

LIFESTREAM AI: Ethical AI for Chronic Disease Prevention

A Research-in-Progress Report

Submitted to the EIU Journal of Action Research

Research Domain: Digital Health Systems | Responsible AI | Preventive Medicine

Submission Category: Research-in-Progress Report

Date: March 2026

[A provisional patent application (No. 63/998,499) covering the innovations described in this work was filed with the United States Patent and Trademark Office on March 6, 2026.]

Abstract

What if the most important moment in a patient's health journey occurs years before diagnosis? Chronic lifestyle-related diseases, including cardiovascular disease, Type 2 diabetes, and metabolic syndrome, account for over 74% of global deaths and are projected to generate a macroeconomic burden of multiple trillions of US dollars in lost output and health expenditures over the coming decades (World Health Organization, 2025; Ferranna et al., 2023). Yet the majority of these conditions are preventable when risk is identified early. LIFESTREAM AI is a research-grade digital health platform designed to address a critical gap in preventive medicine: continuous, clinically integrated, and explainable individual-level risk intelligence. A provisional patent application (No. 63/998,499) was filed with the United States Patent and Trademark Office (USPTO) on March 6, 2026, covering the below platform's core innovations.

Three innovations underpin this work: a four-tier blockchain-enforced consent architecture granting patients granular control over their health data; a carbon-aware federated learning scheduler that defers model training to periods of lower grid carbon intensity with modelling-based estimates suggesting potential emission reductions of 70–75% compared

with a fixed-schedule baseline (conceptual projection, not yet empirically validated). An Explainable AI (XAI) clinical alert prioritisation system using SHapley Additive Explanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) to rank alerts by clinical actionability rather than raw probability, directly addressing clinician alert fatigue.

The platform is designed in compliance with HIPAA, GDPR, and the EU Artificial Intelligence Act (EU AI Act), reflecting a commitment to ethical, transparent, and equitable AI in healthcare.

Keywords: federated learning, explainable AI, digital health, chronic disease prevention, carbon-aware computing.

1. Context and Rationale

Most healthcare systems today are built to respond to illness rather than prevent it. For most healthy adults, doctor visits are few and far between — separated by months, sometimes even years. During those long stretches in between, their health goes largely unmonitored and unrecorded. Research shows the real cost of this gap: conditions such as pre-diabetes and early cardiovascular disease are often diagnosed only after years of underlying metabolic or vascular change have already occurred (American Diabetes Association, 2023). By the time a problem becomes clinically visible, the opportunity for simple, low-cost lifestyle intervention has often already passed.

There is a substantial preclinical period during which metabolic changes precede the clinical diagnosis of type 2 diabetes, while the body is still responsive to intervention. During this window, structured lifestyle interventions have been shown to reduce the risk of progressing to type 2 diabetes by about 58% in high-risk individuals, with benefits sustained over long-term follow-up (Knowler et al., 2002; American Diabetes Association, 2023). What is missing is the practical infrastructure to detect risk early enough and the tools to support individuals in sustaining meaningful behaviour change over time.

Wearable devices and health apps have genuinely changed how people engage with their own bodies. Millions now track their steps, sleep, and heart rate daily. But having data is not the same as understanding it. Most of what these devices collect never reaches a clinician, never feeds into a prediction, and never triggers an intervention. It simply accumulates, unexamined. The research world has moved quickly on AI-powered health monitoring, particularly on making predictions more accurate. However, accuracy alone is insufficient. Data privacy, environmental cost, and clinical interpretability remain under-addressed. These three gaps privacy, sustainability, and explainability are precisely what LIFESTREAM AI is designed to address, as a patent-protected research contribution to a doctoral programme in Business Administration. This research therefore contributes a conceptual framework for designing preventive digital health platforms that simultaneously address patient privacy, environmental sustainability, and clinical explainability.

2. Research Question and Objectives

The central research question guiding this study is:

Can an AI-driven health monitoring platform close the gap between when chronic disease begins and when it is diagnosed — and do so while meeting the competing demands of individual privacy protection, clinical explainability, and environmental sustainability within regulated healthcare systems?

This question further breaks down into 3 research objectives, each corresponding to one of the platform's key innovations:

- To construct a four-tier consent system, secured through blockchain smart contracts, which enables patients to manage the handling of their health data in four distinct contexts, from the usage of personal device only to full clinical integration.
- To develop and test a federated learning scheduling system that delays the transmission of energy-demanding model updates until a time of reduced carbon emissions on the power grid.
- To develop and trial a clinical alert ranking method based on SHapley Additive

Explanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME) to score and rank AI risk alert reports according to their “real-world” actionability, and not based on how high the predicted probability is.

3. Proposed and Ongoing Methodology

3.1 Design Science Research Justification

This study adopts a **Design Science Research (DSR)** methodology, which is well-suited for research that involves building and evaluating technological solutions while also contributing to broader. The central artefact developed in this study is **LIFESTREAM AI**, an AI-driven platform designed to enable continuous, privacy-preserving, and explainable monitoring of chronic disease risk.

DSR is particularly appropriate for this research for three reasons. First, the problem being addressed is inherently socio-technical, requiring a careful balance between technological design and clinical, ethical, and regulatory considerations. Second, the study aims not only to develop a functional system but also to derive design insights that can guide the development of ethical AI systems in healthcare. Third, the iterative nature of DSR aligns closely with the staged development of the platform.

Accordingly, the research follows a three-phase DSR cycle:

- (1) **Design and conceptualisation**, involving architecture definition;
- (2) **Prototype development and validation**, currently in progress; and
- (3) **Clinical evaluation**, planned subject to ethics approval.

Through this approach, the study seeks to deliver both a validated artefact and a structured framework for designing AI-enabled preventive health systems.

3.2 Artefact Architecture

LIFESTREAM AI is conceptualised as a unified research artefact whose architecture operationalises the study's three core objectives across four interrelated components.

3.2.1 Five-Tier Platform Architecture. The platform integrates IoT wearable sensors, on-device edge computing, a federated AI core, a blockchain governance layer, and role-specific decision-support interfaces. At the edge layer, raw physiological data remains on the patient's device; only structured summaries are transmitted, substantially reducing cloud-bound data volume and establishing a foundational privacy layer independent of user consent configuration (Li et al., 2020; Kaissis et al., 2020).

3.2.2 Privacy Architecture: Four-Tier Blockchain Consent. Privacy is operationalised through a four-tier consent architecture enforced via smart contracts on a permissioned blockchain (Hyperledger Fabric). Tiers range from fully local processing (Tier 1, default) through anonymous federated updates (Tier 2), de-identified research sharing (Tier 3), and full clinical integration with a named clinician (Tier 4). All consent decisions are cryptographically recorded and non-modifiable, aligning with GDPR and EU AI Act requirements.

3.2.3 Carbon-Aware Federated Learning Scheduling: A carbon-aware scheduling scenario based on prior work showing substantial variation in the carbon intensity of compute across time and region (Patterson et al., 2021; Dodge et al., 2022). Under this scenario, federated training is restricted to the cleanest quartile of hours, defined by real-time grid carbon intensity. Under this idealised assumption, the average carbon intensity of training could fall to roughly one quarter to one third of the baseline, implying potential emission reductions on the order of 70–75% compared with a naive fixed-schedule system. This remains a conceptual estimate that will be tested in the simulation phase.

3.2.4 XAI-Calibrated Clinical Alert Prioritisation. To address alert fatigue, alerts are ranked using a Clinical Alert Priority Score (CAPS) that combines SHAP-derived feature contributions, which attribute importance values to individual features for each prediction (Lundberg & Lee, 2017), with a clinician-defined Actionability Weight Matrix (AWM) so that alerts are prioritised by clinical relevance rather than raw probability. LIME provides complementary local validation by offering model-agnostic, instance-level explanations for selected alerts (Ribeiro et al., 2016). The AWM is version-controlled via the blockchain

audit layer, supporting transparency and regulatory compliance in how actionability criteria evolve over time.

3.3 Evaluation Plan

The evaluation of LIFESTREAM AI follows three sequential phases consistent with DSR principles of iterative artefact refinement.

Phase 1 — Simulation-Based Feasibility

Using synthetic patient data, this phase evaluates: the carbon savings of the scheduling mechanism against a fixed-schedule baseline; clinician response differences between CAPS-ranked and probability-ranked alerts; and patient comprehension of the four-tier consent model. Findings will inform artefact refinement before prototype deployment.

Phase 2 — Prototype Validation

Internal validation of a functional prototype against defined performance metrics, including clinician alert response rates, alert fatigue reduction, consent comprehension scores, and carbon intensity of federated training cycles. Results will guide iterative refinement in line with the DSR evaluation cycle.

Phase 3 — Clinical Pilot Evaluation

A clinical pilot study with external healthcare partners providing real-world empirical validation and forming the primary basis for theoretical contribution. This phase is subject to Institutional Review Board (IRB) approval and the establishment of clinical partnerships and will proceed following successful completion of Phase 2. No patient data has been collected at any stage to date.

4. Description of Planned and Early Action

At present, LIFESTREAM AI remains at the conceptual design stage. No prototype has yet been built, and no clinical data has been collected. This submission outlines the research direction and invites peer scrutiny of its core assumptions.

A provisional patent application (No. 63/998,499) was filed with the United States Patent and Trademark Office (USPTO) on March 6, 2026, covering the platform's key innovations. The provisional patent application covers the platform's core architectural innovations, specifically the four-tier blockchain-based consent mechanism, the carbon-aware federated learning scheduling approach, and the XAI-calibrated clinical alert prioritisation framework. The research presented in this paper addresses the conceptual design and evaluation strategy of these components rather than their proprietary implementation details. Concurrent literature and prior-art reviews across the three innovation domains are underway to further validate the novelty of the proposed mechanisms.

Following the literature review, the research moves towards a simulation-based feasibility study using synthetic patient data. This is necessary to test several core assumptions whether the carbon scheduling delivers the projected savings, whether clinicians respond differently to CAPS-ranked alerts, and whether patients can navigate a four-tier consent model intuitively.

A clinical pilot is planned as the final phase of this research, subject to ethics approval and clinical partnership development. This phase will be pursued following successful completion of the simulation-based feasibility study.

5. Reflection on Challenges and Learning

Because this is early-stage work rather than a completed study, the challenges identified here are not problems already solved but open design questions that will shape how LIFESTREAM AI evolves. They span technical, clinical, ethical, and governance concerns that must be addressed before any real-world deployment

1. Privacy–Accuracy trade-off and bias in federated learning

Stronger privacy protection via differential privacy, controlled by its key parameter, epsilon (ϵ), the "privacy budget" which typically reduces model accuracy, while

higher ϵ preserves accuracy but weakens privacy guarantees; the appropriate working range (on the order of $\epsilon \approx 1-3$ in clinical contexts) will need to be co-designed with clinicians, ethicists, and patient representatives (Dyda et al., 2021; Mohammadi et al., 2026). The parameter epsilon (ϵ) quantitatively measures how much information about any one individual the model's outputs are allowed to reveal. Smaller ϵ values mean stronger privacy guarantees but require adding more noise and therefore reduce model utility, whereas larger ϵ values allow more accurate models at the cost of weaker privacy protection. Published healthcare AI implementations highlight this trade-off and the risk that populations with lower wearable adoption (e.g., older adults, lower-income groups, or under-resourced regions) become under-represented in federated updates, so the evaluation plan will include explicit per-subgroup equity audits before any clinical deployment (Vayena et al., 2018; World Health Organization, 2021).

2. Defining and governing clinical actionability (CAPS/AWM)

Clinical actionability is context-dependent: what is feasible in a tertiary hospital may not be feasible in a time-constrained primary care consultation. The CAPS formula therefore carries a specific risk if Actionability Weight Matrix values are miscalibrated, for example by over-weighting factors patients cannot realistically change or under-weighting urgent clinical signals; mis-specified weights could reintroduce alert fatigue regardless of SHAP accuracy. This necessitates structured clinical engagement and prospective validation of AWM configurations against retrospective clinical cases, with all revisions captured via the blockchain audit trail to support transparency and regulatory scrutiny.

3. Informed consent and cognitive load in a four-tier model

Digital consent is often poorly understood by patients, particularly when interfaces rely on dense legal language or assume high digital literacy, raising the risk that tier selections do not reflect genuinely informed choices (Norman & Skinner, 2006; Vayena et al., 2018; Torous et al., 2018). In a tiered consent model, this risk is

amplified, and the interface must therefore prioritise plain language, progressive disclosure, comprehension checks, and careful management of cognitive load, especially for older adults and users with lower eHealth literacy (Sweller, 1994; Holden & Karsh, 2010).

4. Engagement and sustained adoption of preventive mHealth tools

Preventive mHealth applications frequently show high early attrition, with engagement falling steeply within the first few months (Krebs & Duncan, 2015; Klasnja & Pratt, 2012). For LIFESTREAM AI, this means risk insights and behavioural nudges (e.g., via push notifications, gamified progress) must feel relevant, trustworthy, and actionable daily, backed by clear data governance and seamless consent renewal.

5. Governance, regulation, and current study status

Governance of the blockchain consent architecture (who operates permissioned nodes, who may update smart contracts, and how immutability interacts with GDPR's right to erasure) remains to be formally specified, as do alignments with the EU AI Act's traceability requirements (European Commission, 2024; World Health Organization, 2021). More broadly, the platform is currently at a conceptual stage: no prototype has been implemented, no clinical data have been collected, and all reported performance figures—including projected emission reductions and improvements in alert response—are modelling-based estimates rather than validated outcomes. Real-world feasibility will depend on future Institutional Review Board approval, clinical partnerships, and stakeholder engagement, so this work should be interpreted as a design and evaluation agenda rather than evidence of demonstrated clinical effectiveness.

Acknowledgements

I acknowledge limited use of Claude (Anthropic) for language editing and writing refinement. All research concepts and technical contributions are my own and I declare no

conflicts of interest. All references have been manually verified against publisher databases to ensure accuracy and traceability.

References

- American Diabetes Association Professional Practice Committee. (2023). Standards of care in diabetes—2023. *Diabetes Care*, 46(Suppl. 1), S1–S291. <https://doi.org/10.2337/dc23-S001>
- Dodge, J., Prewitt, T., Combes, R. T., Odmark, E., Schwartz, R., Strubell, E., Luccioni, A. S., Smith, N. A., DeCario, N., & Buchanan, W. (2022). Measuring the carbon intensity of AI in cloud instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1877–1894). Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533234>
- Dyda, A., Purcell, M., Heymann, D., & Brownstein, J. S. (2021). Differential privacy for public health data: An innovative tool to optimize information sharing while protecting data confidentiality. *Patterns*, 2(12), 100366. <https://doi.org/10.1016/j.patter.2021.100366>
- European Commission. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Ferranna, M., Cadarette, D., & Bloom, D. E. (2023). The macroeconomic impact of noncommunicable diseases. In D. E. Bloom & S. Sevilla (Eds.), *The Oxford handbook of noncommunicable diseases*. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780192896803.013.21>
- Holden, R. J., & Karsh, B. T. (2010). The technology acceptance model: Its past and its future in health care. *Journal of Biomedical Informatics*, 43(1), 159–172. <https://doi.org/10.1016/j.jbi.2009.07.002>
- Kaissis, G. A., Makowski, M. R., Rückert, D., & Braren, R. F. (2020). Secure, privacy-preserving and federated machine learning in medical imaging. *Nature Machine Intelligence*, 2(6), 305–311. <https://doi.org/10.1038/s42256-020-0186-1>
- Klasnja, P., & Pratt, W. (2012). Healthcare in the pocket: Mapping the space of mobile-phone health interventions. *Journal of Biomedical Informatics*, 45(1), 184–198. <https://doi.org/10.1016/j.jbi.2011.08.017>
- Knowler, W. C., Barrett-Connor, E., Fowler, S. E., Hamman, R. F., Lachin, J. M., Walker, E. A., & Nathan, D. M. (2002). Reduction in the incidence of type 2 diabetes with lifestyle intervention or metformin. *New England Journal of Medicine*, 346(6), 393–403. <https://doi.org/10.1056/NEJMoa012512>
- Krebs, P., & Duncan, D. T. (2015). Health app use among US mobile phone owners: A national survey. *JMIR mHealth and uHealth*, 3(4), e101. <https://doi.org/10.2196/mhealth.4924>
- Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765–4774). <https://arxiv.org/abs/1705.07874>

- Lundberg, S. M., Nair, B., Vaithianathan, R., & Lee, S.-I. (2020). Explainable machine-learning predictions for the prevention and treatment of dyslipidaemia. *Nature Biomedical Engineering*, 4(2), 214–226. <https://doi.org/10.1038/s41551-019-0414-8>
- Mohammadi, M., Vejdanihemmat, M., Lotfinia, M., Rusu, M., Truhn, D., Maier, A., & Tayebi Arasteh, S. (2026). Differential privacy for medical deep learning: methods, tradeoffs, and deployment implications. *npj Digital Medicine*, 9, 93. <https://doi.org/10.1038/s41746-025-02280-z>
- Norman, C. D., & Skinner, H. A. (2006). eHealth literacy: Essential skills for consumer health in a networked world. *Journal of Medical Internet Research*, 8(2), e9. <https://doi.org/10.2196/jmir.8.2.e9>
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., & Dean, J. (2021). *Carbon considerations for large neural network training*. arXiv. <https://doi.org/10.48550/arXiv.2104.10350>
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939778>
- Sweller, J. (1994). Cognitive load theory, learning difficulty, and instructional design. *Learning and Instruction*, 4(4), 295–312. [https://doi.org/10.1016/0959-4752\(94\)90003-5](https://doi.org/10.1016/0959-4752(94)90003-5)
- Torous, J., Nicholas, J., Larsen, M. E., Firth, J., & Christensen, H. (2018). Clinical review of user engagement with mental health smartphone apps: Evidence, theory and improvements. *BMJ Evidence-Based Medicine*, 23(5), 187–189. <https://doi.org/10.1136/eb-2018-102891>
- Vayena, E., Blasimme, A., & Cohen, I. G. (2018). Machine learning in medicine: Addressing ethical challenges. *PLOS Medicine*, 15(11), e1002689. <https://doi.org/10.1371/journal.pmed.1002689>
- World Health Organization. (2021). *Global strategy on digital health 2020–2025*. <https://www.who.int/publications/i/item/9789240020924>
- World Health Organization. (2025, September 24). *Noncommunicable diseases* [Fact sheet]. <https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases>