power Research Institute (EPRI), Epoch AI. Scaling Intelligence: The Exponential Growth of AI's Power Needs. Palo Alto (CA): EPRI; 2025 Aug. Product ID: 3002033669.

42. Ramachandran A, Sarabu C, Gupta U, Ghose S, Lee VS. Sustainably Advancing Health AI: A Decision Framework to Mitigate the Energy, Emissions, and Cost of AI Implementation. NEJM Catalyst. 2025 Sep 12.

43. Digital Health Review. Systematic review of greenhouse gas emissions savings from digital health technologies. J Digital Health. 2025; 11:102-115.

44. Connor MJ, et al. Carbon footprint savings from telemedicine implementations in clinical settings. PMC Public Health. 2021; 14:204-210.

45. UCLA Health. Telemedicine use in 2023 reduced monthly carbon emissions. Los Angeles: UCLA; 2025 Apr 22.

46. Aravind Eye Care System. Assisted telemedicine in satellite vision centers: Thiruppuvanam case study. Tamil Nadu: AEH; 2021.

47. Health Care Without Harm (HCWH). GGHH climate learning initiative: Aravind Eye Care System decarbonisation case study. Brussels: HCWH; 2023.

48. Aravind Eye Care System. Environmentally sustainable high-volume eye care model. Madurai: AEH; 2025 Aug.

49. Thiel CL, Schehlein E, Ravilla T, et al. Cataract surgery and environmental sustainability: Waste and lifecycle assessment of phacoemulsification at Aravind Eye Care System. J Cataract Refract Surg. 2017;43(11):1391-1398. Doi: 10.1016/j.jcrs.2017.08.018.

50. van der Scheer J, Lannelongue L, Smith J, et al. Methods to assess the environmental impact of AI tools in the NHS. Cambridge: THIS Institute, University of Cambridge; 2025.

51. Raggio CB, Migliorelli L, Skupien N, et al. Energy-Efficient Federated Learning via Adaptive Encoder Freezing for MRI-to-CT Conversion: A Green AI-Guided Research. arXiv preprint arXiv:2512.03054. 2025 Dec.

52. J Imaging. A Study on Energy Consumption in AI-Driven Medical Image Segmentation. J Imaging. 2025 May 26;11(6):174.

53. Shaukat F, et al. MobileNet-Lite-Health: A Sustainable Edge AI Framework for Resource-Scarce Environments. IEEE Access. 2025/2026.

54. Patterson D, Gonzalez J, Hölzle U, et al. The carbon footprint of machine learning training will plateau, then shrink. Joule. 2022 Oct 19;6(10):2261-2266. Doi: 10.1016/j.joule.2022.08.014.

55. Luccioni S, Gamazaychikov B, Hooker S, et al. Light bulbs have energy ratings — so why can't AI chatbots? Nature. 2024 Aug;632(8026):736-738. doi:10.1038/d41586-024-02680-3.

56. Open MedScience. How medical imaging is going green: The breakthroughs shaping a cleaner, smarter future. Open MedScience. 2024 Oct 6.

57. Health Aff Scholar. Artificial intelligence governance: A comparative study of regional policies. Health Aff Scholar. 2026;3(6): qxaf108.

58. NHS England. Delivering a Net Zero National Health Service strategy progress report. London: NHS England; 2025.

59. Public Policy Projects (PPP) & YewMaker. Accelerating health sector decarbonisation through medicines carbon footprinting. London: PPP; 2026 Feb 24.

60. Union Health Ministry. Launch of SAHI and BODH at India AI Impact Summit 2026. New Delhi: Ministry of Health and Family Welfare; 2026.

61. Wolf I, et al. Potential reduction in healthcare carbon footprint by autonomous AI. npj Digit Med. 2022.

62. Alnafrah I. The Two Tales of AI: Global assessment of environmental impacts from a multidimensional policy perspective. J Environ Manage. 2025.

63. Shehabi A, et al. 2024 United States Data Center Energy Usage Report. Berkeley (CA): Lawrence Berkeley National Laboratory; 2024.

64. Tilmant S, et al. What is the carbon footprint of a 100% digital pathology scenario in France? Pathology. 2026.

65. Kocak B, et al. Radiology AI and sustainability paradox: environmental, economic, and social dimensions. Insights Imaging. 2025.

66. Bignami E, et al. Sustainable AI in medicine: navigating innovation, challenges, and environmental impact. Health Econ Rev. 2025.

67. Bachina S, et al. Sustainable materials for AI technology adoption for energy-efficient patient-centric healthcare solutions. J Educ Health Promot. 2025.

68. Ueda K, et al. Climate change and artificial intelligence in healthcare: Review and recommendations towards a sustainable future. Diagn Interv Imaging. 2024.

69. Alzoubi YI, Mishra A. Green artificial intelligence initiatives: Potentials and challenges. J Clean Prod. 2024.

70. Desislavov R, et al. Compute and Energy Consumption Trends in Deep Learning Inference. Sustain Comput Inform Syst. 2023.

71. Luccioni AS, et al. From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarised Environmental Debate. In: Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT). 2025.

72. Patterson D, et al. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. Computer (IEEE). 2022.

73. de Vries-Gao L. The carbon and water footprints of data centres and what this could mean for AI. Patterns (N Y). 2026.

74. Jegham I, et al. How Hungry is AI? Benchmarking Energy, Water, and Carbon Footprint of LLM Inference. arXiv [Preprint]. 2025.

75. Bloomfield H, et al. Artificial Intelligence in the NHS: Climate and Emissions. J Clim Change Health. 2021.

76. Strubell E, et al. Energy and Policy Considerations for Deep Learning in NLP. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL). 2019.

77. Luccioni AS, Hernandez-Garcia A. Counting Carbon: A Survey of Factors Influencing the Emissions of Machine Learning. arXiv [Preprint]. 2023.

78. Tripathi A, et al. Scaling India's Data Centre Ecosystem: Stakeholder Perspectives on Infrastructure, Energy, and Resilience. New Delhi: Council on Energy, Environment and Water (CEEW); 2026.

79. Moher D, Shamseer L, Clarke M, Ghersi D, Liberati A, Petticrew M, Shekelle P, Stewart LA. Preferred Reporting Items for Systematic Review and Meta-Analysis Protocols (PRISMA-P) 2015 statement. Syst Rev. 2015;4(1):1. doi: 10.1186/2046-4053-4-1

Prachi Thakral D report

ORIGINALITY REPORT

6%	5%	4%	3%
SIMILARITY INDEX	INTERNET SOURCES	PUBLICATIONS	STUDENT PAPERS

PRIMARY SOURCES

1	arxiv.org Internet Source	1%
2	pmc.ncbi.nlm.nih.gov Internet Source	1%
3	www.ceew.in Internet Source	<1%
4	Submitted to City University Student Paper	<1%
5	Submitted to Miami University of Ohio Student Paper	<1%
6	Submitted to Texas A&M University, Central Texas Student Paper	<1%
7	pubmed.ncbi.nlm.nih.gov Internet Source	<1%
8	Submitted to University of Essex Student Paper	<1%
9	www.theresearchinsights.com Internet Source	<1%
10	www.elsevier.es Internet Source	<1%
11	scholarspace.manoa.hawaii.edu Internet Source	<1%
12	Submitted to University of Canberra Student Paper	<1%
13	Submitted to Vrije Universiteit Amsterdam Student Paper	<1%
14	www.azoai.com Internet Source	<1%
15	"Energy Informatics", Springer Science and Business Media LLC, 2026 Publication	<1%

Prachi Thakral

Prachi Thakral D report

Quick Submit

Quick Submit

IIHMR

Document Details

Submission ID

trn:oid:::1:3613470938

Submission Date

Jul 19, 2026, 12:06 PM GMT+5:30

Download Date

Jul 19, 2026, 12:13 PM GMT+5:30

File Name

Dissertation_Research_Project_18.072026.docx

File Size

105.2 KB

41 Pages

17,298 Words

93,165 Characters

*% detected as AI

AI detection includes the possibility of false positives. Although some text in this submission is likely AI generated, scores below the 20% threshold are not surfaced because they have a higher likelihood of false positives.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.